Check for updates

METHOD ARTICLE

REVISED Practical guidelines for Bayesian phylogenetic inference using Markov chain Monte Carlo (MCMC)

[version 4; peer review: 4 approved, 3 approved with reservations]

Joëlle Barido-Sottani[1], Orlando Schwery[2-4], Rachel C. M. Warnock[5], Chi Zhang[6], April Marie Wright[2]

[1]Institut de Biologie de l'ENS (IBENS), École normale supérieure, CNRS, INSERM, Université PSL, Paris, Île-de-France, 75005, France
[2]Department of Biological Sciences, Southeastern Louisiana University, Hammond, Louisiana, 70402, USA
[3]Department of Biological Sciences, Virginia Polytechnic Institute and State University, Blacksburg, Virginia, 24061, USA
[4]Department of Biological Sciences, Louisiana State University, Baton Rouge, Louisiana, 70803, USA
[5]GeoZentrum Nordbayern, Department of Geography and Geosciences, Friedrich-Alexander Universität Erlangen-Nürnberg, Erlangen, Bavaria, 91054, Germany
[6]Key Laboratory of Vertebrate Evolution and Human Origins, Institute of Vertebrate Paleontology and Paleoanthropology, Chinese Academy of Sciences, Beijing, 100044, China

v4 **First published:** 20 Nov 2023, **3**:204
https://doi.org/10.12688/openreseurope.16679.1

Second version: 28 Jun 2024, **3**:204
https://doi.org/10.12688/openreseurope.16679.2

Third version: 05 Aug 2024, **3**:204
https://doi.org/10.12688/openreseurope.16679.3

Latest published: 13 Oct 2026, **3**:204
https://doi.org/10.12688/openreseurope.16679.4

Abstract

Phylogenetic estimation is, and has always been, a complex endeavor. Estimating a phylogenetic tree involves evaluating many possible evolutionary histories that could explain a set of observed data, typically by using a model of evolution. Values for all model parameters need to be estimated as well. Bayesian phylogenetic inference is a state-of-the-art statistical method which allows us to co-estimate a phylogeny, evolutionary parameters, and obtain a detailed representation of the uncertainty associated with the estimates. Bayesian inferences commonly involve complex demographic models for the estimation of a tree, or incorporate data beyond genetic sequences, such as information from the fossil record. All Bayesian phylogenetic tree estimation software use the Markov Chain Monte Carlo (MCMC) algorithm, a method for approximating the posterior distribution of parameters in a mathematical model. While many researchers use MCMC in phylogenetic analyses, interpreting results and diagnosing problems with MCMC remain vexing issues to many biologists. We thus offer an overview of how MCMC is used in Bayesian phylogenetic inference, with a particular emphasis on complex hierarchical models, such as the fossilized birth-death (FBD)

Open Peer Review

Approval Status ✓ ✓ ? ✓ ✓ ?

	1	2	3	4	5	6
version 4 (revision) 13 Oct 2026						
version 3 (revision) 05 Aug 2024					✓ view	? view
version 2 (revision) 28 Jun 2024	✓ view					
version 1 20 Nov 2023	✓ view	✓ view	? view	✓ view		

1. **Daniel Casali**, Universidade de São Paulo, São Paulo, Brazil
2. **Zhirui Hu**, Gladstone Institute of Data Science and Biotechnology, San Francisco, USA
3. **Grégoire Clarté**, University of Helsinki,

model. We will discuss strategies to diagnose common MCMC problems and troubleshoot difficult analyses, in particular convergence issues. We will show how the study design, the choice of models and priors, but also technical features of the inference tools themselves can all be adjusted to obtain the best results. Finally, we will also discuss the unique challenges created by the incorporation of fossil information in phylogenetic inference, and present tips to address them.

Keywords
Bayesian phylogenetic inference, MCMC, troubleshooting, phylogenetic inference software, fossilized birth-death, total-evidence, BEAST2, MrBayes



This article is included in the Horizon 2020 gateway.



This article is included in the Evolution and Ecology gateway.



This article is included in the Marie-Sklodowska-Curie Actions (MSCA) gateway.



This article is included in the Evolutionary Biology collection.

Helsinki, Finland

4. **Sungsik Kong**, University of Wisconsin-Madison, Madison, USA
5. **Bruna M. Farina**, University of Fribourg, Fribourg, Switzerland
   Swiss Institute of Bioinformatics, Lausanne, Switzerland
6. **Peter Beerli**, Florida State University, Tallahassee, USA

Please see the article web page for additional reviewers.

Any reports and responses or comments on the article can be found at the end of the article.

Corresponding author: Joëlle Barido-Sottani (joelle.barido-sottani@m4x.org)

Author roles: Barido-Sottani J: Conceptualization, Investigation, Project Administration, Writing – Original Draft Preparation, Writing – Review & Editing; **Schwery O**: Investigation, Writing – Original Draft Preparation, Writing – Review & Editing; **Warnock RCM**: Conceptualization, Investigation, Writing – Original Draft Preparation, Writing – Review & Editing; **Zhang C**: Investigation, Writing – Original Draft Preparation, Writing – Review & Editing; **Wright AM**: Investigation, Project Administration, Writing – Original Draft Preparation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under the Marie Sklodowska-Curie grant agreement No. 101022928 to JBS. OS was funded by NSF DEB 2045842837 and a Swiss National Science Foundation Postdoc Mobility Fellowship (P500PB 203131). AMW was supported on NSF DEB 2045842 and NSF CIBR 2113425. CZ was funded by the National Natural Science Foundation of China (42172006).
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Barido-Sottani J *et al*. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Barido-Sottani J, Schwery O, Warnock RCM *et al.* **Practical guidelines for Bayesian phylogenetic inference using Markov chain Monte Carlo (MCMC) [version 4; peer review: 4 approved, 3 approved with reservations]** Open Research Europe 2026, **3**:204 https://doi.org/10.12688/openreseurope.16679.4

First published: 20 Nov 2023, **3**:204 https://doi.org/10.12688/openreseurope.16679.1

REVISED **Amendments from Version 3**

We have revised this version to address comments by reviewers Bruna Farina, Peter Beerli and David Černý (see our detailed response to each reviewer in the manuscript comments). Following their suggestions, we have made our introduction and definitions of the Bayesian concepts more rigorous, particularly in terms of the expected output, the convergence and the identifiability. We have also added more historical context and citations to the Bayesian framework, slice sampling and model selection methods. Finally, we have clarified technical details throughout the text.

Any further responses from the reviewers can be found at the end of the article

1. Introduction to MCMC

Phylogenetics has always had a fundamental problem. For any reasonable number of taxa, the number of possible topologies that could connect them quickly scales to be larger than the number of stars in the sky. For instance, even with only 10 taxa there are already more than 34 million rooted phylogenies possible (Cayley, 1889; Felsenstein, 1978; Degnan and Salter, 2005). Despite growing computational power, evaluating all of them is complex and as the number of taxa increases further, it quickly becomes intractable. And yet, increased taxon sampling is crucial to phylogenetic accuracy (Rannala *et al*., 1998; Hillis *et al*., 2003; Heath *et al*., 2008). One computational technique revolutionized our ability to enumerate and evaluate solutions in a Bayesian framework. That technique is Markov chain Monte Carlo (MCMC) (Metropolis *et al*., 1953). Though others exist, MCMC remains the most commonly applied technique for this type of evaluation. Here, we provide a broad level introduction to MCMC, before diving into much more detail in subsequent sections.

To understand MCMC, we must first take a step back and understand mathematical models. In a model, parameters describe what the researcher views as important facets of the process that generated our observed data. For example, in a phylogenetic model of molecular evolution, there may be a parameter governing the ratio at which transitions occur relative to transversions to generate an observed multiple sequence alignment. In most Bayesian models, parameters are usually random (also called stochastic) variables, meaning the value of a parameter is derived from an event with some element of randomness, such as a draw from a probability distribution or a coin flip. In the models we consider, most of the parameters are continuous, meaning they can take any value within their reasonable ranges. The uncertainty of a continuous parameter is described by a probability density function (e.g., a uniform or an exponential distribution), and the probability within a range of values is the area under the curve of the probability density function. For discrete parameters, such as the tree topology, each possible value of the parameter has a probability. We collectively use 'probability distribution' for both discrete and continuous parameters.

In a maximum likelihood (ML) estimation, we try to find the values for all our parameters that maximize the likelihood of the parameters given our data. Good ML solutions can be efficiently estimated through a number of mathematical techniques (Yang, 2014; Nguyen *et al*., 2015). In a Bayesian estimation, we estimate a distribution of the parameters that are plausible under our model given the data. In addition, Bayesian inferences integrate prior distributions, which describe our prior knowledge and understanding about the model and parameters, before having looked at the data. Prior distributions play a major role in phylogenetic inference and will be discussed in more details in the following sections. Bayesian inference thus offers a more complete picture of the results, integrating uncertainty in the results as well as existing information from previous studies (Figure 1). However, it is also more complex, because for many real world scenarios, the posterior distribution of plausible parameters cannot be calculated directly.

MCMC algorithms allow us to find the set of plausible solutions of a Bayesian inference, that is, an estimation of the posterior distribution of the parameters. The algorithm for MCMC sampling most frequently employed in phylogenetic studies is known as the Metropolis-Hastings (MH) algorithm (Metropolis *et al*., 1953; Hastings, 1970) (Figure 2), although other algorithms such as slice sampling (Neal, 2003) also exist and have been used for phylogenetic inference, usually in combination with MH. The general way the MH algorithm works is that a starting set of values is proposed for the parameters. This set is then scored according to some criterion. Then, one or more model parameters are perturbed, or changed. This could be a simple change, like increasing the value of one numerical parameter. In the case of phylogenetics, we often need to use more complex moves to propose new values for non-numeric objects like clades and trees (this will be described in section **Moves/operators/proposals**). This is where the term 'Monte Carlo' comes in: the city Monte Carlo is famous for its casinos and games of chance, which rely on randomness, just like the algorithm does when it perturbs the parameters pseudorandomly (at random within some set of conditions). The new value or set of values proposed will be re-scored according to the evaluation criterion. If it is better, this solution becomes the new current solution from which new moves will be performed. If the score of the proposed value is worse than the current, we still have a chance to accept it – this ensures that we broadly explore the entire parameter space and lower our chances of

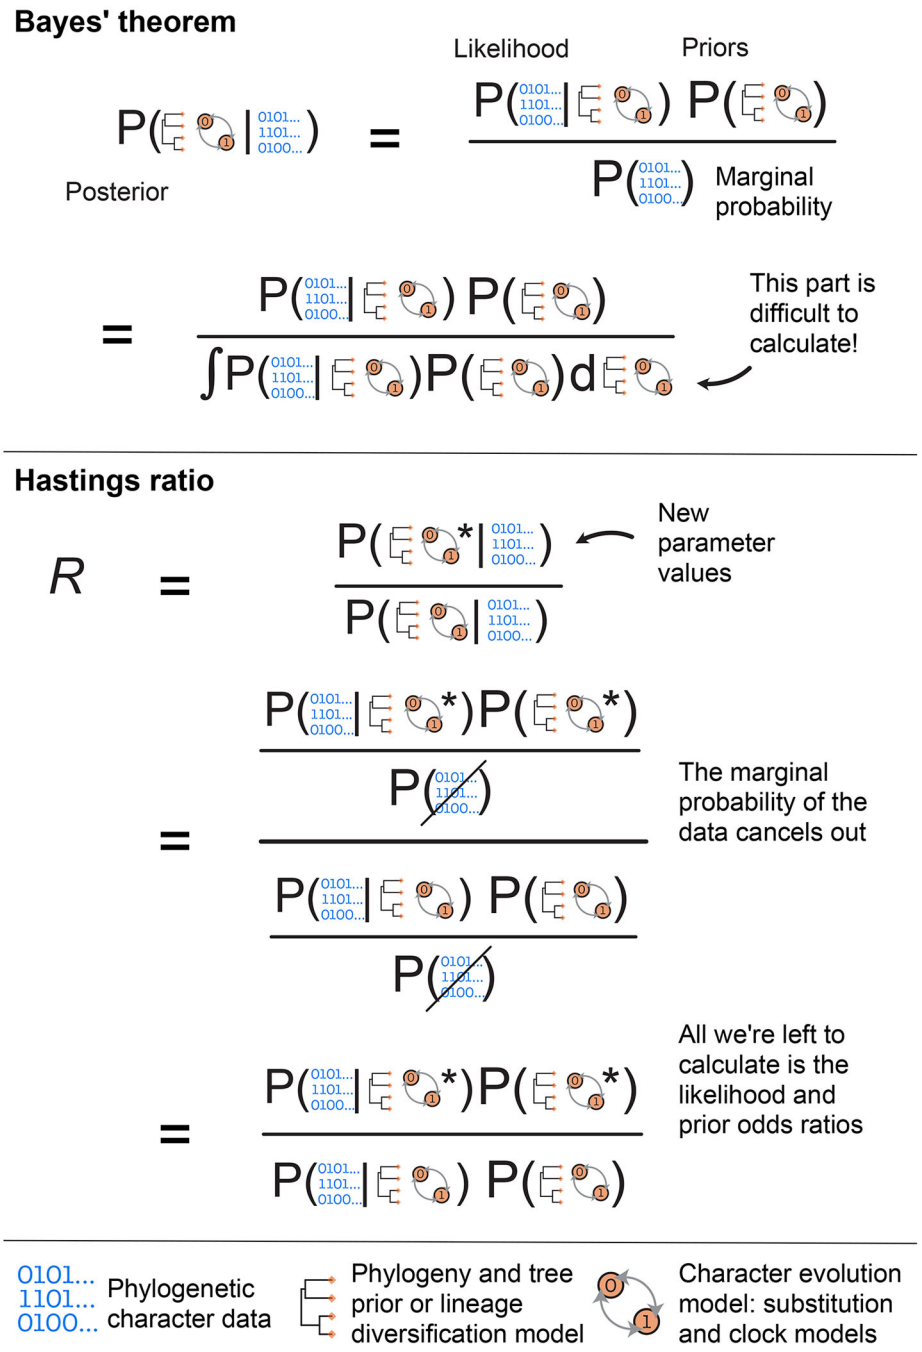

Figure 1. Schematic representation of Bayes' theorem and the Metropolis-Hastings ratio, used in Bayesian inference. The top panel shows Bayes' theorem and the relationship between the posterior, likelihood, priors and the marginal probability of the data. The second line shows an alternative way of writing the marginal probability, which illustrates more explicitly why the marginal probability is difficult to calculate. During MCMC we sample new parameter values (indicated by * in the figure) at each step and compare their posterior probability to the previous set of values using the Metropolis-Hastings or posterior odds ratio, *R*. The middle panel shows the Metropolis-Hastings ratio, and illustrates that since the marginal probability cancels out, we avoid having to calculate it during MCMC. For clarity we have excluded the Hastings part of the ratio; this would add $p(a \mid b)/p(b \mid a)$ to the calculation of *R*, which accounts for the potentially asymmetric probability of moving between *a* and *b*. The bottom panel provides a legend for the main data and model components used in Bayesian phylogenetic inference. The observed data used to infer phylogenies typically comprise discrete morphological characters (shown here) or molecular sequences. During inference the tree topology and branch lengths are evaluated according to a posterior distribution which combines prior and likelihood, and which depends on the type of inference – in particular, if the tree is dated we typically use lineage diversification models, such as birth-death processes. The character evolution model will always include the substitution process and will additionally include the clock model if we aim to infer a dated phylogeny.

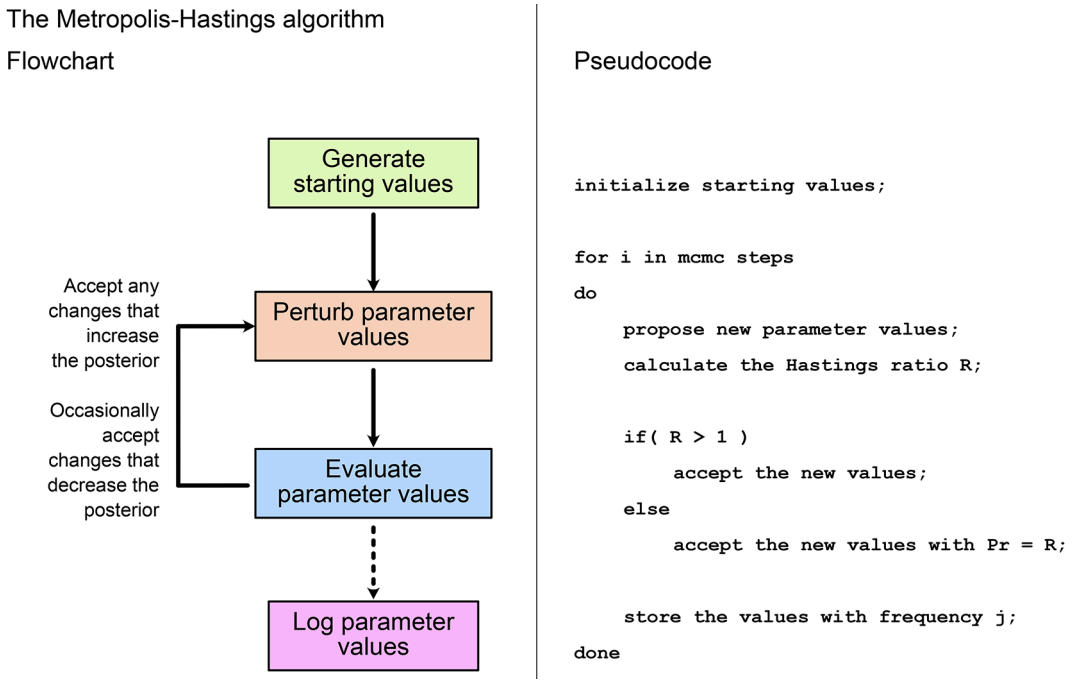

Figure 2. Flowchart and pseudocode showing the main steps in the Metropolis-Hastings algorithm. In the pseudocode, the total number of MCMC steps and the sampling frequency *j* are defined by the user, whereas the Metropolis-Hastings ratio *R* is calculated from the posterior distribution. See Figure 1 and section **The Metropolis-Hastings algorithm** for a full description of the Metropolis-Hastings ratio. Other variants on the MCMC, such as sequential Monte Carlo (SMC, Bouchard-Côté *et al.*, 2014) or Hamiltonian Monte Carlo (HMC, Dinh *et al.*, 2017), also exist but are beyond the scope of this review.

staying stuck in a local optimum. The probability of accepting the proposal depends on the difference of evaluation values between the new and current scores, so that much worse proposals mostly will be discarded. The 'Markov' chain part of the name comes from this being a Markov process, meaning a memoryless process. That is, the new state proposed depends only on the current state, not on the previous states. If a parameter value (or a region of values) has a high score, it will be visited many times in an analysis. In Bayesian phylogenetics, MCMC samples parameter values proportional to their posterior probability. Therefore, if a set of values for model parameters give a good solution according to the evaluation criterion, the MCMC will tend to sample those values and other similar values often. MCMC chains are ergodic, meaning that any state of the chain is theoretically reachable from any other state. To fulfil this criterion of ergodicity, we place moves on every parameter in our model, thereby enabling the algorithm to sample thoroughly. Finally, MCMC is sometimes referred to as a 'simulation' algorithm, which can be confusing. The reason for this is that we are not changing the underlying data, but proposing new values for model parameters to try and improve the fit of the model to the data. Often, this involves drawing parameter values out of a distribution, or scaling parameters in our model – both of these are forms of simulating new values.

Much like Bayesian analysis itself, MCMC was not developed to deal with phylogenetics, or even biological data directly. Those applications came later. Invented in the early 1950's, MCMC was originally used in physics to describe the equilibrium between the liquid and gas phases of a chemical (Metropolis *et al.*, 1953). The original algorithm assumed that moves in the parameter space were always symmetric, in other words that the probability of 2 going from parameter value a to value b has the same probability as going from b to a. Hastings (1970) introduced an additional term to balance potentially asymmetric proposals, leading to the more general version of the algorithm, the MH algorithm. From a humble beginning of trying to model a simple physical system, the MH MCMC algorithm drew the attention of statisticians, who popularized its use across nearly every quantitative discipline. In the following sections, we will discuss how MCMC works for phylogenetic inferences, how to troubleshoot an MCMC inference, and some tips and tricks for MCMC success.

2. MCMC inference applied to phylogenetics

2.1 The basics

Before we can understand MCMC in-depth, we need to discuss some basic information about Bayesian inference. Bayesian inference refers to a statistical framework for evaluating the fit of models and parameters to the observed data,

based on a quantity called the posterior distribution. The posterior distribution is calculated from three quantities: **the likelihood, the prior distribution,** and **the marginal probability** of the data. Bayes' Theorem is shown in Figure 1 and shows the relationship between these three quantities. We will first describe them and how they fit together, then move on to how MCMC is used in their calculation.

2.1.1 The likelihood

The likelihood of the models and parameter values describes how probable the observed data is given those models and values, i.e., how likely it is that those models and values represent the true generating process. If we are only concerned with the highest likelihood given the data, we usually do not need MCMC inference. Many phylogenetic tools can perform ML inference, which finds a set of values for the model parameters that maximize the probability of observing the data.

In a phylogenetic context, the data will usually be our observed molecular sequence alignment and/or morphological character matrix. The model will typically describe the process of evolution that generated these data. In a Bayesian phylogenetic inference, the calculation of the likelihood will include a substitution model, which describes the relative rate of change from one character state to another as well as the frequencies of each character state When inferring a time tree, the likelihood also includes a clock model, which describes the overall rate of change through time and across the tree. The substitution model is normalized so that the total relative rate of substitution equals 1. For example, the simplest substitution models are the Jukes-Cantor model (molecular data; Jukes and Cantor, 1969) and the Mk model (morphological data; Lewis, 2001). These models assume that one parameter describes the process of sequence evolution generating the data, and as a result these models are often referred to as 'all-rates-equal models'. Due to the normalization, this parameter is actually fixed, so that the JC69/Mk model has no free parameter. Many substitution models (such as the Kimura 2-parameter model (Kimura, 1980), the Felsenstein 1981 model (Felsenstein, 1981), the Hasegawa-Kishino-Yano model (Hasegawa *et al*., 1985), and the General Time-Reversible model (Tavaré, 1986)) are more complex, and reflect different assumptions regarding the hypothesized process of sequence change and evolution.

In a Bayesian analysis, the likelihood is one component of the three parts of Bayes' Theorem (Figure 1). It is calculated at each step in the MCMC analysis and is an important part used to estimate the posterior probability distribution given the data. The other important part is the prior.

2.1.2 The prior

A crucial analytical difference between a ML method and a Bayesian one is the presence of a prior. The term prior means that the distribution of the parameters reflects one's belief before observing the data. Each parameter in a Bayesian analysis has a prior probability distribution. For instance, we can set an exponential distribution on a given rate parameter. Under this prior, a rate that is very high is believed to be less likely than one that is very low. This means that rates are expected to be fairly low, but we still allow the possibility that they could be higher.

In Bayes' theorem, the prior and the likelihood are multiplied together, thus proposed parameter values are evaluated based on both the likelihood and the prior distribution. Therefore, if we expected a solution to be unlikely and thus specified a low prior probability for it, that low prior will lower the posterior when being multiplied with the likelihood. Importantly, however, if against our expectations, this solution is strongly supported by the data, the resulting high likelihood may overcome the effect of the low prior and still lead to high posterior support. This is how we can still find solutions which are different from our initial expectations, if the data suggest them. But this also highlights why we have to be careful not to specify priors that are too strict (i.e., that specify the prior probability of reasonable solutions to be 0), and prevent the MCMC from exploring the parameter space the data would favour. In addition, we note that in the context of Bayesian phylogenetics not all parameters of interest will be fully identifiable (that is, we may be able to only estimate a combination of several parameters but unable to fully separate the contribution of each parameter. Examples of this include rates and times, see section **Inferring dated trees and incorporating fossils**). This means that the priors we use in these analyses should be carefully considered and justified.

2.1.3 The marginal probability

The marginal probability of the data, also called the marginal likelihood, is the probability of the data without considering any particular model parameters, but conditioned on the models themselves and the constraints of the prior. Thus it gives the overall likelihood of the chosen model over all possible parameter values. This is usually the most challenging part of the calculation, as calculating the absolute probability of the data averaging over all possible values of the model

parameters is not computationally feasible in many cases. In a typical Bayesian phylogenetic inference, we avoid calculating the marginal probability using the MH algorithm (Figure 1, explained below). However, if we can calculate the marginal probability, it allows us to perform model selection. The marginal probability is typically computed by sampling many different solutions and averaging them for their probability. Different estimation methods have been developed to approximate the marginal likelihood, such as path sampling, which includes thermodynamic integration (Lartillot & Philippe, 2006), the widely used stepping-stone method (Xie *et al*., 2011; Baele *et al*., 2012) or nested sampling (Russel *et al*., 2018). Overall, these methods remain expensive, so recent efforts have been made to develop faster estimators for the marginal likelihood (Wang *et al*., 2023), although this estimator is currently only applicable with a fixed tree topology. Note that prior specificity matters for model selection, and overly-vague priors can cause issues for model selection and parameter estimation, even if the true parameter is included (Zwickl & Holder, 2004).

2.1.4 The posterior

The posterior distribution (posterior for short) is the probability distribution of the model parameters given the data. The posterior can change if the underlying data, model, or prior distributions change. As explained in the previous section, the theoretical posterior (i.e., the exact, 'true' solution) is almost always impossible to calculate directly. Hence we use MCMC to sample a set of parameter values that can approximate the posterior distribution of the parameters (usually called the posterior sample or MCMC sample), using the machinery introduced in section **Implementation of MCMC in phylogenetic inference software**. MCMC is key in Bayesian computation, as it allows us to sample from the posterior distribution. MCMC can even evaluate different potential model solutions through reversible-jump MCMC (Green, 1995), which allows the chain to move between different models (and their associated parameter spaces) during the inference. For instance, reversible-jump MCMC can be used to infer the most adequate substitution model for an alignment (Huelsenbeck *et al*., 2004; Bouckaert & Drummond, 2017). It is important to note that the result of an MCMC inference is the full posterior sample and the distribution of solutions. The distribution not only quantifies a point estimate (such as the mean or median), but also the uncertainty associated with it. In this sense, we typically evaluate the whole set of the posterior sample rather than an individual point from the sample. Along the same lines, it is also good to keep in mind that the MCMC does not search for 'true' parameter values, nor 'optimize' them, despite that being how we talk about it colloquially. It is merely sampling the parameter space to approximate their posterior distributions – and whether or not those reflect the true values also depends on whether the model adequately represents the process underlying the data.

2.1.5 The Metropolis-Hastings algorithm

The MH algorithm enables us to sample from the posterior without having to calculate the marginal probability of the data. The trick is that we use the *posterior odds ratio* or *Metropolis-Hastings ratio (R)* to evaluate how the chain proceeds, i.e., whether we accept the newly proposed values at each iteration. More specifically, this is the ratio of the posterior probabilities for the new values versus the current (parent) values. Since the marginal probability is the same in both cases, it cancels out when we calculate the ratio, meaning we only need to calculate the likelihood and the prior probability for each set of values, shown in Figure 1.

Figure 2 shows the main steps in the MH algorithm. As described in the section **Introduction to MCMC**, we first propose an initial set of values for all model parameters, including the topology (if estimating), and record the likelihood and prior probability associated with these. In each subsequent step, at least one model parameter is perturbed, and again we record the likelihood and prior probability. We evaluate the new values using the Hastings ratio. If $R > 1$, i.e., the new values improve the posterior, these are always accepted and become the updated current values from which the chain proceeds. If $R < 1$, the new values are only accepted with probability $= R$. This means, if the posterior associated with the new values is much lower, there is only a small chance of them being accepted. If the new values are not accepted, then the current values remain unchanged. By following these rules, the algorithm spends most of its time in regions of the parameter space with the highest posterior probability. We repeat the process of perturbation and evaluation until we have a sufficient number of MCMC samples to approximate the posterior. Several online tools demonstrate how this process works on different shapes of posterior space, such as the MCMC robot (Lewis, 2024) or the MCMC demo app (Feng, 2024). We do not need to store the values at every iteration, since consecutive values will be strongly correlated with each other. The best sampling frequency will depend on the size of the moves between iterations, and also on the specific software implementation used. For computational and memory reasons, we usually aim to record between 1,000 and 10,000 posterior samples. We discuss how to assess whether these samples are a good representation of the posterior later in this article.

2.1.6 The posterior sample

The posterior sample is a set of plausible solutions for a given dataset, derived through MCMC analysis. The posterior sample is composed of all recorded steps, which is a subsample of the steps visited by the inference. The distribution of solutions in the posterior sample is, itself, meaningful. Each entry sample in our posterior sample will have a posterior probability, and solutions will be sampled proportional to their posterior. A solution with a good posterior probability will be visited many times, whereas a solution with a poor one will be seldom seen in the posterior sample. How often a solution is sampled out of the total number of samples is often considered a measure of support. For example, a common measure of support for clades on a tree is the posterior probability, which corresponds to the proportion of trees in the posterior sample which contain that specific clade. A nice property of the posterior sample is that it not only provides the joint estimation of all the parameters, but also individual estimations for all the parameters. Indeed, taking only the sampled values for a specific parameter provides the *marginal posterior distribution* of this parameter, which allows us to estimate values for that parameter while integrating over all possible values of the other parameters. This means that all parameters of the inference can be analyzed independently.

2.2 Implementation of MCMC in phylogenetic inference software

2.2.1 Unrooted versus rooted trees

Phylogenetic trees exist in multiple forms. The first important distinction is between unrooted trees, which simply describe the evolutionary relationships of all the samples, and rooted trees, which include an explicit origin or starting point for the evolutionary process. In unrooted trees, internal nodes are not ordered, so it is impossible to know if an internal node is an ancestor or a descendant of another internal node. Another important feature of phylogenies is whether they are dated, i.e., whether their branch lengths are expressed in units of genetic/morphological distance or in units of time. Estimating a dated phylogeny requires a model for the molecular or morphological clock, a prior on the phylogeny, as well as time information to calibrate the tree. This information can be provided directly through the data, if the dataset includes samples from multiple points in time, such as fossil specimens or serially sampled viral sequences. Alternatively, the information can be provided as node calibrations, which provide information directly on the ages of specific nodes of the phylogeny. The prior on the phylogeny can be provided by a model of lineage diversification (usually a birth-death process or a coalescent process (Lambert & Stadler, 2013)), or simply use a uniform distribution on tree topologies (Ronquist *et al.*, 2012).

Dated trees are naturally rooted, as the earliest time point of the tree is obviously the origin of the process. Undated trees can also be rooted, by using one or more outgroup samples, i.e. samples which are known to be outside of the clade for which we want to build a phylogeny (for example, species from a different family or viral sequences from a different strain). In this case, the root is placed at the point in the tree where these outgroups diverge from the most recent common ancestor of the ingroup.

A much wider array of biological questions can be addressed using dated phylogenetic trees (e.g., diversification rate estimation or the application of phylogenetic comparative methods), but inferring dated trees increases the complexity of the analysis, making MCMC inference more challenging. Thus we mainly target this article at analyses which include a molecular or morphological clock as well as time information, although many of the tips detailed here are equally applicable to undated phylogenies.

