

IDENTIFIABILITY OF A MARKOVIAN MODEL OF MOLECULAR EVOLUTION WITH GAMMA-DISTRIBUTED RATES

ELIZABETH S. ALLMAN,^{**} *University of Alaska Fairbanks*

CÉCILE ANÉ,^{***} *University of Wisconsin Madison*

JOHN A. RHODES,^{****} *University of Alaska Fairbanks*

Abstract

Inference of evolutionary trees and rates from biological sequences is commonly performed using continuous-time Markov models of character change. The Markov process evolves along an unknown tree while observations arise only from the tips of the tree. Rate heterogeneity is present in most real data sets and is accounted for by the use of flexible mixture models where each site is allowed its own rate. Very little has been rigorously established concerning the identifiability of the models currently in common use in data analysis, although nonidentifiability was proven for a semiparametric model and an incorrect proof of identifiability was published for a general parametric model (GTR + Γ + I). Here we prove that one of the most widely used models (GTR + Γ) is identifiable for generic parameters, and for all parameter choices in the case of four-state (DNA) models. This is the first proof of identifiability of a phylogenetic model with a continuous distribution of rates.

Keywords: Phylogenetics; identifiability

2000 Mathematics Subject Classification: Primary 60J25

Secondary 92D15; 92D20

1. Introduction

A central goal of molecular phylogenetics is to infer evolutionary trees from DNA or protein sequences. Such sequence data come from extant species at the tips of the tree—the tree of life—while the topology of the tree relating these species is unknown. Inferring this tree helps us understand the evolutionary relationships between sequences.

Phylogenetic data analysis is often performed using Markovian models of evolution: mutations occur along the branches of the tree under a finite-state Markov process. There is ample evidence that some places in the genome undergo mutations at a high rate, while other loci evolve very slowly, perhaps owing to some functional constraint. Such *rate variation* occurs at all spatial scales, across genes as well as across sites within genes. In performing inference, this heterogeneity is accounted for by the use of flexible mixture models where each site is allowed its own rate according to a rate distribution μ . In the context of molecular phylogenetics, the

Received 23 October 2007; revision received 25 November 2007.

* Postal address: Department of Mathematics and Statistics, University of Alaska Fairbanks, PO Box 756660, Fairbanks, AK 99775, USA.

** Email address: e.allman@uaf.edu

*** Postal address: Department of Statistics, University of Wisconsin Madison, Medical Science Center, 1300 University Avenue, Madison, WI 53706, USA. Email address: ane@stat.wisc.edu

**** Email address: j.rhodes@uaf.edu

use of a parametric family for μ is generally considered both advantageous and sufficiently flexible.

The question of identifiability for such a rate-variation model is a fundamental one, as standard proofs of consistency of statistical inference methods begin by establishing identifiability. Without identifiability, inference of some or all model parameters may be unjustified. However, since phylogenetic data is gathered only from the tips of the tree, understanding when one has identifiability of the tree topology and other parameters for phylogenetic models poses substantial mathematical challenges. Indeed, it has been shown that the tree and model parameters are *not* identifiable if the distribution of rates μ is too general, even when the Markovian mutation model is quite simple [13].

The most commonly used phylogenetic model is a *general time-reversible* (GTR) Markovian mutation model along with a gamma distribution family (Γ) for μ . For more flexibility, a class of invariable sites (I) can be added by allowing μ to be the mixture of a gamma distribution with an atom at 0 [4]. Numerous studies have shown that the addition to the GTR model of rate heterogeneity through Γ , I, or both, can considerably improve fit to data at the expense of only a few additional parameters. In fact, when model selection procedures are performed, the GTR+ Γ +I model is preferred in most studies. These stochastic models are the basis of hundreds of publications every year in the biological sciences—over 40 in *Systematic Biology* alone in 2006. Their impact is immense in the fields of evolutionary biology, ecology, conservation biology, and biogeography, as well as in medicine, where, for example, they appear in the study of the evolution of infectious diseases such as HIV and influenza viruses.

The main result claimed in the widely cited paper [11] is the following: the four-base (DNA) GTR + Γ + I model, with unknown mixing parameter and Γ shape parameter, is identifiable from the joint distributions of pairs of taxa. However, the proof given in [11] of this statement is flawed; in fact, two gaps occur in the argument. The first gap is in the use of an unjustified claim concerning graphs of the sort exemplified by Figure 3 of that paper. As this claim plays a crucial role in the entire argument, the statement above remains unproven. The second gap, though less sweeping in its impact, is still significant. Assuming that the unjustified graphical claim mentioned above could be proved, the argument of [11] still uses an assumption that the eigenvalues of the GTR rate matrix be distinct. While this is true for *generic* GTR parameters, there are exceptions, including the well-known Jukes–Cantor and Kimura two-parameter models [4]. Without substantial additional arguments, the reasoning given in [11] cannot prove identifiability in all cases.

Furthermore, bridging either of the gaps in [11] is not a trivial matter. Though we suspect that Rogers' statement of identifiability is correct, at least for generic parameters, we have not been able to establish it by his methods. For further exposition on the nature of the gaps, see Appendix A.

In this paper we consider only the GTR + Γ model, but for characters with any number $\kappa \geq 2$ states, where the case in which $\kappa = 4$ corresponds to DNA sequences. Our main result is the following.

Theorem 1. *The κ -state GTR + Γ model is identifiable from the joint distributions of triples of taxa for generic parameters on any tree with three or more taxa. Moreover, when $\kappa = 4$, the model is identifiable for all parameters.*

The term 'generic' here means for those GTR state distributions and rate matrices which do not satisfy at least one of a collection of equalities to be explicitly given in Theorem 2. Consequently, the set of nongeneric parameters is of Lebesgue measure 0 in the full parameter

space. Our arguments are quite different from those attempted in [11]. We combine arguments from algebra, algebraic geometry, and analysis.

We believe this paper presents the first correct proof of identifiability for any model with a continuous distribution μ of rates across sites that is not fully known. The nonidentifiability of some models with more freely varying rate distributions of rates across sites was established in [13]. That paper also showed identifiability of rate-across-sites models built upon certain group-based models provided that the rate distribution μ was completely known. More recently, Allman and Rhodes [1] proved that tree topologies were identifiable for generic parameters in rather general mixture models with a small number of classes. That result specializes to give the identifiability of trees for the κ -state GTR models with at most $\kappa - 1$ rates-across-sites classes, including the GTR + I model. Identifiability of numerical model parameters for GTR + I is further explored in [2]. There have also been a number of recent works dealing with nonidentifiability of mixture models which are not of the rates-across-sites type; these include [8], [9], [14], and [15].

In Section 2 we define the GTR + Γ model, introduce notation, and reduce Theorem 1 to the case of a three-taxon tree. In Section 3 we use purely algebraic arguments to determine from a joint distribution certain useful quantities defined in terms of the model parameters. In Section 4, in the generic case of certain algebraic expressions not vanishing, an analytic argument uses these quantities to identify the model parameters. Focusing on the important case of $\kappa = 4$ for the remainder of the paper, in Section 5 we completely characterize the exceptional cases of parameters not covered by our generic argument. Using this additional information, in Section 6 we establish identifiability for these cases as well. Finally, in Section 7 we briefly mention several problems concerning identifiability of phylogenetic models that remain open.

2. Preliminaries

2.1. The GTR+rates-across-sites substitution model

The κ -state across-site rate-variation model is parameterized as follows.

