Inferring phylodynamic parameters in bacterial genomic data using Bayesian and hybrid phylogenetic methods

Sebastian Duchene[1]*, David A Duchene[2], Jemma L Geoghegan[3], Zoe A Dyson[1], Jane Hawkey[1], Kathryn E Holt[1]

[1] Department of Biochemistry and Molecular Biology, Bio21 Molecular Science and Biotechnology Institute, University of Melbourne, Parkville, Victoria 3020, Australia

[2] School of Life and Environmental Sciences, University of Sydney, Sydney, NSW 2006, Australia

[3] Department of Biological Sciences, Macquarie University, Sydney, NSW 2109, Australia

* To whom correspondence should be addressed
sebastian.duchene@unimelb.edu.au

Abstract

Background: Recent developments in sequencing technologies make it possible to obtain genome sequences from a large number of isolates in a very short time. Bayesian phylogenetic approaches can take advantage of these data by simultaneously inferring the phylogenetic tree, evolutionary timescale, and phylodynamic parameters (such as population growth rates), while naturally integrating uncertainty in all parameters. Despite their desirable properties, Bayesian approaches can be computationally intensive, hindering their use for outbreak investigations involving genome data for a large numbers of pathogen isolates. An alternative to using full Bayesian inference is to use a hybrid approach, where the phylogenetic tree and evolutionary timescale are estimated first using maximum likelihood. Under this hybrid approach, demographic parameters are inferred from estimated trees instead of the sequence data, using maximum likelihood, Bayesian inference, or approximate Bayesian computation. This can vastly reduce the computational burden, but has the disadvantage of ignoring the uncertainty in the phylogenetic tree and evolutionary timescale.

Results: We compared the performance of a fully Bayesian and a hybrid method by analysing six whole-genome SNP data sets from a range of bacteria. The estimates from the two methods were very similar, suggesting that the hybrid method is a valid alternative for very large data datasets. However, we also found that congruence between these methods is contingent on the presence of strong temporal structure in the data (i.e. clocklike behaviour), which is typically verified using a date-randomisation test in a Bayesian framework. To reduce the computational burden of this Bayesian test we implemented a date-randomisation test using a rapid maximum likelihood method, which has similar performance to its Bayesian counterpart.

Conclusions: Hybrid approaches can produce reliable inferences of evolutionary timescales and phylodynamic parameters in a fraction of the time required for fully Bayesian analyses. As such, they are a valuable alternative in outbreak studies involving a large number of isolates.

Keywords

Bayesian phylogenetics, phylodynamics, molecular clock, bacterial evolution

Background

Genomic data are increasingly used to investigate infectious disease outbreaks caused by microbial pathogens. Recent developments in sequencing technologies have made it possible to obtain data for a very large number of samples, at low cost and within a very short timeframe. Phylogenetic methods can make use of these data to infer their evolutionary dynamics. For example, genome data obtained during the first months of the 2013-2016 Ebola virus epidemic were used to determine the time of origin of the outbreak and the basic reproductive number (R_0) of the circulating strains [1,2]. Some of the key requirements for these inferences are that the data must have sufficient genetic diversity and that they should be a representative sample of the circulating strains.

Serially sampled data are particularly useful because their sampling times can be used to calibrate the molecular clock. This consists of calculating the rate of evolution, which is the amount of genetic change that has accumulated over the sampling time. The rate of evolution is key to infer an evolutionary timescale, typically represented by a phylogenetic tree where the branch lengths correspond to time, known as a chronogram. Some methods assume that the rate of evolution is constant over time, known as a strict molecular clock, but popular Bayesian implementations, such as that in BEAST [3,4], include relaxed-clock models that use a statistical distribution to describe rate variation across time and lineages (reviewed in [5]). Phylodynamic models can be used to estimate the epidemic growth rate (r), R_0, and other parameters [6,7]. Importantly, these models describe the expectation of the distribution of node times in the chronogram. As such, inferences drawn from phylodynamic models rely on accurate estimates of evolutionary rates and timescales. A number of statistical methods are available to assess the robustness of inferences of evolutionary rates and timescales; those that are most widely used are implemented under a Bayesian framework (reviewed in [8]).

Bayesian phylogenetic approaches allow sophisticated evolutionary models to be specified. For example, the evolution of a pathogen during an outbreak can be defined as an exponentially growing population with considerable evolutionary rate variation among lineages; which can be modelled by specifying a nucleotide substitution model, a relaxed-clock model and an exponential-growth tree prior. The parameters for all these models are obtained simultaneously and their estimates correspond to posterior probability distributions, such that their uncertainty is a natural by-product of the analysis. Bayesian methods require specifying a prior distribution for all parameters. Although specifying a prior distribution is not trivial for some parameters, their influence can be assessed by comparing them to the posterior. An advantage of specifying prior distributions is that it is possible to include previous knowledge about the data. As a case in point, a known probability of sampling can be represented with a prior distribution in birth-death models [9].

Whilst Bayesian phylogenetic methods have many desirable properties, analysing large genomic data sets under complex models is often computationally prohibitive (e.g. [10,11]). An alternative to full Bayesian methods is to conduct the analysis in several steps. In this hybrid approach the phylogenetic tree, evolutionary rates and timescales, and demographic parameters are estimated separately.

Phylogenetic trees can be rapidly estimated using various maximum likelihood implementations [12–15]. These methods assume a substitution model, but not a molecular-clock or demographic model, such that the branch lengths of the trees represent the expected number of substitutions per site, and are known as phylograms.

Next, phylograms can be used to estimate evolutionary rates and chronograms, for example, using a recently developed molecular clock method based on least-squares optimisation, called LSD (Least Squares Dating) [16]. LSD is more computationally tractable than Bayesian molecular-clock methods, such that it is feasible to analyse genomic data sets with thousands of samples. Although LSD assumes a strict molecular clock, its accuracy is frequently similar to that obtained using more sophisticated Bayesian clock models [17]. Other non-Bayesian molecular-clock methods have also been developed recently with the purpose of analysing large genomic data sets [18–20].

