

A Variational Approach to Bayesian Phylogenetic Inference

Cheng Zhang*

CHENGZHANG@MATH.PKU.EDU.CN

*School of Mathematical Sciences and Center for Statistical Science
Peking University
Beijing, 100871, China*

Frederick A. Matsen IV

MATSEN@FREDHUTCH.ORG

*Howard Hughes Medical Institute
Computational Biology Program
Fred Hutchinson Cancer Research Center
Seattle, WA 98109, USA
Department of Genome Sciences and Department of Statistics
University of Washington
Seattle, WA 98195, USA*

Editor: Chris Wiggins

Abstract

Bayesian phylogenetic inference is currently done via Markov chain Monte Carlo (MCMC) with simple proposal mechanisms. This hinders exploration efficiency and often requires long runs to deliver accurate posterior estimates. In this paper, we present an alternative approach: a variational framework for Bayesian phylogenetic analysis. We propose combining subsplit Bayesian networks, an expressive graphical model for tree topology distributions, and a structured amortization of the branch lengths over tree topologies for a suitable variational family of distributions. We train the variational approximation via stochastic gradient ascent and adopt gradient estimators for continuous and discrete variational parameters separately to deal with the composite latent space of phylogenetic models. We show that our variational approach provides competitive performance to MCMC, while requiring much fewer (though more costly) iterations due to a more efficient exploration mechanism enabled by variational inference. Experiments on a benchmark of challenging real data Bayesian phylogenetic inference problems demonstrate the effectiveness and efficiency of our methods.

Keywords: Bayesian phylogenetic inference, variational inference, subsplit Bayesian networks, structured amortization

1. Introduction

Bayesian phylogenetic inference is an essential statistical method for modern molecular evolutionary analysis and has been used in a wide range of applications such as genomic epidemiology (Dudas et al., 2017; du Plessis et al., 2021) and conservation genetics (DeSalle and Amato, 2004). Given aligned sequence data (e.g., DNA, RNA or protein sequences) and a model of evolution, Bayesian phylogenetics provides principled approaches to quantify the uncertainty of the evolutionary process in terms of the posterior probabilities of phylogenetic trees (Huelsenbeck et al., 2001). In addition to uncertainty quantification, Bayesian methods

*. Corresponding author.

enable integrating out tree uncertainty in order to get more confident estimates of parameters of interest, such as factors governing the transmission of Ebolavirus (Dudas et al., 2017). Bayesian methods also allow complex substitution models (Lartillot and Philippe, 2004), which are important in elucidating deep phylogenetic relationships (Feuda et al., 2017).

Ever since its introduction to the phylogenetic community in the 1990s, Bayesian phylogenetic inference has been dominated by random-walk Markov chain Monte Carlo (MCMC) approaches (Yang and Rannala, 1997; Mau et al., 1999; Huelsenbeck and Ronquist, 2001; Drummond et al., 2002, 2005). However, this approach is fundamentally limited by the complexities of tree space. A typical MCMC method for phylogenetic inference involves two steps in each iteration: first, a new tree is proposed by randomly perturbing the current tree, and second, the tree is accepted or rejected according to the Metropolis-Hastings acceptance probability. Any such random walk algorithm faces obstacles in the phylogenetic case, in which the high-posterior trees are a tiny fraction of the combinatorially exploding number of trees. Thus, major modifications of trees are likely to be rejected, restricting MCMC tree movement to local modifications that may have difficulty moving between multiple peaks in the posterior distribution (Whidden and Matsen IV, 2015). Although recent MCMC methods for distributions on Euclidean space use intelligent proposal mechanisms such as Hamiltonian Monte Carlo (Neal, 2011), it is not straightforward to extend such algorithms to the composite structure of tree space, which includes both tree topology (discrete object) and branch lengths (continuous positive vector) (Dinh et al., 2017b).

Variational inference (VI) is an alternative approximate inference method for Bayesian analysis which is gaining in popularity (Jordan et al., 1999; Wainwright and Jordan, 2008; Blei et al., 2017). Unlike MCMC methods that sample from the posterior, VI seeks the best candidate from a family of tractable distributions that minimizes a statistical distance measure to the target posterior, usually the Kullback-Leibler (KL) divergence. By reformulating the inference problem into an optimization problem, VI tends to be faster and easier to scale to large data (via stochastic gradient descent) (Blei et al., 2017). However, VI can also introduce a large bias if the variational family of distributions is insufficiently flexible. The success of variational methods, therefore, largely relies on the construction of appropriate tractable variational family of distributions and efficient optimization procedures.

Recent years have witnessed several efforts on VI for Bayesian phylogenetics that focus on continuous parameters (e.g., the branch lengths), where either a fixed tree topology is assumed (Fourment and Darling, 2019; Ki and Terhorst, 2022) or MCMC algorithms (e.g., Gibbs sampling) are employed to handle the uncertainty of tree topologies (Dang and Kishino, 2019). To our knowledge, there have been no previous full variational formulations of Bayesian phylogenetic inference that deal with the joint parameter space. This has been due to the lack of an appropriate family of distributions on phylogenetic tree topologies that can be used as variational distributions. However the prospects for variational inference have changed recently with the introduction of *subsplit Bayesian networks* (SBNs) (Zhang and Matsen IV, 2018), which provide a family of flexible distributions on tree topologies (i.e. trees without branch lengths). SBNs build on previous work (Höhna and Drummond, 2012; Larget, 2013), but in contrast to these previous efforts, SBNs are sufficiently flexible for real Bayesian phylogenetic posteriors (Zhang and Matsen IV, 2018). Other families of distributions over the joint phylogenetic parameter space have also been proposed outside

of the variational framework (Holder et al., 2014; Baele et al., 2016) with the goal of more efficient marginal likelihood estimation.

In this paper, we develop a general variational inference framework for Bayesian phylogenetics. We show that SBNs, when combined with appropriate approximations for branch length distributions, can provide flexible variational approximations over the joint latent space of phylogenetic trees with branch lengths. This, combined with other variational distributions when needed (e.g., distributions of population size parameters), provides a suitable variational family of distributions for Bayesian phylogenetic inference. The variational approximation can be trained using stochastic gradient ascent via efficient unbiased gradient estimators proposed recently for discrete and continuous parameters (Kingma and Welling, 2014; Bornschein and Bengio, 2015; Mnih and Rezende, 2016). We introduce several effective approaches to provide practical subsplit support estimation for SBNs that significantly reduce the number of parameters while maintaining adequate approximation power for tree topology posterior estimation. We also leverage the similarity of local topological structures for a shared parameterization of the branch length distributions over different tree topologies to further reduce the complexity in variational parameterization and ease optimization. With guided exploration in the tree topology space within that support (enabled by SBNs) and joint learning of the branch length distributions across tree topologies (via structured amortization), our variational approach provides competitive performance to MCMC with much fewer (though more costly) iterations. Moreover, the variational approximations can be readily used for downstream statistical analysis such as marginal likelihood estimation for model comparison via importance sampling. We demonstrate the effectiveness and efficiency of our methods on a benchmark of challenging real data Bayesian phylogenetic inference problems. A preliminary version of this work has appeared in Zhang and Matsen IV (2019); in addition to providing a complete exposition, this version provides additional benchmarks and extends our work to an additional class of models (“time trees”) which are very popular for analyzing viral sequences.

2. Background

In this section, we provide a brief overview of Bayesian phylogenetics and variational inference, which are the fundamental ingredients for this paper.

2.1 Bayesian Phylogenetics

The core in Bayesian phylogenetics is a probability model that describes a generative process for a given data set consisting of nucleotide or amino acid sequences from a collection of *taxa* (singular *taxon*, e.g., species). There are two essential ingredients for this probability model and the associated parameters. The first ingredient is a phylogenetic tree which contains a tree topology τ and the associated non-negative branch lengths \mathbf{q} . The tree topology τ is a bifurcating tree whose leaves are labelled by the taxa. It is used to model the evolutionary relationship of the observed species, where the internal nodes represents the unobserved characters (e.g., DNA bases) of the ancestral species. Phylogenetic trees can be *rooted*, in which a root position in the tree is chosen to be the most ancestral part of the tree, or *unrooted*, if only evolutionary relationships and not the direction of evolution is of interest. The second ingredient is a continuous-time Markov model that is used to

describe the transition probabilities of the characters along the branches of the tree. These transition probabilities are assumed to be independent across the tree. In this paper, we focus on the parameter space of phylogenetic trees and assume the Markov evolution model is known (e.g., the Jukes and Cantor (1969) model). The phylogenetic likelihood is, therefore, the function of phylogenetic tree τ, \mathbf{q} that gives the probability of the observed data. Let $\mathbf{Y} = \{Y_1, Y_2, \dots, Y_M\} \in \Omega^{N \times M}$ be the observed sequences (with characters in Ω) of length M over N species. Let $Y_i(v)$ is the observed character of node v at site i . Define an extension of Y_i to the internal nodes to be an assignment of states a^i to the nodes of the tree such that $a_v^i = Y_i(v)$ for every tip node v . The probability of each site observation Y_i is defined as the marginal distribution over the leaves

$$p(Y_i|\tau, \mathbf{q}) = \sum_{a^i} \eta(a_r^i) \prod_{(u,v) \in E(\tau)} P_{a_u^i a_v^i}(q_{uv}) \quad (1)$$

where r is the root node, a^i ranges over all extensions of Y_i to the internal nodes with a_u^i being the assigned character of node u , $E(\tau)$ denotes the set of edges of τ , q_{uv} is the branch length for the edge $(u, v) \in E(\tau)$, and $P_{ij}(t)$ denotes the transition probability from character i to character j across an edge of branch length t given by the Markov model whose stationary distribution is η . For unrooted trees with a time reversible Markov model, (1) also defines a valid marginal probability where r can be any internal node due to the so-called ‘‘pulley principle’’ (Felsenstein, 1981).

Assuming different sites are identically distributed and evolve independently, the likelihood function then takes the following form

$$p(\mathbf{Y}|\tau, \mathbf{q}) = \prod_{i=1}^M p(Y_i|\tau, \mathbf{q}) = \prod_{i=1}^M \sum_{a^i} \eta(a_r^i) \prod_{(u,v) \in E(\tau)} P_{a_u^i a_v^i}(q_{uv}) \quad (2)$$

The phylogenetic likelihood (2) can be evaluated efficiently through the pruning algorithm (Felsenstein, 1981), also known as the sum-product algorithm in general probabilistic graphical models (Strimmer and Moulton, 2000; Koller and Friedman, 2009; Höhna et al., 2014). Given a proper prior distribution with density $p(\tau, \mathbf{q})$ of the tree topology and the branch lengths, the posterior $p(\tau, \mathbf{q}|\mathbf{Y})$ is proportional to the joint density

$$p(\tau, \mathbf{q}|\mathbf{Y}) = \frac{p(\mathbf{Y}|\tau, \mathbf{q})p(\tau, \mathbf{q})}{p(\mathbf{Y})} \propto p(\mathbf{Y}|\tau, \mathbf{q})p(\tau, \mathbf{q}) \quad (3)$$

where $p(\mathbf{Y})$ is the unknown normalizing constant. Bayesian phylogenetics then amounts to properly estimating the phylogenetic posterior in (3), which is typically done via Markov chain Monte Carlo (MCMC) methods.

2.2 Variational Inference

Rather than sampling from the posterior as in MCMC, VI turns the posterior inference problem into an optimization problem. Suppose \mathbf{Y} is the observed data and $p(\mathbf{Y}|\boldsymbol{\theta})$ is the probability model where $\boldsymbol{\theta}$ is the set of model parameters. The objective of VI is to find the closest member from a family of densities over $\boldsymbol{\theta}$ to the posterior $p(\boldsymbol{\theta}|\mathbf{Y})$. Let $q_\phi(\boldsymbol{\theta})$ denotes the density of variational distribution where ϕ is the set of variational parameters.

To measure the “closeness” between $p(\boldsymbol{\theta}|\mathbf{Y})$ and $q_\phi(\boldsymbol{\theta})$, a divergence measure D between distributions is often employed such that $D(q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathbf{Y})) \geq 0$ and $D(q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathbf{Y})) = 0$ iff $q_\phi(\boldsymbol{\theta}) = p(\boldsymbol{\theta}|\mathbf{Y})$. While various divergence measures exist (Amari, 1985, 2009; Minka, 2005; Sason and Verdú, 2016), the most commonly used divergence in VI is the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951)

$$D_{\text{KL}}(q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathbf{Y})) = \int q_\phi(\boldsymbol{\theta}) \log \left(\frac{q_\phi(\boldsymbol{\theta})}{p(\boldsymbol{\theta}|\mathbf{Y})} \right) d\boldsymbol{\theta}.$$

Inference now amounts to solving the following optimization problem

$$\phi^* = \arg \min_{\phi} D_{\text{KL}}(q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathbf{Y})). \quad (4)$$

The best approximation $q_{\phi^*}(\boldsymbol{\theta})$ then serves as a proxy for the exact posterior.

