

ACADEMIC
PRESSAvailable at
WWW.MATHEMATICSWEB.ORG
POWERED BY SCIENCE @ DIRECT®**Journal of
Algorithms**

Journal of Algorithms ●●● (●●●●) ●●●-●●●

www.elsevier.com/locate/jalgor

Performance study of phylogenetic methods: (unweighted) quartet methods and neighbor-joining

Katherine St. John,^{a,*} Tandy Warnow,^b Bernard M.E. Moret,^c
and Lisa Vawter^d

^a *Department of Mathematics & Computer Science, Lehman College and the Graduate Center,
City University of New York, USA*

^b *Department of Computer Sciences, University of Texas at Austin, USA*

^c *Department of Computer Science, University of New Mexico, USA*

^d *Aventis Pharmaceuticals, USA*

Received 16 March 2002

Abstract

We present the results of a large-scale experimental study of quartet-based methods (quartet cleaning and puzzling) for phylogeny reconstruction. Our experiments include a broad range of problem sizes and evolutionary rates, and were carefully designed to yield statistically robust results despite the size of the sample space. We measure outcomes in terms of numbers of edges of the true tree correctly inferred by each method (true positives). Our results indicate that these quartet-based methods are much less accurate than the simple and efficient method of neighbor-joining, particularly for data composed of short to medium length sequences. We support our experimental findings by theoretical results that suggest that quartet-cleaning methods are unlikely to yield accurate trees with less than exponentially long sequences. We suggest that a proposed reconstruction method should first be compared to the neighbor-joining method and further studied only if it offers a demonstrable practical advantage.

© 2003 Published by Elsevier Science (USA).

* Corresponding author.

E-mail addresses: stjoh@lehman.cuny.edu (K. St. John), tand@cs.utexas.edu (T. Warnow),
moret@cs.unm.edu (B.M.E. Moret), lisa_vawter@aventis.com (L. Vawter).

1 **1. Introduction** 1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45


Reconstructing the evolutionary history of a group of taxa is a major research thrust in computational biology. An evolutionary history not only gives relationships among taxa, but also an important tool for determining structural, physiological, and biochemical properties [5,30]. Research on tree reconstruction has focused on reconstructing an evolutionary tree (phylogeny) under various optimization criteria. However, almost all optimization problems of interest to biologists are NP-hard (see [11] for a review), so most biologists use heuristic methods or surrogate optimization criteria.

A popular family of phylogenetic heuristics is based on *quartets*. A quartet is an unrooted binary tree for a quadruple of taxa. For most optimization problems, it is possible to determine the optimal tree on a set of four leaves by analyzing all three possible trees. Quartet-based methods compute a quartet under an optimization criterion for each set of four taxa and then combine the quartets to yield a tree on the full set of taxa. Because there are $\Theta(n^4)$ quartets, many quartet-based methods run in $\Omega(n^5)$ time, which is currently impractical for a hundred or more taxa.

How accurate are quartet-based methods? In biological applications, the true, historical tree cannot be ascertained exactly, which makes assessing the quality of reconstruction methods problematic (one exception is laboratory-created phylogenies for viruses and some bacteria, as illustrated in the study by Hillis et al. [10]). As a consequence, the method of choice for evaluating heuristics has been simulation [9]. In such a simulation, an ancestral biomolecular (DNA, RNA, or amino-acid) sequence is evolved along a “model tree,” producing a synthetic set of biomolecular sequences at the leaves. Phylogenetic reconstruction methods are then assessed based upon how accurately they reconstruct the model tree (the “true” tree). Biologists typically evaluate performance according to the topological accuracy of the reconstructed unrooted tree, because the tree topology is interpreted as the order of past evolutionary events, that is, it yields the relationships among species, genes, or other taxa. (The reconstructed tree is typically unrooted, as determining a root is a very difficult problem of its own.) Topological accuracy is typically measured by the true positives, i.e., the percentage of edges of the true tree found in the reconstructed tree.

Among the distance-based methods (methods that transform input sequences into a distance matrix and then construct the tree from that distance matrix), none is more widely used by biologists than the *neighbor-joining* (NJ) method [27]. Not only is it quite fast ($O(n^3)$ for n taxa [29]), but experimental work has also shown that the trees NJ constructs are reasonably accurate, as long as the rate of evolution is neither too low nor too high. However, there is no comparative study of NJ and quartet-based methods.

We present the results of a detailed, large-scale experimental study of quartet-based methods and NJ under the Jukes–Cantor model of evolution [16]. Our results indicate that, under this model, NJ always outperforms the quartet-based methods we examined, in terms of both accuracy and speed. We suggest that NJ, already the most popular distance-based method, should be used as a minimum standard in the assessment of phylogenetic methods: a proposed method should be compared with NJ and shown to be at least comparable in performance to NJ before it is studied in depth. We also present new theory about convergence rates of quartet-based methods which helps explain our observations.

1 **2. Terminology and review** 12
3 **2.1. Simulations and the Jukes–Cantor model** 34
5 A *model tree* for sequence evolution is a pair, $(T, \{\lambda_e\})$, where T is a rooted unlabeled 5
6 tree and, for each edge e of T , λ_e is the expectation of a Poisson distribution describing the 6
7 number of changes at each site in the sequence along edge e . The *Jukes–Cantor model* [16] 7
8 is the simplest Markov model of biomolecular sequence evolution. In that model, a DNA 8
9 sequence (a string over the set of the four nucleotides: $\{A, C, T, G\}$) at the root evolves 9
10 down a rooted binary tree. The sites (i.e., the positions within the sequences) evolve 10
11 independently and identically, with equal probabilities of transition from one nucleotide 11
12 to any other. A Jukes–Cantor model tree is a model tree in which the site changes take 12
13 place according to the Jukes–Cantor model. 1314
15 **2.2. Measures of accuracy** 1516
17 Let T be the true tree (that used in the model tree) and let T' be a tree produced by 17
18 a reconstruction method, with both T and T' leaf-labeled by a set S of taxa. The edges 18
19 of T' are often called the *reconstructed edges* since the method is trying to reconstruct the 19
20 original tree. The true and inferred trees are compared only with respect to their underlying 20
21 unrooted versions, in part because the reconstructed tree is typically unrooted and in part 21
22 because the topological structure of a tree does not depend on the location of its root; 22
23 therefore, in the remainder of this discussion, all trees are assumed to be unrooted. 2324 For each edge $e \in E(T)$, we define the bipartition π_e induced on S by the deletion of 24
25 the edge e from T . The bipartition π_e can also be written as the *split*, $\{A|B\}$, where A 25
26 consists of all the leaves (that is, elements of S) on one side of the bipartition and where 26
27 B is just $S - A$. (These definitions apply to any tree on the set S of leaves.) Methods for 27
28 reconstructing trees are evaluated according to the degree of topological accuracy obtained, 28
29 by comparing the sets of splits or bipartitions of the two trees. The *true positives* are 29
30 the edges $e \in E(T)$ whose split also occurs in the splits of T' . Figure 1 illustrates this 30
31 concept (note that the trees are drawn as rooted, but are compared only with respect to 31
32 their unrooted versions). 3233
34 **2.3. Statistical performance issues** 3435
36 Under the Jukes–Cantor model, a method M is *statistically consistent* if, for every 36
37 model tree $(T, \{\lambda_e\})$ and every $\varepsilon > 0$, there is a sequence length k (which may depend 3744 Fig. 1. True positives: The true tree is on the left. In the tree on the right, the true positives (with respect to the 44
45 tree on the left) are indicated by solid lines. The other edge of the tree is indicated by dotted lines. 45

1 on M , T , $\{\lambda_e\}$, and ε) such that M recovers the unrooted version of T with probability at 1
 2 least $1 - \varepsilon$, when the method is given sequences (generated under the Jukes–Cantor model 2
 3 on $(T\{\lambda_e\})$) of length at least k . 3

