



Short Communication

Hide and seek: Placing and finding an optimal tree for thousands of homoplasy-rich sequences

Dietrich Radel^a, Andreas Sand^{b,c}, Mike Steel^{a,*}^aBiomathematics Research Centre, University of Canterbury, Christchurch, New Zealand^bBioinformatics Research Centre, Aarhus University, Denmark^cDepartment of Computer Science, Aarhus University, Denmark

ARTICLE INFO

Article history:

Received 8 June 2013

Revised 30 July 2013

Accepted 1 August 2013

Available online 9 August 2013

Keywords:

Phylogenetic tree

Maximum parsimony

Homoplasy

Tree search

ABSTRACT

Finding optimal evolutionary trees from sequence data is typically an intractable problem, and there is usually no way of knowing how close to optimal the best tree from some search truly is. The problem would seem to be particularly acute when we have many taxa and when that data has high levels of homoplasy, in which the individual characters require many changes to fit on the best tree. However, a recent mathematical result has provided a precise tool to generate a short number of high-homoplasy characters for any given tree, so that this tree is provably the optimal tree under the maximum parsimony criterion. This provides, for the first time, a rigorous way to test tree search algorithms on homoplasy-rich data, where we know in advance what the 'best' tree is. In this short note we consider just one search program (TNT) but show that it is able to locate the globally optimal tree correctly for 32,768 taxa, even though the characters in the dataset require, on average, 1148 state-changes each to fit on this tree, and the number of characters is only 57.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

Phylogenetic tree reconstruction methods based on optimization criteria (such as maximum parsimony or maximum likelihood) have long been known to be computationally intractable (NP-hard) (Foulds and Graham, 1982). However, on perfectly tree-like data (i.e. long sequences with low homoplasy), these methods will generally find the optimal tree quickly, even for large datasets. Moreover, when data is largely tree-like, there are good theoretical and computational methods for finding an optimal tree under methods such as maximum parsimony, with an early result more than 30 years ago (Hendy et al., 1980), along with more recent developments (Blalloch et al., 2006; Holland et al., 2005).

So far, it has not been clear whether such methods would be able to find the global 'optimal' tree for homoplasy-rich datasets with large numbers of taxa, particularly when the sequences are short. The traditional view (Sokal and Sneath, 1963) is that homoplasy tends to obscure tree signal, requiring more character data than homoplasy-free data to recover a tree, though contrary opinions that homoplasy can 'help' have also appeared (Källersjö et al., 1999).

A fundamental obstacle arises in trying to answer this question: One usually cannot guarantee in advance that any tree will be

optimal for homoplasy-rich data without first searching exhaustively through tree space, and this precludes datasets involving hundreds (let alone thousands) of taxa. However, a recent mathematical result by Chai and Housworth (2011) can be used to construct large synthetic datasets on many (thousands) of taxa that simultaneously (i) have a high degree of homoplasy, (ii) come with a guaranteed certificate as to what the optimal tree will be under minimum evolution (maximum parsimony), and (iii) have sequence lengths that are much shorter than the number of taxa.

We can thus, for the first time, test existing programs to see how they perform in such settings, as we know ahead of the analysis what the unique optimal tree is. It might be expected that, with many taxa and high homoplasy, finding this uniquely optimal tree would be impossible. However, we show that this is not the case. In particular, one program (TNT) is able to correctly identify the uniquely most parsimonious tree on thousands of taxa, each requiring many changes. In one case, the uniquely most parsimonious tree for 32,768 taxa was successfully found, even though the characters in this dataset required, on average, 1148 state-changes each to fit on this tree, and the number of characters was only 57. This search involved more than 6.1×10^{13} tree rearrangements, but was completed in 12 hours on a common multi-purpose computer.

Our results provide a positive message for molecular phylogenetics on two fronts: (i) globally optimal trees on thousands of taxa can be recovered from short sequences by existing software in reasonable time, and (ii) high levels of homoplasy, rather than

* Corresponding author. Fax: +64 33642587.

