

A ‘Stochastic Safety Radius’ for Distance-Based Tree Reconstruction

Olivier Gascuel¹ · Mike Steel²

Received: 14 November 2014 / Accepted: 23 April 2015
© Springer Science+Business Media New York 2015

Abstract A variety of algorithms have been proposed for reconstructing trees that show the evolutionary relationships between species by comparing differences in genetic data across present-day species. If the leaf-to-leaf distances in a tree can be accurately estimated, then it is possible to reconstruct this tree from these estimated distances, using polynomial-time methods such as the popular ‘Neighbor-Joining’ algorithm. There is a precise combinatorial condition under which distance-based methods are guaranteed to return a correct tree (in full or in part) based on the requirement that the input distances all lie within some ‘safety radius’ of the true distances. Here, we explore a stochastic analogue of this condition, and mathematically establish upper and lower bounds on this ‘stochastic safety radius’ for distance-based tree reconstruction methods. Using simulations, we show how this notion provides a new way to compare the performance of distance-based tree reconstruction methods. This may help explain why Neighbor-Joining performs so well, as its stochastic safety radius appears close to optimal (while its more classical safety radius is the same as many other less accurate methods).

Keywords Tree · Reconstruction · Robustness to random error

✉ Mike Steel
mathmomike@gmail.com

Olivier Gascuel
gascuel@lirmm.fr

¹ Institut de Biologie Computationnelle (IBC) – LIRMM (UMR 5506), CNRS and Université de Montpellier, Montpellier, France

² Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand

1 Introduction

A central task in evolutionary biology is the reconstruction of ‘phylogenetic’ (evolutionary) trees from genetic data sampled from present-day species that describe how these species evolved from a common ancestor. These trees can be estimated from a variety of different types of data, but a common approach involves data that are based on some measure of ‘evolutionary distance’ between species. A variety of fast (polynomial-time) methods have been devised for building a phylogenetic X -tree from an arbitrary distance function d on X . The most popular by far is ‘Neighbor-Joining,’ (NJ) and the paper [25] that described this heuristic algorithm has now been cited more than 39,000 times.

A desirable property of such methods is that when a distance function fits exactly on a binary (fully resolved) tree with branch lengths, the method will return that underlying tree (up to the placement of the root) and its edge lengths. Moreover, when a distance function δ is close to an exact fit on some binary tree T , many methods also come with a guarantee that they will return T when applied to δ .

How close δ needs to be to a ‘tree metric’ d depends crucially on w_{\min} , the smallest interior edge length of T ; a distance-based tree reconstruction method is said to have ‘safety radius’ ρ if the method is guaranteed to return the underlying binary tree T when δ differs from d by less than $\rho \cdot w_{\min}$ on each pair of leaves (a precise definition is given shortly). This notion was introduced by Kevin Attenson 15 years ago in this journal [1], where he established that for NJ, this safety radius is $\rho = \frac{1}{2}$; moreover, this is the largest possible safety radius for any distance-based tree reconstruction method.

While this classical safety radius has provided a precise formal way to compare different tree reconstruction methods, it is not always a good predictor of which method will perform more accurately on simulated (or real) data; for instance, methods that perform well (e.g. NJ) often have a safety radius that is equivalent to those of methods that perform poorly (e.g. the Buneman tree [2]). In this paper, we consider a more relaxed, statistically-based version of the purely combinatorial safety radius, which treats the differences between observed and expected distances as independent random variables. We develop and apply this notion of ‘stochastic safety radius’, derive formal upper and lower bounds, and compare different tree reconstruction methods using it.

1.1 Classical Safety Radius

For a phylogenetic X -tree T with positive edge lengths w , let $d_{(T,w)}$ denote the associated tree metric on X and let w_{\min} denote the minimum interior edge length of T . A method M for reconstructing a phylogenetic X -tree from a dissimilarity map δ on X is said to have a l_{∞} safety radius ρ_n if, for any binary phylogenetic tree T with n leaves we have:

$$\|\delta - d_{(T,w)}\|_{\infty} < \rho_n \cdot w_{\min} \implies M(\delta) = T.$$

Here $\| * \|_{\infty}$ refers to the largest difference between δ and d over all pairs from X . Beginning with the pioneering work of Attenson [1] it is now well known that no method

can have l_∞ safety radius $\rho_n > \frac{1}{2}$ and that certain methods such as NJ and ‘Balanced Minimum Evolution’ (BME) [22] achieve this bound for all n [20]. However, for other methods, the l_∞ safety radius is less than $\frac{1}{2}$ and it can even converge to 0 as n grows [14,20]. We will refer to the following constraint:

$$\|\delta - d_{(T,w)}\|_\infty < \frac{1}{2} \cdot w_{\min}$$

as the *Atteson l_∞ bound*.

Despite the mathematical elegance of these results, there are two problems associated with the l_∞ safety radius approach. Firstly, it is a strict combinatorial condition and the l_∞ metric is extremely sensitive, particularly for large values of n , since it takes only one pair of taxa to have a δ value that conflicts substantially with its d value to result in a violation of the safety radius. A second related point is that simulations show that methods such as NJ often return the correct tree even when the safety radius is violated. One combinatorial approach that goes some way towards addressing this second point was taken in [19] and developed further in [8]. This latter paper also showed that NJ has a l_2 safety radius $\frac{1}{\sqrt{3}} \approx 0.5773$ for trees with $n = 5$ leaves.

A related “edge safety radius” approach was also pioneered by Atteson [1], who showed that ADDTREE [26] is optimal, whereas NJ is not. However, simulations show that NJ performs well (i.e. as well as ADDTREE) regarding the edge safety radius, which somewhat contradicts the theory. Recent results by Bordewich and Mihaescu [3] also indicate that this theory has some shortcomings and produces a ranking among methods (Greedy BME/NJ) which is not observed in practice. There are also examples where the standard safety radius produces strange ranking, e.g. UPGMA/LS methods [14].

2 Stochastic Safety Radius

Let us regard δ as differing from d by a random ‘error’. More precisely, we suppose that:

$$\delta(x, y) = d_{(T,w)}(x, y) + \epsilon_{xy}, \tag{1}$$

where the ϵ_{xy} values are independent normal random variables, each with a mean of 0 and a variance equal to σ^2 . We refer to this simple model as the *random errors model*.

Note that in the context of this random errors model, maximum likelihood estimation (MLE) of a tree is equivalent to the ordinary least squares (OLS) tree-reconstruction method, since the OLS score that this method seeks to minimise is proportional to minus the log of the likelihood function. Several heuristic approaches have been designed to search for the OLS tree, starting with the seminal 1967 papers of Cavalli-Sforza and Edwards [6] and Fitch and Margoliash [9].

