

Original Article

tbea: tools for pre- and post-processing in Bayesian evolutionary analyses

Gustavo A. Ballen^{1,*} and Sandra Reinales²

¹Instituto de Biociências de Botucatu, Universidade Estadual Paulista 'Júlio de Mesquita Filho', Botucatu, SP, Brazil

²Departamento de Botânica, Instituto de Biociências, Universidade de São Paulo, SP, Brazil

*Corresponding author. Instituto de Biociências de Botucatu, Universidade Estadual Paulista 'Júlio de Mesquita Filho', Botucatu, SP, Brazil.

E-mail: gustavo.a.ballen@gmail.com; gustavo.ballen@unesp.br

ABSTRACT

Estimating phylogenies in which branch lengths are expressed in units of absolute time is crucial for testing hypotheses in evolutionary biology. However, bioinformatic tools to pre- and post-process data from Bayesian divergence time estimation analyses are often not easily interoperable, and documenting methodological choices is not a generalized practice. The R package *tbea* is a tool-set to integrate biological, geological, and paleontological information to optimize the specification of models, their parameters, and prior distributions in divergence time estimation analyses. *tbea* implements statistical models to (i) better translate time information in dating sources into the specified calibration densities, (ii) improve comparisons between prior and posterior distributions for parameters of interest, (iii) carry out inference on origination times for a set of distributions, (iv) summarize different distributions into a single one, and (v) improve the reproducibility of divergence time estimation analyses allowing users to document methodological choices. We illustrate the functionalities of *tbea* by carrying out two worked examples, one on the phylogenetic relationships and divergence time estimation of South American Cynodontidae, and the other on the separation time of drainages east and west of the Andes in South America. It is expected that the tools available in *tbea* will be key when estimating events in time from sets of point estimates, and the combination of different posterior densities from the same parameter will be useful in justifying the selection of secondary calibration points, or discussing the timing of biogeographical events when multiple sources are available.

Keywords: biogeography; computational biology; data visualization; divergence time estimation; fossil record; morphology; reproducibility; secondary and tip calibration; total-evidence analysis

INTRODUCTION

In phylogenetic inference, the characters, either molecular or morphological, only contain information about how closely or distantly related the terminals are. Branch lengths, as estimated in standard phylogenetic analysis (e.g. maximum likelihood or Bayesian inference) represent a combination of the expected number of changes per site and the time required to accumulate changes along branches. Therefore, estimating a phylogeny in which branch lengths are in units of absolute time is key in comparative analyses such as modelling morphological evolution, biogeography, or macroevolutionary dynamics (Warnock and Wright 2021). One key reason why time trees are useful lies in the nature of evolutionary models which consider a stochastic process in time, such as the discrete Markov process often used to model evolution (Lewis 2001, Yang 2014).

In a very general form, we can describe a Bayesian divergence time estimation model as

$$P(\Theta | X, C) = \frac{P(X | \theta, \psi_s) P(C | \psi_t) P(\psi_s | \psi_t, \theta) P(\psi_t, \theta)}{P(X, C)} \quad (1)$$

where $P(\Theta | X, C)$ is the posterior distribution of the model parameters, $P(X | \theta, \psi_s)$ is the phylogenetic likelihood of observing the sequence data X , $P(C | \psi_t)$ are the time calibration densities, $P(\psi_s | \psi_t, \theta)$ is the clock model, $P(\psi_t, \theta)$ is the tree process, and $P(X, C)$ is a normalizing constant.

The denominator $P(X, C)$ in Equation (1) cannot be calculated analytically and we need resort to Markov chain Monte Carlo (MCMC) sampling algorithms for approximating the posterior distribution(s) numerically. The result is a posterior sample of

time trees in which branch lengths are in units of absolute time (usually million years, Myr). Node times and branch lengths are parameters in the tree topology ψ , and the posterior age distribution of each node can be summarized from the sample of time trees. When the tree ψ is fixed, we only have the posterior distributions of node ages, but when the tree is coestimated with the divergence times, we also have a posterior distribution of trees, which may or may not share nodes across the different elements of the distribution. A Bayesian analysis of any sort aims to characterize the posterior distributions, and thus the uncertainty on estimated parameters rather than point estimates.

Terminology applied to divergence time estimation has been quite inconsistent in the literature. It has been called molecular clock analysis, tree calibration, molecular dating, tree dating, and clock dating in the massive literature on theoretical and empirical aspects of this family of analyses. It is useful to define here the different terms in use as these allow a more precise description of the elements in the analysis. We propose to use divergence time estimation as the name for the analysis in which branch lengths are separated in rate and time, and results in a time tree where branch lengths represent time. Clock, be it molecular or morphological (Barba-Montoya *et al.* 2021), is a model that describes how the rates change along the tree. Calibration is the specification of time information in the form of a calibration density, which is a probability density function (PDF) that describes the uncertainty of node or tip times in the tree. Dating means telling the age of an object, for instance with the use of geochronological, biostratigraphic, or any other information which allows us to determine the absolute age of fossils. In this sense, none of trees, clocks, or molecules are dated but instead the fossils that are used to construct calibrations. Note that dating is a term widely applied to geochronological research (Rink and Thompson 2015). Under this scheme, molecular (or morphological) clocks as well as calibration densities are components of the complete divergence time estimation model depicted in Equation (1). We use the terms and definitions above throughout the present article.

The most common strategy for using independent temporal evidence to calibrate a phylogeny to absolute time is assigning time constraints to one or more nodes in the phylogeny, usually called node calibrations (Rannala and Yang 2007). This time information comes from a variety of sources, including geochronological estimates from fossils or biogeographical events (primary calibrations; Landis 2017), or time estimates obtained from previous studies (secondary calibrations; Schenk 2016). Node calibrations based on fossil evidence are set using the oldest fossil record of a clade, which establishes a minimum age for that node in the phylogeny (Rannala and Yang 2007, O'Reilly *et al.* 2015). Nevertheless, it is possible to consider that fossils are part of the same diversification process which generates the extant species, and incorporate them as tips alongside living relatives (tip calibration; Ronquist *et al.* 2012, Heath *et al.* 2014). This approach allows for inclusion of all available fossils (Heath *et al.* 2014), but the impact of the taxonomic and the stratigraphic age uncertainty propagates to divergence time estimates and should be incorporated explicitly (O'Reilly *et al.* 2015, Barido-Sottani *et al.* 2019). Several methods are available to estimate time trees using one of these approaches (e.g. Sanderson 2002, Heled and Drummond 2010, Stadler 2010, To *et al.* 2016).

One of the challenges in node calibration strategies is translating the inherent uncertainty in fossil age information into a prior distribution (Warnock *et al.* 2015). This uncertainty usually comes from the litho-, bio-, chemo-, cyclo-, and/or magneto-stratigraphic method using to establish the absolute date of the fossil provenance (Parham *et al.* 2012, O'Reilly *et al.* 2015). There are a variety of probability densities implemented to account for the fossil age uncertainty (e.g. Uniform, Lognormal; Warnock *et al.* 2015). Some programs such as Beast 2 (Bouckaert *et al.* 2019) require users to input parameter values for the node calibration densities, often by eye, aided by plots as is done in *beaut.i*. This process is not reproducible and can be impractical when several calibration points must be set.

On the other hand, when different time estimates for the same event (e.g. a node) are available, summarizing such information to constrain a node age (i.e. secondary node calibration) using all available information is not straightforward. Consequently, empirical studies usually resort to manual specification to find a single and representative distribution, or choose among alternatives using all sorts of arguments such as the most cited estimate for the group of interest, the most recent, or the one with the best taxonomic sampling. Nevertheless, a reproducible method to summarize different sources of information would allow the incorporation of the uncertainty associated with the time estimates for a node, because every study constitutes evidence, as long as the estimates are independent. Tools for summarizing different time estimates can be also applied in biogeography. Empirical approaches to estimating vicariant events in specific groups involve plotting the time estimates along with their uncertainties to visually extract information on when such events are likely to have occurred (Hazzi *et al.* 2018, Pérez-Escobar *et al.* 2019). Another alternative has been to plot the mean of these data as a point estimate for the general time of separation (Smith *et al.* 2014). Nevertheless, a statistical approach would allow us to infer the age and its uncertainty, at which the process responsible for all those vicariance events took place (e.g. separation time of drainages east and west of the Andes in South America), using divergence time estimations for different biotic groups. These tools will impact beyond evolutionary biology as Bayesian phylogenetic methods are also applied to other fields in cultural evolution such as linguistics (Maurits and Griffiths 2014, Hoffmann *et al.* 2021, Neureiter *et al.* 2022, Greenhill 2023), palaeolithic artefacts (Matzig *et al.* 2024), and musicology (Hajic *et al.* 2023, Ballen *et al.* 2024, Hajic *et al.* 2025).

The package *tbea* has tools for improving the reproducibility of the generation of concatenated alignments X , express better the time information in calibration sources onto the specified calibrations $P(C|\psi_t)$, propose adequate initial values for the time tree ψ , improve comparisons between prior and posterior distributions for parameters of interest, in particular node times, and carry out inference on origination times for a set of posterior distributions of the same class in different studies. The package also includes tools for summarizing tree distributions in terms of both topology and branch lengths, as well as two methods for estimating the time of an event where independent estimates $(P_{iN}(\Theta_i | X_i, C_i))$ are available (e.g. multiple estimates of the timing of a biogeographical event). Finally, it implements conflation of densities as a tool for summarizing multiple distributions of the

same parameter into a single one. This tool-set will help researchers in evolutionary biology to improve the reproducibility of their Bayesian analyses, and help them streamline coding strategies in different phases of analysis, from specification to comparisons of results.

