Version dated: March 3, 2017

SOFTWARE FOR SYSTEMATICS AND EVOLUTION

POUMM: An R-package for Bayesian Inference of Phylogenetic Heritability

Venelin Mitov[1,2], Tanja Stadler[1,2]

[1]*Swiss Federal Institute of Technology in Zurich, Switzerland;*

[2]*Swiss Institute of Bioinformatics, Switzerland*

Corresponding authors: Venelin Mitov, Department of Biosystem Sciences and Engineering, Swiss Federal Institute of Technology, Mattenstrasse 26, CH-4058 Basel, Switzerland; E-mail: `vmitov@gmail.com`.

Abstract.—The Phylogenetic Ornstein-Uhlenbeck Mixed Model (POUMM) allows to estimate the phylogenetic heritability of continuous traits, to test hypotheses of neutral evolution versus stabilizing selection, to quantify the strength of stabilizing selection, to estimate measurement error and to make predictions about the evolution of a phenotype and phenotypic variation in a population. Despite this variety of applications, currently, there are no R-packages supporting POUMM inference on large non-ultrametric phylogenetic trees. Large phylogenies of that kind are becoming increasingly available, predominantly in epidemiology, where transmission trees are inferred from pathogen sequences during epidemic outbreaks, but also in some macroevolutionary studies incorporating fossil and contemporary data. In this article, we propose the R-package **POUMM**, providing Bayesian inference of the model parameters on large phylogenetic trees. We describe a novel breadth-first pruning algorithm for fast likelihood calculation, enabling highly parallelizable likelihood calculation on multi-core systems and GPUs. We report simulation-based results proving the technical correctness and performance of the software.

Keywords: PMM, Brownian motion, Ornstein-Uhlenbeck, measurement error, Bayesian inference

The past decades have seen active developement of phylogenetic models of continuous trait evolution, progressing from null neutral models, such as single-trait Brownian motion (BM), to complex multi-trait models incorporating selection, interaction between trait-values and diversification, and co-evolution of multiple traits (O'Meara 2012; Manceau, Lambert, and Morlon 2016). Fitting these models to data has become possible thanks to a growing collection of software packages, many of which written in the R language of statistical computing (R Core Team 2013) and freely available on the Comprehensive R Archive Network (CRAN) (O'Meara 2016).

The **phylogenetic heritability**, introduced with the phylogenetic mixed model (PMM) (Housworth, Martins, and Lynch 2004), measures the proportion of phenotypic variance in a population attributable to heritable factors, such as genes, as opposed to non-heritable factors, such as environment and measurement error. Although the concept of phylogenetic heritability has been applied mostly in the context of the original PMM, i.e. under the assumption of Brownian motion, the same concept applies to any evolutionary model allowing for the estimation of measurement error (ME) (Hansen and Bartoszek 2012). In its simplest form this means adding a white noise error term to the modeled trait-value. Therefore, it comes as a surprise that most recently published R-packages for phylogenetic analysis on large trees have very limited support for estimating ME and, thus, phylogenetic heritability. To give a few examples, the package Rphylopars (Goolsby, Bruggeman, and Ané 2016) allows for the estimation of intraspecies standard error only when the tips in the phylogeny are grouped with several tips per species; diversitree (FitzJohn 2012) only allows for the specification of a parameter `states.sd` through a call to `make.bm` or another function, but does not fit this parameter; geiger (Pennell et al. 2014) allows for fitting a standard error (`SE`), but similarly to diversitree and Rphylopars, does not support likelihood calculation on non-ultrametric trees; GLSME (Hansen and Bartoszek 2012) and RPANDA (Manceau, Lambert, and Morlon 2016) have a rich ME-support for both, ultrametric and non-ultrametric trees, but do not provide fast likelihood calculation on large trees.

Here, we introduce the R-package **POUMM** providing Bayesian inference of phylogenetic heritability for traits evolving under stabilizing selection. Formally, this is an extension of the PMM, replacing the Brownian motion process by an Ornstein-Uhlenbeck process with a single global optimum (Ornstein and Zernike 1919; Uhlenbeck and Ornstein 1930). The package implements a highly parallelizable breadth-first pruning algorithm for fast likelihood calculation on large ultrametric and non-ultrametric trees including polytomies. The same algorithm can be extended to multi-trait scenarios with different model-regimes assigned to different phylogenetic lineages. We present the model, the algorithm and simulation based results for validation of the technical correctness and performance of the software.