Finally, a suite of tools are available to infer phylodynamic parameters from a chronogram, such as that obtained using LSD. For example: TreePar uses maximum likelihood to fit birth-death and skyline models [21]; BEAST2 [4] and RevBayes [22] can fit a range of birth-death, coalescent, and Skyline models using Bayesian inference [7]; and approximate Bayesian computation (ABC) approaches that use tree summary statistics have recently been developed to fit phylogenetic epidemiological models [23,24]. The main disadvantage of these approaches over those that are fully Bayesian is that the estimates are based on a single tree, such that uncertainties in tree topology, branch lengths, and evolutionary rates are ignored. A potential solution is to repeat the analysis using non-parametric bootstrap replicates, but combining the different sources of uncertainty under this framework is not trivial.

Here, we compare the following two methods to infer evolutionary rates and timescales, and demographic parameters:

(i) The fully Bayesian method, implemented in BEAST2, to simultaneously infer the phylogenetic tree, evolutionary timescales and phylodynamic parameters;

(ii) The hybrid method: phylogram inference using maximum likelihood in PhyML v3.1 [14], chronogram inference using LSD v0.3, and estimation of phylodynamics parameters in BEAST2 using Bayesian inference.

To compare the performance of these two methods, we analysed previously published whole genome SNP bacterial data sets of *Mycobacterium tuberculosis* Lineage 2 [25], *Vibrio cholerae* [26], *Shigella dysenteriae* type 1 [11], and *Staphylococcus aureus* ST239 [27]. Because these data sets have small numbers of samples (n=63 for *M. tuberculosis*, n=122 for *V. cholerae*, n=121 for *S. dysenteriae*, and n=74 for *S. aureus*) their analyses are computationally tractable using both approaches. We also demonstrate the unique potential of the hybrid approach by analysing two genomic data sets with larger numbers of sequences, which have been difficult to analyse using a fully Bayesian approach; a global sample of *S. dysenteriae* type 1 (n=329) and *S. dysenteriae* type 1 lineage IV (n = 208) [11].

Results

Estimates of evolutionary rates and timescales

We compared estimates of rates and evolutionary timescales using the full Bayesian approach in BEAST2 and LSD. Because our data consist of SNPs, the rate estimates do not represent genome-wide rates, but SNP substitutions per site per year. In BEAST2 we used both the strict and the uncorrelated lognormal (UCLN [28]) clock models. We found that estimates from these different methods were largely congruent (Fig.1). In all four cases, the 95% credible intervals for the SNP evolutionary rate and age of the root node obtained with BEAST2 overlapped with the 95% confidence intervals obtained for the same parameters with LSD (Fig.1). However, we observed some differences in the genome-wide SNP mean substitution rate estimates, with the UCLN clock model in BEAST2 consistently producing the highest values, followed by the strict clock in BEAST2, and LSD producing the lowest estimate. The largest difference in mean SNP rate estimates was observed in *M. tuberculosis*, with a mean SNP rate of 1.48×10^{-4} (95% credible interval: 4.31×10^{-5} – 3.25×10^{-4}) using the UCLN, 1.31×10^{-4} (95% credible interval: 1.73×10^{-5} – 1.91×10^{-4}) under the strict clock in BEAST2, and 1.10×10^{-5} (95% confidence interval: 1.00×10^{-5} – 3.72×10^{-4}) in LSD (see Fig.1). In contrast we found more congruent mean SNP rate estimates in the *V. cholerae* data set, with estimates of 1.43×10^{-3} (95% credible interval: 1.11×10^{-3} – 1.76×10^{-3}) for the UCLN, 1.28×10^{-3} (95% credible interval: 1.17×10^{-3} – 1.39×10^{-3}) for the strict clock in BEAST2, and 1.16×10^{-3} (95% confidence interval: 9.82×10^{-4} – 1.35×10^{-3}) for LSD. The differences in estimates of the root-node age were similar (Fig.1). In all cases, the posterior distribution under the UCLN was wider than that under the strict clock, which is expected because the UCLN has a larger number of parameters [28].

We investigated two aspects of phylogenetic data that can affect estimates of evolutionary rates; the topological uncertainty and the degree of clocklike variation. We found that the maximum likelihood trees were highly supported, according to local likelihood ratio tests (aLRT) [29] (which ranges from 0 to 1, for low to high branch support, respectively). The median aLRT values across nodes were 0.9 for *M. tuberculosis*, 0.83 for *V. cholerae*, 0.99 for *S. dysenteriae* type 1, and 0.92 for *S. aureus*. Next we investigated the degree of rate variation among lineages by inspecting the coefficient of rate variation, estimated in the UCLN model. This parameter is the standard deviation of branch rates divided by the mean rate. The data are considered to display clocklike behaviour if the distribution for this parameter abuts zero. The *M. tuberculosis* data set, which displayed the largest differences in estimates among the methods, was the only data set to support a strict clock over the UCLN model, whereas the remaining data sets favoured the UCLN model (Fig.1).

Assessing temporal structure using a date-randomisation test

We assessed the reliability of our estimates of evolutionary rate and timescales by conducting a date-randomisation test [30,31]. This test entails repeating the analysis several times after randomising the sampling dates. The resulting rate estimates correspond to the expected values if there is no association

between sampling times and genetic divergence. The data are considered to have strong temporal structure if the rate estimate obtained using the correct sampling times is not contained within the range of values from the randomisations. In a Bayesian context, 10 to 20 randomisations appear to be sufficient [31,32]. We conducted this test in BEAST2 using 20 randomisations and in LSD using 100 randomisations (Fig.2). Interestingly, the results from both tests were congruent, and consistent with visualisations of clock-like behaviour of the data using root-to-tip regressions (Fig.S1). The *M. tuberculosis* data set had no temporal structure with either method (Fig.2): the credible interval of the Bayesian estimate with the correct sampling times overlapped with those from all of the randomisations; using LSD, the estimate with the correct sampling times was around the lower threshold in the program, at 1.00×10^{-5} subs/site/year, which also corresponds to the value obtained for most of the randomisations. The other data sets showed strong temporal structure with both date-randomisation tests: the Bayesian credible intervals using the correct sampling times did not overlap with those from any of the randomisations, and the estimates from LSD using the correct sampling times were not contained within the distributions of the 100 randomisations (Fig.2).