As the exact posterior $p(\boldsymbol{\theta}|\mathbf{Y})$ is only known up to a constant $p(\mathbf{Y})$, the objective in (4) is not computable. In practice, we optimize an alternative function

$$L(\phi) = \int q_\phi(\boldsymbol{\theta}) \log \left(\frac{p(\boldsymbol{\theta}, \mathbf{Y})}{q_\phi(\boldsymbol{\theta})} \right) d\boldsymbol{\theta} = \log p(\mathbf{Y}) - D_{\text{KL}}(q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathbf{Y})) \leq \log p(\mathbf{Y}).$$

This function is called the evidence lower bound (ELBO). As the KL divergence is the gap between the constant $\log p(\mathbf{Y})$ and the ELBO, maximizing the EBLO is equivalent to minimizing the KL divergence.

To complete the specification of the optimization problem in (4), we need to prescribe a variational family of distributions. A common choice is the mean-field variational family, where different components of $\boldsymbol{\theta}$ are mutually independent. Given conditionally conjugate exponential family models for $p(\boldsymbol{\theta}, \mathbf{Y})$, mean-field variational inference can be solved efficiently via coordinate ascent (Bishop, 2006). However, the independence assumption also leads to insufficient approximation accuracy especially when the exact posterior variables are highly correlated. Many more flexible variational families have been proposed (Saul and Jordan, 1996; Bishop et al., 1998; Barber and Wiergerinck, 1999; Rezende and Mohamed, 2015; Kingma et al., 2016; Dinh et al., 2017a), albeit the resulting optimization problems are often more difficult to solve. Furthermore, researchers have also generalized VI beyond conditional conjugate models, and designed generic model-agnostic methods based on stochastic optimization via Monte Carlo gradient estimates (Nott et al., 2012; Paisley et al., 2012; Kingma and Welling, 2014; Rezende et al., 2014; Ranganath et al., 2014). Overall, designing an appropriate variational family that strikes a good balance between expressiveness and complexity in optimization has been one of the key obstacles to the success of VI. In what follows, we first describe an essential ingredient for the construction of a suitable variational family for Bayesian phylogenetics.

3. Subsplit Bayesian Networks

Since the number of tree topologies explodes combinatorially as the number of taxa increases, designing tractable and flexible distributions over the discrete tree topology space turns out to be quite challenging. Subsplit Bayesian networks, as recently proposed by Zhang and Matsen IV (2018), supply a powerful probabilistic graphical model that can provide a flexible

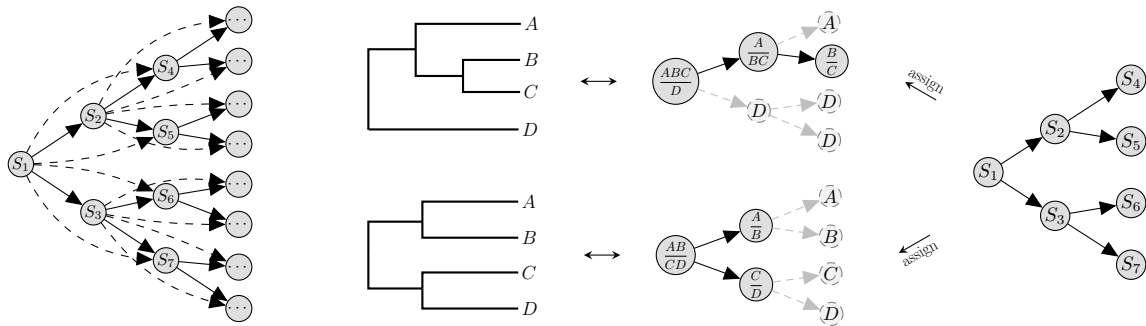


Figure 1: Subsplit Bayesian networks and a simple example for a leaf set of 4 taxa (represented by A , B , C , and D respectively). **(Left)**: General subsplit Bayesian networks. The solid backbone represents the complete and binary tree network $\mathcal{B}_{\mathcal{X}}^*$. The dashed arrows represent the additional dependence for more expressive SBNs. **(Middle(left))**: Examples of phylogenetic trees (rooted) that are hypothesized to model the evolutionary history of the taxa. **(Middle(right))**: The corresponding node assignments for the trees. For ease of illustration, subsplit (Y, Z) is represented as $\frac{Y}{Z}$ in the graph. The *dashed gray subgraphs* represent trivial splitting processes where the subsplits are deterministically assigned, and are used purely to complement the networks such that the overall network has a fixed structure. **(Right)**: The SBN for these examples, which is $\mathcal{B}_{\mathcal{X}}^*$ in this case.

family of distributions over tree topologies. The success of SBNs lies in its proper utilization of the similarity of the local topological structures of trees, i.e., *subsplits*, as defined below:

Definition 1 (Subsplit) *Let \mathcal{X} be a set of labeled leaves that represent the observed taxa. We call a nonempty subset C of \mathcal{X} a clade and denote the set of all clades of \mathcal{X} as $\mathcal{C} = \{C \subset \mathcal{X} : C \neq \emptyset\}$. A nonempty subset of a clade C is called a subclade of C . Let \succ be a total order on \mathcal{C} (e.g., lexicographical order). A subsplit (Y, Z) of a clade C is an ordered pair of disjoint subclades of C such that $Y \cup Z = C$ and $Y \succ Z$.*

The hierarchical structure of a rooted tree topology induces a splitting process from the root node to the tip nodes, where each internal node has a subsplit that represents its local splitting pattern. Following this splitting process, all the topological information of the tree can be naturally converted into a sequence of subsplits. As the tree topology can be easily recovered from these local splitting patterns afterward, we call them the *subsplit decomposition* of the tree. For example, in Figure 1, the top tree has a subsplit decomposition $\{(\{A, B, C\}, \{D\}), (\{A\}, \{B, C\}), (\{B\}, \{C\})\}$ and the bottom tree has a subsplit decomposition $\{(\{A, B\}, \{C, D\}), (\{A\}, \{B\}), (\{C\}, \{D\})\}$. Unlike the full tree topologies that require a combinatorially exploding representation space, subsplit decompositions allow us to represent topological information in a more compact manner since subsplits (and other local splitting patterns) can be shared across different tree topologies, as shown in the following example.




Figure 2: An illustration for Example 1. **(Left)**: A base tree τ_0 with 3-taxon clade assignments for the tip nodes. **(Right)**: The local topologies that can be taken by the subtrees in C_1 .

Example 1 Consider a leaf label set \mathcal{X} that has $3n$ taxa A_1, A_2, \dots, A_{3n} . Let τ_0 be a tree topology with n tips. Now partition \mathcal{X} into n clades with 3 taxa, say $C_1 = \{A_1, A_2, A_3\}$, $C_2 = \{A_4, A_5, A_6\}$, \dots , $C_n = \{A_{3n-2}, A_{3n-1}, A_{3n}\}$ and assign them to the tip nodes of τ_0 . Note that each clade $C_i, 1 \leq i \leq n$ may have 3 different local topologies. Overall, this allows us to represent 3^n tree topologies on \mathcal{X} . On the other hand, the subsplits from these tree topologies are either the subsplits induced by τ_0 (which is $\mathcal{O}(n)$) or the local subsplits from those 3 taxa subtrees in the tips (which is also $\mathcal{O}(n)$). Therefore, we can represent an exponential number of tree topologies with $\mathcal{O}(n)$ subsplits. See Figure 2 for an illustration.

Moreover, with appropriate conditional independence assumptions, the subsplit representations of tree topologies inspire a Bayesian network formulation of distributions over the tree topology space.

Definition 2 (Subsplit Bayesian Network) A subsplit Bayesian network $\mathcal{B}_{\mathcal{X}}$ on a leaf label set \mathcal{X} of size N is a Bayesian network whose nodes take on subsplit values or singleton clade values of \mathcal{X} , and has the following properties.

1. The root node of $\mathcal{B}_{\mathcal{X}}$ takes on subsplits of the entire leaf label set \mathcal{X} .
2. $\mathcal{B}_{\mathcal{X}}$ contains a full and complete binary tree network $\mathcal{B}_{\mathcal{X}}^*$ as a subnetwork.
3. The depth of $\mathcal{B}_{\mathcal{X}}$ is $N - 1$, with the root counting as depth 1.

Figure 1 provides an illustration of the structure of general subsplit Bayesian networks (SBNs), together with a concrete example for a leaf set of 4 taxa. We see that $\mathcal{B}_{\mathcal{X}}^*$ itself is an SBN that is contained in all SBNs, i.e., $\mathcal{B}_{\mathcal{X}}^*$ is the minimum SBN on \mathcal{X} . All SBNs, therefore, share the same set of nodes, and differ only in their conditional independence structures. Moreover, there is a natural universal indexing procedure for the nodes of all SBNs based on $\mathcal{B}_{\mathcal{X}}^*$ as follows. First, we denote the root node as S_1 . Then, for any i , we denote the two children of S_i on $\mathcal{B}_{\mathcal{X}}^*$ as S_{2i} and S_{2i+1} , with S_{2i} being on top of S_{2i+1} , until a leaf node is reached (the right plot in Figure 1). We call the parent nodes in $\mathcal{B}_{\mathcal{X}}^*$ the *natural parent*. Although this configuration of SBNs would take 2^{N-1} nodes in total, only a small portion of

these nodes will be used when it comes to tree probability estimation as the others represent trivial local splitting patterns of tree topologies (the dashed gray subgraphs in the middle plots of Figure 1). In practice, we can use conditional probability sharing in Section 3.1 and subsplit support estimation in Section 4.2 for a compact parameterization of SBNs.

For a rooted tree topology, one can find its unique SBN representation by assigning the subsplit decomposition to the nodes of SBNs according to the same splitting process as before (the middle plots in Figure 1). This way, the subsplit for a node is always a bipartition of a subclade in the subsplit for its natural parent. We refer to this kind of assignments of nodes as *compatible node assignments*.

Definition 3 (Compatible Node Assignment) *For a non-singleton clade X , we say a subsplit (Y, Z) is compatible with X iff $Y \cup Z = X$. For a singleton clade X , we say a singleton clade Y is compatible with X iff $Y = X$. We say a full node assignment $\{S_i = s_i\}_{i \geq 1}$ is compatible if for all interior node assignments s_i that are subsplits (Y_i, Z_i) , the corresponding assignments of the children nodes s_{2i} and s_{2i+1} are compatible with Y_i, Z_i , respectively.*

Clearly, every rooted tree topology corresponds to a unique compatible node assignment. On the other hand, we can also restore a rooted tree given a compatible node assignment. Therefore, we have the following lemma.

Lemma 4 *There is a one-to-one mapping between rooted tree topologies on \mathcal{X} and compatible node assignments.*

See a proof of Lemma 4 in Appendix B. Note that in practice SBNs are often parameterized based on some estimated subsplit supports (see Section 4.2 for more details). Tree topologies that have local splitting patterns that are not covered by those subsplit supports would have zero SBN probabilities. In other words, compatible node assignments would not necessarily lead to a positive SBN probability of the corresponding tree topologies. Note that an alternative formulation of an SBN for a choice of subsplit support is naturally described using a directed acyclic graph; see Jun et al. (2023) for more details.

3.1 Tree Probability Estimation

As Bayesian networks, SBNs can be used to define probabilities for rooted trees based on their corresponding node assignments. Given the subsplit decomposition $\{s_1, s_2, \dots\}$ of a rooted tree τ , where s_1 is the root subsplit and $\{s_i\}_{i > 1}$ are the other subsplits, we define the SBN-induced tree probability of τ as

$$p_{\text{sbn}}(T = \tau) = p(S_1 = s_1) \prod_{i > 1} p(S_i = s_i | S_{\pi_i} = s_{\pi_i}) \quad (5)$$

where π_i denotes the set of indices of the parent nodes of S_i (if we are using the minimal SBN $\mathcal{B}_{\mathcal{X}}^*$ there will only be one parent node). Note that a naive parameterization of SBNs would require conditional probability distributions (CPDs) for all local conditional models $p(S_i | S_{\pi_i})$ that scales exponentially as the number of taxa N increases. To reduce the number of parameters and encourage generalization, *conditional probability sharing* is often assumed

where the same set of CPDs for parent-child subsplit pairs is shared across the SBN network, regardless of their locations (Zhang and Matsen IV, 2018). In practice, this proves effective for a compact parameterization of SBNs (see Section 4.3 for more details). Traditionally, Bayesian networks define distributions over all possible assignments. However, in this setting we only want to consider node assignments that form a tree. To remedy this issue, we therefore impose the following consistency condition for the associated CPDs of SBNs.

