

SUPERTREE ALGORITHMS FOR ANCESTRAL
DIVERGENCE DATES AND NESTED TAXA

**Charles Semple, Philip Daniel, Wim Hordijk,
Roderic D M Page, and Mike Steel**

*Department of Mathematics and Statistics
University of Canterbury
Private Bag 4800
Christchurch, New Zealand*

Report Number: UCDMS2003/18

OCTOBER 2003

SUPERTREE ALGORITHMS FOR ANCESTRAL DIVERGENCE DATES AND NESTED TAXA

CHARLES SEMPLE, PHILIP DANIEL, WIM HORDIJK, RODERIC D. M. PAGE, AND
MIKE STEEL

ABSTRACT. Most supertree methods use just leaf-labelled phylogenetic trees to infer the resulting supertree. In this paper, we describe several new supertree algorithms that extend the allowable information that can be used for phylogenetic inference. These algorithms have been recently implemented and we describe here two illustrative applications.

1. INTRODUCTION

A *supertree* is a rooted phylogenetic tree that is the result of combining a collection of smaller rooted phylogenetic trees on overlapping subsets of species. There are now many techniques ('supertree methods') for constructing supertrees. However, for almost all of these methods, the input is restricted to leaf-labelled phylogenetic trees (without branch lengths or interior node labels or dates), in which case any other related information is ignored. In this paper, we describe several new, and recently implemented, supertree methods which extend the typical input to include some of this additional information.

The new supertree methods divide into two groups depending upon the type of additional information being used as input. The first group allows the input to include ancestral divergence dates which may be either relative or explicit. For example, in this group the input could include information such as whether one particular divergence event on one side of a tree occurred before or after a divergence event on the other side of the tree, or actual time estimates of certain divergence events. The second group of supertree algorithms takes as its input rooted trees in which some of the interior vertices as well as all of their leaves are labelled. This allows the inclusion of nested taxa in the input.

All of the new supertree algorithms described in this paper have been implemented in Java (and are thus platform independent) and are available for applications at

<http://darwin.zoology.gla.ac.uk/~rpage/supertree/>

Date: 15 October 2003.

The second author was supported by the New Zealand Institute of Mathematics and its Applications funded programme *Phylogenetic Genomics* and the first and last authors were supported by the New Zealand Marsden Fund.

The implementation requires the input trees to be given in the Newick format¹. Any tree that is returned by these algorithms is also in the Newick format. All of the algorithms run quickly (that is, require just polynomial time) in the size of the input.

The paper is organised as follows. Each of the new algorithms can be viewed as a variation of BUILD, one of the oldest supertree algorithms. Indeed, the general approach used in each of the algorithms is similar to that used by BUILD. This approach is outlined in the next section. In Sections 3 and 4, we describe the new algorithms for ancestral divergence dates and nested taxa, respectively, and include two applications of these algorithms to some data sets.

We end this section by noting that the formal details of the algorithms described in this paper, including their correctness, are not included here as these details are to appear as chapters [3, 4] of a soon-to-be published book on supertrees [2].

2. THE BUILD APPROACH

Originally designed for other purposes, BUILD [1] is an exact algorithm in that it outputs a tree precisely if the input collection satisfies a particular compatibility criteria. A rooted phylogenetic tree \mathcal{T} *displays* a rooted phylogenetic tree \mathcal{T}' if the label set X' of \mathcal{T}' is a subset of the label set of \mathcal{T} and, up to suppressing degree-two vertices, \mathcal{T}' is a refinement of the minimal rooted subtree of \mathcal{T} that connects the labels in X' . For the purposes of this paper, ‘up to suppressing degree-two vertices’ essentially means ‘overlooking degree-two vertices’. A collection \mathcal{P} of rooted phylogenetic trees is *compatible* if there exists a rooted phylogenetic tree that displays each of the trees in \mathcal{P} . Intuitively, \mathcal{P} is compatible if there is a rooted phylogenetic tree that, up to polytomies, preserves all of the ancestral relationships described by the trees in \mathcal{P} . In particular, if the most recent common ancestor of a and b is a descendant of the most recent common ancestor of a and c in a tree in \mathcal{P} , then this relationship is also preserved in \mathcal{T} .

