

FITTING MODELS OF CONTINUOUS TRAIT EVOLUTION TO INCOMPLETELY SAMPLED COMPARATIVE DATA USING APPROXIMATE BAYESIAN COMPUTATION

Graham J. Slater,^{1,2} Luke J. Harmon,^{3,5} Daniel Wegmann,¹ Paul Joyce,^{4,5} Liam J. Revell,⁶ and Michael E. Alfaro¹

¹Department of Ecology and Evolutionary Biology, University of California, Los Angeles, 610 Charles E. Young Drive East, Los Angeles, California 90095–1606

²E-mail: gslater@ucla.edu

³Department of Biological Sciences, University of Idaho, Moscow, Idaho 83844

⁴Department of Mathematics, University of Idaho, Moscow, Idaho 83844–1103

⁵Initiative in Bioinformatics and Evolutionary Studies (IBEST), University of Idaho, Moscow, Idaho 83844

⁶National Evolutionary Synthesis Center (NESCent), 2024 West Main Street, Suite A200 Erwin Mills Building, Durham, North Carolina 27705

Received October 7, 2010

Accepted September 13, 2011

In recent years, a suite of methods has been developed to fit multiple rate models to phylogenetic comparative data. However, most methods have limited utility at broad phylogenetic scales because they typically require complete sampling of both the tree and the associated phenotypic data. Here, we develop and implement a new, tree-based method called *MECCA* (Modeling Evolution of Continuous Characters using ABC) that uses a hybrid likelihood/approximate Bayesian computation (ABC)-Markov-Chain Monte Carlo approach to simultaneously infer rates of diversification and trait evolution from incompletely sampled phylogenies and trait data. We demonstrate via simulation that *MECCA* has considerable power to choose among single versus multiple evolutionary rate models, and thus can be used to test hypotheses about changes in the rate of trait evolution across an incomplete tree of life. We finally apply *MECCA* to an empirical example of body size evolution in carnivores, and show that there is no evidence for an elevated rate of body size evolution in the pinnipeds relative to terrestrial carnivores. ABC approaches can provide a useful alternative set of tools for future macroevolutionary studies where likelihood-dependent approaches are lacking.

KEY WORDS: Approximate Bayesian computation, Brownian motion, Carnivora, comparative methods, evolutionary rates, incomplete phylogenies.

Evolutionary biologists have long recognized that species richness and trait variation are not evenly distributed among clades. Although constant rate processes provide reasonable null models to explain the distribution of diversity and disparity within and among lineages (e.g., Raup et al. 1973; Garland et al. 1993; Pybus and Harvey 2000; Harmon et al. 2003), shifts in evolutionary rates provide compelling alternative explanations for these pat-

terns. As a result, several phylogenetic comparative methods have been developed in recent years with a focus on testing whether variation in the distribution of continuous phenotypic traits among clades arises under constant or differing rates or patterns of evolution (McPeck 1995; Butler and King 2004; O'Meara et al. 2006; Thomas et al. 2006, 2009; Revell 2008; Revell and Collar 2009).

A wealth of phenotypic data exist in the primary literature following over 200 years of museum collection and taxonomic studies but we currently lack the means to fully leverage this information in broad-scale comparative analyses due to a paucity of complete phylogenetic trees (but see Ackerly and Nyffeler 2004 and Sidlauskas 2007 for methods based on sister lineage comparisons). This is because estimates of rates of trait evolution are explicitly dependent on knowledge of the phylogenetic covariance among species (O'Meara et al. 2006; Ricklefs 2006; Thomas et al. 2006; Revell 2008), information that we lack for incompletely sampled data (Bokma 2010; but see Fitzjohn et al. 2009; Fitzjohn 2010). Approximate Bayesian computation (ABC), a method for fitting models and estimating parameters in a Bayesian framework without likelihoods, provides an alternative that relaxes the necessity of a fully sampled tree and data for inferring rates of evolutionary change on large trees (Bokma 2010). Rather than sampling parameter values based on likelihoods, in ABC, data are simulated under the candidate model using parameter values drawn from their prior distributions (Tavaré et al. 1997). The decision whether to accept or reject the proposed parameter values is based on how close summary statistics of the simulated data come to those of the observed data (for reviews see Joyce and Marjoram 2008; Csillery et al. 2010). As simulation of phylogenetic trees and trait evolution is straightforward, ABC provides a feasible way of estimating rates of continuous trait evolution when the underlying phylogenetic covariance among species is not known (Bokma 2010). A significant advantage of such an approach is that, because data are typically transformed into summary statistics for ABC, we do not require that all species be represented by phenotypic data, only that the sample contains enough randomly sampled species with respect to trait values that the summary statistics adequately describe the distribution of trait values for each unresolved lineage.

In this article we introduce a new, tree-based method called *MECCA* (Modeling Evolution of Continuous Characters using ABC) that integrates recent advances in ABC to estimate posterior distributions for evolutionary rates, diversification rates, and the root state of a continuously distributed phenotypic trait. *MECCA* takes as input an incompletely sampled phylogenetic tree and summary data on the distribution of trait values and species richness within clades. We show here how *MECCA* can be used to estimate the posterior densities of one or more evolutionary rate parameters without knowledge of the underlying, within-clade phylogenetic trees. We then use a novel model selection approach to select among single and multiple Brownian rate models of trait evolution. Finally, we show using simulation that our method has considerable power to detect differences in evolutionary rates from incomplete comparative data. We then apply *MECCA* to an empirical example and show that a model of a constant evolutionary rate of body size evolution among carnivores is preferred over

a model with an elevated rate in aquatic compared to terrestrial carnivores.

Methods

MODEL DEFINITION AND CHOICE OF SUMMARY STATISTICS

Consider a phylogenetic tree, τ , with L terminal lineages, each of which represents some higher level taxon such as a family or order. Each single terminal lineage $i \in \{1, \dots, L\}$ contains r_i species, for which the phylogenetic relationships are not necessarily resolved. Each terminal lineage is also associated with an incompletely sampled dataset D for a phenotypic trait, such as body size. The goal is to estimate the rate of body size evolution in the entire clade based on the set of incompletely sampled comparative data (τ, D). Because we lack an analytical solution for the likelihood of a Brownian diffusion process on an unobserved tree (Bokma 2010), we can instead use a Markov-Chain Monte Carlo (MCMC) without likelihoods approach (hereafter ABC-MCMC, Marjoram et al. 2003) to sample rate parameters from their posterior distribution. For each generation of the ABC-MCMC algorithm, we will replace each terminal lineage in our tree with a simulated clade containing its complete complement of species. We will then simulate trait evolution over the now "completely sampled" tree using candidate Brownian motion parameters. As we discuss in detail below, the decision whether to accept or reject proposed parameters will be made based on comparison of the simulated trait data to our observed data.

We will initially assume that species richness $R = (r_1, \dots, r_L)$ of our L extant clades result from a homogenous stochastic birth-death process, whereas species' trait values evolved under a homogeneous Brownian diffusion process. These are by far the most common and widely used models for comparative data (but see Hansen 1997; Butler and King 2004; Rabosky and Lovette 2008a,b; Harmon et al. 2010 for other models that could potentially be implemented in an ABC framework). The use of these two processes in our model yields four model parameters that must be estimated. Under a birth-death process, we model speciation and extinction rates, λ and μ , respectively (Nee et al. 1994). For a Brownian diffusion process, we require the root state, a , and the Brownian diffusion rate, σ^2 (Felsenstein 1985; Hansen and Martins 1996; O'Meara et al. 2006).