2.2.2 General frameworks

Bayesian phylogenetic inference is often implemented in large software frameworks which group together many different models. These frameworks are generally designed to be modular, with each component of the analysis operating independently from the others. This is made possible by the structure of the models commonly used in phylogenetic inference, which can be defined independently from each other, and by the nature of the MCMC algorithm, which can operate on each parameter independently. The modular structure of the software means that any component, e.g., the substitution model, can be modified easily or extended without having to change anything else. It also means that core parts of the MCMC inference, for instance the MCMC algorithm itself, do not have to be reimplemented when a new model or a new type of data is introduced. In this paper, we chose to focus on BEAST2 (Bouckaert *et al.*, 2014), MrBayes (Huelsenbeck & Ronquist, 2001) and RevBayes (Höhna *et al.*, 2016) as our examples. We highlight these three frameworks as they focus on different use cases. MrBayes is the most straightforward to use and includes most of the common models, whereas BEAST2 includes more advanced models but remains fairly accessible to users through its graphical user interface BEAUti. Lastly, RevBayes is the most flexible and powerful framework, allowing users to implement new types of analyses directly, but it is also complex to use and can be prone to errors.

2.2.3 Moves/operators/proposals

As introduced earlier, MCMC inference relies on moving step by step through the parameter space and recording the state of the model parameters periodically. The recorded parameter states are the MCMC sample. Thus, any MCMC inference software must implement code to advance the chain. Since this code is not part of the model and usually does not depend on the specific choice of analysis, it is generally placed in separate components. In phylogenetic inference tools, these components can be called *proposals*, *moves*, or *operators. Different terms are used in different software implementations*, but they all perform the same function in the inference and they can be thus be treated as synonyms. Examples of some common phylogenetic moves are shown in Figure 3.

Moves are composed of three elements: first is the parameter or parameters they act on, meaning the parameters they change. Some moves only operate on one parameter at a time, while more complex moves can act on several (correlated) parameters at the same time. The second component of a move is the algorithm used to change the value of the parameter(s). These range from basic operations, such as proposing a new value using a sliding window centered on the current value, or scaling the current value of the parameter by a given factor, to much more complex ones such as those used to modify the tree. This component can also be called the 'proposal distribution', as it defines a distribution of

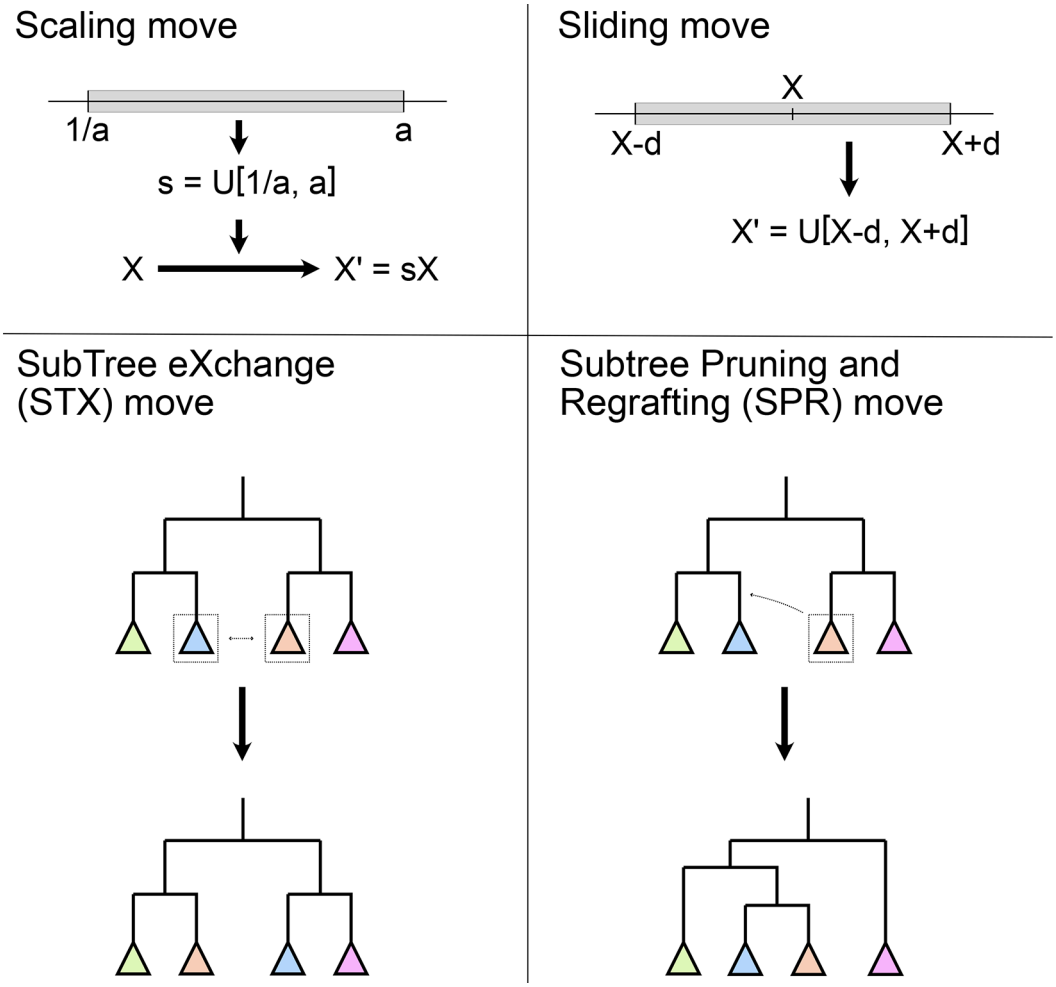

Figure 3. Examples of some common moves used in Bayesian phylogenetic inference. Scaling and sliding moves operate on a numerical parameter (*X*), such as the molecular clock rate, the speciation rate or the age of a fossil, and generally use a uniform distribution (*U*). Scaling moves use a scaling factor *s* chosen uniformly at random in the scaling range defined by the move parameter *a.* Sliding moves pick a new value uniformly at random in a window given by the window size parameter *d.* Subtree exchange (STX) and subtree pruning and regrafting (SPR) moves operate on the tree topology by moving subclades (represented as coloured triangles). A more detailed description of these moves can be found in Drummond *et al.* (2002).

possible values that can be proposed from the current parameter value(s). This proposal distribution depends on the move and its parameters, as well as the current position of the chain, but note that it is completely independent from the prior distribution on the parameter(s). The proposal ratio $p(b|a)/p(a|b)$, where b is the new parameter value and a the old value, is included in the MH ratio and contributes to the acceptance or rejection of the move. Finally, the third component of a move is its weight, which determines the frequency with which it will be used during the actual inference. A move with higher weight will be used more often, which should in principle lead to the corresponding parameters moving more often, and in turn provide more accurate estimates for these parameters. It should be noted that MCMC implementations differ in how weight is applied. Some attempt one move per step in the MCMC chain (e.g., BEAST2 and MrBayes), meaning only one parameter changes at a time and the weights represent the probability for any particular one to be chosen. Others move a whole set of parameters at each step (e.g., RevBayes, by default), with the weights representing how many times a move is attempted on average for a particular parameter during each step. This is also the reason why the number of generations that were run for a given analysis cannot always be compared directly between implementations, as one 'iteration' or 'generation' of the chain may actually imply different numbers of actual parameter moves (or attempted moves). Logically, implementations which use one move per iteration (such as BEAST2) will require many more iterations than implementations which use many different moves per iteration (such as RevBayes, although a single move per iteration schedule can be specified there too (Höhna *et al.*, 2017)).

However, the efficiency of the MCMC inference also depends on the acceptance proportion (also called the acceptance rate) of each move, i.e., the percentage of times that the move is accepted during the MCMC run. A move with a very low acceptance rate will have little impact on the overall inference, even if its weight is high. On the other hand, a very high acceptance rate can indicate that the move is proposing new values that are too close to the original values, which slows down the inference and increases the number of steps needed to properly explore the parameter space. An exception is the Gibbs sampler, which updates a parameter from the conditional distribution of all other parameters, thus has an acceptance ratio of 1, i.e., all proposals are accepted. However, the Gibbs move is seldom used in phylogenetics, as it is typically hard to derive the conditional distributions (but see Neal, 2003; Lartillot, 2006). For MCMC moves operating on a continuous numerical parameter, such as a branch length or evolutionary rate, the highest efficiency is typically achieved when the acceptance proportion is around 0.2 to 0.4 (Yang, 2014, section 7.3–7.4). Software implementations such as MrBayes, BEAST2, and RevBayes, typically provide an automatic tuning mechanism, which is enabled by default and adjusts each operator's configuration to reach the target acceptance proportion, say 0.3. For topological moves or moves which jump between different models, the efficiency is different from that of the more simple moves, and essentially depends on the specific design of the proposal algorithm. As a result, general users cannot easily optimize these moves. Good tree proposals are still under development, there is no perfect proposal to rule them all. In practice, using a collection of moves that make both big and small topological changes is advised. For example, MrBayes combines a nearest neighbour interchange move (NNI, a narrower implementation of subtree exchange (STX)) and two subtree pruning and regrafting (SPR) variants (see Figure 3) to update the tree. Tree moves should usually have much higher weights than the simple moves, as the tree space is tremendous.

The array of available moves in phylogenetic inference can be daunting for users. Luckily, most inference software proposes a default setup for standard analyses, which includes reasonable moves covering all parameters of the analysis. The default selection of moves usually leads to accurate estimates in a reasonable number of steps for most standard analyses, however, it certainly cannot fit all circumstances. We will see in later sections how to diagnose and adjust the move setup to help with misbehaving analyses.

3. Challenges of phylogenetic MCMC inference

As mentioned in the section **Introduction to MCMC**, MCMC was not developed for use in phylogenetics. It was developed for use with physics models, which usually have solely numerical components, often with many observations relative to the number of parameters. The use of MCMC for phylogenetics raises a new set of issues. In a phylogenetic analysis, we are often principally concerned with estimating a non-numeric parameter: the phylogeny itself! Due to the complexity of the underlying biological processes, we also often have high-dimensionality models, which contain a large number of parameters. This can raise serious performance issues for our MCMC inference when exploring the tree space and calculating the posterior probabilities We will now dive into some of these issues, and how MCMC inference has been adapted to work with phylogenetic trees and data.

3.1 Non-numeric data

As explained in the previous section, MCMC relies on perturbing our model parameters through *moves* (Figure 3). For numerical parameters, it is often very easy to perform a move. As shown in Figure 3, slide moves simply change the numeric value of a parameter within a window of a given size. Scale moves make values a bit bigger or smaller, while ensuring negative numbers stay negative and positive ones stay positive. For more complex cases, such as simplexes (sets

of values that must sum to a number, typically one − for instance nucleotide frequencies in a substitution model) or ratios, moves can be designed to ensure the conditions on the parameter are always met.

However, a tree is not a simple number or set of numbers, but a complex structure describing the arrangement of all the samples in a topology. As described in section **Moves/operators/proposals**, to explore the tree space, we thus need to change not only the branch lengths, but also the order and the composition of all subclades of the tree. Tree or topology moves propose rearrangements of the tree topology, and need to adjust or resample the associated branch lengths. Indeed, traversing tree space was a core challenge in developing phylogenetic applications of MCMC. Algorithms for efficiently sampling phylogenetic tree space became available in the late 1990s (Mau & Newton, 1997; Mau *et al.*, 1999; Li, 1996; Rannala & Yang, 1996; Larget & Simon, 1999), when Bayesian approaches for phylogenetics began to appear. However, recent years have seen the development of models to represent more complex evolutionary histories (for instance, involving hybridization or recombination), such as models involving multiple correlated trees or even phylogenetic networks (Kong *et al.*, 2022). Designing good topology moves for these more complex structures remains an issue.

3.2 High-dimensionality models

Biology is complex, and therefore, models to describe the behavior of biological systems will also tend to be complex. Think for a moment about a phylogenetic substitution model, for example, the GTR + Γ model. In this model, each nucleotide (A,C,T,G) has a different frequency, and the rates of substitution between all pairs of nucleotides are different. In addition different sites of the alignment have different overall rates of substitution, modelled by a gamma distribution. Applied in a Bayesian context, the model has many parameters: a tree topology, the branch lengths on the tree, exchangeability rates between nucleotides, equilibrium state frequencies of the nucleotides, the parameter of the gamma distribution representing among-site rate heterogeneity. For even a small tree with few samples, this is many parameters. In addition, some of these parameters may be correlated, for instance the branch lengths of a timed tree and the average clock rate have an inverse relationship. As a result, many posterior spaces in phylogenetic inference are in configurations referred to as "rugged" (Brown & Thomson, 2018), or having mixed areas of high probability ("peaks") and areas of low probabilities ("valleys"). This ruggedness can make it difficult to use MCMC in high-dimensional space. As shown on Figure 2, MCMC will generally refuse to take a step if the proposed solution will be much worse than the current one. Thus the inference can end up trapped in local optima. New computational methods are required to traverse these types of rugged spaces. For example, using proposal algorithms which perturb several correlated parameters at the same time can make it easier to find alternative peaks in the posterior surface. For instance, the up-down operator in BEAST2 will scale both the branch lengths of the tree, and the clock rate simultaneously. In most cases, however, overcoming these difficulties require the developers to implement more efficient MCMC moves, which can be challenging.

In addition to traversal issues, more complex models can also suffer from performance issues in the likelihood calculation itself. A common problem for lineage diversification models such as birth-death processes, for instance, is that we do not observe the parts of the phylogeny which have not been sampled. Thus we are missing a large part of the true evolutionary process. When calculating the prior probability of the phylogeny given the diversification model, we have to account for all possible histories in the unobserved parts of the tree. In more complex models, this calculation will frequently involve numerical integration, which is computationally very expensive and can suffer from numerical instability, meaning that the probability value cannot be estimated for some parameter configurations. Although this issue can be improved by smart implementation of the models (see for instance the work done by Scire *et al.* (2022) on the BEAST2 package BDMM), it represents a fundamental limitations for more complex processes.

3.3 Inferring dated trees and incorporating fossils

Inferring dated trees is substantially more challenging than non-time constrained tree inference. First, phylogenetic characters do not contain information about absolute time, meaning rates and times are not fully identifiable. This means the posterior will be sensitive to the time priors, even when we have a large amount of sequence data (Yang & Rannala, 2006; Rannala & Yang, 2007). Second, it requires the addition of a clock model and uses more complex tree models, usually coalescent or birth-death process models (Thorne *et al.*, 1998; Drummond *et al.*, 2006; Yang & Rannala, 2006). It also requires additional time information. In macroevolutionary phylogenies, this time information generally comes from the fossil record, either in the form of node calibrations, or by directly including fossil specimens in the inference (sometimes called tip calibrations). Tip-calibrated analyses provide a better representation of the uncertainty associated with the phylogenetic placement of fossils, and arguably involve less subjective user choices, such as the choice of the distribution used for node calibrations (Ronquist *et al.*, 2012) (though more objective approaches for generating node calibrations are available, see Wilkinson *et al.*, 2011; Friedman *et al.*, 2013; Matschiner *et al.*, 2017). Further, including fossils directly into the tree presents specific challenges.

There are two main sources of uncertainty associated with fossils that should be considered in Bayesian inference: taxonomic or topological uncertainty and fossil age uncertainty. Inference under the fossilized birth-death (FBD) process can incorporate both phylogenetic and age information (Heath *et al.*, 2014; Stadler, 2010). And because the model incorporates the fossil sampling process explicitly, extinct samples can be recovered as tips or *sampled ancestors* along internal branches. This requires special moves that propose changes to the total number of nodes in the tree, since each sampled ancestor reduces the number of tips by one (Gavryushkina *et al.*, 2014; Heath *et al.*, 2014). In terms of data, we have two alternative options for informing the position of extinct samples within the tree. First, fossils with no character data can be assigned to a node using topological constraints. Constraints can be based on evidence from previous phylogenetic analyses or descriptive taxonomy. Using this strategy, the position of the fossil below the constraint node is sampled using MCMC. The precise position of the fossil cannot be inferred without character data, but the posterior output will reflect the uncertainty associated with fossil placement below the constraint node.

Alternatively, if morphological character data is available for fossil and extant samples, we can use a 'total-evidence' approach. Using this strategy, fossil placement can be sampled using MCMC and the position of taxa with character data can be inferred (Barido-Sottani *et al.*, 2023a; Gavryushkina *et al.*, 2017; Zhang *et al.*, 2016). This approach is conceptually preferable, since it more directly accounts for the phylogenetic uncertainty associated with fossils. In practice, however, character data is not available or limited for most groups (many discrete morphological matrices contain <100 characters) and, unlike DNA, character states can be subjective and uncertain (Wright, 2019). Continuous character data, which might be more readily available for some groups, can also be used in total-evidence analyses (Álvarez-Carretero *et al.*, 2019; Zhang *et al.*, 2024), although the use of this data type for phylogenetic inference needs to be carefully considered (Varón-González *et al.*, 2020).

Fossil age uncertainty is straightforward to incorporate into Bayesian phylogenetic inference using the FBD process. Fossils are dated to within a known geological interval and the bounds of this age range (i.e., the minimum and maximum ages) can be used to inform priors on fossil ages. The age of fossils is then sampled during MCMC, therefore accounting for this uncertainty. This is preferable to fixing fossil ages to a point estimate within the known range of uncertainty, which can lead to erroneous parameter estimates (Barido-Sottani *et al.*, 2019, 2020). In fact, fossil ages can be even be estimated using this approach (Drummond & Stadler, 2016; Barido-Sottani *et al.*, 2023b). Typically, a uniform distribution is used to model the age uncertainty associated with fossils, between the minimum and maximum possible ages based on stratigraphic and radiometric evidence. Although additional information could be used to construct more informative non-uniform priors on fossil ages, in practice, it could be difficult to do this objectively.

Finally, we note that because rates and times are not fully identifiable, the way in which temporal information is included can have a large impact on the posterior results (see Section **Choice of model and priors**). The uncertainty associated with fossil calibrations can be reduced with the addition of phylogenetic character data but can not be entirely eliminated, even when the alignment becomes extremely large (Yang & Rannala, 2006; Inoue *et al.*, 2010). For an excellent illustration of this see dos Reis *et al.* (2016, Figure 1).

4. Troubleshooting tools and techniques

4.1 How do I know if my MCMC is good?

Before we talk about troubleshooting, we first must figure out how we even know if there is anything to troubleshoot. We generally consider an MCMC inference to be complete when it reaches what is termed *convergence*. This is typically when a chain has arrived in its *stationary distribution*, that is, when additional sampling no longer affects the distribution of state values estimated. In plain language, once you are in the stationary distribution, you can do moves and change individual parameters, but the overall distribution of values will not change. The goal is to find this stationary distribution for all the parameters in your analysis. At the very least, users should ensure that the parameters primary interest to their research questions, along with the prior, likelihood and posterior, have converged satisfactorily. The phase before the chain has converged is called *burn-in.* The samples collected during burn-in should be discarded, usually 10–30% of the chain length, only keeping the remaining samples for the parameter estimation. The burn-in proportion is usually set by default. For instance, MrBayes discards the beginning 25% samples as burn-in. However, it is important to check that these values are adapted for each analysis. Setting the burn-in too low will result in biased estimates of the posterior distribution, since we include samples that are not part of the stationary distribution – hence the importance of checking the trace and discarding extra burn-in if needed. Setting the burn-in too high is less damaging with regards to bias in the posterior, but wastes computing time, while removing useful samples from the chain and thus decreasing the amount of information given by the inference.

This sounds easy on the surface, but much ink has been spilled on appropriate ways of diagnosing whether or not our analysis has converged. Assessing convergence is usually done with convergence diagnostics. These are summary statistics that tell the researcher about how the MCMC inference, or chain, has performed and if it has converged. By far, the most commonly used diagnostic in phylogenetics is the *Effective Sample Size*, or ESS, which is calculated for every estimated parameter in the analysis separately (as well as for the likelihood, prior, and posterior).

When we perform MCMC inference, each time we do a move, we draw new values for one or more parameters, then accept or reject these values (Figure 2). This is often called an MCMC step. Different software implementations and models will require different numbers of steps to reach convergence. You might think that the number of steps would be equivalent to the number of samples in the posterior sample. But in an MCMC chain, different steps will be correlated with one another. This is referred to as *autocorrelation*, and is the result of the fact that the parameter values present at step i are used to propose the parameter values for step $i + 1$ (Figure 2). The ESS is specific to a posterior sample and to a given parameter, and describes the number of uncorrelated (independent) samples that would be needed to approximate the posterior distribution of a parameter with similar precision to that posterior sample. It is usually defined as $ESS = N/\tau$, in which N is the number of posterior samples and τ is the autocorrelation time. i.e. the number of steps between independent samples. For continuous parameters, τ can be estimated through numerical methods (Box *et al.*, 2015; Thompson, 2010). Due to autocorrelation, the ESS is typically smaller than the number of recorded samples in the MCMC chain, because the difference between two successive samples is usually quite small. If we were drawing completely independent samples, the difference between sample i and sample $i + 1$ could be quite large (i.e., an independent sample could be drawn from anywhere in parameter space, so a series of such samples may explore the different areas of that space more quickly than when done step by step by an MCMC chain).

In practice, this means that we assess whether an MCMC chain has converged by checking whether the ESS for its parameters is above a certain number, meaning the chain has become stationary for long enough when we account for its autocorrelation. An ESS of over 200 has become the *de facto* standard in biological analyses, though reasons for this are largely arbitrary (but see section **Convenience**). Eventually, we would want the ESS for all parameters to be above the threshold, but there are some more insights to be gained along the way. The ESS of the likelihood and posterior are often treated as the indicator for overall convergence of the analysis, while the individual ESS of each estimated parameter is used to diagnose the specific causes of convergence issues. When inspecting the ESS along with the trace (e.g., in Tracer), they should tell a similar story: parameters with ESS above the threshold should show a stationary trace too, while for parameters with low ESS, the trace may not be stationary yet, and may in fact provide insights as to the nature of the problem. If the trace looks stationary but the ESS is still below the threshold, this may simply imply that it needs to run for a bit longer. Together, ESS and the trace should highlight which parameters struggle to converge, so that they can be targeted for adjustments. However, in practice it is not always that simple, since parameters are interdependent and affect each other's convergence, so this needs to be kept in mind when searching for the culprit. It is also worth noting that certain types of parameters are by nature unlikely to ever converge, because they are not 'proper' parameters but references which do not point to the same parameter in all samples. A prime example for this would be individual branch rates, if the tree is jointly inferred (e.g., in an FBD analysis) – since changes in tree topology change the position and numbering of branches, the trace output for these does not strictly follow the same branch after each step, thus there is no expectation of them converging along with the rest of the chain. In this case, the ESS of the overall variance and mean of branch rates is likely a better indicator of convergence.

It should be noted that the connection between ESS and convergence is merely an assumption: ESS quantifies whether the chain has been stationary long enough to provide a precise estimate of the mean (i.e., low standard error), which is expected to coincide with stationarity as the chain converges. But a chain can get stuck in a local *optimum* due to the presence of several 'peaks' in the posterior (i.e., multimodality). Thus, as the ESS is based purely on samples present in the chain, it can reach acceptable values even if it has been stuck for a long time and failed to reach the highest peak(s). A simple way to check for convergence, in spite of local optima, is to run several different chains for the same analysis. MCMC chains which use the same data, models, and priors are guaranteed to converge on the same distribution (if run long enough), independent of the starting values used. Thus running multiple chains from different starting values and checking if the results obtained match is a good way to assess whether the analysis has converged. Note also that posterior samples from all chains can be combined together in the final result (unless the chains were heated, see below), thus the time spent on the different chains is not wasted.

In the next section, we will discuss software and tools for assessing ESS that were developed for Bayesian phylogenetics, as well as other avenues for understanding convergence issues. Other tools exist that were not developed with phylogenetics in mind, but are nonetheless also very useful, e.g., the R package coda (Plummer *et al.*, 2006). Despite these tools, the assessment of MCMC convergence remains a challenging task in phylogenetic analyses, and future

developments are needed to resolve some of these challenges. A promising example might be the use of L-lag couplings for Markov chains (Biswas *et al.*, 2019) to estimate mixing and convergence jointly across the parameters of the tree (Kelly *et al.*, 2023a, 2023b).

4.2 Tools of the trade

4.2.1 Tracer

Tracer (Rambaut *et al.*, 2018) is one of the most commonly used pieces of software for convergence assessment, due to the ease with which it can be used. A log file of sampled solutions from the MCMC can be read in. In its default view, a list of parameters in the model and their ESS value can be seen, as well as estimates of the value (mean, median, and spread) for each parameter sampled. Tracer automatically flags ESS values below a threshold of 200. Although this threshold value is somewhat arbitrary, it has been widely accepted in current practice as offering a good trade-off between convergence and computational cost of the inference.

The trace panel, however, is most useful for debugging convergence issues (see the next section for some common issues). The trace window shows the values sampled for each parameter over the MCMC run. An example of different traces can be seen in Figure 4. Ideally, the trace will appear as what is often termed the "hairy caterpillar" (Figure 4, last row). This is a sample that is well-converged. This pattern is generated by finding a good solution (or a set of good solutions) and sampling around that solution. Typically when this happens, the run has reached its stationary distribution.

4.2.2 RWTY

Tracer remains the most used software for convergence, but it does not calculate an effective sample size for the most important model parameter – the tree itself. The ESS of the overall posterior or the ESS of parameters tied to the tree, such as the tree height or MRCA ages of specific clades, can be used as indirect signs of the (lack of) convergence of the phylogeny; however, it is preferable to have a direct indicator. The R package RWTY (aRe We There Yet; Nylander *et al.*, 2008; Lanfear *et al.*, 2016; Warren *et al.*, 2017) calculates an approximate ESS of the tree topology, which can provide additional information on the convergence of the tree. Additional graphical outputs can be generated in RWTY, such as treespace plots, which allow the visualization of how an MCMC chain explored parameters during its run.

4.2.3 Convenience

Convenience (Guimarães Fabreti & Höhna, 2022) is an R package that takes a fundamentally different approach to both how to calculate and how to assess ESS than RWTY and Tracer. It can produce visual outputs for convergence assessments, but also can produce simple text outputs stating if a run has converged or not.

ESS is still calculated in convenience. But rather than using an arbitrary threshold, such as an ESS of 200, convenience calculates a minimum threshold for a good ESS based on the standard error of the mean (SEM). The SEM allows a researcher to know how much error there is in the estimate of the posterior mean, compared to the variance of the posterior distribution. For this calculation, the posterior distribution is assumed to be shaped like a normal distribution, so the width of the 95% probability interval of the distribution is approximately equivalent to 4 δ, with δ being the standard deviation. This quantity is the reference used to calculate the threshold. By default, the ESS threshold in convenience is set to 625, which corresponds to an SEM equal to 1% of the interval width. By contrast, the threshold of 200 set by Tracer corresponds to an SEM of 1.77% of the interval width. Although higher ESS values are always better from a convergence point of view, they can also come at considerable computational cost, particularly for more complex analyses. Thus the choice of threshold should be adapted to each situation, for instance by using larger thresholds for critical parts of the inference and lower thresholds for less important estimates.

Convenience also allows the tree convergence to be estimated, by calculating the ESS of splits in the tree. A split represents a particular subclade of the tree, which can be either present or absent in each posterior sample. By calculating the ESS of all splits, we can thus obtain an estimate of the ESS of the tree topology. Finally, the reproducibility of an MCMC run is also considered by convenience. Two MCMC runs of the same analysis can be compared against each other using the Kolmogorov–Smirnov (KS) statistical test (Kolmogorov, 1933; Smirnov, 1939), which tests if two samples were drawn from the same underlying distribution. If your two MCMC chains do not seem to be drawn from the same distribution, then this means your MCMC simulations are not consistently finding the same stationary distribution, meaning that one or both chains has not converged yet. One possible reason for this is the presence of multiple alternative possible solutions, with each chain finding a separate local optimum. Different slices of the same MCMC chain can also be compared against one another using the KS test to assess if the chain is in the process of converging.

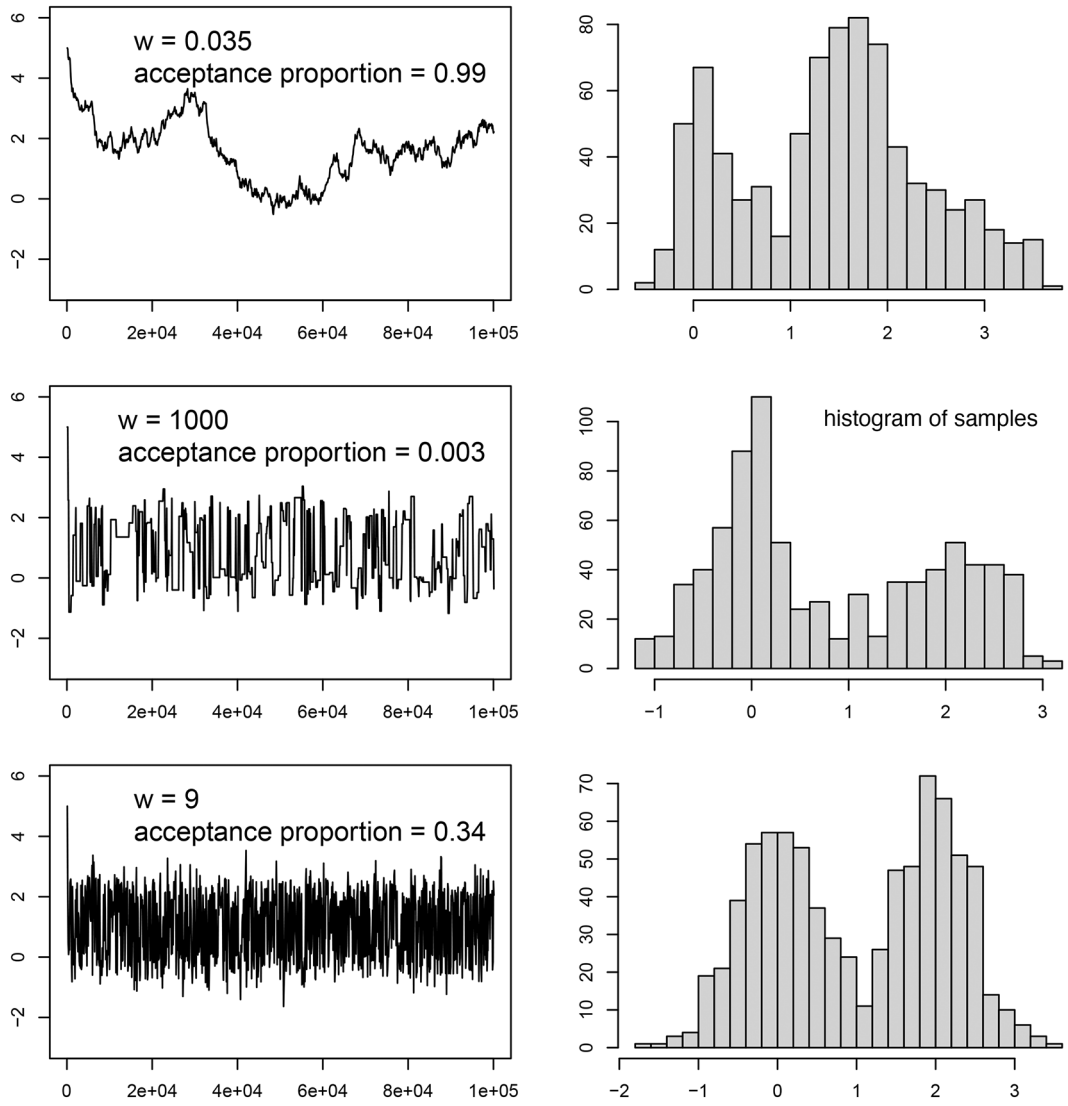

Figure 4. The target distribution is a half-half mixture of two normal distributions, one with mean 0 and standard deviation 0.5, the other with mean 2 and standard deviation 0.5. This distribution is estimated using MCMC with the sliding move (see Figure 3). The starting value is 5.0. The window size (w) is a tuning parameter of the move. For each w value, the left panel shows the trace of the MCMC samples, while the right panel shows the histogram of the MCMC samples (discarding the first 20% samples as burn-in). When *w* is too small, the MCMC performs localized transitions and exhibits slow mixing, as shown in the top panel. On the other hand, too bold moves with large w frequently get rejected and the MCMC is trapped in a state for too long. This can be seen in the middle image, with periods of stasis in the chain. In both cases, the algorithm is inefficient and the posterior distribution is not well estimated from the sample. In the last row, the inference appears to have converged.

