Can We Avoid "SIN" in the House of "No Common Mechanism"?

MIKE STEEL*

*Allan Wilson Centre for Molecular Ecology and Evolution, Biomathematics Research Centre, University of Canterbury,
Private Bag 4800, Christchurch, New Zealand;*
**Correspondence to be sent to: Allan Wilson Centre for Molecular Ecology and Evolution, Biomathematics Research Centre,
University of Canterbury, Private Bag 4800, Christchurch, New Zealand; E-mail: m.steel@math.canterbury.ac.nz.*

In "no common mechanism" (NCM) models of character evolution, each character can evolve on a phylogenetic tree under a partially or totally separate process (e.g., with its own branch lengths). In such cases, the usual conditions that suffice to establish the statistical consistency of tree reconstruction by methods such as maximum likelihood (ML) break down, suggesting that such methods may be prone to statistical inconsistency (SIN). In this paper I ask whether we can avoid SIN for tree topology reconstruction when adopting such models either by using ML or by any other method that could be devised. I show that it is possible to avoid SIN for certain NCM models, but not for others, and the results depend delicately on the tree reconstruction method employed. I also describe the biological relevance of some recent mathematical results for the more usual "common mechanism (CM)" setting. The results are not intended to justify NCM rather to set in place a framework within which such questions can be formally addressed.

SIN in phylogenetics is the tendency of certain tree reconstruction methods to fail to converge on the correct tree topology when applied to increasing quantities of data that evolve under a given model. The phenomenon has been well known for simple methods like maximum parsimony (MP) since the landmark paper of Felsenstein (1978) three decades ago. SIN has contributed to the widespread acceptance of more sophisticated tree reconstruction methods such as ML, corrected distance methods, and Bayesian phylogenetics (Felsenstein 2004; Lemey et al. 2009). These methods are based explicitly on stochastic models of character evolution and for which it is usually possible to establish statistical consistency when the model assumed by the investigator is also the one that generated the data (see, e.g., Chang 1996; Allman and Rhodes 2006; Sober 2008).

A centerpiece of nearly all these models is the assumption that character data (for instance, genetic sequence sites) evolve independently and identically. This "i.i.d." assumption is standard in statistics and implies that each character is described by essentially the same process and that the characters represent a finite random sample of this process. The i.i.d. assumption applies even for mainstream models that allow a distribution

of rates across sites, such as the frequently used "Γ + I" embellishment of the general time reversible (GTR) model. In these models, it is usually assumed that the rate at a site is chosen i.i.d. from a given distribution. Such a "rates across sites" model is subtly different from a "variable site rate" model that assumes that each site has its own particular intrinsic rate (i.e., not chosen i.i.d. from some distribution). Within such a model, the sequence sites may still be independently generated, but they are not identically distributed. In this paper, we will refer to any model in which sites evolves on a fixed tree but where some of the parameters (e.g., branch lengths or site rates) vary arbitrarily from site to site as an NCM model.

If we just consider the frequencies of site patterns, then the two models (rates across sites and variable site rate) can produce (almost) identical data; however, significant differences between the models can become apparent when we come to do tree reconstruction from given sequences. For example, in a ML approach to tree reconstruction, in which one explicitly assumes the variable site rate model, we may wish to estimate a corresponding rate for each site that maximizes the probability of observing the given site pattern—along with a shared underlying set of branch lengths common to all the sites (such an approach was described by Gary Olsen in Swofford et al. 1996). Each rate estimate—one for each site—might later be discarded as a "nuisance parameter" in the search for the underlying tree topology alone (this approach is quite different to doing the "usual" form of ML estimation of a tree topology under a rates-across-sites model). We can ask if such an approach is statistically sound—in particular, can it lead to SIN? What if one allows the branch lengths also to vary from character to character (the more usual form of NCM)? Is ML under this model liable to SIN; if so, can any method reconstruct a tree under this model without SIN? These are the sort of questions I will address. I will also describe the biological relevance of some recent mathematical results concerning tree reconstruction in the more usual CM setting.

First, I outline some of the motivations and concerns surrounding NCM models in phylogenetics. I then discuss statistical consistency in a general setting—first for CM models, where much is known, then for NCM models, where there has been little analysis to date in phylogenetics. In Theorem 1, I present some results in this area and show how the details of the model (and the method) are crucial to whether we are in danger of SIN when working with an NCM model. I also describe different forms of SIN and attempts to measure and manage it. The paper ends with a brief discussion.

Some Reasons for and against NCM

The idea that the evolution of characters in biology might be described by different sets of branch lengths underlies recent attempts to deal with phenomena such as heterotachy (Gaucher and Miyamoto 2005; Phillipé

et al. 2005). However, the idea dates back to the early days of molecular phylogenetics. It is implicit in Walter Fitch's discussion of a covarion model (Fitch 1971) and was discussed more explicitly by Cavender (1981) in reference to his simple two-state Poisson model. In response to the question of whether the probabilities of change should be the same for all characters, Cavender (p. 222) remarked:

> This assumption can and should be removed. It is unacceptable biologically because it says, for example, that an insect species is just as likely to lose (or acquire) wings as a spot of color.

This comment seems reasonable for morphological characters, though even in that setting, we might still expect some correlation in the relative probabilities of character change on a given branch across characters, as it may be more likely to observe changes on branches that correspond to long time intervals between speciation. It is less clear that Cavender's comment should apply to aligned DNA sequence sites, each of which we might view as a random sample from a common process. Nonetheless, different DNA sequence sites may be subject to differing selection pressures and the probability that a site mutation becomes fixed in a population may depend on structural or functional constraints; for example, whether the protein a gene codes for still folds correctly if the substitution changes an amino acid. These constraints may vary with time and across the sequence. So, enforcing an entirely CM model may be too severe. Similar comments apply to other types of genomic data that carry evolutionary signal. In linguistics, a model that allows each character to have its own branch lengths has also been developed for studying language evolution (Warnow et al. 2006).

An additional reason why the NCM approach has received further attention is its relevance to those in the systematics community who advocate the use of maximum parsimony for phylogeny reconstruction (e.g., Farris 2008). This has been justified by an equivalence theorem that demonstrates that MP is the ML estimator of a tree under an NCM model based on a symmetric Poisson process such as the Jukes–Cantor model (Tuffley and Steel 1997). A slightly more streamlined proof of this result has recently been given by Fischer and Thatte (2010), and extensions of this equivalence theorem were described in Steel and Penny (2004, 2005) and, most recently, Fischer and Thatte (2010). This last paper also showed that the original equivalence theorem breaks down if one modifies the Poisson model slightly; either 1) by imposing a molecular clock or 2) by setting an absolute upper bound on the branch lengths.

The significance and implications of the equivalence between MP and ML estimation under NCM have aroused considerable interest (see, e.g., Sober 2004; Farris 2008; Huelsenbeck et al. 2008, 2011). One view is that NCM model is sufficiently general as to capture "truth" and so should be the model of choice, thereby providing a justification for MP (Farris 2008).

An alternative position is that NCM is far too parameter rich and it ignores likely correlations between branch lengths due to shared time frames of speciation intervals. The NCM model required for the formal equivalence between MP and ML under the NCM is also based on a symmetric model of substitution change (such as a Jukes–Cantor model) that reflects a process of drift rather than directional selection.

In the sections that follow our aim is not to defend NCM models, but rather to determine which methods, if any, would allow phylogenetic tree topology to be estimated in a statistically consistent manner were one to adopt various NCM models. We find some interesting contrasts between the NCM and CM settings—for instance, in the latter, ML is consistent if any method is, but in the NCM world, this is no longer true. Also, avoiding SIN in NCM models (by any method) seems to require a "fine balance" in the underlying model. This brings into question the robustness of any consistency results to even slight model misspecification and suggests that other statistical considerations (e.g., bias, efficiency) may override consistency issues.

ML Estimation in General and in Phylogenetics

In this section, I consider a general setting that includes phylogenetic tree reconstruction, and other problems where a discrete parameter (e.g., a tree, network, cluster) is being estimated from discrete data (e.g., DNA sequences, genes) in the presence of unknown additional parameters.

Suppose one has a sequence of observations $u_1, u_2, \ldots$ taking values in a finite set U (the elements of this set can be arbitrary, but I will call them "site patterns" as I will usually be considering aligned DNA sequence sites). Suppose that these observations are generated independently by a model M that has a fixed but unknown discrete parameter a that takes values in some finite set A, alongside other continuous parameters, which may vary from observation to observation. In the phylogenetic setting, A will generally refer to the set of fully resolved tree topologies on a given set of species, and the continuous parameters may refer to branch lengths or other aspects of the substitution model (site rate, transition/transversion ratio, shape parameter for a Γ distribution of rates across sites, etc.).

In all such cases, u_i is generated by a pair (a, θ_i) where θ_i lies in some set $\Theta(a)$, which will be assumed throughout to be an open subset of Euclidean space. For most of this paper, the pair (a, θ) will refer to a fully resolved tree topology (a) on n leaves, together with a collection of associated branch lengths (θ) that lie in the set $\Theta(a) = (0, \infty)^{2n-3}$. Thus, the branches are required to be strictly positive but finite (one may impose further conditions on these branch lengths, such as a molecular clock, or additional parameters, such as those that describe other aspects of the model).

CM and NCM Versions of a Model

In a CM version of a model M, which I will denote by CM-M, it is assumed that all the θ_i values are equal; that is, they take a common value, $\theta \in \Theta(a)$ across the sequence of observations (i.e., as i varies). By contrast, in an NCM version of M, which I will denote by NCM-M, the θ_i can take different values as i varies. Notice, however, that if these θ_i values are assigned randomly and independently from some common distribution (as is the case with most "rates across sites" models in phylogenetics), then this is just a CM version of a slightly more complex model M^* that is derived from M.

