Phylogenetic information complexity: Is testing a tree easier than finding it?

Mike Steel [a,*], Laszlo Székely [b], Elchanan Mossel [c]

[a] Biomathematics Research Centre, Department of Mathematics and Statistics, University of Canterbury, Christchurch, New Zealand
[b] Department of Mathematics, University of South Carolina, Columbia, SC, USA
[c] Statistics Department, UC Berkeley, Berkeley, CA, USA

A R T I C L E   I N F O

A B S T R A C T

Phylogenetic trees describe the evolutionary history of a group of present-day species from a common ancestor. These trees are typically reconstructed from aligned DNA sequence data. In this paper we analytically address the following question: Is the amount of sequence data required to accurately reconstruct a tree significantly more than the amount required to test whether or not a candidate tree was the 'true' tree? By 'significantly', we mean that the two quantities do not behave the same way as a function of the number of species being considered. We prove that, for a certain type of model, the amount of information required is not significantly different; while for another type of model, the information required to test a tree is independent of the number of leaves, while that required to reconstruct it grows with this number. Our results combine probabilistic and combinatorial arguments.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Phylogenetic trees are widely used in evolutionary biology to describe how species have evolved from a shared ancestral species. In the last 25 years, aligned DNA sequence data and related sequences (amino acids, codons, etc.) have been widely used for reconstructing and analysing these trees (Felsenstein, 2003; Semple and Steel, 2003). Tree reconstruction methods usually assume that sequence sites evolve according to some Markov process. The question of how much data is required to reconstruct a phylogenetic tree has been considered by a number of biologists (Churchill et al., 1992; Lecointre et al., 1994; Saitou and Nei, 1986; Wortley et al., 2005) and is topical, as it is not clear whether all trees for all taxa sets could be reconstructed accurately from the available data.

In earlier papers, we have analytically quantified the sequence length required for accurate tree reconstruction when sites evolve i.i.d. under various Markov processes (Daskalakis et al., 2006; Erdös et al., 1999; Mossel and Steel, 2004, 2005; Steel and Székely (2002)). These bounds typically depend on the number of taxa and the properties of the tree—in particular, how close the probability of a change of state ('substitution') on any edge is to 0 or to its upper bound in the model. It is the rate of growth in the sequence with the number of taxa that is of interest here. The growth rate in sequence length required for accurate tree reconstruction has a trivial lower bound growth of $\log(n)$—this

comes from simply comparing the number of binary trees on n leaves with the number of collections of n sequences of given length. What is perhaps surprising is that for certain finite-state Markov processes this primitive rate of growth can be achieved for some models (Daskalakis et al., 2006), given a bound on the substitution probabilities. This $\log(n)$ upper bound on sequence length also applies to a discrete infinite-state model (the 'random cluster model') (Mossel and Steel, 2004), given similar bounds on the substitution probabilities. The $\log(n)$ behaviour for these two models changes to a polynomial dependence on n when the probabilities of state change are allowed to exceed a certain critical value.

In this paper, we address a quite different question: namely if one has both the data and a proposal for a 'true' tree (i.e. the tree that produced the data under the model) we would like to test whether this tree is indeed the true tree, or whether some other tree must have produced the data. The answer provided must be correct with high probability. This concept of testing fits into the Popperian tradition—we would like to be able to refute the hypothesis that a particular tree produced the data, without necessarily exhibiting the tree that did. In statistics, the theory of testing among a discrete set of hypotheses has a long history (see, for example, Wald's paper from 1948, Berger and Wald, 1949). In contemporary phylogenetics, the concept of testing a tree is timely, as various 'tree of life' projects begin to provide detailed, large candidate evolutionary trees, the question of using data to 'test' any such candidate tree arises.

In this paper, we ask whether the information (sequence length) required for these two tasks—reconstructing versus testing—is fundamentally the same, i.e. that it grows at the same rate as a function of the size of the taxon set. Intuitively, testing

* Corresponding author. Tel.: +64 3 3667001x7688; fax: +64 3 3642587.
E-mail addresses: m.steel@math.canterbury.ac.nz (M. Steel), szekely@math.s-c.edu (L. Székely), mossel@stat.berekeley.edu (E. Mossel).

should be 'easier' (require less data) than reconstructing, since for testing, one has additional information, and one is simply asked to make a binary decision. This suggested difference is somewhat analogous to the '$P \neq NP$' conjecture in computational complexity, where any candidate solution for a problem in the class *NP* can be readily verified or refuted (in polynomial time) but it is believed that finding such a solution is fundamentally harder (i.e. not always possible in polynomial time) (Garey and Johnson, 1979). Of course, in our setting, we are dealing with sequence length, not computing time, but the two problems have an analogous flavour.

In the next section, we describe a general framework for discussing these issues, and we exhibit an (abstract) example where the sequence length required to test a discrete parameter is far less than the sequence length required to reconstruct it. Turning to the phylogenetic setting in Section 3, we show that when sequence sites evolve i.i.d. according to a finite-state Markov process then testing requires sequence length growth of rate $\log(n)$—which is the same as reconstructing requires, given bounds on the substitution probabilities. By contrast, for a discrete infinite-state Markov process, the situation is quite different—constant length sequences suffice to test, but order $\log(n)$ length sequences are required for reconstruction. We conclude the paper with some brief comments.

2. Testing versus reconstructing

In this section, we describe definitions and properties of testing and reconstructing in a general setting; we will specialize our approach to the phylogenetic setting in the following section.

Suppose $A = A_n$ and $U = U_n$ are finite sets, and that we have a random variable $X = X_{(a,\theta)}$ taking values in U and whose distribution depends on the discrete parameter $a \in A$, and perhaps some nuisance parameter θ taking values in a set $\Theta(a)$ (in the next section, A will be a set of trees, U will be a set of site patterns and the nuisance parameters will be edge lengths of the tree—all these concepts will be described later). We call $X_{(a,\theta)}$ a *parameterized random variable*, and when the nuisance parameter is either absent or has been specified for each a (so that the distribution of $X_{(a,\theta)}$ just depends on a, in which case we will often write X_a), we refer to a *simply parameterized random variable*. Most of our results concern the simply parameterized setting, except for our second main result, Theorem 3.2, where the further generality is allowed.

Given a sequence $\mathbf{u} = (u_1, u_2, \ldots, u_k)$ of k i.i.d. observations of $X_{(a,\theta)}$, we would like to use $\mathbf{u}$ to identify the discrete parameter a correctly with high probability. This reconstruction task is always possible for sufficiently large values of k provided a weak 'identifiability' condition holds, namely that for all $a \in A$, and $\theta \in \Theta(a)$ we have

$$\inf_{a' \neq a, \theta' \in \Theta(a')} d((a,\theta),(a',\theta')) > 0$$

(see Steel and Székely, 2009), where $d((a,\theta),(a',\theta'))$ denotes the l_1 distance between the probability distribution of the random variables $X_{(a,\theta)}$ and $X_{(a',\theta')}$.