Typical users of `tbea` include evolutionary biologists wanting to use Bayesian methods in the realm of phylogenetics and macroevolution where it is necessary to specify temporal information in the form of calibration priors. Researchers in paleontology and geology may also use functions of this package when interested in format conversion of morphological matrices, summarization of phylogenetic analyses of extinct and/or extant organisms, combination of multiple time estimates in the form of distributions, and event tools for testing the assumption of isochron. Any researcher interested in combining multiple distributions may find the implementation of distribution conflation helpful, even outside the realm of biology as it is a general statistical technique. Finally, researchers applying statistical methods developed in a biological context but later applied to questions in cultural evolution and digital humanities may use the tools herein presented in their own research questions.

IMPLEMENTATION

Three computational techniques are used in several functions of this package: quadratic loss optimization, numerical integration, and metaprogramming. Optimization is used when fitting parameters of interest, as well as during calculation of quantiles. This is achieved by using numerical optimization on a quadratic loss function of the form $\arg\min_x \phi(x) = (f(x) - t)^2$ where $f(x)$ is the target function and t the target value. Optimization stops once $|(f(x) - t)^2| < \delta$ is attained for a given tolerance level δ . Optimization works with both univariate and multivariate functions. This technique is used in the functions `findParams`, `c_trunc-cauchy`, and `quantile_conflation`.

Numerical integration uses an approximation for the area under the curve $\int_a^b f(x)dx$ where a and b can be positive and/or negative infinite as well as real numbers. It approximates the area under the curve using adaptive quadrature in unidimensional functions where multiple subdivisions of such area are approximated and proposed following an adaptive procedure that decides when to stop subdividing the area, once $\epsilon < \delta$ where ϵ is the error and δ is the tolerance, in a similar way as in numerical optimization. This is used for calculating quantiles on arbitrary conflated distributions, as well as for calculating the conflation itself (see below).

Metaprogramming works by writing code that returns code which is in turn evaluated. For instance, suppose a function returns a string which constitutes a valid expression `expr = '1 + 1'`. Then, the value of `expr` can be parsed and evaluated by the interpreter as its input: `eval(expr) => eval('1 + 1') => 2`. The power of such technique is that the user does not need to hard-code expressions for calculating operations that take functions and their evaluations as input, and therefore the output of such evaluations is very flexible. This is especially useful for implementing operations that take functions as input rather than concrete values such as vectors, matrices, or lists. This technique is used when generating function calls using PDFs in `findParams` as well as for conflating distributions.

The package `tbea` is hosted on CRAN (<https://cran.r-project.org/package=tbea>) and development versions are available on GitHub (<https://github.com/gaballench/tbea>). The stable version of the package can be installed with `install.packages('tbea')` and the development version with `remotes::install_github('gaballench/tbea')`.

Dependencies

One of our design choices was to reduce dependencies to a minimum, thus `tbea` only depends at the moment on `ape` (Paradis and Schliep 2019) for phylogenetic tree representation and manipulation, `Rfit` (Kloke and McKean 2012) for extending the possible linear models to be used when inferring the x -intercept of cumulative density functions, and `coda` (Plummer *et al.* 2006) for calculating the highest posterior density interval which is used in crossplots.

Comparison with other packages

Several packages have been developed for phylogenetics and related areas (Gearty *et al.* 2024). Ironically, there are as many standards as software tools, and this is reflected in the multiplicity of approaches to software development applied to evolutionary biology. Multiple functionalities are also implemented in different ways, thus providing a high amount of flexibility to users of phylogenetic methods. Some tools available in this package are already available in a different form in some other packages (e.g. tools for interacting with data matrices, concatenation, and format conversion), while others are completely new to our knowledge (e.g. parameter estimation in distributions with a set of quantiles, confidence intervals on the x -intercept, or the conflation method). We think that having multiple implementations of the same tools is positive rather than redundant as much as the different implementations allow the same tools to be used in different analytical contexts rather than forcing the users to have the same settings, formats, and input, which is the case when fewer tools are available.

Some packages target `Beast` or `Beast2` as the preparation of input, programmatic use, and parsing of output is challenging under certain circumstances. One of the most popular tools, `beauti` (Bouckaert *et al.* 2019), is used as a companion for `Beast` and `Beast2` as it allows to generate the XML input file from a graphical user interface. The package `beastier` (formerly `babette`; Bilderbeek and Etienne 2008), on the other hand, has code for interacting with programs of the `Beast` suite. Although some of our functions can be used to help specify analyses in `Beast2` and parse output logs from the program, our package is not intended to be a wrapper or a gateway to `Beast`, but rather a tool for improving reproducibility when using that program. A similar package, `beastio` (Du Plessis 2020), aims at providing pre- and post-processing tools specifically for `Beast` and `Beast2`.

`EvoPhylo` (Simões *et al.* 2023) targets divergence time estimation analyses using morphological datasets and their clock models, although it can be used for molecular datasets too. The package also pre- and post-processes output from both `Beast2` and `MrBayes` but is restricted to situations where morphological data, in particular for (but not restricted to) completely extinct taxon sampling, is present in the analysis. In a similar way, `Chronospace` (Mongiardino Koch and Milla Carmona 2024) targets

the exploration of time tree parameter distributions and measures of sensitivity when carrying out multiple analyses. Our package has functionality which similarly targets unrooted tree distributions, thus being complementary. *Convenience* (Fabreti and Höhna 2022) is a package for assessing Bayesian convergence in both continuous and discrete (topology) parameters, therefore offering post-analysis tools. *RevGadgets* (Tribble et al. 2022) is tightly integrated with *RevBayes* in a similar way as packages targeting *Beast/Beast2*. It has functionality for assessing estimated distributions similar to *Tracer* but implemented in R, and therefore offers tools similar to our tools for comparing distributions. It also has tools for plotting trees from a number of analyses, thus offering output similar to *FigTree* (Rambaut 2018) but using R code. It is strongly integrated with *ggplot2* whereas our approach is to use *base* instead in order to reduce the amount of dependencies while producing simple yet informative plots. *CladeDate* (Claramunt 2022) is very similar in that it aims at generating calibration densities from the fossil record. However, that package uses data augmentation via sampling from the uncertainty in fossil occurrences in order to generate a distribution of a parameter θ (fitted by using also methods for stratigraphic intervals) representing the origination of the clade. *phruta* (Román-Palacios 2023) has tools for retrieving data from GenBank and then concatenate them for phylogenetic analysis. In contrast, the present package allows us to concatenate matrices including both DNA and morphological data, so that total-evidence divergence time estimation and tree inference can be carried out combining these two data sources. The package *mcmc3r* (dos Reis et al. 2018, Álvarez-Carretero et al. 2019) targets analysis such as model selection using Bayes factors specifically using the program *MCMCTree*, whereas tools in *tbea* can be used for comparing results from independent runs using *MCMCTree*. Also, *mcmc3r* has functions for plotting the minimum and both upper and lower bound calibration densities used by *MCMCTree*, in a similar way as our package aids in plotting *Beast2*'s lognormal calibration density. Finally, *TreeTools* (Smith 2019) implements tools for interacting with output from TNT (Goloboff et al. 2008). Although *tbea* has a function for reading trees in TNT format, its aim is to convert them to newick format rather than to generate an R representation of the tree; therefore, the operations on the TNT tree text representation return also a tree in text representation, but in newick format instead.

Worked examples

We introduce the package and illustrate its functionalities by carrying out two worked examples, one on the phylogenetic relationships and divergence time estimation of South American sabre-tooth characins of the family Cynodontidae, and the other on the separation time of drainages east and west of the Andes in South America. The former targets pre-analysis tools and a couple of functionalities for post-analysis settings, whereas the latter aims at demonstrating the use of post-analysis tools when information from different analyses is available and the goal is to infer a single parameter, in this case, the separation time between these two biogeographical regions.

The South American sabre-tooth characins

We used the South American sabre-tooth characins of the family Cynodontidae for a divergence time estimation analysis. The phylogenetic relationships of the group were revisited by Ballen et al. (2021), although these authors did not provide information about the temporal scale of their evolutionary relationships. We used the morphological and molecular datasets of that study, including a fossil terminal.

The dataset comprises 10 terminals, eight molecular markers, and one morphological partition with 80 characters. We refer the reader to the study of Ballen et al. (2021) for details on the substitution models and other details of the phylogenetic analysis.

Divergence time estimation was carried out using *Beast2* (Bouckaert et al. 2019) with the *SampledAncestors* package (Gavryushkina et al. 2014). An analysis sampling from the prior, that is, ignoring the alignment, was also carried out with the option *sampleFromPrior* in *beauti*. Preparation of the initial tree was carried out using the package *ape* (Paradis and Schliep 2019). Post-analysis tree plotting was carried out in R (R Core Team 2024).