4 The sequence length required by a method is a significant aspect of its performance, 4
 5 because real data sets are of limited length (typically bounded by a few hundred to a 5
 6 few thousand nucleotides). Computational requirements are also important, but it may be 6
 7 possible to wait longer or use more powerful machines, whereas it is not possible to get 7
 8 longer sequences than exist in nature. Consequently, experimental and analytical studies 8
 9 have attempted to bound the sequence lengths required by different phylogenetic methods. 9
 10 The rate at which a method converges to 100% accuracy as a function of the sequence 10
 11 length is called the *convergence rate*. 11

12 2.4. Neighbor-joining 12

13 *Neighbor-Joining (NJ)* was formally described in 1987 [27] and has been a mainstay of 13
 14 phylogeny reconstruction among biologists ever since. NJ is a cubic-time distance-based 14
 15 algorithm that begins by creating a node for every taxa in the input set and making a 15
 16 list of those nodes. It proceeds by repeatedly pairing the two “closest” nodes from the 16
 17 list, adding a new node (“the parent”) with edges to the selected pair, and replacing 17
 18 that pair with a new node. As this process progresses, the list of nodes available to pair 18
 19 shrinks and an acyclic graph is formed by the edges added with the new nodes. This 19
 20 process continues until the list of nodes is empty and a tree on all the nodes has been 20
 21 created. As with all distance-based phylogenetic reconstruction methods, the input to NJ 21
 22 is a dissimilarity matrix $\{d_{ij}\}$, where in practice d_{ij} is corrected, according to assumptions 22
 23 about the stochastic model of evolution underlying the data, in order to account for multiple 23
 24 hits. Even with this correction, however, it does not necessarily follow that two leaves i 24
 25 and j are siblings if d_{ij} is the minimum value (even if d is a tree metric). Therefore, NJ 25
 26 computes a secondary dissimilarity matrix $\{s_{ij}\}$, for which it does follow that i and j 26
 27 are siblings if s_{ij} is minimized and d is a tree metric. At each step, then, NJ chooses the 27
 28 pair of leaves with the smallest s_{ij} distance. NJ eventually returns a binary (i.e., fully resolved) 28
 29 tree and is statistically consistent for the Jukes–Cantor model of evolution, and for any 29
 30 model for which statistically consistent distance corrections exist. (See [6, Section 7.3], 30
 31 for a very readable discussion of NJ, and a sketch of Atteson’s proof [1] of its statistical 31
 32 consistency.) 32
 33 33

34 2.5. Quartet-based methods 34

35 A *quartet* is an unrooted binary tree on four taxa. A quartet induces a unique bipartition 35
 36 of the four taxa and can be denoted by that bipartition. If the taxa are $\{a, b, c, d\} \subseteq S$, we 36
 37 37





Fig. 2. The three possible quartets on four taxa $\{a, b, c, d\}$ and their bipartitions.



quartet $\{ab|cd\}$ is correct
 quartet $\{ae|bc\}$ is an error

Fig. 3. For this model tree, the quartet $\{ab|cd\}$ agrees with the tree, while the quartet $\{ae|bc\}$ is an error with respect to the tree.

can use $\{ab|cd\}$ to denote the quartet that pairs a with b and c with d (see Fig. 2). A quartet $\{ab|cd\}$ agrees with a tree T if all four of its taxa are leaves of T and the path from a to b in T does not intersect the path from c to d in T . Equivalently, $\{ab|cd\}$ agrees with a tree if the homeomorphic subtree induced in T by the four taxa is the quartet itself. The quartet $\{ab|cd\}$ is an error with respect to the tree T if it does not agree with T . Figure 3 illustrates this idea. Let $Q(T) = \{\{ab|cd\} \mid a, b, c, d \in S \text{ and } \{ab|cd\} \text{ agrees with } T\}$. If $Q(T)$ denotes the set of all quartets that agree with T , then T is uniquely characterized by $Q(T)$ and can be reconstructed from T in polynomial time [7].

Quartet-based methods operate in two phases. In the first phase, they construct a set Q of quartets on the different sets of four taxa. A popular approach is to use maximum likelihood (ML), a computationally intensive but statistically sophisticated method [8,22]. In the second phase, they combine these quartets into a tree on the entire set of taxa. In practice, not all quartets are accurately inferred, so it is necessary for quartet methods to handle incorrect quartets. Most optimization problems related to tree reconstruction from quartets are NP-hard. An example of this is the *Maximum Quartet Compatibility* problem [15], which seeks a tree T for a given set Q of quartets such that $|Q(T') \cap Q|$ is maximized.

The methods studied in this paper have no performance guarantees with respect to the *Maximum Quartet Compatibility* problem, although each of them is statistically consistent under the Jukes–Cantor model of evolution. However, with the exception of Quartet Puzzling, all quartet methods we examine do provide guarantees about the edges of the true tree that they reconstruct. These guarantees are expressed in terms of *quartet errors around an edge*, a concept we now define.

Consider an edge e in the true tree T ; its removal defines the split $\{A|B\}$ on the elements of S . Consider those sets of four leaves $\{a, a', b, b'\}$ with $\{a, a'\} \subseteq A$ and $\{b, b'\} \subseteq B$. Let t be a quartet on the leaves $\{a, a', b, b'\}$. Note that there are three such possibilities for t : $\{aa'|bb'\}$ (which agrees with the tree T) and two others: $\{ab'|a'b\}$ and $\{ab|a'b'\}$. The quartet t is said to be an *error around edge e* if t is not $\{aa'|bb'\}$ (i.e., if t does not agree with the tree T); see Fig. 4. Similarly, if T' is a proposed tree, and Q is a set of quartets, then $t \in Q$ is an error around edge $e \in E(T')$ if $t = \{ab|a'b'\}$ or $t = \{ab'|a'b\}$, where π_e is the bipartition $\{A|B\}$.

Two of the methods we study, the Q^* method (also known as the Buneman method) and the Quartet Cleaning methods, can be described in terms of an explicit bound on the number of quartet errors around the edges they reconstruct. We begin with the Q^* method [3]. This method seeks the *maximally resolved* tree T' obeying $Q(T') \subseteq Q$. For example, Fig. 5 contains a list of five quartets, Q . The only bipartition compatible with all

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45




Fig. 4. The quartet $\{c|d|e|f\}$ is in error around each of the dotted edges in the tree.

- $\{ab|cd\}$
- $\{ab|ce\}$
- $\{ab|de\}$
- $\{ae|cd\}$
- $\{bc|de\}$




Fig. 5. Maximally resolved tree T' (right) of a given set of quartets (left).

quartets is $\{ab|cde\}$, so that the maximally resolved tree for these quartets, T' , contains only a single nontrivial edge. The set of quartets induced by T' , $Q(T')$, is a proper subset of the input set of quartets, Q :

$$Q(T') = \{\{ab|cd\}, \{ab|ce\}, \{ab|de\}\} \subset Q.$$

Thus, by definition, there are *no quartet errors* around any edge in the tree T' with respect to the input set of quartets Q . This tree always exists, since the star tree¹ trivially satisfies the constraint on any set of quartets. The Q^* tree is unique and can be constructed in polynomial time. By design, however, the Q^* method is conservative and generally produces very unresolved trees [13].

Quartet-Cleaning (QC) methods [2,4,15] have explicit bounds on the number of quartet errors around each reconstructed edge e . As with the Q^* method, we start with an input set of quartets Q . For each edge e in the true tree T , let q_e be the number of quartets that cross e —that is, all quartets $\{aa'|bb'\}$ where $a, a' \in A$ and $b, b' \in B$ and e induces the split $\{A|B\}$. The error bounds have the form $m\sqrt{q_e}$ where m is a small constant—the larger m , the larger the number of quartet errors around an edge that can be handled. The Q^* method can be viewed as a cleaning method in which $m = 0$. The *global cleaning* method sets $m = 1$, and the *local cleaning* method sets $m = 1/2$. These methods are guaranteed to recover every edge of the true tree for which Q contains a small enough number of quartet errors. The *hypercleaning* method allows m to be an arbitrary integer and thus has the potential to recover more edges. However, its running time is very high—proportional to $n^7 \cdot m^{4m+2}$ —so that it is impractical for m larger than 5.