5. Common issues and proposed resolutions

As we have seen, MCMC analyses are composed of many different parts, which can make it difficult to identify the cause of problems. In this section, we detail some common issues which can affect the convergence of an MCMC inference, or even prevent it entirely from starting. An abbreviated overview of all the issues and resolutions described below can be found in Figure 5. Note that we focus in this article on issues and solutions from the perspective of a user of phylogenetic inference software. However, many issues (such as setups with strong correlations, as mentioned earlier) require further development work, and thus cannot be implemented by the users alone.

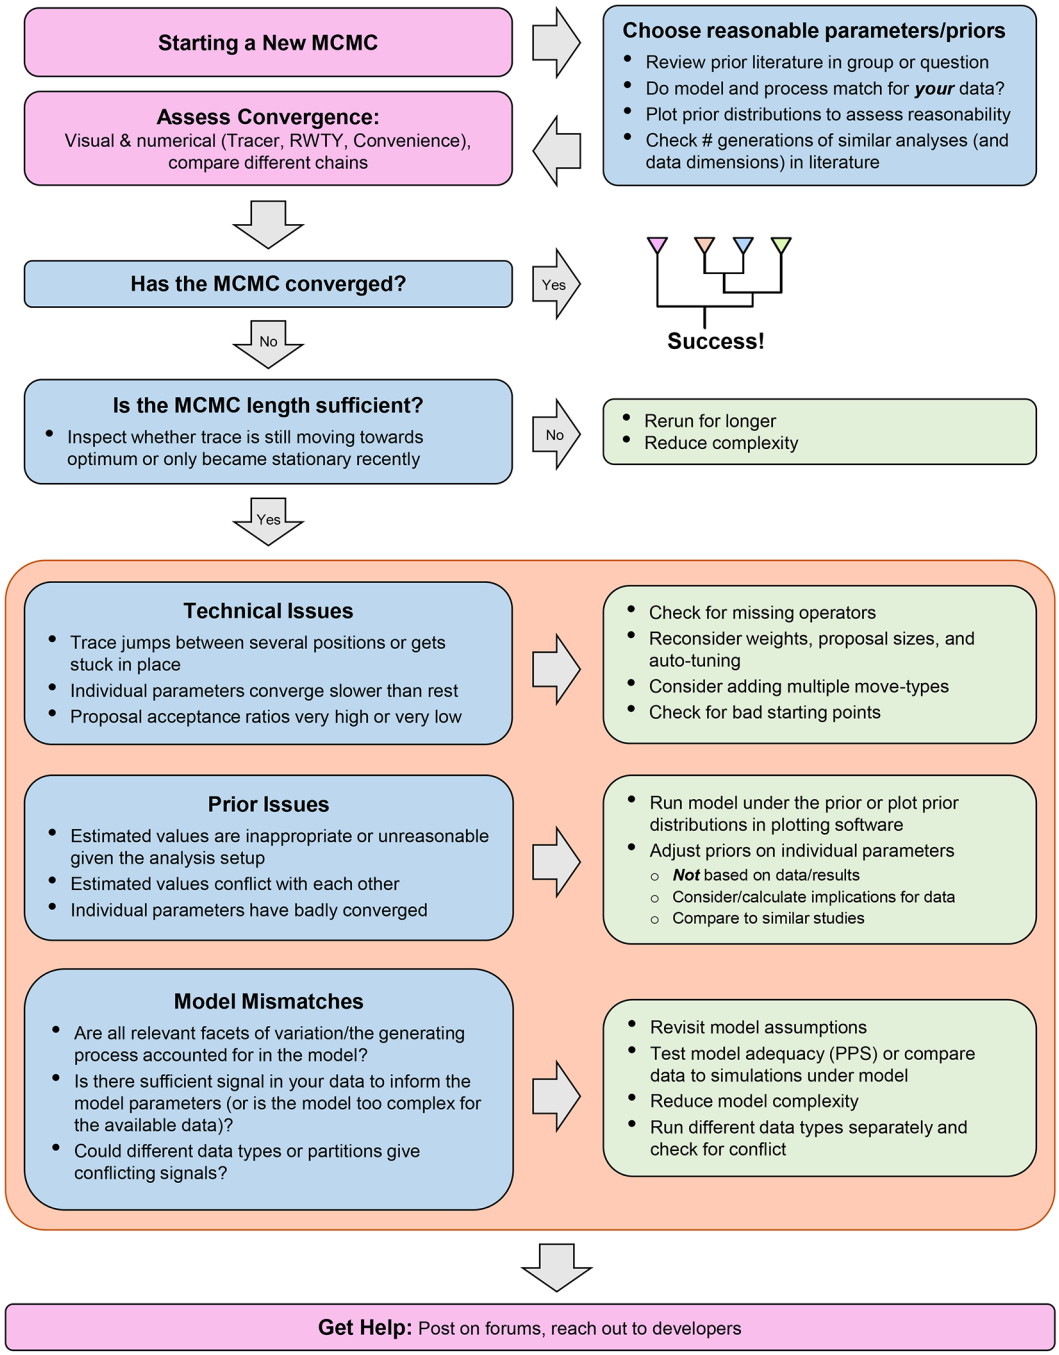

Figure 5. A flowchart to guide users through the MCMC-debugging process, highlighting key points mentioned in the text, with common issues in blue boxes and corresponding resolutions in green. Note that the different types of issues and resolutions within the orange box are not meant to be addressed in the order shown, but represent different avenues for investigating an issue.

5.1 Inference technical setup

5.1.1 Moves/operators/proposals

If an analysis does not converge well, or takes unreasonably long, it is worth checking the operators. Each parameter that is supposed to be estimated by the analysis needs to have at least one operator associated to it, in order to be sampled. If an operator is missing, that parameter will never change from its initial value, which not only means it will not converge, but

also that other parameters can be prevented from converging properly. This issue tends to be more frequent in RevBayes, where there is no default operator setup proposed by the software and the choice of operators is left entirely to the user.

Another possibility is that the weights of the individual operators may need to be reconsidered (i.e., how often a new value should be proposed for the corresponding parameter). In some cases, some parameters are mixing well, and only a few specific ones are causing problems. In this case, it can help to increase the weight of the operators corresponding to badly-estimated parameters, so that more moves are being proposed for them throughout the analysis. Similarly, decreasing the weights of operators corresponding to well-estimated parameters will decrease the amount of computational time spent on proposals for these parameters, without affecting convergence too much.

Alternatively, if changing the weight did not fix the chain's behaviour, we should consider its proposal size (i.e., how far from the current parameter value a proposed new value is). Many proposals, especially proposals on numerical parameters, include a configuration value which affects this size. A proposal size that is too small will make convergence of the corresponding parameter very slow, even at high operator weights (Figure 4, top row), and may even trap the chain on a local optimum. If proposal sizes are too large instead, the chain may 'overshoot' the parameter values at the optimum (i.e., the peak of the posterior) or roam too far from them to converge properly. Overshooting optima for a parameter may also mean that many proposals land on values with worse likelihood scores, which means they get rejected frequently and the chain experiences periods of stasis and thus overall poor mixing (Figure 4, center row). The sampling pattern for that parameter may then also be too 'coarse' to properly capture the peaks and valleys in the likelihood.

A way to catch issues related to proposal-size is to check the final acceptance rates for all operators, as well as the final trace. BEAST2 will even offer suggestions for adjusting the proposal sizes based on acceptance rates at the end of the chain. Other than that, appropriate proposal sizes are not always straightforward for users to determine, but the problem can be alleviated in two ways: by letting the inference software auto-tune them, or by using a combination of proposals operating on the same parameter, but with varying proposal sizes. The latter is particularly helpful if the likelihood surface is very heterogeneous, as the chain then has a variety of step-sizes available, potentially increasing the likelihood that the appropriate one can be proposed. However, auto-tuning should be turned off if this strategy is chosen, so the separate proposal sizes will not change throughout the chain.

If auto-tuning is turned on (which is the default in BEAST2 and MrBayes, and optional in RevBayes), proposal sizes of moves will periodically be adjusted to guarantee a good acceptance ratio (e.g., to match a goal of 0.4). For example, if too many of the recently proposed moves were accepted, this might indicate that their size was too small (i.e., they may be slowly trudging uphill towards an optimum, see Figure 4, top row) and it will thus be increased. If too few proposals were accepted, this might in turn mean their size was too large (i.e., they shoot past values at an optimum into parts of parameter space with lower likelihoods, see Figure 4, center row) and will be decreased.

We generally recommend making use of such tuning features, but urge users not to mistake them for a magic solution to all proposal-size problems. Instead, one should be cautious not to 'mis-tune' or 'over-tune' the analysis. The main considerations when setting up auto-tuning are how often, and for how long to tune. Depending on the implementation, users can specify during what portion of the chain the parameters are tuned (i.e., during a dedicated tuning interval or burn-in phase, or throughout the run), and how often the parameters are being tuned during that interval. Tuning orients itself on the behaviour of the proposals during the chosen tuning intervals. Thus, these intervals need to be representative for the rest of the chain going forward, if the tuned values should be useful. In particular during the early stages of the MCMC (i.e., the burn-in phase), larger proposal sizes may be favoured as the chain moves from parameter values with low likelihood towards the optima, whereas smaller sizes might be favoured when exploring the likelihood surface around the optima. This generally means that proposals should possibly be re-tuned multiple times, to allow for feedback from the new behaviour of the tuned operator, and suggests that longer phases of tuning are needed for chains initialised at naïve starting values, than for those tailored to possibly start closer to the optima. However, if tuning intervals are kept too short, the available information might not be representative for the operator's behaviour, resulting in unnecessary or inappropriate proposal size changes. Furthermore, while continuous tuning throughout the analysis can help account for the different requirements far from the optima versus close, there is a danger to tune towards the current location of the chain, homing in on smaller and smaller proposal sizes and thereby 'trapping' the chain on a local optimum. We would thus prefer to mainly tune during burn-in, and not during the main part of the analysis unless there is evidence that it is necessary. However, using the aforementioned strategy of multiple operators with varying, un-tuned proposal sizes might be a more helpful approach in such a case. Note that these changes can be integrated when running a new chain or when resuming the current one, as proposal configurations do not change the posterior distributions.

It can be difficult to identify which parameters exactly are causing the problem, since they can affect the mixing of others, blurring the picture. In particular, if the tree estimation has not converged, this can affect many other parameters. Often it

is possible to identify the culprits by revisiting how the parameters are causally connected in the model. If available, a look at a schematic representation of the model might help getting more clarity on how different parameters may affect each other's mixing. In BEAST2 or RevBayes, this representation can be obtained directly from the software (through BEAUti in the case of BEAST2, or by printing the model's DAG [directed acyclic graph] in the case of RevBayes). MrBayes has the 'showmodel' command to show the current model, parameters in the model, and priors for the parameters. It also has the 'showmove' command to show the MCMC moves for the parameters.

5.1.2 Starting values

Another problem is the initialisation of the MCMC chain at a 'bad' position. This means that our analysis started at a combination of parameter values that is either very far from those with the highest posterior, or at a combination of values that is implausible or hard to compute given our data. As a result, the analysis may take much longer to converge (since it has to first slowly make its way out of the poorly fitting area of parameter space), or may crash altogether (e.g., because no likelihood could be calculated for conflicting parameter values). Ideally, users will have thought well about the possible values of all parameters and have set the respective prior distributions to favour the most plausible parameter values. However, the initial values are often left as the default (in BEAST2) or are picked at random from the prior (in RevBayes), so the chain can start in an unfavourable part of parameter space, or at an implausible combination of values. For example, we could start with some proposed very short branches along with a very low mutation rate, which could never explain the observed differences between the sequences of taxa. Or the starting values for the speciation and extinction rates could be implausibly high compared to the root age of the tree and its number of taxa.

To combat this, we usually have the option to specify the starting values for each parameter to something we deem reasonable. It may not always be straightforward to know what those values should be for a particular parameter, but beyond trial and error, a few standard options have been established. One possibility is to start at the expected mean of a prior distribution, which would be expected to work well as long as the prior distribution itself is sensible. Reminding oneself of the parameters' biological meaning can also help to come up with a good solution. For example, speciation and extinction rates eventually just determine how many species we expect to arise and die out again over a given time period. Thus, a commonly used starting value for speciation rate is $\lambda = ln\,(nTips/2)/rootAge$, which gives a simple estimate of net diversification (sometimes called the Kendall-Moran estimate; Baldwin and Sanderson, 1998). The starting value for extinction is then sometimes set to $\mu = \lambda/10$ (which is less founded in biology and rather ensures that we start with positive net diversification). Starting values can also be set for non-numerical parameters. Starting trees can be provided which may already be closer to those with a high posterior probability (e.g., a quick maximum likelihood tree or a previously-published estimate) than a randomly drawn tree sample. However, attention has to be paid to the tree not being in conflict with other priors or constraints. For instance, the starting tree needs to be compatible with additional time information such as node calibrations, and with added constraints such as subclades (i.e., monophyletic sets of tips).

It is important to remember that starting values do not have to be spot-on estimates of where the actual values with high posterior probability lie, because after all, the MCMC is expected to sample a variety of solutions. The goal is merely to ensure that we have set a feasible combination of values for the chain to start from. Doing so does not only prevent computational issues (in case of unfeasible parameter combinations), but can also speed up the analysis (because we do not force the chain to trudge through parameter space that is far from the optima anyways, and instead allow it to start exploring feasible solutions). Also, it may prevent issues with the auto-tuning performed by the software. Since auto-tuning usually happens at the beginning of the inference, the behavior of the moves may end up being tuned to suit a different part of parameter space than where the chain eventually should spend most of its time exploring, as described above.

5.1.3 Inference algorithm

A final problem which can happen with MCMC is that the inference gets stuck in local optima for extended periods, which means not only that reaching convergence will take a very long time, but also that there is a risk of assuming the chain has converged when it has not. This issue is more frequent when the likelihood surface of the chosen model is more complex, and particularly for models where a lot of components are correlated (e.g., a multi-species coalescent model containing many different phylogenies which need to be consistent with each other). One sign that this is occurring is the presence of 'plateaus' in the trace, where the parameter value will oscillate in the same place for a lot of consecutive iterations, before jumping suddenly to a very different place in the parameter space. Another sign could be that independent runs generate inconsistent MCMC samples, indicating each run could potentially be sampling a different local optimum.

One potential way to resolve this issue is to use an expansion of the MCMC algorithm called Metropolis-coupled MCMC (MCMCMC or MC^3, Altekar *et al*., 2004). This algorithm relies on running multiple chains, raising the i-th chain to a power of $1/[1 + T(i - 1)]$, in which T is called the 'temperature'. The first chain (called the cold chain) has a power of 1, meaning that it behaves like a regular MCMC chain; this is the chain that we will use to draw samples from the posterior. The other chains have a higher heat (with power < 1), thereby flattening the likelihood surface (Geyer, 1992; Gilks & Roberts, 1996). This means that heated chains can move more quickly and easily through the parameter space, and can reach points of high likelihood that the regular chain has not found yet (Whidden and Matsen, 2015). However, as the heated chains do not evaluate the true posterior probabilities, we discard samples from such chains. Instead, the MCMCMC algorithm will use the heated chains to help explore the parameter space, and will periodically switch states between the cold chain and the heated chains (in an additional MCMC move), allowing the regular chain to 'jump' to new places in parameter space, as long as those new places have a higher posterior than the current ones.

MrBayes (Altekar *et al*., 2004), BEAST2 (Müller & Bouckaert, 2020) and RevBayes currently support the MCMCMC algorithm, and it can be used along with any model or prior choice. MrBayes uses it by default (performing two independent runs and one cold chain and three heated chains per run). While MCMCMC potentially improves the performance of the inference, chain switching can mislead convergence assessment using regular ESS values (Guimarães Fabreti & Höhna, 2022), and poorly chosen temperatures can lead to inefficient mixing and other issues, and should be looked out for (Brown & Thomson, 2018). It is also worth noting that MCMCMC is computationally expensive. For instance, the computing time is eight times longer when running the default 2×4 chains in MrBayes, compared to running a single chain. To speed things up, it is highly recommended to use the parallel version of the program. For example, MrBayes can run the eight chains in parallel, thus the computing time is only slightly longer than running a single chain, with the extra time needed to exchange information among the chains.

5.2 Choice of model and priors

Even with all the technical aspects of the analysis set correctly, we can get convergence problems and faulty behavior of the parameters. Such issues can either stem from unexpected interactions of priors, clashing components of the model, or mismatches of the model with the data. It can at first be challenging to distinguish those. If we do not already have a suspicion as to what the culprit might be (e.g., based on the trace, peculiarities of the data or model), one way to tell whether the issue lies with the analysis setup per se or with the pairing of data and model, is to run the MCMC 'under the prior', producing what is formally called a 'prior predictive distribution'. This can either be done by switching on the option in the program or by using an empty alignment or character matrix (e.g., a column of '?'), so that only values from the prior will be sampled and none of the data is involved. Thus, any remaining issues will be due to problems in the analysis setup, such as conflicting or interacting priors – and vice versa, if there are no such remaining issues, the problem may lie with the data or the model. Running the MCMC inference under the prior is useful not only for troubleshooting potential setup issues, but also for interpreting the results of the actual analysis. The difference between the prior distribution and the full posterior gives an estimate of how much of the signal present in the posterior sample actually comes from the sequence or character data, as opposed to the prior distributions. Note that although fossil ages are technically data, the probability of the tree under the FBD process given the fossil ages is often treated as part of the prior. This is useful for examining the impact of fossil ages or calibrations, in absence of the sequence data, but can also impact model selection and marginal likelihood estimators, as detailed in May and Rothfels (2023).

5.2.1 Priors

The choice of good priors can make a big difference for the success of the MCMC. Of course, coming up with good priors is not trivial, and generally applicable advice is not always available. One difficulty is that priors should be clearly separated from the data. In a Bayesian inference, the probability of the data is accounted for by the likelihood. So, if the priors are also informed by the same data, then the information provided by the data ends up being counted twice by the inference, which will artificially increase its contribution to the posterior. Priors can thus be based on previous studies or biological knowledge, but not on analyses using the current dataset under study. We note that we adopt here a subjective Bayesian perspective, in which a researcher is interpreting the probability of different hypotheses, and updating their belief to the posterior probability using data. Other approaches, such as empirical Bayes in which the prior may be estimated from the data, are also possible. It should be noted that most Bayesian practitioners will use some tools from both schools of thought.

So how do we set priors? It may be tempting to just follow tutorials or use default priors at first, however, we strongly encourage users to think more critically about the implications of the prior choice for each individual analysis. While it is true that developers often design default settings to be a reasonable starting point for most analyses, they are by no means meant to be a one-fits-all solution, and one should not expect them to necessarily be an optimal or even good fit for one's own problem. As an extreme example, the default prior on some parameters in BEAST2 is a so-called improper prior,

such as $1/X$ or Uniform(0, $+\infty$), as the distribution does not integrate to one but to infinity. For example, the prior on the mean clock rate in BEAST2 is set to Uniform(0, $+\infty$), because what constitutes a reasonable value for the clock rate is extremely dependent on both the organism and the timescale of the dataset. Thus it is up to the user to select a reasonable prior distribution for this parameter. Overall, improper priors are sometimes used for the goal of having uninformative priors. However, improper priors are statistically problematic. Although a posterior distribution that includes improper priors can itself be proper, improper priors can lead to issues if the posterior is strongly driven by the prior (and thus is itself improper) or when performing marginalization, for instance for model selection. As a result, we believe that they should be avoided if possible. In general, models are frequently used for applications or in scenarios beyond what the original developers anticipated, so your data or question may be quite different from what the method developers had in mind when choosing the default options. Developers also generally set their defaults as wider than would generally be reasonable, since it is difficult to predict the appropriate settings for specific analyses.

Thinking more carefully about the priors and their implications will go hand in hand with a deeper understanding of the model itself, which should be an additional encouragement to dive into it. The key is to remember that the prior distribution of a parameter represents the plausibility of different values, and that values outside of the prior distribution (i.e., those that are considered implausible) can never be explored by the MCMC inference. In particular, this means that long-tailed prior shapes, such as lognormal or exponential distributions, are often better than uniform distributions, which restrict the range of values which can be tried by the inference. Note also that priors always influence the results of the inference, and that setting very vague priors is not an optimal choice in most circumstances. For instance, in the example of the clock rate prior presented earlier, a prior distribution of Uniform(0, $+\infty$) puts a lot of weight on very high values for that parameter, and will thus encourage the inference to try these values. If the data is not very informative on this particular parameter, this can result in estimated values which are absurdly high from a biological point of view or can lead to convergence issues. A better prior would use our understanding of evolutionary processes to put more weight on biologically plausible values.

When choosing a prior, we thus need to consider what particular parameter values would imply for the data. For instance, substitution rates describe how fast mutations happen in the sequences and become fixed, and thus how much the sequences of the species under consideration could diverge over time.

Overall, in order to identify reasonable priors, we can ask the following questions:

- Have the parameters used in our analysis been estimated in other contexts or for similar datasets?
- What priors have similar studies chosen and how comparable is their data to ours? Note that these priors still need to be critically evaluated, as our understanding of plausible parameter values may have changed since the previous study.
- Does the range of parameter values allowed by the prior make sense given our data and analysis setup, for instance is it consistent with the expected number of substitutions in the alignment or the minimal clade ages?
- Can we do rough calculations to calibrate our prior expectations by, e.g., dividing the number of extant species by the assumed clade age, to get a rough estimate of net diversification?
- Can we obtain estimates for the parameters from sources outside of our dataset, for example using the fossil record to get an idea of how much extinction our focal clade may have experienced? Note that this requires making sure that the parameters chosen actually represent the same quantities between models, which is not always the case. For instance, extinction rates obtained from the fossil record represent a different parameter than death rates used in the fossilized birth-death process (Silvestro *et al.*, 2018; Stadler *et al.*, 2018).

Although this may sound like a lot of work, it is also important to remember that identifying reasonable values for the different parameters, finding previous estimates for comparison, and evaluating the biological implications of the different values will always be needed to interpret the results of the analysis. The main difference in a Bayesian inference compared to other types of inference is that this work has to be performed upfront, rather than after the inference is finished.

For some analyses (especially hierarchical models, see for instance Heath (2012)), users may encounter the use of what is called a *hyperprior*. Put simply, when a prior distribution is specified for a model parameter, the parameters that determine the shape of that prior's distribution can either be set, or instead they can also be estimated, in which case a prior is specified for them too. To distinguish this from a prior on a model parameter, it is then referred to as a hyperprior – a prior

on another prior. For example, when time-calibrating a phylogeny we can specify a gamma distributed prior for our clock rate, but we may have no good sense of what shape that gamma distribution should be. Thus, we can instead specify an exponential prior on each of the shape parameters of the gamma distribution – those exponential hyperpriors are now informing the inference of the gamma prior. Besides making a prior more flexible and general, the benefit of defining hyperpriors this way is that they allow us to represent our uncertainty on what a parameter's prior distribution should be in the model, and estimate it instead of being forced to make a specific assumption. Of course, this assumes that there is enough information in the data to inform that estimate, otherwise it just allows for more uncertainty while adding additional parameters to the model. Consequently, the use of hyperpriors should be applied with caution.

It is generally advisable to plot the specified prior distributions and think about what they imply. Overall, the actual shape of the distribution (lognormal, gamma, etc.) is usually less important than the range of plausible values covered by the distribution (the 5% and 95% quantiles). However, the shape of the distribution affects the weight given to different parts of the range, i.e., whether low values are more likely under the prior than high values. Comparing the distributions for different values by using the visualization tool in BEAUti or plotting them in R is a great way to get a better idea of what is happening. It should be noted that simply looking at a curve may be misleading. Because the area under a certain section of the curve (e.g., a long tail) may still be large, even if the height of that section of the curve looks small. Thus, quantifying how much area is covered by the distribution (such as through quantiles) is still important. In the case of node calibrations, even if each calibration is reasonable by itself, their combination or interaction with the tree prior can restrict the parameter space in unexpected ways (e.g., Heled and Drummond, 2012; dos Reis *et al*., 2012; Warnock *et al*., 2015). Running the analysis under the prior can help spot situations in which the analysis is not specifying parameter distributions that the researcher considers reasonable. Further, the impact of the 'effective priors' can have a tangible impact on the posterior (for further discussion and some examples see Inoue *et al*. (2010); Warnock *et al*. (2012); dos Reis *et al*. (2015); Brown and Smith (2018); Budd and Mann (2024)). For a good illustration of the distinction between user-specified and effective priors, and the impact of applying priors with varying degrees of uncertainty on the posterior see dos Reis *et al*. (2015, Figure 1).

Finally, it can happen that information from previous studies leads to very different or even conflicting results for plausible parameter values in a specific dataset. One possible solution is to set the priors to be very uncertain and let the analysis decide which hypothesis is supported by the data, but this may not always be possible, and it risks creating convergence issues due to using vague priors. A good alternative is to run the same analysis using multiple different prior configurations, which will show how much the results and conclusions are impacted by the prior choices. If different priors yield clearly different results, further investigation is needed to determine what conclusions, if any, can be drawn from the inference. Possible next steps could include running the analysis without data, but under both sets of priors to test whether either setup is strongly biasing the result (and is thus in conflict with the data), or evaluating the biological information that informed the priors to begin with, and whether either of them is backed up more convincingly. If one (or several) parameters show prior sensitivity for both sets, this may indicate that there is not enough signal in the data to properly infer it, and one should explore the possibilities to add data or to simplify the model to avoid estimating this parameter, if feasible. If removing a parameter is not an option, model testing and adequacy testing could be employed to explore the most suitable prior configuration for the data, or model averaging could be used to integrate over the different priors to see what conclusions are still possible despite the lack of signal to inform this parameter. Even when there is no conflict in the literature, evaluating the robustness of the results when using different prior setups provides important information to inform which conclusions can be drawn from the analysis. This is especially important if running the analysis under the prior shows that the prior has a large influence on the results (i.e., the posterior distribution obtained with the data is very similar to the posterior distribution obtained with the prior alone).

5.2.2 Model

When the analysis is set up correctly and priors are reasonable, the cause for convergence problems may lie with the model itself, or how it relates to the data. It may seem daunting to choose between all the different types of models out there. There are a few pieces of software that can help researchers get an idea of plausible models. ClockstaR (Duchêne *et al*., 2014) can be used to choose appropriate relaxed clocks for molecular data. EvoPhylo (Simões *et al*., 2023) can do a similar selection for morphological data partitions. Although model selection can not be used to select between alternative birth-death sampling models because fossil ages are technically considered part of the prior (May & Rothfels, 2023), integrated tools in the Paleobiology Database website (Uhen *et al*., 2023) can also assist in finding reasonable starting parameters for FBD analyses. These tools use established paleontological methods for estimating parameters for speciation, extinction and fossilization rates. Although the rate parameters inferred from fossil versus phylogenies can reflect different processes (Silvestro *et al*., 2018), fossil-based estimates can still provide reasonable starting values, as well as insights into variation in sampling (see section **Data quality issues**). Using these sorts of tools can help with setting priors that have some support from the established literature.

If different data sources are being used for joint analyses, one might want to try running the different data separately in order to confirm whether they might support incompatible solutions. For example, in a total-evidence analysis, molecular and morphological data may each support different tree topologies. So when analyzed jointly, solutions which could increase the likelihood of one type of data will decrease it for the other type, and vice versa, thereby making convergence around one optimum difficult. While Bayesian inference can in theory converge on a multimodal posterior distribution, in practice this may require specific moves, especially if the modes are very distant from each other. The same could apply to other combinations of data sources, e.g., conflicting molecular markers. Running the data for each type/partition separately can help a researcher to determine if the convergence is poor due to methodological issues, or true signal conflict.

Much more fundamentally, the analysis might also just struggle to run or converge because the chosen model is not suitable for the data at all. Carefully revisiting the model's assumptions and how those should manifest in the data is required to judge this, e.g., are there patterns of variation in our data, for instance between different groups, which the model needs to be able to address? An approach specifically designed to judge such model-data mismatches is model adequacy testing. This is done by simulating new data sets from the inferred posterior distributions, an approach termed posterior predictive simulations (PPS). These simulated data sets are then compared to the initial data using summary statistics which capture its relevant characteristics. If the model is inadequate to describe/analyse the variation in the data, that should be revealed through significant differences in the summary statistics between the data and the posterior simulations. These types of tests exist for a variety of phylogenetic models, including substitution models (Bollback, 2002; Brown & ElDabaje, 2009; Lewis *et al.*, 2014; Nielsen, 2002), tree inference using molecular data (Brown, 2014; Duchene *et al.*, 2019; Reid *et al.*, 2014) or morphological data (Mulvey *et al.*, 2024), continuous and discrete trait evolution (Blackmon & Demuth, 2014; Huelsenbeck *et al.*, 2003; Pennell *et al.*, 2015; Slater & Pennell, 2014), and diversification models (Schwery & O'Meara, 2020; Schwery *et al.*, 2023). Approachable examples of such adequacy tests can be found on the RevBayes website (https://revbayes.github.io/tutorials/intro_posterior_prediction/). However, that approach technically would require posterior estimates from a more or less successful MCMC, which would not be available if the analysis keeps crashing, and would likely be uninformative if the MCMC did not converge. To circumvent this one could take an approach akin to approximate Bayesian computation (ABC, Sunnåker *et al.*, 2013): simulating datasets from scratch under the same model and a range of parameters and then comparing them to the original dataset, using the same summary statistics as one would use in the PPS approach. As a way to inform what parameter values should be included (and narrow that range), one could start by drawing from the prior, essentially repurposing prior predictive simulations. Whichever approach is chosen, exploring how the empirical data differs from what is expected under the model allows us to judge the nature of the model-data mismatch.

Finally, it might be worth trying to reduce the complexity of one's model. While it is tempting to make full use of the levels of complexity modern approaches allow us to model, one ought to consider whether there is enough information in the data for the model to work with. Just like any statistical test has sample size requirements to have the power to detect significant differences, these models need the data to have sufficient size and structure/heterogeneity to be able to infer parameter estimates without too much uncertainty. For example, we may want to use a relaxed clock model to account for the possibility that different parts of our tree evolve at different rates. But if we only have one fossil to calibrate our node ages with, or the sequences are not substantially variable, the model has limited information on which to base any rate differences on the tree. As a result, the different branch rates suggested by the MCMC will possibly meander around the parameter space without any receiving overwhelming support. Using a strict clock instead might neglect possible rate heterogeneity, but will at least be able to converge on reasonable estimates given the limited information available.