The algorithms we present in this paper are variations of BUILD. Each algorithm is exact and outputs a tree precisely if the input collection satisfies some particular compatibility criteria. Furthermore, like BUILD, the descriptions of the algorithms take the following form. Each algorithm attempts to construct a tree \mathcal{T} that satisfies the compatibility criteria by constructing the set of clusters of \mathcal{T} . In all cases, this is done by starting with the cluster that is the union of the labels of the trees in \mathcal{P} and successively breaking it down into disjoint subclusters. How the clusters are broken down is determined by an associated graph at each iteration. For each algorithm, this graph as well as the process for breaking up clusters is different. The process continues provided the graph at each iteration satisfies some condition. Eventually, either

- (i) the algorithm outputs a tree that satisfies the compatibility criteria, or
- (ii) outputs a statement indicating that the input collection does not satisfy this criteria.

¹See for example <http://evolution.genetics.washington.edu/phyip/newicktree.html>

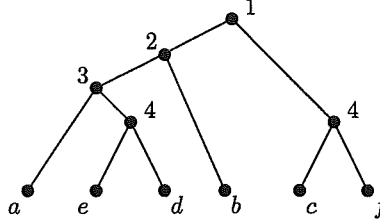


FIGURE 1. A ranked phylogenetic tree.

3. SUPERTREE ALGORITHMS FOR ANCESTRAL DIVERGENCE DATES

We first describe a supertree algorithm that incorporates relative divergence times. This algorithm is called **RANKEDTREE**. An extension of this algorithm to include absolute divergence times or intervals on these times is also possible and this is mentioned at the end of this section.

Essentially, **RANKEDTREE** is an extension of **BUILD** and its input consists of rooted phylogenetic trees as well as information detailing the order in which the divergence events of certain different pairs of species occurred. We call the latter type of input a *relative divergence date* and such information is based, for example, on fossil data or molecular dating techniques. Formally, this type of input takes the form ' $\text{div}(c, d)$ predates $\text{div}(a, b)$ ' which is interpreted as, 'for species a, b, c , and d , the divergence of species c and d predates the divergence of species a and b '.

To include both types of input on a single supertree, we extend the concept of a rooted phylogenetic tree. A *ranked phylogenetic tree* T is a rooted phylogenetic tree in which the interior vertices are assigned a positive integer so that if v_1, v_2 are interior vertices and v_2 is a descendant of v_1 , then the integer assigned to v_1 is less than the integer assigned to v_2 . Such an assignment of positive integers is called a *ranking* of the interior vertices of T . An example of a ranked phylogenetic tree is shown in Fig. 1. Ranking the interior vertices of T in this way corresponds to an ordering of the speciation events associated to these vertices. Note that two different interior vertices may be assigned the same positive integer, in which case, it is inferred that there is no particular ordering on the associated speciation events.

A relative divergence date ' $\text{div}(c, d)$ predates $\text{div}(a, b)$ ' is *preserved* by a ranked phylogenetic tree T if a, b, c, d are leaf labels of T , and the rank assigned to the interior vertex of T corresponding to the most recent common ancestor of c and d is less than the rank assigned to the interior vertex of T corresponding to the most recent common ancestor of a and b . Thus, for example, the ranked phylogenetic tree shown in Fig. 1 preserves the relative divergence date ' $\text{div}(e, b)$ predates $\text{div}(c, f)$ '. A collection \mathcal{P} of rooted phylogenetic trees and a collection \mathcal{D} of relative divergence dates are *compatible* if there is a ranked phylogenetic tree T such that the discrete topology of T displays each of the trees in \mathcal{P} and the ranking of the interior vertices of T preserves all of the relative divergence dates in \mathcal{D} .

The algorithm RANKEDTREE decides whether or not collections of rooted phylogenetic trees and relative divergence dates are compatible. Furthermore, if these collections are compatible, then RANKEDTREE returns a ranked phylogenetic tree that displays each of the rooted phylogenetic trees and preserves each of the relative divergence dates. To illustrate RANKEDTREE, consider the two ranked phylogenetic trees shown in Fig. 2(a) [5, Figure 6] and (b) [9, Figure 3C], each of which is a phylogenetic tree of the cat family. The species labels are the 3-letter abbreviations used in these references. The branch lengths of the source trees have been translated into rankings and added to the interior vertices of these trees. (These branch lengths are also shown.) Observing that species 'LPA', 'PON', and 'CCR' are common to both trees, the ranked phylogenetic tree shown in Fig. 2(c) is the result of applying RANKEDTREE to these two trees. Note that the branch lengths of the resulting ranked phylogenetic tree do not reflect real time.