The final quartet-based method we examined is the best known and the most frequently used by biologists [17,20,26]: the *Quartet-Puzzling (QP)* method [28]. This method computes quartets using an ML-based heuristic and then uses a greedy strategy to construct a tree on which many input quartets are in agreement. QP uses an arbitrary ordering of taxa,

¹ The star tree on n leaves has $n + 1$ nodes and n edges and is composed of a central node to which all n leaves are directly connected.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

1 constructs the optimal quartet on the first four, then inserts each successive taxon in turn, 1
 2 attaching the new leaf to an edge of the current tree so as to optimize a quartet-based 2
 3 score. Because the input ordering of taxa is pertinent, QP uses a large number of random 3
 4 input orderings and computes the *majority consensus* of all trees found. (The majority 4
 5 consensus is the tree that contains all bipartitions that appear in more than half of the trees 5
 6 in the set and is a well-known consensus method among biologists.) Thus QP implicitly 6
 7 seeks to return a tree in which every edge is “well-supported,” in the sense that each edge 7
 8 appears in more than half the trees obtained during the algorithm and has (presumably) 8
 9 many supporting quartets. 9

10 2.6. Previous experimental studies of quartet methods 10

11 Berry et al. conducted experimental studies of various QC methods [2,4]. They evolved 11
 12 sequences on Kimura-Two-Parameter (K2P) model trees,² compared the quartets inferred 12
 13 by various methods with the quartets of the true tree, and determined which edges of the 13
 14 model tree could be reconstructed by their QC method. They varied evolutionary rates 14
 15 and sequence lengths, but only examined trees with 10 taxa. Their results showed that 15
 16 QC methods, especially hypercleaning, outperform the Q^* method with respect to true 16
 17 positives. By design, the QC methods cannot fail to recover an edge that is recovered by the 17
 18 Q^* method. So what is noteworthy in the experiments is that the QC methods *did* succeed 18
 19 in obtaining additional edges. Because the dataset sizes used in these experiments are quite 19
 20 small (only 10 taxa), these results may not generalize to larger numbers of taxa. Willson 20
 21 [31] used 12 taxa and the Jukes–Cantor model in conducting simulation studies to assess 21
 22 the accuracy of quartet inference by various methods, including (local) NJ, ML, maximum 22
 23 parsimony, and variants thereof. He found, as we did, that NJ (using distances corrected for 23
 24 the model of evolution) tended to return better quartets than ML under many conditions. 24
 25 Once again, however, the focus on a small fixed number of taxa limits the significance of 25
 26 the results. Finally, no comparison was made between QC methods and NJ or other tree 26
 27 reconstruction methods. 27
 28 29

30 3. Theoretical bounds on the convergence rates 30

31 We begin with the known upper bounds on the convergence rates of NJ and the Q^* 31
 32 method. Surprisingly, these are identical [1,7], although experimental studies strongly 32
 33 suggest that NJ obtains accurate reconstructions of trees from shorter sequences than Q^* 33
 34 throughout the parameter space of Jukes–Cantor trees [13]. We then discuss upper bounds 34
 35 on quartet cleaning. Experimental results illustrating the tightness of the upper bounds can 35
 36 be found in [13,14] and Section 5.7. 36
 37 38
 39

40 **Theorem 3.1.** *Let $f, g, \varepsilon > 0$ be arbitrary constants with $f < g$. Denote by $Q^*(S)$ the tree 40*
 41 *reconstructed on S by the Q^* method and by $NJ(S)$ the tree reconstructed by NJ. There 41*
 42 *are constants $c_1, c_2 > 0$ such that*
 43 *$\Pr[Q^*(S) \neq NJ(S)] \leq c_1 \varepsilon^{c_2}$ and $\Pr[NJ(S) \neq Q^*(S)] \leq c_1 \varepsilon^{c_2}$ for all S with $|S| \geq c_2$.*
 44

45 ² The K2P model [18] is a slight generalization of the Jukes–Cantor model in which the substitution rates 44
 among nucleotides are defined by two parameters, rather than just one. 45

1 is a constant $c > 0$ such that, for all Jukes–Cantor model trees $(T, \{\lambda_e\})$ on n leaves with
 2 $0 < f \leq \lambda_e \leq g < \infty$ for all $e \in E(T)$, and for a set S of sequences generated randomly
 3 on $(T, \{\lambda_e\})$,
 4

$$5 \quad \Pr[Q^*(S) = \text{NJ}(S) = T] > 1 - \varepsilon$$

6
 7 if the sequence length exceeds $c \log n \cdot e^{O(g \cdot \text{diam}(T))}$, where $\text{diam}(T)$ is the largest number
 8 of edges among all paths in T . (Note that c depends on g and ε .)
 9

10
 11 Since the diameter of an n -leaf tree can be as much as $n - 1$ (and has expected value
 12 in $\Omega(\sqrt{n})$ for random trees [7]), Theorem 3.1 shows that the Q^* and NJ methods will
 13 converge from sequences that grow exponentially in n . While Theorem 3.1 provides only
 14 an upper bound, earlier experimental work shows that the Q^* method performs quite
 15 poorly when g and $\text{diam}(T)$ are both large [14], and that NJ is also affected, although
 16 less severely [13].

17 We now consider the convergence rates of the QC methods. The error bound used in
 18 QC methods is a multiple of $\sqrt{q_e}$, so that the ratio of permitted errors to the number of
 19 quartets around an edge is $m/\sqrt{q_e}$, where m is a constant depending on the choice of QC
 20 method. (Recall that we have $m = 1$ for the global cleaning method and $m = 1/2$ for the
 21 local cleaning method.) Because q_e is $\Omega(n^2)$ and m a small constant, this ratio rapidly
 22 approaches 0 as the number of taxa increases. For example, consider an edge in a 50-taxon
 23 tree producing a 20:30 split. The number of quartets around this edge is 82,650, so that
 24 the bound for local cleaning is only 144; hypercleaning with $m = 5$ brings this bound up
 25 to 1440. Thus, for 50 taxa, even hypercleaning has an error tolerance on some edges that
 26 is less than 2% of the total number of quartets for this edge.
 27

28 The sensitivity of QC methods to errors suggests that, for large n , QC methods will be
 29 close in performance to the Q^* method. As soon as the number of errors around each edge
 30 exceeds the cleaning threshold, a QC method behaves identically to the Q^* method. As
 31 noted earlier, the convergence rate of the Q^* method is bounded from above by a function
 32 that grows exponentially in n , suggesting that the Q^* method might be impractical. If
 33 cleaning methods tend to perform only as well as the Q^* method for large n , then they will
 34 not scale well. In Section 5.7, we present experimental results that support this observation.
 35

36 Consider therefore a hypothetical cleaning method we will call *HypoClean*. This
 37 method is guaranteed to recover an edge e if the number of quartet errors around e is
 38 at most one third of the quartets around the edge—a much more generous bound than that
 39 used in local cleaning or than that used by its authors in hypercleaning. In the following
 40 theorem, we establish a bound on the sequence length that suffices for *HypoClean* to be
 41 accurate on a random Jukes–Cantor tree.