Through the rest of the article we will rely on the following setup. Given is a rooted phylogenetic tree $\mathcal{T}$ with N tips indexed by $1, ..., N$ and a root node, 0 (Fig. 1). Without restrictions on the tree topology, non-ultrametric trees (i.e. tips have different time-distance from the root) and polytomies (i.e. nodes with any finite number of descendants) are accepted. Internal nodes are indexed by the numbers $N+1, ...$ Associated with the tips is a N-vector of observed real trait-values denoted by $\mathbf{z}$. We denote by $\mathcal{T}_i$ the subtree rooted at node i and by $\mathbf{z}_i$ the set of values at the tips belonging to $\mathcal{T}_i$. For any internal node j, we denote by $Desc(j)$ the set of its direct descendants. Furthermore, for any $i \in Desc(j)$, we denote by t_{ji} the length of the edge connecting j with i and by t_{0i} the sum of edge-lengths (time-distance) from the root to i. For two tips i and k, we denote by $t_{0(ik)}$ the time-distance from the root to their most recent common ancestor (mrca), and by τ_{ik} the sum of edge-lengths on the path from i to k (also called phylogenetic/patristic distance between i and k). We use the simbol $\bar{t}$ to denote the mean root-tip distance in the tree. For converting branch-lengths in time-units into absolute time, by convention, the origin of time, 0, is assumed to be at the root, and the time is increasing positively towards the tips.

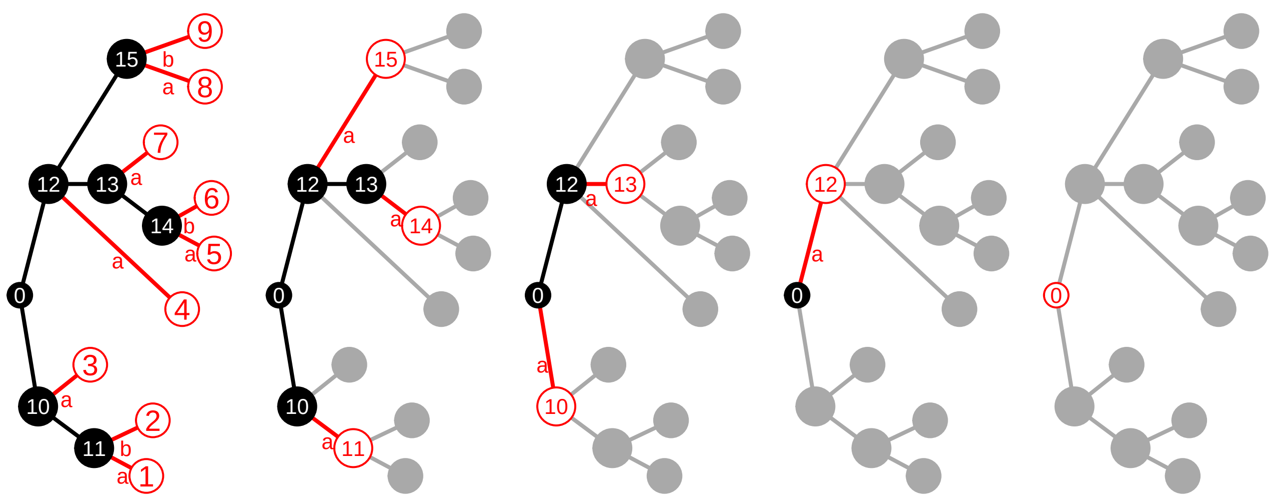

Figure 1: Breadth-first pruning on a tree with $N = 10$ tips. Each tree from left to right depicts one pruning iteration; black: non-tip nodes at a current pruning step; red: tip nodes to be pruned; grey: pruned nodes. Letters 'a' and 'b' next to branches denote the order in which the coefficients a_{ji}, b_{ji}, c_{ji} are added to their parent's a_j, b_j and c_j (algorithm 1).

The Phylogenetic Ornstein-Uhlenbeck Mixed Model

The phylogenetic Ornstein-Uhlenbeck mixed model (POUMM) decomposes the trait value as a sum of a non-heritable component, e, and a genetic component, g, which (i) evolves continuously according to an Ornstein-Uhlenbeck (OU) process along branches; (ii) gets inherited by the branches descending from each internal node. In biological terms, g is a genotypic value (Lynch and Walsh 1998) that evolves according to random drift with stabilizing selection towards a global optimum; e is a non-heritable component, which can be interpreted in different ways, depending on the application, i.e. a measurement error, an environmental contribution, a residual with respect to a model prediction, or the sum of all these. The OU-process acting on g is parameterized by an initial genotypic value at the root, g_0, a global optimum, θ, a selection strength, $\alpha>0$, and a random drift unit-time standard deviation, σ. Denoting by W_t the standard Wiener process (Grimmett and Stirzaker 2001), the evolution of the trait-value, $z(t)$, along a given lineage of the tree is described by the equations:

$$z(t) = g(t) + e \tag{1}$$

$$dg(t) = \alpha[\theta - g(t)]dt + \sigma dW_t, \tag{2}$$

The stochastic differential equation 2 defines the OU-process, which represents a random walk tending towards the global optimum θ with stronger attraction for bigger difference between $g(t)$ and θ (Ornstein and Zernike 1919; Uhlenbeck and Ornstein 1930). The model assumptions for e are that they are iid normal with mean 0 and standard deviation σ_e at the tips. Any process along the tree that gives rise to this distribution at the tips may be assumed for e. For example, in the case of epidemics, a newly infected individual is assigned a new e-value which represents the contribution from its immune system and this value can change or remain constant throughout the course of infection. In the case of macroevolution, e may represent the ecological (non-genetic) differences between species. In particular, the non-heritable component e does not influence the behavior of the OU-process $g(t)$. Thus, if we were to simulate trait values z on the tips of a phylogenetic tree $\mathcal{T}$, we could first simlate the OU-process from the root to the tips to obtain g, and then add the white noise e (i.e. an iid draw from a normal distribution) to each simulated g value at the tips.

The POUMM represents an extension of the phylogenetic mixed model (PMM) (Lynch 1991; Housworth, Martins, and Lynch 2004), since, in the limit $\alpha \to 0$, the OU-process converges to a Brownian motion (BM) with unit-time standard deviation σ. Both, the POUMM and the PMM, define an expected multivariate normal distribution for the trait values at the tips. Note that the trait expectation and variance for a tip i depends on its time-distance from the root (t_{0i}), and the trait covariance for a pair of tips (ij) depends on the time-distance from the root to their mrca ($t_{0(ij)}$), as well as their patristic distance (τ_{ij}) (table 1).

We note that the expressions for the expected variance-covariance matrix of the POUMM are only defined for strictly positive α. We obtain the limit for PMM by noting that $\lim_{\alpha \to 0} \alpha/(1 - e^{\alpha t}) = -1/t$.

Phylogenetic heritability

The **phylogenetic heritability** is defined as the expected proportion of phenotypic variance attributable to g at the tips of the tree, $\sigma^2(g)/\left[\sigma^2(g) + \sigma_e^2\right]$ (Housworth, Martins, and Lynch 2004). This definition is a phylogenetic variant of the definition of broad-sense heritability, H^2, from quantitative genetics (Lynch and Walsh 1998). However, in the case of a trait evolving along a phylogeny, the expected genotypic

Table 1: Population properties at the tips of the phylogeny under POUMM and PMM. μ_i: expected value at tip i; Σ_{ii}: expected variance for tip i; Σ_{ij}: expected covariance of the values of tips i and j; $H^2_{\bar{t}}$: phylogenetic heritability at mean root-tip distance; H^2_{∞}: phylogenetic heritability at long-term equilibrium; H^2_e: time-independent (empirical) phylogenetic heritability.

	POUMM	PMM ($\alpha \to 0$)
Θ:	$<g_0,\, \alpha,\, \theta,\, \sigma,\, \sigma_e>$	$<g_0,\, \sigma,\, \sigma_e>$
$\mu_i(\Theta, \mathcal{T})$:	$e^{-\alpha t_{0i}} g_0 + (1 - e^{-\alpha t_{0i}})\, \theta$	g_0
$\Sigma_{ii}(\Theta, \mathcal{T})$:	$\sigma^2 \frac{\left(1-e^{-2\alpha t_{0i}}\right)}{2\alpha} + \sigma_e^2$	$\sigma^2\, t_{0i} + \sigma_e^2$
$\Sigma_{ij}(\Theta, \mathcal{T})$:	$\sigma^2 \frac{e^{-\alpha \tau_{ij}}\left(1-e^{-2\alpha t_{0(ij)}}\right)}{2\alpha}$	$\sigma^2\, t_{0(ij)}$
$H^2_{\bar{t}}$:	$\frac{\sigma^2\left(1-e^{-2\alpha\bar{t}}\right)}{\sigma^2\left(1-e^{-2\alpha\bar{t}}\right)+2\alpha\sigma_e^2}$	$\bar{t}\sigma^2/(\bar{t}\sigma^2 + \sigma_e^2)$
H^2_{∞}:	$\sigma^2/(\sigma^2 + 2\alpha\sigma_e^2)$	1
H^2_e:	$1 - \sigma_e^2/s^2(\mathbf{z})$	$1 - \sigma_e^2/s^2(\mathbf{z})$

variance, $\sigma^2(g)$, and, therefore, the phylogenetic heritability, are functions of time. The POUMM package reports the following three types of phylogenetic heritability (see table 1 for simplified expressions):