We do not need to use approximate methods to sample birth-death parameters. Given knowledge of a backbone phylogeny and the number of extant species within each terminal lineage, we can sample λ and μ directly from their posterior distributions using MCMC, where candidate diversification parameters λ and μ are accepted or rejected based on $L(\lambda, \mu \mid \tau, R)$, the likelihood of observing both the backbone phylogeny τ and of observing species richness values for each terminal clade given those values (Rabosky et al. 2007). The estimation of extinction

rates from molecular phylogenies has been criticized elsewhere (e.g., Rabosky 2010) but incorporating extinction is important for our purposes as it impacts the relative distribution of branching events in a phylogenetic tree and thus the expected phenotypic disparity among the tips for a given rate of trait evolution (O'Meara et al. 2006). By simulating trees for the unsampled terminal clades conditional on age, species richness, and the sampled diversification parameters (Stadler in press), we obtain a distribution of trees with branching times that are representative of the posterior distribution $\pi(\lambda, \mu \mid \tau, R)$. These simulated trees can then be attached to the backbone phylogeny to create a distribution of “completely sampled” phylogenies.

Although diversification parameters can be sampled directly from their true posterior distributions using likelihoods, this is not possible for Brownian motion parameters with incompletely sampled data (Sidlauskas 2007; Bokma 2010). Using the ABC–MCMC framework, we may instead derive an approximation of the true joint posterior distribution $\pi(a, \sigma^2 \mid \tau, R, D)$ via simulation of trait evolution over our reconstructed, “completely sampled” tree from the previous step. At each generation, we sample values of a and σ^2 from their respective prior distributions and simulate data over the tree. Because likelihoods are unavailable, the decision whether to accept proposed parameters in our ABC–MCMC will instead be made by computing the Euclidean distance, δ , between the simulated data and observed data (Marjoram et al. 2003). If $\delta \leq \delta_{\text{crit}}$, where δ_{crit} is a small, user-defined tolerance specifying how far we are willing to allow the simulated data to be from the observed data, then we will accept the proposed parameters. If $\delta > \delta_{\text{crit}}$, then we reject.

CHOICE OF SUMMARY STATISTICS

Because our trait data and tree are incomplete, we use a set of summary statistics, S , to describe the phenotypic variation present in D for the sampled taxa. In a Brownian diffusion process, the expected mean and variance of a trait across taxa corresponds to the root state a and the product of the path length and diffusion rate σ^2 , respectively. We therefore use vectors of means, $M = (m_1, m_2, \dots, m_L)$, and variances, $V = (v_1, v_2, \dots, v_L)$, of our trait of interest for each terminal lineage as summary statistics. Although not all species need to be represented by phenotypic data, we do require that a sufficient proportion of species have been sampled so that the summary statistics for each clade are adequate descriptors of the total sample. We also assume that this sample is random with respect to trait values. For example, although 30 samples might be sufficient to describe the variation within a clade of 100 species, we do not wish to sample the 30 smallest or 30 largest taxa.

For high dimensional datasets, δ is likely to be very large for most simulations. This will result in low acceptance rates and inefficient mixing of the MCMC chain unless a large

tolerance is used at the expense of precision (Joyce and Marjoram 2008; Leuenberger and Wegmann 2010). For even moderately sized phylogenetic trees, the large number of summary statistics described above ($2L$ where L is the number of unresolved tip clades in the tree) is therefore problematic. To overcome this, we use a partial least squares (PLS) regression transformation of the summary statistics to generate a new, lower dimensional set of summaries, S_{pls} , prior to computing the distance function (Wegmann et al. 2009). PLS is similar to multivariate methods such as principal components analysis in that orthogonal combinations of variables called components are identified that explain successive amounts of covariation in the original data. PLS scores for each component are then computed as the dot product of the summary statistics and their associated component loadings. PLS is particularly advantageous for our purposes because combinations of variables are chosen to maximize the variation explained in a set of response variables. Summary statistics that are good predictors of the Brownian rate, such as trait variances, will therefore receive large loadings on components associated with the rate of trait evolution, whereas those that are poor predictors, such as clade means, will receive small loadings. An additional advantage is that if the trait variance of one or more clades is less informative for the Brownian rate, as might occur if there are only a few species in an old lineage, then those clades will also receive small loadings for PLS components predictive of the rate parameter and their trait variances will be down-weighted when deciding whether to accept or reject in ABC–MCMC.

We determine the number of PLS components to use based on 10,000 calibration simulations generated with parameters drawn from the prior distributions of the model parameters (see below). We first linearize our summary statistics from the calibration simulations using a Box–Cox transformation, resulting in 10,000 sets of $2L$ summaries S_{BoxCox} (Wegmann et al. 2009). PLS analysis is then conducted using these standardized summaries as predictor variables and the set of rates and root states used in those simulations as response variables. Leveling off of root mean square error plots of the parameters predicted by the regression can be used to determine the minimum number of informative PLS components to retain (Wegmann et al. 2009). For simulations under a homogeneous Brownian motion process, we found that the cumulative percentage of total variance in the parameters explained always leveled off at two PLS components. We finally produce a reduced set of summaries S_{pls} by multiplying the set S_{BoxCox} by their associated PLS loadings and summing over each component. For the case of evolution under a single Brownian rate, this reduces the dimensionality of the summary statistics from $2L$ to 2.

CALIBRATION

In regular MCMC samplers, it is possible to initiate the chain in a region of low likelihood because improvements in likelihood,

however small, will move the chain toward its target distribution. In ABC–MCMC, this is much more difficult because proposals are accepted on an absolute rather than relative basis. Therefore, if the chain is initiated in a region of low likelihood and the proposal width and/or distance for acceptance are small, ABC–MCMC samplers can easily become stuck.

To overcome this problem, we use an initial, random sample of simulations, referred to as calibration, to determine several important tuning parameters for the ABC–MCMC (Wegmann et al. 2009). In the calibration simulations, we generated 10,000 complete trees with branching times that are representative of the posterior distribution $\pi(\lambda, \mu | \tau, R)$ as described above. Over each of those complete trees, we simulated trait evolution with a and σ^2 sampled randomly from their prior distributions. Finally, we computed the associated summary statistics for each realization.

Based on these datasets, we first define δ_{crit} such that for 5% of the simulations $\delta < \delta_{\text{crit}}$. We next set proposal ranges for each parameter as uniform widths corresponding to twice their respective standard deviation among those 5% simulations closest to the observed summary statistics. Finally, to avoid the need to discard samples as burn-in, we initiate our ABC–MCMC chains with the parameter values associated with the simulation that resulted in the smallest distance δ . We start the full-MCMC chain used to sample birth–death parameters at the maximum likelihood estimates for λ and μ , computed according to Rabosky et al. (2007). For this chain, we used an unbounded, uniform prior on both λ and μ , and a proposal width of 0.1.