Throughout this paper, we let $N(\mu, \sigma^2)$ denote a normal random variable with mean μ and variance σ^2 . Thus ϵ_{xy} has the distribution $N(0, \sigma^2)$. The following inequality and asymptotic equality (as $x \rightarrow \infty$) are helpful in the results that follow (their proof is in the “Appendix”). For $x > 0$:

$$\frac{1}{2}e^{-x^2/2} < \mathbb{P}(N(0, 1) > x) \sim \frac{e^{-x^2/2}}{x\sqrt{2\pi}}, \tag{2}$$

where \sim denotes asymptotical equivalence as x grows.

2.1 Example: The Four-Taxon Case

With four taxa, most methods, if not all (e.g. ADDTREE [26], NJ [25], UNJ [26], and OLS [16] which is equivalent to MLE within our model) will use the “four-point rule” [27] and select the topology $xy|wz$ if:

$$\delta(x, y) + \delta(w, z) < \min\{\delta(x, w) + \delta(y, z), \delta(x, z) + \delta(y, w)\}.$$

Now, under the random errors model, the three sums

$$\delta(x, y) + \delta(w, z), \delta(x, w) + \delta(y, z), \delta(x, z) + \delta(y, w),$$

constitute three independent normal random variables, each with variance of $2\sigma^2$. Moreover, if the tree generating d has topology $xy|wz$ with an interior edge of length w , then the second and third sum have the same mean, which is larger than the mean of the first sum by $2w$. In particular $\mathbb{P}(\delta(x, w) + \delta(y, z) > \delta(x, y) + \delta(w, z))$ and $\mathbb{P}(\delta(x, z) + \delta(y, w) > \delta(x, y) + \delta(w, z))$ are both equal to

$$\mathbb{P}\left(N\left(2w, 4\sigma^2\right) > 0\right).$$

Consequently, the probability that the correct tree topology $xy|wz$ is selected from δ is at least:

$$1 - 2\mathbb{P}\left(N\left(2w, 4\sigma^2\right) < 0\right) = 1 - 2\mathbb{P}(N(0, 1) < -w/\sigma). \tag{3}$$

Thus, there will be (say) a $\sim 98\%$ probability of correctly inferring the tree topology if σ is $\frac{1}{2}w$ (or less).

It is interesting to compare this to the l_∞ bound of Atteson [1] for methods such as NJ. Recall that this holds when $|d(x, y) - \delta(x, y)| < \frac{w}{2}$ for all six pairs x, y . Now, under the random errors model, the probability of these events all holding is exactly:

$$\left(-w/2 < \mathbb{P}\left(N\left(0, \sigma^2\right) < w/2\right)\right)^6 = (1 - 2\mathbb{P}(N(0, 1) > w/2\sigma))^6.$$

Consequently, for σ set equal to $\frac{1}{2}w$, as above, the probability the l_∞ bound holds is $(1 - 2\mathbb{P}(N(0, 1) > 1))^6 \approx 0.1$, which is much lower than the 98% probability described previously; moreover, in order to ensure the l_∞ bound holds with 98% probability, we would need to reduce σ to around $w/6$.

2.2 Larger Values of n

To extend the above analysis from $n = 4$ to larger values of n , it is useful to allow σ^2 to depend on n (the reason for this becomes clear shortly). Specifically, let us write:

$$\sigma^2 = \frac{c^2}{\log(n)}, \tag{4}$$

for some value $c \neq 0$.

Notice that σ is converging to zero but very slowly (i.e. larger trees require more accuracy, but not a lot more). First we consider what happens with the l_∞ bound as n grows.

Proposition 1 *Under the random errors model, the probability that the Atteson l_∞ bound holds for all pairs x, y in a tree with n leaves converges to 0 for $c > \frac{1}{4} \cdot w_{\min}$ and converges to 1 for $c < \frac{1}{4} \cdot w_{\min}$.*

Proof Let w^* ($= w_{\min}$) be the minimal interior edge length in T . Then:

$$\mathbb{P}(|\delta(x, y) - d_{(T,w)}(x, y)| < w^*/2) = \mathbb{P}\left(N\left(0, \sigma^2\right) \in (-w^*/2, w^*/2)\right),$$

where $N(0, \sigma^2)$ refers to a normal random variable with mean 0, and variance σ^2 given by Eq. (4). Thus:

$$\mathbb{P}(|\delta(x, y) - d_{(T,w)}(x, y)| < w^*/2) = 1 - 2\mathbb{P}\left(N(0, 1) > \frac{w^*}{2c} \sqrt{\log(n)}\right).$$

Now, the l_∞ bound is satisfied precisely when $|\delta(x, y) - d_{(T,w)}(x, y)| < w^*/2$ for all x, y , and so the probability of this l_∞ bound occurring—call it P_∞ —is given by:

$$P_\infty = \left(1 - 2\mathbb{P}\left(N(0, 1) > \frac{w^*}{2c} \sqrt{\log(n)}\right)\right)^{\binom{n}{2}}. \tag{5}$$

Let $x_n = \mathbb{P}\left(N(0, 1) > \frac{w^*}{2c} \sqrt{\log(n)}\right)$. Then the asymptotic equivalence in (2) implies that $x_n \sim \frac{1/\sqrt{2\pi}}{n^\beta \sqrt{\log(n)}}$, where $\beta = \frac{1}{2}(w^*/2c)^2$. Consequently, we have that $P_\infty = (1 - 2x_n)^{\binom{n}{2}}$ converges to 0 and 1 when $\beta < 2$ and $\beta > 2$, respectively, as n grows. These conditions on β translate through to the stated conditions on c in the statement of the proposition. \square

We now define a ‘stochastic safety radius’ that is scaled in such a way as to allow comparisons that are meaningful even as n tends to infinity.

