Automated convergence diagnostic for phylogenetic MCMC analyses

Lars Berling [1,2], Remco Bouckaert [3], and Alex Gavryushkin [1,2]

[1]Biological Data Science Lab, School of Mathematics and Statistics, University of Canterbury, Christchurch, New Zealand
[2]Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand
[3]Center of Computational Evolution, School of Computer Science, University of Auckland, Auckland, New Zealand

August 11, 2023

Abstract

Convergence assessment in Markov chain Monte Carlo (MCMC) based analyses is crucial but challenging, especially so in high dimensional and complex spaces such as the space of phylogenetic trees (treespace). Here we leverage recent advances in computational geometry of the treespace and introduce a method that combines classical statistical techniques and algorithms with geometric properties of the treespace to automatically evaluate and assess convergence of phylogenetic MCMC analyses. Our method systematically evaluates convergence across multiple MCMC chains and achieves high accuracy in detecting convergence of chains over the treespace. Furthermore, our approach is developed to allow for real-time convergence evaluation during the MCMC algorithm run, eliminating any of the chain post-processing steps that are currently required. Our tool therefore improves reliability and efficiency of MCMC based phylogenetic inference methods and makes analyses easier to reproduce and compare. We demonstrate the efficacy of our diagnostic via a well calibrated simulation study and provide examples of its performance on real data sets.

The open source package for the phylogenetic inference framework BEAST2, called ASM, that implements these methods, making them accessible through a user-friendly GUI, is available from `https://github.com/rbouckaert/asm/`. The open source Python package, called tetres, that provides an interface for these methods enabling their applications beyond BEAST2 can be accessed at `https://github.com/bioDS/tetres/`.

berlinglars96@gmail.com
remco@cs.auckland.ac.nz
alex@biods.org

1 Introduction

Bayesian inference via the Markov chain Monte Carlo (MCMC) algorithm is widely used [1–4] to reconstruct the evolutionary relationships among a set of biological sequences. This inference paradigm is based on a statistical model that describes the probability of observing the sequence data given a particular tree topology and model parameters. MCMC algorithms are typically employed to approximate the true probability distribution, with various convergence diagnostic techniques applied to assess the quality of this approximation.

Assessing the convergence of these MCMC algorithms, however, remains a challenging and complex issue, not only in the context of treespaces [5–13] but also the classical Euclidean case [14–20]. The accuracy and precision of results from an analysis are dependent on the convergence of the MCMC chain to the target distribution, making it critical to ensure that convergence has been achieved [21]. In the context of phylogenetic MCMC, the state space is the space of all possible different evolutionary histories, a treespace, along with tree-associated Euclidean parameters. Treespaces, being high-dimensional non-Euclidean spaces, pose the main and unique challenges in analysing phylogenetic data. The complexity of target distributions within this already complex space further adds to the intricacy of this problem [22–27]. Therefore, there is a need for effective methods to evaluate convergence of MCMC chains within the treespace to ensure the reliability of the resulting estimates and evolutionary conclusions.

The most popular tool for phylogenetic MCMC convergence assessment is Tracer [8]. It reads a log file containing a sample of trees and other parameters and produces effective sample size (ESS) estimates and trace plots for all Euclidean parameters associated with the analysis. However, convergence in lower dimensional parameter space does not imply convergence in the much higher dimensional treespace or vice versa [17]. It is therefore crucial to have an assessment of convergence in both the higher dimensional treespace and the Euclidean parameters that are associated with these trees. Generally, classical methods to assess convergence assume a Euclidean parameter space and can therefore not be applied to phylogenetic trees directly.

To address these issues, multiple different approaches to represent and analyse samples of trees as a Euclidean parameter trace have been developed as an aid to assess the tree parameter convergence. The approach developed by Nylander *et al.* [5] uses split frequencies and branch lengths to deploy classical convergence assessment tools on samples of trees. The average standard deviation of split frequencies (ASDSF) [28] is another popular convergence assessment tool based on splits, i.e. the partition of taxa into two subsets. Furthermore, Lanfear *et al.* [6] and Warren *et al.* [7] proposed a novel approach to estimate the ESS of a sample of trees by utilising a vector of tree distances as a continuous parameter trace, referred to as the pseudo ESS. Yet another recently developed approach by Fabreti & Höhna [11] converts trees into sequences based on the absence and presence of splits.

Despite these efforts, the problem of assessing convergence of the sample of

trees remains largely unsolved and overlooked in practise. A promising approach to solve this issue was discussed by Whidden & Matsen [26] in the form of a topological Gelman–Rubin diagnostic. The Gelman–Rubin diagnostic [16, 29], commonly known as the potential scale reduction factor (PSRF), is a widely used convergence assessment tool in MCMC analyses [18]. It estimates convergence by comparing the variance within and between multiple MCMC chains and examining the ratio of these variances, providing a quantitative measure to assess whether the chains have converged to the same target distribution. The *MrBayes* software [1] uses this diagnostic for continuous parameters, while the topological version developed in [26] is however not commonly used.

In this paper, we advance the idea of a tree Gelman–Rubin diagnostic (GR_T), utilising the recently developed ranked nearest neighbour interchange (RNNI) treespace [30–34]. Our approach is structured in two parts: firstly, we develop the GR_T diagnostic for samples of trees. In the context of MCMC, samples of trees should exhibit near-indistinguishability between independent chains if drawn from the same distribution over the treespace. Our presented tree PSRF value quantifies this property of similarity. Building upon this foundation, we integrate our GR_T diagnostic into a comprehensive convergence assessment tool that incorporates continuous parameter trace estimates as well. This Combination of components results in a powerful and robust convergence assessment diagnostic for phylogenetic MCMC analysis.

Our approach advances the state of the art in several different ways. Firstly, our approach is based on the properties of a metric treespace, with geometry based on local tree rearrangements, giving it a strong mathematical and statistical foundation [32]. Secondly, by utilising the first polynomial time computable tree rearrangement based distance [33], we overcome the previous limitations imposed by the computational complexity of such distances. Thirdly, our assessment stands out from all other currently used convergence assessments in that it is fully automated and eliminates the need for visual or manual inspection, resulting in improved reproducibility and further streamlining the automation process.

We structure our paper as follows: first, we introduce our definition of a PSRF value for trees and demonstrate how it serves as the foundation for our GR_T diagnostic. In the Validation section, we provide evidence of the efficacy of our GR_T diagnostic by showing its performance on samples of trees from a well-calibrated simulation study. We then incorporate continuous parameter traces in our convergence assessment tool. In the Results section, we present the performance of our tree-based GR_T diagnostic on the widely used benchmark data sets DS1–DS11. Subsequently, we demonstrate the performance of the more sophisticated version on the same data sets, followed by an application to three real datasets. In conclusion, we highlight the promising implications of the RNNI treespace and the diagnostic for advancing statistical tools in the analysis of samples from phylogenetic time trees.

2 Methods

In this section we introduce our Gelman-Rubin diagnostic tool, which we will refer to as the GR_T diagnostic in the remainder of the paper. We start by presenting a tree potential scale reduction factor (PSRF) value for 2 independent MCMC chains and its generalisation to N independent chains. We also provide a description of how to incorporate this tree PSRF value into our convergence assessment tool. We focus on the treespace based approach and will later in this paper combine it with continuous parameter estimates to form a comprehensive convergence diagnostic tool.

Pseudo-code for the individual algorithms presented is available in the Supplement Section S2 and implementations are available as a BEAST2 (Java) package [35] or as part of a Python package [36].

2.1 Gelman Rubin like diagnostic

For the remainder of the paper we will assume that the length of two (or multiple) chains is the same to simplify notations. We will also assume that if two (or multiple) independent chains sample similar trees for long enough the chain has reached its stationary distribution. This disregards the potential issue of pseudo-convergence, on which we touch upon in Discussion (Section 5).

The Gelman-Rubin diagnostic, sometimes also referred to as the PSRF [29], is one of the most popular tools for diagnosing convergence in MCMC algorithms [18]. The diagnostic is calculated as the ratio of the between-chain variance to the within-chain variance. If the chains have converged to the same stationary distribution, the PSRF should be close to 1.0, i.e. the between-chain variance should be matching the within-chain variance. If the PSRF is significantly greater than 1.0, it indicates that the chains have not converged to the same distribution, or that more iterations are needed to obtain a sample from the stationary distribution. Originally, a tolerance of 0.1 was proposed by Gelman & Rubin [29], meaning that PSRF values larger than 1.1 would be considered as not converged chains. This however has recently been questioned as much too high to yield reasonable estimates by Vats & Knudson [19].

Now we introduce a diagnostic tool, the GR_T diagnostic, for evaluating convergence and termination of phylogenetic MCMC tree inference using a phylogenetic tree metric. A similar approach was developed by Whidden & Matsen [26] using the NP-hard SPR distance, whereas we will be using the efficiently computable RNNI distance [33] that is similarly to SPR based on local tree rearrangements. The variance of a tree t in a set $\mathcal{T}$ is calculated as the sum of its squared RNNI distances to all trees in the set divided by the size of the set.

$$\mathrm{Var}(t)_{\mathcal{T}} = \frac{\Sigma_{t_i \in \mathcal{T}} d(t_i, t)^2}{|\mathcal{T}|}$$

The PSRF value for a tree in either set, given two sets of trees $\mathcal{T}_1$ and $\mathcal{T}_2$, is defined as the square root of the ratio of variances for the tree in the opposite

set of trees to the variance in its own set.

$$\mathrm{PSRF}(t|\mathcal{T}_1, \mathcal{T}_2) = \sqrt{\frac{\mathrm{Var}(t)_{\mathcal{T}_2}}{\mathrm{Var}(t)_{\mathcal{T}_1}}}, t \in \mathcal{T}_1$$

Akin to the originally proposed diagnostic we can interpret the numerator as the between chain variance and the denominator as the within chain variance. The overall PSRF value can be interpreted similarly to the originally proposed diagnostic with values close to 1.0 indicating indistinguishable tree sets and values much larger than 1.0 as different tree sets. In addition, our proposed PSRF value can also be less than 1.0 which is the case if the variance for a tree t in the other tree set is lower. This can be thought of as one tree set being "fully contained" within the perimeter of another tree set. Therefore, a convergence criterion in this case requires specifying a tolerance around the 1.0 value, which we will discuss below.