5.3 Data quality issues

In general, assembling more data leads to more precise and more accurate inferences. For example, previous research has shown that total-evidence studies require ∼300 morphological characters to obtain reliable estimates of tree topology and divergence times in extinct clades of 200 tips (Barido-Sottani *et al.*, 2020). In general, the number of morphological characters needed is likely to depend on the size of the phylogeny and the uncertainty present in the dataset, but at least 100 characters are needed for accurate estimates on realistic tree sizes, even in ideal conditions (Capobianco, 2025). Purely from a performance perspective however, it is important to note that additional data is not necessarily better for the convergence of an MCMC inference. Indeed, adding more data comes with added computational costs, and thus can have a net-negative impact on the performance, especially if the added data is very uncertain or conflicts with the rest of the data or with the chosen models and priors. For instance, Portik *et al.* (2023) built phylogenies using either a complete alignment of nuclear markers, a supersparse matrix of ∼300 genes with large amounts of missing data, or the combination of both. They found that trees obtained using the combined dataset did not significantly differ from the trees obtained using the complete alignment alone. One possible avenue for resolving convergence issues is thus to remove genes or partitions which contain low amounts of information.

Similarly, increasing the number of extant (or fossil) samples in the tree leads to an exponential increase in the number of possible topologies, and so represents an important drag on performance. We typically select a subsample of extant taxa to be included in our analyses. We may assume extant taxa are sampled uniformly at random; but in many cases, they are sampled sparsely by keeping only one living representative per genus or subclade. The diversified sampling scheme has been implemented in the FBD model (Zhang *et al*., 2016) to accommodate such cases, and is available in MrBayes and the BDSKY package for BEAST2 (Zhang *et al*., 2023).

As mentioned above, there are two options for incorporating fossils directly in the phylogeny using the FBD process: assigning fossils to nodes via constraints or using morphological data in a total-evidence framework. Both approaches to positioning fossils present a challenge for MCMC inference, since even with character data, the topological uncertainty associated with fossils is typically large. And when there is a large amount of phylogenetic uncertainty, the posterior can span a broader flatter area, taking more effort to sample and making it harder to reach convergence. The use of very broad constraints (e.g., allowing the fossil to attach anywhere in the tree) in particular can lead to convergence issues, since there is insufficient information to inform the topology or other model parameters. To improve convergence, researchers could use the most precise constraints available, i.e., less inclusive nodes or lower taxonomic divisions, such as genera. In addition, it is possible to set a backbone extant tree, which will fix or strongly restrict the position of extant samples in the phylogeny, leaving only the positions of the fossils and the branch lengths to be estimated. That said, we emphasise that constraints should be implemented with extreme caution, as errors in constraints can lead to inaccurate results (Barido-Sottani *et al.*, 2023a). Having character data for fossils can help improve convergence, as it provides direct information about the topology. If convergence issues persist, provided additional taxonomic information is available, both approaches to fossil placement (the use of character data + constraints) could be combined. If additional taxonomic information or morphological data is unavailable, researchers might need to reconsider the scope of their analyses and the application of the FBD process to the data.

In addition to extant taxa, fossils are also often sampled non-uniformly, with abundant fossils in some strata but rare in others. Extensions of the FBD model can take this into account (see Mulvey *et al*., 2025, for a recent review). For example, rates can vary across time bins (Gavryushkina *et al*., 2014; Zhang *et al*., 2016), or sampling can be modeled as an instantaneous event, akin to fossil Läggerstatten (Magee & Höhna, 2021). To increase the biological realism of the FBD process, researchers might be tempted to incorporate variation in the sampling or diversification processes. This leads to an increase in the number of free parameters and another trade-off between model complexity and data availability that must be considered. Increasing the number of fossils will improve parameter estimation, leading to more accurate and precise estimates of the FBD model parameters, as well as divergence times and topology (provided the model is not strongly violated). However, users should bear in mind that adding fossils also increases overall tree size – each fossil is a tip or potential sampled ancestor, whose position must be sampled using MCMC.

This means that although fossils do not typically burden the computation through the addition of character data, which would increase the cost of calculating the likelihood, they increase the tree space, which will take longer to sample using MCMC. For many broader clades (e.g., mammals, animals, plants), including all fossil occurrences, while desirable, is not feasible. Presently, the maximum tree size that could reasonably be inferred using the FBD process is around 500 samples. One approach to get around this for large clades, or datasets with large numbers of fossil occurrences, is to randomly subsample the fossil data (O'Reilly & Donoghue, 2020). This allows us to obtain a more manageable dataset without violating the sampling process assumptions.

Finally, if age uncertainty is substantial for many or all fossils in your analysis, the MCMC might also take longer to converge. However, compared to analyses that used fixed fossil ages, Barido-Sottani *et al.* (2019) showed using simulations, that incorporating fossil age uncertainty does not make the MCMC inference less efficient, i.e., more iterations are not always required to reach convergence, at least for data sets typical of Cenozoic mammals. This is probably because the use of fixed fossil ages introduces conflict into the tree space, leading to less efficient mixing.

5.4 How long should I run my analysis before giving up?

Some analyses take a long time to converge because it is hard to find high-posterior values in a large parameter space, or because several local optima exist, or because a lot of uncertainty exists around the high-posterior values. Visually inspecting the parameter traces can give indications for this. Are they stabilising around certain values, still showing a trend into a certain direction, or just wildly meandering around? Trends in the trace can indicate that the MCMC is still moving towards high-posterior regions and just requires more time to find them (or perhaps needs to be restarted with starting values further in the suggested direction). But continuously rising or declining parameter values can also be pathological behavior, suggesting misspecified priors or overly strict constraints in related parameters. Over-tuning of moves can also lead to such erratic behavior, e.g., if the step-size for some parameters was tuned to be overly short or long.

Wildly meandering traces could again be an indication of the data not containing enough information for those parameters to be identifiable, or step-sizes to be too long to allow it to settle around a peak in posterior probability.

As mentioned above, a behavior researchers sometimes observe is that an MCMC will appear to stabilize on a set of values, then jump to a completely different likelihood. This can be either an improvement (finding better values) or worse. This can happen because the analysis was previously stuck in a local optimum. That is, a region of parameter space that was good, but not the best in treespace. Thus, exploring this new optimum further is warranted. Or it may be that making a change to one parameter, such as the tree, causes a jump to a worse parameter space for other model parameters. In either event, this raises the question of how we can be sure that convergence has been achieved, and that the chain is not going to suddenly change further. Running multiple MCMC chains is an advisable way to discern between these two scenarios. Many software packages default to using two independent runs. Some, such as RevBayes and MrBayes, allow more to be used. Two to four runs are common in published analyses. Depending on the software used, each run either contains a single MCMC chain (such as BEAST2) or multiple chains (e.g., MrBayes uses MCMCMC by default which executes one cold chain and three heated chains per run).

Overall, the number of steps required to achieve convergence is difficult to estimate, as it will depend on all the components of the analysis, including the specific software used. Searching the literature for similar analyses, both in terms of dataset size and of models used, can provide a reasonable order of magnitude for the number of steps needed. The original publications of the specific model or package used, if available, will also provide estimates for what the original authors believed was a reasonable dataset size. Importantly, inference software all integrate a checkpointing mechanism, so analyses which have not converged can be resumed without losing the work already done. Thus it is not an issue if the initial number of steps is too low. Running several different chains with the same analysis can also be helpful in assessing how far the chain is from convergence. If the posterior distributions obtained by the different chains are largely mismatched, then convergence is likely still very far.

It is not uncommon for users of MCMC inferences to be aghast at the required run time. This is particularly the case when analyses are set up to incorporate too many different factors. Thus, minimizing the complexity of the setup from the start is generally a good idea. Ideally, we would want our analysis to be simple enough to be tractable, but complex enough to capture the relevant aspects of the data to answer our question. Unfortunately, the complexity that strikes that balance is often hard to determine a priori (or may not even exist for some combinations of question and data). While both gradually simplifying an overly complex model or gradually adding complexity to an overly simple model should be feasible strategies, we feel that erring on the side of simplicity may be more advisable. A successfully completed analysis that ends up being overly simplistic provides more information on how to improve it than an overly complex one that fails to run in the first place. Preliminary model testing, such as determining the most suitable substitution model using jModelTest (Posada & Crandall, 1998; Posada, 2008; Darriba *et al*., 2012), can help us narrow down an appropriate range of complexity to start at. Notably, the molecular substitution model in particular may not be the component that adds the most complexity to a phylogenetic analysis 24 (and may ultimately not have much of an impact overall, see Abadi *et al*., 2019; Guimarães Fabreti & Höhna, 2023), although the choice of morphological substitution model remains essential for the accuracy of the estimates.

An important contributor to analysis complexity is the number of partitions, so it is good to consider whether they are all needed, and if some of them can share substitution or clock models. In particular, if you notice in the trace that parameters associated with some partitions are purely driven by the prior, then the data is likely over-partitioned. Similarly, using uncorrelated relaxed clock models increases the number of parameters by a large amount, as each branch of the tree is then associated with its own clock rate. If the dataset contains little time information, then there will be little signal in the data to estimate these rates, which is likely to lead to convergence issues. Luckily, there exist several tools to help determine what number of partitions may be best for a given dataset. For molecular data, the software package PartitionFinder (Lanfear *et al.*, 2012; Lanfear *et al.*, 2017) can similarly be used to find partitions and test for the best substitution models for them. Its output files can be used as input for the aforementioned ClockStaR (Duchêne *et al.*, 2014), to further determine which partitions require different clock models. We have already mentioned how EvoPhylo (Simões *et al.*, 2023) can be used to partition morphological character data for clock models. There are a number of approaches to partitioning morphological data (Casali *et al.*, 2023; Clarke & Middleton, 2008; Rosa *et al.*, 2019). However, conventional approaches used to select among partition schemes for substitution models cannot be used if the size of the Q matrix varies across partition schemes. An alternative approach in these circumstances is to use model adequacy (Mulvey *et al.*, 2024). See also Casali *et al.* (2023) and Khakurel *et al.* (2024) for further details on the topic of partitions.

Beyond partitions, model adequacy testing (e.g., using posterior predictive simulations, PPS, as previously described in section 5.2.2) can also tell us whether our models are of appropriate complexity for the data. If the complexity of the

model does not match that of the data, the differences in the summary statistics between the data and the posterior simulations should show that. However, as mentioned above, unlike preliminary model testing, PPS approaches come into play after the main analysis, as they rely on having the successfully inferred posterior distributions. Thus, starting with a level of complexity that prevents successful completion of an MCMC run will prevent us from using this approach.

Finally, note that using informative priors helps reduce the complexity of the analysis, by reducing the size of the parameter space that needs to be explored by the inference. This is especially true for parameters for which there is little signal in the data, such as the clock rate in an analysis with little time calibration information, or the extinction rate in an analysis with only extant species. This principle is visualized very nicely in Figure 1 of dos Reis *et al.* (2016). For these parameters many different values will result in very similar posterior densities, so the inference can spend a large amount of time exploring a very wide plausible range of values. In this case constraining the values using fairly strict priors will ensure that the inference converges more quickly.

6. Good places to look for help

In addition to the guidance provided in this document, many software-specific resources can help in diagnosing and fixing misbehaving phylogenetic inferences. Bayesian inference frameworks are generally associated with a manual, some tutorials and help repositories which provide guidance on frequently used analyses. Specifically, users can look at the built-in help messages in MrBayes, the Taming the BEAST website (https://taming-the-beast.org) for BEAST2 or the RevBayes website (https://revbayes.github.io/tutorials/) for RevBayes. In particular, the RevBayes tutorial associated with this article (https://revbayes.github.io/tutorials/mcmc_troubleshooting/) illustrates many of the issues introduced here, and presents troubleshooting advice targeted to the specific characteristics of the RevBayes framework.

For more detailed and targeted help, forums such as the BEAST user group (https://groups.google.com/g/beast-users) or the RevBayes user forum (https://groups.google.com/g/revbayes-users) are also available. Making good use of search engines can usually solve most common problems. If the problem appears to be due to a bug in the software (for instance, the inference crashes or returns non-sensical results), filing an issue on the Github repository is the best way to report it. Reporting an issue automatically alerts all developers, and makes the problem visible to other affected users. Note that for BEAST2, each package has a separate repository, so if the problem appears tied to a specific package the issue should be filed on the package repository rather than the general BEAST2 one. Before opening a new issue, you should make sure that the problem has not been reported already by looking through the list of open issues. As a last resort, developers can be contacted directly, although we recommend exploring the above resources first.

Several rules should be kept in mind when requesting help on forums or from tool developers and when filing issues. First, it is generally good to assume that any would-be helper will need to run the analysis themselves in order to identify the issue. Thus, all data, configuration and code files required to reproduce the problem should be included in the help request. The full error message or problematic output should also be included, so helpers can verify that they have correctly reproduced the issue. If possible, simplifying the analysis by removing elements which do not trigger the issue, or comparing the problematic analysis to a similar analysis which worked, will also be very helpful to track down a problem. Finally, detailed information on the computer configuration used (operating system type and version, software version, compiler version if the software was compiled manually, whether the analysis was run on a local machine or computer cluster) should be provided, particularly when the analysis crashes or fails to start.

Ethics and consent

Ethical approval and consent were not required.

Data availability

No data are associated with this article.

Acknowledgements

JBS was supported by funds from the European Union's Horizon 2020 Research and Innovation Programme under the Marie Sk lodowska-Curie grant agreement No. 101022928. OS was funded by NSF DEB 2045842 and a Swiss National Science Foundation Postdoc Mobility Fellowship (P500PB 203131). AMW was supported on NSF DEB 2045842 and NSF CIBR 2113425. CZ was funded by the National Natural Science Foundation of China (42172006).

We thank the Morlon, Uyeda, and Wright research groups for their feedback and suggestions on the manuscript. We also thank Daniel Casali, Zhirui Hu, Grégoire Clarté, Sungsik Kong, Bruna Farina, Peter Beerli, and David Černý for their detailed reviews and comments.

References

Abadi S, Azouri D, Pupko T, *et al.*: **Model selection may not be a mandatory step for phylogeny reconstruction.** *Nat Commun.* 2019; **10**(1): 934.
Publisher Full Text

Altekar G, Dwarkadas S, Huelsenbeck JP, *et al.*: **Parallel metropolis coupled Markov chain Monte Carlo for Bayesian phylogenetic inference.** *Bioinformatics.* 2004; **20**(3): 407–415.
PubMed Abstract | **Publisher Full Text**

Álvarez-Carretero S, Goswami A, Yang Z, *et al.*: **Bayesian estimation of species divergence times using correlated quantitative characters.** *Syst Biol.* 2019; **68**(6): 967–986.
PubMed Abstract | **Publisher Full Text**

Baele G, Lemey P, Bedford T, *et al.*: **Improving the accuracy of demographic and molecular clock model comparison while accommodating phylogenetic uncertainty.** *Mol Biol Evol.* 2012; **29**(9): 2157–2167.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Baldwin BG, Sanderson MJ: **Age and rate of diversification of the Hawaiian silversword alliance (Compositae).** *Proc Natl Acad Sci U. S. A.* 1998; **95**(16): 9402–9406.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Barido-Sottani J, Aguirre-Fernández G, Hopkins MJ, *et al.*: **Ignoring stratigraphic age uncertainty leads to erroneous estimates of species divergence times under the fossilized birth-death process.** *Proc Biol Sci.* 2019; **286**(1902): 20190685.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Barido-Sottani J, Pohle A, De Baets K, *et al.*: **Putting the F into FBD analysis: tree constraints or morphological data?** *Palaeontology.* 2023a; **66**(6): e12679.
Publisher Full Text

Barido-Sottani J, van Tiel NMA, Hopkins MJ, *et al.*: **Ignoring fossil age uncertainty leads to inaccurate topology and divergence time estimates in time calibrated tree inference.** *Front Ecol Evol.* 2020; **8**: 183.
Publisher Full Text

Barido-Sottani J, Żyła D, Heath TA: **Estimating the age of poorly dated fossil specimens and deposits using a total-evidence approach and the fossilized birth-death process.** *Syst Biol.* 2023b; **72**(2): 466–475.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Biswas N, Jacob PE, Vanetti P: **Estimating convergence of markov chains with *L*-lag couplings.** *Adv Neural Inf Process Syst.* 2019: 32.
Reference Source

Blackmon H, Demuth JP: **Estimating tempo and mode of Y chromosome turnover: explaining Y chromosome loss with the fragile Y hypothesis.** *Genetics.* 2014; **197**(2): 561–572.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Bollback JP: **Bayesian model adequacy and choice in phylogenetics.** *Mol Biol Evol.* 2002; **19**(7): 1171–1180.
PubMed Abstract | **Publisher Full Text**

Bouchard-Côté A, Chen M, Kuo L, *et al.*: **SMC (sequential monte carlo) for Bayesian phylogenetics**. Chen M-H, Kuo L, Lewis PO, editors. *Bayesian Phylogenetics: Methods, Algorithms, and Applications.* 2014; pp. 163–185.
Reference Source

Bouckaert RR, Drummond AJ: **bModelTest: Bayesian phylogenetic site model averaging and model comparison.** *BMC Evol Biol.* 2017; **17**(1): 42.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Bouckaert R, Heled J, Kühnert D, *et al.*: **BEAST 2: a software platform for Bayesian evolutionary analysis.** *PLoS Comput Biol.* 2014; **10**(4): e1003537.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Box GEP, Jenkins GM, Reinsel GC, *et al.*: *Time series analysis: forecasting and control.* John Wiley & Sons; 2015.
Reference Source

Brown JM: **Detection of implausible phylogenetic inferences using posterior predictive assessment of model fit.** *Syst Biol.* 2014; **63**(3): 334–348.
PubMed Abstract | **Publisher Full Text**

Brown JM, ElDabaje R: **PuMA: Bayesian analysis of partitioned (and unpartitioned) model adequacy.** *Bioinformatics.* 2009; **25**(4): 537–538.
PubMed Abstract | **Publisher Full Text**

Brown JW, Smith SA: **The past sure is tense: on interpreting phylogenetic divergence time estimates.** *Syst Biol.* 2018; **67**(2): 340–353.
PubMed Abstract | **Publisher Full Text**

Brown JM, Thomson RC: **The behavior of Metropolis-coupled Markov chains when sampling rugged phylogenetic distributions.** *Syst Biol.* 2018; **67**(4): 729–734.
PubMed Abstract | **Publisher Full Text**

Budd GE, Mann RP: **Two Notorious Nodes: A Critical Examination of Relaxed Molecular Clock Age Estimates of the Bilaterian Animals and Placental Mammals.** *Syst Biol*. May 2024; **73**(1): 223–234. ISSN 1063-5157.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Capobianco A: **How many characters are needed to reconstruct a phylogeny?** *Biol Lett.* 2025; **21**(10): 20250288.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Casali DM, Freitas FV, Perini FA: **Evaluating the impact of anatomical partitioning on summary topologies obtained with Bayesian phylogenetic analyses of morphological data.** *Syst Biol.* 2023; **72**(1): 62–77.
PubMed Abstract | **Publisher Full Text**

Cayley A: **A Theorem on Trees.** *Q J Pure Appl Math.* 1889; **23**: 376–378.

Clarke JA, Middleton KM: **Mosaicism, modules, and the evolution of birds: results from a Bayesian approach to the study of morphological evolution using discrete character data.** *Syst Biol.* 2008; **57**(2): 185–201.
PubMed Abstract | **Publisher Full Text**

Darriba D, Taboada GL, Doallo R, *et al.*: **jModelTest 2: more models, new heuristics and parallel computing.** *Nat Methods.* 2012; **9**(8): 772.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Degnan JH, Salter LA: **Gene tree distributions under the coalescent process.** *Evolution.* 2005; **59**(1): 24–37.
PubMed Abstract | **Publisher Full Text**

Dinh V, Bilge A, Zhang C, *et al.*: **Probabilistic path hamiltonian Monte Carlo**. In: *International Conference on Machine Learning PMLR.* 2017; pp. 1009–1018.
Reference Source

dos Reis M, Inoue J, Hasegawa M, *et al.*: **Phylogenomic datasets provide both precision and accuracy in estimating the timescale of placental mammal phylogeny.** *Proc Biol Sci.* 2012; **279**(1742): 3491–3500.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

dos Reis M, Thawornwattana Y, Angelis K, *et al.*: **Uncertainty in the timing of origin of animals and the limits of precision in molecular timescales.** *Curr Biol.* 2015; **25**(22): 2939–2950.
PubMed Abstract | **Publisher Full Text**

dos Reis M, Donoghue PCJ, Yang Z: **Bayesian molecular clock dating of species divergences in the genomics era.** *Nat Rev Genet.* 2016; **17**(2): 71–80.
Publisher Full Text

Drummond AJ, Ho SYW, Phillips MJ, *et al.*: **Relaxed phylogenetics and dating with confidence.** *PLoS Biol.* 2006; **4**(5): e88.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Drummond AJ, Nicholls GK, Rodrigo AG, *et al.*: **Estimating mutation parameters, population history and genealogy simultaneously from temporally spaced sequence data.** *Genetics.* 2002; **161**(3): 1307–1320.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Drummond AJ, Stadler T: **Bayesian phylogenetic estimation of fossil ages.** *Philos Trans R Soc Lond B Biol Sci.* 2016; **371**(1699): 20150129.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Duchene S, Bouckaert R, Duchene DA, *et al.*: **Phylodynamic model adequacy using posterior predictive simulations.** *Syst Biol.* 2019; **68**(2): 358–364.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Duchêne S, Molak M, Ho SYW: **ClockstaR: choosing the number of relaxed-clock models in molecular phylogenetic analysis.** *Bioinformatics.* 2014; **30**(7): 1017–1019.
PubMed Abstract | **Publisher Full Text**

Felsenstein J: **The number of evolutionary trees.** *Syst Zool.* 1978; **27**(1): 27–33.
Publisher Full Text

Felsenstein J: **Evolutionary trees from DNA sequences: a maximum likelihood approach.** *J Mol Evol.* 1981; **17**(6): 368–376.
PubMed Abstract | **Publisher Full Text**

Feng C: **MCMC interactive demo.** 2024.
Reference Source

Friedman M, Keck BP, Dornburg A, *et al.*: **Molecular and fossil evidence place the origin of cichlid fishes long after gondwanan rifting.** *Proc Biol Sci.* 2013; **280**: 20131733.
Publisher Full Text

Gavryushkina A, Heath TR, Ksepka DT, *et al.*: **Bayesian total-evidence dating reveals the recent crown radiation of penguins.** *Syst Biol.* 2017; **66**(1): 57–73.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Gavryushkina A, Welch D, Stadler T, *et al.*: **Bayesian inference of sampled ancestor trees for epidemiology and fossil calibration.** *PLoS Comput Biol.* 2014; **10**(12): e1003919.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Geyer CJ: **Practical Markov Chain Monte Carlo.** *Statist Sci.* 1992; **7**(4): 473–483.
Publisher Full Text

Gilks WR, Roberts GO: **Strategies for improving MCMC.** *Markov Chain Monte Carlo in Practice.* 1996; **6**: 89–114.
Publisher Full Text

Green PJ: **Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination.** *Biometrika.* 1995; **82**(4): 711–732.
Publisher Full Text

Guimarães Fabreti L, Höhna S: **(Luiza Guimarães Fabreti and Sebastian Höhna contributed equally): Convergence assessment for Bayesian phylogenetic analysis using MCMC simulation.** *Methods Ecol Evol.* 2022; **13**(1627): 77–90.
Publisher Full Text

Guimarães Fabreti L, Höhna S: **Nucleotide substitution model selection is not necessary for bayesian inference of phylogeny with well-behaved priors.** *Syst Biol.* 2023; **72**(6): 1418–1432.
Publisher Full Text

Hasegawa M, Kishino H, Yano T: **Dating of the human-ape splitting by a molecular clock of mitochondrial DNA.** *J Mol Evol.* 1985; **22**(2): 160–174.
PubMed Abstract | **Publisher Full Text**

Hastings WK: **Monte Carlo sampling methods using Markov chains and their applications.** *Biometrika*. 1970; **57**(1): 97–109. ISSN 00063444.
Publisher Full Text

Heath TA, Hedtke SM, Hillis DM: **Taxon sampling and the accuracy of phylogenetic analyses.** *J Syst Evol.* 2008; **46**(3): 239.
Reference Source

Heath TA: **A hierarchical bayesian model for calibrating estimates of species divergence times.** *Syst Biol*. 2012; **61**(5): 793–809. ISSN 1063-5157.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Heath TA, Huelsenbeck JP, Stadler T: **The fossilized birth-death process for coherent calibration of divergence-time estimates.** *Proc Natl Acad Sci U. S. A.* 2014; **111**(29): E2957–E2966.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Heled J, Drummond AJ: **Calibrated tree priors for relaxed phylogenetics and divergence time estimation.** *Syst Biol.* 2012; **61**(1): 138–149.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Hillis DM, Pollock DD, McGuire JA, *et al.*: **Is sparse taxon sampling a problem for phylogenetic inference?** *Syst Biol.* 2003; **52**(1): 124–126.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Höhna S, Landis MJ, Heath TA, *et al.*: **RevBayes: Bayesian phylogenetic inference using graphical models and an interactive model-specification language.** *Syst Biol.* 2016; **65**(4): 726–736.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Höhna S, Landis MJ, Heath TA: **Phylogenetic inference using RevBayes.** *Curr Protoc Bioinform.* 2017; **57**(1): 6–16.
Publisher Full Text

Huelsenbeck JP, Ronquist F: **MRBAYES: Bayesian inference of phylogenetic trees.** *Bioinformatics.* 2001; **17**(8): 754–755.
PubMed Abstract | **Publisher Full Text**

Huelsenbeck JP, Nielsen R, Bollback JP: **Stochastic mapping of morphological characters.** *Syst Biol.* 2003; **52**(2): 131–158.
PubMed Abstract | **Publisher Full Text**

Huelsenbeck JP, Larget B, Alfaro ME: **Bayesian phylogenetic model selection using reversible jump Markov chain Monte Carlo.** *MMol Biol Evol.* 2004; **21**(6): 1123–1133.
PubMed Abstract | **Publisher Full Text**

Inoue J, Donoghue PCJ, Yang Z: **The impact of the representation of fossil calibrations on bayesian estimation of species divergence times.** *Syst Biol.* 2010; **59**(1): 74–89.
PubMed Abstract | **Publisher Full Text**

Jukes TH, Cantor CR: **Evolution of protein molecules.** *Mammalian Protein Metabolism.* 1969; **3**: 21–132.
Publisher Full Text

Kelly LJ, Nicholls GK, Ryder RJ, *et al.*: **Traitlab: a Matlab package for fitting and simulating binary tree-like data.** *arXiv preprint arXiv.: 2308.09060.* 2023a.
Publisher Full Text

Kelly LJ, Ryder RJ, Clarté G: **Lagged couplings diagnose Markov Chain Monte Carlo phylogenetic inference.** *Ann Appl Stat.* 2023b; **17**(2): 1419–1443.
Publisher Full Text

Khakurel B, Grigsby C, Tran TD, *et al.*: **The fundamental role of character coding in Bayesian morphological phylogenetics.** *Syst Biol.* 2024: syae033.
PubMed Abstract | **Publisher Full Text**

Kimura M: **A simple method for estimating evolutionary rates of base substitutions through comparative studies of nucleotide sequences.** *J Mol Evol.* 1980; **16**(2): 111–120.
PubMed Abstract | **Publisher Full Text**

Kolmogorov A: **Sulla determinazione empirica di una legge didistribuzione.** *Giorn Dell'inst Ital Degli Att.* 1933; **4**: 89–91.
Reference Source

Kong S, Pons JC, Kubatko L, *et al.*: **Classes of explicit phylogenetic networks and their biological and mathematical significance.** *J Math Biol.* 2022; **84**(6): 47.
PubMed Abstract | **Publisher Full Text**

Lambert A, Stadler T: **Birth-death models and coalescent point processes: the shape and probability of reconstructed phylogenies.** *Theor Popul Biol.* 2013; **90**: 113–128.
PubMed Abstract | **Publisher Full Text**

Lanfear R, Calcott B, Ho SYW, *et al.*: **PartitionFinder: combined selection of partitioning schemes and substitution models for phylogenetic analyses.** *Mol Biol Evol.* 2012; **29**(6): 1695–1701.
PubMed Abstract | **Publisher Full Text**

Lanfear R, Hua X, Warren DL: **Estimating the effective sample size of tree topologies from bayesian phylogenetic analyses.** *Genome Biol Evol.* 2016; **8**(8): 2319–2332.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Lanfear R, Frandsen PB, Wright AM, *et al.*: **PartitionFinder 2: new methods for selecting partitioned models of evolution for molecular and morphological phylogenetic analyses.** *Mol Biol Evol.* 2017; **34**(3): 772–773.
PubMed Abstract | **Publisher Full Text**

Larget B, Simon DL: **Markov chain monte carlo algorithms for the bayesian analysis of phylogenetic trees.** *Mol Biol Evol.* 1999; **16**(6): 750–759.
Publisher Full Text

Lartillot N: **Conjugate Gibbs sampling for Bayesian phylogenetic models.** *J Comput Biol.* 2006; **13**(10): 1701–1722.
PubMed Abstract | **Publisher Full Text**

Lartillot N, Philippe H: **Computing bayes factors using thermodynamic integration.** *Syst Biol.* 2006; **55**(2): 195–207.
Publisher Full Text

Lewis PO: **A likelihood approach to estimating phylogeny from discrete morphological character data.** *Syst Biol.* 2001; **50**(6): 913–925.
PubMed Abstract | **Publisher Full Text**

Lewis PO: **MCMC robot.** 2024.
Reference Source

Lewis PO, Xie W, Chen MH, *et al.*: **Posterior predictive Bayesian phylogenetic model selection.** *Syst Biol.* 2014; **63**(3): 309–321.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Li S: *Phylogenetic tree construction using markov chain monte carlo.* PhD thesis, The Ohio State University; 1996.