- Expectation at the mean root-tip distance : $H^2_{\bar{t}} := \left[\sigma^2 \frac{\left(1-e^{-2\alpha\bar{t}}\right)}{2\alpha}\right] / \left[\sigma^2 \frac{\left(1-e^{-2\alpha\bar{t}}\right)}{2\alpha} + \sigma_e^2\right]$;
- Expectation at equilibrium of the OU-process: $H^2_{\infty} := \lim_{\bar{t}\to\infty} H^2_{\bar{t}}$;
- Empirical (time-independent) version of the heritability based on the sample phenotypic variance $s^2(\mathbf{z})$: $H^2_e := 1 - \sigma_e^2/s^2(\mathbf{z})$.

Algorithm

For a fixed tree, $\mathcal{T}$, the log-likelihood of the observed data is defined as the function:

$$\ell\ell(\Theta) = \ln(f(\mathbf{z}_0|\mathcal{T}, \Theta)), \tag{3}$$

where f denotes a probability density function (pdf) and $\Theta =< g_0, \alpha, \theta, \sigma, \sigma_e >$.

The POUMM package uses a breadth-first variant of the pruning algorithm (Felsenstein 1973). The log-likelihood is calculated by consecutive integration over the unobservable genotypic values, g_i, progressing from the tips to the root. Central for the pruning likelihood calculation is the following theorem, for which we provide a proof in the appendix:

Theorem 1 (Quadratic polynomial representation of the POUMM log-likelihood)**.** *For $\alpha \geq 0$, a real θ and non-negative σ and σ_e, the POUMM log-likelihood can be expressed as a quadratic polynomial of g_0:*

$$\ell\ell(\Theta) = a_0 g_0^2 + b_0 g_0 + c_0, \tag{4}$$

where $a_0 < 0$, b_0 and c_0 are real coefficients. We denote by $u(\alpha, t)$ the function:

$$u(\alpha, t) := \begin{cases} \alpha/(1 - e^{\alpha t}), & \text{for } \alpha > 0 \\ -1/t, & \text{for } \alpha = 0 \end{cases} \tag{5}$$

Then, the coefficients in eq. 4 can be expressed with the following recurrence relation:

1. For $j \in \{1, ..., N\}$ (tips):

$$a_j = -\frac{1}{2\sigma_e^2};\ b_j = \frac{z_j}{\sigma_e^2};\ c_j = -\frac{z_j^2}{2\sigma_e^2} - \ln\sqrt{2\pi\sigma_e^2} \tag{6}$$

2. *For $j > N$ (internal nodes) or $j = 0$ (root):*

$$
\begin{aligned}
a_j &= \sum_{i \in Desc(j)} \frac{a_i u(\alpha, 2t_{ji})}{u(\alpha, 2t_{ji}) - \alpha + \sigma^2 a_i} \\
b_j &= \sum_{i \in Desc(j)} \frac{u(\alpha, 2t_{ji}) \left[2\theta a_i (e^{\alpha t_{ji}} - 1) + b_i e^{\alpha t_{ji}}\right]}{u(\alpha, 2t_{ji}) - \alpha + \sigma^2 a_i} \\
c_j &= \sum_{i \in Desc(j)} \Bigg\{ c_i + \alpha t_{ji} - \frac{0.25\, b_i^2 \sigma^2}{-\alpha + a_i \sigma^2 + u(\alpha, 2t_{ji})} - \\
&\qquad 0.5 \ln \left(\frac{-\alpha + a_i \sigma^2 + u(\alpha, 2t_{ji})}{u(\alpha, 2t_{ji})} \right) + \\
&\qquad \frac{\alpha\theta \left[a_i \theta - (b_i + a_i \theta) e^{\alpha t_{ji}}\right]}{u(\alpha, t_{ji}) + (-\alpha + a_i \sigma^2)\left(1 + e^{\alpha t_{ji}}\right)} \Bigg\}.
\end{aligned}
\tag{7}
$$

It can be shown that current pruning implementations (FitzJohn 2012; Pennell et al. 2014) rely on equivalent formulations of the above theorem. The breadth-first algorithm differs from these implementations in the ordering of algebraic operations so that they can be performed "at once" for groups of tips or internal nodes rather than consecutively for individual nodes in order of depth-first traversal.

Implementation

Before model fitting, the user can choose from different POUMM parametrizations and prior settings (function `specifyPOUMM`). Model fitting is done through a combination of likelihood optimization and adaptive Metropolis sampling (Vihola 2012; Scheidegger 2012). A set of standard generic functions, such as `plot`, `summary`, `logLik`, `coef`, etc., provide means to assess the quality of a fit (i.e. MCMC convergence, consistence between ML and MCMC fits) as well as various inferred properties, such as high posterior density (HPD) intervals.