Definition (*Stochastic safety radius*) For any $\eta > 0$, we will say that a distance-based tree reconstruction method M has η -stochastic safety radius $s = s_n$ if for every binary

phylogenetic X -tree T on n leaves, with minimum interior edge length w_{\min} , and with the distance δ on X described by the random errors model, we have:

$$c < s \cdot w_{\min} \implies \mathbb{P}(M(\delta) = T) \geq 1 - \eta.$$

This is equivalent to the condition:

$$\sigma < s \cdot \frac{w_{\min}}{\sqrt{\log(n)}} \implies \mathbb{P}(M(\delta) = T) \geq 1 - \eta.$$

Proposition 2 *For any method M that has l_∞ safety radius $\rho_n > 0$, and for any $\eta > 0$, there is a value $s = s_n > 0$ so that M has η -stochastic radius (at least) s for all binary trees on n leaves. Moreover, as $n \rightarrow \infty$ we can take s_n arbitrarily close to $\frac{1}{2}\rho_n$.*

Proof If $c \geq s \cdot w_{\min}$ then, from the analogue of (5) (with $\frac{1}{2}w^*$ replaced by $\rho \cdot w^*$), we have:

$$P_\infty \geq \left(1 - 2\mathbb{P}\left(N(0, 1) > \frac{\rho}{s}\sqrt{\log(n)}\right)\right)^{\binom{n}{2}}.$$

Applying the inequality in (2) gives:

$$P_\infty \geq \left(1 - n^{-\rho^2/2s^2}\right)^{\binom{n}{2}} \geq 1 - \binom{n}{2}n^{-\rho^2/2s^2},$$

and the very last term on the right can be made $< \eta$ by selecting $s = s_n$ sufficiently small. Moreover, as $n \rightarrow \infty$, we can take s_n to approach $\frac{1}{2}\rho_n$ for any $\eta > 0$. \square

Notice that this proof just sets a lower bound on the η -stochastic safety radius. This shows that any method having non-zero l_∞ safety radius (e.g. 1/2 for NJ) also has non-zero η -stochastic safety radius, which is roughly equal to half of the l_∞ safety radius (i.e. 1/4 with NJ). In other words, our definition provides non-trivial performance criteria for those methods, with a lower bound that is easily computed from the l_∞ results. However, the bound in Proposition 2 is very severe in that it requires that the l_∞ bound to hold for all pairwise distances. We will see that much better bounds do exist. Moreover, the following definition avoids having to consider the effect of η , which plays a minor role in all calculations.

Definition (*Limiting stochastic safety radius*) We say that a distance-based tree reconstruction method M has a *limiting stochastic safety radius* (LSSR) r if for every $s < r$ and every $\eta > 0$ the η -stochastic safety radius of M is at least s for all binary trees with sufficiently many leaves.

3 Theoretical Results

We first show that the limiting stochastic safety radius of a relatively simple quartet-based approach is considerably larger than the value $\frac{1}{4}$ that is required by Proposition 2

to satisfy the Atteson bound (where $\rho = \frac{1}{2}$). We then present our main theoretical result (Theorem 1), an absolute upper bound on the limiting stochastic safety radius of any distance-based method.

Proposition 3 *There is a distance-based tree reconstruction method which has a limiting stochastic safety radius of $\frac{1}{\sqrt{2}}$.*

Proof We use a result from [21] (see also [18]) which provides a tree reconstruction method that can recover a binary tree with n leaves by asking for the topology of $\Theta(n \log(n))$ quartets (the questions asked are allowed to depend on the answers obtained up to that point). Provided that all these quartet trees are correctly returned, we will infer the correct underlying parent tree. Now, suppose we set $\sigma = sw^*/\sqrt{\log n}$. From Eq. (3), the probability that any particular quartet is correctly inferred is at least:

$$1 - 2\mathbb{P}\left(N(0, 1) < -\frac{w\sqrt{\log(n)}}{w^*s}\right),$$

where w is the length of the interior path of the quartet in the parent tree. Since $w/w^* \geq 1$, Boole’s inequality implies that the probability that every particular quartet is correctly inferred is at least:

$$1 - 2Cn \log(n)\mathbb{P}\left(N(0, 1) < -\frac{\sqrt{\log(n)}}{s}\right), \tag{6}$$

where C is an upper bound constant in the $\Theta(n \log(n))$ construction. Notice that this holds even though the quartet decisions are not (stochastically) independent. The expression in (6) now converges to 1 for any value $s < \frac{1}{\sqrt{2}}$ as $n \rightarrow \infty$, by the asymptotic equivalence in (2). \square

Theorem 1 *No distance-based tree reconstruction method has a limiting stochastic safety radius greater than 1.*

Proof The idea of the proof is to show that MLE (maximum likelihood estimation) cannot allow a safety radius with $c > 1$ on a subset \mathcal{T} of trees (with prescribed branch lengths), from which it will follow that no other method could do so either on that subset \mathcal{T} (and thereby on all binary trees and with variable edge lengths). Here MLE is applied to data generated by just those trees in \mathcal{T} , and MLE is also only estimating trees within \mathcal{T} (and with the given prescribed branch lengths), thus it differs from the more usual MLE (or OLS) approach of finding an arbitrary tree for which an arbitrary assignment of positive branch lengths maximizes the likelihood. The reasons for using this restricted version of MLE will be made clear near the end of the proof.

Consider the binary tree T_{3n} on the leaf set $\mathcal{L} = \bigcup_{i=1}^n \{a_i, a'_i, b_i\}$ of size $3n$, obtained from any fixed binary tree on leaf set $\{1, 2, \dots, n\}$ by replacing each leaf i by the rooted triplet subtree $(a_i, a'_i)b_i$. For T_{3n} assign length 1 to all the interior edges and to the pendant edges that are incident with leaves of type a_i and a'_i (for all i), and assign the length 2 to the pendant edges that are incident with leaves of




Fig. 1 The tree $T_{3n}(\mathbf{t})$.

type b_i (for all i). For a sequence $\mathbf{t} = (t_1, t_2, \dots, t_n)$ where $t_i \in \{-1, 1\}$ let $T_{3n}(\mathbf{t})$ be the tree obtained from T_{3n} by interchanging the leaf labels a'_i and b_i precisely for each i for which $t_i = -1$ (leaving all edge lengths unchanged—thus all cherry pendant edges have length 1, while the non-cherry pendant edges have length 2). Thus $\mathcal{T} = \{T_{3n}(\mathbf{t}) : \mathbf{t} \in \{-1, 1\}^n\}$ is a set of 2^n binary trees, each with a leaf set \mathcal{L} of $3n$ leaves, and with the prescribed branch lengths described (see Fig. 1).