Magee AF, Höhna S: **Impact of k-pg mass extinction event on crocodylomorpha inferred from phylogeny of extinct and extant taxa.** *bioRxiv.* 2021.
Publisher Full Text

Matschiner M, Musilová Z, Barth JMI, *et al.*: **Bayesian phylogenetic estimation of clade ages supports trans-atlantic dispersal of cichlid fishes.** *Syst Biol.* 2017; **66**(1): 3–22.
PubMed Abstract | **Publisher Full Text**

Mau B, Newton MA: **Phylogenetic inference for binary data on dendograms using Markov chain Monte Carlo.** *J Comput Graph Stat.* 1997; **6**(1): 122–131.
Publisher Full Text

Mau B, Newton MA, Larget B: **Bayesian phylogenetic inference via Markov chain Monte Carlo methods.** *Biometrics.* 1999; **55**(1): 1–12.
PubMed Abstract | **Publisher Full Text**

May MR, Rothfels CJ: **Diversification models conflate likelihood and prior, and cannot be compared using conventional model-comparison tools.** *Syst Biol.* 2023; **72**(3): 713–722.
PubMed Abstract | **Publisher Full Text**

Metropolis N, Rosenbluth AW, Rosenbluth MN, *et al.*: **Equation of state calculations by fast computing machines.** *J Chem Phys.* 1953; **21**(6): 1087–1092.
Publisher Full Text

Müller NF, Bouckaert RR: **Adaptive metropolis-coupled MCMC for BEAST 2.** *PeerJ.* 2020; **8**: e9473.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Mulvey LPA, May MR, Brown JM, *et al.*: **Assessing the adequacy of morphological models used in palaeobiology.** *BioRxiv.* 2024: 2024–2001.
Publisher Full Text

Mulvey LPA, Nikolic MC, Allen BJ, *et al.*: **From fossils to phylogenies: exploring the integration of paleontological data into bayesian phylogenetic inference.** *Paleobiology.* 2025; **51**(1): 214–236.
Publisher Full Text

Neal RM: **Slice sampling.** *Ann. Stat.* 2003; **31**(3): 705–767.
Publisher Full Text

Nguyen LT, Schmidt HA, von Haeseler A, *et al.*: **IQ-TREE: a fast and effective stochastic algorithm for estimating maximum-likelihood phylogenies.** *Mol Biol Evol.* 2015; **32**(1): 268–274.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Nielsen R: **Mapping mutations on phylogenies.** *Syst Biol.* 2002; **51**(5): 729–739.
PubMed Abstract | **Publisher Full Text**

Nylander JA, Wilgenbusch JC, Warren DL, *et al.*: **AWTY (Are We There Yet?): a system for graphical exploration of MCMC convergence in Bayesian**

phylogenetics. *Bioinformatics.* 2008; **24**(4): 581–583.
PubMed Abstract | Publisher Full Text

O'Reilly JE, Donoghue PCJ: **The effect of fossil sampling on the estimation of divergence times with the Fossilized Birth-Death process.** *Syst Biol.* 2020; **69**(1): 124–138.
PubMed Abstract | Publisher Full Text

Pennell MW, FitzJohn RG, Cornwell WK, *et al.*: **Model adequacy and the macroevolution of angiosperm functional traits.** *Am Nat.* 2015; **186**(2): E33–E50.
PubMed Abstract | Publisher Full Text

Plummer M, Best N, Cowles K, *et al.*: **CODA: convergence diagnosis and output analysis for MCMC.** *R News.* 2006; **6**(1): 7–11.
Reference Source

Portik DM, Streicher JW, Blackburn DC, *et al.*: **Redefining possible: combining phylogenomic and supersparse data in frogs.** *Mol Biol Evol.* 2023; **40**(5): msad109.
PubMed Abstract | Publisher Full Text | Free Full Text

Posada D: **jModelTest: phylogenetic model averaging.** *Mol Biol Evol.* 2008; **25**(7): 1253–1256.
PubMed Abstract | Publisher Full Text

Posada D, Crandall KA: **ModelTest: testing the model of DNA substitution.** *Bioinformatics.* 1998; **14**(9): 817–818.
PubMed Abstract | Publisher Full Text

Rambaut A, Drummond AJ, Xie D, *et al.*: **Posterior summarization in Bayesian phylogenetics using Tracer 1.7.** *Syst Biol.* 2018; **67**(5): 901–904.
PubMed Abstract | Publisher Full Text | Free Full Text

Rannala B, Yang Z: **Probability distribution of molecular evolutionary trees: a new method of phylogenetic inference.** *J Mol Evol.* 1996; **43**: 304–311.
PubMed Abstract | Publisher Full Text

Rannala B, Huelsenbeck JP, Yang Z, *et al.*: **Taxon sampling and the accuracy of large phylogenies.** *Syst Biol.* 1998; **47**(4): 702–710.
PubMed Abstract | Publisher Full Text

Rannala B, Yang Z: **Inferring speciation times under an episodic molecular clock.** *Syst Biol.* 2007; **56**(3): 453–466.
PubMed Abstract | Publisher Full Text

Reid NM, Hird SM, Brown JM, *et al.*: **Poor fit to the multispecies coalescent is widely detectable in empirical data.** *Syst Biol.* 2014; **63**(3): 322–333.
PubMed Abstract | Publisher Full Text

Ronquist F, Klopfstein S, Vilhelmsen L, *et al.*: **A total-evidence approach to dating with fossils, applied to the early radiation of the Hymenoptera.** *Syst Biol.* 2012; **61**(6): 973–999.
PubMed Abstract | Publisher Full Text | Free Full Text

Rosa BB, Melo GAR, Barbeitos MS: **Homoplasy-based partitioning outperforms alternatives in Bayesian analysis of discrete morphological data.** *Syst Biol.* 2019; **68**(4): 657–671.
PubMed Abstract | Publisher Full Text

Russel PM, Brewer BJ, Klaere S, *et al.*: **Model selection and parameter inference in phylogenetics using nested sampling.** *Syst Biol.* 2018; **68**(2): 219–233.
PubMed Abstract | Publisher Full Text

Schwery O, Freyman W, Goldberg EE: **adequaSSE: model adequacy testing for trait-dependent diversification models.** *bioRxiv.* 2023: 2023–2003.
Publisher Full Text

Schwery O, O'Meara BC: ***BoskR* – testing adequacy of diversification models using tree shape.** *bioRxiv.* 2020.
Publisher Full Text

Scire J, Barido-Sottani J, Kühnert D, *et al.*: **Robust phylodynamic analysis of genetic sequencing data from structured populations.** *Viruses.* 2022; **14**(8): 1648.
PubMed Abstract | Publisher Full Text | Free Full Text

Silvestro D, Warnock RCM, Gavryushkina A, *et al.*: **Closing the gap between palaeontological and neontological speciation and extinction rate estimates.** *Nat Commun.* 2018; **9**(1): 5237.
PubMed Abstract | Publisher Full Text | Free Full Text

Simões TR, Greifer N, Barido-Sottani J, *et al.*: **EvoPhylo: an R package for pre- and postprocessing of morphological data from relaxed clock Bayesian phylogenetics.** *Methods Ecol Evol.* 2023; **14**(8): 1981–1993.
Publisher Full Text

Slater GJ, Pennell MW: **Robust regression and posterior predictive simulation increase power to detect early bursts of trait evolution.** *Syst Biol.* 2014; **63**(3): 293–308.
PubMed Abstract | Publisher Full Text

Smirnov NV: **On the estimation of the discrepancy between empirical curves of distribution for two independent samples.** *Bull Math Univ Moscou.* 1939; **2**(2): 3–11.
Reference Source

Stadler T: **Sampling-through-time in birth-death trees.** *J Theor Biol.* 2010; **267**(3): 396–404.
PubMed Abstract | Publisher Full Text

Stadler T, Gavryushkina A, Warnock RCM, *et al.*: **The fossilized birth-death model for the analysis of stratigraphic range data under different speciation modes.** *J Theor Biol.* 2018; **447**: 41–55.
PubMed Abstract | Publisher Full Text | Free Full Text

Sunnåker M, Busetto AG, Numminen E, *et al.*: **Approximate Bayesian computation.** *PLoS Comput Biol.* 2013; **9**(1): e1002803.
PubMed Abstract | Publisher Full Text | Free Full Text

Tavaré S: **Some probabilistic and statistical problems in the analysis of DNA sequences**. *Some Mathematical Questions in Biology: DNA Sequence Analysis.* 1986; **17**: 57–86.
Reference Source

Thompson MB: **A comparison of methods for computing autocorrelation time.** *arXiv.* 2010.
Publisher Full Text

Thorne JL, Kishino H, Painter IS: **Estimating the rate of evolution of the rate of molecular evolution.** *Mol Biol Evol.* 1998; **15**(12): 1647–1657.
PubMed Abstract | Publisher Full Text

Uhen MD, Allen B, Behboudi N, *et al.*: **Paleobiology database user guide version 1.0.** *PaleoBios.* 2023; **40**(11).
Publisher Full Text

Varón-González C, Whelan S, Klingenberg CP: **Estimating phylogenies from shape and similar multidimensional data: why it is not reliable.** *Syst Biol.* 2020; **69**(5): 863–883.
PubMed Abstract | Publisher Full Text

Wang Y-B, Milkey A, Li A, *et al.*: **Lorad: Marginal likelihood estimation with haste (but no waste).** *Syst Biol.* 2023; **72**(3): 639–648.
PubMed Abstract | Publisher Full Text

Warnock RCM, Yang Z, Donoghue PCJ: **Exploring uncertainty in the calibration of the molecular clock.** *Biol Lett.* 2012; **8**: 156–159.
Publisher Full Text

Warnock RCM, Parham JF, Joyce WG, *et al.*: **Calibration uncertainty in molecular dating analyses: there is no substitute for the prior evaluation of time priors.** *Proc Biol Sci.* 2015; **282**(1798): 20141013.
PubMed Abstract | Publisher Full Text | Free Full Text

Warren D, Geneva A, Lanfear R: **RWTY (R We There Yet): an R package for examining convergence of Bayesian phylogenetic analyses.**2017.
Reference Source

Whidden C, Matsen FA IV: **Quantifying MCMC exploration of phylogenetic tree space.** *Syst Biol.* 2015; **64**(3): 472–491.
PubMed Abstract | Publisher Full Text | Free Full Text

Wilkinson RD, Steiper ME, Soligo C, *et al.*: **Dating primate divergences through an integrated analysis of palaeontological and molecular data.** *Syst Biol.* 2011; **60**(1): 16–31.
PubMed Abstract | Publisher Full Text | Free Full Text

Wright AM: **A systematist's guide to estimating Bayesian phylogenies from morphological data.** *Insect Syst Divers.* 2019; **3**.
PubMed Abstract | Publisher Full Text | Free Full Text

Xie W, Lewis PO, Fan Y, *et al.*: **Improving marginal likelihood estimation for Bayesian phylogenetic model selection.** *Syst Biol.* 2011; **60**(2): 150–160.
PubMed Abstract | Publisher Full Text | Free Full Text

Yang Z: *Molecular evolution: a statistical approach.* Oxford University Press; 2014.
Reference Source

Yang Z, Rannala B: **Bayesian estimation of species divergence times under a molecular clock using multiple fossil calibrations with soft bounds.** *Mol Biol Evol.* 2006; **23**(1): 212–226.
PubMed Abstract | Publisher Full Text

Zhang C, Stadler T, Klopfstein S, *et al.*: **Total-evidence dating under the Fossilized Birth-Death process.** *Syst Biol.* 2016; **65**(2): 228–249.
PubMed Abstract | Publisher Full Text | Free Full Text

Zhang C, Ronquist F, Stadler T: **Skyline fossilized birth–death model is robust to violations of sampling assumptions in total-evidence dating.** *Syst Biol.* 2023; **72**(6): 1316–1336.
PubMed Abstract | Publisher Full Text

Zhang R, Drummond AJ, Mendes FK: **Fast Bayesian inference of phylogenies from multiple continuous characters.** *Syst Biol.* 2024; **73**(1): 102–124.
PubMed Abstract | Publisher Full Text | Free Full Text

Zwickl DJ, Holder MT: **Model parameterization, prior distributions, and the general time-reversible model in Bayesian phylogenetics.** *Syst Biol.* 2004; **53**(6): 877–888.
PubMed Abstract | Publisher Full Text

Open Peer Review

Current Peer Review Status: ✔ ✔ ? ✔ ✔ ? ?

Version 3

Reviewer Report 24 December 2024

https://doi.org/10.21956/openreseurope.19757.r46216

© **2024 Černý D.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

David Černý
[1] The University of Chicago, Chicago, Illinois, USA
[2] The University of Chicago, Chicago, Illinois, USA

In this paper, the authors offer guidelines for setting up and troubleshooting phylogenetic Markov chain Monte Carlo (MCMC) analyses, with a particular emphasis on the inference of time trees using the fossilized birth-death (FBD) process. The manuscript first sets the scene by explaining the basic terminology (move / operator / proposal, convergence / stationarity, etc.), and as such it will serve as a gentle introduction to the topic for those researchers who are hoping to become more informed users of Bayesian phylogenetic methods.

Overall, the manuscript retreads a lot of familiar ground (Fig. 4, showing traces produced by MCMC runs with moves that are too small or too large, feels like it has been recreated in the literature dozens of times now), but there is enough useful detail (for example, about differences between major software packages) to make it worth publishing. I have a lot of minor comments that are listed below, and some of these exemplify what I would consider to be two larger problems. First, the manuscript is not structured as well as it could be, which negatively affects the flow of the text. A lot of this has to do with jumping between different levels of abstraction. I would recommend explaining models first (highest level), then Bayesian inference of model parameters and how it differs from maximum likelihood, followed by concerns that are specific to phylogenetics (e.g., differences between dated and undated trees). Only then would I dive into an explanation of the phylogenetic MCMC, and give tips and recommendations about its use. As it stands, the authors start describing the MCMC before they even explain the meaning of the term “likelihood”, which just doesn’t work very well. There are several passages interspersed throughout the text that anticipate later sections of the manuscript, but look out of place when the reader first encounters them.

My second concern has to do with several problematic instances where the authors seemingly misunderstood the very concepts they are trying to explain – these are likely due to pedagogical concerns about keeping things simple, or just sloppy use of terminology. For example, the authors suggest in several places that the MCMC optimizes parameter values or even searches for the true

values (it does no such thing), they confuse the acceptance ratio with the Hastings ratio (while leaving the actual Hastings ratio unmentioned and unexplained), and the prior distribution with the proposal distribution. I have absolutely no doubt about the authors' expertise, but such relatively basic errors are worrisome, and a manuscript that contains them cannot do the topic justice. Conditional on these errors being fixed, however, I can see the manuscript being useful to many biologists who want to learn more about the topic but are unsure where to start.

Specific comments

Throughout the paper:

(1) Inconsistent American vs. British spellings (behavior vs. behaviour, optimized vs. optimised).

Abstract

(2) I would take issue with the statement that "[m]odern statistical methods" employ models "involving fossil record information". Sure, some do, but most active areas of phylogenetic research have nothing to do with analyzing fossil data. I understand the FBD model is of particular interest to the authors, but here they are perhaps getting a little ahead of themselves.

(3) It seems backwards to mention the MCMC and posterior distributions before mentioning Bayesian analysis.

Section 1

(4) The first example of a model chosen by the authors is perhaps a bit unfortunate: the typical parameterization involves a transversion/transition *ratio*, rather than a separate parameter for transitions and another one for transversions.

(5) The authors again seem to be getting ahead of themselves when they describe model parameters as random variables. That's indeed how they are treated in Bayesian inference, but in maximum-likelihood estimation (and frequentist statistics more generally), every parameter is assumed to have a fixed but unknown value. Since the differences between these two frameworks have not yet been discussed at this point, the authors need to be more careful about not attaching interpretations to model parameters that are only valid under one of them.

(6) "Bayesian inference thus offers a more complete picture of the results, integrating uncertainty in the results as well as existing information from previous studies" – Though I use and have developed Bayesian phylogenetic methods myself, I find this statement to be overselling things in an unacceptable way. First of all, most of the time, we as Bayesian phylogeneticists are desperate to find "flat" or "uninformative" priors (this includes seemingly mechanistic priors like the birth-death process, whose big attraction is actually the fact that it is flat over labeled histories), so the need to incorporate prior information is more of a burden than an advantage. Second, even when we do pick informative priors (and often it's impossible not to), they are almost never based on previous studies. Third, the statement completely ignores the old but important objection that priors are subjective.

(7) "true distribution of plausible parameters" – The authors mean "posterior", not "true". I'm not

sure the notion of a “true distribution” even makes sense to me, since at the end of the day, I do believe there was just one true evolutionary history that generated my data, not a distribution of them (so I’m perhaps a likelihoodist at heart after all).

Section 2.1

(8) Once again, I have to point out that introducing the MCMC before Bayesian analysis seems to me like putting the cart before the horse. It’s really confusing that “likelihood” is only explained several paragraphs after bringing up “maximum likelihood”! I would strongly recommend rearranging this part of the MS.

(9) 2.1.1 – Mentioning clock models and rates only makes sense in the context of inferring time trees; what about time-free / undated trees, where rates and times are confounded as branch lengths? This is yet another instance of concepts not being explained at a high enough level of abstraction when they are first introduced.

(10) Also 2.1.1 – It may be worth pointing out that the single parameter of JC69 and Mk is not free, so there’s actually nothing for us to estimate in these models.

(11) 2.1.2 – “a rate that is very high is believed to be less likely than one that is very short” – This sentence doesn’t work; rates can’t be “short”.

(12) Also 2.1.2 – I expect many readers to not know what it means for parameters to not be identifiable.

(13) 2.1.5 – The nomenclature I’m familiar with is that the acceptance ratio is the product of the posterior odds ratio and the Hastings ratio, so I find the authors’ attempt to conflate these two confusing. The actual Hastings ratio – i.e., p(theta | theta’) / p(theta’ | theta) – doesn’t make an appearance in the authors’ explanation, so what they are describing is actually the Metropolis algorithm rather than the more general Metropolis-Hastings algorithm.

(14) 2.1.6 – The word “subclade” seems unnecessary here. I would just use “clade”.

Section 2.2

(15) Subsection 2.2.1 seems both misplaced and misnamed, since the authors focus on the difference between dated and undated (rather than rooted and unrooted) trees.

(16) Also 2.2.1 – In addition to fossils, it may be good to mention the other common type of time-stamped data: serially sampled viral sequences.

(17) 2.2.2 – This description is accurate but somewhat confusing, since it seems to conflate two kinds of modularity. One is a property of the models used, the other is a feature of software design.

(18) 2.2.3 – While this is the RevBayes default and as such perhaps the only thing the reader needs to know about the program, RevBayes *can* be made to count iterations the MrBayes/BEAST way (Höhna et al. 2017).

(19) Also 2.2.3 – It may be useful to point out that some special moves such as the slice move (Neal 2003) are supposed to have very high acceptance ratios.

Section 3

(20) "This can raise serious performance issues for our MCMC inference, either when exploring the tree space or when calculating the posterior probabilities" – Those are not two separate issues, but two ways of describing the same issue. If you are not doing a good job of exploring the tree space, you are not accurately approximating the posterior probabilities.

(21) 3.1 – The second paragraph tries to reintroduce the topic of tree moves from scratch, despite the fact that they were already described (even including specific examples) several paragraphs above. Here, I would maybe recommend removing the earlier mention and merging it with this latter one to better preserve the flow of the exposition.

(22) 3.2 – "the parameters of the gamma distribution representing among-site rate heterogeneity" – Usually, this distribution only has one free parameter, since its mean needs to equal 1.

(23) 3.3 – The point about tip calibrations being less subjective than node calibrations is an interesting one that made me think. Is this actually the case, or are we just more comfortable placing uniform priors on tip ages compared to node ages? Node calibrations inherently involve more sources of uncertainty, but uncertainty and subjectivity are two different things, and many attempts have been made to design node calibrations in a more objective way (Wilkinson et al. 2011; Friedman et al. 2013; Matschiner et al. 2017; Claramunt 2022).

(24) Also 3.3 – "additional information could be used to construct more informative non-uniform priors on fossil ages" – Yes, one could argue a non-uniform distribution would often be a better representation of our prior knowledge. What if a fossil comes from a Maastrichtian formation, but specifically from the base of the middle member of that formation? What if it is sandwiched between two radiometrically dated horizons, but is closer to one than the other? Etc. But that leads to two problems: (1) the subjectivity is back, big time; and (2) some implementations are not flexible enough to handle non-uniform tip age priors. MrBayes is, and BEAST 2 might be too, but while I don't consider myself a RevBayes beginner, I'm not sure how I'd torture the program to allow them. Here, it might be worth distinguishing between what's possible in principle and what's possible in practice.

Section 4

(25) What I miss from the discussion of effective sample sizes is a clear statement that they only quantify the precision of the inference. They tell you whether you have enough samples to achieve a low standard error of the mean, but as the authors themselves point out, in and of themselves they say nothing about whether these samples come from the posterior distribution. With multiple chains, one would definitely hope that the only way to achieve high ESS values is for these chains to converge on the posterior, but even in that scenario there are certain caveats (e.g., it's not impossible for multiple chains to get stuck in the same region of the parameter space if the posterior is very rugged and their starting values are similar or even the same, as can happen when the user supplies their own initial tree).

(26) 4.1 – This is very pedantic, but the qualifier “typically” in the 3rd sentence of the 1st paragraph is unnecessary. In the context of MCMC, convergence is short for “convergence to the stationary distribution”; there is no other kind.

(27) 4.1 (1st paragraph) – I would take the opposite view with regard to the relative dangers of setting the burnin too low/high. If you set it too low, you can always discard extra samples after the fact, but if you set it too high, you’ve wasted time and CPU cycles and there is nothing you can do.

(28) 4.1 (3rd paragraph) – Is N the number of generations, or the number of moves? This might be a distinction worth drawing given the previous caveat about how different programs count generations in different ways.

(29) 4.1 (4th paragraph) – “the individual ESS of each optimized parameter” – Wrong nomenclature. MCMC doesn’t optimize; it samples.

(30) 4.1 (4th paragraph) – The point about branch rate ESS values being meaningless is an important one, but I would perhaps phrase it in a different way. It’s not that these are parameters that are unlikely to converge, it’s that our diagnostic tools conflate different parameters by giving them the same name (e.g., branch_rates[63] on one topology is a different parameter from branch_rates[63] on another topology). In principle, one could develop tools that do not do this. We could instead extract all unique branches present in the MCMC sample, and compute the ESS of each one’s rate.

(31) 4.1 (5th paragraph) – “are guaranteed to converge” – I would qualify this statement with “when run for long enough”.

(32) 4.1 (5th paragraph) – “in a local optima” – The word “optima” is plural. The corresponding singular is “optimum”.

(33) Caption to Figure 4 – “is doing baby steps” – I’d prefer a more formal phrasing.

(34) 4.2.2 – I would change the punctuation from “[…] of the phylogeny, however it is preferable […]” to “[…] of the phylogeny; however, it is preferable […]”.

(35) Also 4.2.2 – I would also cite Lanfear et al. (2016) here, who first showed how to approximate topological ESS (with RWTY providing an implementation of their algorithms).

(36) 4.2.3 – “This is likely due to one or both chains not having converged yet” – The qualifier “likely” is not needed. The stationary distribution of a correctly constructed MCMC sampler is the posterior distribution of the model, so if it is sampling from a distribution other than the posterior, it is by definition because it has not yet converged to its stationary distribution. The scenario where different chains get stuck at different local optima is not a different type of failure (as implied by the word “also”); rather, it is one possible reason why the chains have not converged to their stationary distribution.

Section 5

(37) 5.1.1 – "Each parameter that is supposed to be estimated by the analysis needs to have at least one operator associated to it, in order to be optimised" – Again, the MCMC doesn't optimize. (It is closely related to an optimization technique called simulated annealing, but that's beside the point here.) I'm surprised the authors would make such a basic mistake after explaining the difference between maximum likelihood and Bayesian inference.

(38) Also 5.1.1 – "so that more moves are being proposed each generation for them" – Putting it this way presupposes the default RevBayes setup with multiple moves per generation; it doesn't accurately describe what increasing the weight does in MrBayes or BEAST.

(39) 5.1.2 (2nd paragraph) – I can think of absolutely no biological justification for setting the extinction rate at 1/10 of the speciation rate. In fact, paleobiology tells us that on average, speciation and extinction rates tend to be very nearly equal (Marshall 2017). This might not make for the best starting values, but in that case, let's not pretend our starting values are motivated biologically rather than by computational convenience.

(40) 5.1.2 (3rd paragraph) – "starting values do not have to be spot-on estimates of where the actual true values lie, because after all, the MCMC is expected to go find those" – No, the MCMC is not expected to find the true values. This is a fundamental misunderstanding of what it does, and one that I find quite disturbing in a paper that takes it upon itself to introduce the technique to others.

(41) 5.1.2 (3rd paragraph) – "instead allow it to start exploring feasible solutions instead" – There is one more "instead" in this sentence than needed.

(42) 5.1.3 – "each run could potentially be sampling different local optima" – Here, I'd use the singular: "each run could potentially be sampling a different local optimum".

(43) 5.2 – "the probability of the tree under the FBD process given the fossil ages is considered part of the prior by BEAST2 and RevBayes" – That's correct for BEAST 2 but not for RevBayes. In the latter, if you run your analysis with the underPrior=TRUE option (available up to v1.2.4) or with the .ignoreAllData() option (which supplanted the former in v1.2.5), fossil ages will be discarded along with the character data. You never want this in an FBD analysis, since you're not learning anything useful by pretending you don't know if there were bears in the Jurassic or trilobites in the Miocene.

(44) 5.2.1 (2nd paragraph) – "Overall, improper priors are statistically problematic and should be avoided if possible" – This is a bit of a controversial statement; authors from the objective Bayesian school (e.g., Jaynes 2003), including some phylogeneticists (Ronquist 2004), do not consider improper priors inherently problematic. There are, after all, improper priors that allow for proper posteriors, and have a clear meaning as the limit of a sequence of proper priors.

(45) 5.2.1. (2nd pragraph) – I'm not sure I understand the authors' point when they say the behavior of a model can only be explored once a new model has been developed. This could use some elaboration, or rephrasing.

(46) 5.2.1 (3rd paragraph) – "The key is to remember that the prior distribution of a parameter represents the probability of those values being proposed during the MCMC" – No, the proposal

distribution is generally not identical to the prior distribution. It can be (in which case the acceptance ratio reduces to the likelihood ratio), but that's a rather crude way of setting up the MCMC that wouldn't get us very far in phylogenetics. This is yet another worrying instance of the paper seemingly not having a firm grasp of the topic it tries to explain.

(47) 5.2.1 (penultimate paragraph) – "The effective prior on a node age in an inference will be the product of the prior set by the tree model, and of all additional calibration times set for the tree" – This is true of BEAST and MrBayes, but not of other Bayesian phylogenetic software such as MCMCTree or RevBayes. Which is a good thing, because this "multiplicative construction" (Heled & Drummond 2012) violates basic rules of probability theory.

(48) 5.2.1 (last paragraph) – "If different priors yield clearly different results, further investigation is needed to determine what conclusions, if any, can be drawn from the inference. Possible next steps could include running the analysis under both sets of priors to test whether either setup is strongly biasing the result" – So when we find that running the analysis under two sets of priors gives us different results, we should… run the analysis under two sets of priors? Again? Surely this doesn't make sense.

(49) 5.2.2 (1st paragraph) – "considered as part of the prior" – The word "as" is unnecessary.

(50) 5.2.2 (1st paragraph) – The passage about using the PBDB to place priors on FBD parameters could definitely be expanded on, especially in light of the previously mentioned caveat that the origination and extinction rates inferred by traditional paleobiology don't coincide with the birth and death rates of the FBD model.

(51) 5.2.2 (2nd paragraph) – "thereby making convergence around an optimal solution impossible" – I think that's the wrong way of putting it. The posterior would be bimodal in such a scenario, but that doesn't in itself mean that the MCMC will have trouble converging on it. For an enlightening exploration of MCMC performance on tree mixtures, see Ronquist et al. (2006).

(52) 5.2.2 (3rd paragraph) – "A good way to circumvent this would be to try and simulate datasets from scratch, based on more or less comparable parameters to the empirical data" – I have no idea what this means. Parameters can't be "compared" to the data. If the alpha parameter of my gamma ASRV model is 0.234 and my sequences look like this: ACCGCGT…, ACTGGGT…, what sort of comparison am I supposed to make, and how did I get that alpha in the first place?

(53) 5.3 – Surely the number of characters needed is a function of tree size? 300 characters will be overkill for 10 tips, but insufficient for 250 tips.

(54) 5.3 (2nd paragraph) – Is FBD with diversified extant sampling available in any package other than MrBayes? If not, this might be worth mentioning.

(55) 5.3 (3rd paragraph) – "assigning all fossils to the root" – To me, that would be an extremely narrow constraint, not a broad one! I guess the authors mean allowing fossils to attach anywhere in the tree (i.e., anywhere within the clade originating with the root node), rather than just along the root branch, but in that case, the description should be rephrased.

(56) 5.3 (4th paragraph) – The piecewise-constant ("skyline") fossil sampling rates described in the

studies cited by the authors are actually not a great way of accommodating what we see in the fossil record, which often consists of a few Lagerstätten punctuated by long periods with almost no samples. The sampling burst model proposed by Magee & Höhna (2021) and implemented in RevBayes is much more promising, though underexplored and underutilized.

(57) 5.4 – I would speak of searching for “high-probability” rather than “optimal” values, to make it clear that the MCMC is not an optimization algorithm.

(58) 5.4 (4th paragraph) – The recommendation of jModelTest feels like a blast from the past given the fact that the substitution model is rarely the biggest contributor to overall model complexity (especially in the case of hierarchical models like FBD analyses with relaxed clocks), and in light of recent studies such as Abadi et al. (2019) and Fabreti & Höhna (2023) which show that the choice of the substitution model often doesn’t matter very much.

Section 6

(59) Please note that the RevBayes user forum has moved from https://groups.google.com/g/revbayes-users to https://github.com/revbayes/revbayes/discussions.

References

(60) The title of Kolmogorov’s 1933 paper is misspelled; it should end with “di distribuzione” (two separate words).

References cited in this review but not in the paper

References

1. Abadi S, Azouri D, Pupko T, Mayrose I: Model selection may not be a mandatory step for phylogeny reconstruction.*Nat Commun*. 2019; **10** (1): 934 PubMed Abstract | Publisher Full Text
2. Claramunt S: CladeDate : Calibration information generator for divergence time estimation. *Methods in Ecology and Evolution*. 2022; **13** (11): 2331-2338 Publisher Full Text
3. Guimarães Fabreti L, Höhna S: Nucleotide Substitution Model Selection Is Not Necessary for Bayesian Inference of Phylogeny With Well-Behaved Priors.*Syst Biol*. 2023; **72** (6): 1418-1432 PubMed Abstract | Publisher Full Text
4. Friedman M, Keck BP, Dornburg A, Eytan RI, et al.: Molecular and fossil evidence place the origin of cichlid fishes long after Gondwanan rifting.*Proc Biol Sci*. 2013; **280** (1770): 20131733 PubMed Abstract | Publisher Full Text
5. Heled J, Drummond AJ: Calibrated tree priors for relaxed phylogenetics and divergence time estimation.*Syst Biol*. 2012; **61** (1): 138-49 PubMed Abstract | Publisher Full Text
6. Höhna S, Landis MJ, Heath TA: Phylogenetic Inference Using RevBayes.*Curr Protoc Bioinformatics*. 2017; **57**: 6.16.1-6.16.34 PubMed Abstract | Publisher Full Text
7. Lanfear R, Hua X, Warren DL: Estimating the Effective Sample Size of Tree Topologies from Bayesian Phylogenetic Analyses.*Genome Biol Evol*. 2016; **8** (8): 2319-32 PubMed Abstract | Publisher Full Text
8. Magee A, Höhna S: Impact of K-Pg Mass Extinction Event on Crocodylomorpha Inferred from Phylogeny of Extinct and Extant Taxa. *bioRxiv*. 2021. Publisher Full Text
9. Marshall CR: Five palaeobiological laws needed to understand the evolution of the living biota.