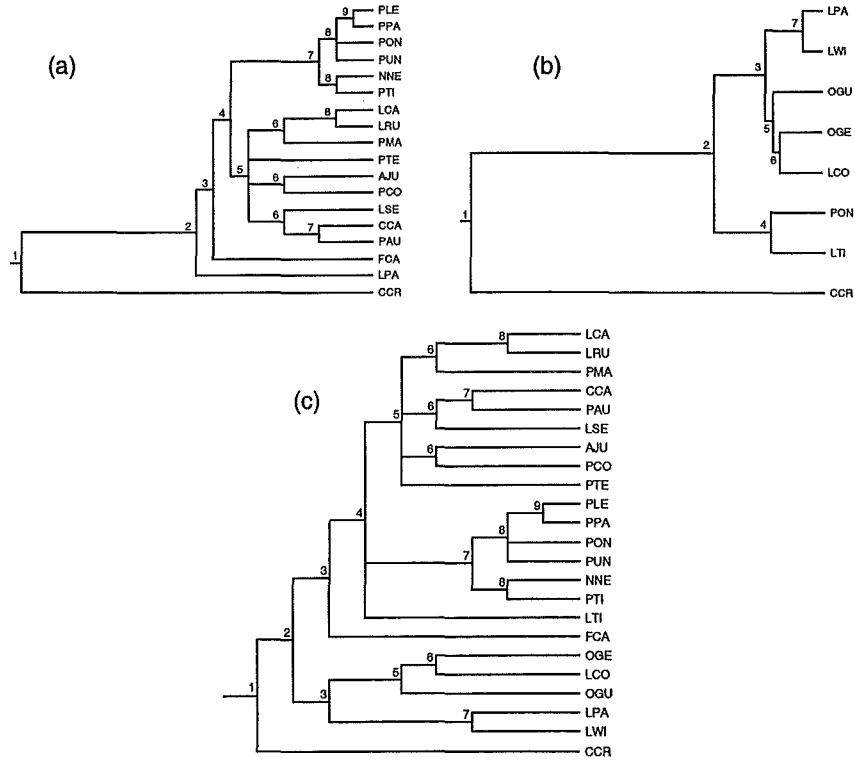


FIGURE 2. An application of RANKEDTREE.

An extension of RANKEDTREE allows for time bounds on speciation events as well as rooted phylogenetic trees and relative divergence dates in its input. A *divergence time bound* for species a and b is either an upper or lower bound (or both) on the number of years ago a and b diverged. To include this information as well as the other information provided by the other inputs, we use a *dated phylogenetic*

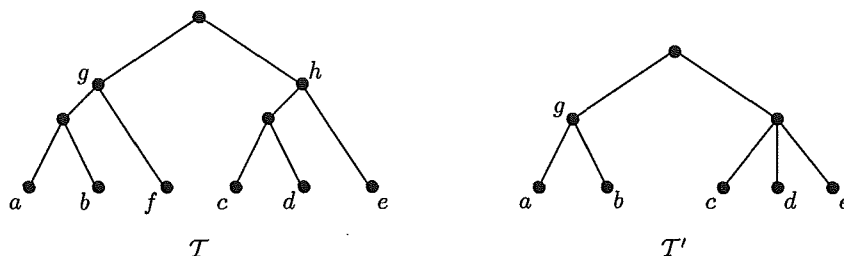


FIGURE 3. Two semi-labelled trees.

tree. Such a tree is similar to that of a ranked phylogenetic tree in that values are assigned to the interior vertices except that these values now represent the number of years ago the corresponding speciation events occurred. Compatibility for this extended input is defined in the obvious way and RANKEDTREE can be modified to solve this compatibility problem also.

4. SUPERTREE ALGORITHMS FOR NESTED TAXA

For supertree algorithms that take as their input collections of rooted phylogenetic trees and only return rooted trees that are leaf-labelled, it is implicit that, as a whole, the leaves of the trees in the input collection represent non-nested taxa. Thus, for example, *Rattus rattus* and ‘mammal’ cannot be represented by two distinct leaves in such a collection as the former is nested inside the latter. This is somewhat limiting in the choice of trees for our input collection. In this section, we describe two supertree algorithms for combining rooted trees in which all of the leaves as well as some of the interior vertices are labelled. These trees are called *rooted semi-labelled trees* and the interior labels of such trees represent taxa at a level higher than that of their descendants. Two semi-labelled trees are shown in Fig. 3.

The two algorithms for combining collections of rooted semi-labelled trees are called SEMI-LABELLEDBUILD and ANCESTRALBUILD. Both algorithms allow a leaf of one of the input trees to represent a taxa that is represented by an interior label of another tree. The motivation for both algorithms came from a problem posed by Page [7].