Let $d^{\mathbf{t}}$ denote the tree metric induced on \mathcal{L} by the tree $T_{3n}(\mathbf{t})$ with its associated branch lengths, and let $\delta^{\mathbf{t}}$ denote the corresponding distances under the random errors model (so $\delta^{\mathbf{t}} = d^{\mathbf{t}} + \epsilon$, for a vector ϵ of independent Gaussians with mean 0 and variance σ^2 (and independent of \mathbf{t})). Notice that we can partition any vector of distances δ on \mathcal{L} into two parts δ_B and δ_W , where δ_B compares leaves between different triplet-subtrees ($B =$ ‘between’) and δ_W compares leaves within given triplet-subtrees ($W =$ ‘within’); formally:

- δ_B is the sequence of δ -values for all pairs $\omega \in B$ where:

$$B = \{(l_i, l_j) : i, j \in \{1, \dots, n\}, i \neq j, l_i \in \{a_i, a'_i, b_i\}, l_j \in \{a_j, a'_j, b_j\}\};$$

- δ_W is the sequence of δ -values for all pairs $\omega \in W$, where:

$$W = \{(l_i, l'_i) : i \in \{1, \dots, n\}, l_i, l'_i \in \{a_i, a'_i, b_i\}, l_i \neq l'_i\}$$

It is useful to partition the ‘within’ pairs further as follows. For each $j \in \{1, \dots, n\}$, let $V(j) = \{(a_j, a'_j), (a_j, b_j)\}$ and let $U(j) = \{(a'_j, b_j)\}$.

A fundamental observation at this point is that the probability distribution of $\delta^{\mathbf{t}}$ on all pairs from B and the pair in $U(j)$ (for each j) does not depend on \mathbf{t} at all. Moreover, for pairs from $V(j)$, the dependence of $\delta^{\mathbf{t}}$ on \mathbf{t} is only via t_j . By this invariance and the independence assumption in the random errors model, for any δ and $\mathbf{t} \in \{-1, 1\}^n$, the probability density function for $\delta^{\mathbf{t}}$ can be written in the following factored way:

$$f(\delta|\mathbf{t}) = \prod_{\omega \in B} f(\delta_\omega) \cdot \prod_{j=1}^n f(\delta(a'_j, b_j)) \cdot \prod_{j=1}^n f(\delta(a_j, a'_j)|t_j) f(\delta(a_j, b_j)|t_j), \quad (7)$$

where, for $\omega \in V(j)$, the terms in the third product are given by:

$$f(\delta(\omega)|t_j) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\delta(\omega) - d_{w_j}(\omega))^2}{2\sigma^2}\right), \tag{8}$$

in which:

$$d_{w_j}(\omega) = \begin{cases} 2, & \text{if } t_j = 1 \text{ and } \omega = (a_j, a'_j) \text{ or } t_j = -1 \text{ and } \omega = (a_j, b_j); \\ 4, & \text{if } t_j = -1 \text{ and } \omega = (a_j, a'_j) \text{ or } t_j = 1 \text{ and } \omega = (a_j, b_j). \end{cases}$$

Notice that from our fundamental observation above, the terms in the first two products appearing in (7) do not depend at all on \mathbf{t} .

Thus, for any δ , the maximum likelihood estimate of \mathbf{t} given δ is the sequence (t_1, \dots, t_n) where, for each $j \in \{1, \dots, n\}$, t_j maximises the product

$$f(\delta(a_j, a'_j)|t_j) \cdot f(\delta(a_j, b_j)|t_j).$$

For such a ML estimate \mathbf{t} , the following inequality must hold for all $j \in \{1, \dots, n\}$:

$$L_j := \frac{f(\delta(a_j, a'_j)|t_j)}{f(\delta(a_j, a'_j)|-t_j)} \cdot \frac{f(\delta(a_j, b_j)|t_j)}{f(\delta(a_j, b_j)|-t_j)} \geq 1. \tag{9}$$

Assume now that $\mathbf{t} = \mathbf{1} = (1, 1, \dots, 1)$. Then if we let $\Delta_j = \delta(a_j, b_j) - \delta(a_j, a'_j)$, application of (8) in (9) simplifies (after some algebra) to the more attractive equation:

$$L_j = \exp\left(\frac{2\Delta_j}{\sigma^2}\right). \tag{10}$$

To this point, δ has been an arbitrary distance. Now, let us further assume that δ is generated on $T_{3n}(\mathbf{1})$ under the random errors model; in other words, $\delta = \delta^{\mathbf{t}}$ for $\mathbf{t} = \mathbf{1}$. We wish to calculate the probability—call it p_n —that MLE will correctly estimate the generating tree $T_{3n}(\mathbf{1})$. By (9) and independence assumptions in the random errors model, this probability p_n satisfies:

$$p_n \leq \mathbb{P}\left(\bigcap_{j=1}^n L_j \geq 1 \mid \mathbf{t} = \mathbf{1}\right) = \prod_{j=1}^n \mathbb{P}(L_j \geq 1 \mid \mathbf{t} = \mathbf{1}). \tag{11}$$

Now, from (10),

$$\mathbb{P}(L_j \geq 1 \mid \mathbf{t} = \mathbf{1}) = \mathbb{P}(\Delta_j > 0 \mid t_j = 1),$$

and Δ_j has a normal distribution with a mean of 2 and a variance of $2\sigma^2$. Thus:

$$\mathbb{P}(\Delta_j > 0 \mid t_j = 1) = 1 - \mathbb{P}\left(Z < -\frac{\sqrt{2}}{\sigma}\right), \tag{12}$$

where $Z = N(0, 1)$ is a standard normal random variable. Substituting (4) and (2) into (12) gives:

$$\mathbb{P}(\Delta_j > 0 | t_j = 1) = 1 - \beta_n,$$

where

$$\beta_n \sim \frac{c}{2\sqrt{\pi \log(n)}} n^{-1/c^2}.$$

Applying this to Eq. (11) gives that the probability p_n that MLE correctly estimates the generating tree satisfies

$$p_n \leq (1 - \beta_n)^n.$$

Straightforward calculus now shows that as $n \rightarrow \infty$, the sequence $(1 - \beta_n)^n$ (and hence p_n) converges to 0 if $c > 1$.