ML under CM and NCM

The ML estimation of a discrete parameter from A under an NCM version of M applied to data $(u_1, \ldots, u_k)$ selects the element $b \in A$ that maximizes

$$L(b|\text{data}) := \sup_{(\theta_1, \ldots, \theta_k) \in \Theta(b)^k} \mathbb{P}[\text{data}|b, (\theta_1, \ldots, \theta_k)]$$

$$= \prod_{i=1}^{k} \sup_{\theta_i \in \Theta(b)} \mathbb{P}[u_i|b, \theta_i], \tag{1}$$

where "sup" in equation (1) refers to supremum (the maximum value either obtained or as a limit). The second equality in equation (1) is justified by the assumption that the observations are independently generated by the model. For ML estimation under the CM version of M, the only difference is that the θ_i values are required to be identical (i.e., $\theta_i = \theta$ for all i).

Given two models M_1 and M_2 (usually, but not necessarily the same model), let us refer to "ML estimation under M_1 applied to M_2-data" as the ML estimation under model M_1 of the discrete parameter in A from data that has been produced under model M_2. Our interest is in determining when this method is statistically consistent (defined shortly) for various M_1, M_2, and if so, what can be said about the sequence length requirements for accurate estimation.

Given two models M_1 and M_2, write $M_1 \subseteq M_2$ if M_1 is a submodel of M_2, that is, M_1 is a special case of M_2 once constraints are placed on its parameters; in particular, for any model M, CM-$M \subseteq$ NCM-M. If M_2 is not contained in M_1, then ML estimation under M_1 applied to M_2-data is often said to be carried out under a "misspecified model"—in this case, one does not generally expect consistency so one is usually more interested in the regular case where the model in which ML is performed includes the model that generates the data, that is, either $M_1 = M_2$ or $M_2 \subseteq M_1$ (one exception occurs in Theorem 1(iv), which provides an instance where ML estimation under a CM model is statistically consistent even when it is applied to data generated under an NCM version of that model).

Basic Models for Character Evolution (N_r)

It will be convenient to describe most of our results for a particular model of character evolution. The simplest

such model assumes that the rates of substitution between each pair of the r character states are equal—this is sometimes referred to as the Neyman r-state model or the "symmetric r-state model"; here, I call it the N_r-model (after the r-state model introduced in Neyman 1971) and it is the same as the "Mk" model for $k = r$ (Lewis 2001). In the special case where $r = 4$, this is the familiar Jukes–Cantor model, whereas $r = 2$ is often referred to as the "Cavendar–Farris–Neyman model." In the N_r model, it is usually (but not always) assumed that the frequency of bases at the root of the tree is the uniform distribution.

I will also consider the limiting case of the N_r model as r becomes large (for a given number of species). This model, denoted here by N_∞, is sometimes called the "Kimura–Crow infinite alleles model" (Kimura and Crow 1964) or the "random cluster model" (Mossel and Steel 2004), and it models the setting where each substitution always results in a new state. Denote the CM and NCM versions of N_r model (r being either finite or infinite) by writing CM-N_r and NCM-N_r, respectively.

SIN for Data Generated under CM Models

In the CM model—either for generating the data or for carrying out ML—we require the θ_i values to all be equal to some common value (call it θ). Note that even if we are not at all interested in estimating the θ values, we often still have to consider their role in any probability calculations; in this case, they are said to be "nuisance parameters."

A method $\mathcal{M}$ for estimating the discrete parameter in A from a sequence of independently generated observations is "statistically consistent" for data generated under a CM model if, for each $a \in A$, and $\theta \in \Theta(a)$, the probability that $\mathcal{M}$ correctly estimates a from $(u_1, \ldots, u_k)$, when each observation u_i is generated independently by the model with parameters (a, θ), converges to 1 as k grows. If this condition fails, the method $\mathcal{M}$ leads to SIN. A related, but slightly different concept of statistical consistency exists, based on the strong (rather than the weak) law of large numbers, but we do not discuss it here.

Two types of SIN are possible in inferring phylogenetic tree topology. The more familiar and stronger form is when the method $\mathcal{M}$ can "positively mislead"—that is, the probability that the method estimates an incorrectly resolved tree converges to 1 as the sequence length grows; this is the type of inconsistency that occurs with MP in the "Felsenstein zone" (the "original SIN" established by Felsenstein 1978).

However, a milder, more venial form of SIN can occur in two ways: 1) if the method becomes unable to decide between the true tree and at least one alternative tree with increasing data or 2) if the method converges with increasing data on a nonresolved tree, of which the true tree is just one resolution. This latter possibility is precisely what can occur with "ancestral maximum likelihood" (AML). In a ML framework, this method optimizes not just the tree topology and its

branch lengths but also a particular set of ancestral sequences and then returns just the tree topology. Mossel et al. (2009) showed that this AML estimation of tree topology applied to CM-N_2 data can converge on the fully unresolved star tree, when the branch lengths of the fully resolved generating tree are in a certain range. Whether AML can lead to the stronger form of SIN of being positively misleading is currently an open question.

Note that either of these two milder forms of SIN is quite different from not having sufficient data to resolve a tree topology (a much more familiar problem for biologists)—I deal with this latter issue in a later section. By contrast, mild SIN requires that the tree will never become fully resolved, no matter how much data we were to obtain.

Topological Aspects of Statistical Estimation

I now describe two conditions ("no-touching" and "kissing") that make accurate estimation of the discrete parameter $a \in A$ simultaneously possible and problematic in the following sense: Given "enough" data, we can be sure to reconstruct a correctly, but we cannot say in advance how large "enough" will be. These two conditions typically hold in the reconstruction of fully resolved phylogenetic trees as well as other related problems. To describe the conditions, two further definitions are required.

Given the model parameters (a, θ), let $p_{(a, \theta)}$ denote the associated probability distribution on site patterns, and let $p(\Theta(a)) := \{p_{(a,\theta)} : \theta \in \Theta(a)\}$, which is a subset of the $|U|$-dimensional simplex of probability distributions on the set U of site patterns. Also, given a subset A of Euclidean space, let $\overline{A}$ denote its (topological) closure. The two conditions can now be stated: For all $a, b \in A$, with $a \neq b$, consider the following:
no-touching ("identifiability") condition:

$$p(\Theta(a)) \cap \overline{p(\Theta(b))} = \emptyset \quad \text{and} \quad (2)$$

kissing condition:

$$\overline{p(\Theta(a))} \cap \overline{p(\Theta(b))} \neq \emptyset. \quad (3)$$

In the phylogenetic setting, where we will often regard Θ as branch lengths, $p(\Theta(a))$ will be all the probability distributions we can obtain on site patterns by varying the branch lengths on the binary tree a over all strictly positive but finite values. The set $\overline{p(\Theta(b))}$ includes not just all probability distributions one can obtain on site patterns by varying the branch lengths on the binary tree b over all strictly positive but finite values but also the limiting distributions as branch lengths tend to zero or to infinity (in all possible combinations). I provide an example (and figure) to illustrate these concepts after some brief remarks.

In general, the no-touching condition (2) alone is sufficient to ensure that ML in the CM setting will consistently reconstruct each discrete parameter in A when the

observations are generated under a CM (Chang 1996; Steel and Székely 2009). The condition holds for many models in molecular systematics, including the general Markov model, a simple Covarion model, and models that exhibit low-parameter rate variation across sites, such as the "GTR + Γ" model (Allman et al. 2008). An outstanding open problem until now has been whether the widely used "GTR + Γ + I" model satisfies the no-touching condition if both the shape parameter and the proportion of invariable sites are unknown. Recent work shows that the answer is "yes" in all (or almost all) cases (Wu and Susko 2010; Chai and Housworth 2011). Shortly, we will describe some models for which the no-touching condition has been shown to fail.

The kissing condition (3) is also relevant to phylogenetics, indeed it applies to all models of character evolution used for inferring tree topology. Any two different trees can produce identical data if the lengths of the interior branches on which the two trees differ shrink to zero and/or the lengths of all (or "most") of the pendant edges grows to infinity ("site saturation"); these phenomena reflect the geometry of tree-space discussed in Kim (2000) and Moulton and Steel (2004), where quite different trees can come arbitrarily close together ("kiss") in terms of the distribution on site patterns they can produce. This means that the sequence length required to reconstruct a tree correctly by any method tends to infinity as the interior branches shrink in length or as the pendant ones grow.

Note that this "tree-space" is related to, but quite different from, the tree space described by Billera et al. (2001), for example, the latter tree space regards two trees of different topologies as becoming infinitely far apart as the length of all their branches grow; however, in the tree space here, they are regarded as becoming closer together because they are tending to produce exactly the same data (random sequences). This is illustrated in the following example.

Example

It is easy to visualize conditions (2) and (3) by means of a simple but instructive example. Consider the three-rooted binary trees on leaf taxa $1, 2, 3$, which have branch lengths that satisfy a molecular clock. Let L denote the length of the interior edge length and l the length of the shorter pendant edge length, so $L + l$ is the length of the longer pendant edge length (Fig. 1a). For the tree $a_1 = 1|23$, the set $\Theta(a_1)$ is the infinite open first quadrant of the plane: $\{(l, L) : L, l > 0\}$.

Now consider the function that assigns to (l, L) the probability distribution on site patterns under some model. For simple models, such as the N_r model, this function can be described as the composition of two continuous one-to-one and onto functions. The first map associates (l, L) with the vector $(1 - e^{-cl}, 1 - e^{-cL})$, where c is a fixed constant (dependent on the model). The image P_1 of this map is the open square $(0, 1) \times (0, 1)$ (Fig. 1b) where representative branch lengths for points

near the 4 corners of P_1 are illustrated. The second map $\xi : P_1 \to (0, 1)^{|U|}$ sends $(x, y) = (1 - e^{-cl}, 1 - e^{-cL})$ to a probability distribution on site patterns that is determined by the branch lengths (l, L) associated with (x, y).