ESTIMATION OF POSTERIOR DISTRIBUTIONS

To account for the necessary loss of precision resulting from ABC approaches, we applied a postsampling regression adjustment to the retained posterior sample from our ABC–MCMC (Beaumont et al. 2002; Csillery et al. 2010). We estimated posterior distributions for a and σ^2 using the ABC–GLM postsampling regression approach proposed by Leuenberger and Wegmann (2010) and implemented in the software package ABCtoolbox (Wegmann et al. 2010). This approach assumes that the likelihood function can be locally approximated by a general linear model (GLM) around the observed values and has been shown to improve posterior estimates substantially over naïve estimates from the ABC samples (Leuenberger and Wegmann 2010).

FINAL MECCA ALGORITHM

We now summarize the combined likelihood/ABC–MCMC algorithm that MECCA uses to estimate $\pi(a, \sigma^2 | \tau, R, D)$, the joint posterior distribution of the Brownian motion parameters a, σ^2 from incompletely sampled data τ, R , and D .

- M1. Perform calibration simulations.
- M2. Linearize the C sets of summaries using a Box–Cox transformation, define PLS components, and transform the

summary statistics. Transform the observed summaries to produce S_{pls} .

- M3. Determine proposal ranges, δ_{crit} and the starting position for an ABC–MCMC chain of length N .
- M4. Set λ and μ to their maximum likelihood estimates. Set $j = 0$.
- M5. If now at λ and μ , propose a move to λ' and μ' based on the transition kernel $q(\lambda, \mu \rightarrow \lambda', \mu')$ and accept with probability $\min(1, \frac{\pi(\lambda', \mu') \cdot L(\lambda', \mu' | \tau, R) \cdot q(\lambda, \mu \rightarrow \lambda', \mu')}{\pi(\lambda, \mu) \cdot L(\lambda, \mu | \tau, R) \cdot q(\lambda', \mu' \rightarrow \lambda, \mu)})$. Otherwise, stay at λ and μ .
- M6. Simulate replacement terminal clades based on τ, λ, μ , and R .
- M7. If now at a and σ^2 , propose parameters a', σ'^2 based on the transition kernel $q(a, \sigma^2 \rightarrow a', \sigma'^2)$.
- M8. Simulate phenotypic data D' with root state a' and rate σ'^2 , compute the summary statistics S' and transform them into S'_{pls} based on the Box–Cox parameters and PLS components computed in M2. Compute the Euclidean distance $\delta = \|S_{pls} - S'_{pls}\|$.
- M9. If $\delta \leq \delta_{\text{crit}}$, accept a', σ'^2 with probability $\min(1, \frac{\pi(a', \sigma'^2) \cdot q(a', \sigma'^2 \rightarrow a, \sigma^2)}{\pi(a, \sigma^2) \cdot q(a, \sigma^2 \rightarrow a', \sigma'^2)})$; otherwise, remain at a, σ^2 , regenerate phenotypic data D' with a, σ^2 , compute the summary statistics S' and transform them into S'_{pls} .
- M10. Store $\lambda, \mu, a, \sigma^2$, and S'_{pls} .
- M11. Increment j . If $j < N$ return to M5.
- M12. Retain the t stored simulations with smallest δ and perform ABC–GLM regression on those samples to estimate the posterior distribution $\pi(a, \sigma^2 | \tau, R, \delta \leq \delta_t)$, where δ_t is the distances associated with the t th simulation closest to S_{pls} .

We use symmetric transition kernels here and so the transition ratios in steps M5 and M9 are equal to 1. Only the prior ratios therefore require computing when determining the acceptance probability.

TWO-RATE MODEL AND HYPOTHESIS TESTING

The approach that we describe above assumes that the evolutionary rate of phenotypic change is homogeneous across lineages. We may, however, also be interested in determining whether the phenotypic diversity observed for some clades evolved under a different rate of trait evolution (McPeck 1995; O'Meara et al. 2006; Thomas et al. 2006, 2009). It is straightforward to modify MECCA to do this with incompletely sampled data. We need only simulate trait evolution in our clade of interest using a different Brownian diffusion rate σ_2^2 with its own prior distribution. With this additional response variable, three PLS components are generally needed to capture the information present in the summaries.

We would also like to assess whether the one- or two-rate model better fits the data. Model selection in ABC is often based on Bayes factors, using acceptance ratios as approximations of the marginal likelihoods for competing models (Beaumont

et al. 2002; Bokma 2010). Leuenberger and Wegmann (2010) suggested instead computing the marginal likelihoods of the observed summary statistics directly from the regression done in ABC-GLM and using these for model selection. This approach, however, requires using the same summary statistics for both models, something that we are unable to do when using PLS-transformed summaries.

Here, we use an alternative, Bayesian approach to assess support for a two-rate model. Our approach makes explicit use of the fact that the one-rate model is nested within the two-rate model and that hypotheses regarding rate differences are typically directional. We compute the posterior support p under the two-rate model for the two evolutionary rates σ_1^2 and σ_2^2 being different and reject the single rate model if this probability is large. To be specific, for a case where we expect σ_2^2 to be larger than σ_1^2 , we would estimate $p = P(\sigma_1^2 < \sigma_2^2 \mid \tau, R, D)$ using samples from the joint posterior distribution of σ_1^2 and σ_2^2 . We generated 10^7 random samples using an MCMC implemented in ABCtoolbox (Wegmann et al. 2010) and thinned it down to 10^5 to reduce correlation among samples.

EMPIRICAL EXAMPLE AND SIMULATION TESTS

The mammalian order Carnivora comprises 286 species distributed among 16 families, and spans five orders of magnitude in body mass (Gittleman and Purvis 1998; Nowak 1999). Most carnivorans are terrestrial but all members of one monophyletic clade, the Pinnipedia (seals, sea lions, and walruses), are semi-aquatic. Unsurprisingly for aquatic mammals that occur in all oceanic regions (Deméré et al. 2003), some species attain extremely large sizes, including the largest extant carnivoran, the Southern elephant seal *Mirounga leonina*, with a body mass of up to 3700 kg (Nowak 1999). We therefore ask whether the dramatic range of body sizes observed in the pinnipeds is due to a clade-specific increase in the rate of body size evolution relative to that of terrestrial carnivorans.

We used the time-calibrated, family level phylogeny of Carnivora (Fig. 1) from Eizirik et al. (2010). Although some families were represented by multiple species (range: 1–8), we pruned the tree down to one representative per family ($L = 16$) and assigned each a species richness value based on the number of recognized taxa in Wozencraft (2005). We also assigned each family a mean and variance for natural log-transformed (ln) body mass data (kg) from the PanTHERIA database (Jones et al. 2009). Not all carnivorans were represented by body mass data in this database (Fig. 1B). However, because data are transformed into summary statistics, we only require that the summary statistics for the available samples adequately describe the distributions for all species, which we assume here they do.

We placed a normal prior on $\ln(\sigma^2)$ (mean = -2.53 , SD = 2). The mean was determined by computing the mean square of

independent contrasts from the family level tree and family mean body masses (Revell et al. 2007). We bounded the rate prior at $\{-4.961845, 4.247066\}$, corresponding to the natural log of half the minimum and twice the maximum rates of body size evolution reported across a range of species by Harmon et al. (2010). The prior on the root state was set as a uniform distribution $U = [-6.9, 3.21]$, corresponding to a range from 1 g to 25 kg. These values span the masses of the smallest living mammals up to over an order of magnitude larger than estimated masses for stem carnivorans (Finarelli and Flynn 2006).