Definition 5 (Consistent Parameterization) *We say the conditional probability $p(S_i|S_{\pi_i})$ is consistent if $p(S_i = s_i|S_{\pi_i} = s_{\pi_i}) > 0$ only when the pair $(S_i = s_i, S_{\pi_i} = s_{\pi_i})$ can be extended to a compatible node assignment of all nodes. We say an SBN is consistently parameterized if all the associated conditional probabilities are consistent.*

The consistency condition restricts the support of SBNs to the compatible node assignments, which together with Lemma 4, guarantee that SBNs provide valid distributions over the rooted tree topology space.

Proposition 6 *Let $\mathcal{T}_{\mathcal{X}}$ be the set of all rooted trees with leaf labels \mathcal{X} . For any consistently parameterized SBN, $\sum_{\tau \in \mathcal{T}_{\mathcal{X}}} p_{\text{sbn}}(T = \tau) = 1$.*

The proof of Proposition 6 is standard and is provided in Appendix B. Based on Proposition 6, we can design a family of distributions over the rooted tree topology space using the conditional independence structures in SBNs (Zhang and Matsen IV, 2018), where $\mathcal{B}_{\mathcal{X}}^*$ serves as a base model (Figure 1, Left).

The SBN framework also generalizes to unrooted trees, which are another popular type of trees in phylogenetic inference that depicts only the relationship between taxa (Figure 3, left). The key idea is to view unrooted trees as rooted trees where the position of the root is treated as a missing variable. More concretely, unrooted trees can be transformed into (and hence be represented as) a set of rooted trees by placing the virtual roots at the edges. On the other hand, rooted trees can also be transformed into unrooted trees by removing the roots. Figure 3 shows a simple example on a four taxon unrooted tree. Through this rooting/unrooting operation, the unrooted topology induces an equivalence relation among rooted trees defined as follows.

Definition 7 (Equivalence Relation) *Let D be the unrooting operator. For any pair of rooted trees τ_1 and τ_2 , we say $\tau_1 \sim \tau_2$ if $D(\tau_1)$ and $D(\tau_2)$ represent the same unrooted tree, that is, $D(\tau_1) = D(\tau_2)$.*

It is easy to verify that \sim defines an equivalence relation in the rooted tree space, and every unrooted tree τ^u corresponds to an equivalence class $[\tau^u] = \{\tau : D(\tau) = \tau^u\}$. We, therefore, define the SBN probability estimates for unrooted trees as the sum of SBN probabilities for the rooted trees in the equivalence class

$$p_{\text{sbn}}(T^u = \tau^u) = \sum_{\tau \in [\tau^u]} p_{\text{sbn}}(\tau). \tag{6}$$

As the equivalence classes naturally partition the entire rooted tree space, (6) also provides a valid probability distribution in the unrooted tree space.

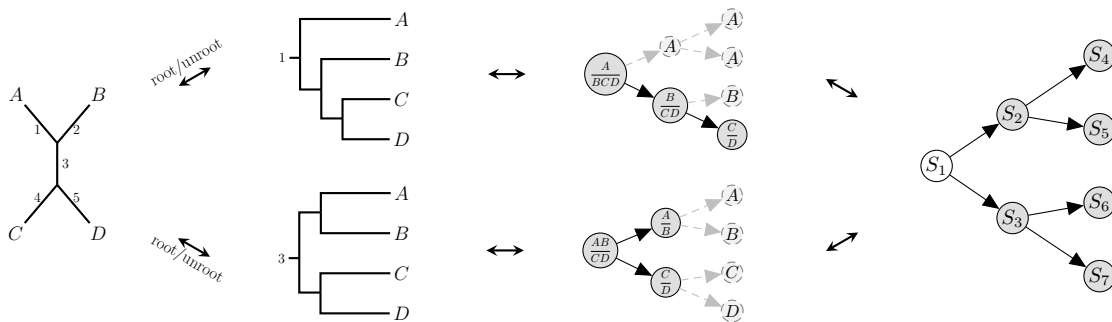


Figure 3: SBNs for unrooted trees. **(Left)**: A simple four taxon unrooted tree. The equivalence class consists of five rooted trees, and each can be obtained by rooting on one of the edges $\{1, 2, 3, 4, 5\}$. **(Middle (left))**: Two exemplary rooted trees in the equivalence class when rooting on edges 1 and 3. **(Middle (right))**: The corresponding SBN assignments for the two rooted trees. **(Right)**: An SBN for the unrooted tree with unobserved root node S_1 .

Proposition 8 Let $\mathcal{T}_{\mathcal{X}}^u$ be the set of all unrooted trees with leaf labels \mathcal{X} . For any consistently parameterized SBN, $\sum_{\tau^u \in \mathcal{T}_{\mathcal{X}}^u} p_{\text{sbn}}(T^u = \tau^u) = 1$.

Proof

$$\sum_{\tau^u \in \mathcal{T}_{\mathcal{X}}^u} p_{\text{sbn}}(T^u = \tau^u) = \sum_{\tau^u \in \mathcal{T}_{\mathcal{X}}^u} \sum_{\tau \in [\tau^u]} p_{\text{sbn}}(\tau) = \sum_{\tau \in \mathcal{T}_{\mathcal{X}}} p_{\text{sbn}}(\tau) = 1.$$

■

Note that each rooted tree in the equivalence class corresponds to an edge (i.e., the rooting position) of the unrooted tree, $[\tau^u]$ therefore can be explicitly represented as $\{\tau_e^{\#} : e \in E(\tau^u)\}$, where $\tau_e^{\#}$ means the resulting rooted tree when the root edge is e . We will use this representation in the sequel. As each edge corresponds to a root subsplit, (6) can also be viewed as the marginal SBN probability over the unobserved root node S_1 (the right plot in Figure 3).

Not only do SBNs provide a family of distributions over the rooted/unrooted tree spaces, they also allow efficient tree probability computation algorithms and tree sampling procedures as we introduce next. All the involved computation complexity is measured with respect to the size of the tree topology, or equivalently, the number of taxa N .

3.2 Tree Probability Computation

In the previous section, we have illustrated how to use SBNs for tree probability estimation. We now show that these SBN-induced tree probabilities can be efficiently computed in linear time $O(N)$ as follows. For rooted trees, we can use a simple tree traversal to collect the required subsplits in (5) and evaluate the product using the conditional probabilities for the corresponding parent-child subsplit pairs in SBNs. This leads to a linear time algorithm




Figure 4: The two-pass algorithm for computing the SBN probabilities of unrooted trees in (7). **(Left)**: The conditional probability for the parent-child subsplit pair appeared in the message updating formulas. The triangles denote subtrees in a simplified form. The subtrees on the other side of node k are ignored. **(Middle)**: The postorder (solid black) and preorder (dashed gray) traversals when node k is chosen as the root node (denoted as a square node). **(Right)**: The root subsplit $(V \cup Z, X \cup Y)$ and parent-child subsplit pairs $(X, Y)|(V \cup Z, X \cup Y)$, $(V, Z)|(V \cup Z, X \cup Y)$ for edge e .

as both the tree traversal and product evaluation take linear time. For unrooted trees, (6) requires computing the SBN probabilities as defined in (5) for all rooted trees in the equivalence class. A naive approach is to run tree traversals for each of these rooted trees and collect the required subsplits $\cup_{e \in E(\tau^u)} \{s_1^{e, \tau^u}, s_2^{e, \tau^u}, \dots\}$, where $\{s_1^{e, \tau^u}, s_2^{e, \tau^u}, \dots\}$ is the subsplit decomposition for τ_e^u (the middle plots in Figure 3). With these subsplits, we can evaluate (6) as follows

$$p_{\text{sbn}}(\tau^u) = \sum_{e \in E(\tau^u)} p(S_1 = s_1^{e, \tau^u}) \prod_{i > 1} p(S_i = s_i^{e, \tau^u} | S_{\pi_i} = s_{\pi_i}^{e, \tau^u}). \quad (7)$$

Although (7) provides a valid computation method, this naive approach has a quadratic time complexity as both the tree traversals and sum-product evaluation takes quadratic time, making it difficult to scale up to large trees.

Fortunately, by leveraging the hierarchical structure of phylogenetic trees, we can design a more efficient linear time algorithm for SBN probability estimation of unrooted trees. Notice that most of the subsplits for rooted trees in the equivalence class and the intermediate partial products involved in (7) are shared due to the hierarchical structure of trees, this allows us to collect all the required subsplits and complete the sum-product evaluation within a two-pass sweep through the unrooted tree, similar to the two-pass algorithm for tree-structured graphical models (Pearl, 1988; Schadt et al., 1998; Willsky, 2002). Without loss of generality, we use the simplest SBN, \mathcal{B}_{χ}^* , as an example and describe our two-pass algorithm as below.

First, initialize all pairwise messages for the edges to one: $M_{i \rightarrow j} = M_{j \rightarrow i} = 1, \forall (i, j) \in E(\tau^u)$. Next, choose a node as the root node (e.g., node k in the middle plot in Figure 4, we will call this a *virtual root*) and make two passes through the tree as follows. In the first

pass, we traverse the tree in a postorder fashion and update the rootward messages when we visit interior node i (except the root node) as follows

$$M_{i \rightarrow \pi_i} = \prod_{j \in \text{ch}(i)} P_{\pi_i}(j \rightarrow i) M_{j \rightarrow i} \quad (8)$$

where $\text{ch}(i)$ means the child nodes of i and $P_{\pi_i}(j \rightarrow i)$ is the conditional probability for the parent-child subsplit pair representing the local splitting pattern of the child node j at node i , given the parent node π_i . More specifically, let (X, Y) be the corresponding subsplit at node j looking away from the parent node i and $(X \cup Y, Z)$ be the corresponding subsplit at node i looking away from the parent node π_i , then $P_{\pi_i}(j \rightarrow i) = p((X, Y) | (X \cup Y, Z))$ (see the left plot in Figure 4). In the second pass, we traverse the tree in a preorder fashion and update the leafward messages when we visit interior node i as follows

$$M_{i \rightarrow k} = \prod_{j \in \text{ne}(i) \setminus \{k\}} P_k(j \rightarrow i) M_{j \rightarrow i}, \quad \forall k \in \text{ch}(i), \quad (9)$$

where $\text{ne}(i)$ means the neighborhood of node i and $P_k(j \rightarrow i)$ is defined similarly as $P_{\pi_i}(j \rightarrow i)$ in (8). See an illustration in Figure 4. In both phases of the algorithm, $M_{i \rightarrow k}$ accumulates the contributions to SBN probabilities from the subtree rooted at node i looking away from neighbor k . After the second pass, for each edge $e = (i, j) \in E(\tau^u)$, we can aggregate the contributions to $p_{\text{sbn}}(\tau^u)$ from the subtrees on both sides of the edge as $M_{i \rightarrow j} M_{j \rightarrow i}$. This, together with the subsplit probability $P_r(e)$ and the parent-child subsplit pair conditional probabilities $P_r(i \rightarrow j), P_r(j \rightarrow i)$ associated with the root edge e (the right plot in Figure 4)¹, allow us to compute the SBN probability for each rooted tree τ_e^\dagger in the equivalence class, and hence compute the marginal SBN probability in (7) as

$$p_{\text{sbn}}(\tau^u) = \sum_{e \in E(\tau^u)} p_{\text{sbn}}(\tau_e^\dagger) = \sum_{e=(i,j) \in E(\tau^u)} P_r(e) P_r(i \rightarrow j) P_r(j \rightarrow i) M_{i \rightarrow j} M_{j \rightarrow i}. \quad (10)$$

Note that the message updates in each iteration of the tree traversals take constant time and the sum-product in (10) takes linear time, the above two-pass algorithm hence enjoys a linear time complexity.

3.3 Tree Sampling

Sampling rooted trees from SBNs is straightforward via ancestral sampling, a method that samples the random variables for the nodes of Bayesian networks in topological order. More specifically, we start by sampling from the root node; then we sample from the child nodes by conditioning their CPDs to the sampled values of their parent nodes. We proceed like this until all nodes have been sampled. According to Lemma 4, the resulting assignment of subsplits uniquely represents a rooted tree topology given that the SBN is consistently parameterized. Sampling unrooted trees from SBNs is almost identical to sampling rooted trees, but with an additional unrooting operation that transforms the sampled rooted trees into unrooted ones. As the rooting/unrooting operation defines the correspondence between

1. These root subsplits and parent-child subsplit pairs associated with the edges turn out to be quite useful. We will revisit them later.

the unrooted trees and their equivalence classes in the rooted tree space, it is easy to show that this sampling procedure is consistent with the SBN probabilities for unrooted trees defined in (6).