For the N_2 model, with a uniform probability on the two states at the root, there are 8 site patterns, $(x, y) = (1 - e^{-2l}, 1 - e^{-2L})$ (i.e., $c = 2$), and the 8 components of $\xi(x, y)$ take just 3 different values according to whether 1) all 3 leaves are in the same state, 2) leaf 1 is in the same state as just one of the other two leaves, or 3) leaf 1 is in a different state to the other two leaves. Using standard Hadamard representation (see, e.g., Semple and Steel 2003), these 3 probabilities, which apply to two-, four-, and two-site patterns, respectively, are as follows: $\frac{1}{8}(1 + \overline{x}^2 + 2\overline{x}^2\overline{y}^2)$, $\frac{1}{8}(1 - \overline{x}^2)$, $\frac{1}{8}(1 + \overline{x}^2 - 2\overline{x}^2\overline{y}^2)$, where $\overline{x} = 1 - x$ and $\overline{y} = 1 - y$.

Moreover, the map ξ extends to $\overline{P_1}$ (the closure of P_1, which is the closed square as shown in Fig. 1c) and $\overline{p(\Theta(a_1))} = \xi(\overline{P_1})$. Similarly, for each of the other two trees, $a_2 = 2|13$ and $a_3 = 3|12$, one has $\overline{p(\Theta(a_i))} = \xi(\overline{P_i})$, $i = 2, 3$ where $\overline{P_2}$ and $\overline{P_3}$ are the corresponding closed squares for the other two trees (Fig. 1d,e). Each point on the bottom boundary of $\overline{P_1}$ (corresponding to $L = 0$) has a corresponding point on the bottom boundary of $\overline{P_2}$ and of $\overline{P_3}$ that induce exactly the same probability distribution on site patterns, and so these three regions "kiss" at each such point (one such shared point is indicated in each region in Fig. 1c,d,e). Thus, we can identify (or glue) these three squares along their bottom boundary (Fig. 1f). Finally, any point on the front boundary (corresponding to $l = \infty$) leads to the same probability distribution on site patterns—for the N_r model, this would simply assign each of the possible site patterns equal probability. Thus, the whole of this Y-shaped part of the complex in Figure 1f is identified to a single point, resulting in the final "paper dart" representation of the tree space shown in Figure 1g (an example of a "closure-finite, weak-topology complex" in topology).

The main point about this complex is that a one-to-one correspondence can be seen between the points on the "paper dart" and the probability distribution on site patterns that can be induced by 3-taxon trees under a molecular clock where the edge lengths can vary from 0 to (actual) infinity. Note that the "spine" of the dart corresponds to the unresolved star tree, whereas the "head" of the dart corresponds to pendant branches of infinite length. The no-touching condition (2) is satisfied because the interior of any one of the 3 wings does not intersect any other wing (even at the boundary of that other wing), whereas the kissing condition (3) holds since the wings all touch each other along the central spine (and also, for a different reason, at the front tip).

Failure of the No-Touching Condition for Certain CM Models

Note that the no-touching condition generally applies to simple models of site substitution in phylogenetics

when A is the set of fully resolved (binary) phylogenetic trees. However, it can collapse in three important cases: The first is if the set A is enlarged to include all phylogenetic trees (binary and nonbinary) on a given set of leaf taxa, because if a has a polytomy, and b is a tree obtained by resolving that polytomy, then

$$p(\Theta(a)) \cap \overline{p(\Theta(b))} = p(\Theta(a)) \neq \emptyset.$$

In words, this condition says that the distribution on site patterns that tree a can generate is just a subset of those that can be generated from tree b in the limit as the branch lengths of edges that resolve the polytomy of a tend to zero (for instance, in Fig. 1, the central "spine" of the paper dart lies in each of the 3 triangles corresponding to the 3 resolved trees). Indeed, even under the CM model, the reconstruction of general (including nonbinary) trees using ML will not be statistically consistent because if the generating tree is nonbinary, the ML tree will typically be a resolution of this tree for any sequence length (though the lengths of the branches that resolve the polytomy will converge to zero in probability as the sequence length k grows, as noted by Chang 1996). It is possible to consistently reconstruct

general (including nonbinary) trees by adopting alternative approaches to tree reconstruction in which edges that are not supported sufficiently strongly for the given sequence length are contracted; one such approach is described by Gronau et al. (2008).

The second situation where the no-touching condition (2) may collapse (even when A is confined to be the set of fully resolved trees) is when we have a phylogenetic mixture, allowing different branch lengths on the same tree, for certain models. In this case, not only can equation (2) fail, but the examples constructed in Matsen et al. (2008) for two-tree mixtures under the N_2 model satisfy the stronger violation condition:

$$p(\Theta(a)) \cap p(\Theta(b)) \neq \emptyset. \tag{4}$$

In the setting of Matsen et al. (2008), $\Theta(a)$ refers to all triples (λ, λ', p) where λ and λ' are assignments of positive but finite branch lengths to tree a, whereas p (respectively $1 - p$) is the probability that the site evolves under the first (respectively second) set of branch lengths under the N_2 model. Thus, equation (4) describes the situation in which two fully resolved trees of differing topology can induce exactly the same
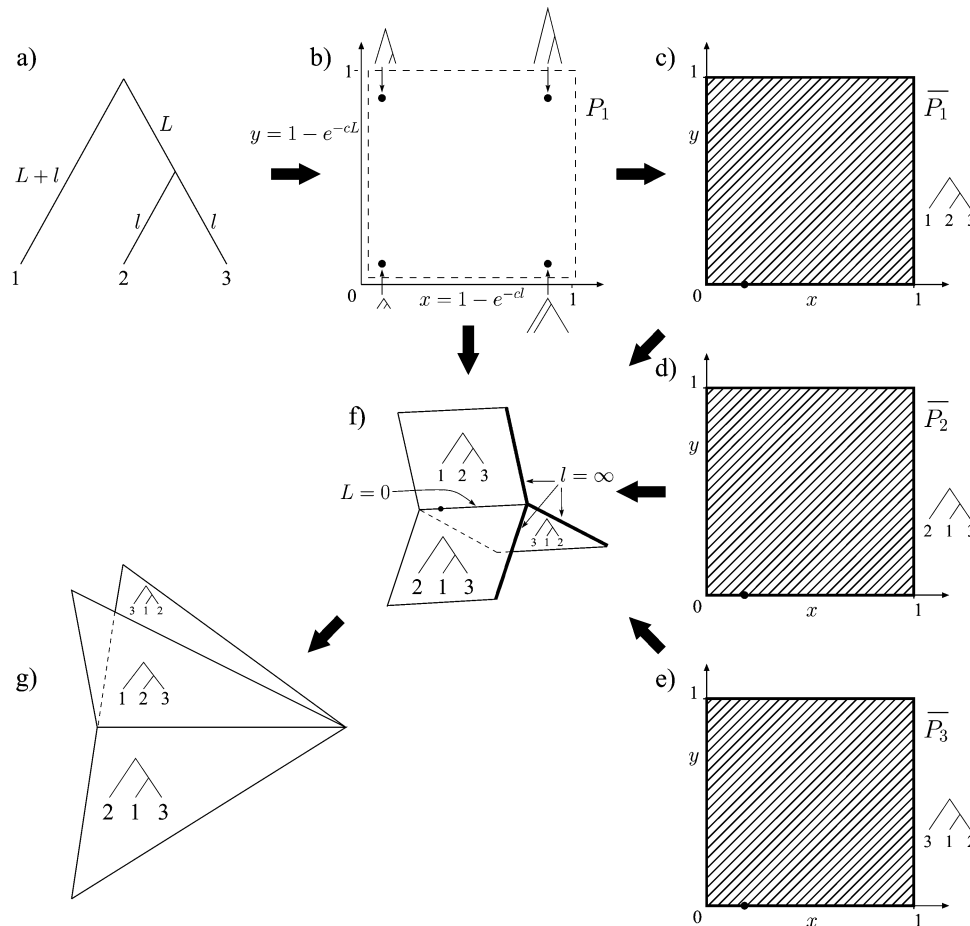

FIGURE 1. The "paper dart" representation of tree space for a Markov process on three taxa subject to a molecular clock (see text for details).

probability distribution on site patterns under their particular mixture processes.

To better visualize condition (4), Figure 2 illustrates a situation where the condition does not apply (three taxa with a molecular clock) and one where it does (four taxa without a molecular clock), both under the N_2 model. Geometrically, to form all possible mixtures of two or more branch lengths on fixed a tree a, we simply add to the unmixed image space $p(\Theta(a))$ all points (i.e., probability distributions) that lie in the "convex hull" of this space (e.g., line segments between any pair of points on $p(\Theta(a))$). For the first case (three taxa and under the N_2 model with a molecular clock, as in Fig. 1), forming mixtures generates no new points, so the set $p(\Theta(a))$ is the same under both the unmixed and the mixed setting (Fig. 2i,ii); in particular, condition (4) does not hold. However, in the second case, $p(\Theta(a))$ is strictly larger in the mixed setting than in the unmixed setting (as shown schematically in moving from Fig. 2iii to Fig. 2iv by the shading). Moreover, the spaces for two different trees (a, b) can be sufficiently entangled geometrically that the shaded regions themselves intersect (Matsen et al., 2008), as illustrated by the central darkly shaded intersection in Figure 2iv; in which case condition (4) holds.

The no-touching condition for general mixture models in phylogenetics has recently been explored by Stefankovic and Vigoda (2007), who established a remarkable duality theorem—either there exists a "linear test" to determine tree topology (such as methods based on linear phylogenetic invariants) or the no-touching condition fails.