1. An unrooted topological tree T with all internal vertices of valence greater than or equal to 3 and with leaves labeled by a_1, a_2, \dots, a_n . These labels represent taxa and the tree represents their evolutionary relationships.
2. A collection of edge lengths $t_e \geq 0$, where e ranges over the edges of T . We require that $t_e > 0$ for all internal edges of the tree, but allow $t_e \geq 0$ for pendant edges, provided that no two taxa are a total-edge-length-distance 0 apart. Thus, if an edge e is pendant, the label on its leaf may represent either an ancestral ($t_e = 0$) or nonancestral ($t_e > 0$) taxon.
3. A distribution vector $\boldsymbol{\pi} = (\pi_1, \dots, \pi_\kappa)$ with $\pi_i > 0$ and $\sum \pi_i = 1$, representing the frequencies of states occurring in biological sequences at all vertices of T .
4. A $\kappa \times \kappa$ matrix $\boldsymbol{Q} = (q_{ij})$ with $q_{ij} > 0$ for $i \neq j$ and $\sum_j q_{ij} = 0$ for each i , such that $\text{diag}(\boldsymbol{\pi})\boldsymbol{Q}$ is symmetric. The matrix \boldsymbol{Q} represents the instantaneous substitution rates between states in a reversible Markov process. We will also assume that some normalization of \boldsymbol{Q} has been imposed, for instance, that $\text{diag}(\boldsymbol{\pi})\boldsymbol{Q}$ has trace -1 .

Note that the symmetry and row summation conditions imply that $\boldsymbol{\pi}$ is a left eigenvector of \boldsymbol{Q} with eigenvalue 0, which in turn implies that $\boldsymbol{\pi}$ is stationary under the continuous-time process defined by \boldsymbol{Q} .

5. A distribution μ , with nonnegative support and expectation $E(\mu) = 1$, describing the distribution of rates among sites. If a site has rate parameter r then its instantaneous substitution rates will be given by rQ .

Letting $[\kappa] = \{1, 2, \dots, \kappa\}$ denote the states, the joint distribution of states at the leaves of the tree T which arises from a rate-across-sites GTR model is computed as follows. For each rate r and edge e of the tree, let $M_{e,r} = \exp(t_e r Q)$. Then, with an arbitrary vertex ρ of T chosen as a root, let

$$P_r(i_1, \dots, i_n) = \sum_{(h_v) \in H} \left(\pi(h_\rho) \prod_e M_{e,r}(h_{s(e)}, h_{f(e)}) \right),$$

where the product is taken over all edges e of T directed away from ρ , edge e has initial vertex $s(e)$ and final vertex $f(e)$, and the sum is taken over the set

$$H = H_{i_1 i_2 \dots i_n} = \{(h_v)_{v \in \text{vert}(T)} \mid h_v \in [\kappa] \text{ if } v \neq a_j, h_v = i_j \text{ if } v = a_j\} \subset [\kappa]^{|\text{vert}(T)|}.$$

Thus, H represents the set of all ‘histories’ consistent with the specified states i_1, \dots, i_n at the leaves, and the n -dimensional table P_r gives the joint distribution of states at the leaves given a site has rate parameter r . Since the Markov process is reversible and stationary on π , this distribution is independent of the choice of root ρ .

Finally, the joint distribution for the GTR + μ model is given by the n -dimensional table

$$P = \int_r P_r d\mu(r).$$

The distribution for the GTR + Γ model is given by additionally specifying a parameter $\alpha > 0$, with μ then specialized to be the Γ -distribution with shape parameter α and mean 1, i.e. with scale parameter $\beta = 1/\alpha$.

2.2. Diagonalization of Q

The reversibility assumptions on a GTR model imply that $\text{diag}(\pi^{1/2})Q\text{diag}(\pi^{-1/2})$ is symmetric, and that Q can be represented as

$$Q = U \text{diag}(0, \lambda_2, \lambda_3, \dots, \lambda_\kappa)U^{-1},$$

where the eigenvalues of Q satisfy $0 = \lambda_1 > \lambda_2 \geq \lambda_3 \geq \dots \geq \lambda_\kappa$ [6], and U is a real matrix of associated eigenvectors satisfying the equivalent statements

$$UU^T = \text{diag}(\pi)^{-1}, \quad U^T \text{diag}(\pi)U = I. \tag{1}$$

Furthermore, the first column of U may be taken to be the vector $\mathbf{1}$.

While the λ_i are uniquely determined by these considerations, in the case in which all the λ_i are distinct the matrix U is determined only up to multiplication of its individual columns by ± 1 . If the λ_i are not distinct, eigenspaces are uniquely determined but U is not.

Our method of determining Q from a joint distribution will proceed by determining eigenspaces (via U) and the λ_i separately. Although the nonuniqueness of U will not matter for our arguments, the normalization determined by (1) will be used to simplify our presentation.

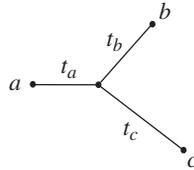


FIGURE 1: The unique three-taxon tree relating taxa a , b , and c with branch lengths t_a , t_b , and t_c .

2.3. Moment generating function

We also use the moment generating function (i.e. essentially the Laplace transform) of the density function for the distribution of rates in our model. As our algebraic arguments will apply to arbitrary rate distributions, while our analytic arguments are focused on Γ distributions, we introduce notation for the moment generating functions in both settings.

Definition 1. For any fixed distribution μ of rates r , let

$$L(u) = L_\mu(u) = E(e^{ru})$$

for $-\infty < u \leq 0$, denote the expectation of e^{ru} . In the special case of Γ -distributed rates, with parameters $\alpha > 0$ and $\beta = 1/\alpha$, let

$$L_\alpha(u) = L_{\Gamma,\alpha} = E(e^{ru}) = \left(1 - \frac{u}{\alpha}\right)^{-\alpha}.$$

Note that L , and in particular L_α , is an increasing function throughout its domain.

2.4. Reduction to the three-taxon case

To prove Theorem 1, it is sufficient to consider only the case of three-taxon trees.

Lemma 1. *If the statements of Theorem 1 hold for three-taxon trees then they also hold for n -taxon trees when $n > 3$.*

Proof. As the generic condition of Theorem 1 is a condition on π and \mathcal{Q} (see Theorem 2, below, for a precise statement), parameters on an n -taxon tree are generic if and only if the induced parameters on all induced three-taxon trees are generic.

If the model on three-taxon trees is identifiable for certain parameters then, from the joint distribution for a tree such as that of Figure 1, we may determine α , \mathcal{Q} , π , and the three edge-lengths t_a , t_b , and t_c . Thus, we may determine the pairwise distances $t_a + t_b$, $t_a + t_c$, and $t_b + t_c$ between the taxa. From an n -taxon distribution, by considering marginalizations to three taxa we may thus determine α , \mathcal{Q} , π , and all the pairwise distances between taxa. From all pairwise distances we may recover the topological tree and all the edge lengths by standard combinatorial arguments, as in [12].

3. Algebraic arguments

We now determine some information that we may obtain algebraically from a joint distribution known to have arisen from the GTR + μ model on a tree T relating three taxa. While in this paper we will only apply the results to the GTR + Γ model, we derive them at their natural level of generality. We therefore denote the moment generating function of the rate distribution by L , with its dependence on μ left implicit.

As marginalizations of the joint distribution correspond to the model on induced trees T' with fewer taxa, we work with trees with one, two, or three leaves.

If T' has only one leaf, it is simply a single vertex, and the distribution of states is therefore π . Thus, π is identifiable from a joint distribution for one or more taxa.

If T' has exactly two leaves, joined by an edge of length $t_e > 0$, then the joint distribution can be expressed as

$$P = \text{diag}(\pi) E(\exp(t_e r Q)) = \text{diag}(\pi) U \text{diag}(L(\lambda_1 t_e), \dots, L(\lambda_\kappa t_e)) U^{-1}.$$

Therefore, diagonalizing $\text{diag}(\pi)^{-1} P$ determines the collection of $L(\lambda_i t_e)$ and the columns of U up to factors of ± 1 . Since L is increasing, we may determine individual $L(\lambda_i t_e)$ by the requirement that

$$1 = L(0) = L(\lambda_1 t_e) > L(\lambda_2 t_e) \geq \dots \geq L(\lambda_\kappa t_e). \tag{2}$$

When the λ_i are distinct, this fixes an ordering to the columns of U . Regardless, we simply make a fixed choice of some U consistent with the inequalities in (2) and satisfying (1). We can further require that this choice of U be made consistently for all two-taxon marginalizations of the joint distribution. Thus, for any tree relating two or more taxa, we may determine the eigenspaces of Q via U and the value $L(\lambda_i d_{jk})$ for each i and pair of taxa a_j and a_k , where d_{jk} is the total edge-length distance between a_j and a_k .