Proposition 9 *Let T^u be an unrooted tree valued random variable defined implicitly by a procedure that samples and deletes the roots of samples from SBNs (via ancestral sampling). Then T^u follows the distribution defined in (6).*

Proof Let τ^u be a realization of T^u . Notice that the set of rooted trees that can be transformed into τ^u via the unrooting operation is exactly the equivalence class $[\tau^u] = \{\tau_e^u : e \in E(\tau^u)\}$. The probability that T^u takes value of τ^u therefore is

$$p(T^u = \tau^u) = \sum_{\tau \in [\tau^u]} p_{\text{sbn}}(\tau) = p_{\text{sbn}}(T^u = \tau^u).$$

■

Note that ancestral sampling for Bayesian networks allows us to sample from SBNs in linear time $O(N)$ and the unrooting operation takes constant time, sampling rooted/unrooted trees from SBNs therefore can be done in linear time $O(N)$.

4. Variational Bayesian Phylogenetic Inference

In this section, we present our variational approach for approximate Bayesian phylogenetic inference. We first introduce the general variational framework with approximating distributions and training objectives. We then discuss some practical and efficient parameterization strategies for our variational approximations. Finally, we introduce several efficient computational techniques, including low-variance stochastic gradient estimators for variational training of the model parameters and a linear algorithm for computing the gradient of SBN parameters which is essential for stochastic gradient estimation.

4.1 A General Framework

As SBNs can provide a family of distributions over tree topologies with efficient probability estimation and sampling algorithms, they stand out as an essential building block to perform variational inference for phylogenetics. Given a family of SBN-based approximating distributions $Q_\phi(\tau)$ over phylogenetic tree topologies, where ϕ denotes the CPDs of SBNs, we can prescribe our variational approximation as

$$Q_{\phi,\psi}(\tau, \mathbf{q}) = Q_\phi(\tau)Q_\psi(\mathbf{q}|\tau)$$

where $Q_\psi(\mathbf{q}|\tau)$ is another family of conditional probability distributions over the branch lengths given the tree topology τ (see more detailed construction in Section 4.3), with ψ being the branch length variational parameters. Variational inference now amounts to finding a candidate of this family that minimizes the Kullback-Leibler (KL) divergence to the exact posterior

$$\phi^*, \psi^* = \arg \min_{\phi, \psi} D_{\text{KL}}(Q_{\phi,\psi}(\tau, \mathbf{q}) \| p(\tau, \mathbf{q}|\mathbf{Y})) \tag{11}$$

which is equivalent to maximizing the evidence lower bound (ELBO)

$$L(\phi, \psi) = \mathbb{E}_{Q_{\phi, \psi}(\tau, \mathbf{q})} \log \left(\frac{p(\mathbf{Y}|\tau, \mathbf{q})p(\tau, \mathbf{q})}{Q_{\phi}(\tau)Q_{\psi}(\mathbf{q}|\tau)} \right) \leq \log p(\mathbf{Y}). \quad (12)$$

While the above single-sample lower bound is commonly used in the VI literature, it is prone to heavily penalize samples that fail to explain the observations. As a result, the variational approximation tends to focus on the high-probability areas of the exact posterior, leading to simple local variational posteriors that underestimate the uncertainty of the exact posteriors². Given that phylogenetic posteriors are often multimodal and require strong exploration ability of the inference algorithms to get good performance, we propose an alternative training objective, a multi-sample lower bound that averages over multiple samples ($K > 1$) when estimating the marginal likelihood (Burda et al., 2016; Mnih and Rezende, 2016) as follows

$$L^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \log \left(\frac{1}{K} \sum_{i=1}^K \frac{p(\mathbf{Y}|\tau^i, \mathbf{q}^i)p(\tau^i, \mathbf{q}^i)}{Q_{\phi}(\tau^i)Q_{\psi}(\mathbf{q}^i|\tau^i)} \right) \leq \log p(\mathbf{Y}) \quad (13)$$

where $Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K}) \equiv \prod_{i=1}^K Q_{\phi, \psi}(\tau^i, \mathbf{q}^i)$. Compared to the single-sample lower bound, the multi-sample lower bounds encourage mode-covering behaviors, and hence facilitate the exploration efficiency in the phylogenetic tree space, especially when getting through the less likely territories between the posterior modes. The multi-sample lower bound gets tighter as the number of samples K increases and converges to the true marginal likelihood in the limit of infinitely many samples (Burda et al., 2016). Note that tighter lower bounds may also deteriorate the training of the variational approximations (Rainforth et al., 2019), and so we suggest using a moderate K in practice (see the experiments in Section 5 and Section 7 for more details). In the rest of the paper, we use the multi-sample lower bounds and refer to them as lower bounds for short.

4.2 Subsplit Support Estimation

As the CPDs of SBNs are associated with all consistent parent-child subsplit pairs, a complete parameterization of SBNs would require an exponentially increasing number of parameters. In practice, however, we find that the parent-child subsplit pairs observed from trees with relatively high posterior probabilities are a small subset of the entire set of subsplit pairs. Therefore, if we can find a sufficiently large collection of subsplits from these favorable trees and restrict the *support* of CPDs (where the conditional probabilities are allowed to be nonzero) accordingly, we can construct SBNs that cover most trees in the high posterior area with an affordable number of parameters. In this section, we explore several easily implemented options for subsplit support estimation on relatively small data sets (that allows us to compare with the exact posteriors), and leave a discussion of other possibilities to the conclusion.

More concretely, we tested four heuristic approaches that provide candidate tree topologies for this kind of subsplit support estimation for SBNs: bootstrapping (BS) based on three commonly used phylogeny reconstruction methods that are (1) neighbor joining (NJ), (2)

2. This is a behavior known as mode-seeking.




Figure 5: A comparison on four heuristic approaches for subsplit support estimation. **(Left)**: The sum of posterior probabilities of trees covered by SBNs with subsplit support estimated by different methods. **(Middle)**: The coverage percentages for root subsplits and subsplit pairs given by different methods. **(Right)**: The coverage efficiencies (the portion of root subsplits/subsplit pairs from the candidate trees that also appear in the ground truth collections) for root subsplits and subsplit pairs given by different methods. The letters R and S in parentheses represent the root subsplit and subsplit pair respectively. The error bars show one standard error based on 10 independent runs.

maximum parsimony (MP) and (3) maximum likelihood (ML); (4) short-run MCMC. The phylogenetic bootstrap performs bootstrapping on the sites of a multiple sequence alignment, which corresponds to the independent-evolution-among-sites assumption embodied in (2); see Felsenstein (1985) for details. All experiments were conducted on a real data set DS1 (Hedges et al., 1990) and the results were averaged over 10 replicates for each approach. Bootstrapping with neighbor joining and maximum parsimony were performed using PAUP* (Swofford, 2001). Bootstrapping with maximum likelihood was run in UFBoot (Minh et al., 2013). All bootstrapping approaches were run with 8,000 replicates. The classical MCMC Bayesian analyses were done in MrBayes (Ronquist et al., 2012), including an extremely long “golden run” of 10 billion iterations (sampled every 1,000 iterations with the first 25% discarded as burn-in) to form the ground truth and short runs of one million iterations (sampled every 100 iterations with the first 20% discarded, yielding 8,000 samples) to collect candidate tree topologies. We assume a uniform prior on the tree topology, an i.i.d. exponential prior with rate parameter 10.0 for the branch lengths and the simple Jukes and Cantor (1969) substitution model in the Bayesian setting.

We used three different criteria to evaluate the performance of different subsplit support estimation methods. The left plot in Figure 5 shows the covered posterior percentages of the resulting SBNs as a function of the number of trees used to estimate the support. The covered posterior percentages were estimated by aggregating the posterior probabilities of those sampled trees from the MCMC golden run that have non-zero SBN probabilities, i.e., are covered by the SBNs. We see that both bootstrapping with the maximum likelihood (ML+BS) and short run MCMC (MCMC_{init}) approaches achieve higher posterior coverage with smaller numbers of trees than the other methods. Bootstrapping with maximum

parsimony (MP+BS) manages to provide comparable posterior coverage, but requires more trees. Trees obtained from bootstrapping with neighbor joining (NJ+BS) are inadequate to cover the high posterior area. The middle plot in Figure 5 shows the coverage percentages of both root subsplits (solid) and subsplit pairs (dashed) for different methods, where the coverage percentages are defined as the portions of the ground truth root subsplits/subsplit pairs that appear in the candidate trees. We see that all methods performed quite well in the case of root subsplits. In the more challenging case of subsplit pairs, bootstrapping with maximum likelihood and short run MCMC, again, outperformed the other methods. The right plot in Figure 5 investigates the coverage efficiency as a function of the number of candidate trees, where the coverage efficiency was estimated in terms of the portion of root subsplits/subsplit pairs from the candidate trees that also appear in the ground truth collections. Clearly, we can see that short run MCMC provides the best coverage efficiency (close to 100%), followed by bootstrapping with maximum likelihood, and bootstrapping with maximum parsimony and neighbor joining performed the worst. Moreover, the coverage efficiency decreases as the number of trees increases for all methods except short run MCMC, so for these other methods this indicates that large coverage percentages may come with large portions of bad root subsplits/subsplit pairs³, especially for low quality candidate trees. Overall, compared to distance-based methods, likelihood-based approaches are more suitable for subsplit support estimation, in terms of posterior coverage, root subsplit and subsplit pair coverage percentage and efficiency. We will use bootstrapping with maximum likelihood and short run MCMC for subsplit support estimation in our experiments. Note that our empirical results here is on DS1, which is a relatively small data set. We leave a more thorough investigation on how these findings trend as tree sizes increase to future work (see Section 8 for more discussion).

4.3 Variational Parameterization

Given an appropriate set of candidate trees acquired via the aforementioned approaches, we can simply traverse these trees and collect root subsplits and parent-child subsplit pairs along the way to form the subsplit support. Now denote the set of root subsplits in the support as \mathbb{S}_r and the set of parent-child subsplit pairs collected during the tree traversals as $\mathbb{S}_{\text{ch|pa}}$. Let ϕ_s be the associated parameter for each element $s \in \mathbb{S}_r \cup \mathbb{S}_{\text{ch|pa}}$. We define the CPDs of SBNs according to the following equations

$$p(S_1 = s_1) = \frac{\exp(\phi_{s_1})}{\sum_{s_r \in \mathbb{S}_r} \exp(\phi_{s_r})}, \quad p(S_i = s | S_{\pi_i} = t) = \frac{\exp(\phi_{s|t})}{\sum_{s \in \mathbb{S}_{|t}} \exp(\phi_{s|t})}$$

where $\mathbb{S}_{|t}$ denotes the set of child subsplits that are compatible with the greater or smaller subclade in parent subsplit t , depending on whether S_i is the top child (i is even) or the bottom child (i is odd) of its parent respectively. Note that this parameterization implicitly assumes conditional probability sharing and is automatically consistent as the root subsplits and parent-child subsplit pairs are collected from tree topologies.

To accommodate the non-negative nature of the branch length parameters in phylogenetic models, we use the Log-normal distribution that has a positive support as our variational

3. This would increase the complexity of the search space and hence makes the training of variational approximations more difficult.

distribution. Furthermore, given a tree topology τ , we assume the branch lengths of τ are independently distributed. This leads to a simple approximate conditional density $Q_\psi(\mathbf{q}|\tau)$ that takes the following form

$$Q_\psi(\mathbf{q}|\tau) = \prod_{e \in E(\tau)} p^{\text{Lognormal}}(q_e \mid \mu(e, \tau), \sigma(e, \tau)). \quad (14)$$

Notice that a naive parameterization of (14) that assumes independent parameters for different tree topologies would require a large number of parameters, especially when the high posterior areas are diffuse. To alleviate this issue, we propose to parameterize the branch length distributions over the tree topology space via certain shared local splitting patterns. More specifically, instead of having separate parameters for each tree topology, we use a shared pool of parameters for these local splitting patterns that maps the tree topologies to their corresponding branch length variational parameters. We provide two concrete examples as follows.