A third situation where the no-touching condition (2) may collapse is when there is a distribution of rates across sites with too many unknown parameters. Indeed, it was shown in Steel et al. (1994), that an even stronger violation than equation (4) is possible, namely all fully resolved trees on a given set of species can induce the same probability distribution on site patterns for appropriately chosen (but positive and finite) branch lengths in an N_2 model and distributions of rates across sites—in other words:

$$\bigcap_{a \in A} p(\Theta(a)) \neq \emptyset,$$

where $\Theta(a)$ is the set of branch lengths and the parameters describing the distribution of rates across sites.

In cases where the no-touching condition (2) fails (i.e., when $p(\Theta(a)) \cap \overline{p(\Theta(a))} \neq \emptyset$ for some pair $a \neq b$), for example, when a model is "overparameterized," we have a useful distinction based on whether or not the overlap of $p(\Theta(a))$ and $\overline{p(\Theta(a))}$ has the same or smaller dimension than $p(\Theta(a))$. In the latter case, although the model fails to satisfy the no-touching condition, it fails only on a subset of $\Theta(a)$ of zero relative volume—in this case, the tree topology is said to be "generically identifiable" under the model. The distinction between global identifiability (i.e., no-touching per se) and generic identifiability is important for trying to decide whether SIN is "theoretically possible but unlikely to occur in practice"

or whether there is a reasonable chance of being in a region of parameter space where we might be unable to distinguish between two competing trees, even from infinite data. The distinction has often been overlooked in earlier studies but is carefully discussed now, particularly as generic identifiability is a notion that sits much more comfortably with current mathematical methods for studying the properties of Markov models based on algebraic geometry and phylogenetic invariants (Allman and Rhodes 2008; Allman et al. 2008).

SIN IN THE NCM SETTING

In the NCM model, the θ_i values may vary in some unknown way. In particular, one does not assume that they are selected i.i.d. from some distribution. By analogy with the CM setting, it is tempting to extend the definition of statistical consistency of a method $\mathcal{M}$ to the NCM setting by the following slight modification: "For each $a \in A$, and sequence $\theta_i \in \Theta(a)$, the probability that M correctly estimates a from $(u_1, \ldots, u_k)$, when each u_i is generated independently by the model with parameters (a, θ_i), converges to 1 as k grows."

However, this condition is too strong: In the tree setting, if the branch lengths grew to infinity (or shrank to zero) sufficiently fast with each observation, then accurate tree reconstruction by any method for any model can be ruled out (by the kissing condition). Nevertheless, there are meaningful notions of statistical consistency in the NCM setting, which generalize the CM definition. Recalling that $\Theta(a)$ is an open subset of Euclidean space and that a "compact" subset of Euclidean space is any subset that is closed and bounded, we will consider the following:

Definition 1. A method $\mathcal{M}$ for reconstructing a discrete parameter in a finite set A will be said to be statistically consistent for data generated by an NCM model if it satisfies the following condition:

> For each $a \in A$, and every compact subset C of $\Theta(a)$, the probability that $\mathcal{M}$ correctly estimates a from $(u_1, \ldots, u_k)$, when each u_i is generated independently by the model with parameters (a, θ_i), where $\theta_i \in C$, converges to 1 as k grows.

This definition is equivalent to the definition of statistical consistency under the CM version of the model if we further insist that all the θ_i values are equal, and in this case, the choice of the compact set C can be restricted to single points in $\Theta(a)$.

Note also that when we perform ML estimation under CM or NCM, we do not require that the θ_i values associated with a lie within any given compact subset C of $\Theta(a)$; they can take any value in $\Theta(a)$.

Which Phylogenetic Models and Methods Can Lead to SIN?

The following main result shows that the issue of statistical consistency under NCM is a delicate one,

depending on the details of the model and the method. The full mathematical proof of the following results is provided in the Appendix (Part (ii) can also be derived from the duality theorem of Stefankovic and Vigoda 2007).

Theorem 1.

i. (Inconsistency of ML for NCM-N_r model) ML estimation of fully resolved tree topology under the NCM-N_r model applied to NCM-N_r data (or even CM-N_r data) is statistically inconsistent for any finite number of states ($r > 1$). Moreover, no tree reconstruction method is statistically consistent for NCM-N_2 data.

ii. (Consistency for the NCM Jukes–Cantor model) In contrast to Part (i), there is a statistically consistent method for inferring fully resolved tree topology from data generated by an NCM-N_4 model.

iiii. (Consistency for NCM models with a molecular clock) Neighbor joining on uncorrected sequence dissimilarity is a statistically consistent method for inferring fully resolved tree topology from data generated by an NCM model where each site evolves under its own GTR process (with its own strictly positive rate matrix and branch lengths) provided that, at each site, the branch lengths are clocklike on the generating tree (e.g., if substitution rate varies arbitrarily from site to site, but at each site is constant across the tree).

iv. (Consistency of ML for NCM-N_∞ model) ML estimation of fully resolved tree topology under the NCM-N_∞ (or even under the CM-N_∞ model) of NCM-N_∞ data is statistically consistent.

MEASURING SIN AND TAKING PRECAUTIONS AGAINST IT

In addition to the topological view of tree reconstruction, described above, there is an equivalent metric view. To explain this, take any continuous distance function d on probability distributions on U (the set of site patterns). For example, one might take the "variational distance" defined by $d(p, q) = \frac{1}{2} \sum_{u \in U} |p(u) - q(u)|$. An alternative way of expressing the no-touching and kissing conditions ((2) and (3)) is to require, for all $a, b \in A$, with $a \neq b$:

$$\inf_{\theta' \in \Theta(b)} d(p_{(a,\theta)}, p_{(b,\theta')}) > 0 \quad \text{and} \quad (5)$$

$$\inf_{\theta \in \Theta(a), \theta' \in \Theta(b)} d(p_{(a,\theta)}, p_{(b,\theta')}) = 0, \quad (6)$$

respectively, where "inf" refers to infimum (the minimal value achieved or in the limit). These are identical conditions to equations (2) and (3), respectively, by standard arguments from analysis based on the Bolzano–Weierstrass theorem.

One advantage of this metric viewpoint is that a strictly positive value in equation (5) not only tells us that ML is consistent in the CM setting but also the

magnitude of this value sets explicit upper and lower bounds on how much data (sequence length) would be required in order to reconstruct the discrete parameter (tree) accurately (Steel and Székely 2002, 2009). The more closely a tree with appropriately chosen branch lengths (or other parameters) can fit the probability distribution on site patterns of a different tree, the more sequence sites it will take to tell which of the two trees produced the data.

The sequence length required for accurate tree reconstruction (under any CM model) also depends on the number of species being classified. Quantifying this relationship is particularly challenging mathematically (see, e.g., Daskalakis et al. 2009). Various optimal or near-optimal results have been established, which usually require developing a new and clever tree reconstruction method (not because such methods are necessarily better than ML estimation but rather because it has been difficult so far to rigorously establish good bounds on the sequence length required for ML to reconstruct a large tree accurately).

Much less is known about the sequence length requirements for tree reconstruction under NCM models. In the case of the N_∞ model (Theorem 1, Part 1(iv)), the sequence lengths required for accurate tree reconstruction from data generated by an NCM version of the model are essentially the same as for the CM version of the model, provided that, in both models, we insist that all edge lengths are bounded between (r, s) where $0 < r \leq s < \frac{1}{2}$. However, it seems entirely possible that for a finite-state N_r model such as the Jukes–Cantor model, the sequence length required for accurate tree reconstruction from NCM-N_4 data will be much larger than for a CM-N_4 model with comparable upper and lower bounds on the branch lengths. If so, this would be another example of where the two models (finite state vs. infinite state) have quite different statistical properties. Two other examples are as follows:

- The sequence length required to resolve a short interior branch of the tree of length ϵ (from the two alternative tree topologies obtained by swapping branches across the edge) grows at the rate $1/\epsilon^2$ for the finite-state model but just at $1/\epsilon$ for the infinite-state model, as $\epsilon \to 0$ (Steel and Székely 2002; Mossel and Steel 2004).

- The sequence length required to determine which of two resolved trees, that classify the same n species, generated the sequence data must grow with n under the finite-state model but can be independent of n for the infinite-state model (Steel et al. 2009).

Returning to the NCM-N_r model of Tuffley and Steel (1997), where branch lengths are allowed to vary freely from site to site, one attempt to avoid the massive overparameterization of this model is to assume that these different branch lengths (between characters and across sites) are assigned randomly (i.i.d.) according to some

common fixed prior distribution. This "Bayesian NCM" model was explored by Huelsenbeck et al. (2008). As the authors noted, this model has an interesting property: If one has an underlying N_r model, then this Bayesian NCM model induces exactly the same probability distribution on site patterns as what might be called the "ultra-CM model" (UCM), where each character has the same branch length, and these branch lengths are the same "across the tree." This formal equivalence between such a tightly constrained model (which would never be used in ordinary phylogenetic practice) and a type of NCM model seems at first a little paradoxical until it is realized that the assumption of common fixed prior distribution on branch lengths (across the characters and across the trees) is a very strong "commonality assumption." The formal equivalence becomes only approximate for more complex substitution models (technically, this is the result of the rate matrix having multiple nontrivial eigenvalues).

A related but less constrained version of this Bayesian NCM model was developed by Wu et al. (2008). In this model, the tree has underlying branch lengths that are common across the characters but can vary across the tree, and each site has an intrinsic rate, which multiplies the branch lengths across the tree. But in contrast to Olsen's model of allowing this per-site rate to be a free parameter to be optimized in ML (as described near the start of this paper), Wu et al. (2008) assume that this rate parameter is selected i.i.d. from a fixed distribution of rates across sites. For this model, when the underlying substitution process is (say) a Jukes–Cantor model, this intermediate-level Bayesian NCM model assigns exactly the same probability distribution on site patterns as a model in which all the sites evolve i.i.d. under a Jukes–Cantor model (with no rate variation across sites). As with the UCM model, the formal equivalence becomes only approximate for more complex substitution models and for the same reasons. However, in this more general setting, Wu et al. (2008) showed how a "log-det" transformation gives a statistically consistent way to establish the tree topology from data generated under this model.