We implemented the breadth-first pruning algorithm in R and in C++ using the library Armadillo (Sanderson and Curtin 2016) through the R-package RcppArmadillo

Algorithm 1: Breadth-first pruning

Data: $\mathcal{T}$, $\mathbf{z}$; α, θ, σ, σ_e
Result: $\max_{g_0} \ell\ell(g_0, \alpha, \theta, \sigma, \sigma_e; \mathbf{z}, \mathcal{T})$
initialization:
for tips $i \in \{1, ..., N\}$, set a_i, b_i, c_i (eq. 6);
for nodes $j > N$ or $j = 0$, set a_j, b_j, c_j to 0;
set $\{< ji >\}$ to the set of edges $< ji >$ in $\mathcal{T}$, where $i \in \{1, ..., N\}$;
while $\{< ji >\} \neq \phi$ **do**
  for $< ji > \in \{ji\}$ **do**
    // vectorized operations
    set $a_{<ji>}$, $b_{<ji>}$, $c_{<ji>}$ to the sub-summands in eq. 7 ;
    add $a_{<ji>}$, $b_{<ji>}$, $c_{<ji>}$ to a_j, b_j, c_j (see branch labels on Fig. 1);
  end
  pruning: set $\mathcal{T}$ to the tree obtained upon removal of $i \in \{< ji >\}$;
  set $\{i\}$ to to the subset of parent nodes in $\{< ji >\}$, which have become tips after the pruning (Fig. 1);
  set $\{< ji >\}$ to the edges leading to $\{i\}$;
end
set $g_0 := -0.5\, b_0/a_0$;
set $\ell\ell(\Theta) := a_0 g_0^2 + b_0 g_0 + c_0$.

(Eddelbuettel and Sanderson 2014). While slightly slower, the R implementation can switch transparently between `double` and `Rmpfr` floating point precision (Maechler 2014), thus, guaranteeing numerical stability in cases of extreme parameter values, trait values or branch lengths.

In addition the POUMM package uses the following third-party R-packages: ape v3.4 (Paradis, Claude, and Strimmer 2004), data.table v1.9.6 (Dowle et al. 2014) and coda v0.18-1 (Plummer et al. 2006), foreach v1.4.3 (Analytics and Weston 2015), ggplot2 v2.1.0 (Wickham 2009) and gsl v1.9-10.3 (Hankin 2006).

Simulations

To validate the correctness of the Bayesian POUMM implmentation, we used the method of posterior quantiles (Cook, Gelman, and Rubin 2006). In this method, the idea is to generate samples from the posterior quantile distributions of selected model parameters (or functions thereof) by means of numerous "replications" of simulation followed by Bayesian parameter inference. In each replication, "true" values of the model parameters are drawn from a fixed prior distribution and trait-data is simulated under the model specified by these parameter values. Then, the posterior quantiles of the "true" parameter values (or functions thereof) are calculated from the corresponding posterior samples generated by the to-be-tested software. By running in parallel multiple independent replications on a fixed prior, it is possible to generate large samples from the posterior quantile distributions of the individual model parameters, as well as any derived quantities. Assuming correctness of the simulations, any statistically significant deviation from uniformity of these posterior quantile samples indicates an error in the to-be-tested software (Cook, Gelman, and Rubin 2006).

In order to test the robustness of the POUMM against model mis-specifications, we extended the above approach by running POUMM inference on data simulated under pure PMM ($\alpha = 0$), and PMM inference on data simulated under POUMM ($\alpha \geq 0$). Simulations scenarios of 2000 replications were run on an ultrametric and non-ultrametric tree ($N = 4000$), using the parametrization $\Theta =< \alpha,\, \theta,\, H_{\bar{t}}^2,\, \sigma_e,\, g_0 >$ and the prior $\Theta \sim \text{Exp}(0.1) \times \mathcal{U}(2, 8) \times \mathcal{U}(0, 1) \times \text{Exp}(1) \times \mathcal{N}(5, 25)$ (Supplementary Text).

Without exception, both, the PMM and POUMM implementation, generate uniformly distributed posterior quantiles for all trees and all relevant parameters when the Bayesian inference has been done on data simulated under the correct simulation mode, i.e. "Simulate BM" for PMM and "Simulate OU" for POUMM. This is confirmed visually by observing the corresponding histograms on Fig. 2, as well as statistically, by a non-significant p-value from a Kolmogorov-Smirnov uniformity test at the 0.01 level. This observation validates the technical correctness of the software.

