An Annealed Sequential Monte Carlo Method for Bayesian Phylogenetics

LIANGLIANG WANG^{1,*}, SHIJIA WANG^{1,*}, AND ALEXANDRE BOUCHARD-CÔTÉ^{2,*}

¹Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby, British Columbia V5A 1S6, Canada and ²Department of Statistics,

University of British Columbia, Vancouver, British Columbia V6T 1Z4, Canada

*Correspondence to be sent to: Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby, British Columbia V5A 1S6, Canada; E-mail: lwa68@sfu.ca; shijiaw@sfu.ca; bouchard@stat.ubc.ca.

> Received 1 June 2018; reviews returned 12 April 2019; accepted 20 April 2019 Associate Editor: Edward Susko

Abstract.—We describe an "embarrassingly parallel" method for Bayesian phylogenetic inference, annealed Sequential Monte Carlo (SMC), based on recent advances in the SMC literature such as adaptive determination of annealing parameters. The algorithm provides an approximate posterior distribution over trees and evolutionary parameters as well as an unbiased estimator for the marginal likelihood. This unbiasedness property can be used for the purpose of testing the correctness of posterior simulation software. We evaluate the performance of phylogenetic annealed SMC by reviewing and comparing with other computational Bayesian phylogenetics, our annealed methods can utilize standard Markov chain Monte Carlo (MCMC) tree moves and hence benefit from the large inventory of such moves available in the literature. Consequently, the annealed SMC method should be relatively easy to incorporate into existing phylogenetic software packages based on MCMC algorithms. We illustrate our method using simulation studies and real data analysis. [Marginal likelihood; phylogenetics; Sequential Monte Carlo.]

INTRODUCTION

The Bayesian paradigm is widely used in systematic biology, principally for the purpose of phylogenetic reconstruction as well as for evaluating the empirical support of evolutionary models (Chen et al. 2014). Both of these tasks, Bayesian phylogenetic reconstruction and model selection, involve an intractable sum over topologies as well as a high-dimensional integral over branch lengths and evolutionary parameters. Consequently, Markov chain Monte Carlo (MCMC) methods have been widely used in the past 20 years to approximate posterior distributions defined over the space of phylogenetic trees (Rannala and Yang 1996).

Despite their success, MCMC phylogenetic methods are still afflicted by two key limitations, hence motivating the need for alternative approximations method for posterior distributions over phylogenetic trees.

Firstly, MCMC methods do not readily take advantage of highly parallel computer architectures. This is problematic in the current context as progress in computational power mostly comes in the form of parallelism gains. Although there are techniques available to parallelize phylogenetic MCMC methods, they are generally not "embarrassingly parallel": for example, parallel Metropolis coupled MCMC (Altekar et al. 2004) may reach a point where the addition of cores actually reduces sampling efficiency (Atchadé et al. 2011).

A second challenge with MCMC-based phylogenetic approximations arises in the context of model selection. By comparing the marginal likelihood Z=p(y), where y denotes observed data, under different models, one can approach scientific questions under the Bayesian framework while naturally taking into account differences in model complexity. More specifically, the ratio $r=p_1(y)/p_2(y)$ of two marginal likelihoods based

on two evolutionary models, $p_1(\cdot), p_2(\cdot)$, can be used to assess the strength of evidence y provides for p_1 (when r > 1) or p_2 (when r < 1). The ratio r is called the Bayes factor (Jeffreys 1935; Lartillot et al. 2006; Oaks et al. 2018). In the context of phylogenetics, the Bayes factor assesses how much support a set of sequencing data provides for one evolutionary model against another one.

Several methods have been proposed to estimate marginal likelihoods based on MCMC methods (Newton and Raftery 1994; Gelman and Meng 1998; Friel and Pettitt 2008, inter alia), including work tailored to the phylogenetic context (Huelsenbeck et al. 2004; Lartillot et al. 2006; Fan et al. 2010; Xie et al. 2010). However these methods all have different drawbacks, see for example the aptly named review, "Nineteen dubious ways to compute the marginal likelihood of a phylogenetic tree topology" (Fourment et al. 2018b). Moreover, one limitation shared by all MCMC-based marginal likelihood estimators is that they are generally biased (in the technical sense of the term as used in computational statistics, reviewed in the theory section of the paper)unless one is able to initialize the MCMC chains to the exact stationary distribution, which in practice is not possible. We argue that in certain scenarios, it can be useful to have unbiased methods. One example we elaborate on is for the purpose of a new test to ascertain correctness of posterior simulation software. Another class of examples comes from the burgeoning field of pseudo-marginal methods (Andrieu and Roberts 2009).

Sequential Monte Carlo (SMC) methods (see Doucet and Johansen 2011 for an accessible introduction to SMC) provide a flexible framework to construct unbiased estimators and past work has shown they can be very efficient in a phylogenetic context (Teh et al. 2008; Görür and Teh 2009; Bouchard-Côté et al. 2012; Görür et al. 2012; Wang et al. 2015; Everitt et al. 2016; Dinh et al. 2017; Smith et al. 2017; Fourment et al. 2018a). One drawback caused by the high degree of flexibility that comes with SMC is that the phylogenetic SMC algorithms developed so far are non-trivial to adapt to existing MCMC-based phylogenetic frameworks. Here, we propose a different construction based on the seminal work of Del Moral et al. (2006), in turn based on annealed importance sampling (AIS) (Neal 2001), which yields an SMC method which is in a sense much closer to standard MCMC, while providing unbiased estimators of the marginal likelihood. The proposed method, which we call phylogenetic annealed SMC, can directly make use of any existing phylogenetic MCMC proposals, a rich literature covering many kinds of phylogenetic trees (Rannala and Yang 1996; Yang and Rannala 1997; Larget and Simon 1999; Mau et al. 1999; Li et al. 2000; Holder and Lewis 2003; Rannala and Yang 2003; Höhna et al. 2008; Lakner et al. 2008; Höhna and Drummond 2012). It is easy to incorporate the proposed annealed SMC into existing phylogenetic software packages that implement MCMC algorithms, such as RevBayes (Höhna et al. 2016) or BEAST (Drummond and Rambaut 2007). At the same time, our method can leverage state-of-the-art advances in the field of adaption of SMC algorithms, making the algorithm fully automated in most cases.

Our implementation of the proposed method is available at https://github.com/liangliangwangsfu/ annealedSMC. All our experimental setups and results are available at https://github.com/shijiaw/ AnnealingSimulation. The algorithms described here are also available in the Blang probabilistic programming language https://github.com/ UBC-Stat-ML/blangSDK, which supports a small but growing set of phylogenetic models.

LITERATURE REVIEW

There is a growing body of work on SMC-based Bayesian phylogenetic inference. Indeed, a powerful feature of the general SMC framework (Del Moral et al. 2006) is that the space on which the distributions π_r are defined is allowed to vary from one iteration to the next. All previous work on SMC methods for phylogenetics has exploited this feature for various purposes reviewed here.

In one direction, several "bottom up" approaches (Teh et al. 2008; Görür and Teh 2009; Bouchard-Côté et al. 2012; Görür et al. 2012; Wang et al. 2015) have been proposed to allow more efficient reuse of intermediate stages of the Felsenstein pruning recursions. For these methods, the intermediate distributions are defined over forests over the observed taxa, and hence their dimensionality increases with *r*. These methods are most effective in clock or nearly clock trees. For general trees, it is typically necessary to perform additional MCMC steps, which makes it harder to use in the context of estimation of marginal likelihoods.

In a related direction, Dinh et al. (2017) and Fourment et al. (2018a) use a sequence of targets where π_r is a tree

over the first *r* tips. This construction is especially useful in scenarios where taxonomic data come in an online fashion.

Another use case of SMC methods in phylogenetics arises from Bayesian analysis of intractable evolutionary models. For example, SMC has been used for Bayesian phylogenetic analysis based on infinite state-space evolutionary models (Hajiaghayi et al. 2014) or for joint inference of transmission networks (Smith et al. 2017).

Finally, a concurrent line of work (Everitt et al. 2016) has explored a combination of reversible jump methods with phylogenetic models.

One drawback of letting the dimensionality of π_r vary with r as all the above methods do, is that it makes it significantly harder to incorporate SMC into existing Bayesian phylogenetic inference packages such as MrBayes (Huelsenbeck and Ronquist 2001), RevBayes, or BEAST. In contrast, in our method the target distributions π_r are all defined over the same space. The annealed SMC framework in this context utilizes Metropolis-Hastings kernels in the inner loop but combines them in a different fashion compared with standard MCMC algorithms, or even compared with parallel tempering MCMC algorithms.

SETUP AND NOTATION

We let *t* denote a phylogenetic tree with tips labeled by a fixed set of operational taxonomic units *X*. The variable *t* encapsulates the tree topology and a set of positive branch lengths. Our methodology is directly applicable to any class of phylogenetic trees where MCMC proposal distributions are available. This includes for example clock trees (Höhna et al. 2008) as well as nonclock trees (Lakner et al. 2008).

We let θ denote evolutionary parameters, for example the parameters of a family of rate matrices such as the general time reversible (GTR) model (Tavaré 1986), or diffusion parameters in the case of continuous traits (Lemey et al. 2010). Again our method is applicable to any situation where MCMC proposals are available for exploring the space of θ . We use $x = (t, \theta)$ to denote these two latent variables.

We let *y* denote observed data indexed by the tips *X* of *t*. We assume a likelihood function p(y|x) is specified such that for any hypothesized tree and parameters, the value p(y|x) can be computed efficiently. This assumption is sometimes called pointwise computation. This is a typical assumption in Bayesian phylogenetics, where this computation is done with some version of Felsenstein pruning (Felsenstein 1973, 1981) (an instance of the Forward-Backward algorithm; Forney 1973).

Finally, let p(x) denote a prior on the parameters and trees, which we assume can also be computed pointwise efficiently. This defines a joint distribution, denoted $\gamma(x) = p(x)p(y|x)$. We ignore the argument *y* from now on since the data are viewed as fixed in a Bayesian analysis context.

2020

We are interested in approximating a posterior distribution on *x* given data *y*, denoted:

$$\pi(x) = \frac{\gamma(x)}{\int \gamma(x') dx'}.$$
(1)

Here the integral and dx' are viewed in an abstract sense and include both summation over discrete latent variables such as topologies and standard integration over continuous spaces.

The denominator can be interpreted as the marginal likelihood under the model specified by the prior and likelihood functions, which we denote by *Z*:

$$Z = p(y) = \int \gamma(x) dx.$$
 (2)

Computation of this quantity, also called the normalization constant or evidence, is the main challenge involved when doing Bayesian model selection.

Other quantities of interest include expectations with respect to the posterior distribution, characterized by a real-valued function of interest f based on which we would like to compute

$$\int \pi(x) f(x) \mathrm{d}x. \tag{3}$$

For example if we seek a posterior clade support for a subset $X' \subset X$ of the leaves X,

$$f(x) = f(t, \theta) = 1[t \text{ admits } X' \text{ as a clade}],$$

where 1[*s*] denotes the indicator function which is equal to one if the Boolean expression *s* is true and zero otherwise.

ANNEALED SMC FOR PHYLOGENETICS

Sequences of Distributions

In standard MCMC methods, we are interested in a single probability distribution, the posterior distribution. However, there are several reasons why we may use a *sequence* of distributions rather than only one.

A first possibility is that we may have an online problem, where the data are revealed sequentially and we want to perform inference sequentially in time based on the data available so far. The distribution at step ris then the posterior distribution conditioning on the first r batches of data. This approach is explored in the context of phylogenetics in Dinh et al. (2017), where a batch of data consists in genomic information for one additional operational taxonomic unit. We do not pursue this direction here but discuss some possibilities for combinations in the discussion.

A second reason for having multiple distributions, and the focus of this work, is to facilitate the exploration of the state space. This is achieved for example by raising the likelihood term to a power ϕ_r between zero and one, which we multiply with the prior

$$\gamma_r(x) = p(y|x)^{\varphi_r} p(x). \tag{4}$$

MCMC may get stuck in a region of the space of phylogenetic trees around the initial value. This may happen for example around a local maximum (mode) in the posterior density. Such a region is sometimes called a "basin of attraction," and no single basin of attraction may be enough to well represent the full posterior distribution. Introducing a series of powered posterior distributions can alleviate this issue. A small value of ϕ_r flattens the posterior and makes MCMC samplers move easily between the different basins of attractions. The samples are initially overly dispersed but are then coerced into the posterior distribution $\pi(x)$ by slowly increasing the annealing parameter ϕ_r .

We do not anneal the prior to ensure that $\gamma_r(x)$ has a finite normalization constant,

$$\int \gamma_r(x) dx = \mathbb{E}_{p(x)}[(p(y|X))^{\phi_r}]$$
$$\leq \left(\mathbb{E}_{p(x)}[p(y|X)]\right)^{\phi_r} = (p(y))^{\phi_r} < \infty,$$

where the first inequality follows from the concavity of $(\cdot)^{\phi_r}$ and Jensen's inequality.

A third scenario is that we may encounter a "tall data" problem, for example biological sequences with a large number of sites. When the number of sites is large, evaluation of the unnormalized posterior $\gamma_r(x)$ defined in Equation (4) is computationally expensive. The idea of data subsampling (Quiroz et al. 2018a, 2018b; Bardenet et al. 2017; Gunawan et al. 2018) could be used to define the sequence of distributions. The construction of the sequence of distributions is described in Appendix 1.

The probability distributions

$$\pi_r(x) = \frac{\gamma_r(x)}{\int \gamma_r(x') dx'}$$
(5)

are therefore well defined and we denote their respective normalization constants by

$$Z_r = \int \gamma_r(x) \mathrm{d}x. \tag{6}$$

If the exponent ϕ_r is zero, then the distribution π_r becomes the prior which is often easy to explore and in fact independent samples can be extracted in many situations. At the other extreme, the distribution at power $\phi_r = 1$ is the distribution of interest.

The intermediate distributions $\{\pi_r\}_{r=1,...,R}$ are defined on a common measurable space $(\mathcal{X}, \mathcal{E})$. The annealed SMC is a generalization of the standard SMC method (Doucet et al. 2001). In standard SMC, the intermediate distributions are defined on a space of strictly increasing dimension.

Basic Annealed SMC Algorithm

We now turn to the description of annealed SMC in the context of Bayesian phylogenetic inference. The algorithm fits into the generic framework of *SMC samplers* (Del Moral et al. 2006): at each iteration, indexed by r=1,2,...,R, we maintain a collection indexed by $k \in \{1,2,...,K\}$ of imputed latent states $x_{r,k}$, each paired with a non-negative number called a weight $w_{r,k}$; such a pair is called a particle. A latent state in our context consists in a hypothesized tree $t_{r,k}$ and a set of evolutionary parameters $\theta_{r,k}$, i.e. $x_{r,k} = (t_{r,k}, \theta_{r,k})$. In contrast to previous SMC methods, $x_{r,k}$ is always of the same data type: no partial states such as forest or trees over subsets of leaves are considered here.

A particle population consists in a list of particles $(x_{r,.}, w_{r,.}) = \{(x_{r,k}, w_{r,k}): k \in \{1, ..., K\}\}$. A particle population can be used to estimate posterior probabilities as follows: first, normalize the weights, denoted after normalization using capital letter, $W_{r,k} = w_{r,k} / \sum_{k'} w_{r,k'}$. Second, use the approximation:

$$\int \pi_r(x) f(x) \mathrm{d}x \approx \sum_{k=1}^K W_{r,k} f(x_{r,k}).$$
(7)

For example, if we seek a posterior clade support for a subset $X' \subset X$ of the leaves X, this becomes

$$\sum_{k=1}^{K} W_{r,k} 1 \text{[sampled tree } t_{r,k} \text{ admits } X' \text{ as a clade]}.$$

The above formula is most useful at the last SMC iteration, r = R, because π_R coincides with the posterior distribution by construction.

At the first iteration, each of the particles' tree and evolutionary parameters are sampled independently and identically from their prior distributions. We assume for simplicity that this prior sampling step is tractable, a reasonable assumption in many phylogenetic models. After initialization, we therefore have a particle-based approximation of the prior distribution. Intuitively, the goal behind the annealed SMC algorithm is to progressively transform this prior distribution approximation into a posterior distribution approximation.