For T with exactly three leaves, let a, b , and c be the taxa labeling them, with edge lengths as in Figure 1, and let X_a, X_b , and X_c denote the character states at these taxa. As in [3], denote by $P^{ab,\gamma}$ the square matrix containing the probabilities

$$P^{ab,\gamma}(i, j) = P(X_b = j, X_c = \gamma \mid X_a = i),$$

which can be computed from the joint distribution. But,

$$P^{ab,\gamma} = E(\exp(rt_a Q) \text{diag}(\exp(rt_c Q)_{\cdot\gamma}) \exp(rt_b Q)),$$

where $\exp(rt_c Q)_{\cdot\gamma}$ is the γ th column of matrix $\exp(rt_c Q)$, so

$$U^{-1} P^{ab,\gamma} U = E(\text{diag}(\exp(rt_a \lambda_1), \dots, \exp(rt_a \lambda_\kappa)) U^{-1} \text{diag}(\exp(rt_c Q)_{\cdot\gamma}) \times U \text{diag}(\exp(rt_b \lambda_1), \dots, \exp(rt_b \lambda_\kappa))).$$

Note that the j th column of

$$\text{diag}(\exp(rt_c Q)_{\cdot\gamma}) U$$

is the same as the γ th column of

$$\text{diag}(U_{\cdot j}) \exp(rt_c Q).$$

Thus, when (i, j) is fixed, the row vector formed by $U^{-1} P^{ab,\gamma} U(i, j)$ for $\gamma = 1, \dots, \kappa$ is

$$\mu^{ij} E(\exp(rt_a \lambda_i) \exp(rt_b \lambda_j) \exp(rt_c Q)), \tag{3}$$

where μ^{ij} is the row vector with

$$\mu^{ij}(k) = U^{-1}(i, k) U(k, j) = \pi(k) U(k, i) U(k, j).$$

Finally, multiplying (3) by U on the right, and setting $v^{ij} = \mu^{ij}U$, we see that the information brought by the triple of taxa $\{a, b, c\}$ amounts to the knowledge of

$$v^{ij} E(\exp(rt_a\lambda_i) \exp(rt_b\lambda_j) \text{diag}(\exp(rt_c\lambda_1), \dots, \exp(rt_c\lambda_\kappa))),$$

i.e. to the knowledge of each

$$E(\exp(rt_a\lambda_i) \exp(rt_b\lambda_j) \exp(rt_c\lambda_k)) = L(t_a\lambda_i + t_b\lambda_j + t_c\lambda_k)$$

for which $v^{ij}(k) \neq 0$.

This motivates the following notation, where, for conciseness, we let $U_{ij} = U(i, j)$. For $i, j, k \in [\kappa]$, let

$$v_{ijk} = \sum_l \pi_l U_{li} U_{lj} U_{lk}.$$

Note that while $v_{ijk} = v^{ij}(k)$, we prefer this new notation since the value of v_{ijk} is unchanged by permuting subscripts:

$$v_{ijk} = v_{ikj} = v_{jik} = v_{jki} = v_{kij} = v_{kji}.$$

Furthermore, since π can be determined from one-taxon marginalizations and U can be determined from two-taxon marginalizations, from a three-taxon distribution we may compute v_{ijk} for all i, j, k .

In summary, we have shown the following.

Proposition 1. *From a distribution arising from the GTR + μ model on the three-taxon tree of Figure 1, we may obtain the following information:*

- (i) π , from one-marginalizations;
- (ii) all matrices U which diagonalize Q as above, and, for all i , the values

$$L(\lambda_i(t_a + t_b)), \quad L(\lambda_i(t_a + t_c)), \quad \text{and} \quad L(\lambda_i(t_b + t_c))$$

from two-marginalizations; and

- (iii) the values $L(\lambda_i t_a + \lambda_j t_b + \lambda_k t_c)$ for all $i, j, \text{ and } k$ such that $v_{ijk} \neq 0$ for some such choice of U .

Note that (ii) can be obtained as a special case of (iii) by taking $j = i$ and $k = 1$, as it is easy to see that $v_{ii1} \neq 0$. We shall also see that $v_{ij1} = 0$ if $i \neq j$, so certainly some of the v_{ijk} can vanish.

One might expect that, for most choices of GTR parameters, all the $v_{ijk} \neq 0$ for $i, j, k > 1$. Indeed, this is generally the case, but for certain choices one or more of these v_{ijk} can vanish. The Jukes–Cantor and Kimura two-parameter and three-parameter models provide simple examples of this for $\kappa = 4$. For these models, we may choose

$$\pi = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right), \quad U = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix},$$

and $v_{ijk} \neq 0$ for $i, j, k > 1$ only when i, j , and k are distinct. While, for the Jukes–Cantor and Kimura two-parameter models, we may make other choices for U ; we can show that these alternative choices of U do not lead to the recovery of any additional information.

Nonetheless, for $\kappa \geq 3$, there is always some genuine three-taxon information available from a distribution, as we now show. Although we do not need the following proposition for the proof of Theorem 1, the method of argument it introduces underlies Section 5, below.

Proposition 2. *With $\kappa \geq 3$, for any choice of GTR parameters, there exists at least one triple $i, j, k > 1$ with $v_{ijk} \neq 0$.*

Proof. Suppose that, for all triples $i, j, k > 1$,

$$v_{ijk} = \sum_l \pi_l U_{li} U_{lj} U_{lk} = 0. \quad (4)$$

From (1) we also know that if $j \neq k$ then

$$v_{1jk} = \sum_l \pi_l U_{lj} U_{lk} = 0. \quad (5)$$

Both of these equations can be expressed more conveniently by introducing the inner product

$$\langle x, y \rangle = x^\top \text{diag}(\boldsymbol{\pi})y.$$

Then with U_i being the i th column of U and W_{jk} being the vector whose l th entry is the product $U_{lj}U_{lk}$, (4) gives the orthogonality statements

$$\langle U_i, W_{jk} \rangle = 0 \quad \text{if } i, j, k > 1,$$

while (5) yields both

$$\langle U_1, W_{jk} \rangle = 0 \quad \text{if } j \neq k \quad \text{and} \quad \langle U_j, U_k \rangle = 0 \quad \text{if } j \neq k.$$

In particular, we see, for $j, k > 1$ and $j \neq k$, that W_{jk} is orthogonal to all U_i , and so $W_{jk} = \mathbf{0}$. Considering individual entries of W_{jk} gives, for every l ,

$$U_{lj}U_{lk} = 0 \quad \text{for all } j, k > 1, j \neq k. \quad (6)$$

Now note that, for any $j > 1$, the vector U_j must have at least two nonzero entries. (This is simply because U_j is a nonzero vector, and $\langle \mathbf{1}, U_j \rangle = 0$ since $U_1 = \mathbf{1}$.) We use this observation, together with (6), to arrive at a contradiction.

First, without loss of generality, assume that the first two entries of U_2 are nonzero. Then, by (6), the first two entries of all the vectors U_3, U_4, \dots must be 0. But then we may assume that the third and fourth entries of U_3 are nonzero, and so the first four entries of U_4, \dots are 0. For the four-state DNA model, this shows that $U_4 = \mathbf{0}$, which is impossible.

More generally, for a κ -state model, we find that $U_k = \mathbf{0}$ as soon as $2(k - 2) \geq \kappa$. Note that, for $\kappa \geq 4$, this happens for some value of $k \leq \kappa$; thus, contradicting the fact that the U_k are nonzero. In the $\kappa = 3$ case the same argument shows that U_3 has only one nonzero entry, which is still a contradiction, since U_3 is orthogonal to $U_1 = \mathbf{1}$. Thus, the proposition is established for a κ -state model with $\kappa \geq 3$.