Generalisation for N independent chains The diagnostic presented by Gelman & Rubin [29] was generalised by Brooks & Gelman [16] to incorporate more than 2 independent chains. For N independent chains our definition of the between chain variance can be modified as follows

$$\widehat{\mathrm{Var}}(t)_{\mathcal{T}_1,\ldots,\mathcal{T}_N} = \frac{1}{|N|}\sum_{j=1}^{N}\frac{\sum_{t_i\in\mathcal{T}_j} d(t_i,t)^2}{|\mathcal{T}_j|}$$

With this new notation of the between chain variance the PSRF value would be computed as

$$\mathrm{PSRF}(t|\mathcal{T}_1,\ldots,\mathcal{T}_N) = \sqrt{\frac{\widehat{\mathrm{Var}}(t)_{\mathcal{T}_1,\ldots,\mathcal{T}_{i-1},\mathcal{T}_{i+1},\ldots,\mathcal{T}_N}}{\mathrm{Var}(t)_{\mathcal{T}_i}}}, t \in \mathcal{T}_i$$

Note that increasing the number of independent chains will significantly increase computational cost as the pairwise distances between all combinations of chains and trees need to be calculated (and stored), i.e. for N independent chains of length M the number of necessary pairwise tree distance computations will be $N\frac{M(M-1)}{2} + \binom{N}{2}M^2 = \binom{NM}{2}$.

Convergence assessment Here we present how we use the previously defined tree PSRF value in our GR_T diagnostic. One difference to the original diagnostic is that we define the value that is used in each iteration, normally the PSRF value, as an average of PSRF values which we will call the GR value in the following. This is because we found that simply assessing the PSRF value is not sufficient (see Supplement Subsection S3.2) and it is important to put each iteration into the context of the full sample of trees from both chains.

Given two MCMC chains and their respective samples of trees $\mathcal{T}_1$ and $\mathcal{T}_2$ up to the i-th sample, i.e. $|\mathcal{T}_1| = |\mathcal{T}_2| = i$, we define the GR value of tree $t_i^1 \in \mathcal{T}_1$ as

$$\mathrm{GR}(t_i^1|\mathcal{T}_1,\mathcal{T}_2) = \frac{1}{i}\sum_{s=0}^{i}\mathrm{PSRF}(t_s|\mathcal{T}_1,\mathcal{T}_2)$$

and analogously for tree $t_i^2 \in \mathcal{T}_2$. When calculating the GR value for sample i we have to newly calculate the PSRF value for each sample $1, \ldots, i$ with the tree set containing all of the i samples and cannot reuse previously calculated PSRF values. However, this does not imply a significant increase in complexity because the required distances can be stored.

Yet another difference lies within the decision to terminate the chain based on this value. In the original diagnostic, it is enough for the PSRF value to fall below the threshold close to 1.0 determined by the tolerance, typically 0.1 or less [19]. Based on experiments and applications we found that in our case a tolerance of 0.05 or less is appropriate, see also Supplement Subsection S3.1. In addition, we also require the chain to continuously sample trees within this threshold around 1.0 until a predetermined ESS estimate for these samples of trees is reached. To calculate this ESS value we use the RNNI distance and the median pseudo ESS variant that uses all trees in a sample as reference trees, see [10]. Moreover, we will refer to this median pseudo ESS using the RNNI distance as the pseudo ESS for the remainder of this paper.

Using these requirements incorporates the concept of tree burn-in within our diagnostic. That is, if a GR value of a tree in either chain is not within the specified tolerance around 1.0 the samples up to this iteration get added to the tree burn-in and subsequently discarded from the set of trees that our assessment returns. Therefore, the diagnostic will return a subset of trees that has a sufficient pseudo ESS estimate (based on a predetermined threshold) and each tree in either chain has GR value within the tolerance around 1.0.

2.2 The ASM package diagnostic

We now present additional tools and concepts that we use in our BEAST2 implementation [35] to enhance the robustness of the previously described diagnostic for practical application by incorporating continuous parameter traces. The first addition is an automated burn–in detection method, loosely inspired by Geweke [14], but to the best of our knowledge no such approach is used anywhere in practise. The second strategy that we incorporate is the concatenation of continuous parameter traces from the individual chains. We conclude the section with a high level description of the resulting comprehensive diagnostic.

Burn-in detection As previously highlighted, our approach draws inspiration from the Geweke diagnostic [14], nevertheless, to the best of our knowledge, our method is entirely novel and distinct from any existing practice. Given a continuous parameter trace we take its 75% tail to calculate a mean μ and standard deviation σ and define a confidence interval as $[\mu-\sigma, \mu+\sigma]$. The algorithm iterates over the whole trace with a sliding window of width $M = 10$ until the window's mean falls within this confidence interval. It then discards everything before this point as burn–in. In practise this diagnostic can be run for a set of specified parameters, which by default we set to be the posterior, likelihood, and prior. The overall burn–in is decided by the maximum value among all of

these traces. The value M for averaging is aiding in making this detection more robust to random outliers as well as temporarily stuck chains.

Concatenation of traces The second strategy we incorporate into our diagnostic involves using ESS thresholds for continuous parameter traces. Rather than applying these thresholds to individual chains, we concatenate the traces before estimating the ESS. This approach enables detecting when two chains are sampling a parameter from different distributions which will result in low ESS estimates of the concatenated trace. Conversely, when the chains sample parameters from the same distribution, the overall run time improves due to the doubling of the sample size. Similarly to the burn-in detection, we enforce this threshold only for the specified subset of priority parameters (posterior, likelihood, and prior by default).

So at the high level our algorithm can be described as follows.

Algorithm 1 High level pseudo code description of our diagnostic

```
tree burn-in ← 0
traces ← {posterior, likelihood, prior}
Start sampling the MCMC chains
while samples arrive from the chains do
    Calculate burn-in of all traces
    Discard burn-in temporarily (for this iteration of the while loop)
    Calculate the GR value of the new tree samples
    if GR of both tree samples is within tolerance around 1.0 then
        Discard tree burn-in from both chains
        Concatenate the samples of the two chains
        if ESS threshold of important traces and trees is reached then
            Stop sampling and return concatenated sample as output
            → Convergence diagnosed
    else
        Update tree burn-in to be set to the current iteration
```

2.3 Implementation

As previously mentioned, we have divided the content of this paper and our implementation into two distinct parts. The first part encompasses the purely tree based diagnostic, using the RNNI distance and is implemented into the time tree statistics python package *tetres* [36]. The main focus of this package is the implementation of statistical tools that are developed primarily over the RNNI treespace. In contrast, the second part is implemented within the BEAST2 framework as the ASM package [35] in the java language.

3 Validation

To validate our diagnostic, we conduct a well-calibrated simulation study. This section is again divided into two parts: firstly, we evaluate the treespace-based approach, and secondly, we repeat the evaluation for our comprehensive diagnostic implemented for BEAST2. For both diagnostics, we assess their performance on the well-calibrated simulation study, utilising regular ESS assessments, and conducting a coverage analysis.

3.1 Well calibrated simulation study

We performed a well calibrated simulation study [37] sampling 50 tip dates randomly from the interval 0 to 1. We use a coalescent tree prior with constant population size log-normal($\mu = -0.79, \sigma = 1.0$) distributed giving a mean of 0.75, an HKY model with kappa log-normal($\mu = 1, \sigma = 1.25$) distributed, and gamma rate heterogeneity with four categories with shape parameter exponentially distributed with mean = 1 and frequencies Dirichlet$(4, 4, 4, 4)$ distributed. Further, gamma is lower bounded by 0.1 [38]. We use a strict clock with uniform$(0, 1)$ prior. Sampling 100 instances from this distribution using MCMC in BEAST2 [4], we get a range of tree heights from 1.03 to 14.74 with mean = 2.81 (note that due to the tips being sampled from 0 to 1, the tree height is lower bounded by 1) and we get a range for the clock rate of 0.01 to 0.99 in our study. With these trees, we sample sequences of 1000 sites using the sequence generator in BEAST2.

3.2 Treespace based GR_T diagnostic

For the following results we used the implementation provided in [36] and MCMC analyses simulated as described in the previous section. This implies that we run our diagnostic on the full set of trees from an MCMC chain with sufficient burn-in discarded initially. For determining convergence, i.e. sufficiently many samples of trees with GR value close enough to 1.0, we used the pseudo ESS as described in the Supplement Subsection S2.1. The evidence provided in the first part of the section establishes the viability of our approach, serving as a proof of concept for our method.

Convergence We compare different settings of our convergence diagnostic on the calibrated simulation study. One of the selected parameters that we compare are different pseudo ESS thresholds for the subset of trees, in this case 200 and 500. We choose these two values because 200 is a general rule of thumb in phylogenetic MCMC analyses [6] and 500 is because of recent discoveries by Magee *et al.* [10] and Fabreti & Höhna [11] which both advocate for a more stringent value higher than 200. In addition, we compare 3 different sizes of an acceptable interval (tolerance) for the GR value: 0.01, 0.02, and 0.05.

As we described in Section 2.1 we consider a pair of two independent chains converged (orange in Figure 1) if we find a large enough subset of trees in both

chains that all have GR values within the respective tolerance around 1.0 and the pseudo ESS value of each subset of trees is at least the selected ESS threshold. If we reach the end of a chain without both of these conditions being met the two independent chains are considered not converged (0 and blue in Figure 1).

We find that for the set of 1000 trees that arise from our simulation setup with a tolerance of 0.05 determines convergence in every case, whereas the more strict tolerances of 0.02 and 0.01 do not always detect convergence. The fundamental concept that underlies well calibrated simulation studies is that an inference method employing the same model as the data's generation should have the capability to accurately retrieve the true parameters. Therefore, these studies if setup properly are considered to return converged samples and are generally used to evaluate new methods and models [37]. Despite this notion, the fact our diagnostic with strict tolerances does not detect convergence might imply that sampling only 1000 trees is sometimes insufficient even for such "nice" data. Therefore, in practice it might be useful to use a more strict tolerance to counter act premature convergence detection, particularly for more complex data sets.

Furthermore, it is important to highlight that the existing methods for assessing MCMC convergence [18] heavily rely on manual inspection guided by "expert" knowledge, which applies broadly to MCMC analyses in various fields, including phylogenetics. It is currently not possible to simulate scenarios that would demonstrate non-convergence without employing two chains on completely independent datasets. Therefore, we are constraint to use these well calibrated simulation studies as our baseline to state that our diagnostic works as expected on datasets that would generally be considered converged analyses. This leaves a negative control as a non trivial open question for future research.