The two tasks that we consider can be summarized, informally, as follows.

Reconstructing. Given $\mathbf{u} \in U^k$, determine with high probability the value $a \in A$ that generated $\mathbf{u}$ (for some $\theta \in \Theta(a)$).

Testing. Given a candidate value $a \in A$, as well as $\mathbf{u} \in U^k$, determine with high probability whether or not $\mathbf{u}$ was generated by (a, θ) (for some $\theta \in \Theta(a)$).

We are interested in determining and comparing the number of i.i.d. samples required to carry out these tasks. Clearly testing is 'easier' than reconstructing (i.e. testing requires a smaller or equal

value of k than reconstructing for the same accuracy), since one can always test by reconstructing and comparing the reconstructed object with the candidate object. Thus we will be particularly interested in determining whether the value of k grows at the same rate with n for these two tasks.

In general a basic lower bound on k is required for reconstructing, and a much weaker one for testing, which shows that testing might require asymptotically much shorter sequences. Before describing these bounds (Proposition 2.2), we formalize the concept of reconstructing and testing.

2.1. Definitions: reconstruction, testing and accuracy

Throughout we will let $(X^1_{(a,\theta)}, \ldots X^k_{(a,\theta)})$ denote a sequence of k i.i.d. observations generated by (a, θ).

A *reconstruction method R* is a randomized estimator of the element $a \in A$ that (along with some θ parameter) generated the sequence of k i.i.d. observations. More formally, a (deterministic) reconstruction method R is a function $R : U^k \to A$ (the value $R(\mathbf{u})$ is the value reconstructed when observing $\mathbf{u} \in U^k$), while, more generally, a randomized reconstruction method R takes a value (also denoted $R(\mathbf{u})$) in A which depends not just on $\mathbf{u}$ but on a random variable that is independent of all the random variables used to generate $\mathbf{u}$. This extra generality arises, for example, with reconstruction methods that have to break ties when performing optimization steps. Then

$$\rho^R_{(a,\theta)} := \mathbb{P}(R((X^1_{(a,\theta)}, \ldots X^k_{(a,\theta)})) = a)$$

is the probability that R will correctly reconstruct a from k i.i.d. samples generated by (a, θ). We say that R has (*reconstruction*) *accuracy* $1 - \varepsilon$ (*for k samples*) if, for all $a \in A$ and $\theta \in \Theta(a)$, we have

$$\rho^R_{(a,\theta)} \geqslant 1 - \varepsilon.$$

Note that, if we let $p_{a,\theta}(u) = \mathbb{P}(X_{(a,u)} = u)$, and for $\mathbf{u} \in U^k$ let[1]

$$p_{a,\theta}(\mathbf{u}) := \prod_{i=1}^{k} p_{a,\theta}(u_i),$$

then

$$\rho^R_{(a,\theta)} = \sum_{\mathbf{u} \in U^k} p_{(a,\theta)}(\mathbf{u}) \cdot \mathbb{P}(R(\mathbf{u}) = a).$$

A *testing process* ψ is a collection $\psi(a, \mathbf{u}) : a \in A, \mathbf{u} \in U^k$ taking values in {true, false} where ψ is either a function that depends just on the pair $(a, \mathbf{u})$ (in which case we say ψ is *deterministic*) or ψ depends on the pair $(a, \mathbf{u})$ and a random variable that is independent of the random variables used to generated $\mathbf{u}$ (which, as for reconstruction, allows extra generality such as breaking ties). In more familiar statistical terms (see for example Castella and Berger, 2002) a testing process is a collection of pairs of matching null and alternative hypothesis tests, one for each element $a \in A$, with each null hypothesis being that a generated the data ($\mathbf{u}$) for some associated nuisance parameter θ, and the matching alternative hypothesis being that the data were generated by an unknown element of A that is different to a.

Let

$$\varepsilon_1, \varepsilon_2 \geqslant 0 \quad \text{and} \quad \varepsilon := \max\{\varepsilon_1, \varepsilon_2\}.$$

We say that a testing process has *accuracy* $1 - \varepsilon_1, 1 - \varepsilon_2$ (*for k samples*) if for all $a \in A$, and $\theta \in \Theta(a)$, the following two conditions hold:

$$\mathbb{P}(\psi(a, (X^1_{(a,\theta)}, \ldots, X^k_{(a,\theta)})) = \text{true}) \geqslant 1 - \varepsilon_1, \tag{1}$$

---

[1] We will write $p_a(\mathbf{u})$ in the case of a simply parameterized random variable.

and for any $b \neq a$, and $\theta' \in \Theta(b)$,

$$\mathbb{P}(\psi(a, (X^1_{(b,\theta')}, \ldots, X^k_{(b,\theta')})) = \text{false}) \geqslant 1 - \varepsilon_2.$$

In other words, ψ returns 'true' with probability at least $1 - \varepsilon_1$ when the discrete parameter $a \in A$ is tested against the data it produced, and ψ returns 'false' with probability at least $1 - \varepsilon_2$ when any other particular element of A is tested against the data. In statistical terms we can regard ε_1 as an upper bound on the type I error rate (the probability of rejecting a true null hypothesis) and ε_2 as an upper bound on the type II error rate (the probability of accepting the null hypothesis when it is false).

Note that any collection of random variables $X_{(a,\theta)}$ has a trivial testing process with accuracy $1 - \varepsilon_1$, $1 - \varepsilon_2$ when $\varepsilon_1 + \varepsilon_2 = 1$; namely for each $a \in A$ and $\mathbf{u} \in U$ let $\psi(a, \mathbf{u}) = \text{true}$ with probability $1 - \varepsilon_1$. To exclude such trivialities we will generally assume that $\varepsilon < \frac{1}{2}$.

If Inequality (1) is strengthened to requiring a type I error of zero:

$$\mathbb{P}(\psi(a, (X^1_{(a,\theta)}, \ldots, X^k_{(a,\theta)})) = \text{true}) = 1$$

for all $a \in A$ and $\theta \in \Theta(a)$ (i.e. $\varepsilon_1 = 0$), we say that the testing process has *strong accuracy* $1 - \varepsilon$.

Given a reconstruction method R, a canonical testing process ψ_R is associated with it, defined as follows:

$$\psi_R(a, \mathbf{u}) = \text{true} \iff R(\mathbf{u}) = a$$

for all $a \in A$ and $\mathbf{u} \in U^k$. The following lemma follows easily from the definitions.