Nat Ecol Evol. 2017; **1** (6): 165 PubMed Abstract | Publisher Full Text
10. Matschiner M, Musilová Z, Barth JM, Starostová Z, et al.: Bayesian Phylogenetic Estimation of Clade Ages Supports Trans-Atlantic Dispersal of Cichlid Fishes.*Syst Biol*. 2017; **66** (1): 3-22 PubMed Abstract | Publisher Full Text
11. Ronquist F: A Broad Look at Tree-Building. *Science*. 2004; **303** (5659): 767-768 Publisher Full Text
12. Ronquist F, Larget B, Huelsenbeck JP, Kadane JB, et al.: Comment on "Phylogenetic MCMC algorithms are misleading on mixtures of trees".*Science*. 2006; **312** (5772): 367; author reply 367 PubMed Abstract | Publisher Full Text
13. Wilkinson RD, Steiper ME, Soligo C, Martin RD, et al.: Dating primate divergences through an integrated analysis of palaeontological and molecular data.*Syst Biol*. 2011; **60** (1): 16-31 PubMed Abstract | Publisher Full Text

Is the rationale for developing the new method (or application) clearly explained?
Yes

Is the description of the method technically sound?
Partly

Are sufficient details provided to allow replication of the method development and its use by others?
Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?
No source data required

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Phylogenetic method development; Bayesian phylogenetic inference; macroevolution; computational paleobiology

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 09 Dec 2025

Joëlle Barido-Sottani

Thank you very much for the review and the suggestions! Please find our detailed response to your comments below (in bold).
Comment: Overall, the manuscript retreads a lot of familiar ground (Fig. 4, showing traces produced by MCMC runs with moves that are too small or too large, feels like it has been

recreated in the literature dozens of times now), but there is enough useful detail (for example, about differences between major software packages) to make it worth publishing. I have a lot of minor comments that are listed below, and some of these exemplify what I would consider to be two larger problems. First, the manuscript is not structured as well as it could be, which negatively affects the flow of the text. A lot of this has to do with jumping between different levels of abstraction. I would recommend explaining models first (highest level), then Bayesian inference of model parameters and how it differs from maximum likelihood, followed by concerns that are specific to phylogenetics (e.g., differences between dated and undated trees). Only then would I dive into an explanation of the phylogenetic MCMC, and give tips and recommendations about its use. As it stands, the authors start describing the MCMC before they even explain the meaning of the term "likelihood", which just doesn't work very well. There are several passages interspersed throughout the text that anticipate later sections of the manuscript, but look out of place when the reader first encounters them.
Response: We are happy to hear that our work provides worthwhile information and appreciate the thoughtful suggestion for what might be an effective alternative structure for this paper, which would probably resonate well with some readers. However, after some consideration, we have decided against a major restructuring, for a couple of reasons. For one, while your suggested alternative structure makes sense conceptually and logically, we believe the current structure is also reasonable and approachable. While there are, as you noted, various instances where the text mentions topics that are only detailed further later on, that primarily happens during the first section, where it serves the purpose to give an overview of the general problem and the relevant concepts - which is something we expect at least some readers to find helpful. And while it is true that we end up jumping between the levels of abstraction, we think that reflects the route one takes when setting up or troubleshooting such models - a process which is usually not following a linear path up or down the levels of abstraction either. It is also not necessarily the intention of a paper like this to only be read from start to end, and so we expect readers to navigate back and forth between the sections relevant to their interest in any case. We did however add some sentences that make more clear to the reader that the first section provides more of an overview over the central issues, while the different elements of it are treated in more detail later on (in addition to the existing references to other relevant sections).

Comment: My second concern has to do with several problematic instances where the authors seemingly misunderstood the very concepts they are trying to explain – these are likely due to pedagogical concerns about keeping things simple, or just sloppy use of terminology. For example, the authors suggest in several places that the MCMC optimizes parameter values or even searches for the true values (it does no such thing), they confuse the acceptance ratio with the Hastings ratio (while leaving the actual Hastings ratio unmentioned and unexplained), and the prior distribution with the proposal distribution. I have absolutely no doubt about the authors' expertise, but such relatively basic errors are worrisome, and a manuscript that contains them cannot do the topic justice. Conditional on these errors being fixed, however, I can see the manuscript being useful to many biologists who want to learn more about the topic but are unsure where to start.

Response: We did indeed try and keep the language simple and approachable, and it seems this made us fall into the trap of employing commonly used, but technically incorrect phrasing (e.g., referring to parameters being optimised by the MCMC). Using these is probably hindering proper understanding of the topic in any case, but we agree that it is even more crucial to avoid this in an article like this. We replaced such instances with more correct wording and added a note on the misconception that the MCMC optimises parameters earlier in the text.

Comment: Throughout the paper:
(1) Inconsistent American vs. British spellings (behavior vs. behaviour, optimized vs. optimised).
Response: Fixed.

Abstract
(2) I would take issue with the statement that "[m]odern statistical methods" employ models "involving fossil record information". Sure, some do, but most active areas of phylogenetic research have nothing to do with analyzing fossil data. I understand the FBD model is of particular interest to the authors, but here they are perhaps getting a little ahead of themselves.
Response: We have edited the abstract to clarify that this is just one example of phylogenetic applications.

(3) It seems backwards to mention the MCMC and posterior distributions before mentioning Bayesian analysis.
Response: We have restructured the abstract to introduce concepts in a more natural order.

Section 1
(4) The first example of a model chosen by the authors is perhaps a bit unfortunate: the typical parameterization involves a transversion/transition *ratio*, rather than a separate parameter for transitions and another one for transversions.
Response: This was meant as a hypothetical, but yes, due to the resemblance to the actual K2P model, we've changed this.

(5) The authors again seem to be getting ahead of themselves when they describe model parameters as random variables. That's indeed how they are treated in Bayesian inference, but in maximum-likelihood estimation (and frequentist statistics more generally), every parameter is assumed to have a fixed but unknown value. Since the differences between these two frameworks have not yet been discussed at this point, the authors need to be more careful about not attaching interpretations to model parameters that are only valid under one of them.
Response: We've edited this to make clear we're discussing Bayesian models.

(6) "Bayesian inference thus offers a more complete picture of the results, integrating uncertainty in the results as well as existing information from previous studies" – Though I use and have developed Bayesian phylogenetic methods myself, I find this statement to be overselling things in an unacceptable way. First of all, most of the time, we as Bayesian

phylogeneticists are desperate to find “flat” or “uninformative” priors (this includes seemingly mechanistic priors like the birth-death process, whose big attraction is actually the fact that it is flat over labeled histories), so the need to incorporate prior information is more of a burden than an advantage. Second, even when we do pick informative priors (and often it’s impossible not to), they are almost never based on previous studies. Third, the statement completely ignores the old but important objection that priors are subjective.
Response: We think this point represents a difference in schools of thought between ourselves and the reviewer. We no doubt agree that the use of priors as we’ve outlined is subjective: at no point do we argue for empirical Bayes or similar. However, we strongly feel that the ability to quantify our subjective prior intuition as a part of the model does indeed present a more complete picture of of the results. In no science is there true objectivity. Down to the process of hypothesis generation and data collection, we are codifying our biases into the research, but are doing so in such a way as to be unquantifiable. As such, uninformative priors, to the extent that they are possible (often it is not, see Zwickl and Holder 2004), still represent a statement that is in its core subjective.

(7) “true distribution of plausible parameters” – The authors mean “posterior”, not “true”. I’m not sure the notion of a “true distribution” even makes sense to me, since at the end of the day, I do believe there was just one true evolutionary history that generated my data, not a distribution of them (so I’m perhaps a likelihoodist at heart after all).
Response: You’re right. Changed.

Section 2.1
(8) Once again, I have to point out that introducing the MCMC before Bayesian analysis seems to me like putting the cart before the horse. It’s really confusing that “likelihood” is only explained several paragraphs after bringing up “maximum likelihood”! I would strongly recommend rearranging this part of the MS.
Response: While we can appreciate this point in general, we disagree with it in this particular instance. In a class-setting, one would present Bayesian analysis first and then introduce MCMC as an associated tool. But the MCMC and its workings are the main topic of this article, hence we start with it and then back up to cover basics of the Bayesian analyses it is used for. Furthermore, the purpose of this paper is primarily to advise on diagnosing and troubleshooting MCMC problems, rather than introducing the approach from scratch (we recapitulate them as a refresher and to ensure the readers and us are on the same page). So arguably, our intended audience, while novices, are not entirely new to those concepts. We are thus confident that the confusion caused by this should be minimal at best.

(9) 2.1.1 – Mentioning clock models and rates only makes sense in the context of inferring time trees; what about time-free / undated trees, where rates and times are confounded as branch lengths? This is yet another instance of concepts not being explained at a high enough level of abstraction when they are first introduced.
Response: We now mention that clock models are only necessary when estimating a time tree. However, Bayesian inference is usually applied to time trees, so we feel like we do not lose in abstraction by introducing the clock model early.

(10) Also 2.1.1 – It may be worth pointing out that the single parameter of JC69 and Mk is not free, so there's actually nothing for us to estimate in these models.
Response: We now mention that substitution models are normalized in a substitution-clock inference, and thus the JC69 and Mk parameters are fixed.

(11) 2.1.2 – "a rate that is very high is believed to be less likely than one that is very short" – This sentence doesn't work; rates can't be "short". **Fixed.**

(12) Also 2.1.2 – I expect many readers to not know what it means for parameters to not be identifiable.
Good catch. Added a short explanation of this term.

(13) 2.1.5 – The nomenclature I'm familiar with is that the acceptance ratio is the product of the posterior odds ratio and the Hastings ratio, so I find the authors' attempt to conflate these two confusing. The actual Hastings ratio – i.e., p(theta — theta') / p(theta' — theta) – doesn't make an appearance in the authors' explanation, so what they are describing is actually the Metropolis algorithm rather than the more general Metropolis-Hastings algorithm.
Thank you for pointing this out! We have fixed this (see also our response to a similar comment from reviewer 5).

(14) 2.1.6 – The word "subclade" seems unnecessary here. I would just use "clade". **Fixed.**

Section 2.2
(15) Subsection 2.2.1 seems both misplaced and misnamed, since the authors focus on the difference between dated and undated (rather than rooted and unrooted) trees.
We agree that beyond the initial mention we don't cover unrooted trees much and focus on dated vs not - we thus changed the subsection's title to 'Types of trees', which covers the whole content of the subsection. However, we disagree that the subsection is misplaced - while it admittedly isn't strictly speaking covering the implementation of the MCMC, it is still crucial background to cover ahead of the following sections.

(16) Also 2.2.1 – In addition to fossils, it may be good to mention the other common type of time-stamped data: serially sampled viral sequences. **Good point, we added those.**

(17) 2.2.2 – This description is accurate but somewhat confusing, since it seems to conflate two kinds of modularity. One is a property of the models used, the other is a feature of software design.
We have clarified here that we are discussing the implementation of the frameworks, but that this implementation is made easier by the modularity of the models implemented.

(18) 2.2.3 – While this is the RevBayes default and as such perhaps the only thing the reader needs to know about the program, RevBayes *can* be made to count iterations the MrBayes/BEAST way (Höhna et al. 2017).
Thanks for reminding us of this option. We have added a note about it, as some readers might find it useful.

(19) Also 2.2.3 – It may be useful to point out that some special moves such as the slice move (Neal 2003) are supposed to have very high acceptance ratios.
Indeed, some moves are based on a Gibbs sampler which has an acceptance ratio of 1. This has been added in the paragraph.

Section 3
(20) "This can raise serious performance issues for our MCMC inference, either when exploring the tree space or when calculating the posterior probabilities" – Those are not two separate issues, but two ways of describing the same issue. If you are not doing a good job of exploring the tree space, you are not accurately approximating the posterior probabilities. **Fixed.**

(21) 3.1 – The second paragraph tries to reintroduce the topic of tree moves from scratch, despite the fact that they were already described (even including specific examples) several paragraphs above. Here, I would maybe recommend removing the earlier mention and merging it with this latter one to better preserve the flow of the exposition.
We added a little more bridge between these two sections. In the first, we explain how these moved work. But the old moves aren't sufficient for new tree types, such as networks, for example. This remains an ongoing problem.

(22) 3.2 – "the parameters of the gamma distribution representing among-site rate heterogeneity" – Usually, this distribution only has one free parameter, since its mean needs to equal 1. **Changed to singular.**

(23) 3.3 – The point about tip calibrations being less subjective than node calibrations is an interesting one that made me think. Is this actually the case, or are we just more comfortable placing uniform priors on tip ages compared to node ages? Node calibrations inherently involve more sources of uncertainty, but uncertainty and subjectivity are two different things, and many attempts have been made to design node calibrations in a more objective way (Wilkinson et al. 2011; Friedman et al. 2013; Matschiner et al. 2017; Claramunt 2022).
Good point. We have clarified that the uncertainty we are referring to here is mainly about the placement of fossil taxa in the tree and added a note with the recommended citations.

(24) Also 3.3 – "additional information could be used to construct more informative non-uniform priors on fossil ages" – Yes, one could argue a non-uniform distribution would often be a better representation of our prior knowledge. What if a fossil comes from a Maastrichtian formation, but specifically from the base of the middle member of that formation? What if it is sandwiched between two radiometrically dated horizons, but is closer to one than the other? Etc. But that leads to two problems: (1) the subjectivity is back, big time; and (2) some implementations are not flexible enough to handle non-uniform tip age priors. MrBayes is, and BEAST 2 might be too, but while I don't consider myself a RevBayes beginner, I'm not sure how I'd torture the program to allow them. Here, it might be worth distinguishing between what's possible in principle and what's possible in practice.

We have added a note with these caveats.
Section 4

(25) What I miss from the discussion of effective sample sizes is a clear statement that they only quantify the precision of the inference. They tell you whether you have enough samples to achieve a low standard error of the mean, but as the authors themselves point out, in and of themselves they say nothing about whether these samples come from the posterior distribution. With multiple chains, one would definitely hope that the only way to achieve high ESS values is for these chains to converge on the posterior, but even in that scenario there are certain caveats (e.g., it's not impossible for multiple chains to get stuck in the same region of the parameter space if the posterior is very rugged and their starting values are similar or even the same, as can happen when the user supplies their own initial tree).
Pointing this out more explicitly would indeed be an improvement. We have done that in the paragraph where we already pointed out that issue and reorganised it accordingly.

(26) 4.1 – This is very pedantic, but the qualifier "typically" in the 3rd sentence of the 1st paragraph is unnecessary. In the context of MCMC, convergence is short for "convergence to the stationary distribution"; there is no other kind. **Fixed.**

(27) 4.1 (1st paragraph) – I would take the opposite view with regard to the relative dangers of setting the burnin too low/high. If you set it too low, you can always discard extra samples after the fact, but if you set it too high, you've wasted time and CPU cycles and there is nothing you can do.
Agreed. It depends on whether we are only concerned with the effect on the posterior vs. the time-efficiency of our analyses too (and whether the burnin is auto-discarded or not). We have updated the paragraph to reflect these points too.

(28) 4.1 (3rd paragraph) – Is N the number of generations, or the number of moves? This might be a distinction worth drawing given the previous caveat about how different programs count generations in different ways.
We have clarified that the ESS is calculated based on the number of posterior samples recorded and so does not depend on the implementation chosen for the generations.

(29) 4.1 (4th paragraph) – "the individual ESS of each optimized parameter" – Wrong nomenclature. MCMC doesn't optimize; it samples. **Changed to "estimated" parameter.**

(30) 4.1 (4th paragraph) – The point about branch rate ESS values being meaningless is an important one, but I would perhaps phrase it in a different way. It's not that these are parameters that are unlikely to converge, it's that our diagnostic tools conflate different parameters by giving them the same name (e.g., branch rates[63] on one topology is a different parameter from branch rates[63] on another topology). In principle, one could develop tools that do not do this. We could instead extract all unique branches present in the MCMC sample, and compute the ESS of each one's rate.
Thank you for this comment. We have clarified that this is not an issue with the parameters themselves, but rather how we represent them in the trace.

(31) 4.1 (5th paragraph) – “are guaranteed to converge” – I would qualify this statement with “when run for long enough”. **Fixed.**

(32) 4.1 (5th paragraph) – “in a local optima” – The word “optima” is plural. The corresponding singular is “optimum”. **Fixed.**

(33) Caption to Figure 4 – “is doing baby steps” – I’d prefer a more formal phrasing.
This has been changed accordingly.

(34) 4.2.2 – I would change the punctuation from “[. . .] of the phylogeny, however it is preferable [. . .]” to “[. . .] of the phylogeny; however, it is preferable [. . .]”. **Fixed.**

(35) Also 4.2.2 – I would also cite Lanfear et al. (2016) here, who first showed how to approximate topological ESS (with RWTY providing an implementation of their algorithms).
Added.

(36) 4.2.3 – “This is likely due to one or both chains not having converged yet” – The qualifier “likely” is not needed. The stationary distribution of a correctly constructed MCMC sampler is the posterior distribution of the model, so if it is sampling from a distribution other than the posterior, it is by definition because it has not yet converged to its stationary distribution. The scenario where different chains get stuck at different local optima is not a different type of failure (as implied by the word “also”); rather, it is one possible reason why the chains have not converged to their stationary distribution.
We have edited this section as suggested.
Section 5

(37) 5.1.1 – “Each parameter that is supposed to be estimated by the analysis needs to have at least one operator associated to it, in order to be optimised” – Again, the MCMC doesn’t optimize. (It is closely related to an optimization technique called simulated annealing, but that’s beside the point here.) I’m surprised the authors would make such a basic mistake after explaining the difference between maximum likelihood and Bayesian inference.
Good point, edited to sampled.

(38) Also 5.1.1 – “so that more moves are being proposed each generation for them” – Putting it this way presupposes the default RevBayes setup with multiple moves per generation; it doesn’t accurately describe what increasing the weight does in MrBayes or BEAST.
Changed wording to apply for either case.

(39) 5.1.2 (2nd paragraph) – I can think of absolutely no biological justification for setting the extinction rate at 1/10 of the speciation rate. In fact, paleobiology tells us that on average, speciation and extinction rates tend to be very nearly equal (Marshall 2017). This might not make for the best starting values, but in that case, let’s not pretend our starting values are motivated biologically rather than by computational convenience.
Agreed, the ’biological meaning’ argument was referring to λ, while μ = λ/10 is merely reporting what is common practice. We have clarified that accordingly now.

(40) 5.1.2 (3rd paragraph) – "starting values do not have to be spot-on estimates of where the actual true values lie, because after all, the MCMC is expected to go find those" – No, the MCMC is not expected to find the true values. This is a fundamental misunderstanding of what it does, and one that I find quite disturbing in a paper that takes it upon itself to introduce the technique to others. **Fixed.**

(41) 5.1.2 (3rd paragraph) – "instead allow it to start exploring feasible solutions instead" – There is one more "instead" in this sentence than needed. **Fixed.**

(42) 5.1.3 – "each run could potentially be sampling different local optima" – Here, I'd use the singular: "each run could potentially be sampling a different local optimum". **Fixed.**

(43) 5.2 – "the probability of the tree under the FBD process given the fossil ages is considered part of the prior by BEAST2 and RevBayes" – That's correct for BEAST 2 but not for RevBayes. In the latter, if you run your analysis with the underPrior=TRUE option (available up to v1.2.4) or with the .ignoreAllData() option (which supplanted the former in v1.2.5), fossil ages will be discarded along with the character data. You never want this in an FBD analysis, since you're not learning anything useful by pretending you don't know if there were bears in the Jurassic or trilobites in the Miocene.
We have removed the specific mention of the software as this is implementation-specific, and as the reviewer notes, can change with different versions. Instead we now explain this issue in general terms.

(44) 5.2.1 (2nd paragraph) – "Overall, improper priors are statistically problematic and should be avoided if possible" – This is a bit of a controversial statement; authors from the objective Bayesian school (e.g., Jaynes 2003), including some phylogeneticists (Ronquist 2004), do not consider improper priors inherently problematic. There are, after all, improper priors that allow for proper posteriors, and have a clear meaning as the limit of a sequence of proper priors.
We have now added more information on improper priors and their potential issues. As we mentioned earlier, we take position here for subjective priors, which should be as informative as possible, whereas the goal of improper priors in objective Bayesian analysis is to be as uninformative as possible.

(45) 5.2.1. (2nd paragraph) – I'm not sure I understand the authors' point when they say the behavior of a model can only be explored once a new model has been developed. This could use some elaboration, or rephrasing.
The point of this section was that models are frequently used for applications or in scenarios beyond what the original developers anticipated, which is another reason relying on default setups can be an issue. We have edited this part to make our reasoning clearer.

(46) 5.2.1 (3rd paragraph) – "The key is to remember that the prior distribution of a parameter represents the probability of those values being proposed during the MCMC" – No, the proposal distribution is generally not identical to the prior distribution. It can be (in which case the acceptance ratio reduces to the likelihood ratio), but that's a rather crude

way of setting up the MCMC that wouldn't get us very far in phylogenetics. This is yet another worrying instance of the paper seemingly not having a firm grasp of the topic it tries to explain.
You are right, we apologize for this mistake and have removed it from the article. We now provide more information on the proposal distribution in the Move section.

(47) 5.2.1 (penultimate paragraph) – "The effective prior on a node age in an inference will be the product of the prior set by the tree model, and of all additional calibration times set for the tree" – This is true of BEAST and MrBayes, but not of other Bayesian phylogenetic software such as MCMCTree or RevBayes. Which is a good thing, because this "multiplicative construction" (Heled & Drummond 2012) violates basic rules of probability theory. The main point we wanted to make in this paragraph is that users can and should visually inspect the effect priors, which can diverge from the user specified priors even without the multiplicative construction, i.e., see Warnock et al. 2015, PRSB.
We have slightly modified the sentence to avoid implying that all software uses the multiplicative approach and added further citations to this paragraph for any readers interested in exploring the topic further. We wished to avoid mentioning specific software implementations here, as the options available across different software are quite complicated and subject to change. (48) 5.2.1 (last paragraph) – "If different priors yield clearly different results, further investigation is needed to determine what conclusions, if any, can be drawn from the inference. Possible next steps could include running the analysis under both sets of priors to test whether either setup is strongly biasing the result" – So when we find that running the analysis under two sets of priors gives us different results, we should. . . run the analysis under two sets of priors? Again? Surely this doesn't make sense.
In this case, we're specifically addressing the circumstance under which you've run the analysis with a few different configurations of priors, and are getting conflicting results. In that case, it's worth looking at what the priors tell you without the data involved. We have edited to make it clearer that we're referring to running under the prior here.

(49) 5.2.2 (1st paragraph) – "considered as part of the prior" – The word "as" is unnecessary.
Fixed.

(50) 5.2.2 (1st paragraph) – The passage about using the PBDB to place priors on FBD parameters could definitely be expanded on, especially in light of the previously mentioned caveat that the origination and extinction rates inferred by traditional paleobiology don't coincide with the birth and death rates of the FBD model.
We have added a note about differences in rates parameters between fossils versus trees but feel that a much longer discussion about this is beyond the scope of this paper.

(51) 5.2.2 (2nd paragraph) – "thereby making convergence around an optimal solution impossible" – I think that's the wrong way of putting it. The posterior would be bimodal in such a scenario, but that doesn't in itself mean that the MCMC will have trouble converging on it. For an enlightening exploration of MCMC performance on tree mixtures, see Ronquist et al. (2006).

While the MCMC can in theory recover bimodal and multimodal posterior distributions, in our experience this has often not been the case, unless mixture models are explicitly accounted for in the analysis. We have edited this section to make it clear that convergence is not impossible, but is likely to encounter issues. (52) 5.2.2 (3rd paragraph) – "A good way to circumvent this would be to try and simulate datasets from scratch, based on more or less comparable parameters to the empirical data" – I have no idea what this means. Parameters can't be "compared" to the data. If the alpha parameter of my gamma ASRV model is 0.234 and my sequences look like this: ACCGCGT.. . , ACTGGGT.. . , what sort of comparison am I supposed to make, and how did I get that alpha in the first place?
We agree that the original wording may be a bit confusing to some, so we edited that paragraph for clarity, leading with the mention of ABC approaches instead.

(53) 5.3 – Surely the number of characters needed is a function of tree size? 300 characters will be overkill for 10 tips, but insufficient for 250 tips.
We have added that the referenced study used trees of 200 tips and now mention that this number is likely to vary based on tree size and dataset uncertainty.

(54) 5.3 (2nd paragraph) – Is FBD with diversified extant sampling available in any package other than MrBayes? If not, this might be worth mentioning.
The availability of diversified sampling has been added to the end of the paragraph.

(55) 5.3 (3rd paragraph) – "assigning all fossils to the root" – To me, that would be an extremely narrow constraint, not a broad one! I guess the authors mean allowing fossils to attach anywhere in the tree (i.e., anywhere within the clade originating with the root node), rather than just along the root branch, but in that case, the description should be rephrased. **Fixed.**

(56) 5.3 (4th paragraph) – The piecewise-constant ("skyline") fossil sampling rates described in the studies cited by the authors are actually not a great way of accommodating what we see in the fossil record, which often consists of a few Lagerst¨atten punctuated by long periods with almost no samples. The sampling burst model proposed by Magee & Höhna (2021) and implemented in RevBayes is much more promising, though underexplored and underutilized.
This would depend on the dataset - the importance of Lagerstätten varies. It might be important for groups such as insects or some vertebrates, but less important for some well preserved marine invertebrates like bivalves or brachiopods. Therefore, we do not wish to say that the skyline model is not useful. However, we have added the recommended citation and added a description of this model.

(57) 5.4 – I would speak of searching for "high-probability" rather than "optimal" values, to make it clear that the MCMC is not an optimization algorithm. **Changed.**

(58) 5.4 (4th paragraph) – The recommendation of jModelTest feels like a blast from the past given the fact that the substitution model is rarely the biggest contributor to overall model complexity (especially in the case of hierarchical models like FBD analyses with relaxed clocks), and in light of recent studies such as Abadi et al. (2019) and Fabreti & Höhna (2023)

which show that the choice of the substitution model often doesn't matter very much.
Thank you for bringing up those studies. Indeed, testing for substitution models was merely meant to be one straightforward example for the idea of running preliminary tests on model components (as it used to be common practice), rather than being an explicit recommendation that this should be the most important test to run (as will surely become clear to readers when we highlight partitioning and adequacy testing in the following paragraphs). But we agree that the relative (un)importance of substitution models in phylogenetic inference should be clarified, so we now do so and cite those two papers. However, we still chose to keep the mention of jModelTest, so the paper can act as a guide to tools and resources which are commonly used in the community, even if some inferences may not benefit from its use.
Section 6

(59) Please note that the RevBayes user forum has moved from https://groups.google.com/g/revbayes-users to https://github.com/revbayes/revbayes/discussions. **Fixed.**
References

(60) The title of Kolmogorov's 1933 paper is misspelled; it should end with "di distribuzione" (two separate words). **Fixed.**

Competing Interests: No competing interests were disclosed.

Reviewer Report 17 December 2024

https://doi.org/10.21956/openreseurope.19757.r46198

© **2024 Beerli P.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Peter Beerli
[1] Florida State University, Tallahassee, Florida, USA
[2] Florida State University, Tallahassee, Florida, USA

Review of Practical guidelines for Bayesian phylogenetic Inference using Markov chain Monte Carlo (MCMC) [version 3]

I enjoyed reading the 'Practical Guidelines'. It offers a good overview of the field of Bayesian inference in phylogenetics and gives good advice on potential problems with complex analyses, highlighting possible ways to find good solutions. I have a few suggestions for improving clarity in some sections. I have added several citations that may better trace the history of the method or give credit to researchers who did similar/same analyses earlier. My biggest concern was Figure 1 and the section describing the MCMC decision ratio, the Metropolis-Hastings ratio. I hope the citations and guidance will help improve this section. All my comments are suggestions, and you

should decide whether to incorporate them or not.

Detailed comments with page and line references:
Abstract:
3rd line: 'solutions' seems vague
line 5: 'evaluated' I would prefer 'estimated' or 'found'
line 7 from bottom: I would prefer direct wording: We offer an overview...

Page 4: Introduction: for the 34 million trees you cite Degnan and Salter 2005,
which is a fine citation, but there were earlier citations, for example, Felsenstein 1978 (there may be even earlier ones, but I do not know them; the groundwork for counting trees was given by A. Cayley in 1889, but he does not elaborate on specific numbers.

Felsenstein, J. (1978) [Ref-1].
Cayley, A. (1889) [Ref-2]

the next sentence: 'It is intractable to evaluate all of them.' This statement is incorrect, evaluating 34 million items is certainly in the realm of current computers (for example, Swoffords's PAUP* can evaluate and enumerate all 10-tip (perhaps up to 12 tips) trees in a parsimony context. It is certainly tedious. How about: 'Evaluating all phylogenies with a number of samples larger than 10 become quickly intractable.'

"That technique is Markov chain Monte Carlo (MCMC)," the next sentences show potential alternatives with citations, it seems odd that you do not cite MCMC here.
Some reviewers suggested to show that alternatives exist. I find showing these "alternatives" that are 'variants' of MCMC in the introduction unnecessary but
could be discussed later towards the end of the tutorial. I suggest to cite Metropolis et al. here. Concerning the citation of the alternative methods given that this is early in the introduction, readers may want to have a citation of the method and not the application in phylogenies, for example, SMC: Doucet et al. 2001 [or perhaps Gordon 1998 associated with introduction of SMC), HMC: Duane et al. [I would remove the other methods and move them somewhere else but still cite Metropolis et al.]

Arnaud Doucet, Nando de Freitas, and Neil Gordon. (2001) [Ref 12]

****Figure 1 and text discussing figure 1: MCMC was introduced by Metropolis et al. and the outline in Figure 2 follows the 'Metropolis algorithm' and should be
labeled the 'Metropolis ratio' but see below: Hastings (1970) generalized the Metropolis algorithm because the original ratio assumes that the proposal distributions for the newly proposed parameter and the old parameters are symmetric, i.e., $g(x|y) = g(y|x)$. In phylogenetics, we do not necessarily guarantee this condition. Modern implementations use the Metropolis-Hastings algorithm to allow for potential asymmetries in the proposal. Figure 1 is labeled 'Hastings ratio,' and a reader needs to know why this is called 'Hastings ratio' because Hastings is not mentioned in the text. I suggest revising the text to include Hastings (1970) and relabeling the 'Hastings ratio' to 'Metropolis-Hastings ratio.'

I also suggest adding context after the Metropolis et al. citation on page 5, along these thoughts: "The original algorithm assumed that the moves in the parameter space are symmetric, going from parameter value a to parameter value b has the same probability as going from b to a. Hastings (1970) introduced an additional term to
balance potentially asymmetric proposals; this led to a general version: the Metropolis-Hastings algorithm. I would remove the sentence, "In this case, all the values ... phylogenetics. And then, in figure 1, use 'Metropolis-Hastings ratio': with minimal change to the figure to resolve the issue, one could simply add in the figure legend: "We left out the Hastings part of the ratio for clarity; it would add g(old|new)/g(new|old) to R."

Citation: W. K. Hastings, 1970. [Ref-4])

page 6 last 3 lines of 2.1.2 use high vs low (not high vs short)

page 6 first column, bottom and second column top: I suggest to cite the authors who defined the methods: path sampling in phylogenetics (Lartillot and Phillipe 2006), Stepping Stone (Xie et al. 2010), nested sampling (Russel et al. 2018 [applies Skilling 2006 to phylogenetics] and very recently, a Partition-Weighted Kernel method (Wang et al., 2023).

Nicolas Lartillot and Herve, 2006 [Ref-5]

Wangang Xie, 2011. [Ref-6].

Yu-Bo Wang, et al., 2023. [Ref-7].

page 6 second column: For instance, reversible jump can be used to infer the most adequate substitution model for an alignment (Huelsenbeck et al. 2004) [they estimate
the best substitution model among 200+ models 13 years before Bouchaert and Drummond]

John P. Huelsenbeck, et al., 2004. [Ref-8].

page 9, second column: Algorithms for efficiently sampling phylogenetic trees ... (citations, see below) this would be a good location to show the early Bayesian phylogenetic methods because they all had to solve the tree-changing problem and this will give a good way to give credit to early work on Bayesian phylogenetics, for example, Li often gets forgotten.

Shuying Li. Phylogenetic tree construction using Markov chain Monte Carlo. Journal of Computational Biology, 3(3):463–474, 1996. doi: 10.1089/cmb.1996.3.463.[Ref - 13]

Bob Mau and M. A. Newton., 1997.[Ref-9].

Bruce Rannala and Ziheng Yang, 1996. [Ref 10].