To formalize this intuition, we use the sequence of distributions introduced in the previous section. The last ingredient required to construct an SMC algorithm is an SMC proposal distribution $K_r(x_{r-1,k}, x_{r,k})$, used to sample a particle for the next iteration given a particle from the previous iteration. Because $x_{r-1,k}$ and $x_{r,k}$ have the same dimensionality in our setup, it is tempting to use MCMC proposals $q_r(x_{r-1,k}, x_{r,k})$ in order to build SMC proposals, for example, subtree prune and regraft moves, and Gaussian proposals for the continuous parameters and branch lengths. Indeed, there are several advantages of using MCMC proposals as the basis of SMC proposals. First, this means, we can leverage a rich literature on the topic (Rannala and Yang 1996; Yang and Rannala 1997; Larget and Simon 1999; Mau et al.

1999; Li et al. 2000; Holder and Lewis 2003; Rannala and Yang 2003; Höhna et al. 2008; Lakner et al. 2008; Höhna and Drummond 2012). Second, it makes it easier to add SMC support to existing MCMC-based software libraries. Third, it makes certain benchmark comparison between SMC and MCMC more direct, as we can then choose the set of moves to be the same for both. On the flip side, constructing MCMC proposals is somewhat more constrained, so some of the flexibility provided by the general SMC framework is lost.

Naively, we could pick the SMC proposal directly from an MCMC proposal, $K_r(x_{r-1,k}, x_{r,k}) = q_r(x_{r-1,k}, x_{r,k})$. However, doing so would have the undesirable property that the magnitude of the fluctuation of the weights of the particles from one iteration to the next, $||W_{r-1}|$. $W_{r,.}$ does not converge to zero when the annealing parameter change $\phi_r - \phi_{r-1}$ goes to zero. This lack of convergence to zero can potentially cause severe particle degeneracy problems, forcing the use of a number of particles larger than what can be realistically accommodated in memory (although workarounds exist, e.g. Jun and Bouchard-Côté 2014). To avoid this issue, we follow Del Moral et al. (2006) and use as SMC proposal the accept-reject Metropolis-Hastings transition probability based on q_r (called a Metropolized proposal), reviewed in Algorithm 1.

Algorithm 1 Accept–reject Metropolis-Hastings algorithm

- Propose a new tree and/or new evolutionary parameters, x^{*}_r ~ q_r(x_{r-1}, ·). ▷
 For example, using a nearest neighbor interchange, and/or a symmetric normal proposal on branch lengths and/or evolutionary parameters.
- 2. Compute the Metropolis-Hastings ratio based on γ_r :

$$\alpha_r(x_{r-1}, x_r^*) = \min\left\{1, \frac{\gamma_r(x_r^*)q(x_r^*, x_{r-1})}{\gamma_r(x_{r-1})q(x_{r-1}, x_r^*)}\right\}$$

3. Simulate
$$u \sim U(0, 1)$$
.

4. if
$$u < \alpha_r(x_{r-1}, x_r^*)$$
 then

U

5. $x_r = x_r^*$. \triangleright Output the proposal x_r^* . 6. else

7. $x_r = x_{r-1}$. \triangleright Output the previous state x_{r-1} .

The key point is that a theoretical argument (reviewed in the Appendix 2) shows that provided that (1) K_r has stationary distribution π_r (which is true by construction, a consequence of using the Metropolis-Hastings algorithm) and (2) we use the weight formula:

$$w_{r,k} = \frac{\gamma_r}{\gamma_{r-1}} (x_{r-1,k}), \tag{8}$$

then we obtain a valid SMC algorithm, meaning that the key theoretical properties expected from SMC hold under regularity conditions, see Theoretical Properties section.

In the important special case where $\gamma_r(x_r)$ is equal to the prior times an annealed likelihood, we obtain

$$v_{r,k} = [p(y|x_{r-1,k})]^{\phi_r - \phi_{r-1}}.$$
(9)

As hoped, the update shown in Equation (9) has the property that weight fluctuations vanish as the annealing parameter difference $\phi_r - \phi_{r-1}$ goes to zero. This will form the basis of the annealing parameter sequence adaptation strategies described in the next section. But for now, assume for simplicity that the number of iterations *R* and the annealing schedule ϕ_r , $r \in \{1,...,R\}$ is prespecified. For example, a simple choice for the annealing parameter sequence (Friel and Pettitt 2008) is $\phi_r = (r/R)^3$, where *R* is the total number of SMC iterations. In this case, the difference between successive annealing parameters is $(3r^2 - 3r + 1)/R^3$. An annealed SMC with a larger value of *R* is computationally more expensive but has a better performance.

In contrast to other SMC algorithms, the annealed SMC algorithm does not require pointwise evaluation of the proposal $K_r(x_{r-1,k}, x_{r,k})$, that is, given $x_{r-1,k}$ and a sampled $x_{r,k}$, we do not need to compute the numerical value of $K_r(x_{r-1,k}, x_{r,k})$ as it does not appear in the weight update formula, Equation (8). This point is important, because for Metropolis-Hastings kernels, pointwise evaluation would require computation of a typically intractable integral under the proposal in order to compute the total probability of rejection. The theoretical justification as to why we do not need pointwise evaluation of K_r is detailed in Appendix 2.

In practice, many proposals are needed to modify different latent variables and to improve mixing. We give in Appendix 3 the list of MCMC proposals we consider. Let q_r^i , i=1,...,M, denote the various proposals, and K_r^i the corresponding Metropolized transition probabilities. We need to combine them into one proposal K_r . To ensure that condition (1) above is satisfied, namely that K_r obeys global balance with respect to π_r , use the following property (Tierney 1994; Andrieu et al. 2003): if each of the transition kernels $\{K^i\}, i=1,...,M$, respects global balance with respect to π , then the cycle hybrid kernel $\prod_{i=1}^{M} K^{i}$ and the *mixture hybrid kernel* $\sum_{i=1}^{M} p_{i} K^{i}$, $\sum_{i=1}^{M} p_{i} = 1$, also satisfy global balance with respect to π . The global balance condition, $\int \pi_r(x) K_r(x, x') dx = \pi_r(x')$, ensures that the Markov chain encoded by K_r admits π_r as a stationary distribution. In practice, the mixture kernel is implemented by randomly selecting K^{i} with probability p_i at each iteration (Andrieu et al. 2003).

We can now introduce in Algorithm 2 the simplest version of the annealed SMC, which alternates between reweighting, propagating, and resampling. Figure 1 presents an overview of the annealed SMC algorithmic framework. In the proposal step, we propose new particles through MCMC moves (typically Metropolis-Hastings moves). Finally, we use resampling to prune particles with small weights. A list of unweighted particles is obtained after the resampling step.

In the annealed SMC algorithm, note that the weighting and proposal steps can be interchanged. This is different from standard SMC algorithms, where in general the proposal has to be computed before weighting. This interchange is possible because in the annealed SMC algorithm, the weighting function,

Equation (8), only depends on particles from the previous iteration and not from those just proposed as in standard SMC algorithms. This flexibility will come handy when designing adaptive schemes.

Algorithm 2 The simplest version of annealed SMC algorithm (for pedagogy)

- 1. Inputs:
- 2. (a) Prior over evolutionary parameters and trees, p(x), where $x = (\theta, t)$;
- 3. (b) Likelihood function p(y|x);
- 4. (c) Sequence of annealing parameters $0 = \phi_0 < \phi_1 < \cdots < \phi_R = 1$.
- 5. **Outputs:** Approximation of the posterior distribution, $\sum_{k} \tilde{W}_{R,k} \delta_{\tilde{x}_{R,k}}(\cdot) \approx \pi(\cdot)$.
- 6. Initialize SMC iteration index: $r \leftarrow 0$.
- 7. Initialize annealing parameter: $\phi_r \leftarrow 0$.
- 8. for $k \in \{1, 2, ..., K\}$ do
- 9. Initialize particles $x_{0,k} \leftarrow (\theta_{0,k}, t_{0,k}) \sim p(\cdot)$.
- 10. Initialize weights to unity: $w_{0,k} \leftarrow 1$.
- 11. **for** $r \in \{1, 2, ..., R\}$ **do**
- 12. **for** $k \in \{1, 2, ..., K\}$ **do**
- 13. Sample particles $\tilde{x}_{r,k} \sim K_r(x_{r-1,k}, \cdot)$; K_r is a π_r -invariant Metropolis-Hastings kernel.

14. Compute unnormalized weights: $w_{r,k} = [p(y|x_{r-1,k})]^{\phi_r - \phi_{r-1}}.$

- 15. **if** *r* < *R* **then**
- 16. **for** $k \in \{1, 2, ..., K\}$ **do**
- 17. Resample the particles: $x_{r,k} \sim \sum_{k'} \tilde{W}_{r,k'} \delta_{\tilde{x}_{r,k'}}(\cdot)$.
- 18. else
- 19. No resampling needed at the last iteration.
- 20. Return the particle population $\tilde{x}_{r,..}, \tilde{W}_{r,..}$

Before moving on to more advanced versions of the algorithm, we provide first some intuition to motivate the need for resampling. Theoretically, the algorithm produces samples from an artificial distribution with state space $\mathcal{X} \times \mathcal{X} \times \cdots \times \mathcal{X} = \mathcal{X}^R$ (this is described in more detail in Appendix 2). However, since we only make use of one copy of \mathcal{X} (corresponding to the particles at the final SMC iteration), we would like to decrease the variance of the state at iteration *R* (more precisely, of Monte Carlo estimators of functions of the state at iteration *R*). This is what resampling for iteration r < R achieves, at the cost of increasing the variance for the auxiliary part of the state space r < R. From this argument, it follows that resampling at the last iteration should be avoided.

When resampling is performed at every iteration but the last, an estimate of the marginal likelihood, p(y), is given by the product of the average unnormalized weights, namely:

$$\hat{Z}_{K} := \prod_{r=1}^{R} \frac{1}{K} \sum_{k=1}^{K} w_{r,k}.$$
(10)



FIGURE 1. An overview of the annealed SMC algorithmic framework for phylogenetic trees. The algorithm iterates the following three steps: (1) compute the weights using samples from the previous iteration, (2) perform MCMC moves to propose new samples, and (3) resample from the weighted samples to obtain an unweighted set of samples.

ADAPTIVE MECHANISMS FOR ANNEALED SMC

We discuss how two adaptive schemes from the SMC literature can be applied in our Bayesian phylogenetic inference setup to improve the scalability and usability of the algorithm described in the previous section. The first scheme relaxes the assumption that resampling is performed at every step, and the second is a method for automatic construction of the annealing parameter sequence. The two mechanisms go hand in hand and we recommend using both simultaneously. The combination yields Algorithm 3 which we explain in detail in the next two subsections.

The two adaptive mechanisms make theoretical analysis considerably more difficult. This is a common situation in the SMC literature. A common work-around used in the SMC literature is to run the algorithm twice, a first time to adaptively determine the resampling and annealing schedules, and then a second independent time using the schedule fixed in the first pass. We call it debiased adaptive annealed SMC.

Measuring Particle Degeneracy Using Relative (Conditional) Effective Sample Size

Both adaptive methods rely on being able to assess the quality of a particle approximation. For completeness,

we provide more background in Appendix 5 on the notions of effective sample size (ESS) and conditional ESS (CESS), a recent generalization which we use here (Zhou et al. 2016). The notion of ESS in the context of IS or SMC is distinct from the notion of ESS in the context of MCMC. The two are related in the sense of expressing a variance inflation compared with an idealized Monte Carlo scheme but they differ in the details. We will assume from now on that ESS refers to the SMC context.

We will use a slight variation of the definition of ESS and CESS where the measures obtained are normalized to be between zero and one. Some tuning parameters of the adaptive algorithms are easier to express in this fashion. We use the terminology relative (conditional) ESS to avoid confusion. Motivated by the analysis of the error of Monte Carlo estimators, the key measure of particle degeneracy needed in the following is the relative CESS:

rCESS(W, u) =
$$\left(\sum_{k=1}^{K} W_k u_k\right)^2 / \sum_{k=1}^{K} W_k u_k^2$$
, (11)

where $W = (W_1, W_2, ..., W_K)$ is a vector of weights of a set of reference weighted particles being updated using a vector of non-negative values $u = (u_1, u_2, ..., u_K)$. What

Algorithm 3 An adaptive annealed SMC algorithm

- Inputs: (a) Prior over evolutionary parameters and trees *p*(*x*), where *x* = (θ, *t*); (b) Likelihood function *p*(*y*|*x*).
- 2. **Outputs:** (a) Approximation *Z* of the marginal data likelihood, $Z \approx p(y) = \int p(dx)p(y|x)$; (b) Approximation of the posterior distribution, $\sum_k \tilde{W}_{R,k} \delta_{\tilde{x}_{R,k}}(\cdot) \approx \pi(\cdot)$.
- 3. Initialize SMC iteration index: $r \leftarrow 0$.
- 4. Initialize annealing parameter: $\phi_r \leftarrow 0$.
- 5. Initialize marginal likelihood estimate: $Z \leftarrow 1$.
- 6. **for** $k \in \{1, 2, ..., K\}$ **do**
- 7. Initialize particles with independent samples: $x_{0,k} \leftarrow (\theta_{0,k}, t_{0,k}) \sim p(\cdot)$.

8. Initialize weights to unity:
$$w_{0,k} \leftarrow 1$$
.

- 9. for $r \in \{1, 2, ...\}$ do
- 10. Determine next annealing parameter: $\phi_r = \text{NextAnnealingParameter}(x_{r-1,.}, w_{r-1,.}, \phi_{r-1})$.
- 11. **for** $k \in \{1, ..., K\}$ **do**
- 12. Compute pre-resampling unnormalized weights: $\tilde{w}_{r,k} = w_{r-1,k}[p(y|x_{r-1,k})]^{\phi_r \phi_{r-1}}$.
- 13. Sample particles $\tilde{x}_{r,k} \sim K_r(x_{r-1,k}, \cdot)$; K_r is a π_r -invariant Metropolis-Hastings kernel.
- 14. **if** $\phi_r = 1$ **then**
- 15. update $Z \leftarrow (Z/K) \cdot \sum_k \tilde{w}_{r,k}$, then return updated *Z* and particle population $\tilde{x}_{r,.}, \tilde{W}_{r,.}$.
- 16. **else**
- 17. **if** particle degeneracy is too severe, i.e. $rESS(\tilde{W}_{r,..}) < \epsilon$ **then**

18. Update marginal likelihood estimate, $Z \leftarrow$ $(Z/K) \cdot \sum_k \tilde{w}_{r,k}.$ Resample the particles. 19. 20. for $k \in \{1, ..., K\}$ do 21. Reset particle weights: $w_{r,k} = 1$. 22. else for $k \in \{1, ..., K\}$ do 23. ⊳ No 24. $w_{r,k} = \tilde{w}_{r,k}; x_{r,k} = \tilde{x}_{r,k}.$ resampling is needed.

W and *u* specifically represent will be explained in the next subsection.

Having a high rCESS value is a necessary but not sufficient condition for a good SMC approximation. If it is low during some of the intermediate SMC iterations, then the ESS at the final iteration may not be representative of the true posterior approximation quality.

Dynamic Resampling

As explained in Basic Annealed SMC Algorithm section , the construction of the proposal guarantees that as the difference $\phi_r - \phi_{r-1}$ goes to zero, the fluctuation of the weights vanishes. In this context (of having small weight updates), resampling at every

iteration is wasteful. Fortunately, SMC algorithms can be modified to forgo a subset of the resampling steps. From a theoretical stand-point, this is achieved by "grouping" the SMC proposals when they are not separated by a resampling round (and grouping similarly the intermediate distributions γ_r). For example, to resample every other round, use a transformed SMC algorithm with proposal $K'_{r/2}(x_r, (x_{r+1}, x_{r+2})) =$ $K_{r+1}(x_r, x_{r+1})K_{r+2}(x_{r+1}, x_{r+2})$, for each even *r*. For convenience, this can be implemented as an algorithm over *R* iterations instead of R/2, with two modifications: first, when resampling is skipped, we multiply the weights; otherwise, we reset the weights to one after resampling. This is implemented in Lines 12 and 21 of Algorithm 3. Second, we only use the weights corresponding to resampling rounds in the estimate of the marginal likelihood (Equation (10)). This is implemented in Lines 15 and 18 of Algorithm 3.