Lemma 2.1. *Given a parameterized random variable $X_{(a,\theta)}$ and an associated reconstruction method R with accuracy $1 - \varepsilon$, the associated testing process ψ_R has accuracy $1 - \varepsilon, 1 - \varepsilon$.*

The following proposition shows that reconstructing in general can require considerably longer sequences than testing.

Proposition 2.2.

 (i) *Suppose a simply parameterized random variable X_a ($a \in A$) takes values in U and has a reconstruction method with accuracy strictly greater than $\frac{1}{2}$ for k samples. Then $|A| \leqslant |U|^k$, and so $k \geqslant \log(|A|) / \log(|U|)$, or, equivalently:*

$$\log(|U|) \geqslant \frac{1}{k} \log(|A|).$$

(ii) *For any $\varepsilon > 0$ and $k = 1$, there exist sets A, U and a simply parameterized random variable X_a (for $a \in A$) taking values in U, and a deterministic testing procedure ψ that has strong accuracy of $1 - \varepsilon$, such that*

$$\log(|U|) = O(\log(\log(|A|))).$$

Proof. Part (i) was established in Steel and Székely (1999) (Theorem 2.1(ii)). For Part (ii), let $U = \{1, \ldots, n\}$ and let A be a collection of subsets of U with the property that for any two elements $a, a' \in A$:

$$\frac{|a \cap a'|}{\min\{|a|, |a'|\}} \leqslant \varepsilon. \tag{2}$$

Consider the following simply parameterized random variable $X_a (a \in A)$ defined by the rule that X_a selects one of the elements of a uniformly at random. Note that X_a takes values in the set U. Consider the deterministic testing process ψ defined by the rule:

$$\psi(a, u) = \begin{cases} \text{true} & \text{if } u \in a, \\ \text{false} & \text{otherwise.} \end{cases}$$

Clearly,

$$\mathbb{P}(\psi(a, X_a) = \text{true}) = 1,$$

and Condition (2) ensures that for $b \neq a$:

$$\mathbb{P}(\psi(a, X_b) = \text{true}) \leqslant \varepsilon,$$

and so ψ has strong accuracy $1 - \varepsilon$. Thus it remains to construct a family A of subsets of $\{1, \ldots, n\}$ satisfying (2) and of cardinality at least e^{n^η} for some $\eta > 0$ (since in that case $\log(|U|) = O(\log(\log(|A|)))$).

The existence of such a large collection can be established by using the probabilistic method (Alon and Spencer, 2008) as follows. Construct N random subsets of $\{1, \ldots, n\}$ by the following process: for each element i of $\{1, \ldots, n\}$ place i in the set with probability $n^{-2/3}$; otherwise, leave that element out. Using standard results on the asymptotic distribution of sums of i.i.d. random variables, the probability p of the event that (i) all of the N sets are of size at least $n^{0.3}$, and (ii) that all pairs of sets intersect in at most $n^{0.2}$ points satisfies (by the subadditivity of probability):

$$p \geqslant 1 - N \exp(-n^{c_1}) - \binom{N}{2} \exp(-n^{c_2})$$

for positive constants c_1, c_2. Now, for $N = e^{n^\eta}$ where $0 < \eta < \min\{c_1, c_2\}$ it holds that $p > 0$ for sufficiently large values of n. In this case, a collection of sets must exist that satisfy Conditions (i) and (ii). Finally, this collection will also satisfy (2) provided n is large enough that $n^{0.2}/n^{0.3} < \varepsilon$. This completes the proof. □

In summary, it is clear that in general, 'testing' can require considerably less information (sequence length) than 'reconstructing.' In the next section we consider what happens in a specific setting that arises in computational evolutionary biology.

We end this section with a technical result concerning testing that is required in the next section.

Lemma 2.3. *Suppose a simply parameterized random variable X_a has a testing process ψ with accuracy $1 - \varepsilon_1$, $1 - \varepsilon_2$ for k samples. Then:*

 (i) $d^{(k)}(a, a') \geqslant 2(1 - \varepsilon_1 - \varepsilon_2)$ *for all $a, a' \in A$ with $a \neq a'$,*

   where $d^{(k)}(a, a') := \sum_{\mathbf{u} \in U^k} |p_a(\mathbf{u}) - p_{a'}(\mathbf{u})|.$

(ii) *Let A' be a proper, nonempty subset of A, and let $a \in A - A'$. Consider the following random variable X' that is simply parameterized by the set $\{a, *\}$, and defined as follows: $X'_a = X_a$, and $X'_* = X_Y$ where Y is selected uniformly at random from the nonempty set A'. Then X' has a reconstruction method with accuracy $1 - \varepsilon$ for k samples (where $\varepsilon := \max\{\varepsilon_1, \varepsilon_2\}$). In particular*

   $$d^{(k)}(a, *) \geqslant 2(1 - 2\varepsilon).$$

Proof. Part (i). First observe that $d^{(k)}(a, a')$ is twice the variational distance between the probability distributions p_a and $p_{a'}$ on U^k, i.e.

$$d^{(k)}(a, a') = 2 \cdot \max_E |\mathbb{P}_a(E) - \mathbb{P}_{a'}(E)|,$$

where maximization is over all events E on U^k. For each $a, a' \in A$, let $E_{a,a'}$ be the event that $\psi(a, (X^1_a, \ldots X^k_a)) = \text{true}$. Then $E_{a,a'}$ has probability at least $1 - \varepsilon_1$ when $a = a'$ and probability at most ε_2 when $a \neq a'$. Consequently,

$$d^{(k)}(a, a') \geqslant 2|\mathbb{P}_a(E_{a,a'}) - \mathbb{P}_{a'}(E_{a,a'})| \geqslant 2(1 - \varepsilon_1 - \varepsilon_2).$$

Part (ii). Let $R : U^k \rightarrow \{a, *\}$ be defined as follows:

$$R(\mathbf{u}) = \begin{cases} a & \text{if } \psi(a, \mathbf{u}) = \text{true}, \\ * & \text{if } \psi(a, \mathbf{u}) = \text{false}. \end{cases}$$

Then, $\rho_a^R = \mathbb{P}(\psi(a, (X_a^1, \ldots, X_a^k)) = \text{true}) \geqslant 1 - \varepsilon_1 \geqslant 1 - \varepsilon$. Moreover

$$\rho_*^R = 1 - \mathbb{P}(\psi(a, (X_Y^1, \ldots, X_Y^k)) = \text{true})$$

and

$$\mathbb{P}(\psi(a, (X_Y^1, \ldots, X_Y^k)) = \text{true}) = \sum_{y \in A'} \mathbb{P}(\psi(a, (X_y^1, \ldots, X_y^k)) = \text{true}) \cdot \frac{1}{|A'|}.$$