The separation of drainages east and west of the Andes

Divergence time estimation data were compiled from the literature that reported separation events between cis- (i.e., east of the Andes) and trans-Andean (i.e., west of the Andes) lineages, be it contemporary species or nodal divergences (Collins and Dubach 2000, Cortés-Ortiz et al. 2003, Dick et al. 2003, Cheviron et al. 2005, Grau et al. 2005, Ribas et al. 2005, 2007, Hardman and Lundberg 2006, Miller et al. 2008, Elias et al. 2009, Patané et al. 2009, Patel et al. 2011, Weir and Price 2011, D'Horta et al. 2013, Říčan et al. 2013, Voss et al. 2013, Abe et al. 2014, Fernandes et al. 2014, Gutiérrez et al. 2014, Machado et al. 2014, Picq et al. 2014, Smith et al. 2014, Hernández Torres 2015, Ruiz-García et al. 2015). A total of 54 data points were compiled along with their credible intervals, most frequently highest posterior density intervals.

The distribution of divergence time data is very asymmetrical with a heavy right tail, the product of several outliers older than 10 Mya (Fig. 1). These outliers affect those methods that make a stronger assumption on the lack of uncertainty in the measurement of time values (e.g. methods based on stratigraphic intervals, and less strongly on those based on cumulative densities).

Coding examples and vignettes

The package *tbea* provides a number of vignettes in addition to the examples in this article. PDF rendered versions of the vignettes can be accessed on CRAN (<https://CRAN.R-project.org/package=tbea>) as well as HTML rendered versions on a dedicated website (<https://gaballench.github.io/tbea>).

RESULTS

Concatenation and total-evidence analysis

Concatenation can mean two different operations when working with partitions of data in phylogenetic analysis. On the one hand, concatenation means pasting together multiple data matrices (i.e.




Figure 1. Descriptive statistics of the divergence time data. The histogram illustrates that the multimodality in the density line is an artefact of discontinuity in the age data, while the boxplot on top of the figure shows the existence of multiple outliers to the right of the distribution. All three methods agree in showing the asymmetry in the distribution with a heavy right tail. Dashed line represents the empirical kernel density.

data partitions) of the same data type (e.g. DNA alignments) into a single one which is in turn analyzed as a single partition, with a single evolutionary model and therefore parameter set. There are numerous issues when combining data partitions in this way, with the fact that different genome regions result in different gene trees is probably the most problematic one. As a consequence, the largest original partition will end up dominating the signal during inference because the parameter set is estimated from the partition as a whole. On the other hand, concatenation means combining different partitions into the same analysis and informing the same tree, although applying independent and potentially different evolutionary models for each one. Independent partitions can now inform shared parameters equally (i.e. branch lengths and topology), while having their independent parameter sets estimated independent from one another. Because of this property, it allows us to use multiple data types as partitions in what has been called total-evidence analyses (Ronquist *et al.* 2012, Heath *et al.* 2014), where we may for instance have some partitions for specific genes, a partition for amino acid data, and yet another one for morphological characters. This form of concatenation does not, however, solve the issue of gene-tree to species-tree discordance (in fact, it can be still very problematic in these scenarios; Mendes and Hahn 2018), but is the only alternative available when multiple data

types are of interest. In this work we use the second meaning of concatenation as it is the relevant one during divergence time estimation using total-evidence approaches.

Tools are available for concatenating DNA matrices (e.g. the R package *apex*; Jombart *et al.* 2017), although few deal with concatenating different data types such as morphological matrices (but see Salinas *et al.* 2024). Our package offers tools designed for including both DNA and morphological matrices (see the vignette *concatenation*), which is necessary, for example, in MrBayes for total-evidence analyses. Often the morphological datasets are compiled in some form of table, for instance a spreadsheet, and the conversion to nexus format for later concatenation to other partitions is a useful feature even for working with programs that support individual files per partition such as RevBayes or iqtrees.

Calibration specification

Calibration densities can be divided into two groups: node and tip calibrations. Node calibrations are PDFs which represent our uncertainty about the age of a given node. Similarly, tip calibrations represent the uncertainty about the age of a terminal in the tree, for instance a fossil terminal for which we also have data in the matrix. Different programs make available different PDFs, with

the Normal, Lognormal, Exponential, and Uniform being popular options. Calibration densities may also represent bounds on an age, that is, lower and upper quantiles of some probability, such as the interval of 95% of the probability. In the case of continuous distributions with support in $(-\infty, \infty)$, $[0, \infty)$, or $(-\infty, 0]$, or, one or both bounds are soft, that is, we allow ages to be outside of that interval with some probability, this being 2.5% on each tail if we use a 95% interval. In other distributions such as the Uniform, bounds can be hard, that is, we say the probability of being outside the interval is 0. It is possible to set distributions with combinations of soft and hard bounds. For instance, the Lognormal distribution with an offset states that the probability of observing ages younger than the offset is 0, whereas the upper bound corresponding to a quantile of 95% will be soft with probability 5%.

We specify a calibration density by picking a given density function and setting its parameter values. Some considerations are important when setting calibration densities; for instance, if we use the uncertainty from the dating of a given fossil as the calibration density, it implies that we mean the fossil age to be also the node age. This may be a strong assumption about the position of the fossil, and then we may want to use the fossil as a lower bound instead (Parham *et al.* 2012). In this case, we may want the fossil to represent a hard bound (i.e. the age of the node of interest cannot be younger than the fossil), and then pick an upper soft bound for an age such that older ages are still possible yet unreliable. Let us suppose we have a node with a fossil age of 50 Mya which we want to set as a hard bound, and then previous estimates of the age of the total group under analysis saying that it is probably younger than 100 Mya, for instance being this the root age. We may pick a distribution with a hard bound at 50 Mya and a soft 95% bound at 100 Mya. Because additional considerations need to be made in this second approach, we will assume that the fossil in the following example is on the node.

Regardless of the specific details supporting our choice of bounds, it is important to use proper densities when specifying calibration priors. Improper densities are those whose area under the curve for the complete support does not integrate to 1.0. In contrast, proper densities integrate to 1.0. Improper densities induce issues during MCMC sampling (Yang 2014) and should be avoided. An example of improper distribution is setting a calibration density as a Uniform with just a minimum hard bound but no maximum bound $U(\min, \infty)$ with support in $[\min, \infty)$, because $\int_{\min}^{\infty} U(t)dt = \infty$.

We specified two calibration prior densities to estimate divergence times for the family Cynodontidae: (i) a node calibration using the oldest record of the genus *Hydrolycus* (early Barrancan, > 41.6 & –40.94 Mya) to calibrate the most-recent common ancestor (MRCA) of *Hydrolycus armatus* and *H. scomberoides* using a Lognormal distribution, and an alternative representation of the same age uncertainty using an Exponential distribution (see vignette `node_calibration_density`). In general, if the uncertainty can be described by either an Exponential or a Lognormal, the former is preferred because it has one parameter instead of two in the latter. The difference between both lies in the flexible position for the mean inside the support of the Lognormal distribution, whereas the mean, 0.025, and 0.975 quantiles are governed by a single parameter in the Exponential distribution.

Node calibration density

When using time information associated with a fossil as calibration density, age uncertainty usually comes from biostratigraphy (e.g. if no radiometric estimate applies closer in the stratigraphic context to the fossil locality). Then, it is possible to represent that uncertainty using a Uniform distribution with first and last time parameters (e.g. 40.94 and 41.6 Mya respectively), or a Lognormal distribution with soft bounds. In the second case, we can find the combination of mean and standard deviation that better describes a distribution whose density cumulative probability values 0.0, 0.5, and 0.975 correspond to the quantiles defined by the Uniform prior above. The function `findParams` finds such combination of parameters for a given pair of quantiles. As the standard Lognormal distribution is defined between $[0, \infty)$, we need to apply an offset towards the minimum age. This will make it possible to relax the rather strong assumption about the node age, because by using a Uniform prior the probability of observing times which fall outside the interval is zero.

Note that the Exponential distribution can be parametrized in terms of the scale ($1/\lambda$, as is the case for *Beast2*) or the rate (λ) as used by R. Therefore, the value to input when defining the prior of an Exponential distribution in *beauti* is the scale, which in our case is $1/0.17 = 5.8$. When plotting the prior in R (e.g. using the function `curve` as above) we need to use λ (i.e. 0.17) rather than the mean (Fig. 2B).

Alternative specification: the L calibration density of MCMCTree

Let us suppose that we want to carry out an analysis using MCMCTree instead (see vignette `l_calibration_mcmctree`). That program uses different densities for specifying a node calibration. One of these, called L, implements a truncated Cauchy distribution (Inoue *et al.* 2010) whose shape can be controlled by the location and scale (p and c). It is used frequently for describing minimum-age calibrations where a maximum is unknown. Although it can give the impression that we only need a minimum age (which can be either soft or hard), this is not the case of an improper distribution as it is controlled by p and c so that its area under the curve is 1.0. It is useful for describing asymmetric calibrations, although care must be taken as specifying both p and c have a profound impact on the shape of the distribution (see e.g. the unnumbered figure in p. 49 of the PAML documentation; Yang 2007). The function `c_truncauchy` helps us specify such density as it estimates the value of c provided that we have both age minimum and maximum, and define some probability pr to be allocated to the right of the maximum age. For instance, our minimum age is say 1 in units of 100 Myr, and our maximum age is 4.93 also in units of 100 Myr. The minimum is a soft bound allowing 0.025 of the density to be allocated to the left of it, whereas 0.975 is the percentile at which we observe the maximum age. The quantity $p = 0.001$ has been chosen so that the mode is closer to the minimum age. Then, we can find the value of c which completes the definition of the L density to be used in MCMCTree, which happens to be around 0.0009.