For $\kappa = 2$, the statement of Proposition 2 does not hold, as is shown by considering the two-state symmetric model with

$$\boldsymbol{\pi} = \left(\frac{1}{2}, \frac{1}{2}\right) \quad \text{and} \quad \mathbf{U} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

However, we can show that this is the only choice of $\boldsymbol{\pi}$ and \mathbf{U} for which $v_{222} = 0$.

4. Identifiability for generic parameters

We now complete the proof of the first statement in Theorem 1, the identifiability of the GTR + Γ model for generic parameters, which is valid for all values of $\kappa \geq 2$. As we now consider only Γ -distributed rates, we use the specialized moment generating function L_α in our arguments.

More precisely, we will establish the following result.

Theorem 2. *For $\kappa \geq 2$, consider those GTR parameters for which there exist some i, j , with $1 < i \leq j$, such that $v_{ijj} \neq 0$. Then restricted to these parameters, the GTR + Γ model is identifiable on three-taxon trees.*

Remark 1. Note that the conditions $v_{ijj} = 0$ are polynomial in the entries of \mathbf{U} and $\boldsymbol{\pi}$. Viewing the GTR model as parameterized by those variables together with the λ_i , then the set of points in parameter space for which $v_{ijj} = 0$ for some i, j with $1 < i \leq j$ forms a proper algebraic variety. Basic facts of algebraic geometry then implies that this set is of strictly lower dimension than the full parameter space. A generic point in parameter space therefore lies off this exceptional variety, and the exceptional points have Lebesgue measure 0 in the full parameter space.

Remark 2. For $\kappa = 2$, identifiability does not hold for the three-taxon tree if the generic condition that $v_{ijj} \neq 0$ for some $1 < i \leq j$ is dropped. Indeed, if $v_{222} = 0$ then, as noted in the last section, $\boldsymbol{\pi}$ and \mathbf{U} arise from the two-state symmetric model. Since there are only two eigenvalues of \mathbf{Q} , $\lambda_1 = 0$ and $\lambda_2 < 0$, the second of these is determined by the normalization of \mathbf{Q} . As the proof of Proposition 1 indicates, the only additional information we may obtain from the joint distribution is the three quantities

$$L_\alpha(\lambda_2(t_a + t_b)), \quad L_\alpha(\lambda_2(t_a + t_c)), \quad \text{and} \quad L_\alpha(\lambda_2(t_b + t_c)).$$

Since these depend on four unknown parameters α, t_a, t_b , and t_c , it is straightforward to see that the parameter values are not uniquely determined.

Our proof of Theorem 2 will depend on the following technical lemma.

Lemma 2. *Suppose that $c \geq a \geq d_1 > 0$ and $c \geq b > d_2 > 0$. Then the equation*

$$d_1^{-\beta} + d_2^{-\beta} - a^{-\beta} - b^{-\beta} - c^{-\beta} + 1 = 0$$

has at most one solution with $\beta > 0$.

Proof. The equation can be rewritten as

$$\left(\left(\frac{c}{d_1}\right)^\beta - \left(\frac{c}{a}\right)^\beta \right) + \left(\left(\frac{c}{d_2}\right)^\beta - \left(\frac{c}{b}\right)^\beta \right) + (c^\beta - 1) = 0. \tag{7}$$

Now a function $g(\beta) = r^\beta - s^\beta$ is strictly convex on $\beta \geq 0$ provided that $r > s \geq 1$, since $g''(\beta) > 0$. If $r = s$ then $g(\beta) = 0$ is still convex. Thus, when viewed as a function of β , the first expression on the left-hand side of (7) is convex, and the second expression is strictly convex. Also, for any $r > 0$, the function $h(\beta) = r^\beta - 1$ is convex, so the third expression in (7) is convex as well. Thus, the sum of these three terms, the left-hand side of (7), is a strictly convex function of β .

But a strictly convex function of one variable can have at most two 0s. Since the function defined by the left-hand side of (7) has one 0 at $\beta = 0$, it therefore can have at most one 0 with $\beta > 0$.

Proof of Theorem 2. For some $j \geq i > 1$, we are given $v_{ijj} \neq 0$. As $v_{ijj} = v_{jij}$, by Proposition 1 we may determine the values

$$\begin{aligned} D_{ijj} &= L_\alpha(\lambda_i t_a + \lambda_j t_b + \lambda_j t_c), \\ D_{jij} &= L_\alpha(\lambda_j t_a + \lambda_i t_b + \lambda_j t_c), \end{aligned}$$

as well as

$$\begin{aligned} C_k &= L_\alpha(\lambda_k(t_a + t_b)), \\ B_k &= L_\alpha(\lambda_k(t_a + t_c)), \\ A_k &= L_\alpha(\lambda_k(t_b + t_c)), \end{aligned}$$

for $k = 1, \dots, \kappa$.

Since L_α is increasing, for any $k > 1$, we can use the values of C_k and B_k to determine which of t_b and t_c is larger. Proceeding similarly, we may determine the relative ranking of t_a , t_b , and t_c . Without loss of generality, we therefore assume that

$$0 \leq t_a \leq t_b \leq t_c$$

for the remainder of this proof. Note, however, that if $t_a = 0$ then $t_b > 0$, by our assumption on the model parameters that no two taxa be a total-edge-length-distance 0 apart.

Observe that

$$L_\alpha^{-1}(D_{ijj}) + L_\alpha^{-1}(D_{jij}) = L_\alpha^{-1}(A_j) + L_\alpha^{-1}(B_j) + L_\alpha^{-1}(C_i),$$

or, using the formula for L_α and letting $\beta = 1/\alpha$,

$$D_{ijj}^{-\beta} + D_{jij}^{-\beta} - A_j^{-\beta} - B_j^{-\beta} - C_i^{-\beta} + 1 = 0. \tag{8}$$

Since $j \geq i > 1$, we have $\lambda_j \leq \lambda_i < 0$. Because L_α is an increasing function and $0 \leq t_a \leq t_b \leq t_c$ with $t_b > 0$, this implies that

$$C_i \geq A_j \geq D_{ijj} \quad \text{and} \quad C_i \geq B_j > D_{jij}.$$

Thus, applying Lemma 2 to (8) with

$$a = A_j, \quad b = B_j, \quad c = C_i, \quad d_1 = D_{ijj}, \quad \text{and} \quad d_2 = D_{jij},$$

we find that β is uniquely determined; so $\alpha = 1/\beta$ is identifiable.

Once α is known, for every k , we may determine the quantities

$$\begin{aligned} \lambda_k(t_a + t_b) &= L_\alpha^{-1}(C_k), \\ \lambda_k(t_a + t_c) &= L_\alpha^{-1}(B_k), \\ \lambda_k(t_b + t_c) &= L_\alpha^{-1}(A_k). \end{aligned}$$

Thus, we may determine the ratio between any two eigenvalues λ_k . As U is known, this determines Q up to scaling. Since we have required a normalization of Q , this means that Q is identifiable. With the λ_k now determined, we can find $t_a + t_b$, $t_a + t_c$, and $t_b + t_c$, and, hence, t_a , t_b , and t_c .

5. Exceptional cases ($\kappa = 4$)

In the previous section identifiability was proved under the assumption that $v_{ijj} \neq 0$ for some $j \geq i > 1$. We now specialize to the case in which $\kappa = 4$, and determine those GTR parameters for which none of these conditions hold. In the subsequent section we will use this information to argue that even in these exceptional cases the GTR + Γ model is identifiable.

Note that while we work only with a four-state model appropriate to DNA, the approach we use may well apply for larger κ , though we should expect additional exceptional subclasses to appear.