E-mail address: mathmomike@gmail.com (M. Steel).

erasing phylogenetic signal, can enhance it in certain settings, in line with [Källersjö et al. \(1999\)](#). However, we caution that our results rely on synthetic data for which the mathematical structure underlying the high levels of homoplasy may have advantages that might not necessarily extend to biological sequence data subject to random site substitutions. This and other caveats are discussed in the concluding comments section.

1.1. Notation

Throughout, we will let n denote the number of taxa, and k the number of characters. For a sequence $D = (c_1, c_2, c_3, \dots, c_k)$ of characters, let $s(c_i, T)$ denote the parsimony score of c_i on an X -tree T , and let $h(c_i, T)$ be the homoplasy score of c_i on T (see [Semple and Steel, 2003](#) for details). In the case of binary characters (excluding the constant character that assigns all taxa the same state), $h(c_i, T) = s(c_i, T) - 1$.

Let $S(D, T) = \sum_{i=1}^k s(c_i, T)$ denote the parsimony score of D on T and $H(D, T) = \sum_{i=1}^k h(c_i, T)$ denote the homoplasy score of D on T . Thus if all the characters in D are binary (and not constant) then $H(D, T) = S(D, T) - k$. Finally, let $H(D) = \min_T H(D, T)$ and the $S(D) = \min_T S(D, T)$ be the homoplasy score and the parsimony score of the most parsimonious tree, respectively.

1.2. Consistency and retention indices

The consistency and retention indices have traditionally been used to measure the amount of homoplasy in a set of characters D . The **consistency index** CI of a set of characters is defined ([Wiley and Lieberman, 2011](#)) as the ratio M/S where $M = \sum_{i=1}^k \min_T s(c_i, T)$ is the sum of the minimum number of steps for each character and S is the sum of the actual number of steps. A set of characters with a consistency index of 1 exhibits no homoplasy, and the consistency index decreases as the amount of homoplasy increases. In the case of binary characters, $M = k$ and S is the best score, and therefore $CI = k/S(D)$. The consistency index unfortunately grows with the number of taxa, making it hard to compare CI values across datasets. To overcome this, the retention index can be used. The **retention index** (RI) is defined ([Wiley and Lieberman, 2011](#)) as the ratio $(G - S)/(G - M)$ where $G = \sum_{i=1}^k \max_T s(c_i, T)$ is the sum of the maximum number of changes for each character c_i on any tree. For binary characters, G corresponds to the sum over all c_i of the size of the smaller portion of the two-partition of taxa determined by c_i .

2. Short binary sequences that have a uniquely most parsimonious tree are ‘noisy’

It was recently shown ([Huber et al., 2005](#)) that for any binary tree T (with any number of leaves), there is a sequence D of just four multi-state characters for which T is the uniquely most parsimonious tree.




Fig. 1. $T' = T \setminus e$; removing an edge with no change for any character does not change the parsimony score.

Moreover, in that setting one also has $H(D, T) = 0$; that is, the characters exhibit no homoplasy on T . When we move to binary characters, however, the situation is very different, as the next result shows.

Theorem 1. For any sequence D of k binary characters that has a uniquely most parsimonious tree on n leaves, we have:

$$H(D) \geq 2n - 3 - k. \tag{1}$$

Furthermore, when $k = \mathcal{O}(\log n)$, the average homoplasy score per character on T tends to infinity as n grows.

Proof. T is a uniquely most parsimonious tree for D if it is the only X -tree that realizes the minimal parsimony score for D . This implies that collapsing any edge of T leads to a tree that is not most parsimonious for D . In particular, if T is a uniquely most parsimonious tree, then it must be a binary phylogenetic tree and so have exactly $2n - 3$ edges.