Instead of specifying in advance the subset of iterations in which resampling should be performed, it is customary in the SMC literature to determine whether to resample in an adaptive fashion (Doucet and Johansen 2011). To do so, the standard approach is to compute a measure of particle degeneracy at every iteration, and to perform resampling only when the particle degeneracy exceeds a predetermined threshold. In Appendix 4, we empirically compare the performance of adaptive annealed SMC algorithm with different resampling thresholds. All our numerical experiments use the multinomial resampling method, but we recommend more advanced schemes such as stratified resampling (Douc and Cappé 2005).

The standard measure of particle degeneracy used for this purpose is called the relative ESS, defined as:

$$\operatorname{ESS}(\tilde{W}_{r,\cdot}) = \left(K \sum_{k=1}^{K} \tilde{W}_{r,k}^2 \right)^{-1}.$$
 (12)

The above formula can be shown to be a special case of rCESS, Equation (11), as follows. Let r^* denote the iteration of the latest resampling round preceding the current iteration r. This implies $W_{r^*,k} = 1/K$ for all k. Plugging in the weight update $u_k = \tilde{w}_{r,k}$ into Equation (11), we obtain

$$CESS(W_{r^*,.},\tilde{w}_{r,.}) = \left(\sum_{k=1}^{K} \frac{1}{K} \tilde{w}_{r,k}\right)^2 / \sum_{k=1}^{K} \frac{1}{K} \tilde{w}_{r,k}^2$$
$$= \frac{1}{K} \frac{\left(\sum_{k=1}^{K} \tilde{w}_{r,k}\right)^2}{\sum_{k=1}^{K} \tilde{w}_{r,k}^2}$$
$$= \left(K \sum_{k=1}^{K} \tilde{W}_{r,k}^2\right)^{-1}.$$

1

Adaptive Determination of Annealing Parameters

Our sequence of intermediate artificial distributions π_r as defined in Equation (4) is determined by the choice of the annealing schedule, { ϕ_r }, or equivalently, by choosing the successive differences $\phi_r - \phi_{r-1}$. Ideally, the sequence of intermediate distributions changes gradually from the prior distribution ($\phi_0=0$) to the posterior distribution ($\phi_R=1$) so that the propagated particles from the current iteration can well approximate the next intermediate distribution.

In practice constructing such a sequence $\{\phi_r\}_{r=1,..,R}$ is difficult and inconvenient. Not only the number of distributions *R* to get a certain accuracy may depend on the number of taxa, the number of sites, and the complexity of the evolutionary model, but also the optimal spacing between consecutive annealing parameters is in general non-regular. To alleviate this, in the following we borrow an adaptive strategy from the Approximate Bayesian Computation literature (Del Moral et al. 2012), also generalized to Bayesian model selection in Zhou et al. (2016).

The adaptive annealing scheme is based on two observations. First, our discrete set of intermediate distributions $\pi_1, \pi_2, ..., \pi_R$ are actually continuously embedded into a continuum of distributions indexed by $\phi \in [0,1]$. Second, in the SMC algorithm presented in Algorithm 2, the weight update, Line 14, depends only on x_{r-1} (whereas in general SMC algorithms, the weight update could depend on both x_{r-1} and x_r ; here it does not because of cancellation explained in Appendix 2). The consequence of the lack of dependence on x_r is that we can swap the order of proposal (Line 13) and particle weighting (Line 14) in Algorithm 2. So instead of computing the weights only for one predetermined annealing parameter ϕ_r , we can search over several tentative values. For each tentative value, we can score the choice using a measure of weight degeneracy applied to the putative weights. Crucially, each choice can be quickly scored without having to propose particles, which is key since proposals are typically the computational bottleneck: in a phylogenetic context, the cost of one proposal step scales linearly in the number of sites whereas the search over ϕ_r proposed in this section has a running time constant in the number of sites and taxa. This is because the search involves fixed values of $p(y|x_{r-1,k})$ cached from the last proposal step, which are exponentiated to different values.

Based on these observations, we select an annealing parameter ϕ such that we achieve a controlled increase in particle degeneracy, namely such that

$$g(\phi) = \alpha g(\phi_{r-1}), \tag{13}$$

where the function $g: [\phi_{r-1}, \infty) \rightarrow [0, 1]$ is defined as

$$g(\phi) = \operatorname{rCESS}\left(W_{r-1,\cdot}, p(y|x_{r-1,\cdot})^{\phi-\phi_{r-1}}\right),$$

and $\alpha \in (0, 1)$ is a tuning parameter, which in practice is close to 1. By construction, $g(\phi_{r-1})=1$, so Equation (13) is equivalent to $g(\phi)=\alpha$.

More precisely, since we want $\phi \in [0,1]$, the annealing parameter adaptation procedure, NextAnnealingParameterm (Algorithm 4), is designed to return $\phi_r = 1$ if $g(1) \ge \alpha$. Otherwise, because there is no closed-form solution for ϕ in Equation (13), we use bisection to solve this one-dimensional search problem in the interval $\phi \in (\phi_{r-1}, 1)$ (Line 7 of Algorithm 4).

We now argue that the search problem in Line 7 of Algorithm 4 always has a solution. Indeed, *g* is a continuous function with, on the left end of the search interval, $g(\phi_{r-1})=1$, and on the right end, $g(1) < \alpha$ (otherwise the algorithm sets $\phi_r = 1$ in Line 5). It follows that there must indeed be an intermediate point ϕ^* with $g(\phi^*)=\alpha$. Note that continuity and the identification of the left end point of the interval is possible thanks to the form of our weight update in Equation (9), hence justifying the earlier informal argument about the need to have the fluctuation of the weights disappearing as $\phi_r - \phi_{r-1}$ goes to zero.

As in the previous section on dynamic resampling, NextAnnealingParameter is again based on relative CESS, but this time, we are interested in the degeneracy of a single iteration, i.e. we do not trace back until the previous resampling step (because the optimization over the annealing schedule can only impact the current iteration). As a corollary, the previous iteration's particles are not always equally weighted, hence the simplification in Equation (12) is not possible here and we use the full formula for relative CESS.

Algorithm 4 Procedure NextAnnealingParameter

- Inputs: (a) Particle population from previous SMC iteration (*x*_{r-1,.}, *w*_{r-1,.}); (b) Annealing parameter φ_{r-1} of previous SMC iteration; (c) A degeneracy decay target α ∈ (0,1).
- Outputs: automatic choice of annealing parameter φ_r.
- 3. Initialize the function *g* assessing the particle population quality associated to a putative annealing parameter φ:

$$g(\phi) = \text{rCESS}\left(W_{r-1,\cdot}, p(y|x_{r-1,\cdot})^{\phi-\phi_{r-1}}\right)$$
$$= \frac{\left(\sum_{k=1}^{K} W_{r-1,k} p(y|x_{r-1,k})^{\phi-\phi_{r-1}}\right)^2}{\sum_{k=1}^{K} W_{r-1,k} p(y|x_{r-1,k})^{2(\phi-\phi_{r-1})}}.$$

4. if $g(1) \ge \alpha$ then 5. return $\phi_r = 1$.

5. retur 6. **else**

7. return $\phi_r = \phi^* \in (\phi_{r-1}, 1)$ such that $g(\phi^*) = \alpha$ via bisection.

The parameter α used in Algorithm 4 encodes the decay in particle population quality that we are aiming for. Based on our experiments we recommend values very close to one. For this reason, we reparameterize the parameter α into $\alpha = 1 - 10^{-\beta}$ and recommend a default

value of $\beta = 5$ as a reasonable starting point. Increasing β improves the approximation accuracy.

Computational Complexity

The computational complexity of annealed SMC is linear in both the number of intermediate distributions R and the number of particles K. Naively, the resampling step scales like $O(K^2)$, but a linear time multinomial resampling algorithm is obtained by generating order statistics via normalization of a Poisson process (Devroye 1986, Section 2.1, p. 214). This technique is well known in the SMC literature (Doucet and Johansen 2011). Alternatively, one can use stratified or systematic resampling (Doucet and Johansen 2011), which provides a simple to implement linear time resampling algorithm.

The memory consumption of annealed SMC is linear in *K* and constant in *R*.

REVIEW OF OTHER MARGINAL LIKELIHOOD ESTIMATION METHODS

For completeness, we review here some alternatives to Equation (10) for estimating marginal likelihoods, which we will compare to SMC from both a theoretical and empirical stand-point.

Stepping Stone

The Stepping Stone (SS) algorithm (Xie et al. 2010) is a method for marginal likelihood estimation. It is widely used via its MrBayes implementation (Huelsenbeck and Ronquist 2001). As with SMC, the SS method introduces a list of annealed posterior distributions connecting the posterior distribution and the prior distribution. We use a notation analogous to SMC, with $\{\pi_d\}_{d=0,1,...,D}$ denoting the intermediate distributions, $\pi_d(x) \propto \gamma_d(x) = p(y|x)^{\phi_d} \pi(x), 0 = \phi_0 < \phi_1 < \phi_2 < \cdots < \phi_D = 1$. The marginal likelihood *Z* can be written as

$$Z \equiv Z_D = Z_0 \prod_{d=1}^D \frac{Z_d}{Z_{d-1}}.$$

We can rewrite the ratio of Z_d and Z_{d-1} as

$$\frac{Z_d}{Z_{d-1}} = \int \frac{\gamma_d(x)}{\gamma_{d-1}(x)} \pi_{d-1}(x) dx.$$
 (14)

The SS method prescribes running several MCMC chains targeting $\pi_{d-1}(x)$ to obtain *N* posterior samples $x_{d-1,1}, x_{d-1,2}, \dots, x_{d-1,N}$, then

$$\frac{\widehat{Z_d}}{Z_{d-1}} = \frac{1}{N} \sum_{i=1}^{N} \{ p(y|x_{d-1,i}) \}^{\phi_d - \phi_{d-1}}.$$
 (15)

The estimator of the marginal likelihood admits the form

$$\widehat{Z_D} = \prod_{d=1}^{D} \frac{1}{N} \sum_{i=1}^{N} \{ p(y|x_{d-1,i}) \}^{\phi_d - \phi_{d-1}}.$$

The number of intermediate distributions is a trade-off between computing cost and accuracy. A larger number of MCMC chains can provide a better approximation for the marginal likelihood, but the computational cost will be higher. To make fair comparison between the marginal likelihood estimators provided by the annealed SMC and SS, we set $K_{SMC}R_{SMC} = N_{SS}D_{SS}$. Another factor that will impact the SS estimator is the choice of annealing parameter sequence $\{\phi_d\}_{d=1,2,...,D}$. In this paper, we use the annealing scheme $\phi_d = (d/D)^{1/a}$ recommended by Xie et al. (2010), where *a* is between 0.2 and 0.4.

Linked Importance Sampling

Stepping stone uses IS to approximate the ratio of marginal likelihoods for two intermediate distributions. However, the IS approximation would be poor if the two successive distributions do not have enough overlaps. Linked IS (Neal 2005) improves the performance of IS by introducing bridge distributions, for example "geometric" bridge: $\gamma_{d-1*d}(x) = \sqrt{\gamma_{d-1}(x)\gamma_d(x)}$. More importantly, Linked Importance Sampling (LIS) provides an unbiased marginal likelihood estimator. The ratio of two marginal likelihoods can be written as

$$\frac{Z_d}{Z_{d-1}} = \frac{Z_{d-1*d}}{Z_{d-1}} \bigg/ \frac{Z_{d-1*d}}{Z_d} = \bigg\{ \int \frac{\gamma_{d-1*d}(x)}{\gamma_{d-1}(x)} \pi_{d-1}(x) dx \bigg\} \bigg/ \\ \bigg\{ \int \frac{\gamma_{d-1*d}(x)}{\gamma_d(x)} \pi_d(x) dx \bigg\}.$$

For d = 1, ..., D, to estimate the ratio Z_d/Z_{d-1} , we first run MCMC targeting $\pi_{d-1}(x)$ to obtain N posterior samples $x_{d-1,1}, x_{d-1,2}, ..., x_{d-1,N}$ (when d = 1, we sample from the prior distribution). Then we sample the initial state of π_d . Two successive MCMC chains $\pi_{d-1}(x)$ and $\pi_d(x)$ are linked by a state $x_{d-1,\mu_{d-1}}$ where index μ_{d-1} is sampled from {1, 2, ..., N} according to the following probabilities:

$$p(\mu_{d-1}|x_{d-1,1:N}) = \frac{\gamma_{d-2*d-1}(x_{d-1,\mu_{d-1}})}{\gamma_{d-1}(x_{d-1,\mu_{d-1}})} \bigg/ \sum_{i=1}^{N} \frac{\gamma_{d-2*d-1}(x_{d-1,i})}{\gamma_{d-1}(x_{d-1,i})}.$$

In case d = 1, the linked state μ_0 is uniformly sampled from the *N* samples of $\pi_0(x)$. Finally, we run MCMC chain $\pi_d(x)$ starting from initial state $x_{d-1,\mu_{d-1}}$ to obtain *N* posterior samples $x_{d,1}, x_{d,2}, \dots, x_{d,N}$. The ratio of two marginal likelihoods can be approximated by

$$\frac{\widehat{Z_d}}{Z_{d-1}} = \frac{\widehat{Z_{d-1*d}}}{Z_{d-1}} \Big/ \frac{\widehat{Z_{d-1*d}}}{Z_d} = \left\{ \frac{1}{N} \sum_{i=1}^N \frac{\gamma_{d-1*d}(x_{d-1,i})}{\gamma_{d-1}(x_{d-1,i})} \right\} \Big/ \\
\left\{ \frac{1}{N} \sum_{i=1}^N \frac{\gamma_{d-1*d}(x_{d,i})}{\gamma_d(x_{d,i})} \right\}.$$

In this article, we use the "geometric" bridge. Hence, the estimator of ratio can be simplified to

$$\frac{\widehat{Z_d}}{Z_{d-1}} = \left\{ \sum_{i=1}^N \{ p(y|x_{d-1,i}) \}^{\frac{\phi_d - \phi_{d-1}}{2}} \right\} / \left\{ \sum_{i=1}^N \{ p(y|x_{d,i}) \}^{\frac{\phi_{d-1} - \phi_d}{2}} \right\}$$

We refer to Appendix 6 for more background on the LIS algorithm.

THEORETICAL PROPERTIES

In this section, we review three theoretical properties of interest, *consistency*, *marginal likelihood estimate unbiasedness*, and *asymptotic normality*, with an emphasis on their respective practical importance.

Properties of Annealed SMC

In the context of SMC algorithms, the first property, *consistency* means that as the number of particles is increased, the approximation of posterior expectations can become arbitrarily close to the true posterior expectation. This makes the approximation in Equation (7) more precise:

$$\sum_{k=1}^{K} W_{r,k} f(x_{r,k}) \to \int \pi_r(x) f(x) dx \quad \text{as } K \to \infty, \qquad (16)$$

provided f satisfies regularity conditions, for example f is bounded, and where convergence of the random variables holds for a set of random seeds having probability one. See for example Wang et al. (2015).

Consistency can be viewed as the "bare minimum" expected from modern SMC algorithms. A more informative class of results consists in central limit theorem equivalents of Equation (16). These results can be used to assess the total variance of Monte Carlo estimators (whereas measures such as ESS described previously are local in nature), see Chan and Lai (2013). However, since numerically stable versions of these methods are still at their infancy (Olsson and Douc 2017), we will focus the remaining on the third property, unbiasedness.

We say an estimator \hat{Z} for a constant *Z* is unbiased if $\mathbb{E}[\hat{Z}]=Z$. Here the expectation is defined with respect to the randomness of the approximation algorithm. This contrasts with the classical statistical definition of unbiasedness in which the randomness comes from the data generation process.

For SMC algorithms, unbiasedness holds in a more restrictive sense compared with consistency. In general:

$$\mathbb{E}\left[\sum_{k=1}^{K} W_{r,k} f(x_{r,k})\right] \neq \int \pi_r(x) f(x) dx, \qquad (17)$$

in other words, repeatedly running SMC with a fixed number of particles but different random seeds and averaging the results does not provide arbitrarily precise approximations (the same negative result holds with MCMC). However, if we restrict our attention to marginal likelihood estimates, remarkably the unbiasedness property does hold (Del Moral et al. 2006), that is for any finite K, \hat{Z}_K as defined in Equation (10) is such that:

$$\mathbb{E}\left[\hat{Z}_{K}\right] = Z = \int \gamma_{r}(x) \mathrm{d}x. \qquad (18)$$

More details on the unbiasedness of the marginal likelihood SMC estimator and other theoretical properties of annealed SMC can be found in Appendix 2, Unbiasedness, Consistency and Central Limit Theorem subsection.