Lemma 3. *For $\kappa = 4$, consider a choice of GTR parameters for which $v_{ijj} = 0$ for all $j \geq i > 1$. Then, up to permutation of the states and multiplication of some columns of U by -1 , the distribution vector π and eigenvector matrix U satisfy one of the two following sets of conditions.*

Case A: $\pi = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ and, for some $b, c \geq 0$ with $b^2 + c^2 = 2$,

$$U = \begin{pmatrix} 1 & c & b & 1 \\ 1 & -c & -b & 1 \\ 1 & -b & c & -1 \\ 1 & b & -c & -1 \end{pmatrix}.$$

Case B: $\pi = (\frac{1}{8}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2})$ and

$$U = \begin{pmatrix} 1 & 2 & \sqrt{2} & 1 \\ 1 & -2 & \sqrt{2} & 1 \\ 1 & 0 & -\sqrt{2} & 1 \\ 1 & 0 & 0 & -1 \end{pmatrix}.$$

Proof. We use the notation of Proposition 2, including the inner product and definition of vectors W_{ij} given in its proof. Orthogonality and lengths will always be with respect to that inner product.

We will repeatedly use the fact that, for i, j with $1 < i \leq j$,

$$\langle W_{jj}, U_i \rangle = v_{ijj} = 0.$$

In particular, setting $j = 4$, we find that W_{44} is orthogonal to U_2, U_3 , and U_4 , and, hence, is a multiple of $U_1 = \mathbf{1}$. This implies that

$$U_4 = (\pm 1, \pm 1, \pm 1, \pm 1),$$

since U_4 has length 1. Without loss of generality, by possibly permuting the rows of U (which is equivalent to changing the ordering of the states in writing down the rate matrix Q) and then possibly multiplying U_4 by -1 , we need now only consider two cases: either

Case A: $U_4 = (1, 1, -1, -1)$, or

Case B: $U_4 = (1, 1, 1, -1)$.

We consider these two cases separately.

Consider case A. Since $U_1 = \mathbf{1}$ and $U_4 = (1, 1, -1, -1)$, the orthogonality of U_1 and U_4 gives

$$\pi_1 + \pi_2 - \pi_3 - \pi_4 = 0.$$

Since $\sum_{i=1}^4 \pi_i = 1$, this tells us that

$$\pi_1 + \pi_2 = \frac{1}{2} \quad \text{and} \quad \pi_3 + \pi_4 = \frac{1}{2}. \tag{9}$$

Now since W_{33} is orthogonal to both U_2 and U_3 , then W_{33} is a linear combination of U_1 and U_4 , and, hence, $W_{33} = (b^2, b^2, c^2, c^2)$. Thus,

$$U_3 = (\pm b, \pm b, \pm c, \pm c).$$

Since U_3 is orthogonal to both U_1 and U_4 , it is orthogonal to their linear combinations, and in particular to $(1, 1, 0, 0)$ and $(0, 0, 1, 1)$. Thus, by permuting the first two entries of U_i , and also permuting the last two entries of U_i , if necessary, we may assume that

$$U_3 = (b, -b, c, -c)$$

with $b, c \geq 0$. This orthogonality further shows that

$$b\pi_1 - b\pi_2 = 0 \quad \text{and} \quad c\pi_3 - c\pi_4 = 0.$$

Thus,

$$\pi_1 = \pi_2 \quad \text{or} \quad b = 0,$$

and

$$\pi_3 = \pi_4 \quad \text{or} \quad c = 0.$$

In light of (9) we have

$$\pi_1 = \pi_2 = \frac{1}{4} \quad \text{or} \quad b = 0,$$

and

$$\pi_3 = \pi_4 = \frac{1}{4} \quad \text{or} \quad c = 0.$$

In any of these cases U_3 has length 1, so

$$b^2(\pi_1 + \pi_2) + c^2(\pi_3 + \pi_4) = 1.$$

Together with (9) this gives

$$b^2 + c^2 = 2.$$

Now since U_2 is orthogonal to U_1, U_3 , and U_4 , we must have

$$U_2 = a \left(\frac{c}{\pi_1}, -\frac{c}{\pi_2}, -\frac{b}{\pi_3}, \frac{b}{\pi_4} \right)$$

for some a , and we may assume that $a > 0$. But the length of U_2 is 1 and U_2 is orthogonal to W_{22} , so

$$\frac{c^2}{\pi_1} + \frac{c^2}{\pi_2} + \frac{b^2}{\pi_3} + \frac{b^2}{\pi_4} = \frac{1}{a^2}, \tag{10}$$

$$\frac{c^3}{\pi_1^2} - \frac{c^3}{\pi_2^2} - \frac{b^3}{\pi_3^2} + \frac{b^3}{\pi_4^2} = 0. \tag{11}$$

If neither b or c is 0, so all $\pi_i = \frac{1}{4}$, then (10) tells us that $a = \frac{1}{4}$, as the statement of the theorem claims.

If $b = 0$ then we already know that $c = \sqrt{2}$ and $\pi_3 = \pi_4 = \frac{1}{4}$. But (11) implies that $\pi_1 = \pi_2$, so these are also $\frac{1}{4}$. We then find from (10) that $a = \frac{1}{4}$, and we have another instance of the claimed characterization of case A. Similarly, if $c = 0$, we obtain the remaining instance.

Consider case B. Since $U_1 = \mathbf{1}$ and $U_4 = (1, 1, 1, -1)$, the orthogonality of U_1 and U_4 implies that

$$\pi_1 + \pi_2 + \pi_3 - \pi_4 = 0.$$

Now W_{33} is orthogonal to U_2 and U_3 , and, hence, it is a linear combination of U_1 and U_4 . Thus, $W_{33} = (b^2, b^2, b^2, c^2)$, so

$$U_3 = (\pm b, \pm b, \pm b, c).$$

But U_3 is orthogonal to both U_1 and U_4 , and, hence, orthogonal to their linear combinations, including $(0, 0, 0, 1)$ and $(1, 1, 1, 0)$. This shows that $c = 0$ and that (possibly by permuting the first three rows of U and multiplying U_3 by -1) we may assume that $U_3 = b(1, 1, -1, 0)$ for some $b > 0$. Orthogonality of U_3 and U_1 then shows that

$$\pi_1 + \pi_2 - \pi_3 = 0.$$

Also W_{22} is orthogonal to U_2 and, hence, it is a linear combination of U_1 , U_3 , and U_4 , so $W_{22} = (d^2, d^2, e^2, f^2)$. Thus,

$$U_2 = (\pm d, \pm d, e, f).$$

However, since U_2 is orthogonal to U_1 , U_3 , and U_4 , it is orthogonal to $(0, 0, 0, 1)$, $(0, 0, 1, 0)$, and $(1, 1, 0, 0)$. Thus, we may assume that $U_2 = d(1, -1, 0, 0)$ with $d > 0$. Finally, orthogonality of U_2 and U_1 implies that

$$\pi_1 - \pi_2 = 0.$$

All the above equations relating the π_i , together with the fact that $\sum_{i=1}^4 \pi_i = 1$ gives

$$\boldsymbol{\pi} = \pi_1(1, 1, 2, 4) = \left(\frac{1}{8}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}\right).$$

We can now determine the U_i exactly, using the fact that they must have length 1, to show that U is as claimed.

6. Identifiability in the exceptional cases ($\kappa = 4$)

We now complete the proof of Theorem 1 by showing identifiability in cases A and B of Lemma 3. We do this by first establishing some inequalities for the eigenvalues of Q that must hold in each of these cases, using the assumption that the off-diagonal entries of Q are positive.

Note that as $U^{-1} = U^T \text{diag}(\boldsymbol{\pi})$, and the entries of $\boldsymbol{\pi}$ are positive, the positivity of the off-diagonal entries of \boldsymbol{Q} is equivalent to the positivity of the off-diagonal entries of the symmetric matrix

$$\tilde{\boldsymbol{Q}} = U \text{diag}(0, \lambda_2, \lambda_3, \lambda_4) U^T.$$

Lemma 4. For $\kappa = 4$, let $0 = \lambda_1 > \lambda_2 \geq \lambda_3 \geq \lambda_4$ denote the eigenvalues of a GTR rate matrix \boldsymbol{Q} . Then the following additional inequalities hold in cases A and B of Lemma 3.