ESS assessment As mentioned before, we did not use any continuous parameter trace ESS estimates to determine convergence but only used the pseudo ESS of the sample of trees using the RNNI distance. Here we check how the ESS estimates for all parameters behave when using our GR_T diagnostic. We find that in this well calibrated setting the ESS estimates are in a similar range as the chosen ESS threshold value. These findings are displayed in Figure 2, where we compare the full MCMC chain versus the subset of tree samples determined by our GR_T diagnostic using a threshold of 200 or 500 respectively.

Moreover, we find that in the case of these well calibrated simulations the pseudo ESS value seems to be a reasonable proxy for all continuous parameter ESS values after appropriately tuning the operator weights. Operators are used to propose new states for a corresponding parameter and the weight of an operator refers to its probability of being selected in each step of the MCMC. When a parameters x has low ESS at the end of an analysis it can be helpful to adjust the weights of operators to improve the sampling for this specific parameter. However, we emphasise that in a real scenario this strategy may not universally succeed due to mixing issues, the complexity of posteriors, or other reasons. This highlights one of the rationales behind our BEAST2 implemen-

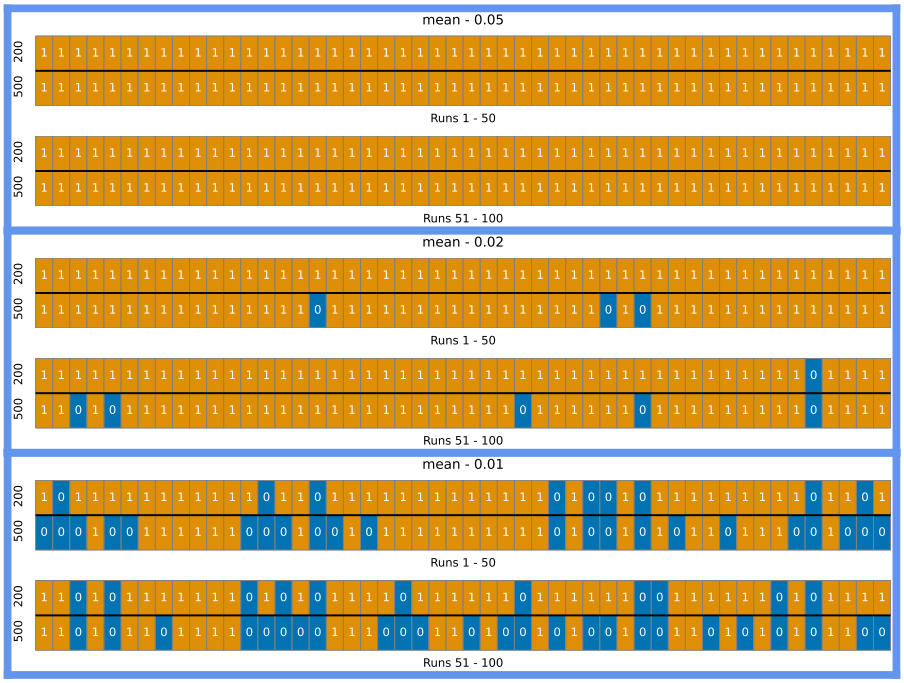

Figure 1: Our convergence diagnostic as implemented in tetres executed on 100 well calibrated simulation studies with 1000 trees each. Each rectangle represents one such simulation, a rectangle labelled 1 (orange) indicates a run diagnosed as convergent, a 0 (blue) indicates that no convergence is diagnosed within the sample. Each row corresponds to the three tolerance values (0.05, 0.02, 0.01). For each of the three different tolerance setups the 100 simulations are split in half, the first and last 50 runs, resulting in 2 heat maps for each tolerance. The "y" axis in each of the 6 heat maps is separating the two ESS threshold values (200 and 500).

tation's utilisation of ESS estimates for the relevant continuous parameter, as elaborated in Subsection 2.1.

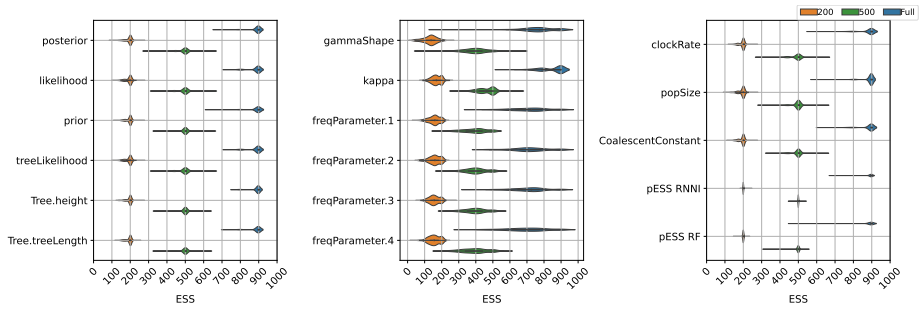

Figure 2: ESS comparison using our diagnostic with a tolerance of 0.05. The plot displays ESS of all parameters, including the pseudo ESS value using the RNNI and the Robinson Foulds (RF) distance, and compares the set of trees returned by a chain (Full) versus the automatically determined portions using our diagnostic with RNNI pseudo ESS threshold values of 200 and 500 respectively.

Coverage analysis tetres Here we preform a coverage analysis where we compare the full tree sets as sampled by the manually configured MCMC chain versus the subset of trees our diagnostic deems to be a convergent sample. The term "manually configured" denotes the best Bayesian phylogenetic practice of selecting the chain length, sampling intervals, and possibly burn–in before starting the MCMC analysis. A coverage analysis for Bayesian MCMC refers to an evaluation of the accuracy and reliability of the posterior inference obtained trough the MCMC sampling methods. By comparing the true parameter values of a simulation with the parameters from the full posteriors (full MCMC chain) and the partial posteriors (subset of the full posterior) produced by our GR_T diagnostic, we are able to assess the accuracy and uncertainty of our method versus a manually set up MCMC run. The coverage is considered appropriate if it implies that the true parameter falls within the intervals with the expected frequency of 95% of the time for a 95% confidence interval. This analysis helps to reveal any potential biases, under- or overestimation of uncertainties, allowing us to assess reliability of our diagnostic in comparison to a manually set up MCMC chain.

The results of this analysis are displayed in Table 1. For a set of 100 independent simulations the 95% highest posterior density (HPD) of the binomial distribution ranges from 91 to 99 inclusive. As indicated by the red values (outside the aforementioned interval) our automatically determined sets of trees perform almost identical to the full simulation setups. Despite the ESS threshold of 200 introducing low coverage in 3 additional parameters we would like to emphasise that this is combined with lower ESS estimates of these parameters

as visualised in Figure 2. Additionally, it is common in practice to disregard nuisance parameters, particularly when their ESS estimates are low. We touch upon this issue again in the next section with two ways to solve it in our BEAST2 implementation. Overall, this outcome gives us confidence that our GR_T diagnostic produces samples that are reliable estimates of the posterior and fully comparable to a manually set up MCMC analysis. As before the results are achieved with a GR tolerance of 0.05. More coverage analyses for different settings, including the use of the median instead of the mean for smoothing of the GR value, can be found in the Supplement, see Subsection S4.1.

	Expected range (91-99)		
Parameter	200	500	Full
Tree.height	93/93	96/96	96/96
Tree.treeLength	94/92	96/95	97/94
clockRate	91/91	90/92	91/93
gammaShape	87/90	93/93	92/92
popSize	90/93	97/96	96/97
CoalescentConstant	91/92	93/92	92/92
kappa	93/94	93/93	94/94
freqParameter.1	95/92	92/92	92/92
freqParameter.2	90/93	94/94	94/95
freqParameter.3	86/87	89/87	87/90
freqParameter.4	94/94	94/93	95/95

Table 1: Coverage of the true value by 95% HPD estimates from 100 independent simulations. Columns correspond to the ESS threshold of 200, 500, and the full sample of trees. The red coloured values indicate coverage outside the expected range between 91 and 99, the two values separated by '/' represent the two independent chains.

3.3 BEAST2 ASM package validation

In the previous section we used a manually set up simulation that had fully run all the predetermined iterations. For the validation of our automated stopping implementation we reran these analyses using our automated convergence diagnostic package ASM and performed a coverage analysis for the resulting samples of trees.

Coverage analysis ASM As we can see the coverage analysis is comparable to the previous table, slight difference are expected as MCMC is a random process. This result again confirms that our diagnostic produces reliable estimates of the posterior distribution.

A minor difference is the coverage of the frequency parameters, which do not quite reach the required 91-99 expected range. However, the frequency

	coverage	mean ESS	min ESS
Tree.height	95	224.70	57.56
Tree.treeLength	94	219.30	95.49
clockRate	91	222.71	66.45
gammaShape	89	124.27	16.35
popSize	91	222.54	95.05
kappa	92	178.38	86.39
freqParameter.1	91	181.88	51.34
freqParameter.2	94	178.08	52.26
freqParameter.3	88	176.19	44.24
freqParameter.4	96	182.90	41.98

Table 2: Coverage of parameters when executing a well calibrated simulation using the ASM package.

parameters in this specific case are nuisance parameters and can therefore be "ingnored". In addition, the table also displays the minimum ESS estimate of these parameters across the 100 simulations, see column min ESS in Table 2. For parameters with very low minimum ESS, such as the frequency parameters, a low coverage is expected because of insufficient number of samples. We want to emphasise that failing the coverage test is not a major concern as there are two ways to circumvent such a problem for any parameter in general using our tool. The first approach is that we can add any parameter in question to the list of important parameters in our diagnostic, which in turn would enforce a sufficiently high ESS estimate and increase the coverage. The second approach is to tune the operator weights in the analysis setup, which will also result in increased ESS estimates.

Furthermore, the inferred parameters of the sample are all highly correlated to the true values, this is visualised for the failed frequency parameter in Figure 3. This visualisation further indicates that there is no major concern about the low coverage parameters.

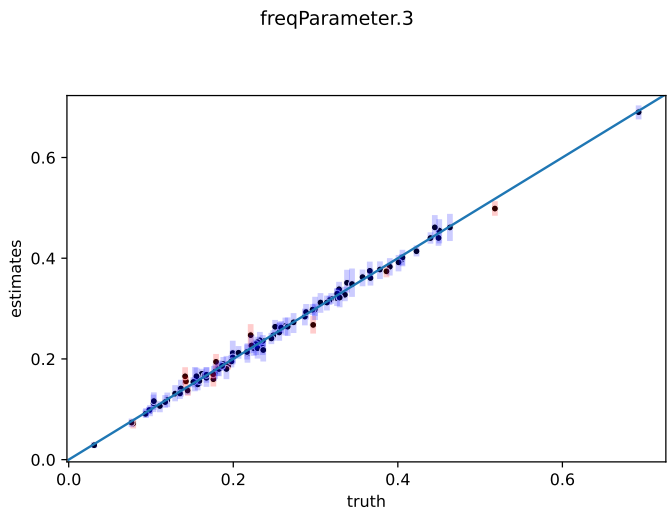

Figure 3: Visualising the coverage test for frequency parameter 3, the x axis displays the true parameter value and the y axis the inferred parameter. Blue boxes indicate that the true value is within the 95%HPD interval.