This shows that we cannot recover $T_{3n}(\mathbf{1})$ with an accuracy bounded away from 0 as n becomes large, by using MLE, if $s > 1$ in the definition of the stochastic safety radius (since the interior edges all have length 1, we have $w^* = 1$ and so $c > 1$ for $s > 1$). Moreover, by symmetry, the same conclusion applies to any of the 2^n trees $T_{3n}(\mathbf{t})$ (there is nothing ‘special’ about $\mathbf{t} = \mathbf{1}$). We now invoke a classic result that MLE is an estimation method that maximises the expected reconstruction accuracy of a discrete parameter when a family of distributions depends on just that parameter, and this parameter is chosen uniformly at random (*c.f.* Theorem 10.3.1 of [5] or Theorem 17.2 of [17])—in our case, the discrete parameter is the vector \mathbf{t} (which determines the tree $T_{3n}(\mathbf{t})$). By symmetry, the reconstruction accuracy of each tree in $T_{3n}(\mathbf{t})$ is the same (and therefore equal to the expected reconstruction accuracy), and since since $\mathcal{T} = \{T_{3n}(\mathbf{t}) : \mathbf{t} \in \{-1, 1\}^n\}$ is a subset of the set of all binary trees on $3n$ leaves with positive branch lengths, it follows that for any distance-based reconstruction method, the limiting stochastic safety radius cannot be larger than 1. \square

4 Simulation Results

We have seen in the previous section that the limiting stochastic safety radius (LSSR) of any algorithm is at most 1 (Theorem 1), and that a simple quartet algorithm [21] has a LSSR value at least $\frac{1}{\sqrt{2}} \approx 0.71$ (Proposition 3). The gap is relatively small between these two bounds and we expect that more sophisticated algorithms have LSSR values that are substantially higher than $\frac{1}{\sqrt{2}}$. In this section, we turn to simulations to study the accuracy of mainstream distance-based methods under the random errors model, with realistic numbers of taxa (previous results are asymptotic). Our goal is to compare these methods and to check how close they come in practice to the 1 bound prescribed by Theorem 1. Notice that this theorem was established using the pronged trees of Fig. 1. These trees are expected to be difficult for two reasons: (1) all internal branches have the same length and thus no branch is easy; (2) a large number ($2n/3$) of taxon pairs

are separated by a single internal branch and are likely to be wrongly exchanged, when trying to infer these trees from noisy data. Not all tree shapes possess this property; for example, in a perfectly balanced tree, all non-cherry taxon pairs are separated by at least *two* internal edges. Thus, we also compare, using simulations, different tree shapes, to establish if some of them are (stochastically) ‘harder’ than the others to reconstruct (depending, perhaps, on the inference method), or whether the opposite is true and all tree shapes seem to be equally difficult.

In the following, we first describe the methods being compared and the comparison criteria, we then study their performance with pronged and other (e.g. perfectly balanced) extreme trees, and, lastly, we use randomly generated tree shapes to obtain average accuracy measures under the random errors model.

4.1 Methods Tested and Comparison Criteria

We ran four standard, fast algorithms, which have been studied within the l_∞ safety radius framework, using FastME implementation (<http://www.atgc-montpellier.fr/fastme/>):

- (1) GME+OLS (Greedy minimum evolution with ordinary least squares) [7] is a greedy algorithm that iteratively adds taxa on a growing tree, minimizing at each step the ordinary least squares (OLS) tree length estimate, in accordance with the OLS version of the minimum evolution principle [24]. The performance of this algorithm was analysed by [20], who showed that its l_∞ safety radius tends to 0 with increasing n .
- (2) UNJ (Unweighted NJ) [10], which is the unweighted (OLS) version of NJ, with an l_∞ safety radius of $\frac{1}{2}$ [1].
- (3) GME+BME [7] which uses the same iterative taxon addition scheme as GME+OLS, but optimizes the balanced version of minimum evolution (BME, [22]) and has an optimal l_∞ safety radius of $\frac{1}{2}$ [20].
- (4) NJ [25] with l_∞ safety radius of $\frac{1}{2}$ [1]. We showed [15] that NJ greedily minimises BME at each agglomeration step (and not the OLS version of minimum evolution, as was originally suggested).

The aim was to see if there is any difference between the algorithmic schemes (taxon addition versus cherry agglomeration), and between the criteria being optimised (BME vs OLS minimum evolution). Because of the OLS-type noise in the random errors model, better accuracy is expected for OLS-based algorithms (UNJ and GME+OLS). On the other hand, the l_∞ safety radius differs widely among these algorithms and converges to 0 for GME+OLS; thus the second aim was to check whether Atteson’s predictions are observed in practice.

In addition, we implemented the quartet method of [21] using the first algorithm described by [18] with $O(n^2)$ time complexity and $O(n \log_2(n))$ quartet queries. The usual four-point rule [26,27] was used to answer the quartet queries. Our aim was to compare the accuracy of this simple algorithm, mostly used for theoretical purposes, to that of algorithms being widely used in phylogenetics, and to check how this algorithm behaves regarding our $\frac{1}{\sqrt{2}}$ bound of Proposition 3.

Table 1 Algorithm accuracy with the limiting noise level ($\sigma = w^*/\sqrt{\log(n)}$) and extreme tree shapes

Tree shape	#taxa	GME+OLS		GME+BME		UNJ		NJ	
		P_c	RF	P_c	RF	P_c	RF	P_c	RF
Pronged (Fig. 1)	12	0.748	0.03111	0.750	0.03088	0.898	0.01266	0.890	0.01377
	30	0.734	0.01103	0.758	0.00992	0.884	0.00444	0.872	0.00488
	90	0.768	0.00303	0.804	0.00243	0.888	0.00128	0.888	0.00128
	270	0.834	0.00069	0.840	0.00063	0.924	0.00029	0.920	0.00030
	810	0.856	0.00019	0.850	0.00020	0.931	0.00009	0.944	0.00007
Caterpillar	90	0.736	0.00360	0.426	0.00937	0.996	0.00005	0.966	0.00039
	270	0.792	0.00095	0.392	0.00348	0.998	0.00001	0.992	0.00003
Balanced	96	0.892	0.00133	0.884	0.00133	0.980	0.00032	0.976	0.00037
	384	0.892	0.00031	0.918	0.00024	0.976	0.00009	0.976	0.00009
Balanced + pronged	72	0.828	0.00281	0.846	0.00243	0.918	0.00130	0.916	0.00136
	288	0.854	0.00055	0.876	0.00047	0.938	0.00022	0.938	0.00022

Two criteria were used to compare algorithm accuracy: (1) the probability P_c of recovering the entire topology of the simulated tree; (2) the normalized bipartition distance (Robinson–Foulds (RF) [23]) between the inferred and simulated trees, which is equal to 0 when both trees define the same bipartitions, and to 1 when they do not share any bipartition in common.

