

Regular Articles

Syst. Biol. 54(4):527–529, 2005
Copyright © Society of Systematic Biologists
ISSN: 1063-5157 print / 1076-836X online
DOI: 10.1080/10635150590947023

Phylogenetic Diversity and the Greedy Algorithm

MIKE STEEL

Allan Wilson Centre for Molecular Ecology and Evolution, Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand;
E-mail: m.steel@math.canterbury.ac.nz

Abstract.—Given a phylogenetic tree with leaves labeled by a collection of species, and with weighted edges, the “phylogenetic diversity” of any subset of the species is the sum of the edge weights of the minimal subtree connecting the species. This measure is relevant in biodiversity conservation where one may wish to compare different subsets of species according to how much evolutionary variation they encompass. In this note we show that phylogenetic diversity has an attractive mathematical property that ensures that we can solve the following problem easily by the greedy algorithm: find a subset of the species of any given size k of maximal phylogenetic diversity. We also describe an extension of this result that also allows weights to be assigned to species [Biodiversity conservation; greedy algorithm; phylogenetic diversity.]

Let \mathcal{T} be an unrooted phylogenetic X -tree, that is, a tree whose leaves comprise the set X of species (or populations) under study, and whose remaining vertices are of degree at least 3. Let λ be the edge-weight function that records, for each edge of \mathcal{T} , its (non-negative) branch length. Given a subset W of X , we can consider the induced phylogenetic W -tree, denoted $\mathcal{T}|W$ that connects just those species in W and its associated edge weighting λ_W , which assigns to each edge e of $\mathcal{T}|W$, the sum of the $\lambda(e)$ values over those edges of \mathcal{T} in the path that corresponds to e . An example is illustrated in Figure 1; formal definitions of these concepts are provided in Semple and Steel (2003).

The *phylogenetic diversity* of W , denoted $PD(W)$, is defined as

$$PD(W) = \sum_e \lambda_W(e)$$

where the summation is over all edges e in the tree $\mathcal{T}|W$. This definition is equivalent to that given by Faith (1992).

For example, in Figure 1, for $W = \{a, b, g, e\}$ we have $PD(W) = 3 + 2 + 4 + 5 + 4 = 18$. Note that $PD(W)$ also depends on (\mathcal{T}, λ) , but to simplify notation we will think of these as fixed. Also, in case $|W| = 1$, we set $PD(W) = 0$. Later we will consider an extension of PD that allows each species to also have a weight assigned to it.

The PD score was formally introduced by Dan Faith in 1992 (Faith, 1992) and has been subsequently applied in areas of biodiversity conservation (see for example Barker, 2002; Mace et al., 2003; Mooers et al., 2005, and the references therein). The score provides some indication of how much feature diversity (sensu Faith, 1992) each possible subset W contains in relation to the entire variation in the tree (by comparing $PD(W)$ to $PD(X) = \sum_e \lambda(e)$) and so may therefore be useful for determining which subsets of species might be best to

conserve when it is not possible to conserve them all. The PD score has also recently been investigated mathematically by Pachter and Speyer (2004) but for a different purpose—namely, to extend the classic result on tree reconstruction from pairwise distances to m -wise values.

In this note we show how to efficiently find (and characterize) subsets of X of given size that have maximal PD score. Clearly, to examine all subsets of X of given size k is not feasible if k is large; however, we show that the sets of any given size that have maximal PD score are precisely the ones that can be built up using a greedy approach starting with the maximal PD sets of size 2. This greedy approach was suggested in Faith’s original paper (Faith, 1992: 5, column 1, paragraph 2), though no claim (or proof) was made that this was anything more than a locally maximal procedure, as we will see it also leads to a global maximization of PD. Using the greedy algorithm described below, it would be straightforward to find a subset of species of size (say) 100 having maximal phylogenetic diversity for that number of species, in an edge-weighted tree on (say) 1000 species. The algorithm is valid for both binary (i.e., fully resolved) phylogenetic trees and nonbinary trees containing polytomies. Also, the algorithm described here solves a more general problem than a related greedy approach described by Nee and May (1997), since we do not assume that the branch lengths satisfy a molecular clock, or impose a particular rooting on the tree.

MAIN RESULT

To state our results we require some further definitions. For $k \geq 1$ let

$$pd_k = \max\{PD(W) : W \subseteq X, |W| = k\}$$

and let

$$PD_k = \{W \subseteq X : |W| = k \text{ and } PD(W) = pd_k\}.$$




FIGURE 1. (Left) An edge-weighted tree. (Right) The induced edge weighted tree on the subset $W = \{a, b, g, e\}$, which is the (unique) set of four species of largest PD score among all sets of four species.

Thus pd_k is the largest possible phylogenetic diversity value across all subsets of species of size k , whereas PD_k is the set of all collections of k species that realize this maximal phylogenetic diversity.