4 Results

In this section we present the results of applying our convergence diagnostic technique to real sequence data sets. We start with the datasets DS1–DS11, which are popular for benchmarking in computational phylogenetics [24–26, 28]. Previous findings on these suggest multi-modality in the treespace which implies that MCMC exploration and convergence within the treespace is challenging [24, 26]. We start by presenting the results of our tree based diagnostic first and find that they are coherent with these previously identified problems. Then, we present the results of applying our BEAST2 ASM package to the same datasets and find that the combination of continuous trace ESS estimates and our GR_T diagnostic perform well on these data sets. Additionally, we demonstrate the effectiveness of our package on three real datasets, revealing no major difficulties in identifying converged MCMC runs.

4.1 Convergence diagnostic on DS1–DS11

For each of these 11 dataset, we set up 6 different BEAST2 analyses to generate samples of fully resolved time trees. Subsequently, our convergence diagnostic as implemented in tetres was applied to these 66 sets of trees. Additionally, we present results based on our BEAST2 package ASM, yielding comparable results, which among other things serves as a validation of the two implementations correctness. Based on these findings, we investigated the number of tree samples required to meet the specified pseudo ESS threshold (in this case 200). This investigation provides further insights into the mixing quality of the tree samples, as measured by pseudo ESS.

Configurations We used the following 6 analysis configurations to generate two independent chains for each of the DS1–DS11 datasets.

- Config 1: 2.5 million samples, 2.5 thousand thinning, 0 pre-burn-in
- Config 2: 2.5 million samples, 2.5 thousand thinning, 1 million pre-burn-in
- Config 3: 5 million samples, 5 thousand thinning, 1 million pre-burn-in
- Config 4: 5 million samples, 5 thousand thinning, 2 million pre-burn-in
- Config 5: 10 million samples, 10 thousand thinning, 0 pre-burn-in
- Config 6: 10 million samples, 5 thousand thinning, 0 pre-burn-in

In Figure 4 we visualise the results of our diagnostic on these data sets. Indicated in blue are runs that, according to our diagnostic, have not reached convergence yet, and could therefore indicate certain convergence problems. The plot also contains the results of three different settings for the tolerance of our diagnostic. The diagnostic performs as expected with lower tolerances being less likely to reach convergence as it is stricter.

Whidden & Matsen [26] found that DS1, DS4, DS5, DS6, DS7, and DS10 are peaky, meaning they contain multiple regions in treespace with high posterior probability which are separated by regions of low posterior probability and therefore harder to sample. In Figure 4 it can be seen that our diagnostic agrees with these findings indicating that DS1, DS4, DS5, DS6, DS7 and DS10 are often unable to reach convergence in our setup. This indicates that these data sets have tricky to sample posterior distribution over treespace.

Our findings also show that DS2, DS3, DS8, DS9, and DS11 are reaching convergence in most of the setup runs. This indicates that the posterior distribution over treespace is easier to sample and hence reaches convergence with fewer samples. Again, this results is consistent with findings of Whidden & Matsen [26], who indicate that the distributions of DS2, DS3, and DS8 do not contain any peaks.

Figure 4: This plot is analogous to the previous Figure 1. Each row corresponds to a tolerance value (0.05, 0.02, 0.01). The 6 configurations for each data set are split in half, resulting in 2 heat maps for each tolerance. Each heat map has two rows which represent the ESS threshold choice (200 or 500). A rectangle labelled 1 (orange) indicates a run diagnosed as converged, a 0 (blue) indicates that no convergence is diagnosed within the range of the sample.

BEAST2 ASM package on DS1–DS11 We tested the BEAST2 implementation on the DS1–DS11 datasets ten times, utilising a strict clock with a fixed rate of 1, HKY substitution model, and Yule tree prior. The stopping criteria employed consist of GR_T statistic with a tolerance of 0.05, requiring pseudo ESS to exceed 200, and the ESS estimate of posterior, likelihood, and prior must also exceed 200. The table Table 3 presents the average chain length, with the log frequency set to 10 thousand. As a result, the average trace and tree log for DS1, for example, contains 246.1 samples.

Data set	Average chain length	Shortest	Longest	Longest/shortest
DS1	2461000	1.99M	3.73M	1.87
DS2	2378000	1.99M	3.28M	1.65
DS3	3435000	2.32M	4.36M	1.88
DS4	4235000	2.00M	9.83M	4.92
DS5	2454000	2.00M	3.41M	1.71
DS6	2902000	2.00M	3.51M	1.76
DS7	3797000	2.34M	4.67M	2.00
DS8	2601000	1.99M	3.53M	1.77
DS9	2697000	1.99M	4.35M	2.19
DS10	2551000	2.21M	3.34M	1.51
DS11	2579000	2.00M	3.57M	1.79

Table 3: Average chain length and fraction of longest divided by shortest chain length over 10 runs for data sets DS1–DS11. Lengths of the 10 individual runs are available in the Supplement Section S5.

We conducted visual inspection of the traces in Tracer [8] for all datasets, confirming convergence to the same posterior distribution. In all cases, there was a maximum difference of 2 between the lowest and highest log posterior, as well as prior and likelihood estimates across the ten runs, indicating that each run effectively converged to the same distribution. The inclusion of an ESS requirement of at least 200 for traces and trees leads to a minimum chain length of 200 samples (1.99 million steps), which, with a log frequency of 10 thousand, results in a minimum chain length of 2 million. The majority of chain lengths in the experiment were only slightly longer than 2 million, except for DS4, which exhibited two outliers just under 10 million (see Supplement Section S5). This variation can be attributed to the inherent stochastic nature of MCMC, causing differences in convergence times.

The presence of chains with "only" 200 samples suggest that the trace ESS criterion limits the stopping process, while the other criteria suggest that the MCMC could be terminated earlier. On the other hand, longer stopping times for some chains were accompanied by much higher ESS estimates for posterior, prior and likelihood, with the GR_T criterion being the primary factor in prolonging the sampling process. These findings indicate that neither criterion can be eliminated, as both are essential in determining the stopping point under

different circumstances or conversely form a bottleneck in certain situations.

Further investigating mixing of trees Due to the nature of MCMC analyses the resulting samples are not independent, hence the idea of an ESS measures to estimate the number of independent samples. Here we are further investigating this non-independence in treespace using our different configurations on the 11 data sets via the median pseudo ESS measure [6, 10]. Specifically, we are investigating how many samples are needed of a converged part of a chain to reach a desired pseudo ESS threshold,

$$\frac{ESST}{|\text{tree-sample}|}.$$

This value is closer to 1 if the MCMC thinning is appropriate, meaning that the samples are basically independent. However, smaller values indicate that the thinning interval is too small and therefore a larger number of samples is needed to reach the desired number of independent samples, i.e. the ESS threshold.

We visualised these findings in Figure 5 and it can be observed that the data sets DS1, DS4, and DS10 have more non-independent samples than others. In contrast, almost all samples in data sets DS2, DS3, and DS8 are independent (fraction close to 1). We also visualised these values for the samples resulting from our BEAST2 implementation in the ASM labelled column, which has very similar values among all datasets except DS3, DS4, and DS7.

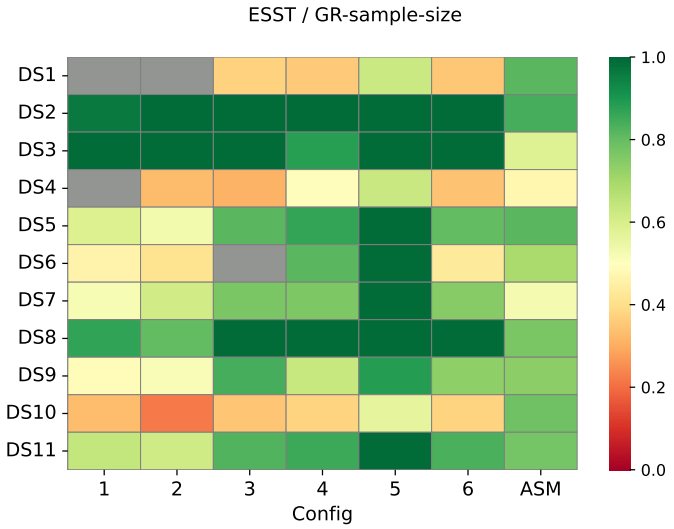

Figure 5: Fraction of ESS threshold over GR_T determined sample size. The grey cells indicate runs that are not considered converged. Dark green indicates that the sample size is the same as the ESS threshold (200 in this case). Accessible version available in the Supplement Figure S1.

We conclude from this investigation that the previously identified "peaky" datasets DS1, DS4, and DS10 display tree mixing that is coherent with these findings. However, we do not find stark differences among any of the datasets using our ASM package. Moreover, we do not find any scenarios that would suggest independent chains would be stuck in different modes. Rather, the datasets that are not diagnosed as converged have insufficient chain setups that sample too few iterations and trees. This results in a simple fix of running the chain longer rather than having to investigate and deal with multimodal tree distributions.

4.2 Real data application of convergence assessment

This section centres on the application of our implemented BEAST2 convergence diagnostic to various real datasets and diverse model combinations. We first analysed a dataset comprising of 63 hepatitis-C virus (HCV) sequences sampled in Egypt in 1993, previously analysed in [39] and [40]. To avoid committing to a specific site model, we used model averaging with bModelTest [41], a strict clock, and Yule skyline [42] tree prior providing a flexible multi epoch pure birth model.

Furthermore, we analysed 36 mammalian species [43] again with bModelTest site model, an optimised relaxed clock [44], and a Yule tree prior together with a number of node calibrations used in [43].

Lastly, we examined an alignment of 257 full COVID–19 genome sequences, a subset from [45], using four partitions as in the original study with a HKY model each, estimated frequencies and relative substitution rate for each partition. Further, we used a strict clock and the BICEPS [41] tree prior, which provides a flexible multi epoch coalescent model. Unlike the previous datasets, this analysis incorporated dated tips.