4.3.1 SPLIT BASED PARAMETERIZATION

A simple choice of such local splitting pattern is the *split*, which is a bipartition (X_1, X_2) of the entire leaf label set \mathcal{X} such that $X_1 \cup X_2 = \mathcal{X}$, $X_1 \cap X_2 = \emptyset$. On a tree topology τ , each edge naturally corresponds to a split, the bipartition that consists of the leaf labels from both sides of the edge. Denote the corresponding split for edge e of topology τ as e/τ . Let $\psi_{e/\tau}^\mu$ and $\psi_{e/\tau}^\sigma$ be the associated mean and variance parameters for the corresponding root subsplit in \mathbb{S}_r ; the split-based parameterization then takes the following form

$$\mu(e, \tau) = \psi_{e/\tau}^\mu, \quad \sigma(e, \tau) = \psi_{e/\tau}^\sigma. \quad (15)$$

4.3.2 PRIMARY SUBSPLIT PAIR BASED PARAMETERIZATION

A potential problem with the above split-based parameterization (15) is that it implicitly assumes the branch lengths in different trees have the same distribution as long as they correspond to the same root subsplit, completely ignoring the variation of the branch length distributions among tree topologies. To capture this variation, one can use a more sophisticated parameterization where the variational parameters μ and σ may depend on other topological structures. More specifically, we use the following local structure for each edge in addition to the root subsplits.

Definition 10 (Primary Subsplit Pair) *Let e be an edge of a phylogenetic tree τ and say the corresponding root subsplit is $e/\tau = (W, Z)$ where $W \cup Z = \mathcal{X}$. Assume that at least one of W and Z , say W , contains more than one leaf label and hence has a subsplit (W_1, W_2) . We call the parent-child subsplit pair $(W_1, W_2)|(W, Z)$ a primary subsplit pair.*

The right plot in Figure 4 provides an example of primary subsplit pairs for a given edge e , where the root subsplit is $(V \cup Z, X \cup Y)$ and the primary subsplit pairs corresponding to the subsplits on either side of e are $(X, Y)|(V \cup Z, X \cup Y)$ and $(V, Z)|(V \cup Z, X \cup Y)$. Denote the primary subsplit pair(s) for edge e of tree τ as $e//\tau$. Notice that primary subsplit pairs are also parent-child subsplit pairs, we can assign additional parameters for each primary

subsplit pair in $\mathbb{S}_{\text{ch|pa}}$ and sum all the associated variational parameters for edge e to form the mean and variance of the corresponding branch length distribution (see Figure 12 in Appendix A for an illustration):

$$\mu(e, \tau) = \psi_{e/\tau}^\mu + \sum_{s \in e//\tau} \psi_s^\mu, \quad \sigma(e, \tau) = \psi_{e/\tau}^\sigma + \sum_{s \in e//\tau} \psi_s^\sigma. \quad (16)$$

Compared to (15), (16) allows contributions from the primary subsplit pairs that contain more topological information and hence could provide more flexible between-tree approximations. It is worth emphasizing that the above structured parameterizations (15) and (16) not only reduce the number of required parameters compared to the naive parameterization that assumes independent parameters for different tree topologies, but also enable joint learning of the branch length distributions across tree topologies due to those shared local splitting patterns.

4.4 Stochastic Gradient Estimators

Now that we have constructed our variational distributions for the phylogenetic posteriors, we can find the best approximation by maximizing the lower bound in (13) via stochastic gradient ascent. However, the naive stochastic gradient estimator obtained by direct differentiation, also known as the score function gradient estimator (Rubinstein and Kreimer, 1983; Glynn, 1990; Williams, 1992; Kleijnen and Rubinstein, 1996), is often of very large variance and is impractical for our purpose. Fortunately, various variance reduction techniques have been introduced in recent years including the control variates (Glynn and Szechtman, 2002; Paisley et al., 2012; Ranganath et al., 2014; Mnih and Gregor, 2014; Mnih and Rezende, 2016) for general latent variables and the reparameterization trick/pathwise derivative (Rubinstein, 1992; Glasserman and Ho, 1991; Pflug, 1996; Glasserman, 2013; Kingma and Welling, 2014; Rezende et al., 2014; Titsias and Lázaro-Gredilla, 2014) for continuous latent variables.

In what follows, we apply these techniques to different components of our latent variables and derive efficient gradient estimators with much lower variance, respectively. In addition, we also consider a closely related stable gradient estimator that is based on an alternative variational objective. More detailed derivations can be found in Appendix C.

4.4.1 THE VIMCO ESTIMATOR

Let $f_{\phi, \psi}(\tau, \mathbf{q}) = \frac{p(\mathbf{Y}|\tau, \mathbf{q})p(\tau, \mathbf{q})}{Q_{\phi}(\tau)Q_{\psi}(\mathbf{q}|\tau)}$. The naive gradient of the lower bound w.r.t. the tree variational parameters ϕ has the form (see a derivation in Appendix C):

$$\nabla_{\phi} L^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \sum_{j=1}^K \left(\hat{L}^K(\phi, \psi) - \tilde{w}^j \right) \nabla_{\phi} \log Q_{\phi}(\tau^j)$$

where $\hat{L}^K(\phi, \psi) = \log \left(\frac{1}{K} \sum_{i=1}^K f_{\phi, \psi}(\tau^i, \mathbf{q}^i) \right)$ and $\tilde{w}^j = \frac{f_{\phi, \psi}(\tau^j, \mathbf{q}^j)}{\sum_{i=1}^K f_{\phi, \psi}(\tau^i, \mathbf{q}^i)}$ are the stochastic lower bound and the self-normalized importance weight, respectively. While the importance weights are bounded, the stochastic lower bound could be problematic since it does not distinguish the samples according to their fitness for the data and assigns the same “learning signal” (Mnih and Gregor, 2014; Mnih and Rezende, 2016) to all of them. Moreover, the

magnitude of this learning signal can be extremely large, especially in the early stage of training when the samples from the variational distribution explain the data poorly. By leveraging these independent and identically distributed multiple samples, Mnih and Rezende (2016) defined a different local learning signal for each sample and proposed a variance-reduced gradient estimate, which leads to the following gradient estimate for the tree variational parameter ϕ :

$$\nabla_{\phi} L^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \sum_{j=1}^K \left(\hat{L}_{j|-j}^K(\phi, \psi) - \tilde{w}^j \right) \nabla_{\phi} \log Q_{\phi}(\tau^j) \quad (17)$$

where

$$\hat{L}_{j|-j}^K(\phi, \psi) := \hat{L}^K(\phi, \psi) - \log \frac{1}{K} \left(\sum_{i \neq j} f_{\phi, \psi}(\tau^i, \mathbf{q}^i) + \hat{f}_{\phi, \psi}(\tau^{-j}, \mathbf{q}^{-j}) \right)$$

is the per-sample local learning signal, with $\hat{f}_{\phi, \psi}(\tau^{-j}, \mathbf{q}^{-j})$ being some estimate of $f_{\phi, \psi}(\tau^j, \mathbf{q}^j)$ for sample j using the rest of the samples (e.g., the geometric mean). This gives the following VIMCO estimator

$$\nabla_{\phi} L^K(\phi, \psi) \simeq \sum_{j=1}^K \left(\hat{L}_{j|-j}^K(\phi, \psi) - \tilde{w}^j \right) \nabla_{\phi} \log Q_{\phi}(\tau^j) \text{ with } \tau^j, \mathbf{q}^j \stackrel{\text{iid}}{\sim} Q_{\phi, \psi}(\tau, \mathbf{q}). \quad (18)$$

4.4.2 THE REPARAMETERIZATION TRICK

The above VIMCO estimator also works for the branch length gradient. However, as the branch lengths are continuous latent variables, we can use an alternative approach, the reparameterization trick, to estimate the gradient. Note that the Log-normal distribution has a simple reparameterization: $q \sim \text{Lognormal}(\mu, \sigma^2) \Leftrightarrow q = \exp(\mu + \sigma\epsilon)$, $\epsilon \sim \mathcal{N}(0, 1)$, we can rewrite the lower bound as follows

$$L^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \log \left(\frac{1}{K} \sum_{j=1}^K \frac{p(\mathbf{Y}|\tau^j, g_{\psi}(\epsilon^j|\tau^j))p(\tau^j, g_{\psi}(\epsilon^j|\tau^j))}{Q_{\phi}(\tau^j)Q_{\psi}(g_{\psi}(\epsilon^j|\tau^j)|\tau^j)} \right)$$

where $g_{\psi}(\epsilon|\tau) = \exp(\boldsymbol{\mu}_{\psi, \tau} + \boldsymbol{\sigma}_{\psi, \tau} \odot \epsilon)$. Then the gradient of the lower bound w.r.t. ψ is (see Appendix C for derivation)

$$\nabla_{\psi} L^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \sum_{j=1}^K \tilde{w}^j \nabla_{\psi} \log f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j)) \quad (19)$$

where $\tilde{w}^j = \frac{f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j))}{\sum_{i=1}^K f_{\phi, \psi}(\tau^i, g_{\psi}(\epsilon^i|\tau^i))}$ is the same self-normalized importance weight as in equation (17). Therefore, we can form the Monte Carlo estimator of the gradient

$$\nabla_{\psi} L^K(\phi, \psi) \simeq \sum_{j=1}^K \tilde{w}^j \nabla_{\psi} \log f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j)) \text{ with } \tau^j \stackrel{\text{iid}}{\sim} Q_{\phi}(\tau), \epsilon^j \stackrel{\text{iid}}{\sim} p_{\epsilon}(\epsilon). \quad (20)$$

Algorithm 1 The variational Bayesian phylogenetic inference (VBPI) algorithm.

- 1: $\phi, \psi \leftarrow$ Initialize parameters
 - 2: **while** not converged **do**
 - 3: $\tau^1, \dots, \tau^K \leftarrow$ Random samples from the current approximating tree distribution
 $Q_\phi(\tau)$
 - 4: $\epsilon^1, \dots, \epsilon^K \leftarrow$ Random samples from the multivariate standard normal distribution
 $\mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 5: $\mathbf{g} \leftarrow \nabla_{\phi, \psi} L^K(\phi, \psi; \tau^{1:K}, \epsilon^{1:K})$ (Use any gradient estimator from Section 4.4)
 - 6: $\phi, \psi \leftarrow$ Update parameters using gradients \mathbf{g} (e.g. SGA)
 - 7: **end while**
 - 8: **return** ϕ, ψ
-

4.4.3 SELF-NORMALIZED IMPORTANCE SAMPLING ESTIMATOR

In addition to the standard variational formulation (11), one can reformulate the optimization problem by minimizing the reversed KL divergence, which is equivalent to maximizing the likelihood of the variational approximation

$$Q_{\phi^*, \psi^*}(\tau, \mathbf{q}), \text{ where } \phi^*, \psi^* = \arg \max_{\phi, \psi} \tilde{L}(\phi, \psi), \tilde{L}(\phi, \psi) = \mathbb{E}_{p(\tau, \mathbf{q} | \mathbf{Y})} \log Q_{\phi, \psi}(\tau, \mathbf{q}). \quad (21)$$

We can use an importance sampling estimator to compute the gradient of the objective

$$\begin{aligned} \nabla_{\phi} \tilde{L}(\phi, \psi) &= \mathbb{E}_{p(\tau, \mathbf{q} | \mathbf{Y})} \nabla_{\phi} \log Q_{\phi, \psi}(\tau, \mathbf{q}) = \frac{1}{p(\mathbf{Y})} \mathbb{E}_{Q_{\phi, \psi}(\tau, \mathbf{q})} \frac{p(\mathbf{Y} | \tau, \mathbf{q}) p(\tau, \mathbf{q})}{Q_{\phi}(\tau) Q_{\psi}(\mathbf{q} | \tau)} \nabla_{\phi} \log Q_{\phi}(\tau) \\ &\simeq \sum_{j=1}^K \tilde{w}^j \nabla_{\phi} \log Q_{\phi}(\tau^j) \text{ with } \tau^j, \mathbf{q}^j \stackrel{\text{iid}}{\sim} Q_{\phi, \psi}(\tau, \mathbf{q}) \end{aligned} \quad (22)$$

with the same importance weights \tilde{w}^j as in (17) (a detailed derivation can be found in Appendix C). This can be viewed as a multi-sample generalization of the wake-sleep algorithm (Hinton et al., 1995) and was first used in the *reweighted wake-sleep* algorithm (Bornschein and Bengio, 2015) for training deep generative models. We therefore call the gradient estimator in (22) the RWS estimator. Like the VIMCO estimator, the RWS estimator also provides gradients for the branch lengths. However, we find in practice the gradient estimator in (20) that uses the reparameterization trick is more useful and often leads to faster convergence, although it uses a different optimization objective. A better understanding of this phenomenon would be an interesting subject of future research.