4.2 Algorithm Accuracy with Pronged and Other Extreme Trees

We used pronged trees (Fig. 1, used in the proof of Theorem 1) with a number of taxa $n = 12, 30, 90, 270$ and 810 . We also used: caterpillar trees with $n = 90$ and 270 ; perfectly balanced trees with $n = 96$ and 384 ; and “balanced+pronged” trees, where each leaf of a perfectly balanced tree is replaced by the same three-taxon tree as in the pronged trees, with $n = 72$ and 288 . Caterpillar and perfectly balanced tree shapes are extreme regarding a number of measurements (e.g. diameter, number of cherries, etc.), and the pronged 3-taxon tree is assumed to make tree inference difficult (see above). In all of these trees, all internal branches had an equal length of 1, which (again) makes tree inference difficult. The length of the external branches was 1 for the caterpillar and balanced trees, and as shown in Fig. 1 for the pronged trees. The pairwise distances were computed and perturbed by an independent and identically distributed normal noise with standard deviation equal to $1/\sqrt{\log(n)}$, i.e. the highest possible noise level regarding Theorem 1, beyond which no algorithm can accurately recover every tree correctly as n grows. The goal was to check if, in these especially difficult conditions, the standard algorithms (e.g. NJ) still show some ability to recover the correct tree. In these conditions, the quartet method had very poor results that are not shown (but see below). For each of the tree shapes and n values, 500 data-sets were generated to obtain average error estimates. The results are summarised in Table 1. We see the following:

- *Pronged trees* are indeed difficult to reconstruct accurately, compared to perfectly balanced trees. However, for all algorithms, the probability of recovery (P_c) increases with n . According to this result, the four tested algorithms could have LSSR equal to 1. However, the algorithms are not equivalent in their performance. We see a clear advantage of the agglomerative scheme (NJ and UNJ) over taxon addition (GME + OLS and GME + BME), a finding that has already been observed in other simulation studies (e.g. [7]). On the other hand, there is no significant difference (considering both P_c and RF) between the algorithms that minimise the OLS version of minimum evolution (GME + OLS and UNJ) and their BME counterparts (GME + BME and NJ, respectively). Notably, we do not see any sign of weakness of GME + OLS, as predicted by its limiting l_∞ safety radius of zero [20].
- *Caterpillar trees* give another view. Again we observe the clear advantage of the agglomeration scheme that obtains nearly perfect results ($P_c \approx 1$), especially with UNJ, which is substantially better than NJ regarding both P_c and RF criteria. This latter finding is expected with such unbalanced trees, where the matrix reduction step is better achieved by UNJ, which accounts for the number of taxa in both agglomerated subtrees, while NJ uses equal weights of $\frac{1}{2}$. Based on these results, NJ and UNJ seem again to have LSSR of (close to) 1. GME-OLS has a lower P_c value, but this increases with n . In the opposite, GME-BME not only has low P_c (< 0.5), but this decreases with n . We have no clear explanation for this poor performance but, based on this result, it is unlikely that GME + BME has LSSR equal to 1.
- *Perfectly balanced trees* confirm again the superiority of the agglomeration scheme, compared to taxon addition. NJ and UNJ are nearly the same, as expected with well-balanced trees (see above). However, although these trees are relatively easy for all algorithms ($P_c \geq 0.9$), we do not see any improvement with larger n values. This questions our previous assumption that LSSR could be equal to 1 for the algorithms tested, unless the convergence towards LSSR is slow.
- *Balanced + pronged trees* are more difficult than balanced trees, as expected. However, for all algorithms, the accuracy increases with n . Algorithm comparisons are consistent with the previous ones: the agglomeration scheme performs better than taxon addition; there is little difference between the OLS and BME versions of algorithms, especially regarding NJ versus UNJ, which are nearly equivalent with such well-balanced trees.

To summarise, NJ and UNJ have remarkably high accuracy with these difficult trees and conditions. These results suggest that these two methods might have an LSSR equal to the optimal value of 1. Note, however, that the performance of NJ and UNJ with perfectly balanced trees lags a little behind, as their accuracy does not seem to improve when n increases; although, this may be because the convergence is slow. The taxon addition scheme is clearly less accurate and the results of GME-BME with caterpillar trees seem to indicate that this algorithm does not have an optimal LSSR of 1. Even though these conclusions are somewhat speculative, as n remained relatively moderate in all of our experiments, these results provide directions for future investigations on the LSSR of mainstream methods.

Table 2 Algorithm accuracy with random trees and a limiting noise level ($\sigma = w^*/\sqrt{\log(n)}$)

#taxa	GME+OLS		GME+BME		UNJ		NJ	
	P_c	RF	P_c	RF	P_c	RF	P_c	RF
10	0.710	0.04714	0.672	0.05542	0.904	0.01457	0.854	0.02171
30	0.778	0.00918	0.800	0.00822	0.948	0.00199	0.936	0.00251
90	0.762	0.00328	0.830	0.00209	0.964	0.00045	0.960	0.00051
270	0.646	0.00164	0.800	0.00080	0.966	0.00014	0.966	0.00014
810	0.498	0.00089	0.798	0.00027	0.970	0.00004	0.970	0.00004

4.3 Algorithm Accuracy with Random Trees

Up to this point, our experiments have used a limited number of extreme tree shapes. In this subsection, we use random trees in the search for other potentially difficult cases and to estimate the average accuracy of tested algorithms under the random errors model.

Tree shapes were generated uniformly at random (the so-called ‘Proportional-to-Distinguishable-Arrangements (PDA) model’) and all branch lengths were set to 1. The number of taxa was $n = 10, 30, 90, 270$ and 810. The pairwise distances were perturbed by an i.i.d. normal noise with standard deviation equal to $1/\sqrt{\log(n)}$, as in previous experiments. Again, these noise level and trees (with equal branch lengths) make tree inference difficult.

The average results over 500 data-sets are displayed in Table 2. NJ and UNJ results are quite consistent with those obtained with extreme tree shapes (cf Table 1). With moderate values of n , UNJ is slightly more accurate than NJ, as expected with the OLS-type noise used in these simulations. With large values of n , UNJ and NJ are nearly perfect ($P_c \approx 1$), and both algorithms become strictly equivalent. Again, the results seem to indicate that both UNJ and NJ could have an optimal LSSR of 1. We cannot exclude that particularly difficult trees do exist, but these must be rare, and in average NJ and UNJ appear to be highly accurate at a noise level that is four times larger than the limit for Atteson’s approach to apply.

With taxon addition, the picture is different. GME+OLS has low accuracy that drops when n increases, a finding which may be seen as consistent with the predictions from [20] using the l_∞ safety radius. GME+BME performs better but its accuracy is not that high and seems to stabilise around 0.8 when n increases. This observation is probably explained by the fact that some trees (e.g. caterpillars, Table 1) are difficult for this algorithm. In both cases, the results in Table 2 seem to indicate that neither GME+BME nor GME+OLS has an optimal LSSR of 1.