Bret Larget and Donald L. Simon, 1999. [Ref- 11].

page 11 "(e.g., in tracer)" You have a section later on Tracer; I would simply remove this () because, at this point, the reader does not know that you mean the tracer application.

page 16 left bottom: The sentence may confuse the reader by using 'temperature' from 1 to 0. Usually, we would consider the Metropolis-Hastings ratio powered with 1/T, where we would set the cold chain at temperature T=1 and also consider hotter chains with higher T, for example, a T=1,000,000 would set that exponent to approximately zero so that all moves will be accepted; any intermediate value will accept more often than the cold chain.

page 16 right 5.2. choice of models and priors "This means removing the likelihood..." A user will not be able to do that, so perhaps suggest that replacing the DNA nucleotide data with all '?' has the same effect as removing the likelihood or setting the log-likelihood to zero.

page 17 top left: this paragraph is correct but also reflects the position of a subjective Bayesian, there are empirical Bayes approaches that estimate priors from the data.

References

1. Felsenstein J: The Number of Evolutionary Trees. *Systematic Zoology*. 1978; **27** (1). Publisher Full Text
2. Cayley, A.: A Theorem on Trees. *Paris: The Quarterly Journal of Mathematics, 23, 376-378*. 1889.
3. Duane S, Kennedy A, Pendleton B, Roweth D: Hybrid Monte Carlo. *Physics Letters B*. 1987; **195** (2): 216-222 Publisher Full Text
4. Hastings W: Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*. 1970; **57** (1). Publisher Full Text
5. Lartillot N, Philippe H: Computing Bayes factors using thermodynamic integration.*Syst Biol*. 2006; **55** (2): 195-207 PubMed Abstract | Publisher Full Text
6. Xie W, Lewis PO, Fan Y, Kuo L, et al.: Improving marginal likelihood estimation for Bayesian phylogenetic model selection.*Syst Biol*. 2011; **60** (2): 150-60 PubMed Abstract | Publisher Full Text
7. Wang YB, Milkey A, Li A, Chen MH, et al.: LoRaD: Marginal likelihood estimation with haste (but no waste).*Syst Biol*. 2023; **72** (3): 639-648 PubMed Abstract | Publisher Full Text
8. Huelsenbeck JP, Larget B, Alfaro ME: Bayesian phylogenetic model selection using reversible jump Markov chain Monte Carlo.*Mol Biol Evol*. 2004; **21** (6): 1123-33 PubMed Abstract | Publisher Full Text
9. Mau B, Newton M: Phylogenetic Inference for Binary Data on Dendograms Using Markov Chain Monte Carlo. *Journal of Computational and Graphical Statistics*. 1997; **6** (1): 122-131 Publisher Full Text
10. Rannala B, Yang Z: Probability distribution of molecular evolutionary trees: a new method of phylogenetic inference.*J Mol Evol*. 1996; **43** (3): 304-11 PubMed Abstract | Publisher Full Text
11. Larget B, Simon D: Markov Chasin Monte Carlo Algorithms for the Bayesian Analysis of Phylogenetic Trees. *Molecular Biology and Evolution*. 1999; **16** (6): 750-759 Publisher Full Text
12. Sequential Monte Carlo Methods in Practice. 2001. Publisher Full Text
13. Shuying Li: Phylogenetic tree construction using Markov chain Monte Carlo /, Ph.D. thesis,The Ohio State University, The Ohio State University. 1996.

Is the rationale for developing the new method (or application) clearly explained?

Yes

Is the description of the method technically sound?

Yes

Are sufficient details provided to allow replication of the method development and its use by others?
Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?
No source data required

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: I develop Bayesian inference methods in population genetics [population genetics, MCMC, Bayesian inference]

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 09 Dec 2025

Joëlle Barido-Sottani

Thank you very much for the review and the suggestions! Please find our detailed response to your comments below (in bold).
Abstract:
3rd line: 'solutions' seems vague **Changed as the whole sentence was restructured.**
line 5: 'evaluated' I would prefer 'estimated' or 'found' **Agreed, changed to 'estimated'.**
line 7 from bottom: I would prefer direct wording: We offer an overview... **Changed to the more direct "We thus offer an overview..."**
Page 4: Introduction: for the 34 million trees you cite Degnan and Salter 2005, which is a fine citation, but there were earlier citations, for example, Felsenstein 1978 (there may be even earlier ones, but I do not know them; the groundwork for counting trees was given by A. Cayley in 1889, but he does not elaborate on specific numbers. **Added.**
the next sentence: 'It is intractable to evaluate all of them.' This statement is incorrect, evaluating 34 million items is certainly in the realm of current computers (for example, Swoffords's PAUP* can evaluate and enumerate all 10-tip (perhaps up to 12 tips) trees in a parsimony context. It is certainly tedious. How about: 'Evaluating all phylogenies with a number of samples larger than 10 become quickly intractable.'
We changed this sentence following the suggestion.
"That technique is Markov chain Monte Carlo (MCMC)," the next sentences show potential alternatives with citations, it seems odd that you do not cite MCMC here. Some reviewers suggested to show that alternatives exist. I find showing these "alternatives" that are

'variants' of MCMC in the introduction unnecessary but could be discussed later towards the end of the tutorial. I suggest to cite Metropolis et al. here. Concerning the citation of the alternative methods given that this is early in the introduction, readers may want to have a citation of the method and not the application in phylogenies, for example, SMC: Doucet et al. 2001 [or perhaps Gordon 1998 associated with introduction of SMC), HMC: Duane et al. [I would remove the other methods and move them somewhere else but still cite Metropolis et al.]
We've added this citation and moved the discussion of other algorithms to figure 2.
****Figure 1 and text discussing figure 1: MCMC was introduced by Metropolis et al. and the outline in Figure 2 follows the 'Metropolis algorithm' and should be labeled the 'Metropolis ratio' but see below: Hastings (1970) generalized the Metropolis algorithm because the original ratio assumes that the proposal distributions for the newly proposed parameter and the old parameters are symmetric, i.e., g(x—y) = g(y—x). In phylogenetics, we do not necessarily guarantee this condition. Modern implementations use the Metropolis-Hastings algorithm to allow for potential asymmetries in the proposal. Figure 1 is labeled 'Hastings ratio,' and a reader needs to know why this is called 'Hastings ratio' because Hastings is not mentioned in the text. I suggest revising the text to include Hastings (1970) and relabeling the 'Hastings ratio' to 'Metropolis-Hastings ratio.'
We have added more context on the ratio as suggested here and in the following legend. We also now discuss the proposal distribution in more details in the move section.
I also suggest adding context after the Metropolis et al. citation on page 5, along these thoughts: "The original algorithm assumed that the moves in the parameter space are symmetric, going from parameter value a to parameter value b has the same probability as going from b to a. Hastings (1970) introduced an additional term to balance potentially asymmetric proposals; this led to a general version: the Metropolis-Hastings algorithm. I would remove the sentence, "In this case, all the values ... phylogenetics.
We have edited this part according to your suggestion.
And then, in figure 1, use 'Metropolis-Hastings ratio': with minimal change to the figure to resolve the issue, one could simply add in the figure legend: "We left out the Hastings part of the ratio for clarity; it would add g(old—new)/g(new—old) to R."
Thank you for pointing this out and for providing such a clean solution. We have implemented these suggestions.
page 6 last 3 lines of 2.1.2 use high vs low (not high vs short) **Fixed.**
page 6 first column, bottom and second column top: I suggest to cite the authors who defined the methods: path sampling in phylogenetics (Lartillot and Phillipe 2006), Stepping Stone (Xie et al. 2010), nested sampling (Russel et al. 2018 [applies Skilling 2006 to phylogenetics] and very recently, a Partition-Weighted Kernel method (Wang et al., 2023).
Thank you for bringing this to our attention. We have added more details to this section and now mention the PWK method.
page 6 second column: For instance, reversible jump can be used to infer the most adequate substitution model for an alignment (Huelsenbeck et al. 2004) [they estimate the best substitution model among 200+ models 13 years before Bouchaert and Drummond]
Fixed.
page 9, second column: Algorithms for efficiently sampling phylogenetic trees . . . (citations, see below) this would be a good location to show the early Bayesian phylogenetic methods because they all had to solve the tree-changing problem and this will give a good way to

give credit to early work on Bayesian phylogenetics, for example, Li often gets forgotten.
Thank you for the suggestion, we have added these references.
page 11 ”(e.g., in tracer)” You have a section later on Tracer; I would simply remove this () because, at this point, the reader does not know that you mean the tracer application.
Fixed.
page 16 left bottom: The sentence may confuse the reader by using ’temperature’ from 1 to 0. Usually, we would consider the Metropolis-Hastings ratio powered with 1/T, where we would set the cold chain at temperature T=1 and also consider hotter chains with higher T, for example, a T=1,000,000 would set that exponent to approximately zero so that all moves will be accepted; any intermediate value will accept more often than the cold chain.
This has been refined as suggested.
page 16 right 5.2. choice of models and priors ”This means removing the likelihood...” A user will not be able to do that, so perhaps suggest that replacing the DNA nucleotide data with all ’ ?’ has the same effect as removing the likelihood or setting the log-likelihood to zero.
This has been refined as suggested.
page 17 top left: this paragraph is correct but also reflects the position of a subjective Bayesian, there are empirical Bayes approaches that estimate priors from the data.
This is a good point, added in the introduction to the prior section: The above represents a subjective Bayesian perspective, in which a researcher is interpreting the probability of different hypotheses, and updating their belief to the posterior probability using data. Other approaches, such as empirical Bayes in which the prior may be estimated from the data, are also possible. It should be noted that most Bayesian practitioners will use some tools from both schools of thought.

Competing Interests: No competing interests were disclosed.

Reviewer Report 02 December 2024

https://doi.org/10.21956/openreseurope.19757.r46205

© 2024 M. Farina B. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Bruna M. Farina
[1] University of Fribourg, Fribourg, Switzerland
[2] Swiss Institute of Bioinformatics, Lausanne, Vaud, Switzerland
[3] University of Fribourg, Fribourg, Switzerland
[4] Swiss Institute of Bioinformatics, Lausanne, Vaud, Switzerland

I appreciate the opportunity to review this manuscript. In this work, the authors provide detailed instructions and guidance on how to better perform Bayesian phylogenetic inference using Markov chain Monte Carlo (MCMC). I really enjoyed how exhaustively (in a good way!) the authors explored the topic, providing very useful tips for both beginners and experts in Bayesian analyses, while keeping the explanations clear and easy to understand.

I only have a few minor comments, listed below, but I want to highlight that these are merely suggestions, and the authors are free to incorporate them or not.

1. **Page 2:** I would suggest including "RevBayes" in the keywords as well.
2. **Page 4:** *"The algorithm for MCMC sampling most frequently employed in phylogenetic studies is known as the Metropolis-Hastings (MH) algorithm (Figure 2), though others exist.".* I recommend mentioning other algorithms and stating—either here or in Section 2.1.5, wherever it flows better—why MH is the most commonly used algorithm in phylogenetics compared to alternatives, like Gibbs sampling.
3. **Page 6: Section 2.1.5 –** see comment 2 above.
4. **Page 8: Section 2.2.2 –** in this section you present the three softwares (BEAST2, RevBayes, and MrBayes). It might be helpful to include a more explicit comparison between them, discussing their advantages or why one might prefer one over the others. For instance, RevBayes is highly flexible and powerful, but it requires some scripting (although there are many tutorials available to help users).
5. **Page 16: Section 5.1.3 –** *"...poorly chosen temperatures can lead to inefficient mixing and other issues, and should be looked out for (Brown & Thomson, 2018)."* I think it is worth briefly mentioning other potential issues might be useful since MCMCMC is MrBayes' default (and widely used).
6. **Page 16 – 18: Section 5.2.1–** Consider defining "hyperprior" as well. The user may encounter this term when running their analyses, so providing a brief definition, along with advice and explanations on why hyperpriors are used in hierarchical models, could be beneficial.
7. Not necessary for understanding the text (it's already super clear), but it would be nice to include a figure illustrating how an informative prior can narrow down the parameter space compared to a vague, uninformative prior.

Is the rationale for developing the new method (or application) clearly explained?
Yes

Is the description of the method technically sound?
Yes

Are sufficient details provided to allow replication of the method development and its use by others?
Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?
No source data required

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Macroevolution, paleobiology

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response 09 Dec 2025

Joëlle Barido-Sottani

Thank you very much for the review and the comments! Please find our detailed response to your comments below (in bold).
1. Page 2: I would suggest including "RevBayes" in the keywords as well. **Added**

2. Page 4: "The algorithm for MCMC sampling most frequently employed in phylogenetic studies is known as the Metropolis-Hastings (MH) algorithm (Figure 2), though others exist.". I recommend mentioning other algorithms and stating—either here or in Section 2.1.5, wherever it flows better—why MH is the most commonly used algorithm in phylogenetics compared to alternatives, like Gibbs sampling.
We now mention slice sampling as an alternative algorithm, but we found very little discussion of why MH is so popular so we left the rest of the text as is.

3. Page 6: Section 2.1.5 – see comment 2 above. **Please see our response above.**

4. Page 8: Section 2.2.2 – in this section you present the three softwares (BEAST2, RevBayes, and MrBayes). It might be helpful to include a more explicit comparison between them, discussing their advantages or why one might prefer one over the others. For instance, RevBayes is highly flexible and powerful, but it requires some scripting (although there are many tutorials available to help users).
Thank you for the suggestion. We have added some information on the different frameworks and what type of use they are best suited for.

5. Page 16: Section 5.1.3 – ". . . poorly chosen temperatures can lead to inefficient mixing and other issues, and should be looked out for (Brown & Thomson, 2018)." I think it is worth briefly mentioning other potential issues might be useful since MCMCMC is MrBayes' default (and widely used).
We have added a note for the computing time of MC3 and a recommendation of using the parallel version of the program.

6. Page 16 – 18: Section 5.2.1– Consider defining "hyperprior" as well. The user may encounter this term when running their analyses, so providing a brief definition, along with advice and explanations on why hyperpriors are used in hierarchical models, could be beneficial.
Thank you for suggesting this addition. We now include a short paragraph explaining what hyperpriors are and how they are used.

7. Not necessary for understanding the text (it's already super clear), but it would be nice to include a figure illustrating how an informative prior can narrow down the parameter space compared to a vague, uninformative prior.

Good point, but instead of adding another figure to this paper, we decided to make use of already existing visualizations, and thus now refer to Figure 1 in dos Reis et al. 2016 in Section 3.3 and to Figure 1 in dos Reis et al. 2015 in Section 5.2.1.

Competing Interests: No competing interests were disclosed.

Version 2

Reviewer Report 29 June 2024

https://doi.org/10.21956/openreseurope.19285.r41769

© 2024 Casali D. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Daniel Casali
[1] Universidade de São Paulo, São Paulo, Brazil
[2] Universidade de São Paulo, São Paulo, Brazil
[3] Universidade de São Paulo, São Paulo, Brazil

I do not have any further comments.

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Systematics (phylogenetics and taxonomy), morphology, and evolutionary biology.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Version 1

Reviewer Report 03 June 2024

https://doi.org/10.21956/openreseurope.18012.r36509

© 2024 Kong S. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Sungsik Kong
[1] University of Wisconsin-Madison, Madison, Wisconsin, USA
[2] University of Wisconsin-Madison, Madison, Wisconsin, USA
[3] University of Wisconsin-Madison, Madison, Wisconsin, USA
[4] University of Wisconsin-Madison, Madison, Wisconsin, USA

Title: Practical guidelines for Bayesian phylogenetic inference using Markov Chain Monte Carlo (MCMC)
Authors: Barido-Sottani J, Schwery O, Warnock RCM, Zhang C, and Wright AM

Summary:
The manuscript by Barrido-Sottani et al. offers practical guidelines for biologists conducting phylogenetic analysis using softwares that implement Bayesian frameworks, namely MrBayes, RevBayes, and BEAST2. The authors focus on Markov chain Monte Carlo (MCMC) that is used in Bayesian phylogenetic inference, and present how to diagnose some common problems and how to troubleshoot them, how to fine-tune parameters for to achieve convergence, and unique challenges when fossil information is incorporated. Overall, the manuscript is in a great shape, very well explained without using unnecessary jargons. I believe this manuscript is well-suited for introductory read for understanding MCMC in Bayesian phylogenetic inference to general audience. Below, I tried to point out any place that can be clarified where a novice to Bayesian phylogenetics (or phylogenetics in general) can possibly get confused.

Comments:
[1] MCMC: Markov Chain Monte Carlo → Markov chain Monte Carlo, unless intended.
[2] Keywords: May be include RevBayes as well?
[3] Abstract: (1) It would be more accurate to say that “estimating a phylogenetic tree involves evaluating many possible hypotheses”, instead of solutions? If the authors meant parameter estimates by the `solutions', it should be clarified.
[4] Introduction to MCMC Paragraph 1: If possible, providing some numbers to demonstrate the vastly large number of possible topologies for given number of taxa n would help readers to admit that it is not feasible to evaluate all possible topologies (even when n is not large). For example, there are < 34 million topologies when n = 10 (Degnan and Salter, 2005, Evolution, 59(1), p.34).
[5] Introduction to MCMC P3: (1) Try to avoid begin a sentence with an abbreviation (i.e., ML, sentence 2). (2) Plus, it would be good idea to provide some references of the mathematical techniques in sentence 2. (3) Also, the concept of prior distribution is briefly introduced here, and it would be nice to direct reader that the concept is described in more detail in the coming section (i.e., 2.1.2?.)
[6] Introduction to MCMC P4: (1) Please mention some examples (or references) of MCMC sampling algorithms other than MH. (2) The authors mention that the one or more model parameters are perturbed, like making a number a little bigger. It would be clearer if stated as “a value of the parameter is increased''? (3) The connection between Monte Carlo and pseudorandom parameter perturbation is not obvious. (4) The concept of ‘parent solution’ should have been introduced, possibly when a ‘starting set of values’ was introduced in the beginning of this paragraph?
[7] Introduction to MCMC P5: Provide reference for the original MCMC in early 1950s?
[8] Section 2.1: I like “The Bayesics.” Just a suggestion, but unless intended, how about listing the

three quantities in the order of apparance in the text (i.e., the likelihood, the prior, and the marginal probability)?
[9] Figure 1: The cartoon illustration of Bayes' theorem looks great. (1) Just make sure to keep the capitalization consistent (Bayes' theorem vs. Bayes Theorem (in text)). (2) It would be more helpful if both top and bottom (i.e., not the second panel) are mentioned in the bolded portion of the caption. (3) Mention the asterisk (∗) represents new parameter values in the caption. (3) I think using either one of Hastings ratio or the posterior odds ratio in the figure itself would remove unnecessary confusion, and mention that both refer to the same thing in the caption.
[10] Section 2.1.1 P1: ML for maximum likelihood is already introduced previously. No need to redo it here.
[11] Section 2.1.2 P1: (1)Use ML since it has been introduced previously. (2) "...than one that is very short → low?"
[12] Section 2.1.4 P2: Provide reference for reversible MCMC. Possibly provide an example for its application (e.g., phylogenetic network inference?).
[13] Section 2.1.5 P2: There is no section called Introduction, but Introduction to MCMC.
[14] Figure 2: I am not sure it it is necessary to include both flowchart and pseudocode. I feel like they are redundant and either one of them should be enough to explain MH. It is up to the authors to choose which one. Furthermore, in the flowchart, it would be more clear to say that the changes that decrease the posterior is accepted at probability of R, instead of vague 'Occasionally'. In pseudocode, possibly add a line somewhere that explains the variable j.
[15] Section 2.2.1 P1: (1) The definition of unrooted trees should be more explicit by including the concept of lack of evolutionary directionality (i.e., unable to identify the ancestor–descent relationship between the nodes). (2) It would also be important to specify the trees where branch lengths represent genetic distance are dated only if the constant molecular clock assumption is met.
[16] Section 2.2.1 P2: (1) Could you define outgroup samples in simple terms? (2) May be use the term 'ingroup' instead of "main clade", because the main clade can be a portion of ingroup taxa.
[17] Section 2.2.3 P1: (1) Are the terms 'moves' and 'operators' synonyms? They are being used interchangeably. If so, it might be a good idea to state that they are the same thing upfront (i.e., move the second last sentence to the top) and stick with one of them throughout the section. (2) I don't think "Scaling move" is not explained in the text at this point, or direct readers to Figure 3.
[18] Section 2.2.3 P2: The last sentence of this paragraph that begins with "This is also the reason..." seems to be very important for the readers when comparing results from different implementations. Could you add a couple sentences with an example?
[19] Section 2.2.3 P3: Two abbreviations STX and SPR were never defined in text. Also, it would be useful to direct readers to some literature where the details of these tree moves are explained.
[20] Section 2.2.3 P4: Please explain what it means by "satisfying" results. This may lead to the subjectivity issue.
[21] Figure 3: (1) Please explicitly define, at least in caption, the notations used in the figure: a, s, d, etc., in scaling and sliding moves, for example. U for Uniform distribution? (2) I thought subtree exchange moves swaps the two points in two different branches in the tree (Vaughan et al,2014) [Ref 1]. Currently, it looks like as if two leaves, blue and orange, are swapped and no branch plays a role in this move. Is this what STX really does? Otherwise, the graphical representation may be misleading and confusing. Similar issue in SPR move.
[22] Section 3: "Posterior values" → "posterior probabilities" for a set of parameter values?
[23] Section 3.1 P1: Refer to Figure 3 when explaining the moves involving numerical change.
[24] Section 3.1 P2: (1) "networks" → phylogenetic networks; It would be good idea to define phylogenetic networks or provide citations where the networks are defined and distinction

between the trees and networks (e.g., Kong S, Et al, 2022 [Ref2] In the last sentence, it might be more accurate to say 'topology moves' rather than "tree moves" since the concept of networks is introduced.
[25] Section 3.2 P1: I feel like the first sentence ("Biology is complex...") is repeating what has been said previously.
[26] Section 3.3: This section is very clear and informative. Beautifully written!
[27] Section 4.1 P1: While the proportion of burn-in might not (or might) directly influence MCMC inference result (particularly if ran long enough), could you elaborate on the consequences of setting the proportion of burn-in either too small or large? For starters, setting this value may seem arbitrary and difficult to decide.
[28] Section 4.3 P1: May be I missed, what do you mean by the autocorrelation time τ? How is it measured?
[29] Figure 4: The description of left and right panels are well explained in the caption. However, it would be great what the different rows represent is described. (i.e., top row being not coverged vs. bottom row being the converged "hairy caterpillar"?
[30] Section 5.1.1 P4: The term "acceptance ratios" was used. Is this synonymous to "acceptance proportion" mentioned in section 2? If so, please mention it or keep these terms consistent.
[31] Section 5.1.1 P7: The abbreviation "DAG" is introduced, but never used again in the manuscript. May be unnecessary.
[32] Section 5.2.2 P2: "...help a researcher determine" → '...help a researcher to determine'?

References

1. Vaughan TG, Kühnert D, Popinga A, Welch D, et al.: Efficient Bayesian inference under the structured coalescent.*Bioinformatics*. 2014; **30** (16): 2272-9 PubMed Abstract | Publisher Full Text
2. Kong S, Pons JC, Kubatko L, Wicke K: Classes of explicit phylogenetic networks and their biological and mathematical significance.*J Math Biol*. 2022; **84** (6): 47 PubMed Abstract | Publisher Full Text

Is the rationale for developing the new method (or application) clearly explained?

Partly

Is the description of the method technically sound?

Yes

Are sufficient details provided to allow replication of the method development and its use by others?

Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?

No source data required

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?

Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Phylogenetic method development; phylogenetic networks; hybridization

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response 29 Jul 2024

Joëlle Barido-Sottani

Thank you very much for the review and the comments! Please find our detailed response (bold text) below.

1. MCMC: Markov Chain Monte Carlo → Markov chain Monte Carlo, unless intended. - **Fixed.**
2. Keywords: May be include RevBayes as well? - **Fixed.**
3. Abstract: (1) It would be more accurate to say that "estimating a phylogenetic tree involves evaluating many possible hypotheses", instead of solutions? If the authors meant parameter estimates by the 'solutions', it should be clarified. - **We have clarified this sentence.**
4. Introduction to MCMC Paragraph 1: If possible, providing some numbers to demonstrate the vastly large number of possible topologies for given number of taxa n would help readers to admit that it is not feasible to evaluate all possible topologies (even when n is not large). For example, there are > 34 million topologies when n = 10 (Degnan and Salter, 2005, Evolution, 59(1), p.34). - **We have added the suggested example.**
5. Introduction to MCMC P3: (1) Try to avoid begin a sentence with an abbreviation (i.e., ML, sentence 2). - **Fixed.**
   (2) Plus, it would be good idea to provide some references of the mathematical techniques in sentence 2. - **Added.**
   (3) Also, the concept of prior distribution is briefly introduced here, and it would be nice to direct reader that the concept is described in more detail in the coming section (i.e., 2.1.2?.) - **As suggested, we now mention that priors are discussed in details in the following sections.**
6. Introduction to MCMC P4: (1) Please mention some examples (or references) of MCMC sampling algorithms other than MH. - **We have added information on slice sampling in this paragraph. As suggested by another reviewer, we also now give examples of alternatives to MCMC in the beginning of the section.**
   (2) The authors mention that the one or more model parameters are perturbed, like making a number a little bigger. It would be clearer if stated as "a value of the parameter is increased"? - **We have edited as suggested.**
   (3) The connection between Monte Carlo and pseudorandom parameter perturbation is not obvious. - **We have edited this part to clarify the connection.**
   (4) The concept of 'parent solution' should have been introduced, possibly when a 'starting set of values' was introduced in the beginning of this paragraph? - **As suggested by another reviewer, we replaced all mentions of 'parent solution' by 'current solution', which we hope makes this part clearer.**
7. Introduction to MCMC P5: Provide reference for the original MCMC in early 1950s? - **We have added this reference.**
8. Section 2.1: I like "The Bayesics." Just a suggestion, but unless intended, how about

listing the three quantities in the order of appearance in the text (i.e., the likelihood, the prior, and the marginal probability)? - **The order of quantities has been changed as suggested.**

9. Figure 1: The cartoon illustration of Bayes' theorem looks great. (1) Just make sure to keep the capitalization consistent (Bayes' theorem vs. Bayes Theorem (in text)). - **The text now consistently uses Bayes' theorem.**
(2) It would be more helpful if both top and bottom (i.e., not the second panel) are mentioned in the bolded portion of the caption. - **We have added a global title for the figure in the caption.**
(3) Mention the asterisk (*) represents new parameter values in the caption. - **Added.**
(4) I think using either one of Hastings ratio or the posterior odds ratio in the figure itself would remove unnecessary confusion, and mention that both refer to the same thing in the caption. - **We have removed the mention of posterior odds ratio from the figure.**
10. Section 2.1.1 P1: ML for maximum likelihood is already introduced previously. No need to redo it here. - **Fixed.**
11. Section 2.1.2 P1: (1)Use ML since it has been introduced previously. (2) "...than one that is very short → low?" - **Fixed.**
12. Section 2.1.4 P2: Provide reference for reversible MCMC. Possibly provide an example for its application (e.g., phylogenetic network inference?). - **We have added a reference and an example of application to the selection of substitution models.**
13. Section 2.1.5 P2: There is no section called Introduction, but Introduction to MCMC. - **Fixed.**
14. Figure 2: I am not sure it it is necessary to include both flowchart and pseudocode. I feel like they are redundant and either one of them should be enough to explain MH. It is up to the authors to choose which one. Furthermore, in the flowchart, it would be more clear to say that the changes that decrease the posterior is accepted at probability of R, instead of vague 'Occasionally'. In pseudocode, possibly add a line somewhere that explains the variable j. - **Our intention is that the flowchart presents the principle of the algorithm, whereas the pseudocode includes the technical details of the calculations performed at each step. As a result, we do not feel that these two parts are redundant and have chosen to leave them both in the figure. We now include some details on the variables in the caption.**
15. Section 2.2.1 P1: (1) The definition of unrooted trees should be more explicit by including the concept of lack of evolutionary directionality (i.e., unable to identify the ancestor–descent relationship between the nodes). - **As suggested, we have added this information about unrooted trees.**
(2) It would also be important to specify the trees where branch lengths represent genetic distance are dated only if the constant molecular clock assumption is met. - **Here we define dated trees as trees with branch lengths in terms of time, so genetic distance trees are undated by definition.**
16. Section 2.2.1 P2: (1) Could you define outgroup samples in simple terms? - **We have added a definition and some examples for outgroup samples.**
(2) Maybe use the term 'ingroup' instead of "main clade", because the main clade can be a portion of ingroup taxa. - **We have removed the references to the "main clade" from this section.**
17. Section 2.2.3 P1: (1) Are the terms 'moves' and 'operators' synonyms? They are being

used interchangeably. If so, it might be a good idea to state that they are the same thing upfront (i.e., move the second last sentence to the top) and stick with one of them throughout the section. - **Yes, the terms "moves", "operators" and "proposals" generally refer to the same components, but are used in different software communities. We have edited the text to make this clearer.**
(2) I don't think "Scaling move" is not explained in the text at this point, or direct readers to Figure 3. - **This reference to scaling moves has been removed.**

18. Section 2.2.3 P2: The last sentence of this paragraph that begins with "This is also the reason..." seems to be very important for the readers when comparing results from different implementations. Could you add a couple sentences with an example? - **We have added more details on why numbers of iterations cannot be directly compared between implementations.**
19. Section 2.2.3 P3: Two abbreviations STX and SPR were never defined in text. Also, it would be useful to direct readers to some literature where the details of these tree moves are explained. - **The abbreviations are now spelled out in the text. We have also added a reference to the original description of these moves in the caption to Figure 3.**
20. Section 2.2.3 P4: Please explain what it means by "satisfying" results. This may lead to the subjectivity issue. - **We have added more details in this sentence.**
21. Figure 3: (1) Please explicitly define, at least in caption, the notations used in the figure: a, s, d, etc., in scaling and sliding moves, for example. U for Uniform distribution? - **We have added these definitions in the caption.**
(2) I thought subtree exchange moves swaps the two points in two different branches in the tree (Vaughan et al,2014) [Ref 1]. Currently, it looks like as if two leaves, blue and orange, are swapped and no branch plays a role in this move. Is this what STX really does? Otherwise, the graphical representation may be misleading and confusing. Similar issue in SPR move. - **The coloured triangles were intended to represent subclades with any number of branches and leaves, rather than just leaves alone. We now make this clear in the figure caption.**
22. Section 3: "Posterior values" → "posterior probabilities" for a set of parameter values? - **Fixed.**
23. Section 3.1 P1: Refer to Figure 3 when explaining the moves involving numerical change. - **Added.**
24. Section 3.1 P2: (1) "networks" → phylogenetic networks; It would be good idea to define phylogenetic networks or provide citations where the networks are defined and distinction between the trees and networks (e.g., Kong S, Et al, 2022 [Ref2]) - **We have edited as suggested and added the reference.**
(2) In the last sentence, it might be more accurate to say 'topology moves' rather than "tree moves" since the concept of networks is introduced. - **Fixed.**
25. Section 3.2 P1: I feel like the first sentence ("Biology is complex...") is repeating what has been said previously. - **We have edited this section to remove the repetition.**
26. Section 3.3: This section is very clear and informative. Beautifully written! - **Thank you !**
27. Section 4.1 P1: While the proportion of burn-in might not (or might) directly influence MCMC inference result (particularly if ran long enough), could you elaborate on the consequences of setting the proportion of burn-in either too small or large? For starters, setting this value may seem arbitrary and difficult to decide. - **We have**

added some details in this section.