Case A: If $bc \neq 0$ then $\lambda_4 > \lambda_2 + \lambda_3$, while if $bc = 0$ then $\lambda_4 > 2\lambda_2$.

Case B: $\lambda_4 > 2\lambda_2$.

Proof. For case A, we compute that

$$\tilde{\boldsymbol{Q}} = \begin{pmatrix} * & -\lambda_2 c^2 - \lambda_3 b^2 + \lambda_4 & -\lambda_2 bc + \lambda_3 bc - \lambda_4 & \lambda_2 bc - \lambda_3 bc - \lambda_4 \\ * & * & \lambda_2 bc - \lambda_3 bc - \lambda_4 & -\lambda_2 bc + \lambda_3 bc - \lambda_4 \\ * & * & * & -\lambda_2 b^2 - \lambda_3 c^2 + \lambda_4 \\ * & * & * & * \end{pmatrix},$$

where the ‘*’ indicates a quantity not of interest. From the positivity of the (1,2) and (3,4) entries of $\tilde{\boldsymbol{Q}}$, we thus know that

$$\lambda_4 > \max(\lambda_2 c^2 + \lambda_3 b^2, \lambda_2 b^2 + \lambda_3 c^2) \geq \frac{(\lambda_2 c^2 + \lambda_3 b^2) + (\lambda_2 b^2 + \lambda_3 c^2)}{2}.$$

Since $b^2 + c^2 = 2$, this shows that $\lambda_4 > \lambda_2 + \lambda_3$. In the case in which $bc = 0$, so $(b, c) = (0, \sqrt{2})$ or $(\sqrt{2}, 0)$, the first inequality gives the stronger statement of the lemma.

For case B,

$$\tilde{\boldsymbol{Q}} = \begin{pmatrix} * & -4\lambda_2 + 2\lambda_3 + \lambda_4 & -2\lambda_3 + \lambda_4 & -\lambda_4 \\ * & * & -2\lambda_3 + \lambda_4 & -\lambda_4 \\ * & * & * & -\lambda_4 \\ * & * & * & * \end{pmatrix}.$$

From the positivity of the off-diagonal entries, we see that

$$\lambda_4 > 2\lambda_3 \quad \text{and} \quad \lambda_4 + 2\lambda_3 > 4\lambda_2.$$

Together, these imply that $\lambda_4 > 2\lambda_2$.

We now return to proving identifiability for the exceptional cases. As in the proof of Theorem 2, we may determine the relative rankings of t_a, t_b , and t_c , and, therefore, assume that

$$0 \leq t_a \leq t_b \leq t_c$$

with $t_b > 0$.

In case A we find that $v_{234} = bc$, so we break that case into the following two subcases.

Subcase A1: $b, c \neq 0$.

Subcase A2: b or $c = 0$.

Consider subcase A1. In this case we find that $v_{ijk} \neq 0$ for all distinct $i, j, k > 1$. Letting

$$D_{342} = L_\alpha(\lambda_3t_a + \lambda_4t_b + \lambda_2t_c),$$

$$D_{423} = L_\alpha(\lambda_4t_a + \lambda_2t_b + \lambda_3t_c),$$

and $A_k, B_k,$ and C_k be as in the proof of Theorem 2, observe that

$$L_\alpha^{-1}(D_{342}) + L_\alpha^{-1}(D_{423}) = L_\alpha^{-1}(A_2) + L_\alpha^{-1}(B_3) + L_\alpha^{-1}(C_4).$$

Setting $\beta = 1/\alpha$ and using the explicit formula for L_α , yields

$$D_{342}^{-\beta} + D_{423}^{-\beta} - A_2^{-\beta} - B_3^{-\beta} - C_4^{-\beta} + 1 = 0. \tag{12}$$

Note that, by Proposition 1, all constants in this equation, except possibly β , are uniquely determined by the joint distribution.

In preparation for applying Lemma 2, we claim that the following inequalities hold:

$$D_{342} \leq A_2, \tag{13}$$

$$D_{423} < B_3, \tag{14}$$

$$D_{342} < C_4, \tag{15}$$

$$D_{423} < C_4. \tag{16}$$

Inequalities (13) and (14) follow easily from the fact that L_α is increasing. For inequality (15), first note that $\lambda_3t_a + \lambda_4t_b + \lambda_2t_c \leq (\lambda_2 + \lambda_3)t_a + \lambda_4t_b$. But Lemma 4 indicates that $\lambda_2 + \lambda_3 < \lambda_4$, so, again using the fact that L_α is increasing, the claim follows. Inequality (16) is similarly shown to hold.

Finally, to apply Lemma 2, let $d_1 = D_{342}$ and $d_2 = D_{423}$. The remainder of the constants in the lemma are chosen in one of three ways, depending on which of $A_2, B_3,$ and C_4 is largest.

- If $C_4 \geq A_2, B_3$ then let $a = A_2, b = B_3,$ and $c = C_4$.
- If $A_2 \geq C_4, B_3$ then let $a = C_4, b = B_3,$ and $c = A_2$.
- If $B_3 \geq C_4, A_2$ then let $a = A_2, b = C_4,$ and $c = B_3$.

Thus, in all subcases, from (12) we find that $\beta > 0$ is uniquely determined.

The remainder of the proof now proceeds exactly as for Theorem 2.

Consider subcase A2 and case B. In both of these cases $v_{224} \neq 0$, so, similarly to the previous case, letting

$$D_{422} = L_\alpha(\lambda_4t_a + \lambda_2t_b + \lambda_2t_c),$$

$$D_{242} = L_\alpha(\lambda_2t_a + \lambda_4t_b + \lambda_2t_c),$$

leads to

$$D_{422}^{-\beta} + D_{242}^{-\beta} - C_4^{-\beta} - A_2^{-\beta} - B_2^{-\beta} + 1 = 0.$$

By Proposition 1 we know that all quantities in this equation, except possibly β , are uniquely determined from the joint distribution.

We also note that the following inequalities hold:

$$A_2 \leq B_2, \tag{17}$$

$$D_{422} \leq A_2, \tag{18}$$

$$D_{242} < B_2, \tag{19}$$

$$D_{242} < C_4. \tag{20}$$

Inequalities (17)–(19) are implied by the fact the L_α is increasing. Inequality (20) will follow from $\lambda_2(t_a + t_c) < \lambda_4 t_a$. However, $\lambda_2(t_a + t_c) \leq 2\lambda_2 t_a < \lambda_4 t_a$ by Lemma 4.

To apply Lemma 2, let $d_1 = D_{422}$ and $d_2 = D_{242}$. In light of inequality (17), we need to assign the remaining constants according to the following two cases only.

- If $C_4 \geq B_2$, let $a = A_2$, $b = B_2$, and $c = C_4$.
- If $B_2 \geq C_4$, let $a = A_2$, $b = C_4$, and $c = B_2$.

In both cases we find that β is uniquely determined, and the proof of identifiability can be completed as in Theorem 2.

Thus, identifiability of the GTR + Γ model when $\kappa = 4$ is established for all cases.

7. Open problems

Many questions remain on the identifiability of phylogenetic models, including those commonly used for data analysis. Perhaps the most immediate one is the identifiability of the GTR + Γ + I model. Despite its widespread use in inference, no proof has appeared that the tree topology is identifiable for this model, much less its numerical parameters. Although our algebraic arguments of Section 3 apply, analogs for GTR + Γ + I of the analytic arguments we gave for GTR + Γ are not obvious. While the Γ rate distribution has only one unknown parameter, Γ +I has two, and this increase in dimensionality seems to be at the heart of the difficulty. Interestingly, empirical studies [16] have also shown that these parameters can be difficult to tease apart, as errors in their inferred values can be highly correlated in some circumstances. Although we conjecture that GTR + Γ + I is identifiable for generic parameters, we make no guess as to its identifiability for all parameters.