Note that in the more usual CM setting, selecting the tree with the maximum posterior probability (using Bayesian methods with, say, exponential priors on branch lengths) is a statistically consistent method for inferring resolved tree topology. Here, one assumes (as with any consistent tree reconstruction method) that the generating tree is fully resolved (with positive finite branch lengths) and the model satisfies the usual no-touching condition (for a short proof of consistency of Bayesian inference in this setting, see Steel 2010). The possibility that such a method might lead to SIN for certain combinations of branch lengths in the generating tree was suggested recently by Kolackzkowski and Thornton (2009); however, the authors have since qualified this claim with on-line commentary on their paper. The statistical properties of NCM have also been investigated recently from standard model-selection approaches such as Akaike's information criterion (AIC) (Holder et al. 2010; Huelsenbeck et al. 2011). It is clear

that NCM can confer higher likelihood scores than a CM model for any data because one has so much flexibility when choosing the nuisance parameters to get a good fit. Model selection techniques such as AIC (as well as Bayesian information criterion and other variants) penalize models that are too parameter rich by subtracting from the log likelihood of the model a term that depends on the number of parameters (Akaike 1973). Under this criterion, it seems unlikely that the full-blown NCM model will ever be favored over CM models under AIC. However, it is not entirely clear that the conditions required to justify the AIC extend rigorously to this NCM setting.

It Is Not SIN if . . .

Theorem 1(i) tells us that, under the NCM-N_r model, the tree reconstruction method NCM-ML (MP) is subject to SIN. Indeed, not only can this method fail to converge on the true tree, but it can be positively misleading (i.e., converge on an incorrect tree). It can be asked whether we can avoid this "mortal SIN" property if our data are sufficiently "convincing" in its support of one tree. To make this idea more precise, suppose that E_k is an event (property of the sequences of length k). Then, we can consider the following approach: If event E_k occurs, then we will select the MP tree (=NCM-ML tree), whereas if E_k fails to occur, we will not return any tree. We would like two properties to hold: 1) the probability of returning an incorrect tree by this conditional method tends to 0 as k grows and 2) the event E_k has a probability that does not converge to zero with increasing k, for certain regions of parameter space (i.e., we do not want to condition on a "miracle" that could never occur in practice). Precisely such a set-up has been described, in a slightly different setting, by Cavender (1978, 1981). For a 4-taxon tree, let E_k be the event that the number of parsimony-informative site patterns that support the most-parsimonious tree exceeds the number of parsimony-informative site patterns that support the second most-parsimonious tree by r_k. Then, provided r_k is at least $(\frac{1}{4}+\delta)k$ for some $\delta > 0$, if we select the MP tree (=NCM-ML tree), then we satisfy property 1. This follows from the results described in Cavender (1978) for the CM setting and their extension to an NCM setting in Cavender (1981). Moreover, in certain portions of parameter space for the 4-taxon tree (short external branches and a long interior branch), property 2 also applies—that is, the event E_k has a probability that is bounded away from zero as k grows.

Can SIN Still Occur If a "No Kissing" Condition Is Enforced?

I conclude this section by pointing out that the kissing condition (3) (or, equivalently equation (6)) is not necessarily the cause of SIN in the NCM setting. To see this, suppose the N_r model is constrained so that all the branch lengths in a tree lie between ϵ and $-\log(\epsilon)$ for some $\epsilon > 0$. Let us call this the N_r^ϵ model. Then, under

the N_r^ϵ model, one has, for different resolved phyloge-netic trees a, b on the same set of species:

$$\inf_{\theta \in \Theta(a), \theta' \in \Theta(b)} d(p_{(a,\theta)}, p_{(b,\theta')}) \geq q > 0, \qquad (7)$$

where $q = q(\epsilon)$ converges monotonically to 0 as $\epsilon \to 0$. Consequently, for $\epsilon > 0$, the kissing condition fails—topologically different trees cannot "look" arbitrarily close through the eyes of data produced under a CM model. However, ML estimation under the NCM-N_r^ϵ model can again be statistically inconsistent. Indeed, suppose one were to take any tree and branch lengths in the interior of a Felsenstein zone for that tree (i.e., a set of branch lengths where MP would converge on an incor-rect tree for data produced under the CM-N_r^ϵ (or NCM-N_r^ϵ model). Then, one can chose $\epsilon > 0$ small enough so that the ML estimate of the tree under the NCM-N_r^ϵ model converges on the wrong tree when applied to the CM-N_r^ϵ data produced from the original tree with its Felsenstein zone branch lengths. The formal proof of this claim is given in the Appendix.

CONCLUDING COMMENTS

Making molecular phylogenetic models more "realis-tic" and thereby capturing more of the complexities of how DNA evolves—across the genome and over differ-ent time scales—usually requires introducing a number of adjustable parameters. If these parameters can be in-dependently estimated from other data, or if they enter into the model in ways that are not problematic for tree inference (as in Theorem 1(ii–iv)), or if they follow some common distribution that is described by few if any un-known parameters), then statistically consistent infer-ence of tree topology is achievable. However, in general, treating branch lengths and other model parameters as unknown quantities can drive reconstruction methods to SIN.

Theorem 1 (Parts (ii–iv)) provides no real endorse-ment for NCM models, but it shows that sweeping assumptions that such models must necessarily lead methods to SIN are incorrect. Such arguments typi-cally proceed as follow: In NCM models, the num-ber of nuisance parameters grows with k and we are unable to estimate them with any precision, thus the usual conditions that suffice for the consistency of ML estimation (Wald 1949) fail and so the method will be inconsistent. All but the final conclusion of this last sen-tence are correct—the failure of a sufficient condition for a statement to be true is not sufficient for the state-ment to be false! Indeed, Theorem 1 provides specific cases where NCM-ML estimation is consistent for cer-tain NCM models.

Even when statistically consistent methods exist for an NCM model, it is still possible that ML can be sta-tistically inconsistent (this contrasts with what happens in the CM setting, where ML is generally consistent if any consistent method exists). This leads to a some-what uncomfortable position for those who wish to

provide some statistical justification for the use of MP as the ML estimator under the NCM model of Tuffley and Steel (1997)—by Theorem 1(i), such a method lives in a state of SIN, yet this could be avoided for this NCM model if one were to renounce MP in favor of a quite different method, such as one based on linear phylogenetic invariants (the method used in the proof of Theorem 1(ii)).

However, this is no strong argument in favor of lin-ear invariants, as they tend to be very inefficient in their ability to extract phylogenetic signal from data (Hillis et al. 1994). A method based on linear phylogenetic in-variants may be guaranteed to converge on the right tree eventually, even under the NCM model, but this may require an astronomical amount of data. By con-trast, methods such as MP appear to be quite efficient at extracting phylogenetic signal when the generating tree branch lengths are some way from those portions of parameter space that lead to inconsistency (Hillis 1996). Thus, although statistical consistency is desirable, it should not override all other considerations—for ex-ample, a powerful method that is consistent in most re-gions of parameter space would generally be preferred over a statistically consistent method that may requires huge amounts of data to converge.

Of course, many of the results in this paper are con-fined to very simple models (such as the Jukes–Cantor); I have chosen to do this for two reasons: First, they are sufficient to demonstrate that even with very simple models, all possibilities (consistency and SIN) are pos-sible given slight tweaks of the assumptions or method; second, the analysis of more complex models is beyond the scope of this paper but would be a worthy objec-tive for future work. Regarding this last point, results in Stefankovic and Vigoda (2007) show that Theorem 1(ii) applies to other models beyond N_4 such as the Kimura 2ST model but not to the Kimura 3ST model.

In summary, the question of whether one is prone to SIN by adopting a particular NCM model and a particu-lar method of inference has a more complex answer than in the CM setting—it depends subtly on the details of the model and on the method. The full-blown general-ity of the NCM of Tuffley and Steel (1997) is unnecessar-ily overparameterized for most data, being a model that was developed to prove a formal equivalence between methods rather than as a model of choice. Far from be-ing a justification of MP, its plethora of ever-growing parameters would surely have not seemed "parsimo-nious" to William of Occam. At the other extreme are simple attempts to include NCM within a Bayesian framework; these avoid SIN, but at the price of forc-ing the NCM model into a CM straightjacket by viewing the parameters as samples from a common underlying prior. Between these extremes, there would seem to be an endless variety of possibilities. The development of carefully constrained yet parameter-rich models, guided by model selection criteria, and which recognize that characters evolve under different processes dependent on their biochemistry, will surely play a significant role in future phylogenetic methodology.

REFERENCES

Akaike H. 1973. Information theory and an extension of the maximum likelihood principle. In: Petrov B., Csaki F., editors. Information theory. Budapest: Akademiai Kiado. p. 267–281.

Allman E.S., Ané C., Rhodes J.A. 2008. Identifiability of a Markovian model of molecular evolution with gamma-distributed rates. Adv. Appl. Probab. 40:228–249.

Allman E.S., Rhodes J.A. 2006. The identifiability of tree topology for phylogenetic models, including covarion and mixture models. J. Comput. Biol. 13:1101–1113.

Allman E.S., Rhodes J.A. 2008. Identifying evolutionary trees and substitution parameters for the general Markov model with invariable sites. Math. Biosci. 211:18–33.

Atteson K. 1999. The performance of neighbor-joining methods of phylogenetic reconstruction. Algorithmica. 25:251–278.

Billera L., Holmes S., Vogtmann K. 2001. Geometry of the space of phylogenetic trees. Adv. Appl. Math. 27:733–767.

Cavender J.A. 1978. Taxonomy with confidence. Math. Biosci. 40: 271–280.

Cavender J.A. 1981. Tests of phylogenetic hypotheses under generalized models. Math. Biosci. 54:217–229.