Inference of phylodynamic parameters

We analysed the data sets using the exponential-growth coalescent model in BEAST2, which has two parameters, the growth rate (r) and the scaled population size (Φ). In the context of pathogen evolution, r determines the speed of spread of the pathogen in the host population, while Φ is proportional to the infected host population size at present. Because these are compound parameters, they cannot be interpreted in an absolute scale without additional information about the size of the infected host population at present [33]. In most cases, the posterior distributions of both parameters were very similar when using either the two clock models in BEAST2 or the hybrid approach, with similar means and uncertainties (Fig.3). Although the three intervals overlapped in *V. cholerae*, *S. dysenteriae*, and *S. aureus*, the mode of the posterior distribution of Φ was highest when using the hybrid approach, followed by that from the strict and the UCLN clock models. The posterior distributions of r were almost identical across methods for the three data sets with temporal signal (Fig.3). The uncertainty in estimates of this parameter did not include 0, except in the case of *V. cholerae*, suggesting that most of these bacterial data sets were undergoing population growth. Interestingly, the *M. tuberculosis* data set, which had no temporal structure, was the only data set to display large differences in estimates among the methods (Fig.3).

Application: analysing large data sets using the hybrid approach

Having demonstrated good performance of the hybrid approach on small data sets with strong temporal signal, we applied it to analyse two published genome-wide SNP data sets whose sample size was prohibitively large to analyse under a full Bayesian framework in the original publication. These data sets consisted of: (i) 329 samples of *S. dysenteriae* type 1 from [11], which included BEAST analysis of a subset of 125 samples; and (ii) 208 samples of lineage IV of *S. dysenteriae* type 1, which was represented by 61 samples in the BEAST analysis in the same study [11]. These three data sets displayed strong temporal

structure according to the date-randomisation test in LSD, with rate estimates that were not contained within the range of estimates from 100 date-randomisations (Fig.4). The SNP rate estimates from LSD were 1.79×10^{-4} (95% confidence interval: 1.59×10^{-4} - 2.19×10^{-4}) SNP subs/site/year for *S. dysenteriae* type 1, and 1.18×10^{-4} (95% confidence interval: 6.32×10^{-5} - 1.65×10^{-4}) SNP subs/site/year for *S. dysenteriae* type 1 Lineage IV (Fig.4). Interestingly, the estimate of *r* for *S. dysenteriae* type 1 lineage IV was nearly an order of magnitude higher than that for the global data set of this bacterium, with a mean of 2.00×10^{-2} for lineage IV compared with 3.4×10^{-3} for the global data set. Importantly, the posterior distributions of *r* for these three data sets did not include zero, indicating epidemic growth (Fig.4).

Computational demands of the Bayesian and the hybrid methods

The hybrid approach was several times faster than the full Bayesian approach. For example, the computation time for each randomisation of the *V. cholerae* data set each was about 2 hours using BEAST2, where as those in LSD took 1.23 seconds (sec). However, a key aspect of the date-randomisation test in LSD is that the tree topology and branch lengths are fixed for all randomisations, where as they are re-estimated for each randomisation in BEAST2. For the *V. cholerae* data set, a complete analysis using the hybrid approach took: 10.06 minutes (min) to infer a maximum likelihood tree in PhyML, 1.23 sec to estimate the evolutionary rate and timescale in LSD, and 5 min to infer *r* and Φ in BEAST2 to obtain effective sample sizes (ESS) of over 200 for all parameters (drawing 1×10^{-7} steps, with 1 minutes per 10^{-6} steps), for a total of about 15 min, and 1/12th of the time required in BEAST2. Analysis of the full *S. dysenteriae* dataset from [11], the largest data set in our study, took 10.6 sec to analyse in LSD and 1 hour infer *r* and Φ BEAST2 (drawing 5×10^{-7} steps, with 1.2 minutes per 10^{-6} steps), for the 329 samples.

Discussion

Our results demonstrate that, as long as a strong temporal signal is present, the hybrid and fully Bayesian method can produce congruent estimates of evolutionary parameters, even in cases where the data display substantial rate variation among lineages. All methods also yielded similar estimates of demographic parameters in data sets with strong temporal signal, indicating the hybrid approach is a reliable alternative to full Bayesian analyses. However, *r* appears to be more robust than Φ to mild differences in estimates of the rate and timescale. This probably occurs because the age of the root-node plays an important role in the population size under the coalescent. In particular, the effective population size, and therefore Φ, are known to scale positively with the age of the root-node [34].

Obtaining congruent estimates between the two methods depends on whether the data meet certain criteria. In practice, it is important to verify that the trees have high branch support and that the data have strong temporal structure. The trees inferred here were highly supported, but it is likely that the hybrid approach will produce misleadingly precise estimates (i.e. with narrow confidence intervals) if branch support is low, because the demographic parameters will still be conditioned on a single, and possibly incorrect, tree obtained in step 1 that does not capture uncertainty in the topology. In contrast, in such

circumstances the Bayesian method will simply integrate over phylogenetic uncertainty and yield wider credible intervals.

We found strong temporal structure to be an important requirement for obtaining congruent estimates between the two methods. Our results show that the date-randomisation test in LSD appears to be as effective as it is in BEAST2, with the advantage of being much less computationally demanding. As a result, it is possible to use a larger number of replicates, which can improve the power of the test. Moreover, the sampling times under a Bayesian analysis of sequentially sampled data are informative about the tree topology. That is, they impose a high prior probability on trees that cluster sequences with similar sampling times, which can render the date-randomisation test unreliable, with an increase in type I error [35]. The date-randomisation test in LSD does not suffer from this problem because sequence data alone, not tip dates, are used to infer the tree topology in maximum likelihood.

Critically, the rates estimated using the date-randomisation in test in LSD are not necessarily unimodal in their distribution. This occurs because a lack of temporal structure usually leads to very low rate estimates, which affects randomisations in LSD and in BEAST2. In the case of LSD, very low values for the rate will correspond to the lower threshold set in the program [17], which we arbitrarily set at 10^{-5} and 10^{-10} subs/site/year, such that most randomisations will have this value. As such, a reasonable approach to interpret the date-randomisation test in LSD is to ensure that the rate estimate with the correct sampling times is higher than those from at least 95% of the randomisations, following the frequentist *p*-value of $\alpha=0.05$.