By assumption, each term in the sum is $\leqslant \varepsilon_2$ and so $\mathbb{P}(\psi(a, (X_Y^1, \ldots, X_Y^k)) = \text{true}) \leqslant \varepsilon_2$. Thus, $\rho_*^R \geqslant 1 - \varepsilon_2 \geqslant 1 - \varepsilon$, as required. By Lemma 2.1 there is a testing procedure for X' that has accuracy at least $1 - \varepsilon$ and so the second claim in Part (ii) now follows by Part (i). $\square$

3. Testing versus reconstruction in phylogenetics

A *phylogenetic* (X-) *tree* is a tree $\mathcal{T}$, whose leaf set X is labelled and whose interior vertices are unlabelled and of degree at least 3. If, in addition, every interior vertex of $\mathcal{T}$ has degree exactly 3 then $\mathcal{T}$ is said to be *binary*. Without loss of generality, we can usually take $X = \{1, \ldots, n\}$.

In this section, we specialize, letting $A = A_n$ be the set of binary phylogenetic trees on leaf set $\{1, \ldots, n\}$, and letting $U = U_n$ be the set of site patterns on the leaf set $\{1, \ldots, n\}$ generated under a Markov process on the tree (this will be defined precisely below). Let $k_t(n) = k_t^{\varepsilon_1, \varepsilon_2}(n)$ denote the sequence length required to test a phylogenetic tree with accuracy $1 - \varepsilon_1$, $1 - \varepsilon_2$ and let $k_r(n) = k_r^\varepsilon(n)$ be the sequence length required to reconstruct a tree with accuracy $1 - \varepsilon_1$, $1 - \varepsilon_2$ (under the same model). We will be interested in the rate at which these quantities grow with n rather than their dependence on the ε values. Moreover, we can provide a lower bound on sequence length growth (Theorem 3.1, Eq. (6)) that is independent of ε_1, ε_2 (provided $\varepsilon = \max\{\varepsilon_1, \varepsilon_2\} < \frac{1}{2}$) while our upper bound (Theorem 3.2) involves ε only in a term that is a multiplicative factor.

For one class of models (finite-state Markov processes with an irreducible rate matrix), we will show (Theorem 3.1) that $k_t(n)$ grows at least logarithmically with n in the simply parametric setting, where each tree has a fixed set of edge lengths. It had already been established earlier that $k_r(n)$ grows at least logarithmically for very general models of sequence evolution (and in some more restricted models and parameter sets, grows at least polynomially) (Daskalakis et al., 2006; Erdös et al., 1999). Thus, in the region where $k_r(n)$ grows at the rate $\log(n)$, $k_t(n)$ grows at the same rate with n as $k_r(n)$.

However, for a closely related Markov process—the 'random cluster' (or 'infinite allele') model which can be used to model rare genomic events, the situation is surprisingly different in one respect. Although reconstruction still requires at least logarithmic (and in a certain range polynomial) number of samples (Mossel and Steel, 2004), testing with strong accuracy of $1 - \varepsilon$ can be achieved with $O(1)$ samples (Theorem 3.2). Moreover, this holds in the more general (non-simply parameterized) setting where nuisance parameters associated with the model are free to vary within constraints. Thus, for this model $k_t(n)$ is constant while $k_r(n)$ grows, so the two growth rates are fundamentally different, in contrast to the finite-state setting described in the previous paragraph.

We will now describe these results, beginning with the finite-state model.

3.1. Testing for a finite-state Markov process requires at least $\log(n)$ length sequences

Finite-state Markov processes on trees underlie many approaches in molecular phylogenetics (see, for example, Felsenstein, 2003). We provide a brief formal description; for more details, the reader may wish to consult Felsenstein (2003) or Semple and Steel (2003).

A finite-state Markov process is a continuous-time Markov process whose state space is some finite set; we will denote the rate (intensity) matrix of this process by S. We assume that S forms a reversible Markov process and we let π denote the equilibrium distribution on the states, determined by S (for background on reversible Markov processes readers may consult Grimmett and Stirzaker, 2001).

Now, suppose we have a phylogenetic X-tree $\mathcal{T}$ for which each edge e has some strictly positive valued 'length' $l(e)$. In this way, we can define a Markov process on $\mathcal{T}$ as follows (cf. Fig. 1(a)). To some vertex w, we assign states according to the distribution π. Then assign states to the remaining vertices of the tree by orienting the edges of $\mathcal{T}$ away from w; for each arc (u, v) for which u has been assigned a state s, assign to v the state obtained by applying the continuous-time Markov process with initial state s for duration $l(e)$. Thus the transition matrix associated to edge e is $\exp(Sl(e))$ and the joint probability distribution on the vertices of $\mathcal{T}$ is independent of the choice of the initial vertex w (by the reversibility assumption). Such a model induces a marginal distribution on the set of site patterns—assignments of states to the elements of X (the leaves of $\mathcal{T}$)—and this constitutes a single sample of the process (the site pattern is a random variable parameterized by the pair $(\mathcal{T}, l)$, where l assigns the lengths to the edges of $\mathcal{T}$).

The main result of this section is the following. Recall that a rate matrix is *irreducible* if the probability of a transition from any one state to any other state in time $\delta > 0$ is strictly positive.

Theorem 3.1. *Suppose we have a finite-state Markov process, with an irreducible rate matrix, on binary phylogenetic trees with leaf set* $\{1, \ldots, n\}$. *Suppose that we generate k i.i.d. site patterns. Then any testing procedure that has accuracy $1 - \varepsilon_1$, $1 - \varepsilon_2$ (with ε_1, $\varepsilon_2 < \frac{1}{2}$) requires k to grow at least at the rate $\log(n)$ in the simply parameterized setting where all edge lengths are equal to a fixed strictly positive value.*

Remark. Theorem 3.1 should be viewed alongside the result of Daskalakis et al. (2006), which shows that tree reconstruction (under the 2-state symmetric Markov model) with high accuracy
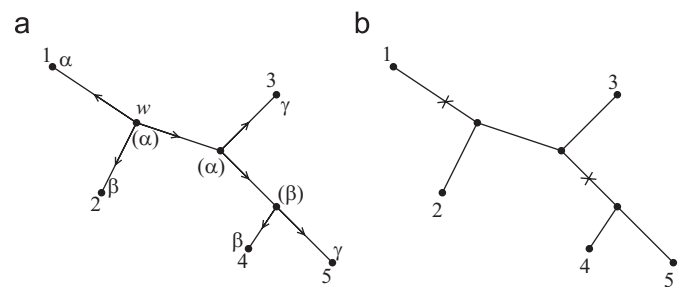

Fig. 1. (a) In a finite-state Markov process, a random state (α) at some vertex w evolves along the arcs of the tree (directed away from w). This gives rise to states at the internal vertices of the tree (in this example (α), (α), (β)) and a site pattern at the leaves (in this case α, β, γ, γ where the ordering corresponds to leaf order 1, 2, 3, 4, 5). (b) In the random cluster model each edge e is independently cut (indicated by $\times$) with an associated probability $p(e)$. In this example two edges were cut resulting in the leaf partition (character) of $\{\{1\}, \{2, 3\}, \{4, 5\}\}$.

is possible for sequences of length order $\log(n)$, even when the edge lengths are not known but constrained to lie within a fixed interval $[f, g]$ for any $f > 0$ and when g is sufficiently small.