Although the notion of unbiasedness has been central to frequentist statistics because its inception, only in the past decade has it started to emerge as a property of central importance in the context of (computational) Bayesian statistics. Traditionally, the main theoretical properties analyzed for a given Monte Carlo method \hat{Z} estimating *Z* was consistency.

With the emergence of pseudo-marginal methods, the bias of Monte Carlo methods is now under closer scrutiny. Pseudo-marginal methods are MCMC methods which replace probability factors in the Metropolis-Hastings ratio by positive unbiased estimators of these probabilities. For example, Andrieu et al. (2010) provide examples where global parameters of state-space models are sampled using an MCMC algorithm where the probability of the data given the global parameters and marginally over the latent states is estimated using an SMC algorithm. We refer the reader to Andrieu and Roberts (2009) for more examples where unbiasedness is used to compose MCMC algorithms in order to attack inference in complex models. In the context of phylogenetic inference, this is useful for Bayesian analysis of intractable evolutionary models, see for example Hajiaghayi et al. (2014).

Another area where unbiasedness can play a role is for checking correctness of Monte Carlo procedures. In contrast to correctness checks based on consistency such as Geweke (2004), which are asymptotic in nature and hence necessarily have false positive rates (i.e. cases where the test indicates the presence of a bug when in fact the code is correct), checks based on unbiasedness can achieve a false positive rate of zero, using the strategy described in the next section.

Using Unbiasedness to Test Implementation Correctness

Typically, the algorithm shown in Algorithm 2 is implemented in a model-agnostic fashion. Hence it is reasonable to assume that we can construct test cases on discrete state spaces. For example, one can use phylogenetic trees with fixed branch lengths, or even simpler models such as hidden Markov models (HMMs). Furthermore, we conjecture that many software defects can be detected in relatively small examples, where exhaustive enumeration is possible, and hence Z can be computed exactly. We can determine sufficient complexity of the examples to use via code coverage tools (Miller and Maloney 1963).

We would like to test if equality of Equation (18)holds for a given implementation. The right-hand side can be computed easily because we assume the example considered is small. To compute analytically the expectation on the left-hand side, we use a method borrowing ideas from probabilistic programming (Wingate et al. 2011), and use an algorithm, called ExhaustiveRandom that automatically visits all possible execution traces τ_i of a given randomized algorithm. The execution trace of a randomized algorithm refers to a realization of all random choices in the algorithm (in the context of SMC, both the resampling steps and the proposal steps). ExhaustiveRandom enumerates all the execution traces while also computing the respective probability p_i of each trace. This is done by performing a depth first traversal of the decision tree corresponding to the randomized algorithm being tested. The number of execution traces grows exponentially fast but this is still a useful tool as very small examples are generally sufficient to reach code coverage.

For each execution trace τ_i , we can also obtain the normalization estimate \hat{z}_i corresponding to that trace, and hence get the value of the left-hand side of Equation (18) as $\sum_i p_i \hat{z}_i$. We used this check via an open source implementation of ExhaustiveRandom (https://github.com/alexandrebouchard/bayonet/blob /1b9772e91cf2fb14a91f2e5e282fcf4ded61ee22/src/main/ java/bayonet/distributions/ExhaustiveDebugRandom. java) to ensure that our software satisfies the unbiasedness property. See the numerical simulation section for details.

Properties of the SS Method

For the stepping stone method, the expected value of Equation (15) depends on the nature of the samples $x_{d-1,1}, x_{d-1,2}, \dots, x_{d-1,N}$. If they are independent, the procedure is unbiased. However, if the samples are obtained from a Markov chain, there are no guarantees that the procedure is unbiased unless the MCMC chain is initialized at the exact stationary distribution. In practice, this is not possible: Xie et al. (2010) use a burned-in MCMC chain, which implies that the chain is asymptotically unbiased, however for any finite number of iterations, a bias remains. Unfortunately, the two main motivations for unbiasedness (pseudo-marginal methods and the correctness checks described earlier) both require unbiasedness to hold for any finite number of Monte Carlo samples; asymptotic unbiasedness is not sufficient.

We show in the numerical simulation section an explicit counterexample where we compute the non-zero bias of the stepping stone method. This motivates the need for implementable unbiased methods, such as the annealed SMC method described in this work.

Comparison of Unbiased Marginal Likelihood Estimators

In Bayesian phylogenetics, the marginal likelihood estimate is generally a very small number. Instead of computing \hat{Z} directly, we compute the logarithm of the marginal likelihood estimate, $\log(\hat{Z})$. For SMC and LIS, although \hat{Z} is an unbiased estimator, taking the logarithm of \hat{Z} introduces bias. Jensen's inequality shows that $\log(\hat{Z})$ is a biased estimator of $\log(Z)$, and is generally underestimated,

$\mathbb{E}[\log(\hat{Z})] \leq \log(\mathbb{E}(\hat{Z})) = \log Z.$

This provides a tool to compare the performance of an unbiased normalization constant estimation method m_1 to another one m_2 . Suppose, we run each method M times with different seeds and a fixed computational budget. Let $L_i = \sum_{j=1}^{M} \log \hat{Z}_{i,j}/M$ denote the average estimate of the log marginal likelihood for the *i*-th method, where $\hat{Z}_{i,j}$ is the estimate with the *j*-th random seed using the *i*-th method. If m_2 is also unbiased then for M large enough, both m_1 and m_2 underestimate $\log \mathbb{E}[Z]$, and the largest L_i is closest to $\log \mathbb{E}[Z]$, which determines the best performing method. If m_2 is not unbiased, then if $L_1 > L_2$ and M is large enough, we can conclude that m_1 is superior (but we cannot confidently order the methods if $L_2 > L_1$).

However, the Monte Carlo counterparts of the orderings should be considered with a pinch of salt because the number of replicates *M* needed may be intractable in some cases.

SIMULATION STUDIES

Simulation Setup and Tree Distance

In order to simulate data sets, we first generated a set of random unrooted trees, including topology and branch lengths, as the reference trees. The tree topology was sampled from a uniform distribution. Each branch length was generated from an exponential distribution with rate 10.0.

Then, for each reference tree, we simulated DNA sequences using the K2P model with parameter $\kappa = 2.0$ (Kimura 1980). Although the main focus of this work is on marginal likelihood estimation, we also performed some benchmarking on the quality of the inferred trees. To do so, we used the majority-rule consensus tree (Felsenstein 1981) to summarize the weighted phylogenetic tree samples obtained from annealed SMC. We measured the distance between each estimated consensus tree to its associated reference tree using three types of distance metrics: the Robinson-Foulds (RF) metric based on sums of differences in branch lengths (Robinson and Foulds 1979), the Kuhner-Felsenstein (KF) metric (Kuhner and Felsenstein 1994), and the partition metric (PM), also known as symmetric difference or topology only RF metric (Robinson and Foulds 1981).



FIGURE 2. Graphical representation of (a) an HMM; (b) transitions between hidden states.



FIGURE 3. Estimates of the log marginal likelihood based on 100 independent random seeds for each configuration. The axis label "nTemps" refers to the number of equally spaced intermediate distributions ("inverse temperatures"). The black horizontal line shows the true value computed using the forward-backward algorithm.

Hidden Markov Models

As discussed when we introduced the unbiasedness correctness test, it is useful to perform some preliminary experiments on finite state models. We used a HMM with a finite latent state (see the graphical representations of a HMM and hidden state transitions in Fig. 2). The variables X_t shown in the figure are unobserved and take on discrete values with a distribution depending on the previous variable X_{t-1} . For each unobserved variable, we define an observed variable Y_t , also discrete, with a conditional distribution depending on X_t . The latent state space in our experiment was set to $\{0, 1, 2, 3, 4\}$ and latent transitions were set uniformly across neighbor integers. The emissions we used take two possible values with conditional probabilities given by (0.2, 0.8), (0.1, 0.9), (0.01, 0.99), (0.2, 0.8), and (0.3, 0.7).The proposals were based on the Gibbs sampler on a single variable. The posterior distribution of interest is over the latent variables X_1, X_2, \dots given the observations Y_1, Y_2, \dots Of course, such a model would not normally be approached using approximate inference methods. Moreover, notice that this is a non-standard way of using SMC for a sequential model where we do not make use of the sequential structure of the model.

We first performed unbiasedness correctness tests on a chain of length two based on three equally spaced annealing parameters (0, 1/2, 1), and observations (0, 1). We first computed the true marginal likelihood, 0.345. Using the method described in the Theoretical Properties section, we computed the exact value of $\mathbb{E}[\hat{Z}]$ by exhaustive enumeration of all execution traces for SMC and the SS method. For SMC with two particles, the ExhaustiveRandom algorithm enumerated 1,992,084 traces resulting in an expectation of 0.3449999999999525. For SS with two MCMC iterations per annealing parameter, the ExhaustiveRandom algorithm enumerated 1,156,288 traces resulting in an expectation of 0.33299145257312235. This supports that SMC is unbiased and provides an explicit counterexample of the bias of the stepping stone method.

Second, we ran experiments on larger versions of the same model, a chain of length 32, as well as with more annealing steps and particles per step. In this regime, it is no longer possible to enumerate all the execution traces so we averaged over 100 realizations of each algorithm instead. The true marginal likelihood can still be computed using a forward-backward algorithm. We show the results in Figure 3.

Comparison of Marginal Likelihood Estimates

In this section, we benchmark the marginal likelihood estimates provided by adaptive annealed SMC (ASMC), debiased adaptive annealed SMC (DASMC),



FIGURE 4. Marginal likelihood (in log scale) estimates for different numbers of taxa with a fixed computational budget.

TABLE 1. Number of annealing parameters (R_{SMC}) for ASMC with K = 1000 and $\beta = 5$ #taxa510152025 R_{SMC} 19323741514260477219

deterministic annealed SMC (DSMC), LIS, and SS. In DASMC, the annealing scheme was determined before running annealed SMC using the same annealing parameters obtained from the ASMC. In DSMC, we used the annealing scheme $\phi_r = (r/R)^3$ with a predetermined *R*.

In the first experiment, we focus on evaluating the marginal likelihood estimates using ASMC, DASMC, LIS, and SS with the same computing budget. We simulated unrooted trees of varying sizes (numbers of taxa): 5, 10, 15, 20, and 25. For each tree, we generated one data set of DNA sequences. Sequence length was set to 100. The execution of each algorithm and setting was repeated 100 times with different random seeds. We used β =5 for adaptive annealed SMC, and the number of particles was set to 1000. In stepping stone and linked IS, we set the total number of heated chains *D* to 50, and the annealing scheme was set to $\phi_d = (d/D)^3$, where d = 1, 2, ..., D. We enforced $K_{SMC}R_{SMC} = N_{SS}D_{SS} = N_{LIS}D_{LIS}$ in order to make the comparisons fair. Information about R_{SMC} are shown in Table 1.

Figure 4 shows the comparison of the performance of the four algorithms in terms of the marginal likelihood in log scale as the number of taxa increases. As described in the Theoretical Properties section, for the unbiased estimators, we have asymptotically that the log of the marginal likelihood should underestimate the marginal likelihood by Jensen's inequality. The results support that ASMC and DASMC can achieve more accurate marginalized likelihood estimates compared with SS and LIS with the same computational cost. The performances of the two SMC algorithms are quite similar, whereas the marginal likelihood estimates provided by LIS and SS are close to each other. In Appendix 9, we describe an experiment comparing ASMC, DASMC, LIS, and SS with a very large value of K. The mean of log marginal likelihood for the four

methods are close (reduction of the gap is expected, because all methods are consistent), whereas ASMC and DASMC still exhibit smaller variance across seeds compared with LIS and SS.

Another experiment was conducted to measure the variability of the marginal likelihood estimates from each algorithm, by comparing the coefficients of variation (CV) for different numbers of taxa with the same setting. The CV is defined as $CV = sd(\tilde{Z})/\mathbb{E}(\tilde{Z})$. We simulated 70 trees, increasing the number of taxa (from 10, 15, 20, 25, 30, 35, 40; 10 trees of each size), and created 10 data sets for each tree. For each data set, we repeated each algorithm 10 times with different random seeds. The upper bound of CV equals $\sqrt{n-1}$, where *n* represents the number of repeats with different random seeds in experiments. We refer to Appendix 7 for the derivation of the upper bound of CV. In our setting, this upper bound is $\sqrt{10-1}=3$. In ASMC, the computational cost was fixed at K = 1000 and $\beta = 5$. In DSMC, we used the same number of particles, and the annealing scheme was set to $\phi_r = (r/R)^3$, where the total number of annealing parameters R was fixed to be the one obtained from running ASMC with K = 1000 and $\beta = 5$ for a tree with 10 taxa.

Figure 5 displays the CV for ASMC and DSMC as a function of the number of taxa. The error bars in the figure represent 95% confidence intervals. The CV of DSMC increases faster than ASMC as the number of taxa gets larger than 15. It gradually converges to the upper bound of CV as the number of taxa reaches 35. The CV of ASMC increases more slowly as the number of taxa increases.

Comparison of Model Selection by Annealed SMC versus SS

In this section, we compare the performance of ASMC and SS on a Bayesian model selection task. We simulated 20 unrooted tree of 10 taxa using a uniform distribution for the tree topology and branch lengths generated from an exponential distribution with rate 10. A total of 60 data sets of DNA sequences of length 500 were generated using each of the simulated tree and the following



FIGURE 5. CV for the marginal likelihood estimates versus the number of taxa for ASMC and DSMC with a fixed number of particles.

TABLE 2. Comparison of model selection by ASMC and SS based on the Bayes factor

| | | Data generated from | | |
|--------|-------|---------------------|-----|-------|
| Method | Model | JC69 | K2P | GTR+Γ |
| ASMC | JC69 | 20 | 0 | 1 |
| | K2P | 0 | 20 | 1 |
| | GTR | 0 | 0 | 18 |
| SS | JC69 | 20 | 0 | 4 |
| | K2P | 0 | 20 | 1 |
| | GTR | 0 | 0 | 15 |

three evolutionary models: JC69, K2P, and GTR+ Γ . The parameter κ in the K2P model was set to 2.0. In the GTR+ Γ model, a symmetric Dirichlet distribution with parameters (10, 10, 10, 10) was used to generate the base frequencies, and a symmetric Dirichlet with parameters (10, 10, 10, 10, 10, 10) was used to generate the GTR relative rates in the rate matrix. The discrete gamma distribution with 4 categories was used to convey among-site rate heterogeneity, with the gamma shape parameter drawn from a Gamma distribution with parameters (2, 3).

For each data set, marginal likelihoods were estimated by ASMC and SS using three evolutionary models, JC69, K2P, and GTR, respectively. In ASMC, we used K = 1000and $\beta = 4$. The total number of iterations in SS was set to the product of the number of particles and number of iterations in ASMC. Table 2 shows the Bayesian model selection results. Both ASMC and SS choose the correct model for all of the 20 data sets generated from the JC69 and K2P model, respectively. For the data generated from GTR+ Γ , SMC chooses the closest model, GTR, 18 times out of 20, whereas SS only chooses GTR 15 times out of 20.