When fitting the PMM to simulations of stabilizing selection (OU), there is a highly

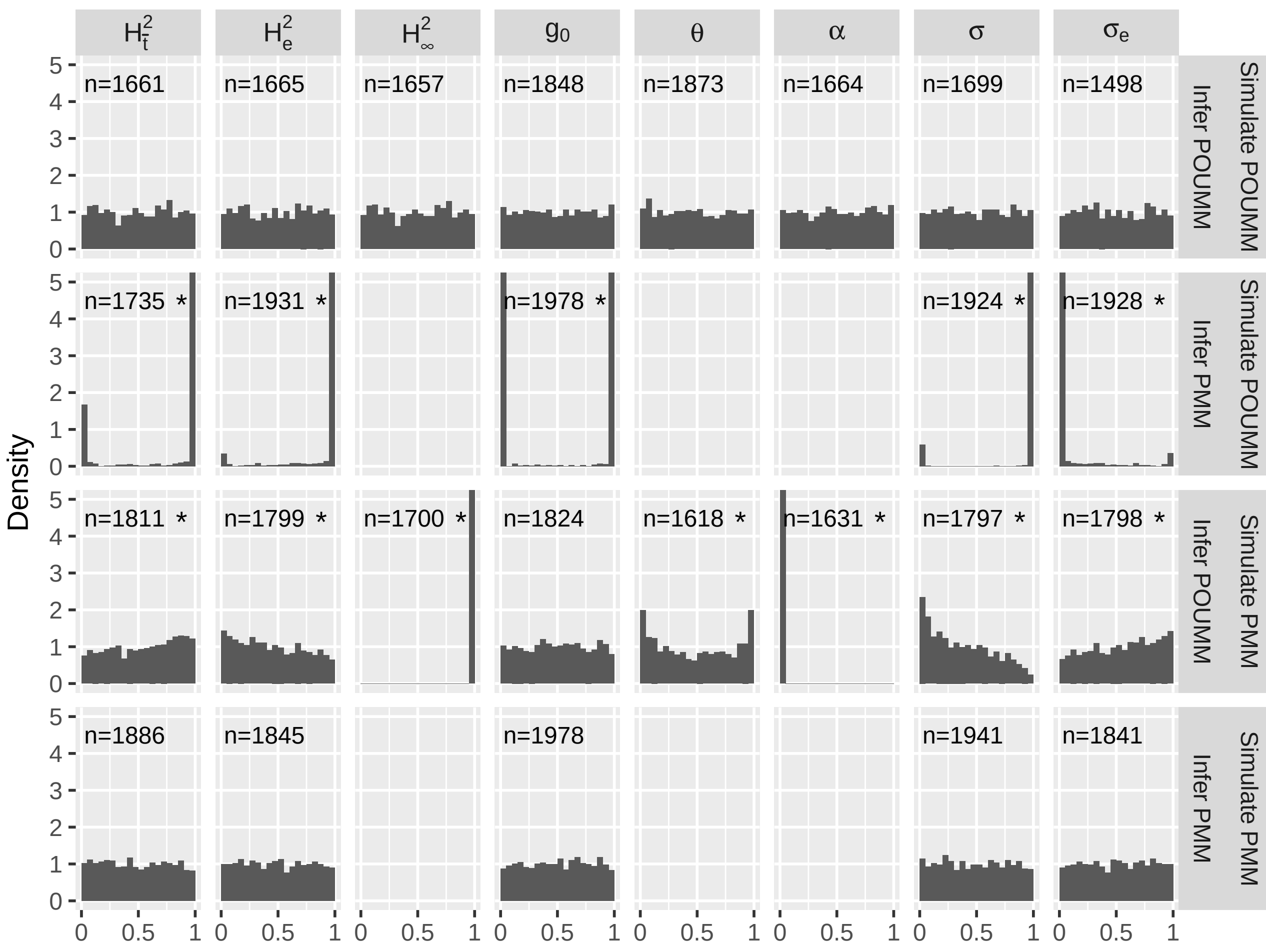

Figure 2: Posterior quantiles from simulation scenarios on a non-ultrametric tree ($N = 4000$). Values tending to 1 indicate that the true value dominates the inferred posterior sample for most of the replications. This means that the model fit tends to underestimate the true parameter. The number n at the top of each histogram denotes the number of replications out of 2000 which reached acceptable MCMC convergence and mixing at the by the one milionth iteration. An asterisk indicates significant uniformity violation (Kolmogorov-Smirnov P-value < 0.01).

significant deviation from uniformity of the posterior quantiles for the parameters $H^2_{\bar{t}}$, g_0, and σ_e and H^2_e. The fact that most posterior quantiles for $H^2_{\bar{t}}$ and H^2_e are at the extremes of the histogram is indicative for a systematic negative or positive bias in the inferred parameters. These results indicate that the PMM can be a very unstable erronous estimator of phylogenetic heritability when the data violates the Brownian motion assumption.