Theorem 1 PD_k consists precisely of those subsets of X of size k that can be built up as follows: Select any pair of species that are maximally far apart (in the edge-weighted tree (T, λ)) and then sequentially add elements of X so as to maximize at each step the increase in PD score.

Example

For the pair (T, λ) in Figure 1, the greedy algorithm will start by selecting $\{a, b\}$, which is the unique element of PD_2 , and then add leaf e and then g to obtain the set $W = \{a, b, e, g\}$, which is the unique element in PD_4 . Thus, for this example, $pd_4 = 18$.

Proof of Theorem 1

The proof relies on first establishing the following fundamental mathematical property of the PD score. Suppose we are given a pair (T, λ) and subsets W, W' of X with $2 \leq |W'| < |W|$. Then there always exists some species $x \in W - W'$ so that:

$$PD(W - \{x\}) + PD(W' \cup \{x\}) \geq PD(W) + PD(W'). \quad (1)$$

To establish this property, we first introduce some notation: we refer to the edges of a phylogenetic tree as *exterior* if they are incident with a leaf (i.e., a species in X) otherwise we say that the edge is *interior*.

Consider the tree $T|W$. We can view the tree T as being obtained from $T|W$ by attaching a set \mathcal{F} of subtrees of T to certain vertices and (subdivisions of) edges of $T|W$. For example, in Figure 1 T is obtained from $T|W$ by attaching the three (one-vertex) subtrees $c, d,$ and f to single vertex subdivisions of the exterior edges of $T|W$ incident with $b, e,$ and g , respectively.

Because $|W'| < |W|$ and $|W'| > 2$ there exists at least one exterior edge e of $T|W$ with the property that any subtree

of \mathcal{F} that attaches to e does not contain any species in W' . Let $x \in W - W'$ be the leaf of T incident with e . We have

$$PD(W) = PD(W - \{x\}) + \lambda_W(e). \quad (2)$$

Now, because e was chosen in $T|W$ so that any subtree in \mathcal{F} that contains a species in W' must attach either to some edge of $T|W$ that is different to e , or to a vertex of $T|W$, we have

$$PD(W' \cup \{x\}) \geq PD(W') + \lambda_W(e). \quad (3)$$

Combining (2) and (3) gives (1).

Now suppose that $W \in PD_k$ and $W' \in PD_{k-1}$. Select a species $x \in W - W'$ to satisfy (1). Then:

$$W' \cup \{x\} \in PD_k \quad \text{and} \quad W - \{x\} \in PD_{k-1} \quad (4)$$

because $PD(W - \{x\}) \leq PD(W)$ with equality precisely if $W - \{x\} \in PD_{k-1}$ and $PD(W' \cup \{x\}) \leq PD(W')$ with equality precisely if $W' \cup \{x\} \in PD_k$, which, combined with (1), gives (4).

Theorem 1 now follows easily from (4) by standard arguments from "greedoid" theory (Korte et al., 1991). Specifically, (4) shows that any element of PD_k (for $k \geq 3$) is obtained from any element of PD_{k-1} by adding a single new element of X , which must by necessity be an element that maximizes the increase in PD score.

An Extension

Suppose we have a function $f: X \rightarrow \mathbb{R}$ and a constant $a > 0$, and we let

$$PD_{f,a}(W) = a \cdot PD(W) + \sum_{x \in W} f(x).$$

For example, in biodiversity conservation, a positive score $f(x) > 0$ might allow for additional incentives to preserve x (for example if it is globally endangered);

alternatively, a negative score $f(x) < 0$ might be the cost associated with conserving species x . This type of modification to PD was investigated by Faith and Walker (1994) in the context of “trade-offs” in biodiversity conservation.

Finding a subset W of size k to maximize $PD_{f,a}(W)$ is easy for any function f , by applying the same greedy approach (starting by finding a set W of size 2 to maximize $PD_{f,a}(W)$, and sequentially adding new species so as to maximize the $PD_{f,a}$ score). That is, Theorem 1 applies if we replace PD by $PD_{f,a}$ for any choice of f and $a > 0$. This follows from the observation that for any $W, W' \subseteq X$ with $x \in W - W'$, the difference

$$\Delta(f, a, x) = PD_{f,a}(W - \{x\}) + PD_{f,a}(W' \cup \{x\}) - (PD_{f,a}(W) + PD_{f,a}(W'))$$

is independent of f and so, by property (1), species x can be chosen so that $\Delta(f, a, x)$ is non-negative (because $a \cdot PD(W) = PD_{0,a}(W)$ for the zero function $0(x) = 0$ for all $x \in X$).