We first investigated the performance and power of *MECCA* for the one- and two-rate carnivore models using 1000 simulated datasets. Trait data were simulated over the carnivore phylogeny, with unsampled terminal clades replaced by trees containing the correct number of species and of the correct stem age (Stadler in press), simulated under maximum likelihood estimates of λ and μ . For each simulation, we drew root state and evolutionary rate parameter values (a , σ_1^2 , and σ_2^2) from their respective prior distributions. We then simulated two datasets per set of parameters—one under a one-rate model and one dataset under a two-rate model.

We checked the number of MCMC generations required to achieve convergence on the target distributions for the Brownian motion parameters for both models. To determine if and when convergence was achieved, we used the R -statistic (Gelman and Rubin 1992) computed for 100 simulated datasets. The R -statistic compares the among- and within-chain variances of each parameter for two or more independent MCMC chains. As chains converge on the same target distributions, R should approach 1. We ran two, independent chains of 150,000 steps for each simulated dataset and computed R at regular intervals from the output. For both rate models, we observed R values below 1.01 in most replicates after 100,000 generations (Fig. 2), indicating acceptable convergence had been achieved. We thus used chains of 100,000 iterations for all other analyses.

We next checked the coverage of the posterior distributions by computing posterior quantiles of the true parameter values for the simulated datasets. If *MECCA* produces unbiased parameter estimates then we should find that, on average, we recover the correct values. For example, the true value of each parameter should be located in the 50% and 95% highest posterior density regions with probabilities 0.5 and 0.95, respectively. This is equivalent to their posterior quantiles (the quantile of the posterior distribution within which the true value falls) being distributed uniformly on $U[0,1]$ (Cook et al. 2006). We computed the posterior quantiles for all 1000 simulated datasets and tested for uniformity using a Kolmogorov–Smirnov test.

We also investigated Type I error rates and power ($1 - \text{Type II error}$) to reject the one-rate carnivore model in favor of the two-rate model. We first plotted posterior probabilities of the two-rate

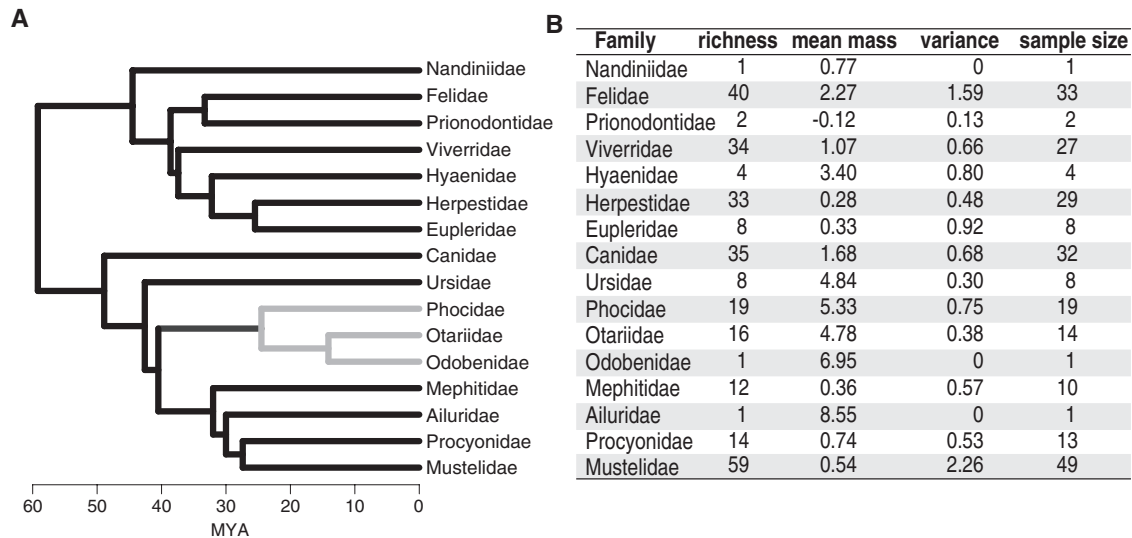


Figure 1. (A) Time-calibrated phylogeny of Carnivora used to estimate rates of trait evolution using *MECCA*. **(B)** Summary data on species richness and means and variances for body size per family. The final column gives the number of species for which body size data were available.

model, $p = P(\sigma_1^2 < \sigma_2^2 \mid \tau, R, D)$, against the proportion of cases for which the two-rate model is true. If *MECCA* produces unbiased estimates of p , we would expect a 1:1 relationship between p and the proportion of cases for which $\sigma_1^2 < \sigma_2^2$. For example, considering all cases analyzed that receive $P = 0.1$, we would expect to find that $\sigma_1^2 < \sigma_2^2$ is true in approximately 10% of cases. If posterior probabilities are low relative to the proportion of true two-rate cases, then *MECCA* is overly conservative and power to detect rate shifts is low. The magnitude of differences in rates is also likely to affect our ability to detect rate shifts. Ideally, posterior probabilities for the two-rate model would be low (~ 0) when the two rates are identical and rise rapidly to values > 0.9 for differences in rate that are greater than zero. Realistically, smaller differences in rate will be more difficult to detect than large differences (e.g., Collar et al. 2005). We investigated how rate

differences affect our model selection abilities by binning samples according to the magnitude of the difference between the two rates and plotting rate differences against posterior probabilities of the two-rate model. We then computed the rate difference required to achieve a “significant” result at the $\alpha = 0.05$ level. We finally checked for the proportion of cases in which we falsely reject the one-rate model, equivalent to Type I error.

After conducting the simulation tests, we ran two separate *MECCA* analyses using the carnivore dataset. For the first, we assumed a single rate of body size evolution for all Carnivora. In the second, we allowed crown pinnipeds (three terminal lineages and one internal branch, Fig. 1A) to evolve under a different rate of body size evolution. We used the same prior distribution for both trait evolutionary rates for the two-rate model. We ran 10,000 calibration steps for both models, from which we computed PLS

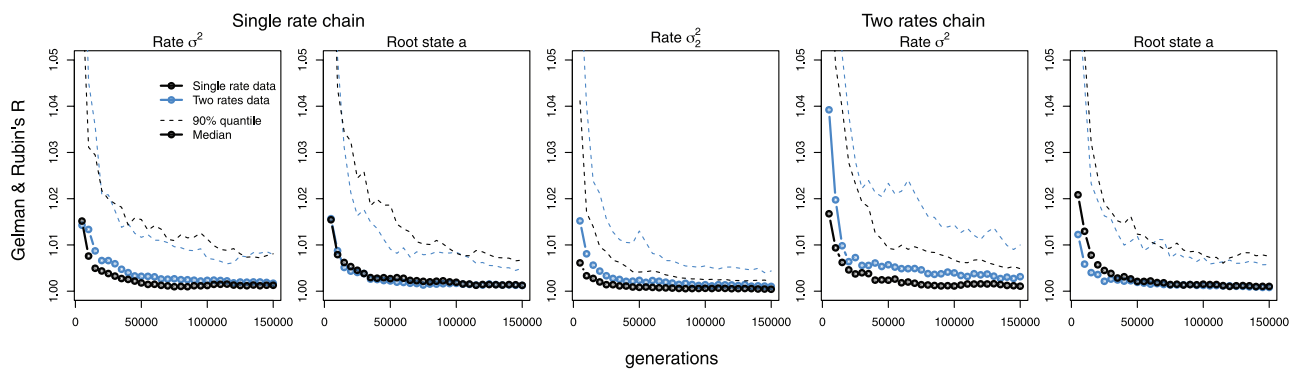


Figure 2. Convergence diagnostics for the single (black) and two-rate (blue) *MECCA* based on two identical chains for the 1000 simulated datasets. Circles indicate median R values at a given generation of the MCMC chain, while dashed lines give the 90% quantiles. R rapidly approaches 1 for all parameters, indicating adequate convergence is achieved within 100,000 generations.