Initial trees

Programs that co-estimate the tree topology and divergence times (e.g. *Beast2*, *RevBayes*, *MrBayes*) first need to initialize the tree in order to start MCMC sampling. Depending on the size




Figure 2. A, calibration density for the node corresponding to the genus *Hydrolycus* as generated by the function `findParams` from the quantiles found in the calibration justification, and plotted with `lognormalBeast`. B, alternative specification of the same calibration density but using an Exponential distribution as generated by the function `findParams`, and plotted with `curve(dexp(x))`.




Figure 3. Calibrated tree of the Cynodontidae which is consistent with the calibration densities and can be used as starting tree in analyses of divergence time estimation using `Beast2`. Calibrations shown are the root calibration density to the left, and the *Hydrolycus* calibration density to the right.

of the analysis (i.e. the number of terminals), such an initial tree can be quite poor and make convergence take longer. It is therefore useful to help initialize the analysis by setting the initial tree, by calculating a reasonable one from which sampling then converges quicker. When performing standard tree inference we do not need to worry about the tree being consistent with times, but this is not the case with divergence time estimation. In the latter case, not just the tree but also the node ages need to be consistent with the calibration densities, and we need to guarantee that before starting the analysis. This is a frequent reason for an analysis in `Beast2` being unable to initialize.

The package `tbea` helps us find initial trees when using the output from `TNT`, a parsimony program, by providing functions

for performing format conversion from `TNT` to standard newick format. It is then possible to use penalized likelihood as implemented in the package `ape` for finding a set of branch lengths which is consistent with our calibrations as shown in the vignette `starting_tree`. [Figure 3](#) gives an example of a tree that was converted from `TNT` format and then set to branch lengths consistent with the calibrations specified in the previous section.

Post-analysis comparisons

Summary of (posterior) tree distributions

After co-estimating the tree topology during MCMC sampling, we end up with two sets of values, one for parameter values and another one with a tree sample. Depending on the type of analysis,

trees can be unrooted as in standard tree inference, or rooted as in divergence time estimation. Different programs implement their own summarizing tools which fit the different design choices and therefore are not necessarily interoperable: tools for *Beast2* do not necessarily work with samples from *MrBayes*, and vice versa. In particular, existing summary tools for tree topology are not well suited for unrooted trees as they were designed with time trees in mind, which are rooted.

There are two steps necessary for summarizing tree samples (see vignette `posterior_tree_distribution`). One is to decide how many distinct topologies (disregarding branch lengths) are in the set of trees and then we need to summarize branch lengths for each of these tree subsets using any sensible summary measure (e.g. mean or median). The functions `topoFreq` and `summaryBrlen` aid in these two tasks respectively. `topoFreq` uses the Robinson–Foulds (RF) distance among trees in order to find clusters which are identical in topology and then bind them together in a list that we can later operate over. The RF distance is defined over unrooted trees, so it is necessary to check whether the input trees are. Unrooting a rooted tree will destroy the root node and then add together both adjacent branch lengths, and this will be reflected when summarizing with `summaryBrlen` using the output of `topoFreq`. This will

return one branch length fewer than the number present in the original sample of rooted trees. Accordingly, this method will only work on sets of unrooted trees, or sets of time trees of fixed topology where we are only interested in finding a summary value for branch lengths. The function `topoFreq` should not be used when both the topology and divergence times are being coestimated (e.g. when the posterior of trees comes from not fixing the topology in *Beast2*, *RevBayes*, and *MrBayes*) but it is safe to use when the inference was done to estimate only the unrooted topology and branch lengths, such as in standard tree inference in *MrBayes*. In the example, we used *MrBayes* for standard tree inference of the Cynodontidae dataset, and then summarized the posterior tree sample using the `topoFreq` and `summaryBrlen` functions (Fig. 4).

Prior–posterior comparisons

It is important to compare prior and posterior distributions of a given parameter in Bayesian analyses because by comparing both sets of distributions we can identify cases where the information content in the likelihood is low, and therefore the posterior is being driven by the prior. This is particularly so for continuous parameters where we can plot the distribution of parameter samples as a kernel density or empirical density. Then we can compare the prior and posterior




Figure 4. Summary of a posterior tree sample from the standard phylogenetic tree inference of the Cynodontidae dataset using *MrBayes*, which comprise nearly 90% of the tree posterior sample. These are labelled by posterior probability and sorted in decreasing order.

densities for a given parameter. Note that we deactivate the data by raising the likelihood function to a power of 0 so that it is 1.0 in every MCMC iteration, and thus sample just from the prior densities. This is made by setting the option `sampleFromPrior` in `beauti`, or using the option `underPrior=TRUE` in the method `mcmc` run of `RevBayes`. We will then run another analysis for which we will have samples from the prior densities to compare with the full analysis which samples from the posterior. We should have now two trace or log files with parameter values.

The package `tbea` provides graphical methods for comparing prior and posterior densities (see the vignette `prior_posterior_comparisons`). For instance, the cross-plot can be used with a pair of log files, one for the prior (the x -axis) and one for the posterior (the y -axis) of the node ages in a divergence time analysis, so that the x, y pairs can represent prior and posterior node age means or medians, and error bars represent the 95% highest posterior density interval (Fig. 5A). This kind of plot has been used in the literature when comparing prior and posterior MCMC samples, as well as when comparing the same kind of estimates coming from different independent runs or types of analysis (e.g. Drummond and Stadler 2016, Álvarez-Carretero *et al.* 2022). This plot allows us to see clearly that the prior node densities have more variance than the posterior ones.

Although visual comparison is useful and advisable, in some situations we may be interested in actually measuring the similarity between two densities, such as when comparing different calibration densities for the same analysis. We may also want to represent the degree of similarity between prior and posterior densities for a given parameter quantitatively. We can define similarity between two densities as the area under the intersection of both density curves: when both density curves are identical, their intersection is 1.0, whereas when they are completely different

and disjoint, the intersection measures 0.0. The function `measuresSimil` achieves this and thus provides a measure of how similar two distributions are. The function can both plot the resulting distributions and their intersection (Fig. 5B) as well as print out its value, or skip the plot and just return the value. When comparing the prior and posterior densities for a calibrated node depicting the genus *Hydrolycus*, we see that their intersection is quite large, 0.83, which may represent strong sensitivity of the posterior to the prior, and thus little impact of the likelihood. This strong sensitivity of the posterior is expected from analyses with few calibration points (Brown and Smith 2018).

Inference of origination time given a set of time estimates

Let X be an arbitrary biogeographical area in geological time bearing a biota B_X . At a given point in time, physical earth processes divide the area in two daughter areas Y and Z in a process commonly known as vicariance:

$$X \rightarrow Y, Z$$

The expectation is that the biota B_X also separates and reacts to this interruption of gene flow by speciation, and thus creates a pattern of repeated sister-species pairs between areas Y and Z (Fig. 6A), although it might also undergo extinction at a given rate, but for simplicity we assumed first that there is no extinction taking place, so that

$$B_X \rightarrow B_Y, B_Z$$

Note that the pattern preserved may also be subject to two obscuring processes: lack of speciation, and immigration from adjacent areas followed or not by speciation, thus creating non- Y ,




Figure 5. A, cross-plot with a comparison between the prior and the posterior of node ages in the divergence time estimation of the Cynodontidae dataset. Note that the prior node densities have more variance than the posterior ones, which is a consequence of using the data through the likelihood function. B, similarity between the prior (red) and posterior (blue) densities of the node *Hydrolycus*. The similarity, which is the value of the intersection (grey) between both distributions, is 0.83.




Figure 6. The inferential model and two possible analytical strategies depicted. A, the separation of an arbitrary biogeographical area X along geological time that produces daughter areas Y and Z ; after a lag in time after vicariance, the biota responds by speciating and the collection of sister pairs distributed in daughter areas inform about the original time of separation t_0 . B, the specification of a stratigraphic interval and with endpoint parameters θ_1, θ_2 , and all the occurrences in time $t_{i:N}$; with such information it is possible to construct confidence intervals at a nominal α value for said parameters, in particular θ_1 . C, the empirical density of the divergence time events that allow us to re-characterize it as a CDF and use regression for estimation of the x -intercept.

Z geographical patterns. We will again ignore these forces for the sake of simplicity. Note that non- Y, Z patterns are uninformative on the process because they may happen before, coeval with, or after the event $X \rightarrow Y, Z$.

Now imagine that the ancestral area X previously consisting of continuous lowlands with a common drainage network now undergoes vicariance, for instance through mountain uplift, thus imposing a physical barrier to drainages that creates two different watersheds, and also creates montane habitats. In sum, these two phenomena block gene flow in aquatic organisms, and promote speciation given that the lowland terrestrial biota is unable to get across mountain regions mostly due to physiological constraints. As the process also cuts the former drainage network creating two or more drainages that lack connectivity, the freshwater biota also loses gene flow, and responds by speciating. However, this process is not instantaneous in time, as both the geological processes will be in the scale of millions of years, and their response in the biota will also take some time after complete area separation. This last phenomenon can be thought of as a lag between actual area separation t_0 and the first speciation event t_1 that will depend on how fast a particular biological group will respond to blocking of gene flow. Then, a number of speciation events in time $t_{i:N}$ later than the true area separation t_0 will take place as a consequence of the vicariance event (Fig. 6A). The vector of speciation events is implicitly assumed to be measured without error because the

probability that either endpoint parameter is actually present within the observed interval is zero following the definition of the confidence intervals as expansions beyond the first and last observed occurrences in time (cf. Equations 2 and 3). Given this inferential issue, outlier removal procedures might be necessary in order to remove suspicious data that are contradicted by the overall distribution of data (e.g. a divergence estimate that is too old in geological time, or based on misleading calibration information).