42 We require the following lemma.
 43

44 **Lemma 3.1.** *The median diameter of all $(2n - 5)!!$ unrooted, leaf-labeled, binary trees on*
 45 *n leaves is $\Theta(\sqrt{n})$.*

1 **Proof.** Penny and Steel [23] gave formulas for the distribution of interleaf distances in 1
2 such trees under the assumption that all $(2n - 5)!!$ such trees are equally likely, obtaining 2

$$3 \mu(D) = 2^{2n} / \binom{2n}{n} \quad \text{and} \quad \sigma^2(D) = 4n - 6 - \mu(D) - \mu^2(D). \quad 3$$

4
5 Since any nondegenerate distribution must have its median within $[\mu - \sigma, \mu + \sigma]$, our 5
6 conclusion follows. \square 6

7
8 **Theorem 3.2.** Let $f, g, \varepsilon > 0$ be arbitrary constants with $f < g$ and denote by $\text{HC}(S)$ the 8
9 tree reconstructed on S by the *HypoClean* method. Then there is a constant c such that, for 9
10 *Jukes–Cantor model trees*, $(T, \{\lambda_e\})$ with $0 < f \leq \lambda_e \leq g < \infty$ for all $e \in E(T)$, and for a 10
11 set S of n sequences generated randomly on T , we have $\Pr[\text{HC}(S) = T] > 1 - \varepsilon$ whenever 11
12 the sequence length exceeds $c \log n \cdot e^{O(g \cdot \sqrt{n})}$, where the constant c depends on g and ε . 12
13

14 **Proof.** Theorem 3.1 shows that quartets of low diameter are more easily reconstructed 14
15 from short sequences than are quartets of high diameter. Assume that we can correctly 15
16 reconstruct the “smallest-diameter” half of the quartets with high probability—we simply 16
17 guess the remaining quartets. We will then correctly reconstruct $2/3$ of the quartets with 17
18 high probability. What sequence length is required for this? Solving the smaller half of 18
19 the quartets is no harder than solving the median-diameter quartets. By Theorem 3.1, 19
20 this latter task is achieved with high probability when the sequence length is at least 20
21 $O(c \log n \cdot e^{O(g \cdot \text{md}(T))})$, where $\text{md}(T)$ is the median diameter of T . By Lemma 3.1, this 21
22 quantity is $\Theta(\sqrt{n})$. Therefore, the sequence length that suffices to reconstruct the true tree 22
23 with high probability using the *HypoClean* method is $O(c \log n \cdot e^{O(g \cdot \sqrt{n})})$. \square 23
24

25 We have established the same form of (exponential) upper bound on the sequence-length 25
26 requirements of all cleaning methods. This upper bound suggests that cleaning methods 26
27 may not scale well—if the upper bound is approached by any cleaning method, that method 27
28 will require very long sequences to ensure high accuracy, yet such sequences may simply 28
29 not be available. 29
30

31 4. Experimental design 32

33 4.1. Overview 34

35 We used Jukes–Cantor model trees with varying numbers of taxa and rates of evolution 35
36 to generate a large number of synthetic datasets of varying lengths. For each dataset 36
37 generated, we computed the NJ and QP trees on the entire dataset and two sets of quartets, 37
38 one based upon (local) ML, Q_{ML} , and one based upon (local) NJ, Q_{NJ} . We then applied 38
39 various cleaning methods to each of the sets Q_{ML} and Q_{NJ} . We use the terms “local” and 39
40 “global” to distinguish between the local application of running a method on each set of 40
41 four leaves to determine the quartet topology versus the global application of running the 41
42 method on the complete set of leaves. We compared quartets of Q_{ML} , of Q_{NJ} , and of the 42
43 reconstructed trees, as well as the reconstructed trees themselves, against the true tree for 43
44 accuracy. 44
45

1 4.2. Model trees 1

2 2

3 3

4 4

5 5

6 6

7 7

8 8

9 9

10 10

11 11

12 12

13 13

14 14

15 15

16 16

17 17

18 18

19 19

20 20

21 21

22 22

23 23

24 24

25 25

26 26

27 27

28 28

29 29

30 30

31 31

32 32

33 33

34 34

35 35

36 36

37 37

38 38

39 39

40 40

41 41

42 42

43 43

44 44

45 45

We randomly generated model tree topologies from the uniform distribution on binary leaf-labeled trees. For each edge of each tree topology, we generated a random number (from the uniform distribution) between 1 and 1000 and used that number as the “length” of the edge. We then scaled each such base model tree by a multiplicative factor, ranging from 10^{-7} to 10^{-3} . This process produces Jukes–Cantor trees with λ_e values ranging from a minimum of 10^{-7} to a maximum of 1. We generated random DNA sequences for the root and used the program Seq-Gen [24] to evolve these sequences down the tree under the Jukes–Cantor model of evolution, thus producing sets of sequences at the leaves, our synthetic datasets.

14 4.3. Statistical considerations 14

15 15

16 16

17 17

18 18

19 19

20 20

21 21

22 22

23 23

24 24

25 25

26 26

27 27

28 28

29 29

30 30

31 31

32 32

33 33

34 34

35 35

36 36

37 37

38 38

39 39

40 40

41 41

42 42

43 43

44 44

45 45

Because the number of distinct unrooted, leaf-labeled trees on n leaves is $(2n - 5)!!$ and because our input space is further expanded by the choice of evolutionary rates, it is not possible to take a fair sample of the entire input space. In order to obtain statistically robust results, we followed the advice of McGeoch [19] and Moret [21] and used a number of runs, each composed of a number of trials (a trial is a single comparison), computed the mean outcome for each run, and studied the mean and standard deviation over the runs of these events. This approach is preferable to using the same total number of samples in a single run, because each of the runs is an independent pseudorandom stream. With this method, one can obtain estimates of the mean that are closely clustered around the true value, even if the pseudorandom generator is not perfect.

27 4.4. Parameter space 27

28 28

29 29

30 30

31 31

32 32

33 33

34 34

35 35

36 36

37 37

38 38

39 39

40 40

41 41

42 42

43 43

44 44

45 45

A critical parameter of our study, one that has not been explored in most prior studies, is the number of input taxa. Previous experimental studies have often been limited to a small number of taxa due to computational problems. However, to resolve phylogenetic trees of interest to biologists, algorithms must scale reasonably, both in terms of topological accuracy and running time, to problems of the size that biologists typically study (20–200 taxa), as well as those they would like to address (a few hundred to several thousand taxa).

Because of the dedicated use of two multiprocessor clusters, we were able to run our test suite for 5, 10, 20, and 40 taxa (full quartet-based methods remain impractical, at least in terms of experimental studies, for large numbers of taxa). Our tests included a selection of eight expected evolutionary rates, from 5×10^{-5} to 5×10^{-1} per tree edge. For each evolutionary rate and problem size, we generated a total of 100 topologies, grouped into 10 runs of 10 trials. All tests were conducted for four sequence lengths: 500, 2000, 8000, and 32,000 (we note that sequence lengths above 1000 are considered long and those above 5000 extremely long; thus our study explores longer sequence lengths than are usually encountered in practice). In all, our study used 16,000 datasets and required many months of computation on the two clusters.

1 4.5. Algorithms 1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

We tested four different phylogenetic reconstruction methods: NJ, local quartet-cleaning for quartets based on (local) NJ, local quartet-cleaning for quartets based on (local) ML, and QP. The code for QP is TREE-PUZZLE, available from their authors at www.tree-puzzle.de; we modified it only by removing its interactive interface. All other code is our own. For quartet-cleaning, our accuracy measurements were made by counting the number of quartets that were in error around each edge. If the error was below the necessary threshold for the given method, then the edge was counted as being correctly reconstructed. For QP and NJ, we counted the number of true positives between the true tree and the tree method constructed. We ran all four algorithms sequentially on a single set of sequences for one trial, stored all data that was generated, then proceeded to the next trial, so that each of the algorithms was run on exactly the same data.

15 4.6. Measurements 15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

Our focus in this study is the accuracy of solutions generated by the various tree reconstruction methods. Because most methods are time-consuming, the running time is briefly addressed; our aim was not to fine-tune implementations, but simply to obtain a rough idea of which methods can be run in a reasonable amount of time on a conventional machine for realistic datasets. We compare running times as gathered on our platforms, all of which are 450 MHz Pentium III machines running Linux.

To assess topological accuracy, we measured the number of true positives (edges of the true tree that appear in the reconstructed tree). For cleaning methods, we measured these values before and after cleaning. For each run of 10 trials, we retained only the mean values. Our results are composed of the means for each set of 10 runs.