Each analysis was run ten times with stopping criterion the same as in Section 4.1, and a log frequency of 10 thousand samples. For each data set, we verified convergence to the same parameter distribution across all chains through visual inspection using Tracer.

In the case of the HCV data, the shortest run took 4.65 millions steps and the longest extended just beyond 19.49 million steps, resulting in a ratio of 4.19 between the longest and shortest run. Similarly, the primate data exhibited a shortest chain length of 2.03 million steps and a longest of 11.81 million, yielding a ratio of 5.82. Chains on the COVID–19 data spanned from 45.6 to 148 million steps, returning a slightly smaller ratio of 3.25.

The notable disparity in convergence rates between fast and slow chains in real data may suggest an approach of starting multiple chains and selecting the fastest to reach convergence in order to accelerate analyses. However, it remains uncertain whether slow converging instances inefficiently sample the state space or merely become trapped in unexplored modes avoided by faster converging chains. Further investigation is required to establish whether this strategy will be suitable.

Table 4: Average, shortest, and longest MCMC chain lengths over 10 runs before automatically stopping. Lengths of all 10 runs can be found in the Supplement Table S7

	HCV	Primates	COVID–19
average	1474500	4340000	77364000
shortest	4650000	2030000	45630000
longest	19490000	11810000	148170000
ratio	4.19	5.82	3.25

5 Discussion

This paper presents a diagnostic tool for the automated assessment of convergence in phylogenetic MCMC analyses, utilising the RNNI treespace and its associated metric. Through an extensive evaluation on well-calibrated simulation studies, we demonstrate the efficacy of our diagnostic. We split the evaluation of our method into two separate parts – first, we demonstrate that our diagnostic tool is capable of producing reliable estimates of the posterior when applied only to samples of trees (without any use of continuous parameters). Then, we extend this tool to incorporate continuous parameter traces. This comprehensive approach enables automatic convergence assessment of all parameters of interest in typical phylogenetic applications. We demonstrate this by applying our method to multiple real data sets.

This approach is a further advance in the area of developing statistical methods over geometrically complex spaces, such as the treespace, which has direct practical implications for phylogenetic MCMC analyses. By eliminating manual visualisation-based steps in setting up MCMC chains, our tool makes these analyses more accessible, reduces variability in the setup process, and contributes improving reproducibility of results obtained using phylogenetic MCMC analysis. In addition, automatic convergence diagnostic is a major stepping stone to enable online phylogenetic inference algorithms [46]. Yet to be explored applications of our tool include comparing tree samples inferred using different evolutionary models and advancing the "convergence" assessment and setup process when using nested sampling approaches [47].

While evaluating and developing this diagnostic we made a few peculiar observations that we believe are worth reporting. First, as we demonstrate on our well calibrated simulation study, a sample of 1000 trees may not consistently guarantee convergence, particularly when tolerance is low. Second, as we analysed DS1–DS11 we found that the convergence assessment was coherent with previous findings on these data sets (specifically multimodality issues implying hard to sample posteriors). However we did not find cases of independent chains being stuck in different modes. All our unconverged runs can be attributed to chains being set up with inadequate lengths; prolonging the chain's runtime would lead to convergence being diagnosed by our diagnostic as it can be seen in Section 4.1.

Another aspect important in practice is the use of ESS estimates, specifically in cases when the log frequency is high and every sample is independent of the previous one. The effective sample size and actual sample size will be very close to each other and in these situations ESS estimates are unstable. We observed cases where the ESS would double by adding a single sample, hence ESS based criteria work better when the number of independent samples is somewhat smaller than the chain length.

We finish this paper by outlining possible future directions as we see them and highlighting open problems that are yet to be addressed in the context of phylogenetic MCMC analyses and convergence assessment. Like we mentioned before, the sample of trees is the most important output of a phylogenetic

MCMC analysis and should therefore be considered as such, especially when assessing convergence. However, it is an unsolved problem whether convergence in the treespace also implies convergence in other low dimensional parameter spaces. Our results in the well calibrated simulation setting hint to this conjecture being true, but it would require major effort to show that this is indeed the case for any real scenario. Therefore, in practice it is still important to use more than one check for convergence which is why our implementation incorporates the use of continuous parameter traces in addition to the treespace based approach.

Another direction of ongoing research into probability distributions over treespaces is related to tree islands or terraces [22–24, 27]. This can also be thought of as multi-modal distributions, where high posterior probability regions in the treespace are separated by valleys of low posterior probability. It has previously been shown [24, 26] that some of the data sets DS1–DS11 suffer from this phenomenon of multi modality, which can lead to false conclusions about an analysis. Our diagnostic is able to detect such problems in these case where independent chains are stuck in different parts of the treespace (see Supplement Subsection S3.1 for more detail).

A problem that is related to multi-modality is pseudo convergence, which is the phenomenon where an MCMC chain is apparently sampling from the stationary distribution but in reality it is stuck somewhere that is not in fact the stationary distribution. This can be the case when it takes much longer than the length of the chain to transition between this pseudo stationary part and the true stationary distribution. For example, when this may happen when the stationary distribution is multi-modal. However, multi-modality does not imply pseudo convergence as sometimes the chain is able to transition between modes, or cross the valleys, and sample from the true stationary multimodal distribution. In fact, in phylogenetics multi-modality is a poorly explored phenomenon, due to geometric complexity of the treespace, with even the notion of a mode being somewhat esoteric. Altogether, the phenomena of multi-modality and pseudo convergence form a compounded challenge.

These problems can potentially cause our diagnostic tools to converge to inaccurate conclusions. For example, if all independent chains pseudo-converged to the same mode, it is fundamentally impossible for a distance based method to diagnose pseudo-convergence. This and the problem of identifying multi-modal tree distributions remain open problems that need to be carefully addressed to deepen our understanding of distributions over the treespace. In this context, the approach developed in this paper provides additional insight into the black box that is phylogenetic MCMC.

Data and software availability

Data used in the experiments can be accessed through `https://github.com/rbouckaert/asm/releases/tag/v0.0.1`. ASM is a BEAST 2 package that provides a user friendly GUI to the methods introduced here. It is open source

and available at `https://github.com/rbouckaert/asm`. The Python package tetres provides an interface to the RNNI based methods introduced here. It is open source and available at `https://github.com/bioDS/tetres`.

Acknowledgement

This work was partially supported by Royal Society Te Apārangi through a Rutherford Discovery Fellowship (UOC1702) and a Marsden grant (21-UOC-057), and by Ministry of Business, Innovation, and Employment of New Zealand through an Endeavour Smart Ideas grant (UOOX1912) and a Data Science Programmes grant (UOAX1932).

We thank Luiz Carvalho and the members of the bioDS lab for their feedback and discussion of the results.

References

1. Huelsenbeck, J. P. & Ronquist, F. MRBAYES: Bayesian inference of phylogenetic trees. *Bioinformatics* **17,** 754–755 (2001).
2. Höhna, S. *et al.* RevBayes: Bayesian phylogenetic inference using graphical models and an interactive model-specification language. *Systematic biology* **65,** 726–736 (2016).
3. Suchard, M. A. *et al.* Bayesian phylogenetic and phylodynamic data integration using BEAST 1.10. *Virus evolution* **4,** vey016 (2018).
4. Bouckaert, R. *et al.* BEAST 2.5: An advanced software platform for Bayesian evolutionary analysis. *PLoS computational biology* **15,** e1006650 (2019).
5. Nylander, J. A., Wilgenbusch, J. C., Warren, D. L. & Swofford, D. L. AWTY (Are We There yet?): A System for Graphical Exploration of MCMC Convergence in Bayesian Phylogenetics. *Bioinformatics* **24,** 581–583. ISSN: 1367-4803. `https://doi.org/10.1093/bioinformatics/btm388` (2022) (Feb. 2008).
6. Lanfear, R., Hua, X. & Warren, D. L. Estimating the Effective Sample Size of Tree Topologies from Bayesian Phylogenetic Analyses. *Genome Biology and Evolution* **8,** 2319–2332. ISSN: 1759-6653. `https://doi.org/10.1093/gbe/evw171` (2022) (Aug. 2016).
7. Warren, D. L., Geneva, A. J. & Lanfear, R. RWTY (R We There Yet): An R Package for Examining Convergence of Bayesian Phylogenetic Analyses. *Molecular Biology and Evolution* **34,** 1016–1020. ISSN: 0737-4038. `https://doi.org/10.1093/molbev/msw279` (2022) (Apr. 2017).
8. Rambaut, A., Drummond, A. J., Xie, D., Baele, G. & Suchard, M. A. Posterior Summarization in Bayesian Phylogenetics Using Tracer 1.7. *Systematic Biology* **67,** 901–904. ISSN: 1063-5157. `https://doi.org/10.1093/sysbio/syy032` (2022) (Sept. 2018).
9. Harrington, S. M., Wishingrad, V. & Thomson, R. C. Properties of Markov Chain Monte Carlo Performance across Many Empirical Alignments. *Molecular Biology and Evolution* **38,** 1627–1640. ISSN: 1537-1719. `https://doi.org/10.1093/molbev/msaa295` (2022) (Apr. 2021).
10. Magee, A. F., Karcher, M. D., Matsen IV, F. A. & Minin, V. N. *How Trustworthy Is Your Tree? Bayesian Phylogenetic Effective Sample Size through the Lens of Monte Carlo Error* Sept. 2021. arXiv: `2109.07629 [q-bio, stat]`. `http://arxiv.org/abs/2109.07629` (2022).
11. Fabreti, L. G. & Höhna, S. Convergence Assessment for Bayesian Phylogenetic Analysis Using MCMC Simulation. *Methods in Ecology and Evolution* **13,** 77–90. ISSN: 2041-210X. `https://onlinelibrary.wiley.com/doi/abs/10.1111/2041-210X.13727` (2022) (2022).
12. Haag, J., Höhler, D., Bettisworth, B. & Stamatakis, A. From Easy to Hopeless—Predicting the Difficulty of Phylogenetic Analyses. *Molecular Biology and Evolution* **39,** msac254 (2022).

13. Kelly, L. J., Ryder, R. J. & Clarté, G. Lagged couplings diagnose Markov chain Monte Carlo phylogenetic inference. *The Annals of Applied Statistics* **17,** 1419–1443 (2023).

14. Geweke, J. F. *Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments* tech. rep. 148 (Federal Reserve Bank of Minneapolis, 1991). `https://ideas.repec.org/p/fip/fedmsr/148.html` (2022).

15. Brooks, S. P. & Roberts, G. O. Assessing convergence of Markov chain Monte Carlo algorithms. *Statistics and Computing* **8,** 319–335 (1998).