All stochastic gradient estimators introduced above can be used in conjunction with stochastic optimization methods such as stochastic gradient ascent (SGA) (Robbins and Monro, 1951) or some of its adaptive variants (e.g. Adam (Kingma and Ba, 2015)) to maximize the lower bounds. See algorithm 1 below for a basic variational Bayesian phylogenetic inference (VBPI) approach. In the experiments, we will use the VIMCO estimator and RWS estimator for the tree topology parameter ϕ and the reparameterization trick for the branch length parameter ψ .

4.5 Efficient Computation of SBN Gradients

In this section, we describe some efficient computational methods for $\nabla_\phi \log Q_\phi(\tau)$, which is an essential ingredient for the stochastic gradient estimators derived in Section 4.4. Similarly to the tree probability computation, the computation of $\nabla_\phi \log Q_\phi(\tau)$ for rooted trees is straightforward and can be done in linear time as follows:

$$\begin{aligned} \nabla_\phi \log Q_\phi(\tau) &= \nabla_\phi \log p(S_1 = s_1) + \sum_{i>1} \nabla_\phi \log p(S_i = s_i | S_{\pi_i} = s_{\pi_i}) \\ &= \sum_{s \in \text{pcsp}(\tau)} \nabla_\phi \log p(s) \end{aligned} \quad (23)$$

where $\{s_i\}_{i \geq 1}$ is the subsplit decomposition of τ and $\text{pcsp}(\tau) = \{s_1\} \cup \{s_i | s_{\pi_i}\}_{i>1}$ is the set of the root subsplit and parent-child subsplit pairs of τ .

For unrooted trees, we have

$$\begin{aligned} \nabla_\phi \log Q_\phi(\tau^u) &= \nabla_\phi \log \sum_{\tau \in [\tau^u]} p_{\text{sbn}}(\tau) = \sum_{\tau \in [\tau^u]} \frac{\nabla_\phi p_{\text{sbn}}(\tau)}{\sum_{\tau \in [\tau^u]} p_{\text{sbn}}(\tau)} \\ &= \sum_{\tau \in [\tau^u]} \frac{p_{\text{sbn}}(\tau)}{\sum_{\tau \in [\tau^u]} p_{\text{sbn}}(\tau)} \nabla_\phi \log p_{\text{sbn}}(\tau) \\ &= \sum_{e \in E(\tau^u)} w_e \nabla_\phi \log p_{\text{sbn}}(\tau_e^\ddagger) \end{aligned} \quad (24)$$

where $w_e = \frac{p_{\text{sbn}}(\tau_e^\ddagger)}{\sum_{e \in E(\tau^u)} p_{\text{sbn}}(\tau_e^\ddagger)}$ is the rooting probability for edge e of τ^u induced by the SBN. Thanks to the efficient two-pass algorithm introduced in Section 3.2, we can compute $\{p_{\text{sbn}}(\tau_e^\ddagger) : e \in E(\tau^u)\}$ and hence the edge rooting probabilities $\{w_e : e \in E(\tau^u)\}$ in linear time. However, a naive implementation of (24) still requires quadratic time since $\nabla_\phi \log p_{\text{sbn}}(\tau_e^\ddagger)$ for each edge takes linear time.

Fortunately, by properly utilizing the structure of the tree topology, (24) can also be computed in linear time. A key observation is that if we expand $\nabla_\phi \log p_{\text{sbn}}(\tau_e^\ddagger)$ as in (23) for each edge e , most of the gradient terms $\nabla_\phi \log p(s)$ for the parent-child subsplit pairs will be shared across different edges. Therefore, we can switch the order of summation and accumulate the coefficients for the gradient terms first as follows:

$$\begin{aligned} \nabla_\phi \log Q_\phi(\tau^u) &= \sum_{e \in E(\tau^u)} w_e \sum_{s \in \text{pcsp}(\tau_e^\ddagger)} \nabla_\phi \log p(s) \\ &= \sum_{s \in \cup_{e \in E(\tau^u)} \text{pcsp}(\tau_e^\ddagger)} \nabla_\phi \log p(s) \sum_{e \in E(\tau^u)} w_e \delta_{s \in \text{pcsp}(\tau_e^\ddagger)}. \end{aligned} \quad (25)$$

Let $w_s = \sum_{e \in E(\tau^u)} w_e \delta_{s \in \text{pcsp}(\tau_e^\ddagger)}$. For the root subsplits and primary subsplit pairs, there is a unique edge e such that $s \in \text{pcsp}(\tau_e^\ddagger)$, and so $w_s = w_e$. For other parent-child subsplit pairs s , there is a corresponding subtree of τ^u that is rooted at the parent node and looks toward the child node and we denote it as $\tau^u|s$ (see the left plot in Figure 4). It turns out that s is a parent-child subsplit pair of τ_e^\ddagger if and only if e is not an edge of the subtree $\tau^u|s$. Therefore, w_s can be obtained from the cumulative probability of the root distribution along

the edges of the tree topology which can be computed via a postorder tree traversal (with a virtual root as in Section 3.2) that accumulates the edge rooting probabilities. Let \bar{w}_e be the cumulative probability at edge e . Depending on whether the virtual root r^v is in the subtree $\tau^u|s$ or not, we can compute w_s as follows:

$$w_s = \begin{cases} \bar{w}_e, & r^v \text{ is in } \tau^u|s \\ 1 - \bar{w}_e + w_e, & \text{otherwise.} \end{cases} \quad (26)$$

Finally, for each edge of the tree topology, we can collect the corresponding root subsplit, primary subsplit pairs and regular parent-child subsplit pairs with all possible orientations induced by the relative position of the edge w.r.t the virtual root, and hence complete the summation over the entire set of $\cup_{e \in E(\tau^u)} \text{pcsp}(\tau_e^\#)$ within a simple tree traversal. This concludes our linear time algorithm for the gradient computation of SBN parameters for unrooted trees.

5. Experiments on Unrooted Phylogenies

Throughout this section we evaluate the effectiveness and efficiency of our variational framework for Bayesian phylogenetic inference on unrooted phylogenetic trees, with comparisons to results that were obtained from intensive MCMC runs. The simplest SBN (the one with a full and complete binary tree structure) was used for the phylogenetic tree topology variational distribution; we have found it to provide sufficiently accurate approximation. The subsplit supports were estimated from either short MCMC runs using MrBayes (Ronquist et al., 2012) or ultrafast maximum likelihood phylogenetic bootstrap using UFBoot (Minh et al., 2013). We investigate the performance of the VIMCO estimator and the RWS estimator with different variational parameterizations for the branch length distributions, while varying the number of samples in the training objective to see how these affect the quality of the variational approximations. All models were implemented in PyTorch (Paszke et al., 2019) with the Adam optimizer (Kingma and Ba, 2015). Results were collected after 200,000 parameter updates. The reported KL divergences are over the discrete collection of phylogenetic tree topologies, from the estimated posteriors $\hat{Q}(\tau)$ (e.g., SBN distributions from VBPI or simple relative frequencies from standard MCMC) to the ground truth posterior $p(\tau|\mathbf{Y})$, and they are computed as follows

$$D_{\text{KL}}\left(p(\tau|\mathbf{Y})\|\hat{Q}(\tau)\right) = \sum_{\tau} p(\tau|\mathbf{Y}) \log\left(\frac{p(\tau|\mathbf{Y})}{\max\{\hat{Q}(\tau), \xi\}}\right), \quad (27)$$

where ξ is a small positive number that ensures that the denominator is positive (We used the machine precision for ξ in our experiments. See other choices of ξ in Appendix D). A low KL divergence means a high quality approximation to the exact posterior of the tree topologies. The code is available at <https://github.com/zcrabbit/vbpi-torch>.

We test the proposed variational Bayesian phylogenetic inference (VBPI) algorithms on 8 real data sets that are commonly used to benchmark phylogenetic MCMC methods (Lakner et al., 2008; Höhna and Drummond, 2012; Larget, 2013; Whidden and Matsen IV, 2015). We concentrate on the most challenging part of the phylogenetic model: joint learning of the tree topologies and the branch lengths. We assume a uniform prior on the tree topology, an i.i.d.

exponential prior with rate parameter 10.0 for the branch lengths and the simple Jukes and Cantor (1969) substitution model. We consider two different variational parameterizations for the branch length distributions as introduced in Section 4.3. In the first case, we use the simple split-based parameterization that assigns parameters to the splits associated with the edges of the trees. In the second case, we assign additional parameters for the primary subsplit pairs (PSP) to better capture the between-tree variation. We formed our ground truth posterior from an extremely long MCMC run of 10 billion iterations (sampled each 1,000 iterations with the first 25% discarded as burn-in) using MrBayes (Ronquist et al., 2012), and gathered the subsplit supports from trees generated by 10 replicates of 10,000 runs of the ultrafast maximum likelihood bootstrap (Minh et al., 2013). Following Rezende and Mohamed (2015), we use a simple annealed version of the lower bound which was found to provide better results. The modified bound is:

$$L_{\beta_t}^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \log \left(\frac{1}{K} \sum_{i=1}^K \frac{[p(\mathbf{Y}|\tau^i, \mathbf{q}^i)]^{\beta_t} p(\tau^i, \mathbf{q}^i)}{Q_{\phi}(\tau^i) Q_{\psi}(\mathbf{q}^i|\tau^i)} \right)$$

where $\beta_t = \min(1, 0.001 + t/100,000)$ is an inverse temperature schedule that goes from 0.001 to 1 after 99,900 iterations. We trained the variational approximations with VIMCO and RWS using 10 and 20 sample objectives and used Adam with the default learning rate 0.001 for both methods. We used an exponential learning rate schedule that decays the learning rate by a factor of 0.75 every 20,000 iterations. To encourage exploration in the tree topology space at the beginning, we initialize the variational parameters ϕ for SBNs to zero (which leads to uniformly distributed CPDs). The same settings are applied to all data sets.

5.1 Tree Topology Posterior Approximation

In this section, we examine the performance of VBPI for tree topology posterior approximation. The left and middle plots in Figure 6 show the resulting KL divergence to the ground truth on DS1 as a function of the number of parameter updates (the right plot will be described below in Section 5.2). The results for methods that adopt the simple split-based parameterization of variational branch length distributions are shown in the left plot. We see that the performance of all methods improves significantly as the number of samples K increases. The middle plot shows the results using PSP for variational parameterization, which clearly indicates that a better modeling of between-tree variation of the branch length distributions is beneficial for all methods with different stochastic gradient estimators and numbers of samples K . Specifically, PSP enables more flexible branch length distributions across tree topologies which makes the learning task much easier, as evidenced by the considerably smaller gaps between the methods. In both cases, we see that RWS learns faster at the beginning while VIMCO performs better in the end. To benchmark the learning efficiency of VBPI, we also compare to MrBayes 3.2.5 (Ronquist et al., 2012), a standard MCMC implementation for phylogenetic analysis on unrooted trees. We ran MrBayes with 4 chains and 10 runs for two million iterations, sampling every 100 iterations. For each run, we computed the KL divergence to the ground truth every 50,000 iterations with the first 25% discarded as burn-in.⁴ For a relatively fair comparison (in terms of the computational

4. Note that the “burn-in” here is only to trim the early samples and is different from the much longer standard burn-in for MCMC runs that ensures convergence.




Figure 6: Performance on DS1. **(Left)**: KL divergence for methods that use the simple split-based parameterization for the branch length distributions. **(Middle)**: KL divergence for methods that use PSP. **(Right)**: Per-tree marginal likelihood estimation (in nats): VBPI vs the Generalized stepping-stone estimator (GSS, described in the text). The number in parentheses specifies the number of samples used in the training objective. KL results are obtained from 10 independent runs. The marginal likelihood estimates for VBPI on each tree were obtained using 1,000 samples and the error bar shows one standard deviation over 100 independent runs.

complexity of likelihood evaluations), we compare 10 (i.e. $2 \cdot 20/4$) times the number of MCMC iterations with the number of 20-sample objective VBPI iterations.⁵ As most MCMC moves can leverage likelihood vector caching for a more efficient implementation, we also provide runtime comparisons where the runtime of VBPI is lower bounded by the likelihood and gradient computation in Appendix D. Although MCMC converges faster at the start, we see that VBPI methods (especially those with PSP) quickly surpass MCMC and arrive at good approximations with much fewer iterations. This is because VBPI updates the approximate distributions of tree topologies (e.g., SBNs) and branch lengths on the fly, which in turn allows guided exploration in the phylogenetic tree space. As a commonly used summarization of tree samples, we also compare the majority-rule consensus trees obtained from VBPI and the ground truth MCMC run. The results on DS1 are presented in Figure 15 in Appendix F. We see that VBPI and MCMC provide the same majority-rule consensus tree.