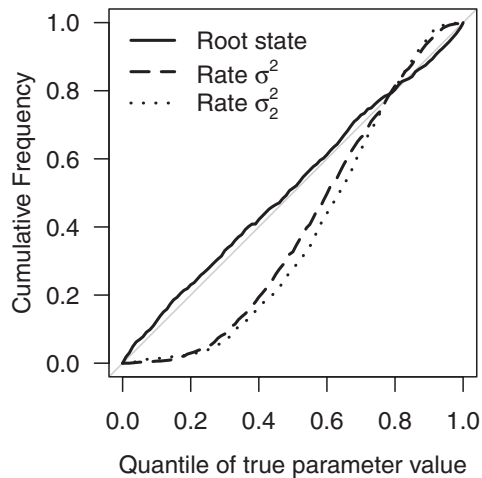


Figure 3. Cumulative density function plot of posterior quantiles for the known parameter values used in 1000 simulated datasets from the carnivore tree. Posterior distributions were estimated from the output of MECCA using ABC–GLM. The distributions are expected to be uniform (light gray line) if posteriors are unbiased. Root state estimates are unbiased ($P = 0.13$). For the rate parameters, quantiles are biased toward larger values, indicating that posterior distributions for these parameters are too wide.

components for transformation of summary statistics and tuning parameters for the MCMC. We then ran 100,000 generations of ABC–MCMC under each model and retained the 5000 samples closest to the observed data for estimation of parameter poste-

rior distributions using ABC–GLM. To perform model selection, we computed the posterior probability that $\sigma_{\text{pinnipeds}}^2$ was greater than $\sigma_{\text{terrestrial carnivores}}^2$ based on 100,000 samples from the joint posterior distribution of evolutionary rates.

Results

SIMULATION TESTS

We ran MECCA on a total of 1000 simulated datasets to assess the performance of our approach. Although estimates of the root state appear well calibrated, rate estimates appear to be conservative, with flatter posterior distributions than expected (Fig. 3). Similarly, we found the posterior support $p = P(\sigma_1^2 < \sigma_2^2 \mid \tau, R, D)$, to be too conservative. For instance, the true proportion of cases for which $\sigma_1^2 < \sigma_2^2$ is up to 20% higher than estimated (Fig. 4A). We used the simulations to improve on model selection performance by treating the estimator p as a summary statistic and attempting to compute the adjusted posterior probability $\hat{p} = P(\sigma_1^2 < \sigma_2^2 \mid p)$. We found that \hat{p} can be adequately approximated as a logistic function of p (Fig. 4A).

For our simulations, we found substantial power to detect differences in the rate at which morphological variation evolves in different clades. For instance, median \hat{p} values are found to be significant at the $\alpha = 0.05$ level if the difference in $\ln(\text{rates})$ was 1.75 or more (Fig. 4B). We finally used our simulations to assess Type I errors when rejecting the one-rate model in favor




Figure 4. Performance and power of MECCA assessed from 1000 simulated datasets. (A) Posterior probabilities P of a two-rate model plotted against the proportion of simulations for which the two-rate model is the true model (simulations are binned by posterior probabilities). The dashed gray line indicates the expected linear relationship. Posterior probabilities greater than 0.5 are underestimated whereas those below 0.5 are slightly overestimated. The blue curve shows the fit of a logistic regression to the same data. (B) The relationship between absolute difference in rates (on the natural log scale) and P for the two-rate model (black solid line = median, 90% quantiles = dashed line). For large differences in rate, we obtain strong posterior support for the two-rate model. However, support rapidly declines for smaller differences in rate. The blue lines show adjusted posterior probabilities, \hat{p} , for the same data after transformation by the logistic regression estimated in A. Use of \hat{p} increases confidence in the two-rate model for smaller differences in rate. (C) Type I error. Note that both axes use a log-scale centered at 0.5. Cumulative distributions of the posterior probabilities that $\sigma_2^2 < \sigma_1^2$ for 1000 simulated datasets generated under the one-rate model. These distributions indicate the Type I error of falsely rejecting the null model (one-rate model). For well-calibrated methods, the Type I error are expected to be uniformly distributed and hence to show a cumulative distribution on the diagonal (gray dashed line). Model choice on the raw posterior probabilities appears to be much too conservative (black line). The transformed posterior probabilities, however, appear to be reasonably calibrated (blue line).

of the two-rate model. Model choice based on p was found to be too conservative as the proportion of cases in which we falsely reject the single-rate model (0.7%) was much below the typical 5% threshold (Fig. 4C). However, model choice based on \hat{p} was much improved, with the one-rate model being falsely rejected in 3.1% of cases (Fig. 4C). Overall these results demonstrate the ability of *MECCA* to reliably discern a two-rate process for modest differences in rates compared to likelihood approaches using completely sampled data (e.g., Collar et al. 2005).

RATES OF BODY SIZE EVOLUTION IN CARNIVORA

We recovered similar estimates for diversification parameters under both single- and two-Brownian rate models as expected. We therefore concatenated posterior samples for diversification rate parameters for summary of results. For ease of interpretation, we have also converted speciation and extinction rates to net diversification ($r = \lambda - \mu$) and turnover ($\epsilon = \mu/\lambda$) rates. We recovered a mode net diversification rate of 0.07, combined with a turnover rate of 0.79. Highest Posterior Density (HPD) intervals were broad for both parameters (r : 95% HPD 0.03–0.10; ϵ : 95% HPD 0.19–0.97). The use of a one- or two-rate model had no effect on root state estimates (Table 1). Root state estimates appear large in comparison to estimates based on the fossil record (e.g., Finarelli and Flynn 2006). Ancestral state reconstructions have been shown to perform poorly when the clade in question experienced a directional trend in trait evolution and fossil data are not included (Finarelli and Flynn 2006; Albert et al. 2009). This result is therefore likely a function of the data, rather than poor performance of our method. Rates estimated from the single- and two-rate models were similar, although the 95% credible range on the pinniped rate was much wider (Table 1). A contour plot of the joint marginal posterior distribution of rates for the two-rate model further suggests that the two rates do not deviate substantially from expectations of a single-rate model (Figs. 5A,B) and

Table 1. Mode and 95% HPD intervals of the marginal posterior distributions for the Brownian rate (σ^2) and root state (a), estimated for the carnivore dataset using *MECCA*. Posteriors were computed using *ABCGLM* on the best 5000 simulations. Note that all estimation for both parameters was done on the natural log scale, but parameters have been transformed here for ease of reading.