4.1. SEMI-LABELLEDBUILD. We say that a rooted semi-labelled tree T *perfectly displays* a rooted semi-labelled tree T' if the label set X' of T' is a subset of the label set of T and, up to suppressing degree-two vertices, T' is the minimal rooted subtree of T that connects the labels in X' . Intuitively, T perfectly displays T' if T preserves all of the ancestral relationships described by T' exactly. In particular, T preserves all of the most recent common ancestor relationships described by T' . A collection \mathcal{P} of rooted semi-labelled trees is *perfectly compatible* if there is a rooted semi-labelled tree T that perfectly displays each of the trees in \mathcal{P} .

For a collection \mathcal{P} of rooted semi-labelled trees, SEMI-LABELLEDBUILD decides whether or not \mathcal{P} is perfectly compatible. Moreover, if \mathcal{P} is perfectly compatible, then SEMI-LABELLEDBUILD returns a rooted semi-labelled tree that perfectly displays each of the trees in \mathcal{P} . Figure 4 shows an application of SEMI-LABELLEDBUILD. The input consists of the two rooted semi-labelled trees shown in Fig. 4(a) and (b). Both input trees describe the evolution of spiders and were obtained from study S1x6x97c14c42c30 in TreeBASE. There are taxa common to both trees and it is of particular interest to note that the taxon *Araneoclada* labels a leaf in the tree in (a), but an interior vertex in the tree in (b). The rooted semi-labelled tree resulting from applying SEMI-LABELLEDBUILD to the two input trees is shown in Fig. 4(c).

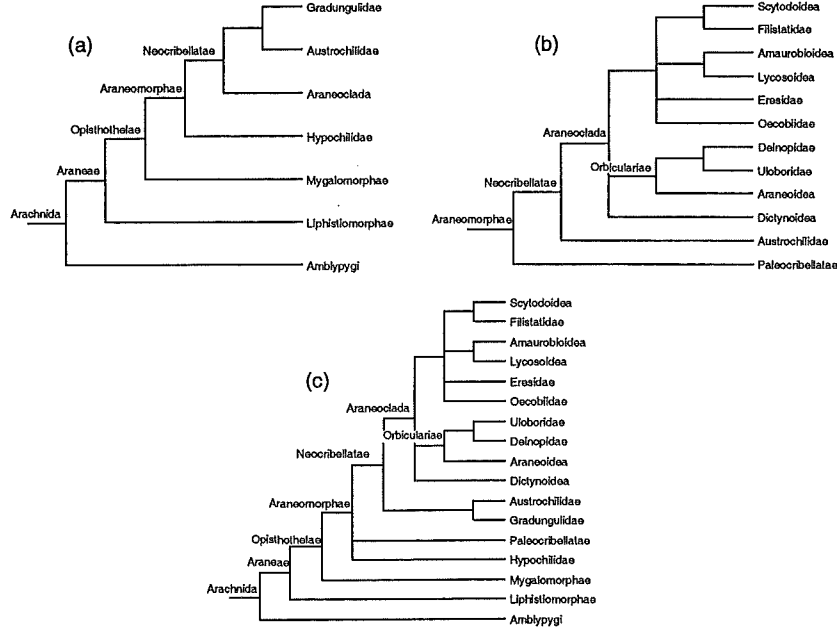


FIGURE 4. An application of SEMI-LABELLEDBUILD.

4.2. ANCESTRALBUILD. The criteria of perfectly displays is very strong as a collection \mathcal{P} of rooted semi-labelled trees is perfectly compatible precisely if there is a rooted semi-labelled tree \mathcal{T} that preserves all of the most recent common ancestor relationships described by the collection. Thus SEMI-LABELLEDBUILD does not allow for the resolution of polytomies. The compatibility criteria and the associated algorithm we describe next relaxes this criteria and allows for the resolution of polytomies.

A rooted semi-labelled tree \mathcal{T} *ancestrally displays* a rooted semi-labelled tree \mathcal{T}' if the following properties hold:

- (i) the label set X' of \mathcal{T}' is a subset of the label set of \mathcal{T} ;

- (ii) up to suppressing degree-two vertices, T' is a refinement of the minimal rooted subtree of T that connects the labels in X' ; and
- (iii) for all labels in X' ,
 - (a) if a is a proper ancestor of b in T' , then a is a proper ancestor of b in T , and
 - (b) if a is neither an ancestor nor a descendant of b in T' , then a is neither an ancestor nor a descendant of b in T .