Comparison of Tree Distance Metrics

In this section, we compare the quality of reconstructed phylogenies using synthetic data. We simulated one unrooted tree, the reference tree, with 50 taxa and then generated one data set of DNA sequences

TABLE 3. Comparison of tree distance metrics using ASMC and MCMC

| Method | R | Κ | Metric | Value |
|--------|----------|-----|----------------------|------------|
| ASMC | 54,876 | 100 | ConsensusLogLL | -72,787.99 |
| | 54,876 | 100 | BestSampledLogLL | -72,826.17 |
| | 54,876 | 100 | PartitionMetric | 0 |
| | 54,876 | 100 | RobinsonFouldsMetric | 0.70623 |
| | 54,876 | 100 | KuhnerFelsenstein | 0.00990 |
| MCMC | 1.0E+07 | | ConsensusLogLL | -72,833.82 |
| | 1.0E+07 | | PartitionMetric | 0 |
| | 1.0E+07 | | RobinsonFouldsMetric | 0.92031 |
| | 1.0E+07 | | KuhnerFelsenstein | 0.03138 |
| MCMC2 | 5.49E+06 | | ConsensusLogLL | -72,784.86 |
| | 5.49E+06 | | PartitionMetric | 0 |
| | 5.49E+06 | | RobinsonFouldsMetric | 0.73,644 |
| | 5.49E+06 | | KuhnerFelsenstein | 0.01066 |

of length 2000 from this tree. The ASMC was run with $\beta = 6$ and K = 100. The MCMC algorithm was initialized with a random tree from the prior distribution. To make a fair comparison, we set the number of MCMC iterations to be no less than $K_{SMC}R_{SMC}$. We discarded 20% of the MCMC chain as "burn-in." Table 3 summarizes the iteration numbers, the log-likelihood of the consensus tree and tree distance metrics from running ASMC and MCMC. Although the computational cost of MCMC is set to about twice as high as ASMC, the log-likelihood of the consensus tree from ASMC is much higher than that from MCMC. In addition, ASMC achieves much lower RF and KF distances to the reference tree. Further, to confirm that both ASMC and MCMC can converge to the same posterior distribution, MCMC was rerun with a better starting value, namely the consensus tree obtained after running ASMC. This run of MCMC is denoted as MCMC2 in Table 3. The computational cost of MCMC2 is set the same as the ASMC algorithm. This time MCMC achieved similar consensus tree log-likelihood and tree distance metrics compared with ASMC, which supports that MCMC is indeed "trapped" in a subspace.

Influence of Number of Threads, β , and K

The runtime of the ASMC is dependent on the number of threads, as well as on the tuning parameters β and *K*. In this section, we focus on investigating the effects of these factors on ASMC. We simulated an unrooted tree with 30 taxa at the leaves, and then generated DNA sequences of length 1500.

Next, Figure 6 displays the computing time versus number of threads for an implementation of ASMC where the proposal step is parallelized. The error bars represent the 95% confidence intervals based on 100 runs. We used K=1000 and β =2 for each number of threads. The results indicate that by increasing the number of cores, the speed of the ASMC algorithm can be increased notably.

In Table 4, we compare the performance of ASMC algorithm as a function of *K*, with β fixed at 5. We chose

 TABLE 4. Comparison of adaptive SMC algorithm with different numbers of particles and β

 K (β = 5)
 $\log(Z)$ β (K = 1000)
 $\log(Z)$

| K(p=3) | $\log(Z)$ | p(K = 1000) | $\log(Z)$ |
|--------|----------------------------------|-------------|----------------------------------|
| 100 | -28,288.5 (-28,283.9, -28,293.7) | 3 | -28,524.8 (-28,466.2, -28,641.0) |
| 300 | -28,283.5(-28,281.1, -28,287.9) | 4 | -28,312.2(-28,304.5, -28,328.8) |
| 1000 | -28,280.5(-28,278.3,-28,283.5) | 5 | -28,280.5(-28,278.3,-28,283.5) |
| 3000 | -28,279.3 (-28,278.1, -28,280.4) | 5.3 | -28,279.5 (-28,278.7, -28,280.5) |



FIGURE 6. Computing time of ASMC using multiple threads.

four different particle values K = 100, 300, 1000, 3000. The marginal likelihood estimates improve as K increases.

We also compared the performance of ASMC algorithm as a function of β , with K = 1000. We selected four distinct β values, $\beta = 3, 4, 5, 5.3$. As expected, the marginal likelihood estimates improve when β increases. The likelihood of the consensus trees and tree distance metrics provided by these two experiments are displayed in Appendix 8. In practice, a value of β close to 5 is recommended as the default value.

Trade-Off between R and K

We conducted an experiment to investigate, for a given amount of computation, the relative importance of R and *K* in improving the quality of the posterior distribution inferred by annealed SMC. We used DSMC with a cubic annealing scheme. We selected values for the tuning parameter K (100,300,1000,3000,10,000), and for each value of *K*, a corresponding value for *R* such that the total computation cost $K \cdot R$ is fixed at 10⁶. We simulated one unrooted tree of 15 taxa, and generated one data set of DNA sequences. Sequence length was set to 300. Figure 7 displays the marginal likelihood estimates, and KF metric provided by DSMC with different K values when the total computational budget $(K \cdot R)$ is fixed. This results indicate that for a given amount of computation, a relatively small K and a large R is optimal. However, the value of K cannot be too small, as an extremely small K necessarily leads to a large Monte Carlo variance.

Analysis of Subsampling SMC

Subsampling SMC, detailed in Appendix 1, can be used to speed up the SMC algorithms at the cost of decreasing the accuracy of estimation. The idea is to divide the data, the sites of biological sequences in our case, into batches, and only use a subset of the data in the intermediate distributions. In this section, we evaluate the impact of the batch size (the number of sites of biological sequences in each batch), denoted b_s , on the speed of the algorithm and the posterior approximation.

In a first experiment, we analyzed the relative computational cost of subsampling SMC with respect to annealed SMC for different batch sizes. We simulated an unrooted tree with 10 taxa, and then generated DNA sequences of length 6000. The annealing parameter sequence ϕ_r , $r=0,1,\ldots,R$, was chosen by running adaptive ASMC using $\beta = 4$ and K = 100. The computational cost in this subsection is measured by the total number of sites involved in computing the unnormalized posterior and the weight update function. For example, in this simulation study, the total number of annealing parameters in adaptive ASMC is 2318, and the number of sites involved in each SMC iteration is 6000. Using the fact that the likelihood for particles evaluated at iteration r-1 can be used to evaluate the weight update function at iteration r, the total cost for ASMC is $2318 \cdot 6000 = 1.39 \times 10^7$. Figure 8 displays the ratio of computational cost (subsampling/annealing) versus the batch size. The cost ratio increases slowly when we increase the batch size from 1 to 100.

We investigated the performance of subsampling SMC with different batch sizes, $b_s = 1, 10, 100, 1000, 6000$ in terms of phylogenetic tree inference. We used K = 100and ran the subsampling SMC algorithm 10 times for each value of b_s . The schedule ϕ_r used to compute the annealing parameter $\psi(s, \phi_r)$ in subsampling SMC was obtained by running adaptive annealed SMC once using $\beta = 4$ and K = 100. Figure 9 displays the performance of the subsampling algorithm with different b_s . As expected, there is a trade-off between the computational cost and accuracy of most metrics: for all metrics except the partition metric, subsampling produces lower quality approximations at a lower cost. However, if the user only require a reconstruction of the tree topology, the partition metric results provide an example where subsampling is advantageous.

Comparison of ASMC and Combinatorial SMC (CSMC)

We compared the performance of the annealed SMC and the combinatorial SMC (CSMC) algorithm (Wang et al. 2015) for three different kinds of trees: clock, relaxed clock, and nonclock. The clock trees were simulated by assuming that the waiting time between two coalescent



FIGURE 7. Performance of deterministic SMC algorithm on a fixed computational budget ($K \cdot R = 10^6$). We select 5 values of K, 100, 300, 1000, 3000, 10, 000, from left to right on X-axis.



FIGURE 8. Ratio of cost (subsampling/annealing) versus the batch size.

events is exponentially distributed with rate 10. The relaxed clock trees were obtained by perturbing the branch length of clock trees. More specifically, we modified each branch of length *l* by adding to it a noise randomly sampled from Unif(-0.3l, 0.3l). The nonclock trees were simulated with uniformly distributed tree topologies and exponentially distributed branch lengths with rate 10.

For each type of phylogenetic tree, we simulated 10 trees with 10 leaves. The JC69 evolutionary model was

used to generate sequences of length 500. Three data sets were generated for each tree. We ran the annealed SMC, DASMC, and CSMC, respectively, for each data set three times with different random seeds. In the annealed SMC, β was set to 5, and *K*=100; in CSMC, the number of particles was set to 100,000.

Figure 10 shows the boxplots of log-likelihood of the consensus trees, three tree distance metrics from the true trees, and computing time (in milliseconds) obtained from running the three algorithms for clock trees (top), relaxed clock trees (middle), and nonclock trees (bottom). Note that we ran CSMC for a longer time to favor this reference method. CSMC performs well for clock trees and relaxed clock trees, whereas the annealed SMC works for all of the three types of trees and clearly outperforms CSMC for nonclock trees.

REAL DATA SETS

We analyzed two difficult real data sets from TreeBASE: M336 and M1809 in Table 1 of Lakner et al. (2008). M336 contains DNA sequences of length 1949 for 27 species. In M1809, there are 59 species and the length of each DNA sequence is 1824. We compared the marginal likelihood estimates, log-likelihood of the



FIGURE 9. Comparison of subsampling SMC algorithms with different size of batch sites.



FIGURE 10. Comparison of adaptive SMC algorithms with CSMC for three types of simulated trees: clock, relaxed clock, and nonclock (from top to bottom).

consensus tree, and tree distance metrics provided by ASMC and MrBayes (with the default setting) with the same computational budget. The reference trees used to compute tree distances are based on at least 6 independent long MrBayes parallel tempering runs provided by Lakner et al. (2008). Convergence to the posterior in these "reference runs" was established with high confidence in that previous work. Note that the comparison handicaps ASMC as the set of tree moves in MrBayes is a superset of those used in ASMC. The evolutionary model we consider in real data analysis is the JC69 model.

Data set M336

We used K = 500 and $\beta = 5.3$ for the ASMC algorithm. The log marginal likelihood estimated from ASMC is -7103.73, which is higher than the log marginal likelihood provided by MrBayes using SS (-7114.04). Table 5 displays the log-likelihood of the consensus tree and tree distance metrics provided by ASMC and MrBayes. In the table, *R* represents the number of annealing parameters in ASMC and the total number of MCMC iterations in MrBayes, respectively. The log-likelihood of the consensus tree estimated from ASMC is slightly lower than MrBayes. The RF and KF metrics

TABLE 5. Comparison of running ASMC and MrBayes for M336 from TreeBASE

| Method | R | Κ | Metric | Value |
|---------|---------|-----|----------------------|----------|
| ASMC | 15,706 | 500 | ConsensusLogLL | -6892.16 |
| | 15,706 | 500 | BestSampledLogLL | -6901.31 |
| | 15,706 | 500 | PartitionMetric | 0 |
| | 15,706 | 500 | RobinsonFouldsMetric | 0.01269 |
| | 15,706 | 500 | KuhnerFelsenstein | 5.55E-06 |
| MrBayes | 8.0E+06 | | ConsensusLogLL | -6889.52 |
| | 8.0E+06 | | PartitionMetric | 0 |
| | 8.0E+06 | | RobinsonFouldsMetric | 0.01832 |
| | 8.0E+06 | | KuhnerFelsenstein | 2.25E-5 |

estimated from MrBayes are slightly higher than ASMC. The majority-rule consensus tree provided by ASMC and MrBayes are identical, and coincide with the reference tree. Figure 11 displays the estimated majority-rule consensus trees and the clade posterior probabilities provided by ASMC and MrBayes. Most clades posterior probabilities provided by ASMC and MrBayes are close. ASMC provides lower posterior support for some clades, which is consistent with the hypothesized superior tree exploration provided by ASMC on a fixed budget.

Data set M1809

We used K = 1000 and $\beta = 5$ for the ASMC algorithm. The log marginal likelihood estimated from ASMC is -37,542.25, the one estimated by MrBayes using SS is -37,335.73. Table 6 displays the tree metrics provided by ASMC and MrBayes. The log-likelihood of the consensus tree provided by ASMC is higher than the one from MrBayes, and PM, RF, KF metrics estimated from ASMC are lower.

CONCLUSION AND DISCUSSION

The annealed SMC algorithm discussed in this article provides a simple but general framework for phylogenetic tree inference. Unlike previous SMC methods in phylogenetics, annealed SMC considers the same state space for all the intermediate distributions. As a consequence, many conventional Metropolis-Hastings tree moves used in the phylogenetic MCMC literature can be utilized as the basis of SMC proposal distributions. Because MCMC tree moves are available for a large class of trees, including nonclock as well as strict and relaxed clock models, the annealed SMC method is automatically applicable to a wide range of phylogenetic models. It should also be relatively easy to incorporate the proposed ASMC into existing phylogenetic software packages that implement MCMC algorithms, such as MrBayes, RevBayes, or BEAST.

The annealed SMC algorithm has two adaptive mechanisms, dynamic resampling and adaptive determination of annealing parameters, to make the algorithm efficient while requiring less tuning. Dynamic resampling based on ESS is a common practice in the SMC literature. Devising the annealing parameter sequence is a relatively newer practice (Del Moral et al. 2012). The annealing parameter sequence can be determined dynamically based on the CESS criterion. Because the particle weights of the current iteration only depend on the previous particles, there is negligible computational cost for finding annealing parameters.

The consistency of annealed SMC discussed in the theoretical results section holds when K goes to infinity. However, K cannot in practice be made arbitrarily large as the memory requirements scale linearly in K. In contrast, increasing the number of intermediate distributions R (in our adaptive algorithm, by increasing β) does not increase memory consumption. We conjecture that consistency for large R but fixed K also holds, in the sense of having the marginal distribution of each particle at the last iteration converging to the posterior distribution. We have explored the relative importance between K and R with fixed computational budgets using simulations. These results suggest that increasing R and K improves the approximation at different rates, with increasing R giving bigger bang to the buck. Assuming that our conjecture on the convergence in R holds true, in the regime of very large R and fixed K, the particle population at the last iteration can be conceptualized as K independent Monte Carlo samples from the true posterior distribution (in particular, we conjecture that the weights will convergence to a uniform distribution and hence to naive Monte Carlo based on independent exact samples). We remind the reader that the power of independent exact Monte Carlo is that the variance does not depend on the dimensionality of the problem. Hence if *R* is sufficiently large, a lower bound for *K* can therefore be obtained by selecting K large enough so that for independent samples X_i with distribution π the variance of the Monte Carlo average $(1/K)\sum_{k=1}^{K} f(X_i)$ is sufficiently small. Here is a concrete example: suppose we have a test function f of interest, for example an indicator function on a fixed clade, with unknown posterior support $p = \int f(x)\pi(x) dx$. We should take *K* large enough so that the Monte Carlo average will have a 95% Monte Carlo confidence interval having a width of no more than say min{p, 1-p}/10. For $p \le 1/2$, this yields $K \ge (p/10)^{-2}(z^*)^2 \operatorname{Var}_{\pi} f \approx 384(1-p)/p$, where $z^* \approx 1.96$ is the 95% critical value. For example, if the clade of interest is believed from a test run to be highly uncertain, $p \approx 1/2$, then in the large *R* regime, at the very minimum K = 400 particles should be used. The value of K should also be sufficiently large to accommodate the number of parallel cores available, and also to ensure that the adaptive annealing scheme is stable (i.e. that a further increase in K results in a qualitatively similar annealing schedule). See also Olsson and Douc (2017) for more sophisticated schemes for estimating the Monte Carlo variance of SMC algorithms.

Importantly, annealed SMC provides an efficient way to estimate the marginal likelihood, which is still a challenging task in Bayesian phylogenetics. We have also reviewed other marginal likelihood estimation methods, including SS and LIS. Our annealed SMC algorithm



FIGURE 11. The majority-rule consensus trees for the M336 data set estimated by (a) ASMC and (b) MrBayes. The numbers on the trees represent the clade posterior probabilities (number 100 is omitted).