Parameter	One-rate model	Two-rate model	
		Background	Pinnipeds
σ^2	0.055 (0.016–0.183)	0.054 (0.011–0.184)	0.053 (0.008–0.302)
a	8.98 (1.90–24.76)	8.09 (1.63–24.75)	

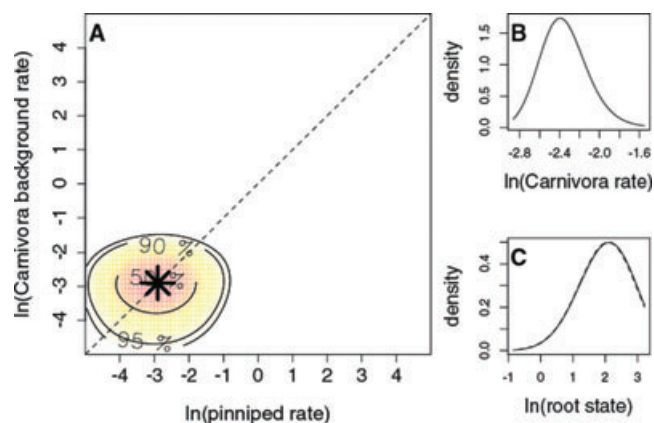


Figure 5. (A) Contour plot of the joint posterior distribution for pinniped and terrestrial carnivore body size evolutionary rates. Hotter colors indicate lower HPD regions. Note that the joint posterior distribution is projected into the entire prior space here. The 50, 90, and 95% highest posterior density regions are marked. The asterisk indicates the mode rate estimated by *MECCA* for the single rate model. (B) the posterior distribution for the evolutionary rate under a single rate model and (C) the posterior distributions for the root state under one (solid line) and two-rate (dashed line) models.

the posterior probability of a model with pinnipeds evolving under a faster rate than terrestrial carnivores is only slightly greater than 0.5 ($P = 0.55$). Adjustment based on the logistic regression fit to the simulated datasets increased the posterior probability of the two-rate model to 0.72. Although three-fourths of the posterior weight support a model where pinnipeds evolve under a faster rate than terrestrial carnivores, a model of a constant rate of body size evolution across all carnivores cannot be rejected.

Discussion

By relaxing the requirement of a completely sampled tree and data, approximate Bayesian methods for estimating evolutionary rates, such as the method presented here, have the potential to dramatically expand the breadth of question that can be asked about the pace of phenotypic evolution on the tree of life. We end here by comparing the performance of *MECCA* to existing methods for inferring rates of trait evolution, as well as considering some future extensions and limitations of ABC approaches in comparative methods.

Our simulations demonstrate that we are able to obtain a great deal of the information contained in raw trait data regarding the Brownian rate and root state by using PLS transformations of means and variances for unsampled clades as summary statistics in ABC. PLS loadings are determined using summary statistics from the calibration simulations and their associated parameter values. By using data simulated on the backbone tree in this way, PLS effectively weeds out summaries from less-informative clades while more heavily weighting summaries from informative clades. This

is essential in the context of an unsampled phylogeny as treating all summaries equally would result in an old, species-poor clade contributing the same amount of weight to the acceptance decision as a young, species-rich clade. In this sense, using PLS-transformed summaries in ABC-MCMC is analogous to incorporating the phylogenetic covariance matrix in likelihood methods. One could imagine replacing our ABC approach with a likelihood equivalent that integrates over uncertainty in the topology of unsampled clades. For example, Fitzjohn et al. (2009) and Fitzjohn (2010) demonstrated that when estimating rates of phenotypic trait evolution, it is possible to integrate over uncertainty in the phylogenetic position of taxa that have been sampled in the phenotypic dataset but are not in the tree. Although these approaches perform well with relatively sparsely sampled clades, they are only computationally practical for clades containing up to several hundred species (Fitzjohn et al. 2009). Our ABC approach performs well with an extremely sparsely sampled phylogeny and can handle trees containing extremely large numbers of species. An analytical solution yielding the likelihood of an observed mean and variance for a given evolutionary rate over an unsampled clade remains the ideal (Bokma 2010). However, our results indicate that approximate approaches to estimating rates of trait evolution on large, sparsely sampled phylogenies provide a compelling alternative.

There are several appealing aspects to the posterior probabilistic approach that we took here. Our approach takes into account the entire joint posterior distribution and therefore accounts for uncertainty in parameter estimation, something that may have a significant impact in evolutionary models based on Brownian motion (e.g., Polly 2001). Perhaps more significantly though, using a posterior probabilistic approach allowed us to take advantage of the reduced data dimensionality provided by PLS-transformed summaries (Wegmann et al. 2009) while also making full use of methods that apply posterior-adjustments to the output of an ABC-MCMC chain (Leuenberger and Wegmann 2010). Moreover, because our model selection does not contrast independently approximated likelihood functions, general criticisms leveled at model selection via Bayes factors in ABC (Robert et al. unpubl. ms.) do not apply to our case. Computation of posterior probabilities based on samples from the joint posterior distribution of the two-rate parameters is straightforward in the case of a two-rate model. For more complicated, multiparameter Brownian models, visualization of the joint posterior distribution (Fig. 5A) will be more challenging but the approach to sampling from the joint space remains the same.

Simulations indicate that we retain considerable power to detect differences in the rate at which morphological disparity is accumulated from incompletely sampled comparative data compared to complete data. Using uncorrected posterior probabilities, P , we found that we would be able to reject a single Brownian rate

model with a high degree of confidence (5% significance level) if the absolute difference between the two rates in natural log space was larger than 2.7 (Fig. 3B), equivalent to the faster rate being around 15 times greater than the background rate. This is comparable power to that found by Collar et al. (2005) and Eastman et al. (in press) for approaches that use likelihood to estimate rates from completely sampled phylogenies. However, model selection in *MECCA* based on \hat{p} was greatly improved relative to these approaches. Here, rate differences of only 1.75, equivalent to a rate ratio of 5.75, were required to accept the two-rate model at the 5% significance level (Fig. 3B). We note that the exact power will be specific to the phylogenetic tree studied and is expected to be greater for larger trees. Nonetheless, these findings suggest that the power to detect rate shifts in our ABC approach based on incomplete trees and summary statistics of trait data for terminal clades is comparable to likelihood methods using complete datasets, even where rate differences are relatively small.

As with any Bayesian approach, the choice of prior distributions can have a large impact on the results obtained. We used a truncated normal prior range for the evolutionary rate, and a bounded uniform distribution for the root state, although other distributions could be envisioned. Our simulation tests, which were based on parameters drawn from the same prior distributions, suggest that the use of a wide, normal prior with upper and lower bounds does not adversely affect parameter estimates when the true values lie far from the mean (Fig. 2A). Others (e.g., Schluter et al. 1997; O'Meara et al. 2006; Thomas et al. 2006) have assumed uninformative, unbounded uniform priors on model parameters when fitting Brownian motion models to comparative data. Unbounded priors can be problematic in ABC because here proposals are accepted on absolute, rather than relative terms. That is, a proposal is accepted only if the simulated data fall within some specified distance of the observed data rather than if the proposal increases the likelihood of the observed data, relative to the current state. Wide priors can therefore occasionally allow ABC-MCMC samplers to become stuck in regions of low likelihood or, if the starting values are not appropriate, failure of the chain to even initiate. We overcame this problem in part by using a calibration step prior to initiating ABC-MCMC. However, because the tolerance for the ABC-MCMC, as well as several tuning parameters, are determined based on the calibration simulations, a wider prior will require a longer calibration to ensure suitable values are chosen. A bounded prior further alleviates some of these issues by limiting sampling to a reasonable range of values. Regardless of the prior distribution used for ABC, the range of values should be appropriate given the data in hand, as in any Bayesian approach. Limiting the width and/or shape of the prior on the root state based on physiological, fossil or some other form of prior information may be wholly appropriate in some contexts. In the absence of prior information, using the range of

tip values however is unlikely to be sufficient (e.g., Polly 2001). Posterior distributions piled up against the upper or lower limit of the prior provide a useful indication that the bounds of the specified prior are inappropriate given the model and data. The units of the data, for example, grams versus kilograms, or millimeters versus meters, and the age of the clade should also be considered when determining the appropriate width of the root state and rate priors.