Conclusions

As shown here, hybrid methods offer an attractive alternative to full Bayesian approaches for genome-scale data sets with very large numbers of samples. The accuracy and precision of both methods are comparable, but hybrid methods can perform an analysis in a about $1/8^{th}$ of the time required for full Bayesian analyses. Nevertheless, some steps of the hybrid method used here require oversimplifications of the evolutionary process. For example, LSD always assumes a strict molecular clock, such that it is impossible to assess among-lineage rate variation or to pinpoint potential biological causes for why lineages have different rates. The choice of whether to use a hybrid method should be made based on what parameters a user wishes to interrogate. In the context of molecular epidemiology, demographic parameters (r and Φ) and divergence time information are of primary interest, all of which appear robust to among lineage rate variation.

In this study, we used a simple demographic model, the exponential-growth coalescent. This model appears to be well suited when outbreak data are sampled at an early stage, but it makes several assumptions, including that the population of susceptible hosts is constant and that there is no population structure [6]. A better understanding of the data used here requires more sophisticated phylodynamic models, such as those that include changes in diversification parameters over time [36], and migration [37].

To this end, our results suggest that harnessing the power of such models and large-scale genome sequencing can be done through hybrid approaches.

Materials and Methods

Data collection

Our bacterial data sets consisted of publically available genome data. We obtained all of our genome-wide SNP alignments from a previous studies [11,25,27,32]. These data sets are freely available online (github.com/sebastianduchene/bacteria_genomic_rates_data).

Phylogenetic analyses under the fully Bayesian approach

We analysed the sequence alignments in BEAST v2.4 using the sampling times for calibration, the GTR+Γ substitution model, the exponential-growth coalescent tree prior, and two clock models; the strict and the UCLN. We used the default priors for all parameters. Our Markov chain Monte Carlo (MCMC) sampling scheme consisted of a chain length of 5×10^{-8} steps, sampling every 10^{-4} steps. We verified that the ESS for all parameters was at least 200. To determine whether the data had temporal structure, we conducted a date-randomisation test by randomising the sampling dates 20 times and repeating the analyses [31].

Phylogenetic analyses using the hybrid approach

We inferred phylogenetic trees using maximum likelihood in PhyML v3.1. We used the GTR+Γ substitution model, and a search strategy that combines the nearest-neighbor interchange and subtree prune and regraft algorithms. To assess branch support, we calculated the aLRT score for each branch. To visually assess temporal structure, we conducted a regression of the root-to-tip distances as a function of the sampling times using TempEst v1.5 [38]. To determine the optimal root in this program we selected the position that maximised R^2.

We analysed the maximum likelihood trees (i.e. phylograms) in LSD v0.3 to infer the SNP evolutionary rate and timescale. We set the sampling times as calibrations, and allowed the program to determine the optimal position of the root. We constrained the branching times of the estimated chronograms such that daughter nodes must be younger than their parent nodes. To obtain an uncertainty around estimates of times and rates, we conducted 100 parametric bootstrap replicates of the branch lengths, as implemented in the program. Therefore, the uncertainty corresponds to the 95% confidence interval of the parametric bootstrap values. We conducted a date-randomisation test 100 times by randomising the sampling times in the ‘date’ file and running LSD each time. In this version of the test, the phylogenetic tree topology and branch lengths are fixed.

We used the chronograms estimated in LSD to infer demographic parameters in BEAST2. This consists in setting the input file to calculate the posterior as the likelihood of the tree given the model parameters multiplied by the priors on the parameters. In the exponential growth coalescent there are two parameters; Φ and r. We used an MCMC chain length of 1×10^{-7} sampling every 10^{-4} steps, and we verified that all parameters had ESS values of at least 200.

Abbreviations

LSD, Least-squares dating; ABC, Approximate Bayesian Computation; UCLN, uncorrelated lognormal clock; aLRT, local Likelihood ratio test for branch support; MCMC, Markov chain Monte Carlo.

Declarations

Acknowledgements

SD was supported by a McKenzie fellowship from the University of Melbourne. ZAD is funded by strategic award #106158 from the Wellcome Trust of Great Britain. KEH is supported by fellowship #1061409 from the NHMRC of Australia.

Author contributions

SD, DD, JLG, and KEH conceived and designed the experiments. SD, JLG, JH and ZD analysed the data. SD wrote the manuscript with input from all the authors.

Availability of supporting data

The data set analysed in this paper are available online at:
github.com/sebastianduchene/bacteria_genomic_rates_data

Competing interest

The authors declare no competing interests.

Ethics and consent to participate

Not applicable.

Availability of data and materials

All the software used in this study is freely available and open source. The data are all available online
github.com/sebastianduchene/bacteria_genomic_rates_data

366 **References**

367

368 1. Gire SK, Goba A, Andersen KG, Sealfon RSG, Park DJ, Kanneh L, et al. Genomic surveillance elucidates
369 Ebola virus origin and transmission during the 2014 outbreak. Science (80-.). American Association for the
370 Advancement of Science; 2014;345:1369–72.

371 2. Holmes EC, Dudas G, Rambaut A, Andersen KG. The evolution of Ebola virus: Insights from the 2013–
372 2016 epidemic. Nature. Nature Research; 2016;538:193–200.

373 3. Drummond AJ, Suchard MA, Xie D, Rambaut A. Bayesian phylogenetics with BEAUti and the BEAST 1.7.
374 Mol. Biol. Evol. 2012;29:1969–73.

375 4. Bouckaert R, Heled J, Kühnert D, Vaughan T, Wu C-H, Xie D, et al. BEAST 2: a software platform for
376 Bayesian evolutionary analysis. PLOS Comput. Biol. Public Library of Science; 2014;10:e1003537.

377 5. Ho SYW, Duchêne S. Molecular-clock methods for estimating evolutionary rates and time scales. Mol.
378 Ecol. 2014;23:5947–75.

379 6. Volz EM, Koelle K, Bedford T. Viral phylodynamics. PLOS Comput. Biol. Public Library of Science;
380 2013;9:e1002947.

381 7. du Plessis L, Stadler T. Getting to the root of epidemic spread with phylodynamic analysis of genomic
382 data. Trends Microbiol. Elsevier; 2015;23:383–6.

383 8. Rieux A, Balloux F. Inferences from tip-calibrated phylogenies: a review and a practical guide. Mol. Ecol.
384 2016;25:1911–24.