For computational reasons, standard software packages for phylogenetic inference implement a discretized Γ distribution [17], rather than the continuous one dealt with in this paper. While results on continuous distributions are suggestive of what might hold in the discrete case, they offer no guarantee. It would therefore also be highly desirable to have proofs of the identifiability of the discretized variants of GTR + Γ and GTR + Γ + I, either for generic or all parameters. Note that such results might depend on the number of discrete rate classes used, as well as on other details of the discretization process. So far, the only result in this direction is that of [1] on the identifiability of the tree parameter for generic numerical parameter choices when the number of rate classes is less than the number of observable character states (e.g. at most three rate classes for four-state nucleotide models, or at most 60 rate classes for 61-state codon models). As the arguments in that work use no special features of a Γ distribution, or even of an across-site rate-variation model, we suspect that stronger claims should hold when specializing to a particular form of a discrete-rate distribution.

Finally, we mention that beyond [1], almost nothing is known on identifiability of models with other types of heterogeneity, such as covarion-like models and general mixtures. As these are of growing interest for addressing biological questions [5], [7], [10], much remains to be understood.

Appendix A. The gaps in Rogers' proof

Here we explain the gaps in the published proof of Rogers [11] that the GTR + Γ + I model is identifiable. Since that paper has been widely cited and accepted as correct, our goal is to clearly indicate where the argument is flawed, and illustrate, through some examples, the nature of the logical gaps.

We emphasize that we do not *prove* that the gaps in the published argument cannot be bridged. Indeed, it seems most likely that the GTR + Γ + I model is identifiable, at least for generic parameters, and it is possible a correct proof might follow the rough outline of [11]. However, we have not been able to complete the argument Rogers attempts. Our own proof of the identifiability of the GTR + Γ model presented in the body of this paper follows a different line of argument.

We assume that the reader of this appendix will consult [11], as pinpointing the flaws in that paper requires rather technical attention to the details in it.

A.1. Gaps in the published proof

There are two gaps in Rogers' argument which we have identified. In this section we indicate the locations and nature of these flaws, and in subsequent ones we elaborate on them individually.

The first gap in the argument occurs roughly at the break from page 717 to page 718 of [11]. To explain the gap, we first outline Rogers' work leading up to it. Before this point, properties of the graph of the function $v^{-1}(\mu(x))$ have been carefully derived. An example of such a graph, for particular values of the parameters α , a , π , and p occurring in the definitions of v and μ , was shown in Figure 2 of the paper. For these parameter values and others, Rogers carefully and correctly showed that, for $x \geq 0$, the graph of $v^{-1}(\mu(x))$

1. is increasing;
2. has a single inflection point, where the graph changes from convex to concave (i.e. the concavity changes from upward to downward); and
3. has a horizontal asymptote as x tends to ∞ .

Although Rogers outlined other cases for different ranges of the parameter values, he highlighted the case when these three properties held.

At the top of page 718 of the paper, Figure 3 is presented, plotting the points whose coordinates are given by the pairs $(v^{-1}(\mu(\tau_1\lambda_i)), v^{-1}(\mu(\tau_2\lambda_i)))$ for all $\lambda_i \geq 0$. Here $\tau_2 > \tau_1$ are particular values, while α , a , π , and p are given the values leading to Figure 2. Rogers pointed out the following.

As in Figure 2, the graph [of Figure 3] has an inflection point, is concave upwards before the inflection point, and is concave downwards after the inflection point.

Then he claimed that 'Similar graphs will be produced for any pair of path distances such that $\tau_2 > \tau_1$.' However, he gave no argument for this claim. As the remainder of the argument strongly used the concavity properties of the graph of his Figure 3 (in the second column on page 718 the phrase '... as shown by Figure 3' appears), without a proof of this claim the main result of the paper is left unproved.

Judging from the context in which it is placed, a more complete statement of the unproved claim would be that, for any values of α , a , π , and p resulting in a graph of $v^{-1}(\mu(x))$ with the geometric properties of Figure 2 and any $\tau_2 > \tau_1$, the graph analogous to Figure 3 has a single inflection point. As no argument is given to establish the claim, we can only guess what the author intended for its justification. From what appears earlier in the paper, it seems likely that the author believed that the three geometric properties of the graph in Figure 2 enumerated above implied the claimed properties of Figure 3. However, that is definitely not the case, as we will show in Section A.2, below.

Note that we do not assert that the graphs analogous to Figure 3 for various parameter values are not as described in [11]. While plots of them for many choices of parameter values certainly suggest that Rogers' claim holds, it is of course invalid to claim a proof from examples. Moreover, with four parameters α , a , π , and p to vary, it is not clear how confident we should be of even having explored the parameter space well enough to make a solid conjecture. In light of the example we give in Section A.2, justifying Rogers' claim would require a much more detailed analysis of the functions ν and μ than Rogers attempted.

If this first gap in the proof were filled, a second problem would remain. Though less fundamental to the overall argument, this gap would mean that identifiability of the model would be established for *generic* parameters, but that there might be exceptional choices of parameters for which identifiability failed. ('Generic' here can be taken to mean for all parameters except those lying in a set of Lebesgue measure 0 in parameter space. More informally, for any reasonable probability distribution placed on the parameter space, randomly chosen parameters will be generic.)

Although the origin of this problem with nongeneric parameters is clearly pointed out by Rogers, it is open to interpretation whether he attempted to extend the proof to all parameter values at the very end of the paper. However, as the abstract and introductory material of [11] make no mention of the issue, this point at the very least seems to have escaped many readers attention.

This gap occurs because the published argument requires that the nonzero eigenvalues of the GTR rate matrix Q be three *distinct* numbers. On page 718, at the conclusion of the main argument, it is stated that

Therefore, if the substitution rate matrix has three distinct eigenvalues, the parameters of the I+ Γ rate heterogeneity will be uniquely determined.

The author then goes on to point out that, for the Jukes–Cantor and Kimura two-parameter models, this assumption on eigenvalues is violated, but '[f]or real data sets, however, it is unlikely that any two or all three of the eigenvalues will be exactly identical.'

Leaving aside the question of what parameters one might have for a model which fits a real data set well, Rogers here clearly indicated that his proof of identifiability up to this point omitted some exceptional cases. In the concluding lines of the paper he pointed out that these exceptional cases can be approximated arbitrarily closely by parameters with three distinct eigenvalues. While this is true, such an observation cannot be used to argue that the exceptional cases are not exceptional, as we will discuss below in Section A.3. It is unclear whether the concluding lines of [11] were meant to 'fill the gap' or not.

Of course, one might not be too concerned about exceptional cases. Indeed, if the first flaw were not present in his argument, then Rogers' proof would still be a valuable contribution in showing that for 'most' parameter values identifiability held. One might then look for other arguments to show identifiability also held in the exceptional cases. Nonetheless, it is disappointing that the exceptional cases include models such as the Jukes–Cantor and Kimura two-parameter models that are well known to biologists and might be considered at least reasonable approximations of reality in some circumstances.

A.2. A counterexample to the graphical argument

It seems that the origin of the first flaw in Rogers' argument is in a belief that the three enumerated properties he proved were exhibited in his Figure 2 result in the claimed properties of his Figure 3. In this section we show this implication is not valid, by exhibiting a function

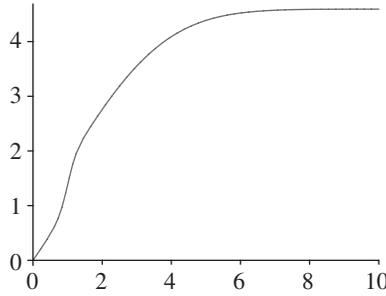


FIGURE 2: The graph $y = f(x)$.