4.4 Algorithm Accuracy with Various Noise Levels

We also ran simulations where the noise level varied around the limiting value used in previous experiments ($\sigma = rw^*/\sqrt{\log(n)}$, $r = 1$), in order to study the sharpness

Table 3 Algorithm accuracy with various noise levels

r (noise level)	#taxa	GME+OLS		GME+BME		UNJ		NJ	
		P_c	RF	P_c	RF	P_c	RF	P_c	RF
3/4	30	0.986	0.00051	0.992	0.00029	1.0	0.0	1.0	0.0
	90	0.974	0.00029	0.986	0.00016	1.0	0.0	1.0	0.0
	270	0.982	0.00007	0.996	0.00001	0.998	<0.00001	0.998	<0.00001
1	30	0.778	0.00918	0.800	0.00822	0.948	0.00199	0.936	0.00251
	90	0.762	0.00328	0.830	0.00209	0.964	0.00045	0.960	0.00051
	270	0.646	0.00164	0.800	0.00080	0.966	0.00014	0.966	0.00014
4/3	30	0.316	0.04170	0.314	0.04177	0.670	0.01414	0.606	0.01696
	90	0.204	0.01972	0.240	0.01696	0.634	0.00556	0.616	0.00590
	270	0.054	0.01099	0.122	0.00776	0.592	0.00202	0.560	0.00217

Table 4 Accuracy of the quartet method with various noise levels

r (noise level)	#taxa	3/5		3/4	
		P_c	RF	P_c	RF
30		0.902	0.556	0.556	0.09170
90		0.946	0.479	0.479	0.10795
270		0.952	0.370	0.370	0.12478

of our bounds with realistic numbers of taxa. Our goal was to check whether the accuracy improves substantially for taxon addition when r is less than 1, and whether the agglomeration algorithms (NJ and UNJ) still show some ability to recover the correct tree when r is larger than 1. We also studied the performance of the quartet method with (relatively) low r values. We used the same random trees (with all branch lengths equal to $1 = w^*$), as in the previous subsection, but used $r = 3/5, 3/4$ and $4/3$ instead of $r = 1$.

The average results over 500 data sets are displayed in Table 3 for the standard algorithms. We see that our previous conclusions are confirmed when r differs from 1: the agglomeration scheme performs better than taxon addition; there is little difference between both minimum evolution versions; however, UNJ is slightly better than NJ (e.g. see $r = 4/3$, with both P_c and RF). We also see that the 1 bound prescribed by Theorem 1 is rather sharp with moderate number of taxa: with $r = 3/4$, both taxon addition algorithms have high accuracy for all n values, and NJ and UNJ are nearly perfect. Conversely, with $r = 4/3$, the accuracy of all algorithms drops dramatically, especially that of the taxon addition scheme where P_c approaches 0 with $n = 270$. This confirms that the limiting optimal bound of Theorem 1, obtained with special (“pronged”) trees, is robust and found again, at least qualitatively, with random trees and realistic n values.

Table 4 displays the results of the quartet method for $r = 3/5$ and $3/4$ (the results with $r = 1$ are quite poor and are not shown). Again, the stochastic radius framework and our bounds of Proposition 3 ($LSSR \geq 1/\sqrt{2} \approx 0.71$) and Theorem 1 ($LSSR \leq 1$) have a good predictive accuracy. With $r = 3/5$, the accuracy is high and increases

with n , while we observe the opposite with $r = 3/4$, which seems to indicate that the quartet method has a LSSR close to our $1/\sqrt{2}$ bound. Note, however, that when r approaches $\frac{1}{\sqrt{2}}$ from below (e.g. $r = 2/3$ or $r = 0.70$), the accuracy is not that high and does not increase with the moderate values of n used in these simulations. This is most likely explained by the very slow convergence of the bound in Eq. (6), combined with asymptotic equivalence (2), when r is close to $\frac{1}{\sqrt{2}}$.

5 Discussion

Our simulation results show that the stochastic radius framework introduced in this paper has a good predictive capacity and seems to be robust. The optimal bound ($r = 1$) of Theorem 1, which was obtained with special “pronged” trees, seems to apply to a large variety of tree topologies (no topology is easy or else these are quite rare). Moreover, when the noise level decreases below $r = 1$, the accuracy rises for all algorithms and all values of n . The behaviour of the quartet method (with moderate n values) is also consistent with the (limiting) prediction of Proposition 3. We thus believe that the LSSR approach will show a high capacity in predicting algorithm performance in realistic conditions, a property which does not hold in several cases with the l_∞ safety radius, as noted in the Introduction.

An important outcome of the simulations is that the ability to recover single branches may still be high, even when the probability of recovering the entire tree drops due to high noise level; for example, with $r = 4/3$ and $n = 270$, the probability that any given branch is correct is higher than ~ 0.99 for all standard algorithms (Table 3, RF values). This strongly suggests the development of stochastic ‘edge radius’ approaches (analogous to the classical non-stochastic concept considered also by Atteson [1]) which would account for the length of the branch being considered, and thus will not use the worst-case approach used here in several places, where all branches have the same length. In other words, the tree T may have some very short edges, however, provided a given edge is not too short, then we may be able to recover the corresponding split with high accuracy, even if the entire tree cannot be reconstructed. Finer algorithm analyses should follow from such a framework.

Our study also suggests two further theoretical questions that would be worth investigating in future. Firstly, it would be of interest to analytically calculate the precise limiting stochastic safety radius of NJ and other standard methods; in particular, to determine if it takes the value 1 or some number less than this.

It would also be of interest to study the stochastic safety radius of distance-based tree reconstruction methods for the more general class of models in which the ϵ_{xy} values have a multivariate normal distribution, with means of 0 and a covariance matrix Σ (the OLS model we considered in this paper assumes that Σ is the diagonal matrix with each diagonal entry equal to σ^2). In such models, the variance of ϵ_{xy} would typically increase with the path length between leaves x and y in the tree (the weighted least squares (WLS) assumption [9]), while the covariance for two pairs of taxa would typically increase with the total length of the shared branches that are present in both paths connecting each pair (the generalized least-squares (GLS) assumption [4]). Such extended models provide a more accurate approximation to the complex distribution

of corrected distances under site substitution models of Jukes–Cantor type (*c.f.* Section 6.2 of [12]), which explains the superior performance of (GLS-based) BIONJ [11] over NJ on distances from simulated sequence data. Note, however, that UNJ (the OLS version of NJ) was shown to be fairly accurate with sequence data [13], which demonstrates the robustness of our simple OLS-type error model used in this paper. Moreover, analysing the stochastic safety radius of different methods under this extended (GLS) model will be less straightforward, since we have exploited the independence assumption freely in many of the arguments above.