16. Brooks, S. P. & Gelman, A. General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics* **7,** 434–455. ISSN: 1061-8600. `https://www.tandfonline.com/doi/abs/10.1080/10618600.1998.10474787` (2022) (Dec. 1998).

17. Gill, J. Is partial-dimension convergence a problem for inferences from MCMC algorithms? *Political Analysis* **16,** 153–178 (2008).

18. Roy, V. Convergence diagnostics for markov chain monte carlo. *Annual Review of Statistics and Its Application* **7,** 387–412 (2020).

19. Vats, D. & Knudson, C. Revisiting the Gelman–Rubin Diagnostic. *Statistical Science* **36,** 518–529. ISSN: 0883-4237, 2168-8745. `https://projecteuclid.org/journals/statistical-science/volume-36/issue-4/Revisiting-the-GelmanRubin-Diagnostic/10.1214/20-STS812.full` (2022) (Nov. 2021).

20. Vehtari, A., Gelman, A., Simpson, D., Carpenter, B. & Bürkner, P.-C. Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of MCMC (with discussion). *Bayesian analysis* **16,** 667–718 (2021).

21. Brooks, S., Gelman, A., Jones, G. & Meng, X.-L. *Handbook of markov chain monte carlo* (CRC press, 2011).

22. Hillis, D. M., Heath, T. A. & John, K. S. Analysis and visualization of tree space. *Systematic biology* **54,** 471–482 (2005).

23. Sanderson, M. J., McMahon, M. M. & Steel, M. Terraces in phylogenetic tree space. *Science* **333,** 448–450 (2011).

24. Höhna, S. & Drummond, A. J. Guided tree topology proposals for Bayesian phylogenetic inference. *Systematic biology* **61,** 1–11 (2012).

25. Larget, B. The estimation of tree posterior probabilities using conditional clade probability distributions. *Systematic biology* **62,** 501–511 (2013).

26. Whidden, C. & Matsen IV, F. A. Quantifying MCMC Exploration of Phylogenetic Tree Space. *Systematic Biology* **64,** 472–491. ISSN: 1063-5157. `https://doi.org/10.1093/sysbio/syv006` (2022) (May 2015).

27. Jombart, T., Kendall, M., Almagro-Garcia, J. & Colijn, C. treespace: Statistical exploration of landscapes of phylogenetic trees. *Molecular ecology resources* **17,** 1385–1392 (2017).

28. Lakner, C., Van Der Mark, P., Huelsenbeck, J. P., Larget, B. & Ronquist, F. Efficiency of Markov chain Monte Carlo tree proposals in Bayesian phylogenetics. *Systematic biology* **57,** 86–103 (2008).

29. Gelman, A. & Rubin, D. B. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science* **7,** 457–472. ISSN: 0883-4237. `https://www.jstor.org/stable/2246093` (2022) (1992).

30. Gavryushkin, A. & Drummond, A. J. The space of ultrametric phylogenetic trees. *Journal of theoretical biology* **403,** 197–208 (2016).

31. Gavryushkin, A., Whidden, C. & Matsen IV, F. A. The combinatorics of discrete time-trees: theory and open problems. *Journal of mathematical biology* **76,** 1101–1121 (2018).

32. Collienne, L., Elmes, K., Fischer, M., Bryant, D. & Gavryushkin, A. Discrete coalescent trees. *Journal of Mathematical Biology* **83,** 60 (2021).

33. Collienne, L. & Gavryushkin, A. Computing nearest neighbour interchange distances between ranked phylogenetic trees. *Journal of Mathematical Biology* **82,** 8 (2021).

34. Berling, L., Collienne, L. & Gavryushkin, A. Estimating the mean in the space of ranked phylogenetic trees. *bioRxiv,* 2023–05 (2023).

35. Bouckaert, R. *BEAST2 ASM pacakge* `https://github.com/rbouckaert/asm/`.

36. Berling, L. *Time tree statistics python package* `https://github.com/biods/tetres/`.

37. Bouckaert, R., Carvalho, L. M. F., Mendez, F. H. K. & Swanepoel, C. *BEAST2 developer manual* `https://github.com/rbouckaert/DeveloperManual/`.

38. Bouckaert, R. R. OBAMA: OBAMA for Bayesian amino-acid model averaging. *PeerJ* **8,** e9460 (2020).

39. Drummond, A. J., Rambaut, A., Shapiro, B. & Pybus, O. G. Bayesian coalescent inference of past population dynamics from molecular sequences. *Molecular biology and evolution* **22,** 1185–1192 (2005).

40. Stadler, T., Kühnert, D., Bonhoeffer, S. & Drummond, A. J. Birth–death skyline plot reveals temporal changes of epidemic spread in HIV and hepatitis C virus (HCV). *Proceedings of the National Academy of Sciences* **110,** 228–233 (2013).

41. Bouckaert, R. R. & Drummond, A. J. bModelTest: Bayesian phylogenetic site model averaging and model comparison. *BMC evolutionary biology* **17,** 1–11 (2017).

42. Bouckaert, R. R. An efficient coalescent epoch model for Bayesian phylogenetic inference. *Systematic Biology* **71,** 1549–1560 (2022).

43. Yang, Z. & Yoder, A. D. Comparison of likelihood and Bayesian methods for estimating divergence times using multiple gene loci and calibration points, with application to a radiation of cute-looking mouse lemur species. *Systematic biology* **52,** 705–716 (2003).

44. Douglas, J., Zhang, R. & Bouckaert, R. Adaptive dating and fast proposals: Revisiting the phylogenetic relaxed clock model. *PLoS computational biology* **17,** e1008322 (2021).

45. Douglas, J. *et al.* Phylodynamics reveals the role of human travel and contact tracing in controlling the first wave of COVID-19 in four island nations. *Virus evolution* **7,** veab052 (2021).

46. Bouckaert, R., Collienne, L. & Gavryushkin, A. *Online Bayesian Analysis with BEAST 2* May 2022. `https://www.biorxiv.org/content/10.1101/2022.05.03.490538v1` (2022).

47. Russel, P. M., Brewer, B. J., Klaere, S. & Bouckaert, R. R. Model selection and parameter inference in phylogenetics using nested sampling. *Systematic biology* **68,** 219–233 (2019).

S1 Accessible data visualisation

A accessible version of tree mixing plot Figure 5.

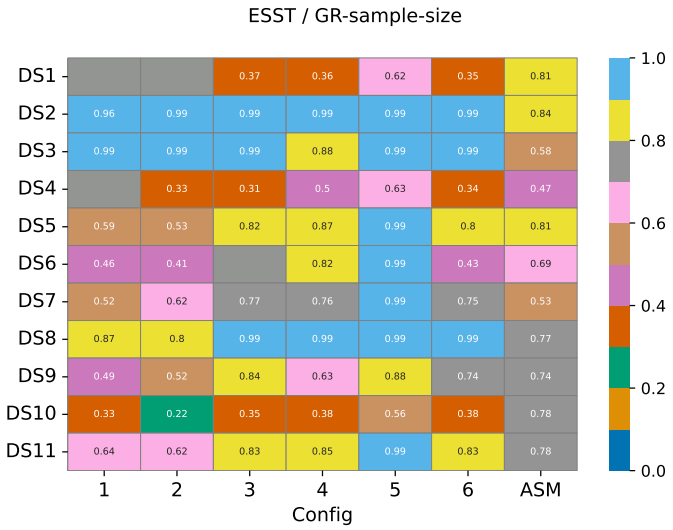

Figure S1: Accessible version of Figure 5. Labels in each rectangle display the value of the fraction.

S2 Pseudo Code of presented algorithms

In this section, we present the pseudo code for essential algorithms utilized in our diagnostic. We begin by introducing all the preliminary algorithms and conclude with a comprehensive pseudo code of our diagnostic, offering more algorithmic depth than the version presented in the main paper.

S2.1 An effective sample size measure for a set of trees

In this manuscript we use the following definition of pseudo ESS [6, 10], which is the computation of an ESS measure for a sample of trees given a tree metric. We first reduce the set of trees to a pairwise tree distance matrix, with $D_{i,j} = d_{RNNI}(t_i, t_j)$ being the distance between the trees t_i and t_j and D_i denoting the i-th row of this matrix. In our experience, considering a single focal tree leads to pseudo ESS estimates that are rather temperamental, sometimes getting stuck at low values depending on the choice of focal tree. Therefore, we consider all trees in the set as focal tree. The pseudo ESS is then defined as the median over the ESS estimate of every row.

$$\text{pESS}(\mathcal{T}) = \text{median}(\{\text{ESS}(D_i | \forall t_i \in \mathcal{T})\}) \tag{1}$$

We can use this as a criterium for convergence by making sure pESS($\mathcal{T}$) is at least as high as the target ESS. Algorithm S1 provides more details.

Algorithm S1 Pseudo tree ESS calculation,

```
procedure PSEUDO-ESS(𝒯)
    ▷ This function needs access to the distance matrices
    ESS–list ← ∅
    for all t ∈ 𝒯 do
        distances ← {d_RNNI(t, s), ∀s ∈ 𝒯}
        ▷ ESS is a continuous parameter ESS estimator
        ESS–list.add(ESS(distances))
    if median(ESS–list) ≥ target ESS then
        return True
    else
        return False                          ▷ No convergence diagnosed
```

S2.2 An effective sample size measure for a traces

A common convergence check is to make sure ESSs for traces of interest, in particular posterior, prior and likelihood, exceed a certain threshold value. Algorithm S2 shows how this can be automated as a convergence criterion.

Algorithm S2 Trace ESS calculation

```
procedure TRACE-ESS(traces)
    ▷ This function needs access to the traces
    for all t ∈ traces do
        seq ← remove burn-in from trace 1 & 2 and concatenate
        ▷ ESS is a continuous parameter ESS estimator
        ESS–list.add(ESS(seq))
    if minimum(ESS–list) ≥ target ESS then
        return True
    else
        return False                                  ▷ No convergence diagnosed
```

S2.3 Tree PSRF calculation

The following Algorithm S3 presents the calculation of the PSRF value for trees. Its input consists of two sets of trees $\mathcal{T}_1, \mathcal{T}_2$ and an index t that is the index of a sample in $\mathcal{T}_1$. It returns the PSRF value for the tree in $\mathcal{T}_1$ at the given index t.