In particular, this extension allows us to determine easily whether or not any given species x in X is in every set (or in any set) of maximal phylogenetic diversity among all subsets of X of size k , without having to check them all exhaustively. To achieve this, let $pd_k(f, a)$ denote the maximal value of $PD_{f,a}(W)$ over all subsets W of X of size k and let us compare $pd_k, pd_k(+\delta_x, 1)$ and $pd_k(-\delta_x, 1)$ where

$$\pm\delta_x(x') = \begin{cases} \pm 1, & \text{if } x = x'; \\ 0, & \text{otherwise.} \end{cases}$$

It is easily checked that $pd_k(+\delta_x, 1) - pd_k = 1$ iff x lies in at least one set in PD_k and $pd_k(-\delta_x, 1) - pd_k < 0$ iff x is in every set in PD_k .

More generally, for any function $f: X \rightarrow \mathbb{R}$ we can also determine easily whether or not any given species x in X is in every (or in any) set of size k of maximal $PD_{f,a}$ score by comparing $pd_k(f, a)$ and $pd_k(f \pm \delta_x, a)$.

Other applications are also possible—for example, if f takes negative values, and if we wish to find a largest subset W of X for which $PD_{f,a}(W)$ is positive, then one can simply use the greedy algorithm repeatedly to find the largest value k for which $pd_k(f, a) > 0$.

REMARKS

An alternative representation of phylogenetic diversity is given as follows:

$$PD(W) = \sum_{x,y \in W} \mu_{\mathcal{T}}(x, y) d_{(\mathcal{T}, \lambda)}(x, y) \tag{5}$$

where $\mu_{\mathcal{T}}(x, y) = \prod_{v \in p(\mathcal{T}|W, x, y)} (d_{\mathcal{T}|W}(v) - 1)^{-1}$ (here $p(\mathcal{T}|W, x, y)$ is the set of nonleaf vertices of $\mathcal{T}|W$ that lie on the path connecting x and y and $d_{\mathcal{T}|W}(v)$ is the number of edges of $\mathcal{T}|W$ incident with v) and $d_{(\mathcal{T}, \lambda)}(x, y)$ is the sum of the edge weights across the path in \mathcal{T} connecting x and y . This follows from a representation of $PD(W)$ when $W = X$ that was described for binary phylogenetic trees by Pauplin (2000), and generalized to arbitrary phylogenetic trees in Semple and Steel (2004). Equation (5) may be useful in ecological studies where a phylogenetic tree is well established, but its branch lengths are not. The representation (5) separates the topological features of the tree (the term $\mu_{\mathcal{T}}(x, y)$) from the metric properties $d_{(\mathcal{T}, \lambda)}$, so the well-established topology of the tree suffices to determine the $\mu_{\mathcal{T}}(x, y)$ values. If reasonable estimates of evolutionary distance between pairs of species are known (but not exactly, otherwise the branch lengths could be accurately recovered), then these can be used as estimates of the $d_{(\mathcal{T}, \lambda)}(x, y)$ values.

ACKNOWLEDGEMENTS

I thank Gary Barker and Dan Faith for helpful discussion on the biological aspects of phylogenetic diversity during a NZ Landcare workshop (March 2004), and to Andreas Dress for some helpful background on the greedy algorithm. I also thank the New Zealand Marsden Fund (UOC310) for supporting this research.

REFERENCES

Barker, G. M. 2002. Phylogenetic diversity: A quantitative framework for measurement of priority and achievement in biodiversity conservation. *Biol. J. Linn. Soc.* 76:165–194.
 Faith, D. P. 1992. Conservation evaluation and phylogenetic diversity. *Biol. Conserv.* 61:1–10.
 Faith, D. P., and P. A. Walker. 1994. DIVERSITY: A software package for sampling phylogenetic and environmental diversity. Reference and user’s guide. v 1.0. CSIRO Division of Wildlife and Ecology, Canberra, Australia.
 Korte, B., L. Lovász, and R. Schrader. 1991. Greedoids, algorithms and combinatorics. Springer, Berlin.
 Mace, G. M., J. L. Gittleman, and A. Purvis. 2003. Preserving the tree of life. *Science* 300:1701.
 Mooers, A. O., S. B. Heard, and E. Chrostowski. 2005. Evolutionary heritage as a metric for conservation. Pages 120–138 in *Phylogeny and conservation* (A. Purvis, T. L. Brooks, and J. L. Gittleman, eds.). Cambridge University Press, Cambridge.
 Nee, S., and R. M. May. 1997. Extinction and the loss of evolutionary history. *Science* 278:692–694.
 Pachter, L., and D. Speyer. 2004. Reconstructing trees from subtree weights. *Appl. Math. Lett.* 17:615–621.
 Pauplin, Y. 2000. Direct calculation of a tree length using a distance matrix. *J. Mol. Evol.* 51:41–47.
 Semple, C., and M. Steel. 2003. *Phylogenetics*. Oxford University Press, Oxford.
 Semple, C., and M. Steel. 2004. Cyclic permutations and evolutionary trees. *Adv. Appl. Math.* 32:669–680.

First submitted 10 June 2004; reviews returned 2 October 2004;
 final acceptance 2 October 2004
 Associate Editor: Dan Faith