Proof of Theorem 3.1. Let a be the tree shown in Fig. 2, and with all the edge lengths in this tree set equal to any fixed positive value. With a view to applying Lemma 2.3, let A' denote the set of m trees consisting of precisely those trees obtained from a by interchanging the leaf labels $3i - 1$ and $3i$ for one value of $i = 1, \ldots, m$ (and keeping the edge lengths fixed). Suppose there exists a testing process of accuracy $1 - \varepsilon_1, 1 - \varepsilon_2$ for phylogenetic trees in this simply parameterized setting. Then, by Lemma 2.3(ii) we have

$$d^{(k)}(a, *) \geqslant 2(1 - 2\varepsilon), \tag{3}$$

where $\varepsilon = \max\{\varepsilon_1, \varepsilon_2\} < \frac{1}{2}$. Suppose we generate k sites i.i.d. under parameter a or $*$ (in the case of $*$, we select the random element of A' and then generate k sites i.i.d. using that element). Let C be the random vector variable that specifies the sequences (over the state space), each of length k, occurring at the unlabelled internal vertices of the model tree a in Fig. 2. Since this tree has $3m$ internal vertices, there will be $3m$ such sequences described by C. Note that C has the same probability distribution for any element b in A' as it does for a, and so, in particular, we have

$$\mathbb{P}_a(C = c) = \mathbb{P}_*(C = c)$$

for all choices of c. Thus

$$
\begin{aligned}
d^{(k)}(a, *) &= \sum_{\mathbf{u} \in U^k} |\mathbb{P}_a(\mathbf{u}) - \mathbb{P}_*(\mathbf{u})| \\
&= \sum_{\mathbf{u}} \left| \sum_c (\mathbb{P}_a(\mathbf{u}|c) - \mathbb{P}_*(\mathbf{u}|c)) \cdot \mathbb{P}_a(C = c) \right|
\end{aligned}
$$

and so

$$
\begin{aligned}
d^{(k)}(a, *) &\leqslant \left(\max_c \sum_{\mathbf{u}} |\mathbb{P}_a(\mathbf{u}|c) - \mathbb{P}_*(\mathbf{u}|c)| \right) \cdot \sum_c \mathbb{P}_a(C = c) \\
&= \max_c \sum_{\mathbf{u}} |\mathbb{P}_a(\mathbf{u}|c) - \mathbb{P}_*(\mathbf{u}|c)|.
\end{aligned}
\tag{4}
$$

We will establish the following crucial inequality: for any c:

$$\sum_{\mathbf{u} \in U^k} |\mathbb{P}_a(\mathbf{u}|c) - \mathbb{P}_*(\mathbf{u}|c)| < \tau^k / \sqrt{m} \tag{5}$$

for a constant $\tau > 1$. Then, combining (3)–(5) gives

$$\tau^k / \sqrt{m} > 2(1 - 2\varepsilon)$$

and so (since $n = 3m + 2$),

$$k \geqslant \frac{\frac{1}{2} \log\left(\frac{n-2}{3}\right) + \log(2(1 - 2\varepsilon))}{\log(\tau)} \geqslant \frac{\frac{1}{2} \log\left(\frac{n-2}{3}\right)}{\log(\tau)} \tag{6}$$

that is, k must grow at least logarithmically with n (for any $\varepsilon < 1/2$) as claimed by the theorem.
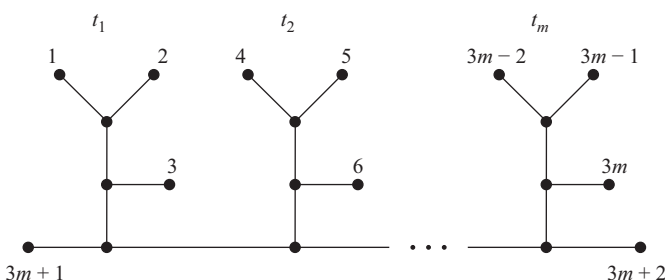

Fig. 2. The generating tree in the proof of Theorem 3.1.

Thus, to establish the theorem, it suffices to justify Inequality (5).

For any c maximizing $\sum_{\mathbf{u}} |\mathbb{P}_a(\mathbf{u}|c) - \mathbb{P}_*(\mathbf{u}|c)|$, let us denote $\mathbb{P}_x(\mathbf{u}|c)$ by $Q_x(\mathbf{u})$ for $x = a, *$ and $b \in A'$. Then

$$\sum_{\mathbf{u}} |Q_a(\mathbf{u}) - Q_*(\mathbf{u})| = \sum_{\mathbf{u}} \left| Q_a(\mathbf{u}) - \frac{1}{m} \sum_{b \in A'} Q_b(\mathbf{u}) \right|. \tag{7}$$

We can rewrite the expression on the right-hand side of (7) as

$$\sum_{\mathbf{u}} \left| Q_a(\mathbf{u}) \left(1 - \frac{1}{m} \sum_{b \in A'} \frac{Q_b(\mathbf{u})}{Q_a(\mathbf{u})} \right) \right| \leqslant \sum_{\mathbf{u}} Q_a(\mathbf{u}) \cdot \left| \frac{1}{m} \sum_{b \in A'} \left(1 - \frac{Q_b(\mathbf{u})}{Q_a(\mathbf{u})} \right) \right|. \tag{8}$$