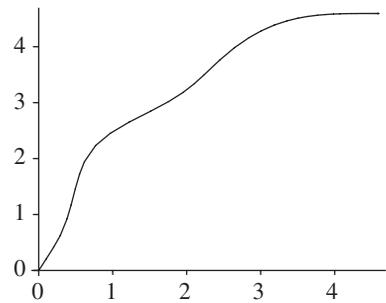


FIGURE 3: The points $(f(x), f(2x))$.

whose graph has the three properties, but when the graph analogous to Figure 3 is constructed, it has multiple inflection points.

Let

$$f(x) = \int_0^x \exp\left(\exp(-10(t - 1)^2) - \frac{(1 - t)^2}{10}\right) dt.$$

Then $f(0) = 0$ and

$$f'(x) = \exp\left(\exp(-10(x - 1)^2) - \frac{(1 - x)^2}{10}\right),$$

so $f'(x) > 0$ and f is increasing. Furthermore, we see that $f'(x)$ decays quickly enough to 0 as x tends to ∞ , so that $f(x)$ has a horizontal asymptote as x tends to ∞ .

To see that $f(x)$ has a single inflection point where the graph passes from convex to concave, it is enough to show that $f'(x)$ has a unique local maximum and no local minima. But this would follow from $g(x) = \ln(f'(x))$ having a unique local maximum and no local minima. Since

$$g(x) = \exp(-10(x - 1)^2) - \frac{(1 - x)^2}{10},$$

and the two summands here have unique local maxima at $x = 1$ and no local minima, g must have these features as well. Thus, f exhibits the enumerated properties of Rogers' Figure 2. For comparison, we plot f in our Figure 2.

The analog of Figure 3 for the function f would show the points $(f(\tau_1 x), f(\tau_2 x))$. If we choose $\tau_1 = 1$ and $\tau_2 = 2$, we obtain the graph shown in our Figure 3. Obviously, the curve

in Figure 3 has multiple—at least three—inflection points. Although we will not give a formal proof here that this curve has multiple inflection points, it is not difficult to do so.

A.3. Identifiability for generic parameters versus all parameters

The second gap in Rogers' argument arises because it is possible to have identifiability for generic parameters, but not for all parameters. Even if identifiability of the generic parameters had been proved, then one cannot easily argue that identifiability must hold for the nongeneric, exceptional cases as well. To illustrate this, we give a simple example.

Consider the map $\phi: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, defined by

$$\phi(a, b) = (a, ab).$$

Here a and b play the roles of 'parameters' for a hypothetical model, whose 'joint distribution' is given by the vector-valued function ϕ .

Suppose that (x, y) is a particular distribution which arises from the model (i.e. is in the image of ϕ), and we wish to find a and b such that $\phi(a, b) = (x, y)$. Then, provided that $x \neq 0$ (or, equivalently, $a \neq 0$), it is straightforward to see that a and b must be given by the formulae

$$a = x, \quad b = \frac{y}{x}.$$

Thus, for generic a and b (more specifically, for all (a, b) with $a \neq 0$), this hypothetical model is identifiable.

Note, however, that if $(x, y) = (0, 0)$, the situation is quite different. From $x = 0$ we see that we must have $a = 0$. But since $\phi(0, b) = (0, 0)$, we find that all parameters of the form $(0, b)$ lead to the same distribution $(0, 0)$. Thus, these exceptional parameters are *not* identifiable. Therefore, we have identifiability precisely for all parameters in the two-dimensional ab -plane *except* those lying on the one-dimensional line where $a = 0$. These exceptional parameters, forming a set of lower dimension than the full space, have Lebesgue measure 0 within it.

Note that even though there are parameter values arbitrarily close to the exceptional ones $(0, b)$ which are identifiable (for instance, (ε, b) for any small $\varepsilon \neq 0$), it is invalid to argue that the parameters $(0, b)$ must be identifiable as well.

This example shows that even if the first flaw in the argument of [11] were repaired, the approach outlined there will at best give identifiability for generic parameters. The final lines of that paper are not sufficient to prove identifiability for all parameter values.

Obviously, the function ϕ given here could not really be a joint distribution for a statistical model, since the entries of the vector $\phi(a, b)$ do not add to one, nor are they necessarily nonnegative. However, these features can be easily worked into a more complicated example. If one prefers a less contrived example then instances of generic identifiability of parameters but not full identifiability occur in standard statistical models used outside of phylogenetics (for instance, in latent class models). We have chosen to give this simpler example to highlight the essential problem most clearly.

Acknowledgements

ESA and JAR thank the Institute for Mathematics and its Applications and the Isaac Newton Institute, where parts of this work were undertaken, for their hospitality and funding. Work by ESA and JAR was also supported by the National Science Foundation (DMS 0714830).

References

- [1] ALLMAN, E. S. AND RHODES, J. A. (2006). The identifiability of tree topology for phylogenetic models, including covarion and mixture models. *J. Comput. Biol.* **13**, 1101–1113.
- [2] ALLMAN, E. S. AND RHODES, J. A. (2008). Identifying evolutionary trees and substitution parameters for the general Markov model with invariable sites. *Math. Biosci.* **211**, 18–33.
- [3] CHANG, J. T. (1996). Full reconstruction of Markov models on evolutionary trees: identifiability and consistency. *Math. Biosci.* **137**, 51–73.
- [4] FEISENSTEIN, J. (2004). *Inferring Phylogenies*. Sinauer Associates, Sunderland, MA.
- [5] GASCUEL, O. AND GUIDON, S. (2007). Modelling the variability of evolutionary processes. In *Reconstructing Evolution: New Mathematical and Computational Advances*, eds O. Gascuel and M. Steel, Oxford University Press, pp. 65–107.
- [6] HORN, R. A. AND JOHNSON, C. R. (1985). *Matrix Analysis*. Cambridge University Press.
- [7] KOLACZKOWSKI, B. AND THORNTON, J. (2004). Performance of maximum parsimony and likelihood phylogenetics when evolution is heterogeneous. *Nature* **431**, 980–984.
- [8] MATSEN, F. A. AND STEEL, M. A. (2007). Phylogenetic mixtures on a single tree can mimic a tree of another topology. *Syst. Biol.* **56**, 767–775.
- [9] MATSEN, F. A., MOSSEL, E. AND STEEL, M. (2008). Mixed-up trees: the structure of phylogenetic mixtures. To appear in *Bull. Math. Biol.*
- [10] PAGEL, M. AND MEADE, A. (2004). A phylogenetic mixture model for detecting pattern-heterogeneity in gene sequence or character-state data. *Syst. Biol.* **53**, 571–581.
- [11] ROGERS, J. S. (2001). Maximum likelihood estimation of phylogenetic trees is consistent when substitution rates vary according to the invariable sites plus gamma distribution. *Syst. Biol.* **50**, 713–722.
- [12] SEMPLE, C. AND STEEL, M. (2003). *Phylogenetics* (Oxford Lecture Ser. Math. Appl. **24**). Oxford University Press.
- [13] STEEL, M. A., SZÉKELY, L. AND HENDY, M. D. (1994). Reconstructing trees from sequences whose sites evolve at variable rates. *J. Comput. Biol.* **1**, 153–163.
- [14] ŠTEFANKOVIČ, D. AND VIGODA, E. (2007). Phylogeny of mixture models: robustness of maximum likelihood and non-identifiable distributions. *J. Comput. Biol.* **14**, 156–189.
- [15] ŠTEFANKOVIČ, D. AND VIGODA, E. (2007). Pitfalls of heterogeneous processes for phylogenetic reconstruction. *Syst. Biol.* **56**, 113–124.
- [16] SULLIVAN, J., SWOFFORD, D. L. AND NAYLOR, G. J. P. (1999). The effect of taxon sampling on estimating rate heterogeneity parameters of maximum-likelihood models. *Molec. Biol. Evolution* **16**, 1347–1356.
- [17] YANG, Z. (1994). Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: approximate methods. *J. Molec. Evol.* **39**, 306–314.