Discussion

A main advantage of a breadth-first approach with respect to to depth-first pruning implementations (e.g. diversitree (FitzJohn 2012) and geiger (Pennell et al. 2014)) is that most of the algebraic calculations are done on vectors instead of single numbers. Contemporary computer architectures and languages such as Matlab and R are optimized for vector operations. Therefore, an implementation of breadth-first pruning written in R is nearly as fast as an analogous (breadth-first-) or a depth-first implementation written in C++ (Supplementary Materials). Moreover, on multi-core systems, a breadth-first implementation can be easily parallelized by linking to OpenMP- or GPU-accelerated libraries.

The OU process has been applied as a model for stabilizing selection in macro-evolutionary studies (LANDE 1976; Felsenstein 1988; Hansen 1997; Harmon et al. 2010). Most of theses studies and the accompanying software packages assume that the whole trait evolves according to an OU process, usually disregarding the presence of a biologically relevant non-heritable component or of a measurement error with a-priori known variance (FitzJohn 2012). When modelling species trait evolution, a non-heritable ecological contribution may be well justified and may in fact be important to understand the full evolutionary process. When modeling pathogen evolution, the branching points in the tree represent transmission events, and the environmental contribution is the contribution of the host immune system. Thus, for pathogens, it is crucial to incorporate e in the model in order to quantify the importance of host- versus pathogen factors in trait formation (Alizon et al. 2010; Shirreff et al. 2013).

The idea to infer phylogenetic heritability assuming that g follows an OU process along the phylogeny has so far been discouraged mainly for interpretational and practical reasons: (i) in biology, individuals get selected based on their whole trait-values z, rather than the genotypic component g (unless e is simply measurment noise); (ii) small ultrametric macro-evolutionary trees do not contain sufficient signal for a simultaneous inference of the OU-and environmental variance (Housworth, Martins, and Lynch 2004). We argue that modeling an OU process on z rather than g comes at the cost of additional parameters and reduced statistical power, because it necessitates to account for jumps in z at the branching points as well as the unobserved speciation/transmission events along the tree. Conversely, assuming that the OU process acts directly on g rather than z is mathematically more convenient, because it allows the inference of a single continuous OU-process along the tree, while adding e only at the tips of the tree.

Finally, our simulations suggest that the POUMM could make a suitable estimator of phylogenetic heritability when the trait is subject to stabilizing selection, but also, tends to be more robust than PMM towards model mis-specification (Fig. 2). Thus, the POUMM R-package should provide a useful tool for furture phylogenetic analysis in epidemiology and macro-evolution.

Supplementary Material

The proof of theorem 1 and further details on the simulation setup can be found in an online appendix. A performance benchmark for the breadth-first pruning algorithm is provided in the supplementary file `CompareOUPackages.html`. The user manual for the POUMM package is provided in the package vignette.

Funding

V.M. and T.S. thank ETH Zürich for funding. T.S. is supported in part by the European Research Council under the 7th Framework Programme of the European Commission (PhyPD: Grant Agreement Number 335529).

Acknowledgements

We thank Dr. Krzysztof Bartoszek for valuable insights on the Ornstein-Uhlenbeck process.

References

Alizon, Samuel, Viktor von Wyl, Tanja Stadler, Roger D Kouyos, Sabine Yerly, Bernard Hirschel, Jürg Böni, et al. 2010. "Phylogenetic approach reveals that virus genotype largely determines HIV set-point viral load." *PLoS Pathogens* 6 (9): e1001123.

Analytics, Revolution, and Steve Weston. 2015. "Provides Foreach Looping Construct for R [R package foreach version 1.4.3]."

Cook, Samantha R, Andrew Gelman, and Donald B Rubin. 2006. "Validation of Software for Bayesian Models Using Posterior Quantiles." *Journal of Computational and Graphical Statistics* 15 (3): 675–92.

Dowle, Matthew, T Short, S Liangolou, and A Srinivasan. 2014. "data.table: Extension of data.frame," July, 9.

Eddelbuettel, Dirk, and Conrad Sanderson. 2014. "RcppArmadillo - Accelerating R with high-performance C++ linear algebra." *Computational Statistics & Data Analysis* 71: 1054–63.

Felsenstein, J. 1973. "Maximum-likelihood estimation of evolutionary trees from continuous characters." *American Journal of Human Genetics* 25 (5): 471–92.

———. 1988. "Phylogenies And Quantitative Characters." *Annual Review of Ecology and Systematics* 19 (1): 445–71.