Finally, we remark that, although the variance σ^2 in the random errors model that allows accurate tree reconstruction needs to reduce with increasing n , the explicit dependence seems to involve only a very modestly declining function of n , namely $1/\log(n)$. However σ^2 is also proportional to shortest interior edge length w_{\min} , and this typically will also decrease with n . For instance, if edges weights are simulated using an i.i.d. exponential distribution, then w_{\min} declines at the rate $1/n$. Similarly, if we increase the sampling of species within a given clade, then w_{\min} is likely to decrease at least as fast. These comments concern the accurate reconstruction of the entire tree—for the more modest goal of accurately reconstructing just a given edge of interest in the tree, the question of determining how the stochastic ‘edge radius’ (mentioned above) might depend on n would be a worthy topic for further investigation.

Acknowledgments MS thanks the Allan Wilson Centre and the NZ Marsden Fund for supporting this work. We thank the two anonymous reviewers for a number of helpful suggestions.

Appendix: Proof of (2)

Substituting $t = x + u, u \geq 0$ in $\mathbb{P}(N(0, 1) > x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$ gives:

$$\mathbb{P}(N(0, 1) > x) = e^{-x^2/2} \int_0^\infty \frac{1}{\sqrt{2\pi}} e^{-xu} e^{-u^2/2} du < e^{-x^2/2} \int_0^\infty \frac{1}{\sqrt{2\pi}} e^{-u^2/2} du,$$

where the second inequality is from $e^{-xu} < 1$ for all $x, u > 0$. Since the last term on the right is $\frac{1}{2}$, we get the inequality in (2). Turning to the asymptotic relationship, consider:

$$\lim_{x \rightarrow \infty} \frac{\frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-t^2/2} dt}{\frac{1}{x\sqrt{2\pi}} e^{-x^2/2}}. \tag{13}$$

Since the numerator and denominator limits are both zero, we can apply L’Hôpital’s rule. Straightforward calculus (using the fundamental theorem of calculus for the numerator) establishes that the limit in (13) equals 1. □

References

1. Atteson, K.: The performance of neighbor-joining methods of phylogeny reconstruction. *Algorithmica* **25**(2–3), 251–278 (1999)
2. Berry, V., Gascuel, O.: Inferring evolutionary trees with strong combinatorial evidence. *Theor. Comput. Sci.* **240**(2), 271–298 (1997)

3. Bordewich, M., Mihaescu, R.: Accuracy guarantees for phylogeny reconstruction algorithms based on balanced minimum evolution. In: Moulton, V., Singh, M. (eds.) Proceedings of WABI 2010, 10th international workshop on algorithms in bioinformatics, volume 6293 of LNBI, pp. 250–261. Springer (2010)
4. Bulmer, M.: Use of the method of generalized least-squares in reconstructing phylogenies from sequence data. *Mol. Biol. Evol.* **8**, 868–883 (1991)
5. Casella, G., Berger, R.L.: *Statistical Inference*. Duxbury Press, Belmont (1990)
6. Cavalli-Sforza, L.L., Edwards, A.W.F.: Phylogenetic analysis: models and estimation procedures. *Am. J. Hum. Genet.* **19**, 223–257 (1967)
7. Desper, R., Gascuel, O.: Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle. *J. Comput. Biol.* **9**, 687–706 (2002)
8. Eickmeyer, K., Huggins, P., Pachter, L., Yoshida, R.: On the optimality of the neighbor-joining algorithm. *Algorithms Mol. Biol.* **3**, 5 (2008)
9. Fitch, W.M., Margoliash, E.: Construction of phylogenetic trees. *Science* **155**, 279–284 (1967)
10. Gascuel, O.: Concerning the NJ algorithm and its unweighted version, UNJ. In: Mirkin, B., McMorris, F.R., Roberts, F.S., Rzhetsky, A. (eds.) *Mathematical Hierarchies and Biology*, pp. 149–170. American Mathematical Society, Providence (1997)
11. Gascuel, O.: "BIONJ: an improved version of the NJ algorithm based on a simple model of sequence data. *Mol. Biol. Evol.* **14**(7), 685–695 (1997)
12. Gascuel, O.: Data model and classification by trees: the minimum variance reduction (MVR) method. *J. Classif.* **17**, 69–99 (2000)
13. Gascuel, O.: On the optimization principle in phylogenetic analysis and the minimum evolution criterion. *Mol. Biol. Evol.* **17**(3), 401–405 (2000)
14. Gascuel, O., McKenzie, A.: Performance analysis of hierarchical clustering algorithms. *J. Classif.* **21**, 3–18 (2004)
15. Gascuel, O., Steel, M.: Neighbor-Joining revealed. *Mol. Biol. Evol.* **23**(11), 1997–2000 (2006)
16. Gascuel, O., Levy, D.: A reduction algorithm for approximating a (nonmetric) dissimilarity by a tree distance. *J. Classif.* **13**, 129–155 (1996)
17. Guiasu, S.: *Information Theory with Applications*. McGraw-Hill, New York (1977)
18. Kannan, S.K., Lawler, E.L., Warnow, T.J.: Determining the evolutionary tree using experiments. *J. Algorithms* **21**, 26–50 (1996)
19. Mihaescu, R., Levy, D., Pachter, L.: Why neighbor-joining works. *Algorithmica* **54**(1), 1–24 (2009)
20. Pardi, F., Guillemot, S., Gascuel, O.: Robustness of phylogenetic inference based on minimum evolution. *Bull. Math. Biol.* **72**, 1820–1839 (2010)
21. Pearl, J., Tarsi, M.: Structuring causal trees. *J. Complex.* **2**, 60–77 (1986)
22. Pauplin, Y.: Direct calculation of a tree length using a distance matrix. *J. Mol. Evol.* **51**, 41–47 (2000)
23. Robinson, D.R., Foulds, L.R.: Comparison of phylogenetic trees. *Math. Biosci.* **53**, 131–147 (1981)
24. Rzhetsky, A., Nei, M.: Theoretical foundation of the minimum-evolution method of phylogenetic inference. *Mol. Biol. Evol.* **10**, 1073–1095 (1993)
25. Saitou, N., Nei, M.: The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Mol. Biol. Evol.* **4**, 406–425 (1987)
26. Sattath, S., Tversky, A.: Additive similarity trees. *Psychometrika* **42**, 319–345 (1997)
27. Zarestkii, K.: Reconstructing a tree from the distances between its leaves (In Russian). *Uspehi Matematicheskikh Nauk* **20**, 90–92 (1965)