TABLE 6. Comparison of running ASMC and MrBayes for M1809 from TreeBASE

| Method | R | Κ | Metric | Value |
|---------|------------------------------------------------|--------------------------------------|----------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| ASMC | 17,639 17,639 17,639 17,639 17,639 | 1000 1000 1000 1000 1000 | ConsensusLogLL BestSampledLogLL PartitionMetric RobinsonFouldsMetric KuhnerFelsenstein | -36,972.513 -36,991.443 2.0 0.13741 3.95E-4 |
| MrBayes | 1.76E+07 1.76E+07 1.76E+07 1.76E+07 | | ConsensusLogLL PartitionMetric RobinsonFouldsMetric KuhnerFelsenstein | -36,996.13 16.0 0.513285 0.01137 |

enjoys advantageous theoretical properties. The main property that justifies the use of the annealed SMC is the unbiasedness of its marginal likelihood estimate. In addition, the unbiasedness of the marginal likelihood estimate can be used to test implementation correctness of the algorithm. Our simulation studies have shown that ASMC can give a similar marginal likelihood estimate as the one obtained from the ASMC with the same but deterministic annealing parameter sequence (debiased ASMC). With the same computing budget, ASMC has been demonstrated to result in more accurate estimates. Moreover, the ASMC algorithm requires less tuning than the other methods considered. Both LIS and SS need a predetermined annealing parameter sequence, which is often inconvenient to choose in practice. ASMC leads to a more stable estimate for the marginal likelihood compared with the other methods considered. MCMC moves often come with tuning parameters. For example, proposal distributions typically have a bandwidth parameter which needs to be tuned (Roberts et al. 1997). To improve the performance of annealed SMC, it would be possible to use automatic tuning of proposal distributions within SMC algorithms, as proposed in Zhou et al. (2016).

A second future direction would be to investigate modifications in the specification of the sequence of intermediate distributions. For example, Fan et al. (2010) proposed an alternative to the prior distribution to replace π_0 . The same choice could be used within our framework. In another direction, it may be possible to combine the construction of Dinh et al. (2017) with ours to handle online problems via standard moves: instead of integrating the new taxon with all of its sites unannealed (which requires specialized proposals), it may be beneficial to anneal the newly introduced site.

In terms of empirical comparisons, it would be interesting to expand the set of metrics, models and data sets used to compare the algorithms. For example, in addition to the tree distance metrics used in this article, geodesic tree distance (Billera et al. 2001) is also an important metric to compare distances between phylogenetic trees. The GTP software (Owen and Provan 2011) allows easy calculation of the geodesic tree distance.

We have investigated the subsampling SMC algorithm for "tall data" phylogenetic problems. The annealing parameters $\psi(s, \phi_r)$ of the subsampling SMC is derived from the annealing parameter sequence ϕ_r of the adaptive ASMC without subsampling. This choice was made for implementation convenience, and there is no reason why the two optimal sequences of distributions should coincide. To improve the algorithm performance, one direction is therefore to design an adaptive scheme tailored to the subsampling version. One challenge is that pointwise evaluation of the adaptation function $g(\phi)$ is more expensive in the subsampling setup, with a cost that grows with ϕ . Bayesian optimization might be useful in this context. Another use of the subsampling arises in situations where the sampling algorithm is not constructed using an accept-reject step. For example, conjugate Gibbs sampling on an augmented target distribution (Lartillot 2006) is used by PhyloBayes (Lartillot et al. 2009) to efficiently sample the evolutionary model parameters. It is not clear how the conjugate Gibbs sampling step can be modified to accommodate the annealed distribution in Equation (4) for $\phi < 1$. On the other hand, conjugate sampling is directly applicable to intermediate distributions that consist in taking subsets of sites. This sequence of distributions could be used to handle conjugate Gibbs sampling not only in annealed SMC but also in the context of parallel tempering or any other sequence of measure-based method. One last line of work is to combine control variates to annealed SMC to reduce the variance of the likelihood estimator, a general strategy that has been very successful in other subsampling work (Bardenet et al. 2017).

Funding

This research was supported by Discovery Grants of A. Bouchard-Côté and L. Wang from the National Science and Engineering Research Council and a Canadian Statistical Sciences Institute Collaborative Research Team Project. Research was enabled in part by support provided by WestGrid (www.westgrid.ca) and Compute Canada (www.computecanada.ca).

ACKNOWLEDGMENTS

We would like to thank the reviewers and editor for the constructive feedback, as well as Arnaud Doucet and Fredrik Ronquist for helpful discussions. Part of this work was done while visiting the Department of Statistics at Oxford.

APPENDIX 1

Construction of Intermediate Distributions for Subsampling SMC

Let us decompose the unnormalized posterior distribution as

$$\gamma_1(x) = p(x) \prod_{s=1}^{\#S} p(y_s|x),$$

where *x* refers to the phylogenetic tree and evolutionary parameter of interest, *s* is an index for one batch of sites from a biological sequence, and #*S* represents the total number of batches. Each batch contains one or more sites of the biological sequence; we denote the number of sites in each batch by b_s .

Consider the annealing parameter sequence $0 = \phi_0 < \phi_1 < \cdots < \phi_R = 1$. We define the sequence of intermediate distributions for subsampling as follows:

$$\gamma_{\phi_r}(x) = p(x) \prod_{s=1}^{\#S} p(y_s|x)^{\psi(s,\phi_r)},$$

where

$$\psi(s,\phi_r) = \begin{cases} 1 & \text{if } \phi_r \ge s/\#S, \\ 0 & \text{if } \phi_r \le (s-1)/\#S, \\ \#S \cdot \phi_r - (s-1) & \text{otherwise.} \end{cases}$$

The subsampling SMC algorithm is a more general version of the annealed SMC algorithm. If we define #S=1 in $\gamma_1(x)$, then the sequence of intermediate distributions of subsampling SMC is exactly the same as the intermediate distributions of the annealed SMC. In this case, the computational cost of subsampling SMC is exactly the same as the annealed SMC. Another extreme case is that #S=n, in which case we sequentially incorporate the sites of sequence one by one.

Appendix 2

Theoretical Foundations of Annealed SMC

In this section, we review the construction of Del Moral et al. (2006), which is the basis for our work. See also Wang et al. (2015) for a similar construction tailored to a phylogenetic setup.

The corresponding sequence of un-normalized distributions are denoted by $\{\gamma_r\}_{1,...,R}$. The annealed SMC can be obtained by defining an auxiliary sequence of distributions that admit the distribution of interest, $\pi_r(x_r)$, as the marginal of the latest iteration

$$\tilde{\pi}_r(\mathbf{x}_r) = \pi_r(\mathbf{x}_r) \prod_{j=1}^{r-1} L_j(x_{j+1}, x_j),$$

where $L_j(x_{j+1}, x_j)$ is an auxiliary "backward" Markov kernel with $\int L_j(x_{j+1}, x_j) dx_j = 1$. We never sample from L_j , rather its role is to allow us to derive weight updates that yield a valid SMC algorithm.

The idea is then to apply standard SMC (i.e. SMC for product spaces such as state space models) to this auxiliary sequence of distributions, $\tilde{\pi}_1, \tilde{\pi}_2, ..., \tilde{\pi}_R$. The resulting sampler has a weight update given by

$$w(x_{r-1}, x_r) \propto \frac{\tilde{\pi}_r(x_r)}{\tilde{\pi}_r(x_{r-1})} \frac{1}{K_r(x_{r-1}, x_r)} = \frac{\pi_r(x_r)L_{r-1}(x_r, x_{r-1})}{\pi_{r-1}(x_{r-1})K_r(x_{r-1}, x_r)},$$

which is different from the one in a standard SMC.

When K_r satisfies global balance with respect to π_r , a convenient backward Markov kernel that allows an easy evaluation of the importance weight is

$$L_{r-1}(x_r, x_{r-1}) = \frac{\pi_r(x_{r-1})K_r(x_{r-1}, x_r)}{\pi_r(x_r)}.$$

This choice is a properly normalized backward kernel, $\int L_{r-1}(x_r, x_{r-1}) dx_{r-1} = 1$: this follows from the assumption that K_r satisfies global balance with respect to π_r . With this backward kernel, the incremental importance weight becomes

$$w(x_{r-1}, x_r) = \frac{\gamma_r(x_r)}{\gamma_{r-1}(x_{r-1})} \cdot \frac{L_{r-1}(x_r, x_{r-1})}{K_r(x_{r-1}, x_r)}$$
$$= \frac{\gamma_r(x_r)}{\gamma_{r-1}(x_{r-1})} \cdot \frac{\pi_r(x_{r-1})K_r(x_{r-1}, x_r)}{\pi_r(x_r)}$$
$$\cdot \frac{1}{K_r(x_{r-1}, x_r)} = \frac{\gamma_r(x_{r-1})}{\gamma_{r-1}(x_{r-1})}.$$

General Estimates of Marginal Likelihood

In Basic Annealed SMC Algorithm section, we describe the estimator for marginal likelihood in a simplified setting, that is without adaptation. Here, we describe the marginal likelihood estimator in full generality.

Recall that we denote the marginal likelihood by *Z* for simplicity. With a slight abuse of notation, we use $K_r(x_{r-1}, \cdot)$ to denote the proposal distribution for x_r in this section.

Let us start by rewriting the normalization constant of the first intermediate distribution as

$$Z_1 = \int \frac{\gamma_1(x_1)}{K_1(x_1)} K_1(x_1) dx_1 = \int w_1(x_1) K_1(x_1) dx_1$$

where $K_1(\cdot)$ is the proposal distribution for x_1 . Correspondingly, an estimate of Z_1 is

$$Z_{1,K} = \frac{1}{K} \sum_{k=1}^{K} w_{1,k}.$$

Similarly, we can rewrite the ratio of the normalization constants of two intermediate distributions as

$$\begin{aligned} \frac{Z_r}{Z_{r-1}} &= \frac{\int \gamma_r(x_r) dx_r}{Z_{r-1}} = \frac{\int \gamma_r(x_r) dx_r}{\gamma_{r-1}(x_{r-1})/\pi_{r-1}(x_{r-1})} \\ &= \int \frac{\gamma_r(x_r)}{\gamma_{r-1}(x_{r-1})} \pi_{r-1}(x_{r-1}) dx_r \\ &= \int \frac{\gamma_r(x_r)}{\gamma_{r-1}(x_{r-1})K_r(x_{r-1},x_r)} \\ &= \pi_{r-1}(x_{r-1})K_r(x_{r-1},x_r) dx_r \\ &= \int w_r(x_r)\pi_{r-1}(x_{r-1})K_r(x_{r-1},x_r) dx_r. \end{aligned}$$

Straightforwardly, an estimate of Z_r/Z_{r-1} is provided by

$$\frac{\widehat{Z_r}}{Z_{r-1}} = \frac{1}{K} \sum_{k=1}^{K} w_{r,k}$$

Since the estimate of the marginal likelihood can be rewritten as

$$Z \equiv Z_R = Z_1 \prod_{r=2}^R \frac{Z_r}{Z_{r-1}},$$

an estimate of the marginal likelihood Z is

$$\hat{Z}_{R,K} = \prod_{r=1}^{R} \left(\frac{1}{K} \sum_{k=1}^{K} w_{r,k} \right)$$
$$= \prod_{r=1}^{R} \left(\frac{1}{K} \sum_{k=1}^{K} \{ p(y|x_{r-1,k}) \}^{\phi_r - \phi_{r-1}} \right), \quad (19)$$

which can be obtained from an SMC algorithm readily. If resampling is not conducted at each iteration r, an

alternative form is provided by

$$\widehat{Z_{R,K}} = \prod_{j=1}^{t_{R-1}+1} \left(\sum_{k=1}^{K} W_{n_{j-1},k} \prod_{m=n_{j-1}+1}^{n_j} \{ p(y|x_{m-1,k}) \}^{\phi_m - \phi_{m-1}} \right), \quad (20)$$

where n_j is the SMC iteration index at which we do the *j*th resampling, t_{R-1} is the number of resampling steps between 1 and R-1.

Unbiasedness, Consistency and Central Limit Theorem for Annealed SMC

Here, we provide more information on the theoretical properties discussed in Properties of Annealed Sequential Monte Carlo section.

Theorem 1 (*Unbiasedness*): For fixed $0 = \phi_0 < \phi_1 < \cdots < \phi_R = 1$, $\hat{Z}_{R,K}$ is an unbiased estimate of *Z*,

$$\mathbb{E}(\hat{Z}_{R,K})=Z.$$

This result is well known in the literature, although many statements of the result are specialized to SMC for state-space models (Doucet and Johansen 2011). The results in *Theorem 7.4.2* of Del Moral (2004) provide a very general set of conditions which includes the annealed SMC algorithm presented here. However, the theoretical framework in Del Moral (2004) being very general and abstract, we outline below an alternative line of argument to establish unbiasedness of phylogenetic annealed SMC.

First, by the construction reviewed in Theoretical Foundations of Annealed SMC section, we can transform the sequence of distributions on a fixed state space, $\pi_r(x_r)$, into a sequence of augmented distributions $\tilde{\pi}_r(x_r)$ on a product space admitting $\pi_r(x_r)$ as a marginal. We now apply Theorem 2 of Andrieu et al. (2010) with the distribution $\pi_n(x_{1:n})$ in this reference set to $\tilde{\pi}_r(x_r)$ in our notation. To be able to use Theorem 2, we only need to establish the "minimum assumptions" 1 and 2 in Andrieu et al. (2010). Assumption $\overline{1}$ is satisfied by the fact that valid MCMC proposals are guaranteed to be such that $q(x, x') > 0 \iff q(x', x) > 0$. Assumption 2 holds because we use multinomial resampling. Next, since the conditions of Theorem 2 of Andrieu et al. (2010) hold, we have the following result from the proof of Theorem 2 in Appendix B1 of Andrieu et al. (2010):

$$\frac{\tilde{\pi}^{N}(k,\bar{x}_{1},...,\bar{x}_{P},a_{1},...,a_{P-1})}{q^{N}(k,\bar{x}_{1},...,\bar{x}_{P},a_{1},...,a_{P-1})} = \frac{\hat{Z}^{N}(\bar{x}_{1},...,\bar{x}_{P})}{Z}$$

and hence

$$Z \tilde{\pi}^{N}(k, \bar{x}_{1}, ..., \bar{x}_{P}, a_{1}, ..., a_{P-1}) = \hat{Z}^{N}(\bar{x}_{1}, ..., \bar{x}_{P}) q^{N}(k, \bar{x}_{1}, ..., \bar{x}_{P}, a_{1}, ..., a_{P-1}).$$

Now taking the integral on all variables $k, \bar{x}_1, ..., \bar{x}_P, a_1, ..., a_{P-1}$ with respect to the reference measure μ associated to $\bar{\pi}^N$, we obtain:

$$Z \int \tilde{\pi}^{N}(k, \bar{x}_{1}, ..., \bar{x}_{P}, a_{1}, ..., a_{P-1}) d\mu(k, \bar{x}_{1}, ..., \bar{x}_{P}, a_{1}, ..., a_{P-1})$$
$$= \int \hat{Z}^{N}(\bar{x}_{1}, ..., \bar{x}_{P}) q^{N}(k, \bar{x}_{1}, ..., \bar{x}_{P}, a_{1}, ..., a_{P-1})$$
$$d\mu(k, \bar{x}_{1}, ..., \bar{x}_{P}, a_{1}, ..., a_{P-1}).$$

The left-hand side is just *Z* because $\tilde{\pi}^N$ is a density with respect to μ . For the right-hand side, note that q^N is the law of the full set of states produced by the particle filter, hence the right-hand side is just $\mathbb{E}[\hat{Z}_{R,K}]$ in our notation. This concludes the proof.

Next, we discuss consistency results. In the SMC literature, they are generally available both in the L^2 convergence and almost sure convergence flavors. We cover the L^2 case here and refer to Del Moral et al. (2006) for almost sure consistency results.

Theorem 2 (*Consistency*): Assume there is a constant *C* such that $|f| \le C$ and $w_{r,k} \le C$ almost surely. For a fixed ϕ_r (r=1,...,R), the annealed SMC algorithm provides asymptotically consistent estimates:

$$\sum_{k=1}^{K} W_{r,k} f(x_{r,k}) \to \int \pi_r(x) f(x) dx \text{ as } K \to \infty,$$

where the convergence holds in the L^2 norm sense.

The result can be deduced from Proposition 5 in Wang et al. (2015) as follows. Assumption 3 in Wang et al. (2015) holds because MCMC proposals satisfy $q(x,x') > 0 \iff q(x',x) > 0$. Assumption 4 holds because the support of the prior coincides with support of the posterior.

Note that the assumption that the weights are bounded is not valid for general tree spaces. However, if the branch lengths are assumed to be bounded then the space is compact and the assumption therefore holds in that setting.

Finally, we turn to the central limit theorem.