30 5. Experimental results 30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

Except for runs on 5 taxa, the standard deviations we observed remained consistent at 1–2% of the mean; with 5 taxa, standard deviations were (as expected) larger, reaching 10–15% of the mean. In all of our figures, QCNJ and QCML denote quartet-cleaning of quartets derived by local NJ and by local ML, respectively.

37 5.1. Estimating quartets 37

38

39

40

41

42

43

44

45

The technique used to construct the set Q of quartets provided to quartet-based methods can have a significant impact on the performance of these methods. The phylogenetics community has generally expected that (local) ML would produce more accurate quartets than (local) NJ. We therefore compared local ML and local NJ in terms of the quartet sets, Q_{NJ} and Q_{ML} , that they computed. As a reference point, we also examined how global NJ performed in terms of the trees it induced on each fourtuple of leaves from the global NJ tree. Figure 6 shows the proportion of true positives in each of the sets of quartets.

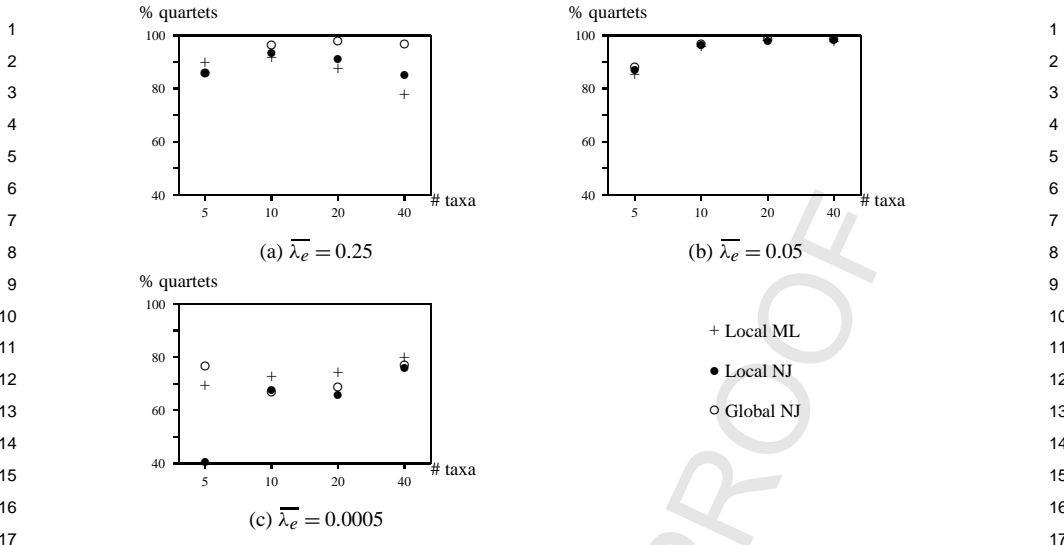


Fig. 6. Percentages of quartets computed by local ML, local NJ and induced by global NJ that agree with the true tree for various numbers of taxa and a sequence length of 500. Here $\bar{\lambda}_e$ refers to the expected number of events on a random edge in the model tree.

The relative performance of local NJ and local ML (NJ and ML applied to each quadruple of leaves to estimate the quartets) is interesting. At the highest rates of evolution (see Fig. 6(a)), except for 5-taxon trees, local NJ slightly outperforms local ML, but this gap increases with increasing numbers of taxa. At the second highest rate of evolution (see Fig. 6(b)), they are indistinguishable up to 40 taxa. However, at the lowest rate of evolution (see Fig. 6(c)), local ML slightly outperforms local NJ, although the gap decreases with increasing numbers of taxa. Overall, while the choice of local NJ vs. local ML does influence the results, our data do not allow us to establish a preference for one over the other: neither ML nor NJ dominates the other in terms of accuracy, but each has a range in which it yields slightly better quartet estimations.


A comparison between these sets of quartets and the quartets obtained by using global NJ (i.e., the quartets induced by the NJ tree) is also interesting. At the lowest rate of evolution (Fig. 6(c)), except for 5-taxon trees, local ML is superior to global NJ and both are superior to local NJ; however, the gap between the three ways of computing quartet trees narrows with increasing number of taxa. At the middle rate (Fig. 6(b)), the methods are indistinguishable (up to 40 taxa), while at the highest rate (Fig. 6(a)), global NJ is clearly superior, and the gap between global NJ as a quartet method and the two other quartet methods increases with increasing numbers of taxa. Thus, for high rates of evolution (and potentially for all large enough trees), the best quartet estimator may simply be global NJ—i.e., compute the NJ tree and use its quartets.

In terms of the quality of the quartets obtained, the best accuracy was obtained at the second highest rate of evolution. At the lowest rate of evolution, only 1 in 2000 sites changes on average around each edge, so that, for a sequence length of 500, roughly 25% of the edges have changes on them. Thus, although it may be possible to guess an edge


1 accurately, the best possible reconstruction at the lowest rate will only yield about 75%
 2 of the edges—approximately what the best performing method (local NJ) obtains. At the
 3 highest rate the accuracy starts to decrease with more than 10 taxa. A decrease in accuracy
 4 with increasing numbers of taxa for a fixed sequence length is predicted by theory (if
 5 only for information-theoretic reasons); hence, even for the lower rates of evolution, as the
 6 number of taxa increases, the accuracy of the quartet estimations should decrease.

7
 8 **5.2. Two measures of accuracy: quartets and edges**
 9

10 Although the standard measure of accuracy is the number of true edges in the
 11 reconstructed tree, the percentage of correctly inferred quartets has also been used as a
 12 surrogate [4]. However, correlation between correct quartets and edges of the true tree
 13 returned by a method has not been shown. We address this issue by examining the
 14 performance of QP and global NJ with respect to both criteria. Figures 7 and 8 make it
 15 clear that edge accuracy is a more demanding criterion than quartet accuracy, and should
 16 therefore be used to assess performance of phylogenetic reconstruction methods. Both
 17 global NJ and QP can return trees with 80% of quartets correct, but only 20% of edges
 18 correct. Worse yet, both methods, except when the percentage of correct quartets is close
 19 to 100%, can return fewer than 80% of the true tree edges (in the case of *QP*, some such
 20



21
 22
 23
 24
 25
 26
 27
 28
 29
 30
 31 Fig. 7. Percent of true tree edges recovered by global NJ for various $\bar{\lambda}_e$ as a function of the percentage of correct
 32 induced quartets for 40 taxa and two sequence lengths.



33
 34
 35
 36
 37
 38
 39
 40
 41
 42
 43
 44 Fig. 8. Percent of true tree edges recovered by QP for various $\bar{\lambda}_e$ as a function of the percentage of correct
 45 induced quartets for 40 taxa and two sequence lengths.

1 trees had only 60% of the true tree edges). Because failure to obtain at least 90 or 95% of
 2 the edges can be unacceptable to systematists, quartet-based measures of accuracy are not
 3 acceptable surrogates for measures of accuracy based on true tree edges.

5.3. Sensitivity to input quality

Methods that operate by estimating quartets and then combining them into a single
 tree can be greatly affected by the quality of the input quartets. Figure 9 shows how QC
 methods respond to input quality. QC methods, as well as the other quartet methods we
 study, require a larger fraction of correct input quartets than the fraction of true tree edges
 that they return.

5.4. Scaling of methods with increasing numbers of taxa

Theory predicts that the accuracy of methods will eventually degrade as the number
 of taxa increases while sequence length and average edge length (the expected number of
 changes for a random site on each edge) are held fixed. Figure 10 shows the edge accuracy
 achieved by all six methods as a function of the number of taxa for a sequence length of 500
 and for three different average edge lengths. Figure 11 shows the same set of results for a
 sequence length of 2000. All methods decrease in accuracy as the number of taxa increases,




Fig. 9. Percentage of correct input NJ quartets vs. true tree edges for Q*NJ and QCNJ for sequence length fixed at 2000, with each graph showing runs for all numbers of taxa and all average edge lengths.