385 9. Stadler T, Kouyos R, von Wyl V, Yerly S, Böni J, Bürgisser P, et al. Estimating the basic reproductive
386 number from viral sequence data. Mol. Biol. Evol. SMBE; 2012;29:347–57.

387 10. Wong VK, Baker S, Pickard DJ, Parkhill J, Page AJ, Feasey NA, et al. Phylogeographical analysis of the
388 dominant multidrug-resistant H58 clade of Salmonella Typhi identifies inter-and intracontinental
389 transmission events. Nat. Genet. Nature Research; 2015;47:632–9.

390 11. Njamkepo E, Fawal N, Tran-Dien A, Hawkey J, Strockbine N, Jenkins C, et al. Global phylogeography
391 and evolutionary history of Shigella dysenteriae type 1. Nat. Microbiol. 2016;1:16027.

392 12. Stamatakis A. RAxML version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies.
393 Bioinformatics. Oxford Univ Press; 2014;30:1312–3.

394 13. Nguyen L-T, Schmidt HA, von Haeseler A, Minh BQ. IQ-TREE: a fast and effective stochastic algorithm
395 for estimating maximum-likelihood phylogenies. Mol. Biol. Evol. SMBE; 2015;32:268–74.

396 14. Guindon S, Dufayard J-F, Lefort V, Anisimova M, Hordijk W, Gascuel O. New algorithms and methods
397 to estimate maximum likelihood phylogenies: assessing the performance of PhyML 3.0. Syst. Biol.
398 2010;59:307–21.

399 15. Price MN, Dehal PS, Arkin AP. FastTree 2–approximately maximum-likelihood trees for large
400 alignments. PLoS One. Public Library of Science; 2010;5:e9490.

401 16. To T-H, Jung M, Lycett S, Gascuel O. Fast dating using least-squares criteria and algorithms. Syst. Biol.
402 Oxford University Press; 2016;65:82–97.

403 17. Duchêne S, Geoghegan JL, Holmes EC, Ho SYW. Estimating evolutionary rates using time-structured
404 data: a general comparison of phylogenetic methods. Bioinformatics. 2016;

405 18. Kumar S, Hedges SB. Advances in time estimation methods for molecular data. Mol. Biol. Evol. 2016;

406 19. Volz EM, Frost SDW. Scalable relaxed clock phylogenetic dating. Virus Evol. Oxford University Press;
407 2017;3.

408 20. Sagulenko P, Puller V, Neher R. TreeTime: maximum likelihood phylodynamic analysis. bioRxiv. Cold
409 Spring Harbor Laboratory; 2017;153494.

410 21. Stadler T. Mammalian phylogeny reveals recent diversification rate shifts. Proc. Natl. Acad. Sci. National
411 Acad Sciences; 2011;108:6187–92.

412 22. Höhna S, Landis MJ, Heath TA, Boussau B, Lartillot N, Moore BR, et al. RevBayes: Bayesian
413 phylogenetic inference using graphical models and an interactive model-specification language. Syst. Biol.
414 Oxford University Press; 2016;65:726–36.

415 23. Poon AFY. Phylodynamic inference with kernel ABC and its application to HIV epidemiology. Mol. Biol.
416 Evol. SMBE; 2015;32:2483–95.

417 24. Saulnier E, Alizon S, Gascuel O. Assessing the accuracy of Approximate Bayesian Computation
418 approaches to infer epidemiological parameters from phylogenies. PLOS Comput. Biol. Cold Spring Harbor
419 Labs Journals; 2017;13:e1005416.

420 25. Eldholm V, Monteserin J, Rieux A, Lopez B, Sobkowiak B, Ritacco V, et al. Four decades of
421 transmission of a multidrug-resistant Mycobacterium tuberculosis outbreak strain. Nat. Commun.
422 2015;6:7119.

423 26. Devault AM, Golding GB, Waglechner N, Enk JM, Kuch M, Tien JH, et al. Second-pandemic strain of
424 Vibrio cholerae from the Philadelphia cholera outbreak of 1849. N. Engl. J. Med. 2014;370:334–40.

425 27. Baines SL, Holt KE, Schultz MB, Seemann T, Howden BO, Jensen SO, et al. Convergent adaptation in
426 the dominant global hospital clone ST239 of methicillin-Resistant Staphylococcus aureus. MBio.
427 2015;6:e00080-15.

428 28. Drummond AJ, Ho SYW, Phillips MJ, Rambaut A. Relaxed phylogenetics and dating with confidence.
429 PLOS Biol. 2006;4:699–710.

430 29. Anisimova M, Gascuel O. Approximate likelihood-ratio test for branches: A fast, accurate, and powerful
431 alternative. Syst. Biol. Oxford University Press; 2006;55:539–52.

432 30. Ramsden C, Holmes EC, Charleston MA. Hantavirus evolution in relation to its rodent and insectivore
433 hosts: no evidence for codivergence. Mol. Biol. Evol. 2009;26:143–53.

434 31. Duchêne S, Duchêne DA, Holmes EC, Ho SYW. The performance of the date-randomization test in
435 phylogenetic analyses of time-structured virus data. Mol. Biol. Evol. 2015;32:1895–906.

32. Duchêne S, Holt KE, Weill F-X, Le Hello S, Hawkey J, Edwards DJ, et al. Genome-scale rates of evolutionary change in bacteria. Microb. Genomics. Microbiology Society; 2016;2.

33. Boskova V, Bonhoeffer S, Stadler T. Inference of epidemiological dynamics based on simulated phylogenies using birth-death and coalescent models. PLOS Comput Biol. Public Library of Science; 2014;10:e1003913.

34. Rosenberg NA, Nordborg M. Genealogical trees, coalescent theory and the analysis of genetic polymorphisms. Nat. Rev. Genet. Nature Publishing Group; 2002;3:380.

35. Murray GGR, Wang F, Harrison EM, Paterson GK, Mather AE, Harris SR, et al. The effect of genetic structure on molecular dating and tests for temporal signal. Methods Ecol. Evol. Wiley Online Library; 2015;7:80–9.

36. Stadler T, Kühnert D, Bonhoeffer S, Drummond AJ. Birth–death skyline plot reveals temporal changes of epidemic spread in HIV and hepatitis C virus (HCV). Proc. Natl. Acad. Sci. National Acad Sciences; 2013;110:228–33.