As we lack direct evidence of the event $X \rightarrow Y, Z$, we aim at using information preserved in the statistical distribution of events $t_{i:N}$ to estimate the time at which the vicariance event took place, t_0 in the notation above. Herein we propose two alternative ways of looking at the statistical structure of these phenomena in time: first, as a stratigraphic interval formed by the unobserved origination of the vicariance event and its end (perhaps unnecessary), along with the collection of events in time that resulted from it (Fig. 6B), and second, as the distribution of these events in time (as well as its cumulative distribution) (Fig. 6C).

Different estimates of the same parameter are often of interest in meta-analyses as well as the means of specifying secondary node calibrations. However, how to summarize such information is not straightforward (Hill 2011, Hill and Miller 2011). Here, we implemented three methods in which estimates from different distributions measuring the same parameter can be summarized in a single

distribution: (i) using stratigraphic intervals (see vignette `stratigraphic_intervals`), (ii) using x -intercept confidence intervals (see vignette `x_intercept`), and (iii) using conflation of distributions (see vignette `conflation`). As an example, we estimated the age of the separation of drainages east and west of the Andes by using the first two approaches, and then show a hypothetical instance of secondary calibration using conflation.

Stratigraphic intervals

The vicariant model presented earlier shows an interesting distribution of events in time: they start with the origination time, and then happen for some time until the last component of the biota responds to the separation event. This pattern of events in time can be described by a stratigraphic interval. This is a collection of events in time for which the unobserved origination and extinction times are unknown. All the information we have comes from the set of occurrences in time (Fig. 7). These models have been extensively used in paleobiology and some flexible implementations have become recently available (Ballen 2025). Although defined for species in the fossil record, we propose extending their meaning to something more general, such as sets of general occurrences in time. This also includes their original applications to sets of fossil occurrence times for a given species but also encompass sets of biogeographical occurrences such as vicariance events preserved in the phylogenies of multiple groups and triggered by a paleogeographical event, for instance drainage separation due to orogeny. It is important to note that nothing in the mathematical formulation of stratigraphic interval models is unique to species and therefore its application can be expanded without loss of correctness. The only difference in interpretation is that preservation parameters are now describing the distribution pattern of events in time.

Frequentist approaches provide ways to estimate the confidence intervals on the endpoints of a stratigraphic interval, that is an estimate for the point of appearance and disappearance of a given taxon in geological time. These methods are relevant here

since we can think of a biogeographical events in the same way as taxa appear and become extinct. This approach has so far not been used in historical biogeography and seems to fit the main goal of determining the final point of connection between biogeographical areas. Marshall (2010) summarized different approaches for estimating confidence intervals under certain sampling regimes, assumptions and confidence levels, ranging from classical to Bayesian estimators. Ballen (2025) presented an overview of the different methods for stratigraphic intervals with some additions posterior to the work of Marshall. The package `tbea` implements two of these methods, the constant-preservation estimator of Strauss and Sadler (1989) and the distribution-free estimator of Marshall (1994).

The first method provides estimators where one- (θ_1) or two-parameter (θ_1, θ_2) cases are constructed adding and subtracting a portion α of the magnitude of the observed stratigraphic interval $[y, z]$. Such a constant depends on the confidence level wanted (e.g. 95%).

$$y - \alpha(z - y) < \theta_1 < y \quad (2)$$

$$z < \theta_2 < z + \alpha(z - y) \quad (3)$$

The constant α can be calculated either exactly (one-parameter case) or iteratively (two-parameter case).

The method of distribution-free confidence intervals was developed without assuming any particular distribution of underlying gap sizes, and instead calculates quantiles of gap size between occurrence points for an ordered vector of gap sizes. The major cost of relaxing the assumption of a parametric approach in comparison with the constant-preservation method is that such intervals are larger, whereas the distribution-free approach requires much more data for constructing intervals of the same width as the former (e.g. 95%).

The method works by constructing a confidence interval that has a level of confidence C (i.e. C th percentile) for a γ confidence probability with the information from N gaps.




Figure 7. Stratigraphic interval. Suppose we have an Ammonite species through time. The species originates at time θ_1 and becomes extinct at time θ_2 . Along this interval, it leaves preserved occurrences at times $\tau_1, \dots, \tau_6$. The underlying distribution here is Uniform, which is governed by the preservation potential, and therefore the realizations of the preservation process. We estimate the origination and extinction times, as well as the preservation parameters from the data. The horizontal axis represents time from present (in Myr) and vertical axis represents density in the underlying distribution.

The lower bound is then the i th gap in the ordered vector applied to $\min(\text{vector}) - x_i$ where x_i is the magnitude of such i th element in the vector:

$$\gamma > \sum_{x=0}^x \binom{N}{x} C^x (1-C)^{N-x} \quad \text{for } x \leq N \quad (4)$$

The upper limit can be constructed with the formula:

$$(1-\gamma) < \sum_{x=0}^x \binom{N}{x} C^x (1-C)^{N-x} \quad \text{for } x \leq N \quad (5)$$

The method makes use of the binomial distribution to find the quantiles that satisfy the condition of being smaller than γ and from these picking the largest (in the case of the lower bound) and the smallest (for the upper bound). It is noteworthy that if only one x satisfies Equations (4) or (5), no lower or upper confidence interval can be constructed. Also, if $x = N$ in Equation (5), no upper confidence interval can be constructed since there is no $(x+1)$ th gap in the vector. Marshall (1994) noted that inference using his method is only valid as long as there is no trend in gap size (i.e. a trend of increase or decrease in gap size as we go along the time data). It is possible to use the Mann–Kendall test for a monotonic trend to test such assumption (Mann 1945, Kendall 1948). Because the method works with quantiles of the sorted gap sizes, these are first calculated from the sorted vector of times, but then internally sorted when using the method ‘Marshall94’; however, whether there is any monotonic trend must be assessed on the unsorted version. We use the Mann–Kendall test on the times vector prior to applying the function.

These two methods are available in the function `stratCI` which allows us to pick one of two methods: ‘`strauss-sadler89`’ or ‘`marshall94`’.

As Marshall (1994) noted, the method requires a large number of gaps to calculate both bounds for the 95% confidence interval; for our dataset with just 42 gaps (i.e. 43 times), we can find both bounds at the 80% confidence level (or lower), and thus claim that the lower and upper bounds for the 95% confidence interval are 7.07–7.2 Mya. When we raise the confidence level to 95%, we can only infer the lower bound, thus being able just to claim that the 95% confidence interval is at least as old as 7.14 Mya. The higher the confidence level for finding the bounds on the 95% confidence interval, the better, in a similar way as calculating a 95% confidence interval is better than calculating a 70% confidence interval.

x-Intercept inference

Let us go back to the vicariance model presented earlier. One consequence of the separation event is that the biota does not necessarily respond immediately but instead with some lag depending on the lability to interruption of gene flow imposed by the separation event. For instance, in drainage separation via orogeny we expect that aquatic organisms will respond faster than terrestrial and volant ones. There is always some lag, small or not, of the accumulation pattern of vicariant events in time.

These separation events can be described as a distribution regardless of its generator process, which will follow a PDF $f(t)$, with an associated cumulative density function (CDF) $P(X < x)$.

We can fit an empirical CDF and then use a linear regression model. The only case where a CDF would be truly linear would be in the case of a Uniform distribution where

$$P(X \leq x) = \frac{x}{b-a} \quad \text{for } x \in [a, b]$$

which implies that the empirical data in time follow a Uniform distribution. For any other case, a fitted CDF would be non-linear, and thus a linearization should be applied to it. However, it is also possible to use robust non-parametric regression in order to avoid the issues with non-linearity of the CDF. Such a regression model with an additional parameter x -intercept may be thought of as an estimator for the initial time t_0 given the trend present in the linear model (Fig. 8).

Estimation of uncertainty regions such as confidence intervals on the x -intercept can be even more useful than the point estimate itself. Such regions can be constructed by several methods, for instance using bootstrap or analytical formulae for the predictive region or confidence intervals.

The linear model has the form

$$y = mx + b$$

where m is the slope and b is the y -intercept. However, we are not interested in any of these parameters but instead in a third one,




Figure 8. Empirical cumulative distribution of divergence time events. Above, plot of the empirical density; below, model fit of the cumulative distribution. The interrupted line represents the regression model on the cumulative values, and the red diamond represents the x -intercept.

the x -intercept, which is defined as the point in x where $y = 0$. This parameter can be defined by letting $y = 0$ in the equation and thus

$$x_{y=0} = -\frac{b}{m}$$

allows us to find the point at which the data began to be generated.