5.2 Lower Bound and Marginal Likelihood Estimation

To evaluate the overall approximation accuracy of VBPI for phylogenetic posteriors, we report the standard single-sample evidence lower bounds (ELBO) for variational approximations trained via the VIMCO gradient estimator in Table 1⁶. As before, we see that the ELBOs are all significantly improved with the PSP parameterization, in terms of both the estimated

5. The extra factor of $2/4$ is because the likelihood and the gradient can be computed together in twice the time of a likelihood computation (see Schadt et al. (1998)) and we run MCMC with 4 chains.

6. Note that this is different from the multi-sample lower bounds that we used as the training objectives.

Table 1: Variational lower bound and marginal likelihood estimation of different methods across 8 benchmark data sets for Bayesian phylogenetic inference. The marginal likelihood estimates of all variational methods were obtained via importance sampling using 1,000 samples, and the results (in units of nats) were averaged over 100 independent runs with standard deviation in parentheses. We ran the annealed importance sampling stepping-stone estimator (SS) (Xie et al., 2011) in MrBayes using default settings, with 4 chains for 10,000,000 iterations and sampled every 100 iterations. The results were averaged over 10 independent runs with standard deviation in parentheses. SS is not a variational method thus there are no entries for it in the upper part of the table.

data sets	(#Taxa,#Sites)	VIMCO(10)	VIMCO(20)	VIMCO(10)+PSP	VIMCO(20)+PSP	SS
Variational Lower Bound (nats)						
DS1	(27, 1949)	-7112.39(1.44)	-7112.60(2.10)	-7111.17(0.99)	-7111.57(1.10)	-
DS2	(29, 2520)	-26369.89(0.95)	-26369.80(0.94)	-26369.44(0.55)	-26369.57(0.75)	-
DS3	(36, 1812)	-33736.81(0.34)	-33736.94(0.45)	-33736.70(0.36)	-33736.82(0.37)	-
DS4	(41, 1137)	-13333.00(0.77)	-13333.25(0.84)	-13332.32(0.43)	-13332.70(0.45)	-
DS5	(50, 378)	-8219.00(0.24)	-8129.30(0.32)	-8218.44(0.19)	-8218.74(0.21)	-
DS6	(50, 1133)	-6729.78(0.74)	-6730.40(0.77)	-6729.21(0.36)	-6730.24(0.97)	-
DS7	(59, 1824)	-37335.44(0.13)	-37335.81(0.14)	-37335.25(0.12)	-37335.52(0.15)	-
DS8	(64, 1008)	-8659.45(0.65)	-8661.09(1.21)	-8655.55(0.40)	-8655.96(0.52)	-
Marginal Likelihood (nats)						
DS1	(27, 1949)	-7108.48(0.26)	-7108.46(0.23)	-7108.41(0.17)	-7108.41(0.16)	-7108.42(0.18)
DS2	(29, 2520)	-26367.73(0.10)	-26367.74(0.09)	-26367.73(0.09)	-26367.73(0.09)	-26367.57(0.48)
DS3	(36, 1812)	-33735.13(0.12)	-33735.12(0.12)	-33735.12(0.13)	-33735.12(0.11)	-33735.44(0.50)
DS4	(41, 1137)	-13330.02(0.31)	-13330.01(0.26)	-13329.96(0.24)	-13329.95(0.20)	-13330.06(0.54)
DS5	(50, 378)	-8214.78(0.56)	-8214.77(0.53)	-8214.67(0.44)	-8214.64(0.46)	-8214.51(0.28)
DS6	(50, 1133)	-6724.55(0.56)	-6724.54(0.52)	-6724.45(0.43)	-6724.41(0.43)	-6724.07(0.86)
DS7	(59, 1824)	-37332.16(0.51)	-37332.16(0.38)	-37332.05(0.33)	-37332.03(0.31)	-37332.76(2.42)
DS8	(64, 1008)	-8652.93(1.09)	-8651.33(0.87)	-8650.66(0.56)	-8650.64(0.53)	-8649.88(1.75)

means and variances. Although having a larger number of samples K in the training objective has been shown to facilitate topological exploration for better tree topology posterior estimation, we see that it does not necessarily improve the overall posterior approximation in terms of the ELBO (as observed and analyzed in Rainforth et al. (2019)). However, as we see next, tighter lower bounds indeed bias the variational approximations towards providing better marginal likelihood estimates, especially for the more challenging learning without the PSP parameterization.

The variational approximations provided by VBPI can be readily used to perform importance sampling for phylogenetic inference (see more details in Appendix G). Thanks to the structured amortization for the branch length distributions, VBPI can provide variational approximations on branch lengths for any tree topology that is covered by the subsplit support. We, therefore, can use them to estimate the marginal likelihood of tree topologies via importance sampling. The right plot in Figure 6 compares VBPI using VIMCO with 20-sample objective and PSP to the state-of-the-art generalized stepping-stone (GSS) (Fan et al., 2011) algorithm for estimating the marginal likelihood of trees in the 95% credible set of DS1. For GSS, we used 50 power posteriors and for each power posterior we ran




Figure 7: Scalability of VBPI and MCMC on Influenza Data. **(Left)**: KL divergence after 100,000 iterations. **(Middle)**: KL divergence after 200,000 iterations. **(Right)**: KL divergence after 400,000 iterations. KL results were obtained from 10 independent runs.

1,000,000 MCMC iterations, sampling every 1,000 iterations with the first 10% discarded as burn-in. The reference distribution (a distribution that matches the posterior more closely than the prior distribution; Fan et al., 2011) for GSS was obtained from an independent Gamma approximation using the maximum a posteriori estimate. The GSS estimates were computed using the phylogenetic package `physher` (Fourment and Holmes, 2014) (<https://github.com/4ment/physher>). We see that the estimates given by VBPI closely match GSS estimates (the scatter points are close to the orange diagonal line) and have small variance in general. Similarly, we can use VBPI to estimate the marginal likelihood of the data (i.e., the model evidence) via importance sampling. Table 1 (the lower part) also shows the marginal likelihood estimates using different VIMCO approximations and one of the state-of-the-art methods, the stepping-stone (SS) algorithm (Xie et al., 2011), which unlike GSS, is able to also integrate out tree topologies. For each data set, all methods provide estimates for the same marginal likelihood, with better approximation leading to lower variance. We see that VBPI using 1,000 samples is already competitive with SS using 100,000 samples and provides estimates with much less variance (hence more reproducible and reliable). Increasing the number of samples K is helpful for providing stabler marginal likelihood estimates, especially for learning without the PSP parameterization. With the extra flexibility enabled by PSP, the demand for larger number of samples used in the training objective has been greatly alleviated, making it possible to achieve high quality variational approximations with less samples.

5.3 Benchmarking Scalability Using Influenza Data

To benchmark the scalability of VBPI compared to MCMC, we curated a set of influenza data with an increasing number, i.e. 25, 50, 75, 100, of nested hemagglutinin (HA) sequences. The sequences were obtained from the NIAID Influenza Research Database (IRD) (Zhang et al., 2017) through the web site at <https://www.fludb.org/>, downloading all complete HA

Table 2: Marginal likelihood estimation of VBPI and SS across 4 benchmark data sets with nested hemagglutinin sequences. The marginal likelihood estimates of VBPI and SS were obtained and the results were presented the same way as in Section 5.2.

Methods	Subset size 25	Subset size 50	Subset size 75	Subset size 100
VBPI	-13378.38(0.06)	-18615.40(0.16)	-23681.85(0.27)	-28556.96(0.36)
SS	-13378.23(0.24)	-18615.82(1.57)	-23647.14(13.25)	-28176.80(47.16)

sequences that passed quality control, which were then subset to H7 sequences, and further downsampled using the Average Distance to the Closest Leaf (ADCL) criterion (Matsen et al., 2013). The purpose of this curation was to get sequence subsets with a suitable amount of phylogenetic uncertainty to benchmark the scalability of VBPI against MCMC, which are representative of moderately-dense sampling of viral sequences. The sequence subsets are available in the `vbpi-torch` GitHub repository. As before, we formed our ground truth posteriors from extremely long MCMC runs of 10 billion iterations (sampled each 1,000 iterations with a burn-in rate of 25%) using MrBayes. For VBPI, we used $K = 10$ samples in the training objective with PSP branch length parameterization and the VIMCO gradient estimator for training. The subsplit supports for VBPI were estimated based on 10 independent short MCMC runs of 1,000,000 iterations (sampled each 100 iterations with the first 25% discarded) respectively. For a relatively fair comparison to MCMC, as before, we ran MrBayes with 4 chains and 5 (i.e. 2·10/4) times the number of iterations when compared to VBPI, and sampled every 100 iterations with the first 25% discarded as burn-in. Figure 7 shows the KL divergence of the estimated posteriors from VBPI and MCMC to the ground truth at 100,000, 200,000, and 400,000 VBPI iterations. We see that VBPI scales much better than MCMC as the number of sequences increases, especially for a small computational budget (100,000 iterations). The number of iterations that VBPI takes for convergence appears to be quite steady across different sequence subsets, while MCMC requires many more, especially for large number of sequences. We also tested the performance of VBPI and SS on marginal likelihood estimation and the results are presented in Table 2. We see that VBPI and SS provides similar marginal likelihood estimates when the subset size is small. As the subset size increases, SS tends to provide higher marginal likelihood estimates than VBPI. However, the variances of these estimates also increases drastically, indicating the decreasing reliability of those marginal likelihood estimates. In contrast, the variances of VBPI estimates are much smaller and more stable. Note that SS has been found to overestimate the marginal likelihood when diffuse priors are used (Baele et al., 2016) due to the notorious difficulty of MCMC sampling from vague priors. Here we are following standard practice for unrooted Bayesian phylogenetics by using a uniform distribution on tree topologies, which certainly counts as a diffuse prior. Given that the SBN approximations are close to the tree topology posteriors (KL divergence is small, see Figure 7) and the support of Lognormal distributions covers all positive real numbers, the smaller variances of VBPI estimates here would imply more reliable estimates. Also, we did

observe in our experiments that the estimates given by SS are often dominated by some “large” values from one of the multiple runs as it uses the `logmeanexp` function to summarize the results from multiple runs. That may explain why SS tends to overestimate the marginal likelihood when the variance is large. As importance sampling scales exponentially in the KL divergence between the proposal and the target (Chatterjee and Diaconis, 2018), this challenge is also present when using variational approximations for importance sampling. However, it should be less of a problem in our setting as the divergence between variational approximation and the posterior is much smaller than that between the prior and the posterior for SS; we leave a more thorough investigation for future work.

6. VBPI for Time Measured Phylogenies

The VBPI framework introduced in Section 4 is mainly designed for unrooted phylogenetic models. When it comes to population genetic inference, rooted time-measured phylogenetic trees (also known as *time trees*) are adopted to allow estimation of divergence times between sequences and the underlying population dynamics (Pybus et al., 2000; Drummond et al., 2005; Sagulenko et al., 2018). The “branch lengths” of a time tree, therefore, correspond directly to elapsed time and each node of the tree is associated with a calendar time that reflects its known or inferred date⁷. Compared to unrooted trees, time trees are often subject to more complicated constraints induced by the additional sampling times for the data sequences and the evolutionary direction of the tree topologies. As a result, the priors used for time trees are more complicated than those for unrooted trees, and the variational approximations need to be more carefully designed to handle these constraints efficiently. In this section, we will describe several essential ingredients that extend the VBPI framework to time tree models. We start with a brief introduction to the coalescent model that serves as a popular prior for time trees, followed by two typical coalescent priors that allow efficient computation of the coalescent likelihoods. We then discuss a tree-based reparameterization for constructing variational approximations to the constrained posteriors of time trees, and appropriate variational parameterization techniques for branch length distributions over tree topologies.