Chai J., Housworth E. 2011. On Rogers's proof of identifiability for the GTR + Γ + I model. Syst. Biol. in press.

Chang J.T. 1996. Full reconstruction of Markov models on evolutionary trees: identifiability and consistency. Math. Biosci. 137:51–73.

Daskalakis C., Mossel E., Roch S. 2009. Evolutionary trees and the Ising model on the Bethe lattice: a proof of Steel's conjecture. Probab. Theor. Related Fields. in press. doi:10.1007/s00440-009-0246-2.

Farris J.S. 2008. Parsimony and explanatory power. Cladistics. 24: 825–847.

Felsenstein J. 1978. Cases in which parsimony or compatibility methods will be positively misleading. Syst. Zool. 27:401–410.

Felsenstein J. 2004. Inferring phylogenies. Sunderland (MA): Sinauer Associates.

Fischer M., Thatte B. 2010. Revisiting an equivalence between maximum parsimony and maximum likelihood methods in phylogenetics. Bull. Math. Biol. 72:208–220.

Fitch W. 1971. Rate of change of concomitantly variable codons. J. Mol. Evol. 1:84–96.

Gaucher E., Miyamoto M. 2005. A call for likelihood phylogenetics even when the process of sequence evolution is heterogeneous. Mol. Phylogenet. Evol. 37:928–931.

Grimmett G., Stirzaker D. 2001. Probability and random processes. 3rd ed. Oxford: Oxford University Press.

Gronau I., Moran S., Snir S. 2008. Fast and reliable reconstruction of phylogenetic trees with very short edges. In: Shang-Hua Teng, editor. SODA: ACM-SIAM Symposium on Discrete Algorithms. Philadelphia (PA): Society for Industrial and Applied Mathematics. p. 379–388.

Hillis D.M. 1996. Inferring complex phylogenies. Nature. 383:130–131.

Hillis D.M., Huelsenbeck J.P., Swofford D.L. 1994. Hobgoblin of phylogenetics? Nature. 369:363–364.

Holder M., Lewis P.O., Swofford D. 2010. The Akaike information criterion will not chose the no common mechanism model. Syst. Biol. 59:477–485.

Huelsenbeck J., Alfaro M., Suchard M. 2011. Biologically-inspired phylogenetic models strongly outperform the no-common-mechanism model. Syst. Biol. in press.

Huelsenbeck J., Ané C. Larget B., Ronquist F. 2008. A Bayesian perspective on a non-parsimonious parsimony model. Syst. Biol. 57:406–419.

Kim J. 2000. Slicing hyperdimensional oranges: the geometry of phylogenetic estimation. Mol. Phyl. Evol. 17:58–75.

Kimura M., Crow J. 1964. The number of alleles that can be maintained in a finite population. Genetics. 49:725.

Kolackzkowski B., Thornton J.W. 2009. Long-branch attraction bias and inconsistency in Bayesian phylogenetics. PLos One. 4:e7891.

Lake J. 1987. A rate-independent technique for analysis of nucleic acid sequences: evolutionary parsimony. Mol. Biol. Evol. 4: 167–191.

Lemey P., Salemi M., Vandamme A.-M. 2009. The Phylogenetic handbook: a practical approach to phylogenetic analysis and hypothesis testing. 2nd ed. New York: Cambridge University Press.

Lewis P. 2001. A likelihood approach to estimating phylogeny from discrete morphological character data. Syst. Biol. 50:913–925.

Matsen F.A., Mossel E., Steel M. 2008. Mixed-up trees: the structure of phylogenetic mixtures. B. Math. Biol. 70:1115–1139.

Matsen F.A., Steel M. 2007. Phylogenetic mixtures on a single tree can mimic a tree of another topology. Syst. Biol. 56:767–775.

Mossel E., Roch S., Steel M. 2009. Shrinkage effect in ancestral maximum likelihood. IEEE/ACM Trans. Comput. Biol. Bioinform. 6: 126–133.

Mossel E., Steel M. 2004. A phase transition for a random cluster model on phylogenetic trees. Math. Biosci. 187:189–203.

Moulton V., Steel M. 2004. Peeling phylogenetic 'oranges'. Adv. Appl. Math. 33:710–727.

Neyman J. 1971. Molecular studies of evolution: a source of novel statistical problems. In: Gupta S., Yackel J., editors. Statistical decision theory and related topics. New York: Academic Press. p. 1–27.

Phillipé H., Zhou Y. Brinkmann H. Rodrigue N., Delsuc F. 2005. Heterotachy and long-branch attraction in phylogenetics. BMC Evol. Biol. 5:50.

Schulmeister S. 2004. Inconsistency of maximum parsimony revisited. Syst. Biol. 53:521–528.

Semple C., Steel M. 2003. Phylogenetics. Oxford: Oxford University Press.

Sober E. 2004. The contest between likelihood and parsimony. Syst. Zool. 53:6–16.

Sober E. 2008. Evidence and evolution: the logic behind the science. Cambridge: Cambridge University Press.

Steel M. 2010. Consistency of Bayesian inference of resolved phylogenetic trees. arXiv:1001.2864.

Steel M., Penny D. 2004. Two further links between MP and ML under the poisson model. Appl. Math. Lett. 17:785–790.

Steel M., Penny D. 2005. Parsimony, phylogeny and genomics. In: Albert V., editor. Maximum parsimony and the phylogenetic information in multi-state characters. Oxford: Oxford University Press. p. 163–178.

Steel M., Székely L.A. 2002. Inverting random functions (ii): explicit bounds for discrete maximum likelihood estimation, with applications. SIAM J. Discrete Math. 15:562–575.

Steel M., Székely L., Mossel E. 2009. Phylogenetic information complexity: is testing a tree easier than finding it? J. Theor. Biol. 258:95–102.

Steel M.A., Fu Y.X. 1995. Classifying and counting linear phylogenetic invariants for the Jukes-Cantor model. J. Comput. Biol. 2:39–47.

Steel M.A., Székely L. 2009. Inverting random functions (iii): discrete MLE revisited. Ann. Comb. 13:373–390.

Steel M.A., Székely L., Hendy M. 1994. Reconstructing trees when sequence sites evolve at variable rates. J. Comput. Biol. 1:153–163.

Stefankovic D., Vigoda E. 2007. Phylogeny of mixture models: robustness of maximum likelihood and non-identifiable distributions. J. Comput. Biol. 14:156–189.

Swofford D.L., Olsen G.J., Waddell P.J., Hillis D.M. 1996. Phylogenetic inference. Chapter 11. In: Hillis D.M., Moritz C., Marble B.K.,

editors. Molecular systematics. 2nd ed. Sunderland (MA): Sinauer Associates.

Tuffley C., Steel M. 1997. Links between maximum likelihood and maximum parsimony under a simple model of site substitution. B. Math. Biol. 59:581–607.

Wald A. 1949. A note on the consistency of the maximum likelihood estimate. Ann. Math. Stat. 20:595–600.

Warnow T., Evans S.N., Ringe D., Nakhleh L. 2006. A stochastic model of language evolution that incorporates homoplasy and borrowing. Chapter. 7. In: Forster P., Renfrew C., editors. Phylogenetic methods and the prehistory of languages. Cambridge: McDonald Institute for Archaeological Research. p. 75.

Wu J., Susko E. 2010. Rate variation need not defeat phylogenetic inference through pairwise comparisons. J. Theor. Biol. 263:587–589.

Wu J., Susko E., Roger A.J. 2008. An independent heterotachy model and its implications for phylogeny and divergence time estimation. Mol. Phylogenet. Evol. 46:801–806.

APPENDIX: TECHNICAL DETAILS

The following lemmas are required in the proof of Theorem 1.

Lemma 1. A method $\mathcal{M}$ for estimating the discrete parameter $a \in A$ from sequences of observations in U is statistically consistent under an NCM-M model if it satisfies the following property: For each $a \in A$, there is a nested family $\mathring{\Theta}_k(a), k = 1, 2, \ldots$, of increasing open subsets of Euclidean space with $\Theta(a) = \bigcup_{k=1}^{\infty} \mathring{\Theta}_k(a)$, so that the following condition holds: the probability that $\mathcal{M}$ correctly estimates element a from $(u_1, \ldots, u_k)$, whenever each u_i is generated by (a, θ_i), with $\theta_i \in \mathring{\Theta}_k(a)$, converges to 1 as k grows.

Proof. Suppose C is a compact subset of $\Theta(a)$. Then, $C \cap \mathring{\Theta}_k(a), k \geq 1$, is an open cover of C. Since C is compact, C is equal to the union of finitely many of the sets $C \cap \mathring{\Theta}_k(a)$, $k \geq 1$, are nested, a value of $k=k_1$ exists for which $C \subset \mathring{\Theta}_{k_1}(a)$. By the hypothesis of the lemma, the event that $\mathcal{M}$ correctly returns any a from $(u_1, \ldots, u_k)$ when each u_i is generated by (a, θ_i), where $\theta_i \in \mathring{\Theta}_k(a)$, has probability that converges to 1 as k grows. Since $C \subseteq \mathring{\Theta}_{k_1} \subseteq \mathring{\Theta}_k$ for $k \geq k_1$, restricting θ_i to lie in C ensures that the probability $\mathcal{M}$ correctly returns a from $(u_1, \ldots, u_k)$ when each u_i is generated by (a, θ_i), where $\theta_i \in C$ also converges to 1 as k grows. □

Lemma 2. (Azuma's inequality) There are several variants of this inequality (see, e.g., Grimmett and Stirzaker 2001) here I give a special case of a more general version. Suppose that $X_1, \ldots, X_k$ are independent random variables taking values in some set S and that $Y=f(X_1, \ldots, X_k)$ where $f: S^k \rightarrow \mathbb{R}$ is any function with the property that $|f(y_1, \ldots, y_k) - f(y'_1, \ldots, y'_k)| \leq c$ whenever $y'_i = y_i$ for all but one value of i. Then, $\mathbb{P}(Y - \mathbb{E}[Y] \geq x)$ and $\mathbb{P}(Y - \mathbb{E}[Y] \leq -x)$ are each less or equal to $\exp(-x^2/2c^2k)$.