FitzJohn, Richard G. 2012. "Diversitree: comparative phylogenetic analyses of diversification in R." *Methods in Ecology and Evolution* 3 (6): 1084–92.

Goolsby, Eric W, Jorn Bruggeman, and Cécile Ané. 2016. "Rphylopars: fast multivariate phylogenetic comparative methods for missing data and within-species variation." *Methods in Ecology and Evolution* 8 (1): 22–27.

Grimmett, Geoffrey, and David Stirzaker. 2001. *Probability and Random Processes*. Oxford University Press.

Hankin, RKS. 2006. "Special functions in R: introducing the gsl package." *R News*.

Hansen, Thomas F. 1997. "Stabilizing Selection and the Comparative Analysis of Adaptation." *Evolution; International Journal of Organic Evolution* 51 (5): 1341–51.

Hansen, Thomas F, and Krzysztof Bartoszek. 2012. "Interpreting the evolutionary regression: the interplay between observational and biological errors in phylogenetic comparative studies." *Systematic Biology* 61 (3): 413–25.

Harmon, Luke J, Jonathan B Losos, T Jonathan Davies, Rosemary G Gillespie, John L Gittleman, W Bryan Jennings, Kenneth H Kozak, et al. 2010. "Early bursts of body size and shape evolution are rare in comparative data." *Evolution; International Journal of Organic Evolution* 64 (8): 2385–96.

Housworth, Elizabeth A, Emília P Martins, and Michael Lynch. 2004. "The phylogenetic mixed model." *The American Naturalist* 163 (1): 84–96.

LANDE, R. 1976. "Natural-Selection and Random Genetic Drift in Phenotypic

Evolution." *Evolution; International Journal of Organic Evolution* 30 (2): 314–34.

Lynch, Michael. 1991. "Methods for the Analysis of Comparative Data in Evolutionary Biology." *Evolution; International Journal of Organic Evolution* 45 (5): 1065–80.

Lynch, Michael, and Bruce Walsh. 1998. *Genetics and Analysis of Quantitative Traits*. Sinauer Associates Incorporated.

Maechler, Martin. 2014. "Rmpfr: R MPFR - Multiple Precision Floating-Point Reliable."

Manceau, Marc, Amaury Lambert, and Helene Morlon. 2016. "A unifying comparative phylogenetic framework including traits coevolving across interacting lineages." *Systematic Biology*, December, syw115.

Ornstein, L S, and F Zernike. 1919. "The theory of the Brownian Motion and statistical mechanics." *Proceedings of the Koninklijke Akademie Van Wetenschappen Te Amsterdam* 21 (1/5): 109–14.

O'Meara, Brian C. 2012. "Evolutionary Inferences from Phylogenies: A Review of Methods." *Dx.doi.org* 43 (1): 267–85.

———. 2016. "CRAN Task View: Phylogenetics, Especially Comparative Methods," December.

Paradis, E, J Claude, and K Strimmer. 2004. "APE: Analyses of Phylogenetics and Evolution in R language." *Bioinformatics* 20 (2): 289–90.

Pennell, Matthew W, Jonathan M Eastman, Graham J Slater, Joseph W Brown, Josef C Uyeda, Richard G FitzJohn, Michael E Alfaro, and Luke J Harmon. 2014. "geiger v2.0: an expanded suite of methods for fitting macroevolutionary models to phylogenetic trees." *Bioinformatics* 30 (15): 2216–8.

Plummer, Martyn, Nicky Best, Kate Cowles, and Karen Vines. 2006. "CODA:

Convergence Diagnosis and Output Analysis for MCMC." *R News* 6: 7–11.

R Core Team. 2013. "R: A Language and Environment for Statistical Computing."

Sanderson, Conrad, and Ryan Curtin. 2016. "Armadillo: a template-based C++ library for linear algebra." *Journal of Open Source Software* 1 (2).

Scheidegger, Andreas. 2012. "adaptMCMC: Implementation of a generic adaptive Monte Carlo Markov Chain sampler."

Shirreff, George, Samuel Alizon, Anne Cori, Huldrych F Günthard, Oliver Laeyendecker, Ard van Sighem, Daniela Bezemer, and Christophe Fraser. 2013. "How effectively can HIV phylogenies be used to measure heritability?" *Evolution, Medicine, and Public Health* 2013 (1): 209–24.

Uhlenbeck, G E, and L S Ornstein. 1930. "On the Theory of the Brownian Motion." *Physical Review* 36 (5): 823–41.

Vihola, Matti. 2012. "Robust adaptive Metropolis algorithm with coerced acceptance rate." *Statistics and Computing* 22 (5): 997–1008.

Wickham, Hadley. 2009. "ggplot2 - Elegant Graphics for Data Analysis." *Use R.*