In particular, consider the following random variable:

$$Z_b := 1 - \frac{Q_b(\xi)}{Q_a(\xi)}$$

for each $b \in A'$, where ξ is the random element of U^k generated by the probability distribution Q_a. Then the expression on the right-hand side of (8) can be rewritten as

$$\mathbb{E}\left(\left| \frac{1}{m} \sum_{b \in A'} Z_b \right| \right),$$

where expectation is taken with respect to the probability distribution Q_a on U^k. Now, for all $b \in A'$ we have

$$\mathbb{E}(Z_b) = \sum_{\mathbf{u} \in U^k} Q_a(\mathbf{u}) \cdot \left(1 - \frac{Q_b(\mathbf{u})}{Q_a(\mathbf{u})} \right) = \sum_{\mathbf{u} \in U^k} (Q_a(\mathbf{u}) - Q_b(\mathbf{u})) = 1 - 1 = 0,$$

and the Z_b are independent random variables (since we have conditioned on a particular value $C = c$, and the Markov property). Thus, by Jensen's inequality:

$$\left(\mathbb{E}\left(\left| \frac{1}{m} \sum_{b \in A'} Z_b \right| \right) \right)^2 \leqslant \mathbb{E}\left(\left(\frac{1}{m} \sum_{b \in A'} Z_b \right)^2 \right) = \frac{1}{m^2} \sum_{b \in A'} \mathbb{E}(Z_b^2). \tag{9}$$

Now, since $\mathbb{E}(Z_b) = 0$, we have $1 + \mathbb{E}(Z_b^2) = \mathbb{E}((Z_b + 1)^2) = \mathbb{E}(Q_b(\xi)^2 / Q_a(\xi)^2)$ and so

$$\mathbb{E}(Z_b^2) = -1 + \mathbb{E}\left(\frac{Q_b(\xi)^2}{Q_a(\xi)^2} \right). \tag{10}$$

Now, let us write $\mathbf{u} = (u_1, \ldots, u_k)$ and $c = (c_1, \ldots, c_k)$, where u_i is the assignment of states to the leaf set $\{1, \ldots, n\}$ by the i-th generated site pattern, and c_i is the assignment of states to the $3m$ interior vertices by the process that generated the i-th site pattern. By the Markov property, the ratio $\mathbb{P}_b(u_i|c_i)/\mathbb{P}_a(u_i|c_i)$ consists of a product of terms in the numerator and denominator (indexed over the interior vertices of the trees b and a, respectively) and each numerator term cancels a corresponding denominator term, except for two terms that correspond to the two pendant edges incident with the leaves $3i - 1$ and $3i$ whose interchange converts tree a into tree b. Consequently, by the irreducibility condition on the Markov process, the following inequality holds:

$$\frac{\mathbb{P}_b(u_i|c_i)}{\mathbb{P}_a(u_i|c_i)} \leqslant \tau \tag{11}$$

for some absolute constant $\tau > 1$ dependent only on the rate matrix S and the (equal) value of the common edge length. Thus, by independence, $Q_b(\mathbf{u})/Q_a(\mathbf{u}) \leqslant \tau^k$, and so (10) gives

$$\mathbb{E}(Z_b^2) \leqslant -1 + \tau^{2k} < \tau^{2k}.$$

Consequently, by (9)

$$\mathbb{E}\left(\left|\frac{1}{m}\sum_{b\in A'}Z_b\right|\right)\leqslant\sqrt{\frac{1}{m^2}|A'|\tau^{2k}}<\tau^k/\sqrt{m},$$

and so (by (7) and (8)) we have verified Inequality (5), and thereby completed the proof. □

3.2. An O(1) test for the random cluster model

In this section, a *character* (on X) denotes an arbitrary partition $\{\alpha_1,\ldots,\alpha_m\}$ of X into any number of disjoint subsets.

In the *random cluster model*, one has a phylogenetic X-tree $\mathcal{T}$ and each edge e has an associated probability $p(e)$ that the edge of $\mathcal{T}$ is cut. These cuts are performed independently across the tree, resulting in a (generally disconnected) graph and the leaves in each connected component form the blocks of a resulting random character on X. Thus $\mathcal{T}$ along with the $p(e)$ values (called 'substitution probabilities') provide a well-defined probability distribution on characters on X (see Fig. 1(b)).

This is the same distribution on characters as one obtains in the limit as $s\to\infty$ from a finite-state Markov process on $\mathcal{T}$ that has an $s\times s$ rate matrix in which all its off-diagonal entries are equal, and where one considers the character on X whose blocks are the sets of leaves of the same state. Notice, however, that if we apply Inequality (6) that furnishes the $\log(n)$ lower bound growth rate for k in Theorem 3.1 to this particular family of Markov processes, we obtain a lower bound on k that converges to a constant as the number of states s goes to infinity sufficiently fast with n (since for a $s\times s$ rate matrix with all its off-diagonal entries equal, τ grows at the rate s^2). This opens the possibility of testing a tree under the random cluster model with $O(1)$ characters, and we will see that this is indeed possible.

We can view the random cluster model as a type of infinite-state Markov process, and it is known in population genetics as the *infinitely many alleles model* (Kimura and Crow, 1964). The model is relevant for describing evolution in settings where transitions generally lead to states that have not appeared elsewhere in the tree, such as with gene order re-arrangement, or other rare genomic events (as described in Steel and Penny, 2005, Section 9.8, see also Rokas and Holland, 2000; Semple and Steel, 2002); the model has also been investigated for modelling language evolution (Warnow et al., 2006).

Given a character χ on X and a phylogenetic X-tree $\mathcal{T}$, let $\mathcal{T}_\alpha$ denote the minimal subtree of $\mathcal{T}$ connecting the leaves of $\mathcal{T}$ in block α. Then χ is said to be *homoplasy-free* on $\mathcal{T}$ if the collection of trees $\mathcal{T}_\alpha:\alpha\in\chi$ is vertex-disjoint. Given a sequence $\mathcal{C}=(\chi_1,\chi_2,\ldots,\chi_k)$ of characters, consider the following deterministic testing process ψ_H on phylogenetic trees:

$$\psi_H(\mathcal{C},\mathcal{T})=\begin{cases}\text{true} & \text{if }\chi_i\text{ is homoplasy-free on }\mathcal{T}\text{ for }i=1,\ldots,k,\\ \text{false} & \text{otherwise.}\end{cases}$$

Theorem 3.2. *Under the random cluster model on binary phylogenetic trees with leaf set $\{1,\ldots,n\}$, suppose that we generate k i.i.d. characters, where the substitution probability $p(e)$ on any edge e lies in some fixed interval $[a,b]$ where $0<a\leqslant b<\frac{1}{2}$. Then ψ_H is a testing process with strong accuracy of $1-\varepsilon$ whenever the number of characters is at least k, where*

$$k=\gamma^{-1}\cdot\log(\varepsilon^{-1}),$$

and where $\gamma=a\cdot(1-2b)^4/(1-b)^4$. This holds for all values of n.