Proof of Theorem 1

Part (i): ML estimation under the NCM-N_r model applied to any sequence of r-state characters returns the same tree(s) as MP (Tuffley and Steel 1997). This latter method was shown to be statistically inconsistent for CM-N_2 data (Felsenstein 1978) and, more generally, for CM-N_r data for $r \geq 2$ by later authors (see Schulmeister 2004 and the references therein) even for trees on four species. Since the CM-N_r model is just a submodel of the NCM-N_r model, both assertions in the first claim of Part (i) follow. Specifically, we can take a to be any resolved binary tree on 4 leaves, and $\Theta(a)$ to be $(0, \infty)^5$ and select the θ_i values all to be equal to a choice of branch lengths $\theta \in (0, \infty)^5$ for which MP (and thereby ML under NCM-N_r) converges on an incorrect tree.

The proof of the second claim, that concerning the NCM-N_2 model, follows directly from the examples in Matsen and Steel (2007).

For Parts (ii–iv), I will establish the statistical consistency of various methods by establishing the property described in Lemma 1.

Part (ii): The proof relies on the existence of certain linear phylogenetic invariants for the Jukes–Cantor model (the existence of such invariants for models that include the Jukes–Cantor was described by Lake 1987). In particular, from theorem 1 (part 5) of Steel and Fu (1995), any binary phylogenetic tree T has an associated function L_T of the site pattern frequencies, such that

(i) $L_T(\mathbf{p}) = 0$, where $\mathbf{p}$ is the probability vector of site patterns generated by T under any assignment of branch lengths,

(ii) for any binary phylogenetic tree T' that is different from T but has the same leaf set, we have $L_{T'}(\mathbf{p}) \geq f_T(u, v) > 0$,

where u is the shortest branch length, v is the largest branch length, and f is a continuous function that has the following 2 properties:

- for all $u > 0$, f is monotone decreasing in v and converges to 0 as v tends to infinity and

- for all $v > 0$, f is monotone increasing in u and converges to 0 as u tends to 0.

Although these are all the properties of f required for the rest of the proof, I provide an explicit description of f as follows:

For a binary tree T on 4 leaves (a quartet tree) with topology $ij|kl$, a linear invariant L_T for the Jukes–Cantor model was described in Steel and Fu (1995) with property (i) described in the previous paragraph and for any binary phylogenetic tree T' that is different from T but has the same leaf set, we have $L_{T'}(\mathbf{p}) \geq \exp\left(-\frac{4}{3}S\right) \cdot \left(1 - \exp\left(-\frac{8}{3}l\right)\right)$, where S is the sum of the lengths of the 4 pendant edges and l is the length of the interior edge. For any binary tree T (with 4 or more leaves), select any collection Q of quartet trees that define T (i.e., for which T is the unique tree that displays those quartet trees) and take the sum of the linear invariants just described to give a linear invariant L_T. Notice that L_T satisfies both conditions (i) and (ii) in the previous paragraph by taking $f(u, v)$ to be the function:

$\exp\left(-\frac{4}{3}(2n-4)v\right)\cdot\left(1-\exp\left(-\frac{8}{3}u\right)\right)$, where n is the number of species (so the number of edges of T that lie the pendant edges of any induced quartet tree is at most $(2n-4)$). This is the function f promised.

Returning to the proof of Part (ii), for a site s, let $\mathbf{X}^s$ be the 4^n-dimensional vector, indexed by site patterns, that takes the value 1 for the site pattern observed at site s and 0 otherwise, and let $\hat{\mathbf{x}}^{(k)} = \frac{1}{k}\sum_{s=1}^{k}\mathbf{X}^s$. Consider the following tree reconstruction method $\mathcal{M}$: select the binary phylogenetic tree T that minimizes $L_T(\hat{\mathbf{x}}^{(k)})$.

We will apply Lemma 1 by ensuring that the branch lengths in the generating tree at site $i \in \{1,\dots,k\}$ lie between ϵ_k and L_k, where these two sequences converge monotonically to zero and to infinity (respectively), sufficiently slowly with k.

To this end, suppose k sites evolve on a fully resolved tree T under an NCM-N_4 model. Let $\mathbf{p}^s = \mathbb{E}[\mathbf{X}^s]$, the vector of probabilities of the different site patterns at site s, and let $\mathbf{p}^{(k)} := \frac{1}{k}\sum_{s=1}^{k}\mathbf{p}^s$. By the invariant property of L_T, we have $L_T(\mathbf{p}^s) = 0$ for all s, and since L_T is linear, it follows that

$$\mathbb{E}[L_T(\hat{\mathbf{x}}^{(k)})] = L_T(\mathbb{E}[\hat{\mathbf{x}}^{(k)}]) = L_T(\mathbf{p}^{(k)}) = 0 \quad \text{for all } k \geq 1. \quad (8)$$

Similarly, for any fully resolved phylogenetic tree T' that is different from T but has the same leaf set, we have

$$\mathbb{E}[L_{T'}(\hat{\mathbf{x}}^{(k)})] = L_{T'}(\mathbf{p}^{(k)}) > f(\epsilon_k, L_k). \quad (9)$$

From the continuity of f and its other listed properties, we can allow ϵ_k to tend to 0 and L_k to tend to infinity sufficiently slowly (with increasing k) that the following condition is satisfied:

$$\lim_{k\to\infty} k \cdot f^2(\epsilon_k, L_k) \to \infty. \quad (10)$$

Now, since the $\mathbf{X}^s$:$s=1,\dots,k$ are independent random variables, the Azuma inequality combined with (8) and (10) gives $\lim_{k\to\infty}\mathbb{P}\left(L_T(\hat{\mathbf{x}}^{(k)}) > \frac{1}{2}f(\epsilon_k, L_k)\right) = 0$, whereas for any alternative fully resolved phylogenetic tree T' ($\neq T$), equations (9) and (10) give $\lim_{k\to\infty}\mathbb{P}\left(L_{T'}(\hat{\mathbf{x}}^{(k)}) < \frac{1}{2}f(\epsilon_k, L_k)\right) = 0$. Combining these two last equations gives: $\lim_{k\to\infty}\mathbb{P}(L_T(\hat{\mathbf{x}}^{(k)}) < L_{T'}(\hat{\mathbf{x}}^{(k)})) = 1$, and so

$$\lim_{k\to\infty}\mathbb{P}(L_T(\hat{\mathbf{x}}^{(k)}) < L_{T'}(\hat{\mathbf{x}}^{(k)}) \quad \text{for all } T' \neq T) = 1.$$

By Lemma 1, this implies that method $\mathcal{M}$ is statistically consistent under the model described.

Part (iii): Let $d_{ij}^{(k)}$ denote the proportion of the k sites on which species i, j disagree and let $\mu_{ij}^{(k)} = \mathbb{E}[d_{ij}^{(k)}]$. Thus, $d_{ij}^{(k)} = \frac{1}{k}\sum_{s=1}^{k}\xi_s^{ij}$, where ξ_s^{ij} takes the value 1 if sequences i and j differ at site s, and 0 otherwise. By the standard theory of reversible r-state Markov processes, combined with the molecular clock hypothesis, for any two species x, y, we can write

$$\mathbb{E}[\xi_s^{xy}] = \left(1 - \sum_{i=1}^{r}\pi_{s,i}^2\right) + \sum_{j=1}^{r-1}\alpha_{s,j}e^{-2\beta_{j,s}t_{xy}}, \quad (11)$$

where

- $\pi_{s,i}$ is the vector of equilibrium base frequency of base i at site s,

- $-\beta_{s,j}$ are the nonzero eigenvalues of the GTR rate matrix at site s,

- the coefficients $\alpha_{s,j}$ are positive (and determined by the eigenvalues of the GTR matrix at site s, along with the $\pi_{s,i}$ values),

- t_{xy} is the time from when species x and y diverged in the tree to the present.

Consequently, if the generating tree T resolves the triplet of species i, j, l as the rooted tree $ij|l$, then

$$\mathbb{E}[\xi_s^{il}] = \mathbb{E}[\xi_s^{jl}] > \mathbb{E}[\xi_s^{ij}], \text{ and } \mathbb{E}[\xi_s^{il}] - \mathbb{E}[\xi_s^{ij}] \geq g(u_s, v_s), \quad (12)$$

where g is a continuous function that has the same properties as f in the previous proof and where

- u_s is the sum of the branch lengths on the path between the least common ancestor of i, l and the least common ancestor of i, j at site s times the substitution rate at site s,

- v_s is the sum of the branch lengths between the root and any leaf multiplied by the largest magnitude of any eigenvalue of the GTR matrix at site s.

An explicit description of the function g is as follows: From equation (11), we have $\mathbb{E}[\xi_s^{il}] - \mathbb{E}[\xi_s^{ij}] = \sum_{j=1}^{r-1}\alpha_{s,j}(\exp(-2\beta_{j,s}t_{il}) - \exp(-2\beta_{j,s}t_{ij}))$, and, using the identity $e^{-x} - e^{-y} = e^{-y}(e^{y-x} - 1) \geq e^{-y}(y - x)$ for $0 < x < y$, we have

$$\mathbb{E}[\xi_s^{il}] - \mathbb{E}[\xi_s^{ij}] \geq 2\sum_{j=1}^{r-1}\alpha_{s,j}\beta_{j,s}\exp(-2\beta_{j,s}t_{il})\cdot(t_{il} - t_{ij}).$$

Now the term $\sum_{j=1}^{r-1}\alpha_{s,j}\beta_{j,s}$ is the substitution rate at site s, and so we can set $g(u, v) = 2u_s e^{-2v_s}$.