Fig. 10. Percentage of edges correct vs. number of taxa for sequences of length 500 and various $\bar{\lambda}_e$.




Fig. 11. Percentage of edges correct vs. number of taxa for sequences of length 2000 and various $\bar{\lambda}_e$.

even though both NJ and QP show an initial increase. QC provides a distinct improvement over the Q^* method, whether the quartets are computed using local ML or local NJ. QCML and QCNJ are very close in performance, although QCNJ slightly outperforms QCML; similarly Q^*NJ slightly outperforms Q^*ML . Of the five quartet methods, QP is the best throughout the range of parameters studied, but global NJ completely dominates it (and the other quartet methods we study).

5.5. A comparison between Q^* and QC

QC can be seen as an improvement to the Q^* method, because Q^* does not permit errors around any reconstructed edges, while QC reconstructs every edge around which there is a bounded number of errors. In Figs. 10 and 11, we showed performance for different rates of evolution as the number of taxa varies, which gives evidence that QC methods return additional true edges under many conditions. In Fig. 12, we explore the relative improvement in edge recovery obtained on local NJ or local ML quartets by using a QC method instead of the Q^* method. Curiously, the improvement obtained in terms of quartet accuracy is less satisfactory, never averaging more than 2% for low rates of evolution and for large number of taxa at high rates of evolution. QC provides the largest improvement when almost all input quartets are correct; indeed, this is what the theory about QC suggests. In particular, the most improvement occurs at a high rate of evolution—not our highest rate, but our second highest rate, when the error rate in input quartets is also lowest.

5.6. The effects of sequence length

Although sequence length and rate of evolution have a strong effect on the absolute performance of phylogenetic methods, the relative ordering of NJ, QP, and QCNJ is almost constant throughout our experiments: NJ is the best followed by QP, and then by QCNJ. Figure 13 presents data for 40 taxa at three different rates of evolution, for sequence lengths varying from 500 (a typical length) up to 32,000 (a quite large length). Note that all methods increase in accuracy with increasing sequence length (as expected since all methods are statistically consistent under the Jukes–Cantor model).

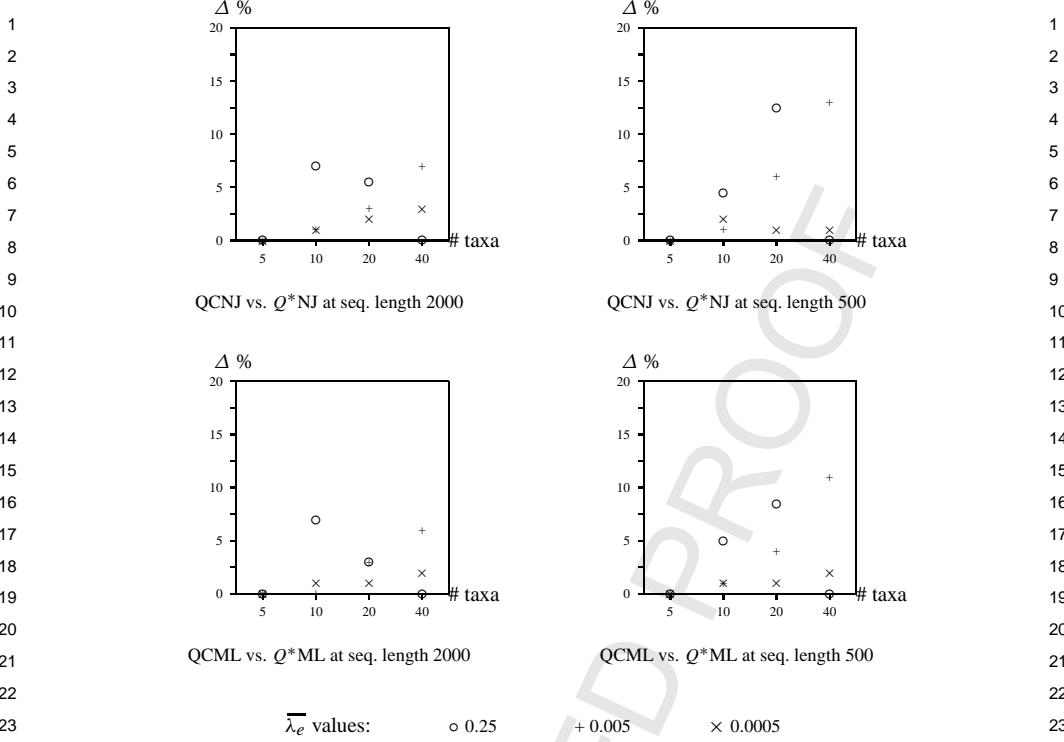


Fig. 12. QC vs. Q*: cleaning-induced improvement for NJ and ML in the percentage of tree edges that agree with the true tree. Δ is the additional percentage of true tree edges obtained by using quartet cleaning.

5.7. Experimental bounds on sequence length

Theorem 3.1 provides only an upper bound on the sequence length sufficient for accurate reconstruction by the Q* method; no theoretical lower bound is known for the necessary sequence length. Using the same experimental set-up as before, we measured the sequence length required to recover accurately all of the edges at least 90% of the time for global and local quartet cleaning (with neighbor-joining determining the topology of the quartets). We generated a tree (under the distributions described above), evolved sequences down the tree of length 500, and used those sets of sequences as inputs to the methods. If a method fails to recover all of the edges of 9 out of 10 of the trees, the sequence length is increased by 500, and we repeat the reconstruction. We stop this process when a method has been successful, or the sequence length reaches 32,000 (a very large length). We note that, at all evolutionary rates, NJ outperforms QCNJ in the sequence lengths needed to reconstruct the edges 90% of the time (see Fig. 14).

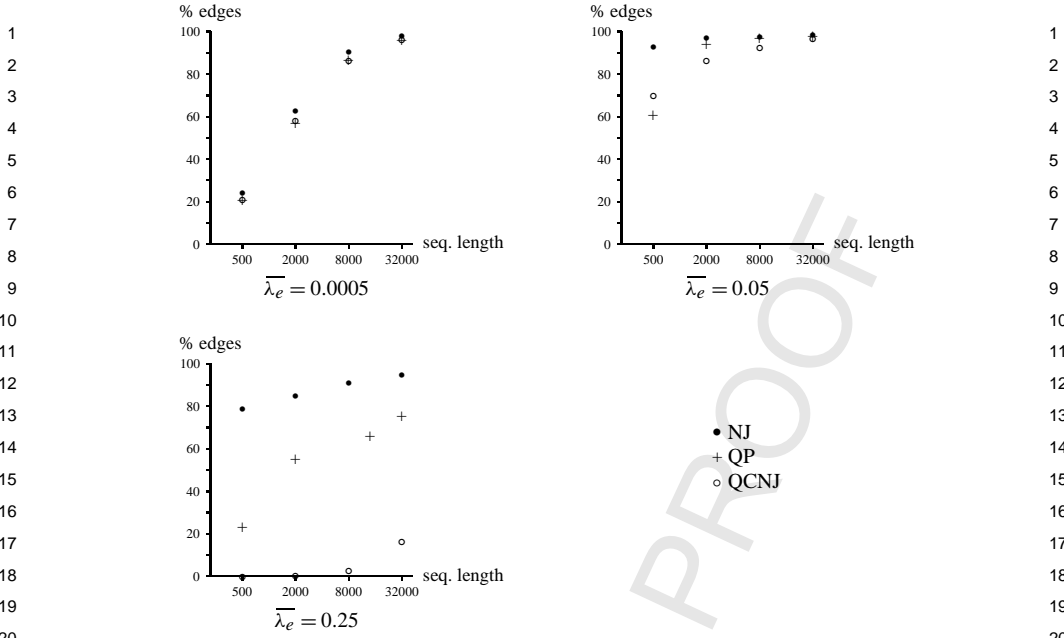


Fig. 13. Accuracy of various methods as a function of sequence length for 40 taxa.