To prove Theorem 3.2 we first require a combinatorial lemma.

Lemma 3.3. *Suppose $\mathcal{T}_1$, $\mathcal{T}_2$ are two binary phylogenetic X-trees, with $\mathcal{T}_1\neq\mathcal{T}_2$, and $|X|>3$. There exist induced rooted phylogenetic*

subtrees $\mathcal{T}_A$, $\mathcal{T}_B$, on leaf sets A and B, respectively, where A, B are disjoint, nonempty subsets of X, such that:

(i) $\mathcal{T}_A$ and $\mathcal{T}_B$ are present as pendant subtrees in $\mathcal{T}_1$ and $\mathcal{T}_2$; and
(ii) the root of $\mathcal{T}_A$ and of $\mathcal{T}_B$ are adjacent to a common vertex in $\mathcal{T}_2$ but not in $\mathcal{T}_1$.

Proof. The proof is by induction on $|X|$. For $|X|=4$, the result is easily seen to hold. Suppose the lemma holds for $|X|=n-1$ where $n\geqslant 5$, and that $|X|=n$. If $\mathcal{T}_2$ has a cherry (a pair of elements $\{x,y\}$ of X that label leaves that are adjacent to a common vertex) that is not also a cherry of $\mathcal{T}_1$, then we can take $A=\{x\}$, $B=\{y\}$ and the claim in the theorem holds. Otherwise, every cherry of $\mathcal{T}_2$ is a cherry of $\mathcal{T}_1$ and since $\mathcal{T}_2$ has at least one cherry (say $\{x,y\}$) we may replace $\mathcal{T}_1$, $\mathcal{T}_2$ with the pair of trees $\mathcal{T}'_1$, $\mathcal{T}'_2$ obtained by deleting the leaves labelled $\{x,y\}$ and their incident edges from each tree, and assigning each newly created leaf vertex the label $v_{x,y}$. Note that $\mathcal{T}'_1$, $\mathcal{T}'_2$ are binary phylogenetic trees, with a leaf label set $X-\{x,y\}\cup\{v_{x,y}\}$, and that $\mathcal{T}'_1\neq\mathcal{T}'_2$ (otherwise it is easily seen that $\mathcal{T}_1=\mathcal{T}_2$). Thus we may apply the induction hypothesis to the pair $\mathcal{T}'_1$, $\mathcal{T}'_2$. Given the sets A, B for this pair that meet the requirements stated in the lemma, we then replace any occurrence of $v_{x,y}$ in A or B by the elements x, y—the resulting pair of sets now satisfies the requirements stated in the lemma. This completes the proof. □

Proof of Theorem 3.2. Suppose that $\mathcal{T}_1$ is the binary phylogenetic X-tree that generates the characters. Then $\psi_H(\mathcal{F},\mathcal{T}_1)=$ 'true' with probability 1, since the event χ is homoplasy-free on $\mathcal{T}_1$ has probability 1 for any character χ that evolves on $\mathcal{T}_1$ under the random cluster model. Now, suppose that $\mathcal{T}_2$ is a binary phylogenetic X-tree that is different from $\mathcal{T}_1$. By Lemma 3.3, $\mathcal{T}_2$ and $\mathcal{T}_1$ both share the same pair of pendant subtrees $\mathcal{T}_A$ and $\mathcal{T}_B$ for which the roots are adjacent in $\mathcal{T}_2$ but not in $\mathcal{T}_1$, as illustrated in Fig. 3. Now consider the evolution of one of the characters χ_i under the random cluster model on $\mathcal{T}_1$. Let α, β
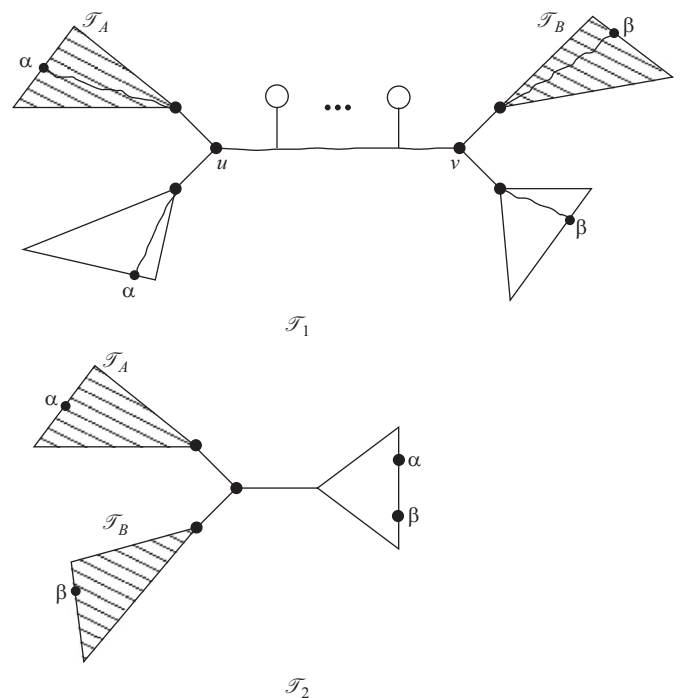

Fig. 3. Canonical representation of two different binary phylogenetic trees in the proof of Theorem 3.2.

denote, respectively, the blocks of the character at vertices u, v of $\mathscr{T}_1$, and consider the conjunctive event $E = \bigcap_{i=1}^{5} E_i$ in which:

(E_1) $\alpha \neq \beta$;
(E_2) at least one leaf of $\mathscr{T}_A$ is present in α;
(E_3) at least one leaf of the other subtree of $\mathscr{T}_1$, which is incident with u and does not contain v, is present in α;
(E_4) at least one leaf of $\mathscr{T}_B$ is present in β; and
(E_5) at least one leaf of the other subtree of $\mathscr{T}_1$, which is incident with v and does not contain u, is present in β.