Several methods have been proposed for estimating the uncertainty on the x -intercept estimate, from the intuitive inverse regression $X \sim Y$ to bootstrapping methods. According to [Seber and Lee \(2013, pp. 146–147\)](#) this problem has long been discussed in the statistical literature without much consensus, although some proposals are of interest due to theoretical appeal, ease of implementation, or simplification of assumptions. Uncertainty on the estimate of the x -intercept can be obtained through Taylor series simplification for symmetric intervals. Another approach by [Draper and Smith \(1998, pp. 83–86\)](#) allows the calculation of asymmetric confidence intervals by projecting the values of the confidence region around the linear model on the x -axis by manipulation of the curves describing the confidence region, which gives

$$CI_{X_0} = \hat{X}_0 + \frac{(X_0 - \bar{X})g \pm (ts/b_1)\{[(\hat{X}_0 - \bar{X})^2 S_{XX}] + (1-g)/n\}^{1/2}}{1-g}$$

where $g = t^2 s^2 / (b_1^2 / S_{XX})$, b_1 is the slope, S_{XX} is the sum of squares of X , t is the t -statistic with $n-2$ degrees of freedom and at $1-\alpha/2$ significance level, s^2 is the standard deviation, $\hat{X}_0$ is the estimation

of the x -intercept, and $\bar{X}$ is the mean of X . Here we need to fit a linear model in order to use the curves for the confidence region, and we can use classical least square models or robust non-parametric regression ([Kloke and McKean 2014](#)).

It is also possible to use bootstrapping for estimating the variance and thus the confidence interval in asymmetric cases such as this one. Instead of using the estimates from a single fitted model on the coordinates of the empirical CDF, this method fits multiple models with random subsamples in to construct a set of estimated values of the x -intercept and this provides a confidence interval.

The package `tbea` implements estimation of the confidence interval for the x -intercept using the method of [Draper and Smith \(1998\)](#) as well as estimation via bootstrap (see vignette `x_intercept`). Both require a linear model which can be set to one of the following: either ordinary least squares or a robust linear model ([Kloke and McKean 2014](#)).

Both methods produce similar results ([Fig. 9B](#)). Although the method based on robust regression produced a confidence interval wider than the one based on ordinary least squares, the bootstrap confidence interval estimate was as wide as the two former combined. As the bootstrapping based on robust estimates is of more general application, it seems more adequate for these situations. It is noteworthy that these inferences are based on the assumption that the events in time are measured without error. However, this assumption relaxes due to the effect of re-sampling and uncertainty in the estimates. According to both estimators, the separation between cis- and trans-Andean areas took place between ca. 4.3 and 5.8 Mya.




Figure 9. Estimation of the time of separation between cis- and trans-Andean drainages. A, based on the methods of stratigraphic intervals of [Strauss and Sadler \(1989\)](#) and [Marshall \(1994\)](#). B, based on the methods of CDF fitting using the estimator of [Draper and Smith \(1998\)](#) with two different approaches at model fitting, one using least squares traditional regression, and the other using robust regression ([Kloke and McKean 2014](#)) and bootstrapping techniques. Note that confidence intervals are near but not bound to be older than the oldest occurrence in methods based on CDF fitting, whereas they are forced to be older than the oldest occurrence in estimates based on stratigraphic intervals. This property of both methods can help to decide which is more reasonable to apply in specific circumstances, for example when the ages of the oldest occurrences are not very reliable.

These methods have produced results in line with those based on stratigraphic intervals. However, they allow x -intercepts and consequently their confidence intervals that are not necessarily older than the oldest divergence time estimation event. Thus, they are robust to the assumption of measurement without error that is implicit in stratigraphic interval methods. Nevertheless, they are still sensitive to the presence of strong outliers as preliminary analyses showed that the regressions on these data with outliers strongly affected the robust regressions.

Conflation of distributions

Combining different distributions into a single one is a useful tool for summarizing different sources of information. Given a set of N distributions describing the same parameter, we aim at building a general distribution which incorporates information from the individual components. Multiple alternatives are available, and as with x -intercept estimation, the problem of summarizing a set of distributions has several possible answers (Genest and Zidek 1986). Hill (2011) discussed these methods and proposed conflation as a better and general way of combining distributions describing the same parameter, and introduced the untary operator ‘&’ as representing conflation. Let $Q = \&(P_1, \dots, P_N)$ be the conflation of N densities describing the same time event τ :

$$Q(\tau; \Theta) = \frac{\prod_{i=1}^N P(\tau_i; \Theta_i)}{\int_{-\infty}^{\infty} \prod_{i=1}^N P(x; \Theta_i) dx} \quad (6)$$

Among the useful properties of conflation, we have that Q is also a distribution, which can be used for calculating probabilities. Also, the conflation combines the variance of each distribution, so that intuitively densities with more uncertainty (i.e. higher variance) will have a lower impact on the conflated distribution than others with lower variance, thus weighting the informativeness of all individual distributions appropriately.

Conflation proves useful when building secondary calibrations where more than two densities are available for the same time event. Suppose that we have a biogeographical event with time τ , for which three different studies provide three independent densities, for instance, posterior probabilities for a node which corresponds to the biogeographical event. The best summary of these independent pieces of information is a conflation $Q(\tau) = \&(P_1(\tau), P_2(\tau), P_3(\tau))$, which can be used as a secondary calibration point for a new divergence time estimation analysis as it is itself a PDF. Suppose that $P_1 = N(10, 1)$, $P_2 = N(13, 0.5)$, and $P_3 = N(14.4, 1.7)$. We can calculate the conflation of these three distributions.

Here, three Normals were used to approximate the posterior distribution of a given divergence time (Fig. 10), and their conflation is depicted as a black line. It is evident that the conflated distribution shows more resemblance to the component with lowest variance (blue distribution), although incorporating the information from all of them. This distribution can in turn be used as a secondary calibration point in a subsequent divergence time estimation analysis, where a single node calibration must be used for a given node.

The function `conflate` is used for building a conflation given arbitrary PDFs available in, that is, any function `dDIST`




Figure 10. Conflation of three distributions representing imaginary posterior densities approximated by Normals of the same divergence time in million years (red, blue, and dark green lines). The conflated distribution of these three represents the information coming from each proportional to their variance (black line).

(`dnorm`, `dexp`, etc.). It uses a constructor function `density_fun` that helps to specify the functions without need to calculate any object beforehand. The function `quantile_conflation` can then be used on its output for calculating quantiles on such distribution, for instance, for approximating a Normal distribution on these quantiles for calibration specification.

DISCUSSION

Bayesian phylogenetic methods have had a dramatic impact in evolutionary biology due in part to their ability to integrate multiple sources of information and provide a natural way of characterizing uncertainty in parameter estimation. One of the challenges in empirical analyses is reproducibility, although some programs have internal tools for improving it. Even though model specification is often carried out through the graphical user interface program `beauti`, `Beast2` will return the complete XML specification of the analysis as a commented header of the log file containing the parameter values from the MCMC sampling. The downside of this strategy is that the XML specification is very challenging to understand by readers (although convenient from a development point of view). Both `MrBayes` and `RevBayes` support interactive sessions and the models can be specified piece by piece. However, their highest reproducibility standard is reached through the use of scripts, where the complete model is documented in code. `MCMCTree` is a special case where interactive mode is restricted to entering the path to the control file. The complete analysis needs to be specified in the control file and therefore it qualifies as a script for practical purposes. The consequence is that any analysis specified in scripts is easier to understand and reproduce. Nonetheless, all these strategies are unable to document how the users chose the specific models, their parameters, and their distributions (e.g. priors). The pre-analysis tools in the present package aim at allowing the user not just to find the

best specification of distributions and model components, but also to document them as R scripts that can be shared alongside the analysis scripts. The same applies for post-analysis tools such as the comparison between model priors and posteriors, which are often verbally mentioned to have been done but not represented graphically in publications.

Techniques for estimating the origin time for a collection of events can be applied beyond the vicariant case presented herein. It is possible to use the same approach for a collection of events such as biotic interchanges (e.g. the Great American Biotic Interchange across the Panamá Isthmus) or colonizations (e.g. the invasion of South America by Gharials during the Cenozoic), which are, in the end, collections of events happening in time but with a single origination time. Finally, the tools for estimating events in time from sets of point estimates, as well as the combination of different posterior densities from the same parameter, are expected to improve the transparency of research choices when, for example, justifying the choice of secondary calibration points, or discussing the timing of biogeographical events when multiple sources are available for them. An additional advantage of this toolset is to integrate information from multiple areas of biological and earth sciences. They allow us to make a better use of geological and paleontological information for the specification of priors, as well as for estimating events in time, as some of the models (e.g. stratigraphic intervals) have been transferred here from paleobiology to the analysis of Bayesian estimates. Distribution conflation or estimation of parameters in stratigraphic intervals using estimates from divergence time estimation analyses, in turn, can be helpful for research in paleogeography and tectonics as they provide an independent source to contrast the predictions made with the aid of geological tools. In the end, both areas are concerned with events in geological time.

ACKNOWLEDGEMENTS

We thank the feedback by students at the Universidad de Antioquia in Colombia during a graduate course on phylogenetics in 2021, as well as the feedback from students in the graduate course Efficient Divergence Time Estimation using MCMCTree at Universidade Estadual Paulista in 2024, and help by Klauss Schliep when implementing metaprogramming code for numerical optimization. We thank both Uwe Ligges and Benjamin Altmann for their feedback during submission to CRAN which helped improve the documentation of the package. We thank T. Simões for feedback on functionalities of `Evophylo`. We thank the excellent and detailed feedback from the editor T. Guillerme and two anonymous reviewers which helped us improve the manuscript.