Let T be the uniquely most parsimonious tree for $D = \{c_1, c_2, \dots, c_k\}$, and fix a most parsimonious reconstruction \bar{c}_i of each character c_i on T (thus \bar{c}_i is an assignment of states to the vertices of T that extends the leaf assignment c_i). Suppose that T has an edge $e = \{u, v\}$ for which the $\bar{c}_i(u) = \bar{c}_i(v)$ for all $i \in \{1, \dots, k\}$. Then for the tree $T' = T \setminus e$ obtained from T by collapsing the edge $e = \{u, v\}$ (a simple example, where $k = 2$ is shown in [Fig. 1](#)), we have: $S(D, T') = S(D, T)$, which is a contradiction, since T is assumed to be the only most parsimonious tree for D . So, for every edge $e = \{u, v\}$ of T , we must have $\bar{c}_i(u) \neq \bar{c}_i(v)$ for at least one $i \in \{1, \dots, k\}$. For an edge $e = \{u, v\}$ of T and a character c_i of D , let:

$$I(e, c_i) = \begin{cases} 1, & \text{if } \bar{c}_i(u) \neq \bar{c}_i(v); \\ 0, & \text{otherwise.} \end{cases}$$

Then, from our argument above, $\sum_{i=1}^k I(e, c_i) \geq 1$ for each edge e in the edge set E of T , and $s(c_i, T) = \sum_{e \in E} I(e, c_i)$ for every c_i in D . Thus we have:

$$S(D, T) = \sum_{i=1}^k s(c_i, T) = \sum_{i=1}^k \sum_{e \in E} I(e, c_i) = \sum_{e \in E} \sum_{i=1}^k I(e, c_i) \geq 2n - 3,$$

as T has $|E| = 2n - 3$ edges. Hence, since every c_i is binary and T is the most parsimonious tree, we get $H(D) = S(D) - k \geq 2n - 3 - k$.

For the second claim, note that $\frac{H(D)}{k} \geq \frac{2n-3}{k} - 1 \rightarrow \infty$ as $n \rightarrow \infty$, when $k = \mathcal{O}(\log n)$. \square

3. An explicit construction

We applied the construction described in [Chai and Housworth \(2011\)](#), which allows the construction, for each integer value of $p \geq 2$, of a set D of $n = 2^p$ taxa sequences with $k = 4p - 3$ characters

each and with a uniquely most parsimonious tree with the parsimony score $S(D) = 2n - 3$ (see Fig. 2). The consistency index for all these datasets is given by: $CI = \frac{k}{S(D)} = \frac{4p-3}{2^{p+1}-3}$, which converges exponentially fast to 0 when p (and thereby the number of taxa) grows towards infinity. Moreover, for calculating the retention index (RI) we have: $G = \frac{n}{4}(k-1) + \frac{n}{2} = 2^p(p - \frac{1}{2})$. To see this, recall that G is the sum over all characters c_i of the size of the smaller portion of the partition of taxa determined by c_i , and note that in the construction illustrated in Fig. 2(a) the last character of the first block always contains $n/2$ taxa in state 1 and the remaining characters have $n/4$ taxa in state 1. For the homoplasy per character, $h(D)$, we have $h(D) = H(D)/k = \frac{2n-3-k}{k} = \frac{2^{p+1}-3}{4p-3} - 1$, which tends to infinity as p (and thereby the number of taxa) goes to infinity (as stated in Theorem 1). We note that this matches the bound given in Theorem 1, and thus shows that this bound can be realized.




Fig. 2. The balanced tree construction based on Chai and Housworth (2011). (a) The number of taxa can be doubled by repeating each taxon and then adding four new characters at the end of each sequence, repeating the pattern 1000, 0100, 0010, 0001. (b) The tree topology for eight taxa, where a, b, \dots, h are labeled with the 1st, 2nd, \dots , 8th row in column two of (a). The interior nodes can be labeled such that each edge has a change for exactly one character, and the parsimony score therefore equals the number of edges.