Theorem 3 (*Central Limit Theorem*): Under the integrability conditions given in Theorem 1 of Chopin et al. (2004), or Del Moral (2004), section 9.4, pages 300–306,

$$K^{1/2} \left[\sum_{k=1}^{K} W_{r,k} f(x_{r,k}) - \int \pi_r(x) f(x) dx \right] \to N(0, \sigma_r^2(f))$$

as $K \to \infty$,

where the convergence is in distribution. The form of asymptotic variance $\sigma_r^2(f)$ depends on the resampling scheme, the Markov kernel K_r and the artificial backward kernel L_r . We refer readers to Del Moral et al. (2006) for details of this asymptotic variance.

177

APPENDIX 3

MCMC Proposals for Bayesian Phylogenetics

In this article, we used the proposals q_r^i defined as follow:

- 1. q_r^1 : the *multiplicative branch proposal*. This proposal picks one edge at random and multiply its current value by a random number distributed uniformly in [1/*a*,*a*] for some fixed parameter *a* > 1 (controlling how bold the move is) (Lakner et al. 2008).
- 2. q_r^2 : the global multiplicative branch proposal that proposes all the branch lengths by applying the above multiplicative branch proposal to each branch.
- 3. q_r^3 : the *stochastic NNI proposal*. We consider the nearest neighbor interchange (NNI) (Jow et al. 2002) to propose a new tree topology.
- 4. q_r^4 : the *stochastic NNI proposal with resampling the edge* that uses the above NNI proposal in (3) and the multiplicative branch proposal in (1) for the edge under consideration.
- 5. q_r^5 : the *Subtree Prune and Regraft (SPR) move* that selects and removes a subtree from the main tree and reinserts it elsewhere on the main tree to create a new tree.

Note that here, we only describe the MCMC kernels for phylogenetic trees. To sample evolutionary parameters θ , one can use simple proposals such as symmetric Gaussian distributions, or more complex ones, see for example Zhao et al. (2016).

APPENDIX 4

Comparison of Resampling Strategies

In this section, we compare the performance of adaptive annealed SMC with three different resampling thresholds. The first resampling threshold is $\epsilon_1 = 0$. In this case, the particles are never resampled. The second resampling threshold is $\epsilon_2 = 1$, in which case resampling is triggered at every iteration. The third resampling threshold is $\epsilon_3 = 0.5$. The resampling method we used in all our experiments was the multinomial resampling scheme. We simulated one unrooted tree of 15 taxa, and generated one data set of DNA sequences of length 200. The tree simulation setup was the same as in Simulation Studies section. We ran adaptive annealed SMC algorithm 20 times with the three resampling thresholds described above. We used $rCESS_r = 0.99999$ and K = 100. Figure A.1 demonstrates the advantage of resampling triggered by a threshold of $\epsilon_3 = 0.5$ over the other two choices by displaying the marginal likelihood estimates, log-likelihood of the consensus tree and tree metrics provided by adaptive annealed SMC using three different ϵ 's. The log marginal likelihood estimate and log-likelihood of the consensus tree provided by adaptive annealed SMC using $\epsilon_3 = 0.5$ are higher and admit smaller variation. The PF, RF, and KF metrics provided by adaptive annealed SMC using $\epsilon_3 = 0.5$ are lowest. Therefore the threshold 0.5 has been used in the rest of the article.

APPENDIX 5

Review of Particle Degeneracy Measures

The two adaptive schemes in ASMC, adaptively conducting resampling and the automatically construction of the annealing parameter sequence, rely on being able to assess the quality of a particle approximation. For completeness, we provide some background in this section on the classical notation of ESS and of CESS, a recent generalization which we use here (Zhou et al. 2016). The notion of ESS in the context of importance sampling (IS) or SMC is distinct from the notion of ESS in the context of MCMC. The two are related in the sense of expressing a variance inflation compared with an idealized Monte Carlo scheme but they differ in the details. We will assume from now on that ESS refers to the SMC context.

We will also use a slight variation of the derivation of ESS and CESS where the measures obtained are normalized to be between zero and one (some hyperparameters of the adaptive algorithms are easier to express in this fashion). We use the terminology relative (conditional) ESS to avoid confusion.

The fundamental motivation of (relative and/or conditional) ESS stems from the analysis of the error of Monte Carlo estimators. Recall that for a given function of interest f (think for example of f being an indicator function on a clade),

$$I = \int \pi_r(\mathrm{d}x) f(x) \approx \sum_{k=1}^K W_{r,k} f(x_{r,k}) =: \hat{I}.$$

The quantity on the right-hand side is a random variable (with respect to the randomness of the SMC algorithm), \hat{I} , and we can think about it as an estimator of the deterministic quantity on the left-hand side, I. Moreover the right-hand side is a real-valued random variable, so we can define its mean square error, which can be further decomposed as a variance term and a squared bias term. For SMC algorithms, the variance term dominates as the number of particles goes to infinity (Del Moral 2004). For this reason, we are interested in estimates of the variance of \hat{I} across SMC random seeds, Var_{SMC}[\tilde{I}]. However, the variance of I depends on the choice of function f, which is problem dependent, and we would like to remove this dependency. The first step is to consider a notion of relative variance, comparing to the variance we would obtain from a basic Monte Carlo scheme \hat{I}^* relying on iid exact samples $x_1^{\star}, \dots, x_K^{\star} \sim \pi$, $\operatorname{Var}_{\mathrm{MC}}[\hat{I}^*] = \operatorname{Var}[\frac{1}{K}\sum_k f(x_k^{\star})].$



FIGURE A.1. Comparison of three resampling thresholds, 0, 0.5, and 1.

To make further progress, we will make approximations of the ratio $Var_{MC}[\hat{I}^*]/Var_{SMC}[\hat{I}]$.

To understand these approximations, let us start with a simplified version of Algorithm 3, where the function NextAnnealingParameter returns the value 1.0. In this setting, no resampling occurs, and the algorithm reduces to an IS algorithm (more specifically, it reduces to a single iteration of the AIS algorithm; Neal 2001). IS is easier to analyze because the individual particles are independent and identically distributed, allowing us to summarize the behavior based on one particle, say k=1. If we assume further that (A1) $\gamma_i = \pi_i$, that is, the normalization constant is one, then a classical argument by Kong (1992) based on the Delta method yields

$$\frac{\operatorname{Var}_{MC}[I^*]}{\operatorname{Var}_{SMC}[\hat{I}]} = \frac{\operatorname{Var}_{MC}[I^*]}{\operatorname{Var}_{\pi_0}[\hat{I}]} \approx \frac{1}{1 + \operatorname{Var}_{\pi_0}[\tilde{w}_{1,1}]},$$
$$= \frac{1}{\mathbb{E}_{\pi_0}[(\tilde{w}_{1,1})^2]},$$

where we used the fact that in this simple setting the distribution of one proposed particle is just π_0 , so $Var_{SMC}[\cdot] = Var_{\pi_0}[\cdot]$ and in the last line,

$$\mathbb{E}_{\pi_0}[\tilde{w}_{1,1}] = \int \pi_0(x_{0,1}) \frac{\pi_1(x_{0,1})}{\pi_0(x_{0,1})} dx_{0,1} = 1.$$

In general, assumption (A1) does not hold, that is the normalization constant is not one, so for a general onestep AIS algorithm we get instead the approximation:

$$\frac{\operatorname{Var}_{\mathrm{MC}}[\hat{I}^*]}{\operatorname{Var}_{\mathrm{SMC}}[\hat{I}]} \approx \left(\mathbb{E}_{\pi_0} \left[\left(\frac{\pi_1(x_{0,1})}{\pi_0(x_{0,1})} \right)^2 \right] \right)^{-1} \\ = \left(\mathbb{E}_{\pi_0} \left[\left(\frac{\gamma_1(x_{0,1})/Z_1}{\gamma_0(x_{0,1})/Z_0} \right)^2 \right] \right)^{-1} \\ = \left(\frac{Z_1}{Z_0} \right)^2 / \mathbb{E}_{\pi_0} \left[\left(\frac{\gamma_1}{\gamma_0}(x_{0,1}) \right)^2 \right].$$

Generalizing the notation of this section into a general SMC setup, π_1 here plays the role of the current iteration,

and π_0 , of the previous iteration. However, since π_0 is not known in this case, we plug-in a particle approximation $\hat{\pi}_0 = \sum_{k=1}^{K} W_{0,k} \delta_{x_{0,k}}$ to get:

$$\frac{\operatorname{Var}_{\mathrm{MC}}[\hat{I}^*]}{\operatorname{Var}_{\mathrm{SMC}}[\hat{I}]} \approx \left(\frac{Z_1}{Z_0}\right)^2 / \mathbb{E}_{\hat{\pi}_0} \left[\left(\frac{\gamma_1}{\gamma_0}(x_{0,1})\right)^2 \right]$$
$$= \left(\frac{Z_1}{Z_0}\right)^2 / \sum_{k=1}^K W_{0,k} \left(\frac{\gamma_1}{\gamma_0}(x_{0,k})\right)^2.$$

The effect of this additional approximation is that it makes our estimator over-optimistic, by ignoring the error of the approximation $\hat{\pi}_0$ of π_0 . It is nonetheless a useful tool to assess the degradation of performance over a small number of SMC iterations.

Finally, since the ratio of normalization constants is also unknown, we also need to estimate it. Based on a particle approximation of Equation (14), we obtain:

$$\frac{\operatorname{Var}_{\mathrm{MC}}[\hat{I}^*]}{\operatorname{Var}_{\mathrm{SMC}}[\hat{I}]} \approx \left(\sum_{k=1}^{K} W_{0,k} \frac{\gamma_1}{\gamma_0}(x_{0,k})\right)^2 / \sum_{k=1}^{K} W_{0,k} \left(\frac{\gamma_1}{\gamma_0}(x_{0,k})\right)^2.$$

This quantity is called the relative CESS (rCESS), Equation (11). Having a high rCESS value is a necessary but not sufficient condition for a good SMC approximation. If it is low during some SMC iteration, especially an iteration close to the final iteration, then with high probability most of the particles will have very small or zero weights, which will lead to a collapse of the quality of the annealed SMC algorithm.

Comparison of rESS and rCESS

In earlier work on adaptive SMC methods, the function NextAnnealingParameter was implemented using a different criterion based on rESS instead of rCESS. Later, Zhou et al. (2016) argued that rCESS was more appropriate. Here, we confirm that this is also the case in a phylogenetic context. We provide two experiments. In the first experiment, we simulated one unrooted tree of 10 taxa, and generated one data set of DNA sequences



FIGURE A.2. Comparison of rCESS (left) and rESS (right) in terms of rESS as a function of r (top) and $\phi_r - \phi_{r-1}$ as a function of ϕ_r (bottom).

and each sequence has length 100. The setup of tree simulation is the same as Simulation Studies section. We ran adaptive annealed SMC algorithm in two schemes: (1) $rCESS_r = 0.99$; (2) $rESS_r = 0.99$. The one based on *rESS_r* only differs in the way NextAnnealingParameter is implemented; shown in Algorithm 5. We used K =1000 particles. Resampling of particles was triggered when rESS < 0.5. Figure A.2 demonstrates the advantage of using rCESS over rESS in adaptive annealed SMC. The annealing parameter difference $(\phi_r - \phi_{r-1})$ increases smoothly in the rCESS scheme, whereas in the rESS scheme there are big gaps in annealing parameter increment after doing resampling, then the consecutive annealing parameter change decreases gradually until the next resampling time. The number of iterations R for adaptive annealed SMC using rESS is much larger than using rCESS.

In our second experiment, we compared the performance of adaptive annealed SMC using rCESS and rESS in terms of tree metrics and marginal likelihood. We simulated one unrooted tree of 15 taxa, and generated one data set of DNA sequences. Each sequence has length 200. The tree simulation setup is the same as Simulation Studies section. We ran adaptive annealed SMC algorithm 20 times with $rCESS_r = 0.999$ and $rESS_r = 0.978$, respectively. The number of particles was set to K = 500. Under this setting, the computational costs of the two schemes are quite similar. The numbers of annealing parameters (*R*) selected via rCESS and

Algorithm 5 Alternative NextAnnealingParameter procedure (suboptimal)

- Inputs: (a) Particle population from previous SMC iteration (*x*_{r-1,·}, *w*_{r-1,·}); (b) Annealing parameter φ_{r-1} of previous SMC iteration; (c) A degeneracy decay target α∈(0,1).
- 2. **Outputs:** automatic choice of annealing parameter ϕ_r .
- Initialize the function *ğ* assessing the particle population quality associated to a putative annealing parameter φ:

$$\tilde{g}(\phi) = \left(\sum_{k=1}^{K} W_{r-1,k} p(y|x_{r-1,k})^{\phi-\phi_{r-1}}\right)^{2} / \sum_{k=1}^{K} (W_{r-1,k} p(y|x_{r-1,k})^{\phi-\phi_{r-1}})^{2},$$
(1) > $g(\phi_{r-1})$ then

4. if $\tilde{g}(1) \ge \alpha \tilde{g}(\phi_{r-1})$ then

5. return $\phi_r = 1$.

- 6. else
- 7. return $\phi_r = \phi^* \in (\phi_{r-1}, 1)$ such that $\tilde{g}(\phi^*) = \alpha \tilde{g}(\phi_{r-1})$ via bisection.

rESS are similar. Figure A.3 displays the marginal likelihood estimates, consensus likelihood and tree metrics provided by adaptive SMC using rCESS and

179

method 📥 rCESS (0.999) 📥 rESS (0.978)



FIGURE A.3. Comparison of adaptive annealed SMC using rCESS and rESS in terms of estimating the log marginal likelihood, the log-likelihood of the consensus tree, tree distance metrics, and the number of SMC iterations (*R*).

rESS. The log marginal likelihood estimates and log consensus likelihoods provided by adaptive SMC using rCESS are higher and have lower variability. The PF, RF, and KF metrics provided by the two schemes are quite close, whereas the metrics provided by rCESS scheme have lower variability.

APPENDIX 6

Estimates of Marginal Likelihood from LIS

We described the LIS procedure as follows:

- 1. Sample an index v_0 randomly from $\{1, 2, ..., N\}$, and sample $x_{0,v_1} \sim \pi_0(\cdot)$.
- 2. For d=0,1,...,D, sample *N* states from π_d as follows:
 - (a) If d > 0: sample an index v_d from $\{1, 2, ..., N\}$, and set $x_{d,v_d} = x_{d-1*d}$.
 - (b) For $k = v_d + 1, ..., N$, sample $x_{d,k}$ from the forward kernel $x_{d,k} \sim K_d(x_{d,k-1}, \cdot)$.
 - (c) For $k = v_d 1, ..., 1$, sample $x_{d,k}$ from the backward kernel $x_{d,k} \sim L_d(x_{d,k+1}, \cdot)$.
 - (d) If d < D, sample μ_d from $\{1, 2, ..., N_d\}$ according to the following probabilities:

$$p(\mu_d|x_d) = \frac{\gamma_{d-1*d}(x_{d,\mu_d})}{\gamma_d(x_{d,\mu_d})} \bigg/ \sum_{k=1}^{N_d} \frac{\gamma_{d-1*d}(x_{d,k})}{\gamma_d(x_{d,k})},$$

and set x_{d*d+1} to x_{d,μ_d} .

3. Compute the likelihood estimate

$$\hat{Z}_{LIS} = \prod_{d=1}^{D} \left[\frac{1}{N} \sum_{k=1}^{N} \frac{\gamma_{d-1*d}(x_{d-1,k})}{\gamma_{d-1}(x_{d-1,k})} \right] \frac{1}{N} \sum_{k=1}^{N} \frac{\gamma_{d-1*d}(x_{d,k})}{\gamma_{d}(x_{d,k})} .$$

Note that if the backward kernel is reversible, then the forward kernel is the same as backward kernel. In this article, we use the MCMC kernel as backward and forward kernels in LIS.

Appendix 7

Derivation of Upper Bound of CV

$$CV = \frac{sd(\hat{Z})}{\mathbb{E}(\hat{Z})}$$

= $\frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{Z}_{i} - \frac{1}{n}\sum_{i=1}^{n}\hat{Z}_{i})^{2}}}{\frac{1}{n}\sum_{i=1}^{n}\hat{Z}_{i}}$
= $\sqrt{n}\sqrt{\sum_{i=1}^{n}(\frac{\hat{Z}_{i}}{\sum_{i=1}^{n}\hat{Z}_{i}} - \frac{1}{n})^{2}}$

For non-negative Z_i , the CV is maximized when $\frac{\hat{Z}_i}{\sum_{i=1}^n \hat{Z}_i} = 1$ for some *i*, and 0 for the rest. In this extreme case, \hat{Z}_i is much larger than the rest. The upper bound of the CV can be simplified to $\sqrt{n-1}$.