Fig. 14. The sequence length required to accurately reconstruct all of the edges 90% of the time. For low rates of evolution (left), QCNJ and NJ perform comparably. At higher rates of evolution (right), NJ consistently requires shorter sequences to reconstruct the true tree.

6. Discussion

6.1. Quality of quartets

The technique used to construct the set Q of input quartets provided to quartet-based methods can have a significant impact on the performance of these methods. The

1 phylogenetics community has generally expected that local ML would produce more 1
 2 accurate quartets than other local quartet inference methods. However, in our studies, 2
 3 neither local ML nor local NJ dominates the other as a quartet estimator; instead, local 3
 4 ML outperforms local NJ only for the lowest rates of evolution, whereas local NJ clearly 4
 5 outperforms local ML for higher rates. Because our observations differ from the received 5
 6 wisdom in the field, we offer the following possible explanation. In earlier studies [12], 6
 7 the performance of local ML and local NJ as quartet estimators was studied by explicitly 7
 8 simulating evolution on 4-taxon trees. Here, we have simulated evolution on larger trees 8
 9 and then looked at the quartets defined by these larger trees. Good performance on quartets 9
 10 drawn from a large tree is not the same as good performance on quartets drawn from a very 10
 11 different sample space. While it is possible to sample 4-taxon model trees so as to produce 11
 12 the same kind of quartets we gave as input to our methods, the studies in [12] did not use 12
 13 such a sampling strategy. 13

14 6.2. Robustness of quartet methods to quartet errors 15

16
 17 How robust are quartet-based methods with respect to errors in Q ? The Q^* method is 17
 18 the least robust. QC methods provide some error tolerance, sufficient to recover additional 18
 19 true edges even under high rates of evolution and for moderate numbers of taxa. However, 19
 20 both of these methods are inferior to QP in terms of error tolerance, even though QP also 20
 21 fails to get a good estimation of the true tree when the input set of quartets has over 5% of 21
 22 errors (for $n = 40$). Finally, in our experiments, NJ was always at least as accurate as QP 22
 23 and nearly always much better. Thus, the reason quartet methods fail to recover good trees 23
 24 is *not* that the input distance matrix is too noisy for any method to recover a good estimate 24
 25 of the true tree. 25

26 6.3. Running times 27

28
 29 NJ was clearly the fastest method tested. QC and QP methods must compute all $\Theta(n^4)$ 29
 30 quartets and hence must take $\Omega(n^4)$ time. ML-based methods also construct quartets 30
 31 through expensive estimation methods, the running time of which increases sharply with 31
 32 increasing sequence length. Thus QCML and QP were by far the slowest of the methods 32
 33 tested, slow enough that running them on more than a hundred taxa appears infeasible at 33
 34 present. With default settings, QP takes more than 200 days of computation to analyze ten 34
 35 runs of ten trials each for a single set of parameters on 80 taxa with a sequence length 35
 36 of 500. In contrast, NJ dispatches the same analysis in about 30 minutes. 36

37 6.4. Comparison between methods 38

39
 40 Our experiments clearly establish a linear order of accuracy for the methods we studied 40
 41 (except under very low rates of evolution): NJ (applied globally) is the preferred method, 41
 42 with QP second, the QC methods significantly behind QP, and the Q^* methods somewhat 42
 43 behind the QC methods. The particular technique used to infer quartets also has an 43
 44 influence on the quality of the trees obtained by the quartet methods, with QCNJ often 44
 45 better than QCML and Q^* NJ often better than Q^* ML (at least for large enough trees with 45

1 moderate to high rates of evolution). Furthermore, global NJ requires significantly shorter 1
2 sequences to reconstruct the trees than the quartet methods we studied. 2

3 7. Conclusions and open questions 3

4
5 Why does global NJ outperform the quartet methods throughout the parameter space we 5
6 examined (except on some 5-taxon trees)? As Fig. 14 shows, the actual convergence rates 6
7 for both global NJ and QCNJ appear much better than exponential, suggesting that our 7
8 upper bounds on the convergence rate are loose for both cleaning methods and NJ. Yet the 8
9 same figure also shows clearly that the convergence rate of NJ is much better than that of 9
10 QCNJ. The sharp degradation in accuracy that we see in cleaning methods with increasing 10
11 numbers of taxa suggests that their convergence rate, while perhaps subexponential, is 11
12 asymptotically poor. In contrast, global NJ (and, to a lesser extent, QP) degrades far more 12
13 gracefully, and only when the rate of evolution is close to saturation. The good performance 13
14 of QP as a quartet method does not seem to result from its use of ML-based quartets, 14
15 since by that reasoning QCML should demonstrate a comparable improvement over QCNJ 15
16 (which it does not). One reason for the better behavior of QP could be the manner in which 16
17 it combines quartets. We suspect that the issue is partly that the Q^* and QC methods place 17
18 too stringent a requirement on the edges; by comparison the QP method places no absolute 18
19 restriction. Thus, we suspect that the ability of global NJ and QP to handle noisy input 19
20 data lies in the specific techniques each uses to construct trees and the fact that neither 20
21 places strict bounds on errors. This in itself may help explain why QP outperforms the 21
22 other quartet methods we studied, but it does not explain why global NJ outperforms QP. 22
23 We conjecture that methods which operate by combining quartets do not make use of all 23
24 available information: we suggest that quartet-based methods may be impeded by their 24
25 very structure, in having to decide the tree based on quadruples of taxa, without reference 25
26 to the other taxa. 26
27 27
28 28

29 These observations suggest that quartet methods, if they are to be competitive with 29
30 global NJ, need to be flexible in combining quartets into a single tree on the full set 30
31 of taxa. Because of the lack of correlation between quartet accuracy and edge accuracy, 31
32 seeking to solve the quartet compatibility problem may not produce the best trees either. 32
33 Therefore, quartet methods with good performance (reaching or improving upon global 33
34 NJ's performance) will require both more flexibility and greater sophistication than the 34
35 current quartet methods. 35

36 Another experimental study of quartet-based methods [25] compared QP with variants 36
37 of global NJ, ML, and maximum parsimony, on 12-taxon trees. They noted poor per- 37
38 formance by QP with respect to the other methods studied, which they attribute to poor 38
39 weighting of the quartets (pointing out how difficult it is to decide how much weight or 39
40 confidence to give each quartet). This study, along with our results, suggest that some 40
41 flexibility in weighting quartets could improve the accuracy of quartet-based methods. 41
42 We conclude with the following comments about algorithm design and performance 42
43 studies in phylogenetics. From the perspective of experimental performance studies 43
44 and algorithm design, global NJ should be regarded as a universal lowest common 44
45 denominator in phylogeny reconstruction algorithms. Its speed makes it easy to use 45

1 under all circumstances; its topological accuracy makes it an acceptable starting point for 1
2 tree reconstruction in biological practice. We suggest that a proposed method should be 2
3 compared with NJ and abandoned if it does not offer a demonstrable advantage over NJ 3
4 for substantial subproblem families. 4

5 6 7 Acknowledgments 7

8 9 The authors thank the anonymous referees for their useful comments in improving 9
10 the readability of this paper and suggesting improvements that led to the experiments 10
11 in Section 5.7. The authors thank the Department of Computer Science at University of 11
12 New Mexico for hosting KS, TW, and LV in Summer 2000. KS thanks the Department of 12
13 Computer Sciences and the Texas Institute of Computational and Applied Mathematics at 13
14 the University of Texas for hosting her during the 2000–2001 academic year. 14

15 This work was supported in part by National Science Foundation grants ACI 00-81404 15
16 (BM), CDA 95-03064 (BM), CCR 94-57800 (TW), DEB 01-20709 (BM and TW), DEB 16
17 01-211651 (KS), EIA 99-73874 (KS), EIA 01-13095 (BM), EIA 01-13654 (TW), EIA 01- 17
18 21377 (BM), and EIA 01-21680 (TW), and by the David and Lucile Packard Foundation 18
19 (TW). 19