Thus, if we let $\mu_{ij}^{(k)} = \mathbb{E}[d_{ij}^{(k)}]$, then

$$\mu_{il}^{(k)} = \mu_{jl}^{(k)} \quad \text{and} \quad \mu_{il}^{(k)} - \mu_{ij}^{(k)} > g(\epsilon_k, L_k), \quad (13)$$

where $\epsilon_k = \min\{u_s : 1 \leq s \leq k\}$ and $L_k = \max\{v_s : 1 \leq s \leq k\}$.

As in the previous proof, by the continuity of g and its other listed properties, we can allow ϵ_k to tend to 0 and L_k to tend to infinity sufficiently slowly (with increasing k) that $\lim_{k\to\infty}k \cdot g^2(\epsilon_k, L_k) \to \infty$. Then by Azuma's inequality:

$$\lim_{k\to\infty}\mathbb{P}\left(\max_{ij}|d_{ij}^{(k)} - \mu_{ij}^{(k)}| \geq \frac{1}{2}g(\epsilon_k, L_k)\right) = 0. \quad (14)$$

Note that, by equation (13), the μ values are additive on T and each interior edge has a branch length of at least $g(\epsilon_k, L_k)$. We can thus invoke the "safety radius" result

of Atteson (1999), which guarantees that neighbor joining applied to the matrix of $d_{ij}^{(k)}$ values, for all pairs i, j, will return T provided that each pairwise distance $d_{ij}^{(k)}$ differs from $\mu_{ij}^{(k)}$ by at most $\frac{1}{2}g(\epsilon_k, L_k)$. This last event has probability converging to 1 as k grows by equation (14) and so Part (iii) now follows from Lemma 1.

Part (iv): For any data consisting of a sequence of characters on a set of species, the only phylogenetic trees that have positive likelihood under the NCM-N_∞ model are those on which the data are homoplasy free (i.e., require no reverse or convergent evolutionary events). Thus, it suffices to show that for k characters generated by an NCM-N_∞ model on T, the probability that T is the only phylogenetic tree for the given species on which these characters are homoplasy free converges to 1 as $k \to \infty$. Following Warnow et al. (2006), it suffices to show that the following event E_k has probability converging to 1 as k grows: E_k is the event that for each induced quartet tree $ab|cd = T|\{a, b, c, d\}$ of T, at least one of the k characters assigns the same state to a and b, and the same state to c and d, and with these two states being different. By the independence assumption between changes on different edges in the N_∞ model, and by the Bonferroni inequality, we have

$$\mathbb{P}(E_k) \geq 1 - \binom{n}{4} \cdot \prod_{s=1}^{k}(1 - p_s q_s^4), \qquad (15)$$

where p_s (respectively q_s) is the smallest substitution probability on an edge (respectively the largest substitution probability on a path) for the process that generates site s. Thus, provided that the branch lengths at site s are bounded between (ϵ_k, L_k) where ϵ_k converges to 0 sufficiently slowly and that L_k converges to infinity sufficiently slowly (with increasing k), then $\lim_{k \to \infty} \mathbb{P}(E_k) = 1$, by equation (15). Part (iv) now follows from Lemma 1.

Proof that ML under an ϵ-Constrained NCM Can Be Statistically Inconsistent

For $\mathbf{u} = (u_1, \ldots, u_k) \in U^k$ and a fully resolved tree a, let $L_a(\mathbf{u})$ be the log of the ML value of the data $\mathbf{u}$ under the NCM-N_r model. From Tuffley and Steel (1997), we have

$$L_a(\mathbf{u}) = -(l(\mathbf{u}, a) + k) \cdot \log(r), \qquad (16)$$

where $l(\mathbf{u}, a)$ is the parsimony score of $\mathbf{u}$ on tree a. Similarly, for $\epsilon > 0$, let $L_a^\epsilon(\mathbf{u})$ be the log of the ML value of the data $\mathbf{u}$ under the constrained NCM-N_r model on tree a in which each branch length is required to lie between ϵ and $-\log(\epsilon)$. Clearly, $L_a^\epsilon(\mathbf{u}) \leq L_a(\mathbf{u})$.

Consider the following way to "prune" branch lengths in any tree c, which associates to each vector of branch lengths θ a corresponding set of branch lengths $\theta^{(\epsilon)}$ that satisfy the ϵ constraint: For each branch length shorter than ϵ reset that branch length to ϵ, and for each branch length larger than $-\log(\epsilon)$ reset it to $-\log(\epsilon)$. This transformation $\theta \mapsto \theta^{(\epsilon)}$ enjoys the following property for the N_r model: For any site pattern $u \in U$,

we have

$$\mathbb{P}(u|c, \theta^{(\epsilon)}) \geq \mathbb{P}(u|c, \theta) - O(\epsilon),$$

where $\mathbb{P}(u|c, \theta^{(\epsilon)})$ is the probability of generating u on tree c with branch lengths $\theta^{(\epsilon)}$ and where $O(\epsilon)$ is a term that depends just on ϵ and the number of leaves in the tree, and which tends to zero as $\epsilon \to 0$. It follows that

$$\frac{1}{k}L_c^\epsilon(\mathbf{u}) \geq \frac{1}{k}L_c(\mathbf{u}) - O(\epsilon). \qquad (17)$$

Now, by elementary algebra:

$$\frac{1}{k}(L_a^\epsilon(\mathbf{u}) - L_b^\epsilon(\mathbf{u})) = \Delta_1 + \Delta_2 + \Delta_3, \qquad (18)$$

where

$$\Delta_1 = \frac{1}{k}(L_a^\epsilon(\mathbf{u}) - L_a(\mathbf{u})) \leq 0,$$

$$\Delta_2 = \frac{1}{k}(L_a(\mathbf{u}) - L_b(\mathbf{u})) \quad \text{and} \quad \Delta_3 = \frac{1}{k}(L_b(\mathbf{u}) - L_b^\epsilon(\mathbf{u})).$$

Now, for any two trees a, b, equation (16) gives

$$\Delta_2 = \left(\frac{l(\mathbf{u}, b)}{k} - \frac{l(\mathbf{u}, a)}{k}\right) \cdot \log(r). \qquad (19)$$

If $\mathbf{u}$ is generated by a CM-N_r on a with branch lengths in the interior of the Felsenstein zone (a region of branch lengths for tree a where MP converges on an incorrect tree), then, for tree b having a different topology from a the term $\frac{l(\mathbf{u}, b)}{k} - \frac{l(\mathbf{u}, a)}{k}$ in equation (19) converges in probability to a negative constant $-C$ (the actual value of which is dependent on the branch lengths used in the Felsenstein zone setting).

Regarding Δ_3, we can apply inequality (17) for $c = b$ and select ϵ sufficient small (but strictly positive) so that (i) the branch lengths used in the Felsenstein zone setting are all greater than ϵ and less that $-\log(\epsilon)$ and (ii) the $O(\epsilon)$ term in equation (17) is less than $\frac{1}{2}C\log(r)$ and so, for all $k \geq 1$ and all $\mathbf{u} \in U^k$:

$$\frac{1}{k}(L_b(\mathbf{u}) - L_b^\epsilon(\mathbf{u})) \leq \frac{1}{2}C\log(r). \qquad (20)$$

Thus, from equations (18) and (20), we have

$$\frac{1}{k}(L_a^\epsilon(\mathbf{u}) - L_b^\epsilon(\mathbf{u})) \leq \Delta_2 + \frac{1}{2}C\log(r), \qquad (21)$$

for $\epsilon > 0$ sufficiently small.

Since Δ_2 converges in probability to $-C\log(r)$ with increasing k, it follows from equation (21) that the probability that $\frac{1}{k}(L_a^\epsilon(\mathbf{u}) - L_b^\epsilon(\mathbf{u}))$ is negative when $(u_1, \ldots, u_k)$ is generated under the CM-N_r model tends to 1 as k grows. That is, ML estimation under an NCM-N_r^ϵ model is statistically inconsistent, for data generated under the CM-N_r^ϵ model (or an NCM-N_r^ϵ model) when the branch lengths lie within the Felsenstein zone for tree a, and ϵ is chosen sufficiently small (relative to those branch lengths).

Steel M. 2010. Can we avoid "SIN" in the house of "no common mech-
anism"? Systematic Biology 60:96–109.

Figure 2 was not printed in Steel (2010). The figure
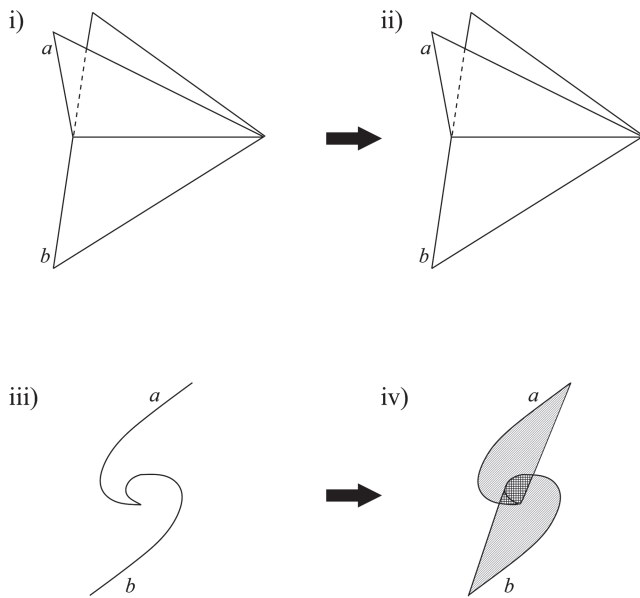appears below. The publisher apologizes for this error.

FIGURE 2. A Case (i and ii) where condition (4) fails, and one (iii
and iv) where it holds (see text for details).