Algorithm S3 PSRF calculation for a tree

```
procedure PSRF(T^1, T^2, t)
    ▷ This function needs access to the distance matrices
    ▷ assume k is the length of both treesets, t is index of tree
    v_in ← Σ_{i=1}^{N} d(T_t^1, T_i^1)^2
    v_bt ← Σ_{i=1}^{N} d(T_t^1, T_i^2)^2
    PSRF_t ← sqrt(v_bt / v_in)
    return PSRF_t
```

S2.4 GR_T diagnostic

This allows us to formulate a stopping criterion presented in Algorithm S4 based on the comprehensive GR_T diagnostic that we developed. While, for simplicity, we previously excluded the possibility of traces having varying lengths, it is noteworthy that in practice, traces can indeed have different lengths, and burn–in periods between chains can vary significantly. This practical consideration adds to the algorithms' complexity, leading us to omit it from the pseudo code to enhance readability.

Algorithm S4 Automated convergence assessment

```
procedure GR–DIAGNOSTIC(ESS threshold, tolerance)
    ▷ Assuming N is the length of the individual chains
    psrf_1 ← mean {PSRF(trees_1, trees_2, x) ∀x ∈ {trees_1}}
    psrf_2 ← mean {PSRF(trees_2, trees_1, x) ∀x ∈ {trees_2}}
    if both psrf_{1,2} within tolerance around 1.0 then
        return True                                  ▷ Convergence diagnosed
    else
        tree burn–in ← N
    return False                                     ▷ No convergence diagnosed
```

S2.5 Burn–in detection

The following Algorithm S5 estimates the burn-in for a trace of continuous parameters. Its input is a trace $x_1, x_2, \ldots, x_n$ of length n and it returns the number of samples that should be discarded as burn–in from the beginning of this trace.

Algorithm S5 Burn–in detection

```
procedure DETECT–BURN(x_1, x_2, ..., x_n)
    M ← 10                                           ▷ Moving average
    μ ← mean(x_⌊0.75*n⌋, ..., x_n)
    v ← std(x_⌊0.75*n⌋, ..., x_n)
    accept ← [μ − v, μ + v]
    for i ∈ [x_M, ..., x_⌊0.75*n⌋] do
        if mean({x_{i−M}, ..., x_i}) ∈ accept then
            return i
    return ⌊0.75 * n⌋
```

S2.6 Updating matrices

This is an auxiliary algorithm for the GR_T diagnostic, and is the main additional computational cost at every sampling step of the MCMC. Adding a tree to a distance matrix implies that the distance calculation with complexity n^2, where n is the number of taxa, has to be executed N times, with N being the number of samples of the MCMC that have been stored so far (excluding this newest sample). In addition to adding the trees from each chain to its respective distance matrix we also need to keep track of the in–between distances, which adds additional N distance calculations for either tree. Therefore, the added complexity of our diagnostic at every iteration is $O((4N + 1) * n^2)$.

In practice it is feasible to impose an upper limit on the size of the stored distance matrices. Upon reaching this user-defined limit, every second row and column is removed from the matrix, and not further considered in pseudo-

ESS and PSRF calculations (Algorithms S1 and S3). Moreover, future updates exclusively apply to even–numbered trees, with odd numbered trees disregarded. Upon reaching this size limitation k times, only trees numbered i are considered, where i is divisible by 2^k.

Algorithm S6 Updating distance matrices

```
procedure UPDATE.MATRICES(t^1, t^2, D^1, D^2, D^{1,2})
    D^1 add t^1 column
    D^2 add t^2 column
    D^{1,2} add t^1 row and t^2 column
    if sizes of D^1, D^2, D^{1,2} reach limitation then
        delete every second row and column from D^1, D^2, D^{1,2}
```

S2.7 Automatically stopping MCMC

The complete automatically stopping MCMC runs two independent MCMC chains, as well as a complete diagnostic tool (Methods section) executed in parallel with the MCMC analyses. The algorithm, detailed in Algorithm S7, is invoked for each iteration of the sampling process to determine the convergence status of the chains. This assessment considers the number of samples, ESS estimates for continuous parameters and trees, and the PSRF values.

Algorithm S7 Have the chains converged?

```
procedure CONVERGED(ESS threshold, tolerance)
    D^1, D^2, D^{1,2}          ▷ Distance matrices for chains, including in–between
distances
    traces = [posterior, likelihood, prior]          ▷ Important parameters
    tree burn–in ← 0
    while (t^1, t^2) arrive from sampling do
        UPDATE.MATRICES(t^1, t^2, D^1, D^2, D^{1,2})
        for i ∈ {1, 2} do
            burn_i ← max_traces {DETECT-BURN(chain_i(trace))}
            Discard trace and tree burn-in temporarily for the next calculations
            ▷ This affects distances matrices and continuous traces
        if GR–DIAGNOSTIC(ESS threshold, tolerance) == True
            and PSEUDO-ESS(ESS threshold, tolerance) == True
            and TRACE-ESS(ESS threshold, tolerance) == True then
                Stop sampling
                return Concatenated samples without burn-in and tree burn-in
```

S3 GR value related decisions

In this section we provide insight into why we choose the tolerance of the GR value to be 0.05 or lower. In addition we showcase why the "smoothing" or rather the use of an average function for the GR value is essential.

S3.1 Choice of tolerance

The tolerance of 0.05 is chosen based on the following result, which despite only showing one specific experiment was done for multiple different analyses and datasets. The following plots display the GR value for a fully run analysis, i.e. two independent chains sampled the tree sets $\mathcal{T}_1 = \{t_1^1, \ldots, t_N^1\}$ and $\mathcal{T}_2 = \{t_1^2, \ldots, t_N^2\}$ then the values plotted on the y axis are

$$\mathrm{GR}(t_i^1|\mathcal{T}_1, \mathcal{T}_2) = \frac{1}{N}\sum_{s=0}^{N} \mathrm{PSRF}(t_s|\mathcal{T}_1, \mathcal{T}_2)$$

and

$$\mathrm{GR}(t_i^2|\mathcal{T}_2, \mathcal{T}_1) = \frac{1}{N}\sum_{s=0}^{N} \mathrm{PSRF}(t_s|\mathcal{T}_2, \mathcal{T}_1)$$

for each sample i (x axis).

Such an analysis is visualised in Figure S3 were it is visible that the GR values for almost all samples fall into the interval $(0.98, 1.02)$. The density, displayed in the right plot is very high and showcases how narrow this interval is. In Figure S3, on the other hand, we visualised the same values but using 2 samples of trees from different simulations. It can be observed that these GR values are much higher, have very big variances, and the two samples of trees have very distinct values. This gives us confidence that, due to its sensitivity, our diagnostic is accurately diagnosing convergence problems in treespace. These convergence problems include slowly mixing chains, chains that are stuck or sampling different regions of treespace, or simply chains that are not sampling similar enough trees.

S3.2 Why smoothing is necessary?

Here we present why smoothing of the PSRF values, which results in the GR value, is necessary for our diagnostic. The smoothing takes place whenever the GR value is calculated and we compared using the mean

$$\mathrm{GR}(t_i^1|\mathcal{T}_1, \mathcal{T}_2) = \frac{1}{i}\sum_{s=1}^{i} \mathrm{PSRF}(t_s|\mathcal{T}_1, \mathcal{T}_2)$$

or median

$$\mathrm{GR}(t_i^1|\mathcal{T}_1, \mathcal{T}_2) = \mathrm{median}\left(\{\mathrm{PSRF}(t_s|\mathcal{T}_1, \mathcal{T}_2) \forall s \in \{1, \ldots, i\}\}\right)$$

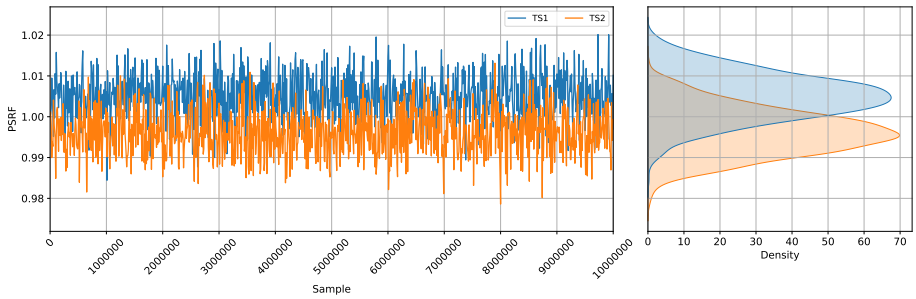

Figure S2: GR value for two independent chains on the same dataset. The left part of the plot displays these values for each sample in the chain and the plot on the right visualises the density of these values for each chain. Notably, the density of these values is very high and the variance around the value 1 is very low.

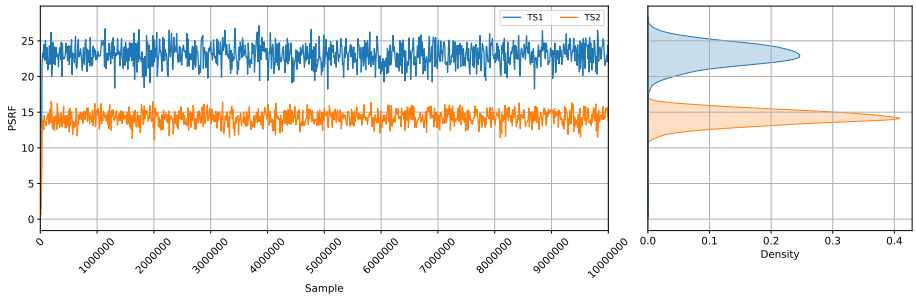

Figure S3: GR value for two independent chains on different datasets. The left displays the GR values for both chains and the right visualises the density for both chains. Unlike the previous plot there is a very wide gap between the two chains and the absolute value is comparably very high, the variance is also very high.

In Figure S4 we display the PSRF values without smoothing, i.e. at every iteration i (x axis) the value $\text{PSRF}(t_i|\mathcal{T}_1, \mathcal{T}_2)$ is plotted (y axis). It can be seen that these values have a very high variance and absolute value when compared to the previously presented GR values in Figure S2.