37. Kühnert D, Stadler T, Vaughan TG, Drummond AJ. Phylodynamics with migration: A computational framework to quantify population structure from genomic data. Mol. Biol. Evol. Society for Molecular Biology and Evolution; 2016;33:2102–16.

38. Rambaut A, Lam TT, Carvalho LM, Pybus OG. Exploring the temporal structure of heterochronous sequences using TempEst (formerly Path-O-Gen). Virus Evol. The Oxford University Press; 2016;2:vew007.

Figure legends

Figure 1. Estimates of evolutionary rate, time to the most recent common ancestor, and the coefficient of rate variation of the UCLN. The distributions in blue correspond to the posterior under the strict molecular clock model in BEAST2, while those in yellow are for the UCLN. The red solid line is the estimate from LSD, and the dashed lines correspond to the 95% confidence interval. The grey histograms show the posterior of the coefficient of rate variation under the UCLN, a measure of clocklike behaviour.

Figure 2. Date randomisation test using LSD and BEAST2. The left column shows histograms of the rate estimates with randomised sampling times in LSD (grey). The red line corresponds to the estimate using the correct sampling times. The right column shows the date randomisation test in BEAST2. The grey bars denote the 95% credible intervals of substitution rate estimates from the randomisations. The black lines correspond to the 95% credible interval of the rate estimates using the correct sampling times. The circles correspond to the mean value. The *x*-axis in the left column and the *y*-axis in the right column are in logarithmic scale.

Figure 3. Posterior estimates of demographic parameters, *Φ and r* using the full Bayesian and hybrid approaches. The red histograms correspond to the estimates from the hybrid approach, where the likelihood is calculated on a fixed tree. The yellow and blue histograms correspond to the posterior estimates using the UCLN and strict molecular clock models in BEAST2.

Figure 4. Date randomisation test in LSD and estimates of demographic parameters for large data sets using the hybrid approach. The grey histograms correspond to rate estimates from the randomisations, while the red lines correspond to the estimates using the correct sampling times. The red histograms correspond to the posterior distribution of parameters *Φ* and *r*.

Supplementary material legends

Fig.S1. Root-to-tip regression for all data sets. The blue points correspond to tips in the tree. The black line represents the linear regression of root-to-tip distance as a function of the sampling time. The root-to-tip distance is measured by fitting the root of the tree that maximises R^2.

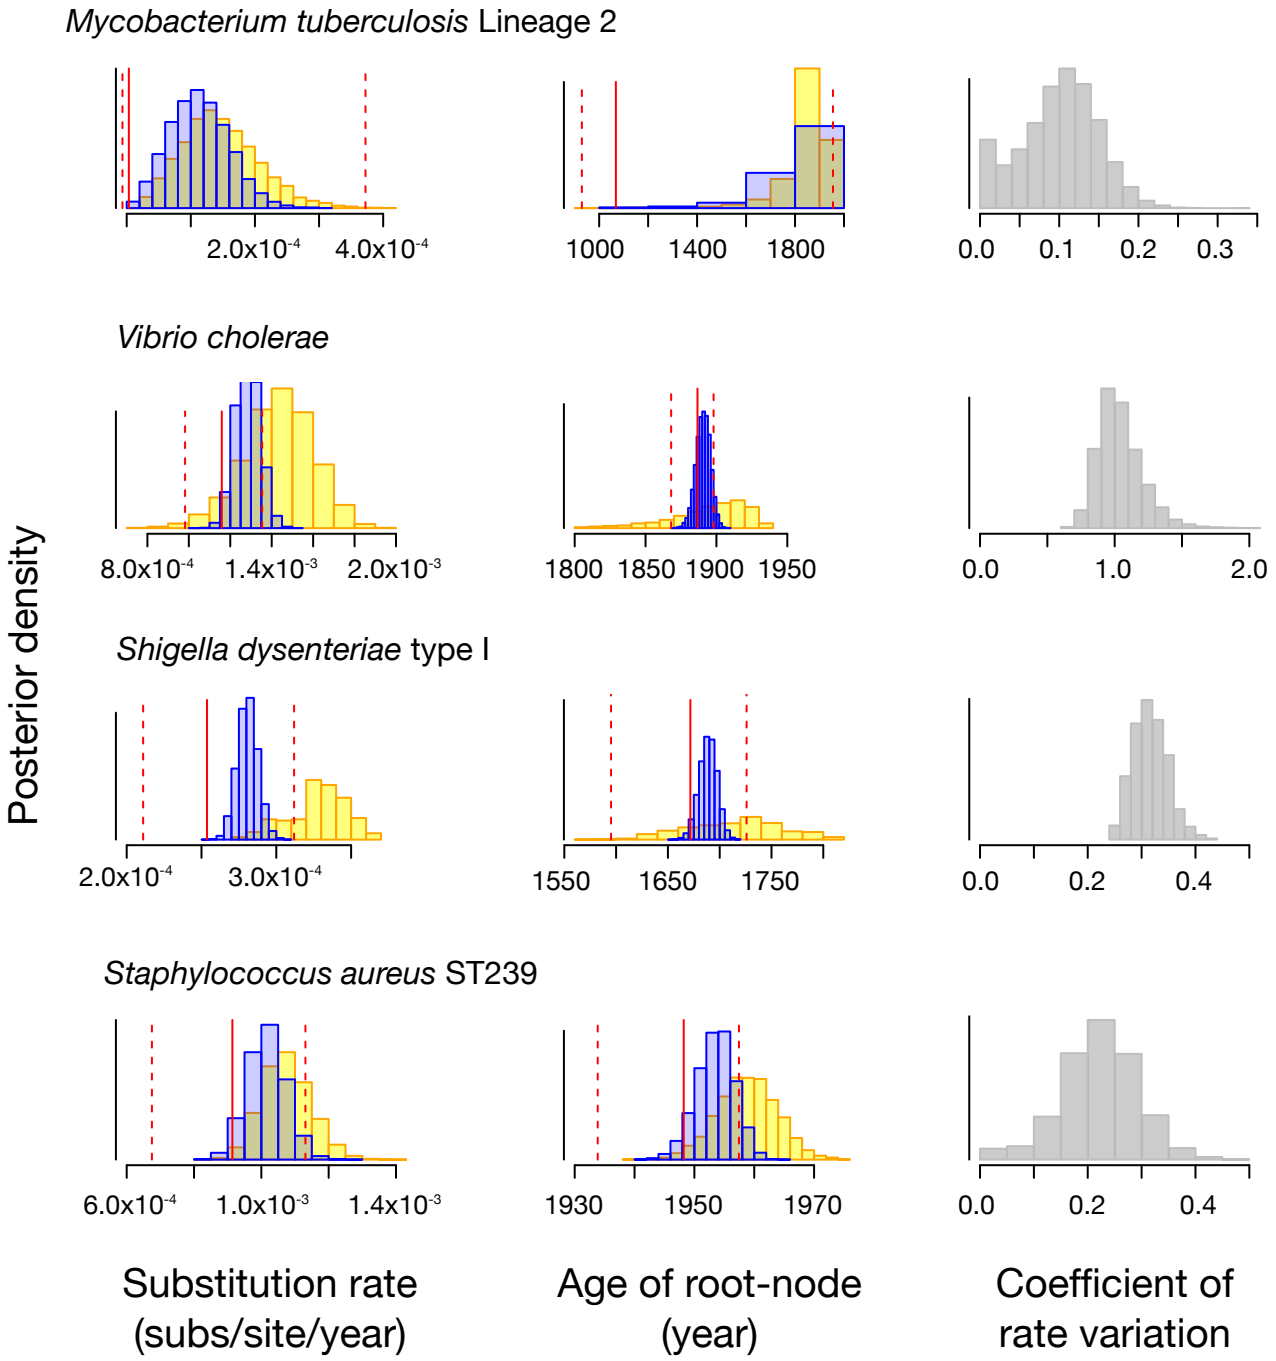
Mycobacterium tuberculosis Lineage 2
2.0x10-4
4.0x10-4
1000
1400
1800
0.0
0.1
0.2
0.3
Vibrio cholerae
8.0x10-4
1.4x10-3
2.0x10-3
1800
1850
1900
1950
0.0
1.0
2.0
Shigella dysenteriae type I
2.0x10-4
3.0x10-4
1550
1650
1750
0.0
0.2
0.4
Staphylococcus aureus ST239
6.0x10-4
1.0x10-3
1.4x10-3
1930
1950
1970
0.0
0.2
0.4
Posterior density
Substitution rate (subs/site/year)
Age of root-node (year)
Coefficient of rate variation