4. Results

To test how well TNT (Goloboff et al., 2008) could recover phylogenies with high amounts of homoplasy, we generated datasets with $n = 8, 16, 32, \dots, 32,768$ taxa according to the construction by Chai and Housworth (2011), and ran TNT version 1.1 64 bit (May 2012, Linux 64 version) to recover the phylogenies by maximum parsimony. The search heuristics of TNT involve a number of techniques, including tree rearrangement operations. Furthermore, these heuristics can be guided by the user, based on knowledge about the input by setting multiple command options. But as we were mainly interested in the default performance, we did not use these. However, to speed up the computations, we used the option “`xmult = level x;`” (where x is a number between 0 and 10). The authors of TNT recommend using level 0–2 for easy datasets, 3–5 for medium, and 6–10 for hard datasets. These experiments are summarized in Table 1. Each experiment was performed in a single thread on a dual Intel® Xeon® CPU (3.07 GHz, six cores each) computer running openSUSE v. 11.x.

In Table 1 we first of all note that the most parsimonious tree was successfully reconstructed in all experiments. It was found almost instantly for datasets with up to $n = 128$ taxa (using “`xmult = level 1;`”). For datasets with up to $n = 16,384$ taxa, the most parsimonious trees were found within two hours (using “`xmult = level 3;`”). And even for the dataset with $n = 32,768$ taxa, the most parsimonious tree topology was identified, although this computation took approximately twelve hours (using “`xmult = level 4;`” which is used for medium-difficulty datasets, and here it involved 6.1×10^{13} tree rearrangements; on lower `xmult` settings, which involve less intensive searching, sub-optimal solutions were found).

In another set of experiments (results not shown) we tested TNT on datasets which were built using another construction by Chai and Housworth (2011) such that the uniquely most parsimonious tree for each dataset is a caterpillar tree. Thus these experiments tested TNT’s performance on trees whose shape lies at the opposite pole of tree space. The uniquely most parsimonious trees were also successfully reconstructed across a selection of these earlier analyses, although the reconstruction was significantly more time consuming (e.g. 2 h 32 min for $n = 8192$). This was unexpected, and the reason for it is unclear to us, but it may be a property of TNT’s search heuristic.

5. Concluding comments

The mathematical foundation provided by the Chai–Housworth construction opens the door to a unique experiment that has been impossible until now: searching for the most parsimonious tree in a dataset involving sequences on large numbers of taxa and with

Table 1
Summary of experiments on the balanced tree topology. The uniquely most parsimonious tree was successfully reconstructed in all experiments.

#Taxa n	#Chars k	Exact optimal score (S)	Best score found	Difference	Homoplasy per char (h)	CI	RI	Time
8	9	13	13	0	0.44444	0.69231	0.63636	0:00:00
16	13	29	29	0	1.23077	0.44828	0.62791	0:00:00
32	17	61	61	0	2.58824	0.27869	0.65354	0:00:00
64	21	125	125	0	4.95238	0.16800	0.68580	0:00:00
128	25	253	253	0	9.12000	0.09881	0.71747	0:00:00
256	29	509	509	0	16.55172	0.05697	0.74617	0:00:00
512	33	1021	1021	0	29.93939	0.03232	0.77124	0:00:05
1024	37	2045	2045	0	54.27027	0.01809	0.79280	0:00:21
2048	41	4093	4093	0	98.82927	0.01002	0.81121	0:01:36
4096	45	8189	8189	0	180.97778	0.00550	0.82694	0:04:18
8192	49	16,381	16,381	0	333.30612	0.00299	0.84043	0:19:25
16,384	53	32,765	32,765	0	617.20755	0.00162	0.85207	1:41:26
32,768	57	65,533	65,533	0	1148.70175	0.00087	0.86218	11:59:15

high homoplasy, for which we know in advance what the most parsimonious tree is. It was not at all clear whether existing phylogenetic programs would be able to locate this most parsimonious tree in such a large tree space (when $n = 2^{15}$, the search space contains more than 10^{140000} trees), yet at least one program (TNT) was able to do so. It was not the intention of this short note to compare different parsimony programs on this test dataset, but that would surely be a reasonable project for future work.