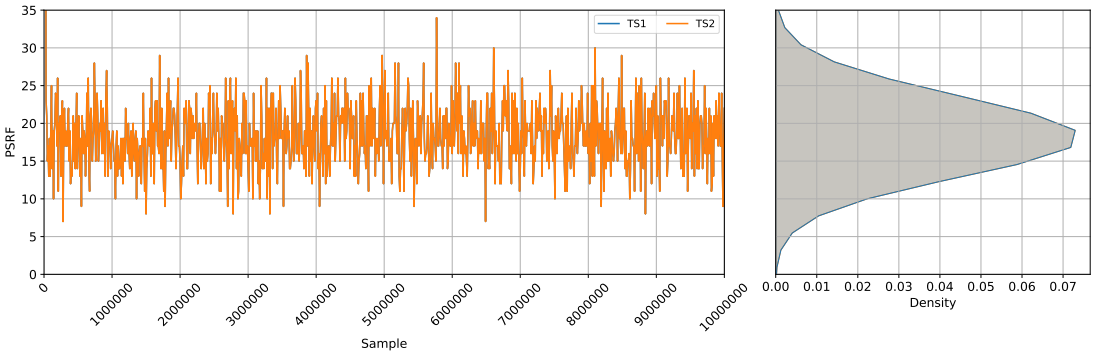

Figure S4: Displaying only the PSRF value at each iteration on the x axis. These have very high absolute values and a large variance (y axis).

S4 Validation

In this section we present further results on validating our method on well calibrated simulation studies. We compare the use of the median smoothing instead of the mean. Moreover, we provide more results in situation with different GR tolerances. We also present more coverage tables of different analyses.

S4.1 Comparing mean and median convergence

In Figure S5 the mean (left column) average is compared to the median (right column) average to smooth out the PSRF values to form the GR value. There are no significantly big differences among the two averages, the median is as expected a slightly more conservative measurement. Therefore a couple of runs that are diagnosed as converged under the mean average are not converged when using the median average.

Coverage analyses The following tables include coverage analysis when using different tolerances. Note, that not all of the 100 runs are considered converged under different settings, therefore each column has individual number of total simulations and corresponding confidence intervals.

Table S1: mean–0.02

	ESS threshold, total runs, confidence interval		
Paramter	200, 99, (90-98)	500, 92, (84-91)	Full, 100, (91-99)
Tree.height	92/91	88/88	96/96
Tree.treeLength	93/92	87/86	97/94
clockRate	89/91	83/85	91/93
gammaShape	88/89	85/84	92/92
popSize	93/91	89/89	96/97
CoalescentConstant	91/89	85/84	92/92
kappa	90/95	85/86	94/94
freqParameter.1	91/91	85/85	92/92
freqParameter.2	89/91	85/87	94/95
freqParameter.3	86/85	81/76	87/90
freqParameter.4	95/93	84/84	95/95

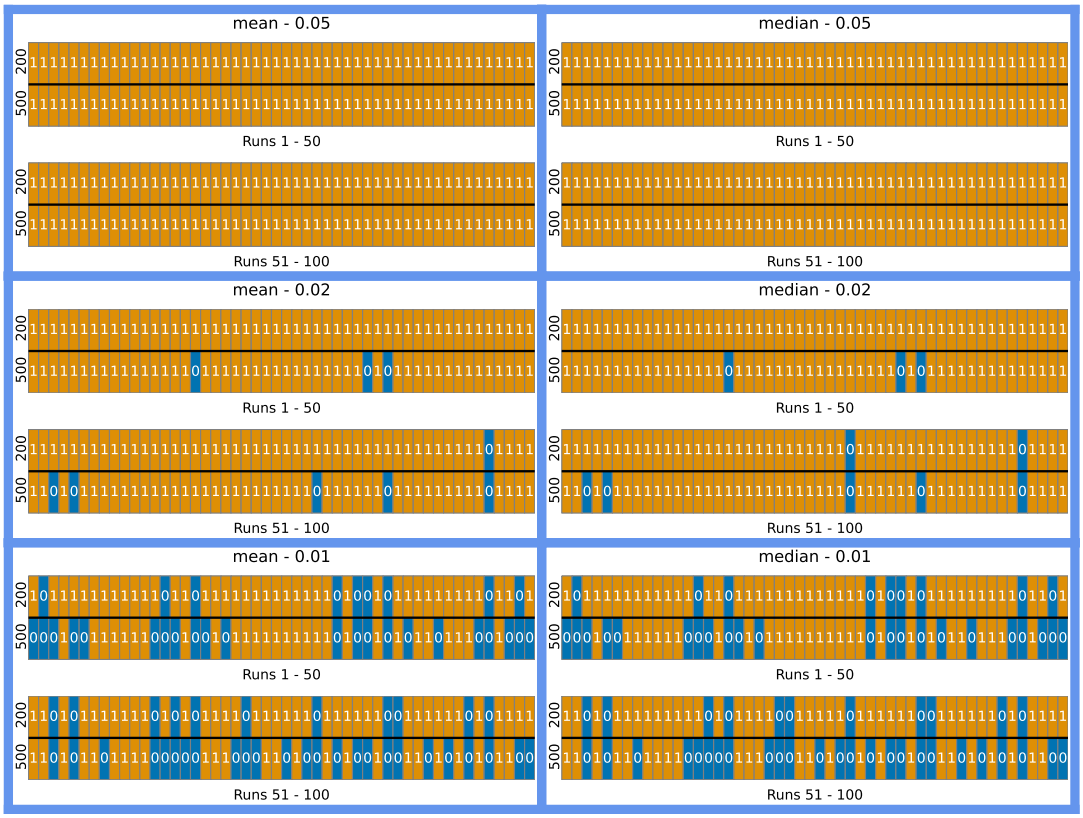

Figure S5: Showcasing convergence detection on the well calibrated simulation study as described in the main papers Validation section. This plot compares the mean and median averages for smoothing the PSRF values to calculate the GR value. There are no significant differences between the two columns of the figure. See main paper Figure 1 for more detail.

Table S2: mean–0.01

	ESS threshold, total runs, confidence interval		
Paramter	200, 80, (72-79)	500, 53, (48-53)	Full, 100, (91-99)
Tree.height	74/76	51/50	96/96
Tree.treeLength	76/75	49/50	97/94
clockRate	73/72	49/49	91/93
gammaShape	72/71	49/49	92/92
popSize	76/75	51/50	96/97
CoalescentConstant	73/73	48/48	92/92
kappa	73/75	49/49	94/94
freqParameter.1	72/74	49/49	92/92
freqParameter.2	73/76	48/49	94/95
freqParameter.3	69/70	45/45	87/90
freqParameter.4	73/74	48/49	95/95

S4.2 Median smoothing coverage for well calibrated simulation study

The following tables display the coverage analyses when using the median smoothing option for the GR value. Because of no notable difference to the version using the mean, the coverage tables are also very similar. Note, that not all of the 100 runs are considered converged under different settings, therefore each column has individual number of total simulations and corresponding confidence intervals.

Table S3: median–0.05

	ESS threshold, total runs, confidence interval		
Paramter	200, 100, (91-99)	500, 100, (91-99)	Full, 100, (91-99)
Tree.height	92/94	95/96	96/96
Tree.treeLength	94/92	96/95	97/94
clockRate	91/92	90/92	91/93
gammaShape	87/90	92/93	92/92
popSize	90/94	97/95	96/97
CoalescentConstant	91/91	93/91	92/92
kappa	92/94	93/93	94/94
freqParameter.1	94/91	92/92	92/92
freqParameter.2	93/92	94/94	94/95
freqParameter.3	86/87	89/88	87/90
freqParameter.4	95/95	94/94	95/95

Table S4: median–0.02

	ESS threshold, total runs, confidence interval		
Paramter	200, 98, (89-97)	500, 92, (84-91)	Full, 100, (91-99)
Tree.height	91/90	88/88	96/96
Tree.treeLength	91/92	88/86	97/94
clockRate	89/89	83/84	91/93
gammaShape	86/89	85/82	92/92
popSize	92/93	89/89	96/97
CoalescentConstant	90/88	85/84	92/92
kappa	89/92	85/86	94/94
freqParameter.1	90/90	85/85	92/92
freqParameter.2	89/90	85/87	94/95
freqParameter.3	85/83	81/77	87/90
freqParameter.4	93/92	85/84	95/95

Table S5: median–0.01

	ESS threshold, total runs, confidence interval		
Paramter	200, 80, (72-79)	500, 53, (48-53)	Full, 100, (91-99)
Tree.height	75/76	51/50	96/96
Tree.treeLength	75/75	49/50	97/94
clockRate	73/73	49/49	91/93
gammaShape	73/72	49/49	92/92
popSize	77/77	51/50	96/97
CoalescentConstant	72/72	48/48	92/92
kappa	73/74	49/49	94/94
freqParameter.1	73/75	49/49	92/92
freqParameter.2	73/75	48/49	94/95
freqParameter.3	69/70	45/45	87/90
freqParameter.4	72/73	48/50	95/95

S5 Beast ASM package on DS1–DS11 and other data sets

The following tables display the individual run lengths of 10 independent chains on DS1–DS11 using the ASM package, each chain has a sample frequency of $10,000$.

Table S6: MCMC chain lengths over 10 runs before automatically stopping

	DS1	DS2	DS3	DS4	DS5	DS6
shortest	1990000	1990000	2320000	2000000	2000000	2000000
	2000000	1990000	2470000	2020000	2020000	2360000
	2000000	1990000	2860000	2300000	2130000	2540000
	2000000	2000000	2940000	2310000	2130000	2610000
	2080000	2120000	3460000	2710000	2260000	2870000
	2290000	2360000	3550000	3100000	2280000	2970000
	2430000	2560000	4080000	4100000	2620000	3330000
	3020000	2620000	4120000	4780000	2620000	3390000
	3070000	2870000	4190000	9200000	3070000	3440000
longest	3730000	3280000	4360000	9830000	3410000	3510000

	DS7	DS8	DS9	DS10	DS11
shortest	2340000	1990000	1990000	2210000	2000000
	2970000	1990000	2300000	2240000	2150000
	3510000	2110000	2310000	2250000	2190000
	3750000	2270000	2320000	2440000	2260000
	3770000	2270000	2390000	2480000	2390000
	3990000	2360000	2440000	2490000	2630000
	4240000	2770000	2480000	2520000	2690000
	4320000	3210000	3090000	2750000	2780000
	4410000	3510000	3300000	2790000	3130000
longest	4670000	3530000	4350000	3340000	3570000

The following tables display the individual run lengths of 10 independent chains on HCV, primate and COVID-19 alignments using the ASM package, each chain has a sample frequency of $10,000$.

Table S7: MCMC chain lengths over 10 runs before automatically stopping j

	HCV	Primates	COVID-19
shortest	4650000	2030000	45630000
	6430000	2370000	52130000
	12420000	2800000	53410000
	14560000	2910000	53940000
	14750000	3370000	63740000
	17940000	3450000	73740000
	18500000	4030000	79830000
	19310000	4600000	99210000
	19410000	6030000	103840000
longest	19490000	11810000	148170000