Mycobacterium tuberculosis Lineage 2

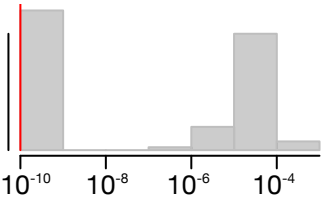

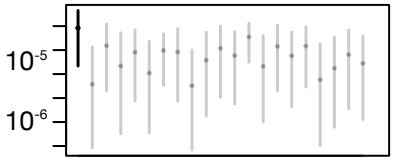

Vibrio cholerae

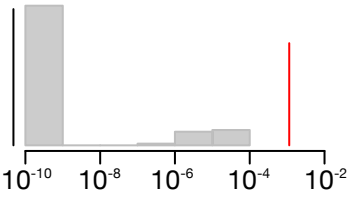

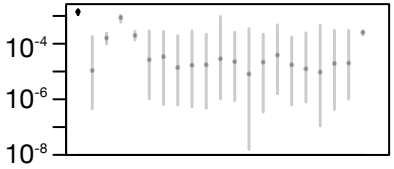

Frequency

Substitution rate (subs/site/year)

Shigella dysenteriae type I

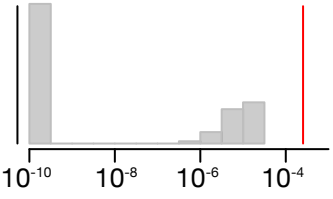

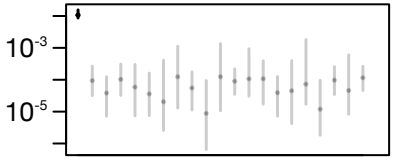

Staphylococcus aureus ST239

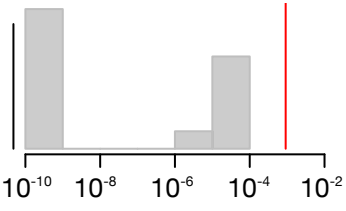

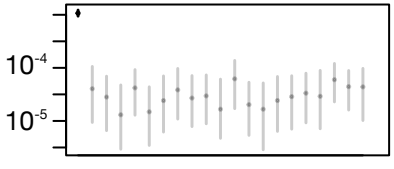

Substitution rate (subs/site/year)