To illustrate ancestrally displays and compare it with perfectly displays, in Fig. 3, T ancestrally displays T' , but T does not perfectly display T' .

One can think of ancestrally displays as preserving all of the ancestor-descendant relationships. However, observe that it does not preserve the most recent common ancestor relationships. For example, if the most recent common ancestor of a and b is c in T' , then, although c is an ancestor of both a and b in T , and neither a nor b is an ancestor of each other in T , c need not be the most recent common ancestor of a and b in T . A collection \mathcal{P} of rooted semi-labelled trees is *ancestrally compatible* if there is a rooted semi-labelled tree T that ancestrally displays each of the trees in \mathcal{P} . The algorithm ANCESTRALBUILD determines the ancestral compatibility of a collection of rooted semi-labelled trees, in which case, it outputs a rooted semi-labelled tree that ancestrally displays each of the trees in this collection.

5. CONCLUSION

Supertree methods have attracted much interest recently, particularly in the light of well-funded ‘Tree of Life’ initiatives, and studies that have combined large numbers of trees to construct phylogenies on hundreds, or even thousands of species. This has led to some vigorous argument both for and against the use of supertree (versus supermatrix) approaches for phylogeny reconstruction, as well as the emergence of some new techniques as alternatives to the standard MRP (matrix recoding with parsimony) approach. In this paper we have describes some further new methods, that allow for additional types of information to be incorporated—information that would be difficult to include directly using a traditional MRP analysis.

It is important to note that all of the algorithms described in this paper are ‘all-or-nothing’ algorithms. Each algorithm either returns a supertree with certain desirable properties relative to the input or returns a statement indicating that there is no such supertree. In practice, this limits their use. However, such algorithms are important first steps in developing supertree algorithms that always return a supertree and whose input includes information that goes beyond leaf-labelled phylogenetic trees. Indeed, for each of the algorithms described in this paper, we believe that there is a MINCUTSUPERTREE type algorithm (see [6] and [8]) that overcomes this limitation.

REFERENCES

- [1] Aho, A. V., Sagiv, Y., Szymanski, T. G., and Ullman, J. D. (1981). Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM Journal on Computing*, **10**, 405-421.
- [2] Bininda-Emonds, O. R. P., ed. *Phylogenetic supertrees: combining information to reveal the Tree of Life*, in preparation.
- [3] Bryant, D., Semple, C., and Steel, M. Combining evolutionary trees with ancestral divergence dates. In O. Bininda-Emonds (ed.), *Phylogenetic supertrees: combining information to reveal the Tree of Life*, Computational Biology Series, Kluwer, in press.
- [4] Daniel, P. and Semple, C. Supertree algorithms for nested taxa. In O. Bininda-Emonds (ed.), *Phylogenetic supertrees: combining information to reveal the Tree of Life*, Computational Biology Series, Kluwer, in press.
- [5] Janczewski, D. N., Modi, W. S., Stephens, J. C., and O'Brien, S. J. (1995). Molecular Evolution of Mitochondrial 12S RNA and Cytochrome b Sequences in the Pantherine Lineage of Felidae. *Molecular Biology and Evolution*, **12**, 690-707.
- [6] Page, R. D. M. (2002). Modified mincut supertrees. In R. Guigó and D. Gusfield (eds.), *Proceedings of the Second International Workshop on Algorithms in Bioinformatics (WABI 2002)*, pp.537-552, Springer.
- [7] Page, R. D. M. Taxonomy, Supertrees, and the Tree of Life. In O. Bininda-Emonds (ed.), *Phylogenetic Supertrees*, Computational Biology Series, Kluwer, in press.
- [8] Semple, C. and Steel, M. (2000). A supertree method for rooted trees. *Discrete Applied Mathematics*, **105**, 147-158.
- [9] Slattery, J. P., Johnson, W. E., Goldman, D., O'Brien, S. J. (1994). Phylogenetic Reconstruction of South American Felids Defined by Protein Electrophoresis, *Journal of Molecular Evolution*, **39**, 296-305.

BIOMATHEMATICS RESEARCH CENTRE, DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF CANTERBURY, CHRISTCHURCH, NEW ZEALAND

E-mail address: {c.semple, pjd62, w.hordijk, m.steel}@math.canterbury.ac.nz

DEEB, IBLS, GRAHAM KERR BUILDING, UNIVERSITY OF GLASGOW, GLASGOW G12 8QP, UNITED KINGDOM

E-mail address: r.page@bio.gla.ac.uk