Also, the question of whether the homoplasy present in this data is a good proxy for 'noise' in biological data is difficult to determine – the construction by [Chai and Housworth \(2011\)](#) does have an obvious pattern and structure, so it might be argued that finding a tree for such data may be inherently easier than for data for which the homoplasy comes about through random processes. However, there is currently no way to guarantee what the maximum parsimony tree would be for random data, though a conjecture (Conjecture 1.3.1 in [Albert \(2005\)](#)), if established, would provide one. Nevertheless, we find it surprising that one can find a uniquely most parsimonious tree on more than 3×10^4 taxa with just 57 characters that require, on average, more than 1000 substitutions to fit on the best tree.

An interesting question for further theoretical work would be to determine whether or not the Chai–Housworth data allows non-optimal maximum parsimony trees that have locally optimal parsimony scores under tree re-arrangement operations. If local optimal trees do not exist, it would provide a basis to better understand the results reported here.

A further task that would be worthy of study would be to investigate the influence of different tree shapes, beyond the symmetric branching trees and the caterpillar trees considered here, using the general construction given by [Chai and Housworth \(2011\)](#). This is, however, again beyond the scope of this short note.

Acknowledgments

We thank the *Allan Wilson Centre for Molecular Ecology and Evolution* for helping fund this research. We acknowledge use of the program TNT, which is made available with the sponsorship of the Willi Hennig Society. We also thank the two anonymous reviewers for their helpful comments.

References

- Albert, V.A. (Ed.), 2005. *Parsimony, Phylogeny, and Genomics*. Oxford University Press, Oxford.
- Blelloch, G.E., Dhamdhere, K., Halperin, E., Ravi, R., Schwartz, R., Sridhar, S., 2006. Fixed parameter tractability of binary near-perfect phylogenetic tree reconstruction. In: *Automata, Languages and Programming*. Springer, pp. 667–678.
- Chai, J., Housworth, E.A., 2011. On the number of binary characters needed to recover a phylogeny using maximum parsimony. *Bulletin of Mathematical Biology* 73 (6), 1398–1411.
- Foulds, L.R., Graham, R.L., 1982. The Steiner problem in phylogeny is NP-complete. *Advances in Applied Mathematics* 3 (1), 43–49.
- Goloboff, P.A., Farris, J.S., Nixon, K.C., 2008. TNT, a free program for phylogenetic analysis. *Cladistics: The International Journal of the Willi Hennig Society* 24 (5), 774–786.
- Hendy, M., Foulds, L., Penny, D., 1980. Proving phylogenetic trees minimal with 1-clustering and set partitioning. *Mathematical Biosciences* 51 (1), 71–88.
- Holland, B., Huber, K., Penny, D., Moulton, V., 2005. The minmax squeeze: Guaranteeing a minimal tree for population data. *Molecular Biology and Evolution* 22 (2), 235–242.
- Huber, K.T., Moulton, V., Steel, M., 2005. Four characters suffice to convexly define a phylogenetic tree. *SIAM Journal on Discrete Mathematics* 18 (4), 835–843.
- Källersjö, M., Albert, V.A., Farris, J.S., 1999. Homoplasy increases phylogenetic structure. *Cladistics* 15 (1), 91–93.
- Semple, C., Steel, M., 2003. *Phylogenetics*. Oxford Lecture Series in Mathematics and Its Applications, vol. 24. Oxford University Press.
- Sokal, R.R., Sneath, P.H., 1963. *Principles of Numerical Taxonomy*. WH Freeman & Co.
- Wiley, E.O., Lieberman, B.S., 2011. *Phylogenetics: Theory and Practice of Phylogenetic Systematics*. Wiley-Blackwell.