28. Section 4.3 P1: Maybe I missed, what do you mean by the autocorrelation time τ? How is it measured? - **We have added a definition of τ and information on how it is calculated.**
29. Figure 4: The description of left and right panels are well explained in the caption. However, it would be great what the different rows represent is described. (i.e., top row being not converged vs. bottom row being the converged "hairy caterpillar") - **We have expanded the caption to Figure 4 to better describe the different scenarios.**
30. Section 5.1.1 P4: The term "acceptance ratios" was used. Is this synonymous to "acceptance proportion" mentioned in section 2? If so, please mention it or keep these terms consistent. - **We have changed the word ratio to rate, and we now clarify that acceptance proportion and acceptance ratio are used interchangeably in the literature.**
31. Section 5.1.1 P7: The abbreviation "DAG" is introduced, but never used again in the manuscript. May be unnecessary. - **We have chosen to keep the abbreviation as it is used frequently in other articles discussing RevBayes.**
32. Section 5.2.2 P2: "...help a researcher determine" → '...help a researcher to determine'? - **Fixed.**

Competing Interests: No competing interests were disclosed.

Reviewer Report 12 March 2024

https://doi.org/10.21956/openreseurope.18012.r37734

© 2024 Clarté G. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Grégoire Clarté
[1] University of Helsinki, Helsinki, Finland
[2] University of Helsinki, Helsinki, Finland
[3] University of Helsinki, Helsinki, Finland
[4] University of Helsinki, Helsinki, Finland

This paper summarises the current knowledge about Bayesian phylogenetic inference with a particular focus on numerical methods. The paper is particularly welcome as this problem is particularly complex and frequent in numerous applications. No new method is presented, but the article offers a nice presentation of the most common method and their implementation. It contains also a general workflow for carrying Bayesian Phylogenetical Inference.

Overall, I think some choices are strange (for example the strange notations for parameters and observations in Fig. 1). These choices must be explained by the readers having little mathematical

background. Nevertheless, I think this is not a good idea, as the article should aim at giving the mathematical tools needed to the reader, and I don't think ideographic representation gives better understanding than proper definitions (I needed some time to understand the images). Introducing proper notations would also allow to discuss the notion of ``non informative priors'' and more importantly conjugated priors that would help the choice of priors. It can also be beneficial to allow the users to communicate with the computational statisticians that developed the tools.

For example, in 2.1.6. I think the article would benefit from mentioning the standard results of Markov Chain theory (that is ergodicity of the chain). This in return would be of interest when mentioning the different moves that can be used in MCMC algorithm, and could give insight as for possible reasons of failures.

The part about the different problems of MCMC methods is well written and contains all the important information in my opinion.

I disagree with section 5.1.2, or I think the phrasing is strange: the choice of a fixed initial value can be misleading (see later the case of poorly ergodic chains because of multimodality). I think the choice of a random initialisation from the prior along with repeated runs should be the correct recommendation (along with more careful observation of convergence).

I also have some doubts on the diagnostic section: this problem is one of the most difficult in phylogenetics, for the reason that the parameters of the model (that include the topology) cannot be all summarised by one ESS. Overall all the existing methods are inefficient (but once again that's a review of the numerical methods that exists, and these inefficient methods are the only ones implemented). It would be interesting to the reader to see what are the possible issues that arises from using the ESS in the phylogenetic case (which itself has to be defined, as the author mention). Nevertheless, I am happy to see the part about the trace observation (which should be the first check run by any user of an MCMC method).

I am a little surprised the authors do not mention earlier the most simple method to detect multimodality: running several chains in parallel. If the chains converge to different trees, that indicates that further study is required. In particular the use of standard MCMC such as the ones presented in the paper (MrBeast, etc.) will not be able to handle these problems. In my field this is a very frequent case of failure of the MCMC. This is mentioned later in the paper but it would be beneficial to discuss that earlier. On the question of detecting convergence, the remarks of 5.4 can have links with the results on MCMC convergence on trees developed following the works of Biswas, N., Jacob, P. E. and Vanetti, P. (2019) *Estimating convergence of Markov chains with L-lag couplings* (NeurIPS 7389–7399). These methods are not implemented in the current software but they have the advantage of being theoretically sound, and have been applied to phylogenetics. Of course, this is out of the scope of the article if we consider that the users will only make use of the already implemented tools.

A problem I can see, is that the article mentions problems that cannot be solved in most of the software mentioned unless the user has a good knowledge of them. For example, if there are mixing problems for some parameters, this can be a problem of correlation between the parameters, and a move in the joint space would be recommended, which is unfortunately not possible to implement easily in most cases.

Finally, I am wondering if mentioning more complex but more efficient methods such as SMC for trees (for example the works from Bouchard-Côté) can be interesting. My idea is that if there are problems with the standard methods, the user should move on to more complex numerical methods, although they are not implemented in a user friendly way.

As a conclusion, I would say that the article is good but requires some polishing depending on the goals the authors have. It obviously is tailored for non-statisticians that are using Bayesian phylogenetics tools. It does a great job at presenting the problems, the tools, and the possible solutions (I am happy to see the recommendation to discuss complex problems on the forums) but I still think stronger maths would be beneficial.

Minor comments:

- "2.1 The Bayesics" I'm not sure the pun is needed in this kind of article (also, it can be difficult to understand for non native english speakers that would pronounce Bayes differently).

- I have philosophical problems with the notion of "choosing" the prior, as the prior is not chosen by the user but just exists out of the general knowledge (as is described by the paper). I think the author could mention repeated experiments with different priors to ensure the prior effect is negligible with respect to the results.

- I use more often marginal likelihood than marginal probability to designate the integrate of the likelihood times the prior, I don't know what are the terms used in the phylogenetics community.

Is the rationale for developing the new method (or application) clearly explained?

Yes

Is the description of the method technically sound?

Yes

Are sufficient details provided to allow replication of the method development and its use by others?

Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?

Yes

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?

Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Computational bayesian phylogenetics, computational statistics, numerical

methods in bayesian statistics.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 30 May 2024

Joëlle Barido-Sottani

Thank you very much for the review and suggestions! Please find our detailed response below (in bold). Overall, I think some choices are strange (for example the strange notations for parameters and observations in Fig. 1). These choices must be explained by the readers having little mathematical background. Nevertheless, I think this is not a good idea, as the article should aim at giving the mathematical tools needed to the reader, and I don't think ideographic representation gives better understanding than proper definitions (I needed some time to understand the images). Introducing proper notations would also allow to discuss the notion of "non informative priors" and more importantly conjugated priors that would help the choice of priors. It can also be beneficial to
allow the users to communicate with the computational statisticians that developed the tools.
**We have chosen to retain the pictures used in place of mathematical notation in Figure 1 (although, as described in response to reviewer 2, we have made the image clearer and we now include a figure legend with descriptions of the components). Our reasoning for this is that phylogenetics is perhaps unusual in that it draws a lot of users who have no formal training in mathematics, who often have limited practice or experience reading equations in any context. Formal notation appears in many fundamental textbooks and papers that introduce Bayesian phylogenetics and MCMC (e.g., Ziheng Yang's (2006) Computational Molecular Evolution book or Joseph Felsenstein's (2004) Inferring phylogenies book). These resources can be challenging at first for those unfamiliar with mathematical notation and new to phylogenetics. Our intention is to provide a reference that is both complementary to these resources and more accessible for beginners. This perspective is based on our experience as individual researchers
and our experience teaching Bayesian phylogenetics in a wide range of contexts (especially, as part of the RevBayes or BEAST2 workshops which are aimed at empirical researchers).**
We have included some discussion on the use of non-informative or improper priors and the potential issues these can cause (section 5.2 Choice of model and priors). However, we chose not to discuss the use of conjugate priors, as it is not often possible to use these in the context of Bayesian phylogenetics. A regular phylogenetic software user would rarely (if ever) encounter this term. For example, in 2.1.6. I think the article would benefit from mentioning the standard results of Markov Chain theory (that is ergodicity of the chain). This in return would be of interest when mentioning the different moves that can be used in MCMC algorithm, and could give insight as for possible reasons of failures.
We have added some information on ergodicity and its implications on the MCMC

moves. The part about the different problems of MCMC methods is well written and contains all the important information in my opinion.
I disagree with section 5.1.2, or I think the phrasing is strange: the choice of a fixed initial value can be misleading (see later the case of poorly ergodic chains because of multimodality). I think the choice of a random initialisation from the prior along with repeated runs should be the correct recommendation (along with more careful observation of convergence).
We have edited this to make it clearer. We aren't trying to argue for one method of doing this over another, simply to state the major classes of solutions found in the literature. To this section, we've added some explanation about how starting point may affect the inference, and how to detect sensitivity to starting point. I also have some doubts on the diagnostic section: this problem is one of the most difficult in phylogenetics, for the reason that the parameters of the model (that include the topology) cannot be all summarised by one ESS. Overall all the existing methods are inefficient (but once again that's a review of the numerical methods that exists, and
these inefficient methods are the only ones implemented). It would be interesting to the reader to see what are the possible issues that arises from using the ESS in the phylogenetic case (which itself has to be defined, as the author mention). Nevertheless, I am happy to see the part about the trace observation (which should be the first check run by any user of an MCMC method).
Standard convergence diagnostics in phylogenetics involve calculating one ESS per parameter that is involved in the analysis, rather than one single ESS for the whole model, as the initial text might have implied. Some of the confusion may have resulted from the later mention of the topology ESS as calculated by Convenience, since that comes closest to the idea of one single ESS for the whole analysis. We have thus edited the text to make that clearer. We have furthermore added a brief section detailing which ESS to check and what to pay attention to. I am a little surprised the authors do not mention earlier the most simple method to detect multimodality: running several chains in parallel. If the chains converge to different trees, that indicates that further study is required. In particular the use of standard MCMC such as the ones presented in the paper (MrBeast, etc.) will not be able to handle these problems. In my field this is a very frequent case of failure of the MCMC. This is mentioned later in the paper but it would be beneficial to discuss that earlier.
We now mention the use of multiple chains in the convergence assessment section. We have edited this part to emphasize that certain issues such as multimodality can only be detected through this method. On the question of detecting convergence, the remarks of 5.4 can have links with the results on MCMC convergence on trees developed following the works of Biswas, N., Jacob, P. E. and Vanetti, P. (2019) Estimating convergence of Markov chains with L-lag couplings (NeurIPS 7389–7399). These methods are not implemented in the current
software but they have the advantage of being theoretically sound, and have been applied to phylogenetics. Of course, this is out of the scope of the article if we consider that the users will only make use of the already implemented tools.
We agree that while this may be a nascent method, it might be useful for readers to be aware of it and explore the possibilities of its use. We have thus added a mention of it to the convergence section, along with two papers by Kelly et al, who already apply the method of Biswas et al. in a phylogenetic context. A problem I can see, is that the

article mentions problems that cannot be solved in most of the software mentioned unless the user has a good knowledge of them. For example, if there are mixing problems for some parameters, this can be a problem of correlation between the parameters, and a move in the joint space would be recommended, which is unfortunately not possible to implement easily in most cases.
We have tried to cover as many as possible solutions available in the software that users can adjust or fine-tune across appropriate sections. We now clearly state that our focus in on solutions which can be implemented by the user without additional development, but that many issues do require changes in the software itself. The problem of correlation between parameters is particularly hard to deal with, typically requiring the developers to implement efficient MCMC moves to overcome it. Nevertheless, we now mention the up-down operator in BEAST2 which will scale both the branch lengths of the tree, and the clock rate simultaneously. Finally, I am wondering if mentioning more complex but more efficient methods such as SMC for trees (for example the works from Bouchard-Côté) can be interesting. My idea is that if there are problems with the standard methods, the user should move on to more complex numerical methods, although they are not implemented in a user friendly
way.
We added references to other algorithms when introducing MCMC. These are very interesting and worth discussing, but out-of-scope for us. As a conclusion, I would say that the article is good but requires some polishing depending on the goals the authors have. It obviously is tailored for non-statisticians that are using Bayesian phylogenetics tools. It does a great job at presenting the problems, the tools, and the possible solutions (I am happy to see the recommendation to discuss complex problems on the forums) but I still think stronger maths would be beneficial. Minor comments:
"2.1 The Bayesics" I'm not sure the pun is needed in this kind of article (also, it can be difficult to understand for non native english speakers that would pronounce Bayes differently).
Thank you for pointing out the valid concern about non-native speakers. In consultation with the non-native English speakers among the authors, we opted to remove the pun here. I have philosophical problems with the notion of "choosing" the prior, as the prior is not chosen by the user but just exists out of the general knowledge (as is described by the paper). I think the author could mention repeated experiments with different priors to ensure the prior effect is negligible with respect to the results.
We disagree with the idea that the prior should be negligible. If the prior is reflecting sound empirical information, that information should be incorporated in the analysis. But more to the point, any analysis contains many choices. Whether it is a choice of methodology, or the choice of null and alternative hypotheses, these choices are often consequential. Thus, it's our feeling that as with any other choice, consequentiality is less important than justifiability. We have kept the phrasing of "choice" when discussing the priors, but we now mention testing different priors to explore their impact on the results and conclusions. We also now emphasise that in some contexts in Bayesian phylogenetics the parameters of interest are not fully identifiable (this is especially true for divergence time estimation) and that as a result the priors need to very carefully considered (sections 2.1.2 and 3.3). I use more often marginal likelihood than marginal probability to designate the integrate of the likelihood times the prior, I don't know what are the terms used in the phylogenetics community.

We edited the first mention of the term to state that both terms are used interchangeably.

Competing Interests: No competing interests were disclosed.

Reviewer Report 12 March 2024

https://doi.org/10.21956/openreseurope.18012.r37731

© 2024 Hu Z. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Zhirui Hu
[1] Gladstone Institute of Data Science and Biotechnology, San Francisco, California, USA
[2] Gladstone Institute of Data Science and Biotechnology, San Francisco, California, USA
[3] Gladstone Institute of Data Science and Biotechnology, San Francisco, California, USA
[4] Gladstone Institute of Data Science and Biotechnology, San Francisco, California, USA

The authors provide a very thorough introduction of Bayesian analysis, MCMC for posterior inference in Bayesian phylogenetics, common issues and troubleshooting of model choices and running MCMC. The article is well-structured, informative, easy to read and technically sound. In general, I think it would be more helpful to include more examples and figures to illustrate issues and troubleshooting techniques, and/or an example illustrating the entire process from model design, MCMC convergence diagnostics and model selection or update.

Besides, there are a few minor issues in the article:

- page 4, “between the new and parent scores”, “current” is more commonly used than “parent” to indicate the current state in the Markov chain.
- Figure 1, the cartoon illustration for phylogeny and alignment etc. is interesting but too small to read. Also, summation sign instead of integral could be used for summing over discrete variables, i.e. phylogenetic trees.
- Page 7, “a minimum of 10,000 posterior samples”, I think choosing the number of MCMC samples should depend on how correlated samples are or ESS.
- Section 2.2.3, “Scaling move the components”, is it a typo?
- Pg 9, section 3.1, the author mentioned that traversing tree space was largely solved in the 1990s but later mentioned challenges in phylogenetic inference. I think the authors need to clarify in which situation the problem was solved.
- Section 3.2, adding a figure might be helpful to illustrate the posterior space. “when calculating the prior probability of the phylogeny…”, should it be “posterior probability”?

Also, what is a diversification model? Maybe add some reference here.

- Figure 4 is good to illustrate different types of trace plot, but the authors could add more explanations on Figure 4. What are the problems of the first two trace plots?
- Section 5.1.1 tuning step size is very important in MCMC. The author could provide a table or list to summarize pros and cons of large/small step size and tuning step size.
- Pg 16, posterior predictive simulations is very useful and it would be helpful if the authors can provide a toy example of it. “If the model is adequate to describe/analyse the variation in the data,...”, is it a typo? Should it be “inadequate”?
- Pg 17, any reference for ABC?

Is the rationale for developing the new method (or application) clearly explained?
Yes

Is the description of the method technically sound?
Yes

Are sufficient details provided to allow replication of the method development and its use by others?
Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?
No source data required

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Bayesian statistics, phylogenetics

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response 30 May 2024

Joëlle Barido-Sottani

Thank you very much for the review and the comments! Please find our response to the other issues below (in bold). In general, I think it would be more helpful to include more examples and figures to illustrate issues and troubleshooting techniques, and/or an example illustrating the

entire process from model design, MCMC convergence diagnostics and model selection or update.
We feel that specific examples of the troubleshooting process will be too software-dependent to be helpful in the manuscript, however we have written a RevBayes tutorial illustrating many of the issues and techniques outlined here. The tutorial is now mentioned in section 6. Besides, there are a few minor issues in the article: page 4, "between the new and parent scores", "current" is more commonly used than "parent" to indicate the current state in the Markov chain. **Fixed.** Figure 1, the cartoon illustration for phylogeny and alignment etc. is interesting but too small to read. Also, summation sign instead of integral could be used for summing over discrete variables, i.e. phylogenetic trees.
**We have reorganised Figure 1 to make everything larger, added a legend, and added colour to distinguish the data vs. the model (+ tree). We have also expanded the legend to include more information. We opted not to add the summation for discrete variables, since the tree is also comprised of continuous
branch lengths. We believe the main point we are trying to make here — that the marginal probability of the data must take into account all possible parameter values and trees — is accurately conveyed without making this expression more complicated.** Page 7, "a minimum of 10,000 posterior samples", I think choosing the number of MCMC samples should depend on how correlated samples are or ESS.
We have edited this section to make it clearer that the number of recorded samples is due to space and memory constraints, and that a higher number of samples is not indicative of the convergence of the chain, since as you note samples are correlated. Section 2.2.3, "Scaling move the components", is it a typo?
It was indeed a typo, thank you for pointing this out. Pg 9, section 3.1, the author mentioned that traversing tree space was largely solved in the 1990s but later mentioned challenges in phylogenetic inference. I think the authors need to clarify in which situation the problem was solved.
This was unclear indeed. This has been clarified to "Algorithms for efficiently sampling phylogenetic tree space became available in the late 1990s" Section 3.2, adding a figure might be helpful to illustrate the posterior space.
We believe that adding a figure would make this section too long. However, we have added in the introduction links to several online tools which allow users to gain a better intuition for posterior spaces and how the MCMC algorithm works. "when calculating the prior probability of the phylogeny. . . ", should it be "posterior probability"?
We have removed the word "prior" as the probability of the phylogeny can be part of the prior or the likelihood depending on the analysis. Also, what is a diversification model? Maybe add some reference here.
Following a suggestion by Reviewer 1, we now briefly introduce diversification models in section 2.1.1 and give a reference. We hope that this clarifies the text here. Figure 4 is good to illustrate different types of trace plot, but the authors could add more explanations on Figure 4. What are the problems of the first two trace plots?
We have added clarifying comments to the caption of this, pointing to diagnostic features of the problems shown. Section 5.1.1 tuning step size is very important in MCMC. The author could provide a table or list to summarize pros and cons of large/small step size and tuning step size.
We agree that these are important elements of running an MCMC, but feel that our

treatment of the pros and cons of them in this section would be sufficient. We thought that adding an extra table only to summarize this information again would not add all that much clarity. Instead, since we are already demonstrating the outcomes of different stepsizes and tuning parameters in Figure 4, we have added a reference to the figure in the respective parts of this section. Pg 16, posterior predictive simulations is very useful and it would be helpful if the authors can provide a toy example of it.
We agree that toy examples would be helpful, but we believe that this would be covered better by software-specific tutorials. We now link to such a tutorial on the RevBayes website. "If the model is adequate to describe/analyse the variation in the data,. . . ", is it a typo? Should it be "inadequate"?
This was indeed a typo and we have changed it to "inadequate". Pg 17, any reference for ABC? **Added.**

Competing Interests: No competing interests were disclosed.

Reviewer Report 21 February 2024

https://doi.org/10.21956/openreseurope.18012.r36973

© 2024 Casali D. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Daniel Casali
[1] Universidade de São Paulo, São Paulo, Brazil
[2] Universidade de São Paulo, São Paulo, Brazil
[3] Universidade de São Paulo, São Paulo, Brazil
[4] Universidade de São Paulo, São Paulo, Brazil

Dear editor and authors,

The manuscript prepared by Barido-Sottani et al., entitled "Practical guidelines for Bayesian phylogenetic inference using Markov Chain Monte Carlo (MCMC)", begins providing a concise, but also detailed, review of Markov Chain Monte Carlo (MCMC) sampling procedure, widely used in Bayesian phylogenetic inferences. The main focus of the text then moves to assessing the practical performance of this tool, particularly with respect to convergence issues commonly encountered by researchers.

Despite being a common problem, failing to achieve convergence in Bayesian analyses is anything but a trivial matter, and, until now, to the best of my knowledge, there has been no study that delves so exhaustively and directly into this issue. The manuscript, therefore, constitutes an invaluable contribution to the field of study, useful not only to undergraduate and graduate students embarking on these analyses for the first time, but also to more experienced users.

The article spans from more technical issues, such as improving the operators/movements used

to propose new values during the MCMC sampling progression, to other practical aspects, such as defining well-behaved priors, selecting and critically evaluating models applied in inferences, and understanding dataset characteristics that can lead to performance issues. Useful guidelines are provided on how to initiate an analysis with good chances to converge, as well as how to address common issues like sampling from multiple optima, chain mixing problems, among others. In addition, it provides directions for asking for help in cases where all other solutions have failed. In sum, this article constitutes a very useful source of information for all those interested in the practical aspects of running phylogenetic Bayesian analyses.

Below, I make a few minor suggestions that I believe could enhance the article's utility for its target audience. However, I emphasize that these are only small changes, and I leave it to the authors' discretion to incorporate all or any of them.

1. Page 4. "...this ensures that we explore the entire parameter space and do not stay stuck in a local optimum." **A somewhat pedantic observation here: The *entire* parameter space cannot be ensured to be explored. Perhaps replace by something like: "we broadly explore the parameter space"?**
2. Page 4. "Invented in the early 1950's, MCMC was originally used in physics to describe equilibrium between the liquid and gas phases of a chemical." **I recommend a citation here, for readers interested in the history of the method itself.**
3. Page 5. "...a substitution model, which describes the relative rate of change from one character to another, and...". **I suggest to add also: "...and the relative rate of change among character states.**
4. **P**age 5. "Under this prior, a rate that is very high is believed to be less likely than one that is very [- short] **low**.
5. Page 6. "Different estimation methods have been developed to approximate the marginal likelihood, such as path sampling (Baele et al., 2012)[Ref-7] or nested sampling (Russel et al., 2018), but they remain expensive". **Even though, technically, stepping stones sampling is a kind of path sampling procedure, I would specifically mention stepping stones separately here as well, since is the most widely used method for estimate marginal likelihoods. Also, in the same sentence, I consider that it would be informative to indicate the term "marginal likelihood" is synonymous with the other term more consistently used in the paper, marginal probability.**
6. Page 7. "...We typically record the state of the chain with a frequency that results in a minimum of 10,000 posterior samples.". **But probably less than that, if we perform many moves per generation, as in RevBayes?**
7. Page 7. "Another important feature of phylogenies is whether they are dated, i.e., whether their branch lengths are expressed in units of genetic/**morphological** distance or in units of time." and "Thus we mainly target this article at analyses which include a molecular **/morphological** clock...".
8. Page 7. Estimating a dated phylogeny requires a model for the molecular or morphological clock, **a model of lineage diversification,** as well as time information to calibrate the tree.
9. Page 7. Thus, Scaling **moves** the components designed to advance the chain **and** are a core part of any MCMC inference software.
10. Page 10. "In practice, however, character data is not available or limited for most groups..." **As an alternative here, continuous morphological characters could be used in total-evidence analyses (e.g., Álvarez-Carretero et al. (2019)[Ref-1], Zhang et al. (2021)[Ref-2]in press). These could be more readily available for some taxonomic groups, although the performance of including these characters need to be carefully**

considered (Varón-González et al. (2020)[Ref-3].

11. Page 12*. 5.1 Inference technical setup.* **Here I missed some mentioning of MCMCMC and the use of heated chains or adjusting chain temperature values to try to improve convergence. I guess this is a central topic in the current subject, that should be briefly mentioned by the authors.**
12. Page 12.” If an operator is missing…”. **Maybe the authors could emphasize here that this is a more relevant issue (as far as I can see) in RevBayes, in which we are "freer" to customize the inclusion of operators (or anything else, basically!).**
13. Page 14. “..and with added constraints such as [- monophyletic] subclades.”. **Clades (or subclades) are monophyletic by definition.**
14. **P**age 15. *5.2.1 Priors*. **Here I missed some advice on avoiding the use of improper priors, as in some of the Beast2 default settings. These improper priors can also lead to convergence issues sometimes.**
15. Page 18. ” …partition morphological character data”. **Although only weakly related to the main theme of the paper, I think other methods of morphological data partitioning and what we know about their performances (which is little, if compared to dna…) could be briefly mentioned here, to give a broader picture to the reader (e.g., Clarke & Middleton (2008)[Ref-4], Rosa et al. (2019)[Ref-5], Casali et al. (2023)[Ref-6]).**

My best regards,
Daniel Casali

P.S. The first three questions presented in the peer review form (all answered “YES”), actually do not apply to this study, because no new method is proposed. The paper, although a methods paper, is more of a review and a practical guide for troubleshooting problems in Bayesian phylogenetic analyses.

References

1. Álvarez-Carretero S, Goswami A, Yang Z, Dos Reis M: Bayesian Estimation of Species Divergence Times Using Correlated Quantitative Characters.*Syst Biol*. 2019; **68** (6): 967-986 PubMed Abstract | Publisher Full Text
2. Zhang R, Drummond A, Mendes F: Fast Bayesian inference of phylogenies from multiple continuous characters. *bioRxiv*. 2021. Publisher Full Text
3. Varón-González C, Whelan S, Klingenberg CP: Estimating Phylogenies from Shape and Similar Multidimensional Data: Why It Is Not Reliable.*Syst Biol*. 2020; **69** (5): 863-883 PubMed Abstract | Publisher Full Text
4. Clarke JA, Middleton KM: Mosaicism, modules, and the evolution of birds: results from a Bayesian approach to the study of morphological evolution using discrete character data.*Syst Biol*. 2008; **57** (2): 185-201 PubMed Abstract | Publisher Full Text
5. Rosa BB, Melo GAR, Barbeitos MS: Homoplasy-Based Partitioning Outperforms Alternatives in Bayesian Analysis of Discrete Morphological Data.*Syst Biol*. 2019; **68** (4): 657-671 PubMed Abstract | Publisher Full Text
6. Casali DM, Freitas FV, Perini FA: Evaluating the Impact of Anatomical Partitioning on Summary Topologies Obtained with Bayesian Phylogenetic Analyses of Morphological Data.*Syst Biol*. 2023; **72** (1): 62-77 PubMed Abstract | Publisher Full Text
7. Baele G, Li WL, Drummond AJ, Suchard MA, et al.: Accurate model selection of relaxed molecular clocks in bayesian phylogenetics.*Mol Biol Evol*. 2013; **30** (2): 239-43 PubMed Abstract | Publisher Full Text

Is the rationale for developing the new method (or application) clearly explained?
Yes

Is the description of the method technically sound?
Yes

Are sufficient details provided to allow replication of the method development and its use by others?
Yes

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?
No source data required

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Systematics (phylogenetics and taxonomy), morphology, and evolutionary biology.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response 30 May 2024

Joëlle Barido-Sottani

Thank you very much for the review and the comments! Please find our detailed response (bold text) below. 1. Page 4. ". . . this ensures that we explore the entire parameter space and do not stay stuck in a local optimum." A somewhat pedantic observation here: The entire parameter space cannot be ensured to be explored. Perhaps replace by something like: "we broadly explore the parameter space"? **Fixed.** 2. Page 4. "Invented in the early 1950's, MCMC was originally used in physics to describe equilibrium between the liquid and gas phases of a chemical." I recommend a citation here, for readers interested in the history of the method itself.
Good point, we are now citing Metropolis et al. (1953). 3. Page 5. ". . . a substitution model, which describes the relative rate of change from one character to another, and. . . ". I suggest to add also: "...and the relative rate of change among character states.
We have clarified this to: "substitution model, which describes the relative rate of change from one character state to another as well as the frequencies of each character state", as in most substitution models, the relative rate is based on equilibrium frequencies. 4. Page 5. "Under this prior, a rate that is very high is believed to be less likely than one that is very [- short] low. **Fixed.** 5. Page 6. "Different estimation

methods have been developed to approximate the marginal likelihood, such as path sampling (Baele et al., 2012)[Ref-7] or nested sampling (Russel et al., 2018), but they remain expensive". Even though, technically, stepping stones sampling is a kind of path sampling procedure, I would specifically
mention stepping stones separately here as well, since is the most widely used method for estimate marginal likelihoods. Also, in the same sentence, I consider that it would be informative to indicate the term "marginal likelihood" is synonymous with the other term more consistently used in the paper, marginal probability.
Thanks, we have changed the wording to now mention stepping stone sampling as a popular type of path sampling, and clarify that marginal probability and marginal likelihood are synonymous when introducing the concepts. 6. Page 7. ". . . We typically record the state of the chain with a frequency that results in a minimum of 10,000 posterior samples.". But probably less than that, if we perform many moves per generation, as in RevBayes?
We have edited this section to make it clearer that the number of recorded samples is due to space and memory constraints, and that the sampling frequency will indeed be dependent on the specific software used. 7. Page 7. "Another important feature of phylogenies is whether they are dated, i.e., whether their branch lengths are expressed in units of genetic/morphological distance or in units of time." and "Thus we mainly target this article at analyses which include a molecular/morphological clock. . . ". **Fixed.** 8. Page 7. Estimating a dated phylogeny requires a model for the molecular or morphological clock, a model of lineage diversification, as well as time information to calibrate the tree.
We have edited this section to mention diversification models, but note that dated phylogenies can also be estimated without such a model (e.g. assuming a uniform prior on topologies and some continuous distribution on the branch lengths). 9. Page 7. Thus, Scaling moves the components designed to advance the chain and are a core part of any MCMC inference software. **Fixed.** 10. Page 10. "In practice, however, character data is not available or limited for most groups. . . " As an alternative here, continuous morphological characters could be used in total-evidence analyses (e.g., Álvarez-Carretero et al. (2019)[Ref-1], Zhang et al. (2021)[Ref-2]in press). These could be more readily available for some taxonomic groups, although the performance of including these characters need to be carefully considered (Varón-González et al. (2020)[Ref-3]. **Added.** 11. Page 12. 5.1 Inference technical setup. Here I missed some mentioning of MCMCMC and the use of heated chains or adjusting chain temperature values to try to improve convergence. I guess this is a central topic in the current subject, that should be briefly mentioned by the authors.
Thank you for pointing out this missing topic. We have now added a section on MCMCMC. 12. Page 12." If an operator is missing. . . ". Maybe the authors could emphasize here that this is a more relevant issue (as far as I can see) in RevBayes, in which we are "freer" to customize the inclusion of operators (or anything else, basically!).
We now mention that this issue is particularly relevant for users of RevBayes. 13. Page 14. "..and with added constraints such as [- monophyletic] subclades.". Clades (or subclades) are monophyletic by definition.
We have clarified this sentence. 14. Page 15. 5.2.1 Priors. Here I missed some advice on avoiding the use of improper priors, as in some of the Beast2 default settings. These improper priors can also lead to convergence issues sometimes.
We have added advice on improper priors. 15. Page 18. " . . . partition morphological character data". Although only weakly related to the main theme of the paper, I think other

methods of morphological data partitioning and what we know about their performances (which is little, if compared to dna. . .) could be briefly mentioned here, to give a broader picture to the reader (e.g., Clarke & Middleton (2008)[Ref-4], Rosa et al. (2019)[Ref-5], Casali et al. (2023)[Ref-6]). **Added.**

Competing Interests: No competing interests were disclosed.