In this implementation, we assumed time invariant rates for both diversification and trait evolution. However, the great flexibility of ABC approaches is that, provided that it is possible to simulate data, parameter values can be sampled under any models for which likelihood expressions are or are not available (e.g., Rabosky 2010; Bokma 2010). It is straightforward, provided informative summary statistics exist, to accommodate a more diverse range of models, such as early burst (Harmon et al. 2010), Ornstein-Uhlenbeck (Hansen 1997; Butler and King 2004), evolutionary trend (Hunt 2006) or single-shift along an internal edge (e.g., Thomas et al. 2006, 2009; Revell et al. in press) models of trait evolution. One could also use ABC to fit models where diversification rates depended on character states (e.g., Maddison et al. 2007). A particularly tractable extension would be to simultaneously infer rates of trait evolution for unsampled clades in conjunction with identification of diversification rate shifts using the MEDUSA approach (Rabosky et al. 2007; Alfaro et al. 2009). Differences in diversification parameters among lineages have the potential to greatly influence the inference of trait evolutionary rates in the ABC context because higher extinction rates will tend to result in branching events that are biased toward the recent, and thus lower expected trait disparity (Pie and Weitz 2005; O'Meara et al. 2006). It should be noted that such an approach would not rescue *MECCA* from producing a false positive result if the putative shift in the rate of morphological trait evolution were associated with an unsampled clade that suffered high rates of extinction relative to other lineages. This is because a clade that diversified under high extinction rates will be “tippy” and have lower phenotypic variance among species than a clade with the same age and species diversity that diversified under the same net rate of diversification but with low extinction. If we lack information regarding the true tree shape and fit a model that assumes constant extinction, we would bias ourselves toward finding that the tippy clade showed a decrease in the rate of phenotypic evolution over time. Although a general problem associated with incompletely sampled data, including information on crown clade ages, based on molecular or paleontological data, would be a useful way of dealing with this issue.

In conclusion, ABC methods have great potential in phylogenetic comparative biology, particularly where models of trait evolution can be simulated but analytic expressions cannot be solved. Although a complete phylogeny and phenotypic dataset

remain the ideal for comparative analysis, we have shown here how an ABC approach can be used to infer rates of evolution from large, incompletely sampled data. By applying *MECCA* to a backbone phylogeny of carnivores, we have been able to show that there is no evidence for a faster rate of body size evolution in the aquatic pinnipeds compared to terrestrial carnivore species, despite the considerable power of *MECCA* to detect differences in evolutionary rates in this context, as we found through simulations. The larger body masses found in the pinnipeds might thus just be the result of the stochastic nature of Brownian processes. Alternative explanations include a distinct evolutionary optimum among pinnipeds (e.g., Butler and King 2004), or rapid evolution of body size in the stem lineage of pinnipeds, followed by relatively stable evolution since that time (Simpson 1953; Thomas et al. 2006, 2009; Revell et al. in press), rather than a clade-wise elevated rate. The ABC approach presented here can be readily extended to contrast such different evolutionary models, even at phylogenetic levels for which we are unlikely to have complete trees in the near future.

ACKNOWLEDGMENTS

We thank F. Bokma, J. Brown, J. Eastman, M. Pennell, S. Price, M. Suchard, A. Mooers, Associate Editor G. Hunt, and an anonymous reviewer for discussion of ideas and comments on the manuscript. This work is funded by NSF grant DEB 0918748 to MEA and NSF DEB 0919499 to L.J.H. L.J.H. was supported the National Evolutionary Synthesis Center (NSF EF-0423641). DW was funded via a Searle Scholar Program award to John Novembre.

LITERATURE CITED

- Ackerly, D. D., and R. Nyffeler. 2004. Evolutionary diversification of continuous traits: phylogenetic tests and application to seed size in the California flora. *Evol. Ecol.* 18:249–272.
- Albert, J. S., D. M. Johnson, and J. H. Knouft. 2009. Fossils provide better estimates of ancestral body size than do extant taxa in fishes. *Acta Zool.* 90:357–384.
- Alfaro, M. E., F. Santini, C. Brock, H. Alamillo, A. Dornburg, D. L. Rabosky, G. Carnevale, and L. J. Harmon. 2009. Nine exceptional radiations plus high turnover explain species diversity in jawed vertebrates. *Proc. Natl. Acad. Sci. USA* 106:13410–13414.
- Beaumont, M. A., W. Y. Zhang, and D. J. Balding. 2002. Approximate Bayesian computation in population genetics. *Genetics* 162:2025–2035.
- Bokma, F. 2010. Time, species and separating their effects on trait variance in clades. *Syst. Biol.* 59:602–607.
- Butler, M. A., and A. A. King. 2004. Phylogenetic comparative analysis: a modeling approach for adaptive evolution. *Am. Nat.* 164:683–695.
- Collar, D. C., T. J. Near, and P. C. Wainwright. 2005. Comparative analysis of morphological diversity: does disparity accumulate at the same rate in two lineages of centrarchid fishes? *Evolution* 59:1783–1794.
- Cook, S. R., A. Gelman, and D. B. Rubin. 2006. Validation of software for Bayesian models using posterior quantiles. *J. Comput. Graph. Stat.* 15:675–692.
- Csillery, K., M. G. B. Blum, O. E. Gaggiotti, and O. Francois. 2010. Approximate Bayesian computation (ABC) in practice. *Trends Ecol. Evol.* 25:410–418.
- Deméré, T. A., A. Berta, and P. J. Adam. 2003. Pinnipedimorph evolutionary biogeography. *B Am. Mus. Nat. Hist.* 13:32–76.