APPENDIX 8

Tuning of β and K

In Figure A.4, we compare the performance of ASMC algorithm as a function of *K*, with β fixed at 5. We used four different numbers of particles *K* = 100, 300, 1000, 3000. Both the marginal likelihood estimate and tree metrics improve as *K* increases. Figure A.5 displays the performance of ASMC algorithm as a function of β , with *K* = 1000. *R* is the total number of SMC iterations. We used five distinct β values, β = 3,4,4.3,5,5.3. The marginal likelihood estimates and tree metrics improve as β increases; they tend to be stable after β reaches 5. A larger value of β can improve the performance of ASMC more significantly than an increase in *K*.



FIGURE A.4. Comparison of adaptive SMC algorithm with different numbers of particles, from left to right K = 100, 300, 1000, 3000.



FIGURE A.5. Comparison of adaptive SMC algorithm with different β , from left to right $\beta = 3, 4, 4.3, 5, 5.3$. Here, *R* is the total number of SMC iterations.

APPENDIX 9

Comparison of ASMC, DASMC, LIS, and SS for large K

In this experiment, we focus on evaluating the marginal likelihood estimates using ASMC, DASMC, LIS, and SS with a shared, large computational budget. We simulated an unrooted tree of 4 taxa, generated one data set of DNA sequences of length 10. Every algorithm for each data set was repeated 50 times with different random seeds. We set $\beta=2$ and K=200,000. The setup of DASMC, LIS, and SS is the same as Comparison of Marginal Likelihood Estimates section.

Figure A.6 shows the comparison of the performance of the four algorithms in terms of the marginal likelihoods in the log scale. The mean log marginalized likelihood estimates provided by ASMC, DASMC, LIS, and SS are quite close. The variance of estimates for ASMC and DASMC is smaller than LIS and SS.

References

- Altekar G., Dwarkadas S., Huelsenbeck J., Ronquist F. 2004. Parallel Metropolis coupled Markov chain Monte Carlo for Bayesian phylogenetic inference. Bioinformatics 20:407–415.
- Andrieu C., de Freitas N., Doucet A., Jordan M.I. 2003. An introduction to MCMC for machine learning. Mach. Learn. 50:5–43.



FIGURE A.6. Comparison of marginal likelihood (in log scale) provided by ASMC, DASMC, LIS, and SS with a fixed computational budget when K = 200000.

- Andrieu C., Doucet A., Holenstein R. 2010. Particle Markov chain Monte Carlo methods. J. R. Stat. Soc. Series B Stat. Methodol. 72:269–342.
- Andrieu C., Roberts G.O. 2009. The pseudo-marginal approach for efficient Monte Carlo computations. Ann. Stat. 37:697–725.
- Atchadé Y.F., Roberts G.O., Rosenthal J.S. 2011. Towards optimal scaling of metropolis-coupled Markov chain Monte Carlo. Stat. Comput. 21:555–568.

- Bardenet R., Doucet A., Holmes C. 2017. On Markov chain Monte Carlo methods for tall data. J. Mach. Learn. Res. 18:1515–1557.
- Billera L.J., Holmes S.P., Vogtmann K. 2001. Geometry of the space of phylogenetic trees. Adv. Appl. Math. 27:733–767.
- Bouchard-Côté A., Sankararaman S., Jordan M.I. 2012. Phylogenetic inference via sequential Monte Carlo. Syst. Biol. 61:579–593.
- Chan H.P., Lai T.L. 2013. A general theory of particle filters in hidden Markov models and some applications. Ann. Stat. 41:2877–2904.
- Chen M.-H., Kuo L., Lewis P.O. 2014. Bayesian phylogenetics: methods, algorithms, and applications. Boca Raton, FL: CRC.
- Chopin N. 2004. Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. Ann. Stat. 32:2385–2411.
- Del Moral P. 2004. Feynman-Kac formulae: genealogical and interacting particle systems with applications. New York: Springer.
- Del Moral P., Doucet A., Jasra A. 2006. Sequential Monte Carlo samplers. J. R. Stat. Soc. Series B Stat. Methodol. 68:411–436.
- Del Moral P., Doucet A., Jasra A. 2012. An adaptive sequential Monte Carlo method for approximate Bayesian computation. Stat. Comput. 22:1009–1020.
- Devroye L. 1986. Non-uniform random variate generation. New York: Springer-Verlag.
- Dinh V., Darling A.E., Matsen IV F.A. 2017. Online Bayesian phylogenetic inference: theoretical foundations via Sequential Monte Carlo. Syst. Biol. 67:503–517.
- Douc R., Cappé O. 2005. Comparison of resampling schemes for particle filtering. In: ISPA 2005. Proceedings of the 4th international symposium on image and signal processing and analysis, 2005. IEEE, Zagreb, Croatia, Croatia. p. 64–69.
- Doucet A., de Freitas N., Gordon N. 2001. Sequential Monte Carlo methods in practice. New York: Springer.
- Doucet A., Johansen A.M. 2011. A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later. In: Crisan, D. and Rozovskii, B., editors. The Oxford Handbook of Nonlinear Filtering. New York: Oxford University Press. pp. 656–704.
- Drummond A., Rambaut A. 2007. BEAST: Bayesian evolutionary analysis by sampling trees. BMC Evol. Biol. 7:214.
- Everitt R.G., Culliford R., Medina-Aguayo F., Wilson D.J. 2016. Sequential Bayesian inference for mixture models and the coalescent using sequential Monte Carlo samplers with transformations. Technical report, arXiv:1612.06468.
- Fan Y., Wu R., Chen M.-H., Kuo L., Lewis P.O. 2010. Choosing among partition models in Bayesian phylogenetics. Mol. Biol. Evol. 28:523– 532.
- Felsenstein J. 1973. Maximum likelihood and minimum-steps methods for estimating evolutionary trees from data on discrete characters. Syst. Biol. 22:240–249.
- Felsenstein J. 1981. Evolutionary trees from DNA sequences: a maximum likelihood approach. J. Mol. Evol. 17:368–376.
- Forney G.D. 1973. The viterbi algorithm. Proc. IEEE 61:268-278.
- Fourment M., Claywell B.C., Dinh V., McCoy C., Matsen I., Frederick A., Darling A.E. 2018a. Effective online Bayesian phylogenetics via sequential Monte Carlo with guided proposals. Syst. Biol. 67:490– 502.
- Fourment M., Magee A.F., Whidden C., Bilge A., Matsen IV F.A., Minin V.N. 2018b. 19 dubious ways to compute the marginal likelihood of a phylogenetic tree. Technical report, arXiv:1811.11804.
- Friel N., Pettitt A.N. 2008. Marginal likelihood estimation via power posteriors. J. R. Stat. Soc. Series B Stat. Methodol. 70:589–607.
- Geĺman A., Meng X.-L. 1998. Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. Stat. Sci. 13:163–185.
- Geweke J. 2004. Getting it right. J. Am. Stat. Assoc. 99:799-804.
- Görür D., Boyles L., Welling M. 2012. Scalable inference on Kingman's coalescent using pair similarity. J. Mach. Learn. Res. 22:440–448.
- Görür D., Teh Y.W. 2009. An efficient sequential Monte Carlo algorithm for coalescent clustering. In: Koller D., Schuurmans D., Bengio T., Bottou L. editor. Advances in neural information processing systems. Curran Associates, Inc. p. 521–528.

- Gunawan D., Kohn R., Quiroz M., Dang K.-D., Tran M.-N. 2018. Subsampling sequential Monte Carlo for static Bayesian models. Technical report, arXiv:1805.03317.
- Hajiaghayi M., Kirkpatrick B., Wang L., Bouchard-Côté A. 2014. Efficient continuous-time Markov chain estimation. In: Eric P. Xing, Tony Jebara, editors. Proceedings of the 31st International Conference on Machine Learning (ICML), Vol. 32. Bejing, China: PMLR. p. 638–646.
- Höhna S., Defoin-Platel M., Drummond A. 2008. Clock-constrained tree proposal operators in Bayesian phylogenetic inference. In: 8th IEEE international conference on bioinformatics and bioengineering, Athens, Greece. p. 1–7.
- Höhna S., Drummond A.J. 2012. Guided tree topology proposals for Bayesian phylogenetic inference. Syst. Biol. 61:1–11.
- Höhna S., Landis M.J., Heath T.A., Boussau B., Lartillot N., Moore B.R., Huelsenbeck J.P., Ronquist F. 2016. Revbayes: Bayesian phylogenetic inference using graphical models and an interactive model-specification language. Syst. Biol. 65:726–736.
- Holder M., Lewis P.O. 2003. Phylogeny estimation: traditional and Bayesian approaches. Nat. Rev. Gen. 4:275.
- Huelsenbeck J.P., Larget B., Alfaro M.E. 2004. Bayesian phylogenetic model selection using reversible jump Markov chain Monte Carlo. Mol. Biol. Evol. 21:1123–1133.
- Huelsenbeck J.P., Ronquist F. 2001. MRBAYES: Bayesian inference of phylogenetic trees. Bioinformatics 17:754–755.
- Jeffreys H. 1935. Some tests of significance, treated by the theory of probability. In: Mathematical proceedings of the Cambridge Philosophical Society, Vol. 31. Cambridge University Press. p. 203– 222.
- Jow H., Hudelot C., Rattray M., Higgs P.G. 2002. Bayesian phylogenetics using an RNA substitution model applied to early mammalian evolution. Mol. Biol. Evol. 19:1591–1601.
- Jun S.-H., Bouchard-Côté A. 2014. Memory (and time) efficient sequential Monte Carlo. In: Proceedings of the 31st International Conference on Machine Learning (ICML), Vol. 31. Bejing, China: PMLR. p.514–522.
- Kimura M. 1980. A simple method for estimating evolutionary rates of base substitutions through comparative studies of nucleotide sequences. J. Mol. Evol. 16:111–120.
- Kong A. 1992. A note on importance sampling using standardized weights. The University of Chicago. Department of Statistics, Technical report 348.
- Kuhner M.K., Felsenstein J. 1994. A simulation comparison of phylogeny algorithms under equal and unequal evolutionary rates. Mol. Biol. Evol. 11:459–468.
- Lakner C., van der Mark P., Huelsenbeck J.P., Larget B., Ronquist F. 2008. Efficiency of Markov chain Monte Carlo tree proposals in Bayesian phylogenetics. Syst. Biol. 57:86–103.
- Larget B., Simon D. 1999. Markov chain Monte Carlo algorithms for the Bayesian analysis of phylogenetic trees. Mol. Biol. Evol. 16: 750–759.
- Lartillot N. 2006. Conjugate Gibbs sampling for Bayesian phylogenetic models. J. Comput. Biol. 13:1701–1722.
- Lartillot N., Lepage T., Blanquart S. 2009. PhyloBayes 3: a Bayesian software package for phylogenetic reconstruction and molecular dating. Bioinformatics 25:2286–2288.
- Lartillot N., Philippe H., Lewis P. 2006. Computing Bayes factors using thermodynamic integration. Syst. Biol. 55:195–207.
- Lemey P., Rambaut A., Welch J.J., Suchard M.A. 2010. Phylogeography takes a relaxed random walk in continuous space and time. Mol. Biol. Evol. 27:1877–1885.
- Li S., Pearl D.K., Doss H. 2000. Phylogenetic tree construction using markov chain monte carlo. J. Am. Stat. Assoc. 95:493–508.
- Mau B., Newton M., Larget B. 1999. Bayesian phylogenetic inference via Markov chain Monte Carlo. Biometrics 55:1–12.
- Miller J.C., Maloney C.J. 1963. Systematic mistake analysis of digital computer programs. Commun. ACM 6:58–63.
- Neal R.M. 2001. Annealed importance sampling. Stat. Comput. 11:125– 139.
- Neal R.M. 2005. Estimating ratios of normalizing constants using linked importance sampling. Technical report, arXiv:math/0511216.

- Newton M.A., Raftery A.E. 1994. Approximate Bayesian inference with the weighted likelihood bootstrap. J. R. Stat. Soc. Series B Methodol. 56:3–48.
- Oaks J.R., Cobb K.A., Minin V.N., Leaché A.D. 2018. Marginal likelihoods in phylogenetics: a review of methods and applications. Technical report, arXiv:1805.04072.
- Olsson J., Douc R. 2017. Numerically stable online estimation of variance in particle filters. Technical report, arXiv:1701.01001.
- Owen M., Provan J.S. 2011. A fast algorithm for computing geodesic distances in tree space. IEEE/ACM Trans. Comput. Biol. Bioinform. 8:2–13.
- Quiroz M., Kohn R., Villani M., Tran M.-N. 2018a. Speeding up MCMC by efficient data subsampling. J. Am. Stat. Assoc. 1–13.
- Quiroz M., Tran M.-N., Villani M., Kohn R. 2018b. Speeding up MCMC by delayed acceptance and data subsampling. J. Comput. Graph. Stat. 27:12–22.
- Rannala B., Yang Z. 1996. Probability distribution of molecular evolutionary trees: a new method of phylogenetic inference. J. Mol. Evol. 43:304–311.
- Rannala B., Yang Z. 2003. Bayes estimation of species divergence times and ancestral population sizes using DNA sequences from multiple loci. Genetics 164:1645–1656.
- Roberts G.O., Gelman A., Gilks W.R. 1997. Weak convergence and optimal scaling of random walk Metropolis algorithms. Ann. Appl. Probab. 7:110–120.
- Robinson D., Foulds L. 1981. Comparison of phylogenetic trees. Math. Biosci. 53:131–147.
- Robinson D.F., Foulds L.R. 1979. Comparison of weighted labelled trees. In: Horadam A.F., Wallis W. D., editors. Combinatorial mathematics VI. Berlin: Springer. p. 119–126.
- Smith R.A., Ionides E.L., King A.A. 2017. Infectious disease dynamics inferred from genetic data via sequential Monte Carlo. Mol. Biol. Evol. 34:2065–2084.

- Tavaré, S. 1986. Some probabilistic and statistical problems in the analysis of DNA sequences. Lect. Math. Life Sci. 17:57–86.
- Teh Y.W., Daume III H., Roy D.M. 2008. Bayesian agglomerative clustering with coalescents. In: Platt J. C., Koller D., Singer Y., Roweis S. T., editors. Advances in neural information processing systems. 20. Vancouver, B.C., Canada: Curran Associates, Inc. p. 1473–1480.
- Tierney L. 1994. Markov chains for exploring posterior distributions. Ann. Stat. 22:1701–1762.
- Wang L., Bouchard-Côté A., Doucet A. 2015. Bayesian phylogenetic inference using a combinatorial sequential Monte Carlo method. J. Am. Stat. Assoc. 110:1362–1374.
- Wingate D., Goodman N., Stuhlmueller A., Siskind J.M. 2011. Nonstandard interpretations of probabilistic programs for efficient inference. In: Shawe-Taylor J., Zemel R.S., Bartlett P.L., Pereira F., Weinberger K.Q., editors. Advances in neural information processing systems 24. Granada Spain: Curran Associates, Inc. p. 1152–1160.
- Xie W., Lewis P.O., Fan Y., Kuo L., Chen M.-H. 2010. Improving marginal likelihood estimation for Bayesian phylogenetic model selection. Syst. Biol. 60:150–160.
- Yang Z., Rannala B. 1997. Bayesian phylogenetic inference using DNA sequences: a Markov Chain Monte Carlo method. Mol. Biol. Evol. 14:717–724.
- Zhao T., Wang Z., Cumberworth A., Gsponer J., Freitas N.D., Bouchard-Côté, A. 2016. Bayesian analysis of continuous time Markov chains with application to phylogenetic modelling. Bayesian Anal. 11:1203– 1237.
- Zhou Y., Johansen A.M., Aston J.A. 2016. Toward automatic model comparison: an adaptive sequential Monte Carlo approach. J. Comput. Graph. Stat. 25:701–726.