AUTHOR CONTRIBUTIONS

Gustavo A. Ballen (Conceptualization [Equal], Data curation [Equal], Formal analysis [Equal], Funding acquisition [Equal], Investigation [Equal], Methodology [Equal], Project administration [Equal], Resources [Equal], Software [Equal], Supervision [Equal], Validation [Equal], Visualization [Equal], Writing—original draft [Equal], Writing—review & editing [Equal]) and Sandra Reinales (Conceptualization [Equal], Data curation [Equal], Formal analysis [Equal], Funding acquisition [Equal], Investigation

[Equal], Methodology [Equal], Project administration [Equal], Resources [Equal], Software [Equal], Supervision [Equal], Validation [Equal], Visualization [Equal], Writing—original draft [Equal], Writing—review & editing [Equal]).

CONFLICT OF INTEREST

The authors declare no conflicts of interest.

FUNDING

This work was funded by FAPESP processes 2014/11558-5, 2016/02253-1, and 2023/07838-1 to G.A.B., 2020/13433-6 to Cláudio de Oliveira, and 2020/09442-0, 2020/10206-9 and 2022/01533-1 to S.R.

DATA AVAILABILITY STATEMENT

The R package `tbea` is hosted on CRAN (<https://cran.r-project.org/package=tbea>) and development versions are available on GitHub (<https://github.com/gaballench/tbea>). The stable version of the package can be installed with `library('tbea')` and the development version with `remotes::install_github('gaballench/tbea')`. The stable version herein described is v.1.7.0, with DOI 10.32614/CRAN.package.tbea.

REFERENCES

- Abe KT, Mariguela TC, Avelino GS *et al.* Systematic and historical biogeography of the Bryconidae (Ostariophysi: Characiformes) suggesting a new rearrangement of its genera and an old origin of Mesoamerican ichthyofauna. *BMC Evolutionary Biology* 2014;**14**:152.
- Álvarez-Carretero S, Goswami A, Yang Z *et al.* Bayesian estimation of species divergence times using correlated quantitative characters. *Systematic Biology* 2019;**68**:967–86.
- Álvarez-Carretero S, Tamuri AU, Battini M *et al.* A species-level timeline of mammal evolution integrating phylogenomic data. *Nature* 2022;**602**:263–7.
- Ballen GA. A flexible Bayesian method for estimating stratigraphic intervals and their co-occurrence in time. *bioRxiv*, 2025, pages 2025–02, preprint: not peer reviewed.
- Ballen GA, Moreno-Bernal JW, Jaramillo C. The fossil record of Sabretooth characins (Teleostei: Characiformes: Cynodontinae), their phylogenetic relationships and palaeobiogeographical implications. *Journal of Systematic Palaeontology* 2021;**19**:1679–92.
- Ballen GA, Mühlová KH, Hajič Jr J. What did the dove sing to Pope Gregory? Ancestral melody reconstruction in Gregorian chant using Bayesian phylogenetics. *bioRxiv* 2024, pages 2024-08, preprint: not peer reviewed.
- Barba-Montoya J, Tao Q, Kumar S. Molecular and morphological clocks for estimating evolutionary divergence times. *BMC Ecology and Evolution* 2021;**21**:83.
- 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. *Proceedings. Biological Sciences* 2019;**286**:20190685.
- Bilderbeek RJ, Etienne RS, babette: Beauti2, Beast2 and tracer for R. *Methods in Ecology and Evolution* 2008;**9**:2034–40.
- Bouckaert R, Vaughan TG, Barido-Sottani J *et al.* Beast 2.5: an advanced software platform for Bayesian evolutionary analysis. *PLoS Computational Biology* 2019;**15**:e1006650.

- Brown JW, Smith SA. The past sure is tense: on interpreting phylogenetic divergence time estimates. *Systematic Biology* 2018;**67**:340–53.
- Cheviron ZA, Hackett SJ, Capparella AP. Complex evolutionary history of a Neotropical lowland forest bird (*Lepidothrix coronata*) and its implications for historical hypotheses of the origin of Neotropical avian diversity. *Molecular Phylogenetics and Evolution* 2005;**36**:338–57.
- Claramunt S. Cladate: Calibration information generator for divergence time estimation. *Methods in Ecology and Evolution* 2022;**13**:2331–8.
- Collins AC, Dubach JM. Biogeographic and ecological forces responsible for speciation in *Ateles*. *International Journal of Primatology* 2000;**21**:421–44.
- Cortés-Ortiz L, Bermingham E, Rico C *et al.* Molecular systematics and biogeography of the Neotropical monkey genus, *Alouatta*. *Molecular Phylogenetics and Evolution* 2003;**26**:64–81.
- D'Horta FM, Cuervo AM, Ribas CC *et al.* Phylogeny and comparative phylogeography of *Sclerurus* (Aves: Furnariidae) reveal constant and cryptic diversification in an old radiation of rain forest understorey specialists. *Journal of Biogeography* 2013;**40**:37–49.
- Dick CW, Abdul-Salim K, Bermingham E. Diversification of a widespread tropical rain forest tree. *The American Naturalist* 2003;**162**:691–703.
- dos Reis M, Gunnell GF, Barba-Montoya J *et al.* Using phylogenomic data to explore the effects of relaxed clocks and calibration strategies on divergence time estimation: Primates as a test case. *Systematic Biology* 2018;**67**:594–615.
- Draper NR, Smith H. *Applied Regression Analysis*. New York: Wiley Interscience, 1998.
- Drummond AJ, Stadler T. Bayesian phylogenetic estimation of fossil ages. *Philosophical Transactions of the Royal Society B: Biological Sciences* 2016;**371**:20150129.
- Du Plessis L. Beastio: Functions for Pre-and Post-processing of BEAST and BEAST2 Files, 2020. <https://github.com/laduplessis/beastio/>
- Elias M, Joron M, Willmott K *et al.* Out of the Andes: patterns of diversification in clearwing butterflies. *Molecular Ecology* 2009;**18**:1716–29.
- Fabreti LG, Höhna S. Convergence assessment for Bayesian phylogenetic analysis using MCMC simulation. *Methods in Ecology and Evolution* 2022;**13**:77–90.
- Fernandes AM, Wink M, Sardelli CH *et al.* Multiple speciation across the Andes and throughout Amazonia: The case of the spot-backed antbird species complex (*Hylophylax naevius*/*Hylophylax naevioides*). *Journal of Biogeography* 2014;**41**:1094–104.
- Gavryushkina A, Welch D, Stadler T *et al.* Bayesian inference of sampled ancestor trees for epidemiology and fossil calibration. *PLoS Computational Biology* 2014;**10**:e1003919.
- Gearty W, O'Meara B, Berv J, Ballen GA, Ferreira D, Lapp H, Schmitz L, Smith MR, Upham NS, Nations JA. Cran task view: Phylogenetics, 2024. <https://CRAN.R-project.org/view=Phylogenetics>
- Genest C, Zidek JV. Combining probability distributions: a critique and an annotated bibliography. *Statistical Science* 1986;**1**:114–35.
- Goloboff PA, Farris JS, Nixon KC. TNT, a free program for phylogenetic analysis. *Cladistics* 2008;**24**:774–86.
- Grau ET, Pereira SL, Silveira LF *et al.* Molecular phylogenetics and biogeography of Neotropical piping guans (Aves: Galliformes): *Pipile Bonaparte*, 1856 is synonym of *Aburria* Reichenbach, 1853. *Molecular Phylogenetics and Evolution* 2005;**35**:637–45.
- Greenhill SJ. Language phylogenies: Modelling the evolution of language. In: *The Oxford Handbook of Cultural Evolution*. Oxford: Oxford University Press, 2023.
- Gutiérrez EE, Anderson RP, Voss RS *et al.* Phylogeography of *Marmosa robinsoni*: insights into the biogeography of dry forests in northern South America. *Journal of Mammalogy* 2014;**95**:1175–88.
- Hajic Jr. J, Ballen GA, Muhlova KH, Vlhova-Worner H. Towards building a phylogeny of Gregorian chant melodies. In: *Proceedings of the 24th International Society for Music Information Retrieval Conference*. pp. 571–8. 2023.
- Hajic Jr. J, Lanz V, Ballen GA. Genome of melody: Applying bioinformatics to study the evolution of Gregorian chant. *EcoEvoRxiv* 2025, pages 1–27.
- Hardman M, Lundberg JG. Molecular phylogeny and a chronology of diversification for “phractocephaline” catfishes (Siluriformes: Pimelodidae) based on mitochondrial DNA and nuclear recombination activating gene 2 sequences. *Molecular Phylogenetics and Evolution* 2006;**40**:410–8.
- Hazzi NA, Moreno JS, Ortiz-Movliav C *et al.* Biogeographic regions and events of isolation and diversification of the endemic biota of the tropical andes. *Proceedings of the National Academy of Sciences of the United States of America* 2018;**115**:7985–90.
- Heath TA, Huelsenbeck JP, Stadler T. The fossilized Birth-Death Process: a coherent model of fossil calibration for divergence time estimation. *Proceedings of the National Academy of Sciences* 2014;**111**:E2957–E2966.
- Heled J, Drummond AJ. Bayesian inference of species trees from multilocus data. *Molecular Biology and Evolution* 2010;**27**:570–80.
- Hernández Torres CL. Species delimitation, phylogenetics, and biogeography of the catfish genus *Rhamdia* Bleeker (Heptapteridae) of Central America and the Trans-Andean Region of Colombia. Ph.D. Thesis, The University of Southern Mississippi, 2015.
- Hill TP. Conflations of probability distributions. *Transactions of the American Mathematical Society* 2011;**363**:3351–72.
- Hill TP, Miller J. How to combine independent data sets for the same quantity. *Chaos (Woodbury, N.Y.)* 2011;**21**:033102.
- Hoffmann K, Bouckaert R, Greenhill SJ *et al.* Bayesian phylogenetic analysis of linguistic data using Beast. *Journal of Language Evolution* 2021;**6**:119–35.
- Inoue J, Donoghue PCJ, Yang Z. The impact of the representation of fossil calibrations on Bayesian estimation of species divergence times. *Systematic Biology* 2010;**59**:74–89.
- Jombart T, Archer F, Schliep K *et al.* apex: Phylogenetics with multiple genes. *Molecular Ecology Resources* 2017;**17**:19–26.
- Kendall MG. *Rank Correlation Methods*. Griffin, 1948.
- Kloke JD, McKean JW. Rfit: rank-based estimation for linear models. *The R Journal* 2012;**4**:57–64.
- Kloke JD, McKean JW. *Nonparametric Statistical Methods Using R*, 2nd ed. Boca Raton, FL: CRC Press, 2014.
- Landis MJ. Biogeographical dating of speciation times using paleogeographically informed processes. *Systematic Biology* 2017;**66**:128–44.
- Lewis PO. A likelihood approach to estimating phylogeny from discrete morphological character data. *Systematic Biology* 2001;**50**:913–25.
- Machado LF, Leite YLR, Christoff AU *et al.* Phylogeny and biogeography of tetralophodont rodents of the tribe Oryzomyini (Cricetidae: Sigmodontinae). *Zoologica Scripta* 2014;**43**:119–30.
- Mann HB. Nonparametric tests against trend. *Econometrica* 1945;**13**:245–59.
- Marshall CR. Confidence intervals on stratigraphic ranges: Partial relaxation of the assumption of randomly distributed fossil horizons. *Paleobiology* 1994;**20**:459–69.
- Marshall CR. Using confidence intervals to quantify the uncertainty in the end-points of stratigraphic ranges. *The Paleontological Society Papers* 2010;**16**:291–316.
- Matzig DN, Marwick B, Riede F *et al.* A macroevolutionary analysis of European late upper palaeolithic stone tool shape using a bayesian phylogenetic framework. *Royal Society Open Science* 2024;**11**:240321.
- Maurits L, Griffiths TL. Tracing the roots of syntax with Bayesian phylogenetics. *Proceedings of the National Academy of Sciences of the United States of America* 2014;**111**:13576–81.
- Mendes FK, Hahn MW. Why concatenation fails near the anomaly zone. *Systematic Biology* 2018;**67**:158–69.
- Miller MJ, Bermingham E, Klicka J *et al.* Out of Amazonia again and again: episodic crossing of the Andes promotes diversification in a lowland forest flycatcher. *Proceedings of the Royal Society B: Biological Sciences* 2008;**275**:1133–42.
- Mongiardino Koch N, Milla Carmona P. Chronospaces: an r package for the statistical exploration of divergence times promotes the assessment of methodological sensitivity. *Methods in Ecology and Evolution* 2024;**15**:1822–33.
- Neureiter N, Ranacher P, Efrat-Kowalsky N *et al.* Detecting contact in language trees: a Bayesian phylogenetic model with horizontal transfer. *Humanities and Social Sciences Communications* 2022;**9**:1–14.