20 21 22 References 22

- 23 24 [1] K. Atteson, The performance of the neighbor-joining method of phylogeny reconstruction, *Algorith-* 24
mica 25 (2–3) (1999) 251–278. 25
- 26 [2] V. Berry, D. Bryant, T. Jiang, P. Kearney, M. Li, T. Wareham, H. Zhang, A practical algorithm for recovering 26
the best supported edges of an evolutionary tree, in: *Proc. 11th Ann. ACM–SIAM Symp. Discrete Algs.* 27
SODA 2000, SIAM Press, 2000, pp. 287–296. 27
- 28 [3] V. Berry, O. Gascuel, Inferring evolutionary trees with strong combinatorial evidence, *Theoret. Comput.* 28
Sci. 240 (2) (2000) 271–298. 29
- 30 [4] V. Berry, T. Jiang, P. Kearney, M. Li, T. Wareham, Quartet cleaning: improved algorithms and simulations, 30
in: *Proc. Europ. Symp. Algs., ESA99*, in: *Lecture Notes in Comput. Sci.*, Vol. 1643, 1999, pp. 313–324. 31
- 32 [5] J.K. Chambers, L.E. Macdonald, H.M. Sarau, R.S. Ames, K. Freeman, J.J. Foley, Y. Zhu, M.M. McLaughlin, 32
P. Murdock, L. McMillan, J. Trill, A. Swift, N. Aiyar, P. Taylor, L. Vawter, S. Nahed, P. Szekeres, G. Hervieu, 33
C. Scott, J.M. Watson, A.J. Murphy, E. Duzic, C. Klein, D.J. Bergsma, S. Wilson, G.P. Livi, A G protein- 33
coupled receptor for Uridine 5'-diphosphoglucose (UDP-glucose), *J. Biol. Chem.* 275 (2000) 10767–10771. 34
- 35 [6] R. Durbin, S. Eddy, A. Krogh, G. Mitchison, *Biological Sequence Analysis: Probabilistic Models of Proteins* 35
and Nucleic Acids, Cambridge Univ. Press, 1988. 36
- 37 [7] P.L. Erdős, M. Steel, L. Székely, T. Warnow, A few logs suffice to build almost all trees—I, *Random* 36
Structures Algorithms 14 (1997) 153–184. 37
- 38 [8] J. Felsenstein, Evolutionary trees from DNA sequences: A maximum likelihood approach, *J. Mol. Evol.* 17 38
(1981) 368–376. 39
- 39 [9] D.M. Hillis, Approaches for assessing phylogenetic accuracy, *Syst. Biol.* 44 (1995) 3–19. 39
- 40 [10] D.M. Hillis, J.J. Bull, M.E. White, M.R. Badgett, I.J. Molineux, Experimental phylogenies: generation of a 40
known phylogeny, *Science* 255 (1992) 589–592. 41
- 42 [11] D.M. Hillis, C. Moritz, B. Mable, *Molecular Systematics*, Sinauer, Boston, 1996. 42
- 43 [12] J. Huelsenbeck, D. Hillis, Success of phylogenetic methods in the four-taxon case, *Syst. Biol.* 42 (3) (1993) 43
247–264. 44
- 44 [13] D. Huson, S. Nettles, K. Rice, T. Warnow, S. Yooseph, The hybrid tree reconstruction method, *ACM J. Exp.* 44
Algorithms 4 (5) (1999), <http://www.jea.acm.org/1999/HusonHybrid/>. 45

- 1 [14] D. Huson, S. Nettles, T. Warnow, Disk-covering, a fast converging method for phylogenetic tree 1
2 reconstruction, *J. Comp. Biol.* 6 (1999) 369–386. 2
- 3 [15] T. Jiang, P.E. Kearney, M. Li, A polynomial-time approximation scheme for inferring evolutionary trees 3
4 from quartet topologies and its application, *SIAM J. Comput.* 30 (6) (2001) 1942–1961. 4
- 5 [16] T.H. Jukes, C. Cantor, *Mammalian Protein Metabolism*, Academic Press, New York, 1969, pp. 21–132. 5
- 6 [17] P.J. Keeling, M.A. Luker, J.D. Palmer, Evidence from beta-tubulin phylogeny that microsporidia evolved 6
7 from within the Fungi, *Mol. Biol. Evol.* 17 (2000) 23–31. 7
- 8 [18] M. Kimura, A simple method for estimating evolutionary rate of base substitutions through comparative 8
9 studies of nucleotide sequences, *J. Mol. Evol.* 16 (1980) 111–120. 9
- 10 [19] C.C. McGeoch, Analyzing algorithms by simulation: variance reduction techniques and simulation 10
11 speedups, *ACM Comput. Surveys* 24 (1992) 195–212. 11
- 12 [20] B. Mishof, C.L. Anderson, H. Hadrys, A phylogeny of the damselfly genus *Calopteryx* (Odonata) using 12
13 mitochondrial 16s rDNA markers, *Molecular Phylog. Evol.* 15 (2000) 5–14. 13
- 14 [21] B.M.E. Moret, Towards a discipline of experimental algorithmics, in: M. Goldwasser, D.S. Johnson, C.C. 14
15 McGeoch (Eds.), *Proc. 5th & 6th DIMACS Implementation Challenges*, American Mathematical Society, 15
16 Providence, 2003, available at <http://www.cs.unm.edu/~moret/dimacs.ps>. 16
- 17 [22] G.J. Olsen, H. Matsuda, R. Hagstrom, R. Overbeek, fastDNAm1: A tool for construction of phylogenetic 17
18 trees of DNA sequences using maximum likelihood, *Comput. Appl. Biosci.* 10 (1994) 41–48. 18
- 19 [23] D. Penny, M.A. Steel, Distributions of tree comparison metrics—some new results, *Syst. Biol.* 42 (1993) 19
20 126–141. 20
- 21 [24] A. Rambaut, N.C. Grassly, Seq-Gen: An application for the Monte Carlo simulation of DNA sequence 21
22 evolution along phylogenetic trees, *Comput. Appl. Biosci.* 13 (1997) 235–238. 22
- 23 [25] V. Ranwez, O. Gascuel, Quartet-based phylogenetic inference: improvements and limits, *Mol. Biol.* 23
24 *Evol.* 18 (6) (2001) 1103–1116. 24
- 25 [26] F. Rodrigues-Trelles, L. Alarcon, A. Fontdevila, Molecular evolution and phylogeny of the *buzzatii* complex 25
26 (*D. repleta* group): a maximum likelihood approach, *Mol. Biol. Evol.* 17 (2000) 1112–1122. 26
- 27 [27] N. Saitou, M. Nei, The neighbor-joining method: A new method for reconstructing phylogenetic trees, *Mol.* 27
28 *Biol. Evol.* 4 (1987) 406–425. 28
- 29 [28] K. Strimmer, A. von Haeseler, Quartet puzzling: a maximum likelihood method for reconstructing tree 29
30 topologies, *Mol. Biol. Evol.* 13 (1996) 964–969. 30
- 31 [29] J.A. Studier, K.J. Keppler, A note on the neighbor-joining method of Saitou and Nei, *Mol. Biol. Evol.* 5 31
32 (1981) 729–731. 32
- 33 [30] P.G. Szekeres, A.I. Muir, L.D. Spinage, J.E. Miller, S.I. Butler, A. Smith, G.I. Rennie, P.R. Murdock, L.R. 33
34 Fitzgerald, H. Wu, L.J. McMillan, S. Guerrero, L. Vawter, N.A. Elshourbagy, J.L. Mooney, D.J. Bergsma, 34
35 S. Wilson, J.K. Chambers, Neuromedin U is a potent agonist at the orphan G protein coupled receptor FM3, 35
36 *J. Biol. Chem.* 275 (2000) 20247–20250. 36
- 37 [31] S.J. Willson, An error-correcting map for quartet can improve the signals for phylogenetic trees, *Mol. Biol.* 37
38 *Evol.* 18 (3) (2001) 344–351. 38
- 39 39
40 40
41 41
42 42
43 43
44 44
45 45