Under the random cluster model, these five events are independent (by the assumption that the cuts on edges are made independently) and so

$$\mathbb{P}(E) = \mathbb{P}\left(\bigcap_{i=1}^{5} E_i\right) = \prod_{i=1}^{5} \mathbb{P}(E_i) \geqslant \gamma,$$

since $\mathbb{P}(E_1) \geqslant a$, and, by Lemma 2.1 of Mossel and Steel (2004), $\mathbb{P}(E_i) \geqslant (1 - 2b)/(1 - b)$ for $i \in \{2, \ldots, 5\}$. Now, $\psi_H(\mathscr{F}, \mathscr{T}_2) = $ 'false' whenever event E occurs. If we evolve k independent characters under the assumptions stated in the theorem, then the probability that E occurs at least once is at least $1 - (1 - \gamma)^k$, and this probability is at least $1 - \varepsilon$ when $k \geqslant \gamma^{-1} \log(\varepsilon^{-1})$, by virtue of the inequality $1 - (1 - x)^y \geqslant 1 - e^{-xy}$. This completes the proof. □

4. Concluding comments

The reader may be curious as to where our proof for the $\log(n)$ lower bound on sequence length for testing under the finite-state model breaks down for the random cluster model. The crucial distinction is that the random cluster model fails to satisfy condition (11) required in the proof for Theorem 3.1. That is, in the random cluster model, some characters have a positive probability on some trees but have zero probability on other trees. Indeed it has been shown that, for any binary phylogenetic tree $\mathscr{T}$ with n leaves, there is a set of just *four* characters such that $\mathscr{T}$ is the only tree for which these characters have strictly positive probability Huber et al. (2005). Thus, in contrast to finite-state models, under the random cluster model each tree can be reconstructed from $O(1)$ characters (using, say maximum likelihood estimation), provided these characters are carefully selected. However, if the characters randomly evolve under the random cluster model (rather than being carefully chosen) then, as mentioned earlier, we are likely to need at least $\log(n)$ such characters for accurate tree reconstruction (Mossel and Steel, 2004).

Note also that Theorem 3.2 can be extended to a setting in which the substitution probabilities vary from character to character, provided they all lie in some fixed interval $[a, b]$ where $0 < a \leqslant b < \frac{1}{2}$. If we generate k characters independently (but not necessarily identically) in this more general setting, testing the true tree using ψ_H will return 'true' with probability 1, while testing any other tree will return 'false' with probability at least $1 - \varepsilon$ provided k satisfies the lower bound described in Theorem 3.2.

Although testing for the finite-state Markov model can require the same $\log(n)$ growth rate in sequence length required for reconstructing, there is a related task where $O(1)$ sequence length suffices. This is for *teasing a tree*, where one is given sequences of length k and a set of two trees—one of which is the tree that generated the data, and one is asked to identify which of the two trees generated the data. For the symmetric 2-state Markov process and under suitable restrictions on the substitution probabilities, it was shown in Steel and Székely (2007) that sequences of length $O(1)$ (i.e. independent of n) suffice to correctly

solve (with high probability) the teasing problem on binary phylogenetic trees with n leaves.

Finally, we have considered reconstruction only in the part of the parameter range (on the substitution probabilities on the edges of the tree) where reconstruction requires logarithmic length sequences. Outside of this region, it is known that polynomial-length sequences can be required, both for the finite state Markov model (Daskalakis et al., 2006) and for the random cluster model (Mossel and Steel, 2004). It may be of interest in future work to determine the sequence length required for testing in these portions of parameter space.

References

Alon, N., Spencer, J., 2008. The Probabilistic Method. Wiley-Interscience, New York.
Berger, A., Wald, A., 1949. On distinct hypotheses. Ann. Math. Stat. 20 (1), 104–109.
Castella, G., Berger, R.L., 2002. Statistical Inference, second ed. Duxbury, Pacific Grove, CA.
Churchill, G., von Haeseler, A., Navidi, W., 1992. Sample size for a phylogenetic inference. Mol. Biol. Evol. 9 (4), 753–769.
Daskalakis, C., Mossel, E., Roch, S., 2006. Optimal phylogenetic reconstruction. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006), pp. 159–168.
Erdös, P.L., Steel, M.A., Székely, L.A., Warnow, T., 1999. A few logs suffice to build (almost) all trees (part 1). Random Struct. Algorithms 14 (2), 153–184.
Felsenstein, J., 2003. Inferring Phylogenies. Sinauer Press.
Garey, M.R., Johnson, D.S., 1979. Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman and Company, New York.
Grimmett, G., Stirzaker, D., 2001. Probability and Random Processes, third ed. Oxford University Press, USA.
Huber, K., Moulton, V., Steel, M., 2005. Four characters suffice to convexly define a phylogenetic tree. SIAM J. Discrete Math. 18 (4), 835–843.
Kimura, M., Crow, J., 1964. The number of alleles that can be maintained in a finite population. Genetics 49, 725–738.
Lecointre, G., Philippe, H., Van Le, H.L., Le Guyader, H., 1994. How many nucleotides are required to resolve a phylogenetic problem? The use of a new statistical method applicable to available sequences. Mol. Phyl. Evol. 3 (4), 292–309.
Mossel, E., Steel, M., 2004. A phase transition for a random cluster model on phylogenetic trees. Math. Biosci. 187, 189–203.
Mossel, E., Steel, M., 2005. How much can evolved characters tell us about the tree that generated them? In: Gascuel, O. (Ed.), Mathematics of Evolution and Phylogeny. Oxford University Press, Oxford, pp. 384–412.
Rokas, A., Holland, P.W.H., 2000. Rare genomic changes as a tool for phylogenetics. Trends Ecol. Evol. 15 (11), 454–459.
Saitou, N., Nei, M., 1986. The number of nucleotides required to determine the branching order of three species, with special reference to the human-chimpanzee-gorilla divergence. J. Mol. Evol. 24, 189–204.
Semple, C., Steel, M., 2003. Phylogenetics. Oxford University Press, Oxford.
Semple, C., Steel, M., 2002. Tree reconstruction from multi-state characters. Adv. Appl. Math. 28 (2), 169–184.
Steel, M., Penny, D., 2005. Maximum parsimony and the phylogenetic information in multi-state characters. In: Albert, V. (Ed.), Parsimony, Phylogeny and Genomics. Oxford University Press, Oxford, pp. 163–178.
Steel, M.A., Székely, L.A., 1999. Inverting random functions. Ann. Comb. 3, 103–113.
Steel, M.A., Székely, L.A., 2002. Inverting random functions II: explicit bounds for the discrete maximum likelihood estimation, with applications. SIAM J. Discrete Math. 15 (4), 562–575.
Steel, M.A., Székely, L.A., 2009. Inverting random functions III: discrete maximum likelihood revisited. Ann. Comb., in press (preprint arXiv:math/0608273).

Steel, M.A., Székely, L.A., 2007. Teasing apart two trees. Comb. Probab. Comput. 16, 903–922.

Warnow, T., Evans, S.N., Ringe, D., Nakhleh, L., 2006. A stochastic model of language evolution that incorporates homoplasy and borrowing. In: Phylogenetic Methods and the Prehistory of Languages, McDonald Institute for Archaeological Research, pp. 75–.

Wortley, A.H., Rudall, P.J., Harris, D.J., Scotland, R.W., 2005. How much data are needed to resolve a difficult phylogeny? Case study in Lamiales. Syst. Biol. 54 (5), 696–709.