- O'Reilly JE, Dos Reis M, Donoghue PC. Dating tips for divergence-time estimation. *Trends in Genetics* 2015;**31**:637–50.
- Paradis E, Schliep K. ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. *Bioinformatics (Oxford, England)* 2019;**35**:526–8.
- Parham JF, Donoghue PCJ, Bell CJ *et al.* Best practices for justifying fossil calibrations. *Systematic Biology* 2012;**61**:346–59.
- Patané JSL, Weckstein JD, Aleixo A *et al.* Evolutionary history of *Ramphastos* toucans: molecular phylogenetics, temporal diversification, and biogeography. *Molecular Phylogenetics and Evolution* 2009;**53**:923–34.
- Patel S, Weckstein JD, Patané JSL *et al.* Temporal and spatial diversification of *Pteroglossus Araçaris* (Aves: Ramphastidae) in the neotropics: constant rate of diversification does not support an increase in radiation during the Pleistocene. *Molecular Phylogenetics and Evolution* 2011;**58**:105–15.
- Pérez-Escobar OA, Lucas E, Jaramillo C *et al.* The origin and diversification of the hyperdiverse flora in the Chocó biogeographic region. *Frontiers in Plant Science* 2019;**10**:1328.
- Picq S, Alda F, Krahe R *et al.* Miocene and Pliocene colonization of the Central American Isthmus by the weakly electric fish *Brachyhyopomus occidentalis* (Hypopomidae, Gymnotiformes). *Journal of Biogeography* 2014;**41**:1520–32.
- Plummer M, Best N, Cowles K *et al.* Coda: convergence diagnosis and output analysis for mcmc. *R News* 2006;**6**:7–11.
- R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing, 2024.
- Rambaut A. Figtree v1.4.4, 2018. <https://github.com/rambaut/figtree>
- Rannala B, Yang Z. Inferring speciation times under an episodic molecular clock. *Systematic Biology* 2007;**56**:453–66.
- Ribas CC, Gaban-Lima R, Miyaki CY *et al.* Historical biogeography and diversification within the Neotropical parrot genus *Pionopsitta* (Aves: Psittacidae). *Journal of Biogeography* 2005;**32**:1409–27.
- Ribas CC, Moyle RG, Miyaki CY *et al.* The assembly of montane biotas: linking Andean tectonics and climatic oscillations to independent regimes of diversification in *Pionus* parrots. *Proceedings. Biological Sciences* 2007;**274**:2399–408.
- Říčan O, Pálek L, Zardoya R *et al.* Biogeography of the Mesoamerican Cichlidae (Teleostei: Heroini): colonization through the GAARlandia land bridge and early diversification. *Journal of Biogeography* 2013;**40**:579–93.
- Rink WJ, Thompson JW. *Encyclopedia of Scientific Dating Methods*. Berlin: Springer, 2015.
- Román-Palacios C. The phrta R package: increasing access, reproducibility and transparency in phylogenetic analyses. *Methods in Ecology and Evolution* 2023;**14**:2284–99.
- Ronquist F, Klopstein S, Vilhelmsen L *et al.* A total-evidence approach to dating with fossils, applied to the early radiation of the Hymenoptera. *Systematic Biology* 2012;**61**:973–99.
- Ruiz-García M, Vásquez C, Sandoval S *et al.* Phylogeography and spatial structure of the lowland tapir (*Tapirus terrestris*, Perissodactyla: Tapiridae) in South America. *Mitochondrial DNA* 2015;**1736**:1–9.
- Salinas NR, Eshel G, Coruzzi GM *et al.* Bad2matrix: Phylogenomic matrix concatenation, indel coding, and more. *Applications in Plant Sciences* 2024;**12**:e11604.
- Sanderson MJ. Estimating absolute rates of molecular evolution and divergence times: a penalized likelihood approach. *Molecular Biology and Evolution* 2002;**19**:101–9.
- Schenk JJ. Consequences of secondary calibrations on divergence time estimates. *PLoS One* 2016;**11**:e0148228.
- Seber GAF, Lee AJ. *Linear Regression Analysis*, 2nd ed. Hoboken, NJ: Wiley Interscience, 2013.
- 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 in Ecology and Evolution* 2023;**14**:1981–93.
- Smith BT, McCormack JE, Cuervo AM *et al.* The drivers of tropical speciation. *Nature* 2014;**515**:406–9.
- Smith MR. TreeTools: create, modify and analyse phylogenetic trees. Comprehensive R Archive Network. R package version 1.16.1, 2019. <https://doi.org/10.32614/CRAN.package.TreeTools>
- Stadler T. Sampling-through-time in birth–death trees. *Journal of Theoretical Biology* 2010;**267**:396–404.
- Strauss D, Sadler PM. Classical confidence intervals and Bayesian probability estimates for ends of local taxon ranges. *Mathematical Geology* 1989;**21**:411–27.
- To T-H, Jung M, Lycett S *et al.* Fast dating using least-squares criteria and algorithms. *Systematic Biology* 2016;**65**:82–97.
- Tribble CM, Freyman WA, Landis MJ *et al.* Revgadgets: an R package for visualizing Bayesian phylogenetic analyses from RevBayes. *Methods in Ecology and Evolution* 2022;**13**:314–23.
- Voss RS, Hubbard C, Jansa SA. Phylogenetic relationships of new world porcupines (Rodentia, Erethizontidae): implications for taxonomy, morphological, and biogeography. *American Museum Novitates* 2013;**3769**:1–36.
- Warnock RC, Parham JF, Joyce WG *et al.* Calibration uncertainty in molecular dating analyses: there is no substitute for the prior evaluation of time priors. *Proceedings of the Royal Society B: Biological Sciences* 2015;**282**:20141013.
- Warnock RCM, Wright AM. *Understanding the Tripartite Approach to Bayesian Divergence Time Estimation*. Elements of Paleontology. Cambridge: Cambridge University Press, 2021.
- Weir JT, Price M. Andean uplift promotes lowland speciation through vicariance and dispersal in *Dendrocincla* woodcreepers. *Molecular Ecology* 2011;**20**:4550–63.
- Yang Z. Paml 4: a program package for phylogenetic analysis by maximum likelihood. *Molecular Biology and Evolution* 2007;**24**:1586–91.
- Yang Z. *Molecular Evolution: A Statistical Approach*. Oxford: Oxford University Press, 2014.