- Eastman, J. M., M. E. Alfaro, P. Joyce, A. L. Hipp, and L. J. Harmon. in press. A novel comparative method for identifying shifts in the rate of character evolution on trees. *Evolution*, doi:10.1111/j.1558-5646.2011.01401.x
- Eizirik, E., W. J. Murphy, K. P. Koepfli, W. E. Johnson, J. W. Drago, R. K. Wayne, and S. J. O'Brien. 2010. Pattern and timing of diversification of the mammalian order Carnivora inferred from multiple nuclear gene sequences. *Mol. Phylogenet. Evol.* 56:49–63.
- Felsenstein, J. 1985. Phylogenies and the comparative method. *Am. Nat.* 125:1–15.
- Finarelli, J. A., and J. J. Flynn. 2006. Ancestral state reconstruction of body size in the Caniformia (Carnivora, Mammalia): the effects of incorporating data from the fossil record. *Syst. Biol.* 55:301–313.
- FitzJohn, R. G. 2010. Quantitative traits and diversification. *Syst. Biol.* 59:619–633.
- FitzJohn, R. G., W. P. Maddison, and S. P. Otto. 2009. Estimating trait-dependent speciation and extinction rates from incompletely resolved phylogenies. *Syst. Biol.* 58:595–611.
- Garland, T., A. W. Dickerman, C. M. Janis, and J. A. Jones. 1993. Phylogenetic analysis of covariance by computer simulation. *Syst. Biol.* 42:265–292.
- Gelman, A., and D. B. Rubin. 1992. Inference from iterative simulation using multiple sequences. *Stat. Sci.* 7:457–472.
- Gittleman, J. L., and A. Purvis. 1998. Body size and species-richness in carnivores and primates. *Proc. R. Soc. Lond. B* 265:113–119.
- Hansen, T. F. 1997. Stabilizing selection and the comparative analysis of adaptation. *Evolution* 51:1341–1351.
- Hansen, T. F., and E. P. Martins. 1996. Translating between microevolutionary process and macroevolutionary patterns: the correlation structure of interspecific data. *Evolution* 50:1404–1417.
- Harmon, L. J., J. B. Losos, T. J. Davies, R. G. Gillespie, J. L. Gittleman, W. B. Jennings, K. H. Kozak, M. A. McPeck, F. Moreno-Roark, T. J. Near, et al. 2010. Early bursts of body size and shape evolution are rare in comparative data. *Evolution* 64:2385–2396.
- Harmon, L. J., J. A. Schulte, A. Larson, and J. B. Losos. 2003. Tempo and mode of evolutionary radiation in iguanian lizards. *Science* 301:961–964.
- Hunt, G. 2006. Fitting and comparing models of phyletic evolution: random walks and beyond. *Paleobiology* 32:578–601
- Jones, K. E., J. Bielby, M. Cardillo, S. A. Fritz, J. O'Dell, C. D. L. Orme, K. Safi, W. Sechrest, E. H. Boakes, C. Carbone, et al. 2009. PanTHERIA: a species-level database of life history, ecology, and geography of extant and recently extinct mammals. *Ecology* 90:2648.
- Joyce, P., and P. Marjoram. 2008. Approximately sufficient statistics and Bayesian computation. *Stat. Appl. Genet. Mol.* 7, article 26.
- Leuenberger, C., and D. Wegmann. 2010. Bayesian computation and model selection without likelihoods. *Genetics* 184:243–252.
- Maddison, W. P., P. E. Midford, and S. P. Otto. 2007. Estimating a binary character's effect on speciation and extinction. *Syst. Biol.* 56:701–710.
- Marjoram, P., J. Molitor, V. Plagnol, and S. Tavaré. 2003. Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* 100:15324–15328.
- McPeck, M. A. 1995. Testing hypotheses about evolutionary change on single branches of a phylogeny using evolutionary contrasts. *Am. Nat.* 145:686–703.
- Nee, S., R. M. May, and P. H. Harvey. 1994. The reconstructed evolutionary process. *Phil. Trans. R. Soc. Lond. B.* 344:305–311.
- Nowak, R. M. 1999. Walker's mammals of the world. The Johns Hopkins Univ. Press, Baltimore.
- O'Meara, B. C., C. Ane, M. J. Sanderson, and P. C. Wainwright. 2006. Testing for different rates of continuous trait evolution using likelihood. *Evolution* 60:922–933.
- Pie, M. R., and J. S. Weitz. 2005. A null model of morphospace occupation. *Am. Nat.* 166:E1–E13.
- Polly, P. D. 2001. Paleontology and the comparative method: ancestral node reconstructions versus observed node values. *Am. Nat.* 157:596–609.
- Pybus, O. G., and P. H. Harvey. 2000. Testing macro-evolutionary models using incomplete molecular phylogenies. *Proc. R. Soc. Lond. B* 267:2267–2272.
- Rabosky, D. L. 2010. Extinction rates should not be estimated from molecular phylogenies. *Evolution* 64:1816–1824.
- Rabosky, D. L., and I. J. Lovette. 2008a. Density-dependent diversification in North American wood warblers. *Proc. R. Soc. Lond. B* 275:2363–2371.
- . 2008b. Explosive evolutionary radiations: decreasing speciation or increasing extinction through time? *Evolution* 62:1866–1875.
- Rabosky, D. L., S. C. Donnellan, A. L. Talaba, and I. J. Lovette. 2007. Exceptional among-lineage variation in diversification rates during the radiation of Australia's most diverse vertebrate clade. *Proc. R. Soc. Lond. B* 274:2915–2923.
- Raup, D. M., S. J. Gould, T. J. M. Schopf, and D. S. Simberloff. 1973. Stochastic models of phylogeny and evolution of diversity. *J. Geol.* 81:525–542.
- Revell, L. J. 2008. On the analysis of evolutionary change along single branches in a phylogeny. *Am. Nat.* 172:140–147.
- Revell, L. J., and D. C. Collar. 2009. Phylogenetic analysis of the evolutionary correlation using likelihood. *Evolution* 63:1090–1100.
- Revell, L. J., L. J. Harmon, R. B. Langerhans, and J. J. Kolbe. 2007. A phylogenetic approach to determining the importance of constraint on phenotypic evolution in the neotropical lizard *Anolis cristatellus*. *Evol. Ecol. Res.* 9:261–282.
- Revell, L. J., D. L. Mahler, P. R. Peres-Neto, and B. D. Redelings. in press. A new phylogenetic method for identifying exceptional phenotypic diversification. *Evolution*, doi:10.1111/j.1558-5646.2011.01435.x
- Ricklefs, R. E. 2006. Time, species, and the generation of trait variance in clades. *Syst. Biol.* 55:151–159.
- Schluter, D., T. Price, A. O. Mooers, and D. Ludwig. 1997. Likelihood of ancestor states in adaptive radiation. *Evolution* 51:1699–1711.
- Sidlauskas, B. 2007. Testing for unequal rates of morphological diversification in the absence of a detailed phylogeny: a case study from characiform fishes. *Evolution* 61:299–316.
- Simpson, G. G. 1953. The major features of evolution. Columbia University Press, New York, USA.
- Stadler, T. in press. Simulating trees on a fixed number of species. *Syst. Biol.*, doi:10.1093/sysbio/syr029
- Tavaré, S., D. J. Balding, R. C. Griffiths, and P. Donnelly. 1997. Inferring coalescence times from DNA sequence data. *Genetics* 145:505–518.
- Thomas, G. H., R. P. Freckleton, and T. Szekely. 2006. Comparative analyses of the influence of developmental mode on phenotypic diversification rates in shorebirds. *Proc. R. Soc. Lond. B* 273:1619–1624.
- Thomas, G. H., S. Meiri, and A. B. Phillimore. 2009. Body size diversification in *Anolis*: novel environment and island effects. *Evolution* 63:2017–2030.
- Wegmann, D., C. Leuenberger, and L. Excoffier. 2009. Efficient approximate Bayesian computation coupled with Markov chain Monte Carlo without likelihood. *Genetics* 182:1207–1218.
- Wegmann, D., C. Leuenberger, S. Neuenschwander, and L. Excoffier. 2010. ABCtoolbox: a versatile toolkit for approximate Bayesian computations. *BMC Bioinform.* 11:116.
- Wozencraft, W. C. 2005. Order carnivora. Pp. 532–628 in D. E. Wilson and D. M. Reeder, eds. *Mammal species of the world: a taxonomic and geographic reference*. The Johns Hopkins Univ. Press, Baltimore.

Associate Editor: G. Hunt