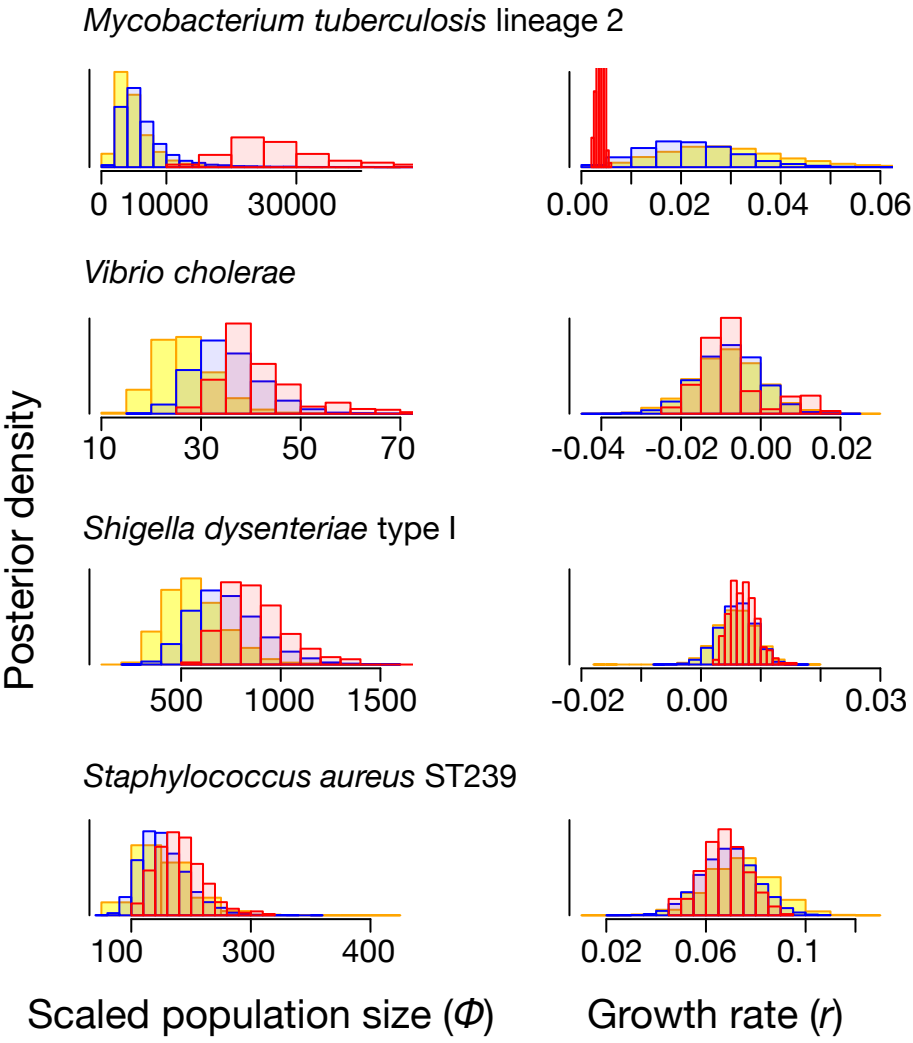
Mycobacterium tuberculosis lineage 2
0
10000
30000
0.00
0.02
0.04
0.06
Vibrio cholerae
10
30
50
70
-0.04
-0.02
0.00
0.02
Shigella dysenteriae type I
500
1000
1500
-0.02
0.00
0.03
Staphylococcus aureus ST239
100
300
400
0.02
0.06
0.1
Posterior density
Scaled population size (Φ)
Growth rate (r)

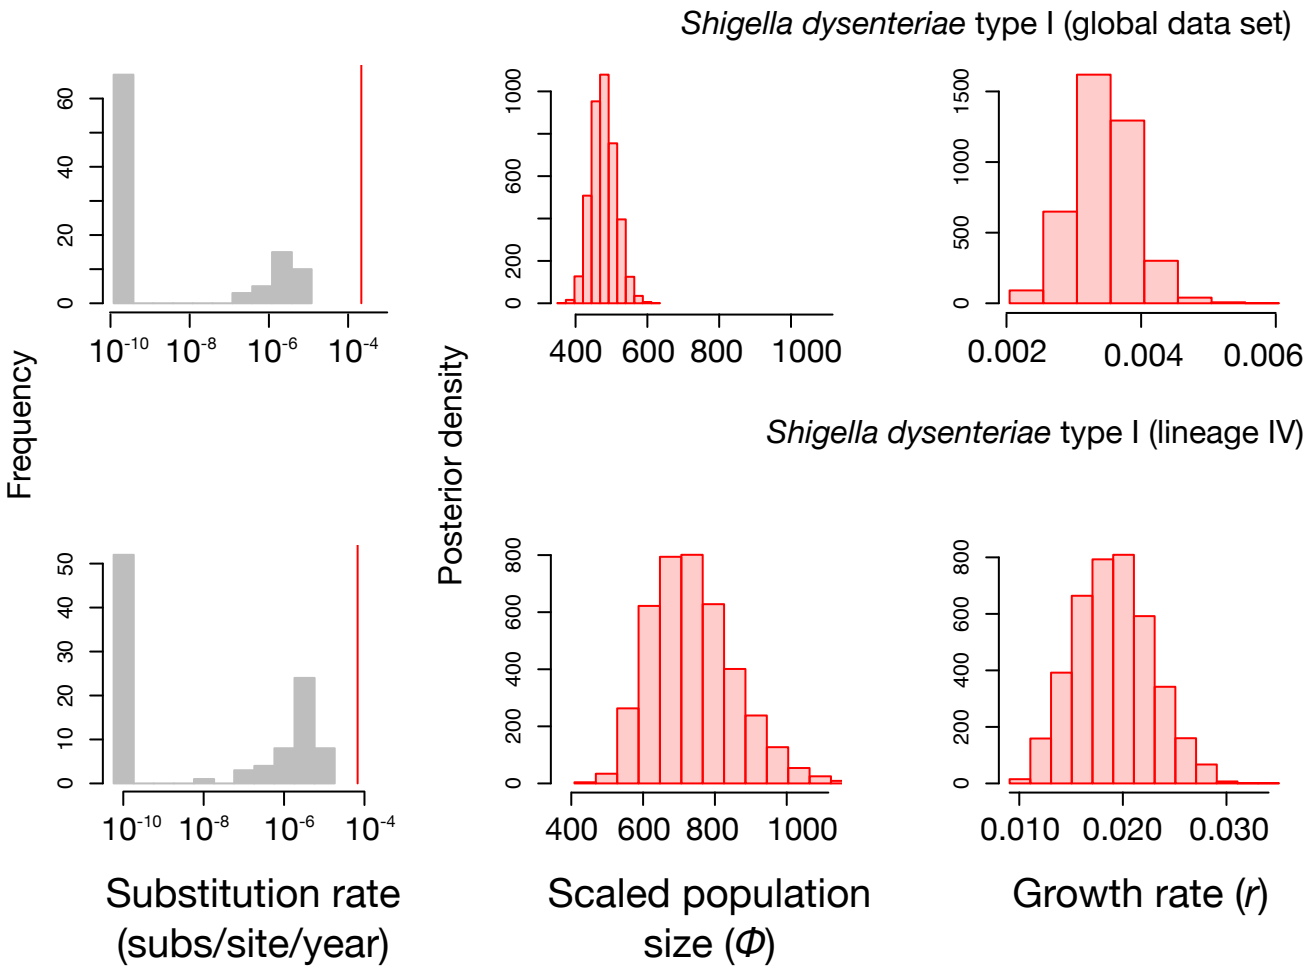
Shigella dysenteriae type I (global data set)
Shigella dysenteriae type I (lineage IV)
Frequency
Posterior density
Substitution rate (subs/site/year)
Scaled population size (Φ)
Growth rate (r)