6.1 Coalescent Background

Kingman’s coalescent model (Kingman, 1982) is a commonly used prior for time trees, providing a stochastic process that produces genealogies relating the observed molecular sequences. Given a random population sample of N sequences, the process starts at a sampling time $t = 0$ and proceeds backward in time as t increases, merging lineages one pair at a time until the time to the most recent common ancestor (TMRCA) of the sample is reached. Those merging times are called *coalescent times*. The coalescent model is parameterized in terms of the *effective population size*, a dynamic model parameter that controls the way that different lineages merge along time. Let $N_e(t)$ be the time-dependent effective population size. Suppose we observe a time tree of N sequences sampled at time 0, and the internal nodes have coalescent times $t_{N-1} < \dots < t_1$. The joint density of the

7. These rescaled branch lengths will be translated back into standard branch lengths that reflect the amount of molecular evolution via molecular clock models. We will revisit it later in Section 6.3.

coalescent times $t_N = 0 < t_{N-1} < \dots < t_1$ is obtained by multiplying the conditional probabilities for each coalescent event (Griffiths and Tavaré, 1994)

$$p(\mathbf{t}|N_e(t)) = \prod_{k=2}^N \frac{A_k}{N_e(t_{k-1})} \exp\left(-\int_{t_k}^{t_{k-1}} \frac{A_k}{N_e(t)} dt\right), \quad (28)$$

where $\mathbf{t} = (t_1, \dots, t_N)$ and $A_k = \binom{k}{2}$. The sampling time may be different as well, see more details on coalescent models with *heterochronous* sampling times in Appendix E. As the node times \mathbf{t} are often called node heights, we will use node heights in the sequel.

6.2 Coalescent Priors

The coalescent densities introduced above require appropriate forms for $N_e(t)$ to be properly estimated. There are many prior distributions to model the evolution of the population size dynamics over time, known as coalescent priors (Pybus et al., 2000; Minin et al., 2008; Gill et al., 2012). In this work, we consider the following two commonly used coalescent priors.

6.2.1 CONSTANT POPULATION SIZE

This model assumes that the population size has remained constant through time, that is,

$$N_e(t) = N_e = \exp(\gamma_e), \quad t \geq 0$$

where γ_e is the only parameter. In Bayesian settings, we may assume that γ_e has a $\mathcal{N}(\mu_0, \sigma_0^2)$ prior.

6.2.2 SKYRIDE

As the true parametric form for $N_e(t)$ is usually unknown, non-parametric coalescent priors are often adopted to provide flexible modeling of complicated demographic trajectories. Let $t_{k_{N-1}} < t_{k_{N-2}} < \dots < t_{k_1} = t_1$ denote the coalescent times for the internal nodes. The Skyride (Minin et al., 2008) is a non-parametric model that assumes $N_e(t)$ can change its value only at coalescent times,

$$N_e(t) = \exp(\gamma_n), \quad t_{k_n} < t \leq t_{k_{n-1}}, \quad n = 2, \dots, N$$

where $\boldsymbol{\gamma} = (\gamma_2, \dots, \gamma_N)$ are model parameters. Moreover, to encourage the smoothness of the population size dynamics over time, Minin et al. (2008) introduces a Gaussian Markov random field (GMRF) prior that penalizes the differences between successive components of $\boldsymbol{\gamma}$ as

$$p(\boldsymbol{\gamma}|\lambda) \propto \lambda^{\frac{N-2}{2}} \exp\left(-\frac{\lambda}{2} \sum_{n=2}^{N-1} \frac{(\gamma_{n+1} - \gamma_n)^2}{\delta_n}\right) \quad (29)$$

where λ is the overall precision of the GMRF and $\{\delta_n\}_{n=2}^{N-1}$ are temporal smoothing coefficients. As no prior knowledge about the smoothness of the effective population size trajectory (i.e. function in terms of time) is available, a relatively uninformative gamma prior is often assigned to the GMRF precision parameter λ ,

$$p(\lambda) = \frac{b^a}{\Gamma(a)} \lambda^{a-1} e^{-b\lambda}. \quad (30)$$

Note that the gamma prior (30) is a conjugate prior to the GMRF model (29), we therefore can integrate out λ to obtain the marginal prior for γ (Gelman et al., 2003),

$$p(\gamma) \propto \frac{b^a}{\Gamma(a)} \cdot \frac{\Gamma(a + \frac{N-2}{2})}{\left(b + \frac{1}{2} \sum_{n=2}^{N-1} \frac{(\gamma_{n+1} - \gamma_n)^2}{\delta_n}\right)^{a + \frac{N-2}{2}}}. \quad (31)$$

These piecewise constant population size trajectories allow easy computation of the coalescent densities introduced in Section 6.1 which we denote as $p(\mathbf{t}|\gamma)$ in the sequel.

6.3 Posterior Approximation and Variational Parameterization

As the node heights \mathbf{t} on a rooted time tree τ measure the duration of evolution in terms of time, molecular clock models are often employed to transform the node heights to branch lengths that reflect the amount of evolution in terms of the expected number of substitutions per site. For the sake of simplicity, we use the strict clock model that assumes a constant evolutionary rate r across different branches, leaving more flexible clock models (e.g. Drummond and Suchard (2010)) to future work. Given the evolutionary rate r , the branch length from the parent node i to the child node j is $q_{ij} = r(t_i - t_j)$. We, therefore, denote the branch lengths \mathbf{q} as a function of \mathbf{t} and r : $\mathbf{q} = \mathbf{q}(\mathbf{t}, r)$. With appropriate choices of the coalescent prior $p(\gamma)$ as described in Section 6.2, and the rate prior $p(r)$, the posterior distribution can be written as

$$p(\tau, \mathbf{t}, \gamma, r | \mathbf{Y}) \propto p(\mathbf{Y} | \tau, \mathbf{q}(\mathbf{t}, r)) p(\mathbf{t} | \gamma) p(\gamma) p(r). \quad (32)$$

Similarly as before, we prescribe a family of variational distributions as

$$Q_{\phi, \psi, \omega, \nu}(\tau, \mathbf{t}, \gamma, r) = Q_{\phi}(\tau) Q_{\psi}(\mathbf{t} | \tau) Q_{\omega}(\gamma | \tau) Q_{\nu}(r | \tau) \quad (33)$$

where $Q_{\phi}(\tau)$ represents a family of SBN-based distributions over rooted tree topologies, and $Q_{\psi}(\mathbf{t} | \tau)$, $Q_{\omega}(\gamma | \tau)$, $Q_{\nu}(r | \tau)$ are the conditional distributions for \mathbf{t} , γ , r , respectively, that are assumed to be independent given the tree topology τ .

Compared to other parameters, constructing appropriate conditional distributions for the node heights needs more careful consideration. Unlike the branch lengths for unrooted trees, the node heights \mathbf{t} for rooted time trees are subject to certain constraints induced by the sampling times and tree topologies. More specifically, the heights of the parent nodes are required to be greater than the heights of their children and the heights of the tip nodes are given by the corresponding sampling times. To cope with these constraints, we reparameterize the node heights as follows (similar reparameterizations have also been used before (Yang, 2007; Fourment and Darling, 2019; Ji et al., 2023)). First, we initialize the heights (and the lower bounds) of the tip nodes to the sampling times. Then, we run a postorder traversal and compute the lower bounds for the node heights when we visit internal node i via

$$t_i^{\text{lb}} = \max\{t_j^{\text{lb}} : j \in \text{ch}(i)\}.$$

These lower bounds handle the constraints from the sampling times. When the sampling times are all 0, we can simply set these lower bounds to 0. Now, we traverse the tree in a

preorder fashion and reparametrize the node heights using a non-negative parameter T for the root height and parameters θ as follows. We first visit the root node and

$$t_1 = t_1^{\text{lb}} + T, \quad T > 0. \quad (34)$$

When we visit internal node i , we set the node height to (recalling that π_i is the parent of i)

$$t_i = t_{\pi_i} - \theta_i(t_{\pi_i} - t_i^{\text{lb}}), \quad 0 < \theta_i < 1. \quad (35)$$

Let \mathcal{C} denote the set of indices for the internal nodes. The remaining constraints for T and θ can be completely removed with the following transforms

$$T = \exp(\alpha_1), \quad \theta_i = \frac{1}{1 + \exp(-\alpha_i)}, \quad \forall i \in \mathcal{C} \setminus \{1\}, \quad (36)$$

where $\{\alpha_i \in \mathbb{R} : i \in \mathcal{C}\}$ are unconstrained parameters for node heights. This completes our reparameterization of the conditional node height distributions.

Thanks to the hierarchical structure of the tree topology, the probability density function of the transformed node heights is computable given a tractable density for the base distribution of $\alpha|\tau$ (Fourment and Darling, 2019; Ji et al., 2023). First, the transform defined in (34) and (35) has a triangular Jacobian matrix whose determinant can be readily computed as $\prod_{i \in \mathcal{C} \setminus \{1\}} (t_{\pi_i} - t_i^{\text{lb}})$. Second, transform (36) has a diagonal Jacobian matrix whose determinant is $T \cdot \prod_{i \in \mathcal{C} \setminus \{1\}} \theta_i(1 - \theta_i)$. Due to the change of variable formula, we have

$$Q_\psi(\mathbf{t}|\tau) = Q_\psi(\alpha|\tau) \left(\prod_{i \in \mathcal{C} \setminus \{1\}} (t_{\pi_i} - t_i^{\text{lb}}) \cdot T \cdot \prod_{i \in \mathcal{C} \setminus \{1\}} \theta_i(1 - \theta_i) \right)^{-1} \quad (37)$$

where $Q_\psi(\alpha|\tau)$ is the base distribution.

A simple choice for the base distribution is the diagonal Normal distribution

$$Q_\psi(\alpha|\tau) = \mathcal{N}(\alpha_r | \mu(\tau), \sigma^2(\tau)) \cdot \prod_{e \in E(\tau)} \mathcal{N}(\alpha_e | \mu(e, \tau), \sigma^2(e, \tau))$$

where α_r is the parameter associated with the root height parameter T . Similarly as before, we can use the similarity of local splitting patterns to parameterize our node height variational approximations across different rooted tree topologies. More specifically, we can use the root subsplit for the root height parameter α_r . For the other node height parameters $\{\alpha_e\}_{e \in E(\tau)}$, we can follow the evolutionary direction of the rooted trees and use the clades (i.e., sets of descendant leaf nodes' labels) instead. Given a rooted tree τ , let r/τ denote the root subsplit and $e|\tau$ denote the clade for edge e . Let $\mathbb{S}_r, \mathbb{S}_c$ be the sets of root subsplits and clades in the support respectively. We assign parameters for μ, σ for each element in $\mathbb{S}_r \cup \mathbb{S}_c$ and parameterize the mean and variance of $\alpha|\tau$ as follows:

$$\mu(\tau) = \psi_{r/\tau}^\mu, \quad \sigma(\tau) = \psi_{r/\tau}^\sigma, \quad \mu(e, \tau) = \psi_{e|\tau}^\mu, \quad \sigma(e, \tau) = \psi_{e|\tau}^\sigma. \quad (38)$$

Again, more sophisticated parameterizations can be designed by including more tree-dependent local structures. For example, for the root height distribution, we can introduce

additional parameters for the associated primary subplit pairs; for the other node height distributions, we can introduce additional parameters for the leafward primary subplit pairs and the subsplits for the clades. All the associated variational parameters are then summed up to form the mean and variance of the corresponding node height distribution. This parameterization can be viewed as a straightforward adaption of the primary subplit parameterization presented in Section 4.3.2 and we refer to it as the primary subplit parameterization in the sequel.

Compared to the node height parameter α , the population size parameter γ and rate parameter r are often much less related to the tree topology τ when a strict clock model is assumed. For the sake of simplicity, we remove the dependency of γ and r on τ from our variational approximation and assume $Q_\omega(\gamma)$ and $Q_\nu(r)$ take diagonal Normal and Log-normal densities respectively. The best variational approximation from our variational family (33) then can be found by maximizing the multi-sample lower bound

$$L^K(\phi, \psi, \omega, \nu) = \mathbb{E}_{Q_{\phi, \psi, \omega, \nu}(\tau^{1:K}, \mathbf{t}^{1:K}, \gamma^{1:K}, r^{1:K})} \log \left(\frac{1}{K} \sum_{i=1}^K \frac{p(\mathbf{Y}|\tau^i, \mathbf{q}^i)p(\mathbf{t}^i|\gamma^i)p(\gamma^i)p(r^i)}{Q_\phi(\tau^i)Q_\psi(\mathbf{t}^i|\tau^i)Q_\omega(\gamma^i)Q_\nu(r^i)} \right)$$

as before, using stochastic gradient estimators introduced in Section 4.4.

7. Experiments on Time Trees

In this section, we analyzed two data sets to evaluate the accuracy of the posterior approximations provided by the proposed VBPI algorithms on time tree models. As before, we focus on the most challenging part: joint learning of the tree topology, node heights and the population size parameters, and assume the simple Jukes and Cantor (1969) substitution model. The traditional MCMC analyses were done using BEAST (Suchard et al., 2018), a commonly used MCMC implementation for phylogenetic analysis on time trees. In the first example, we analyzed a set of heterochronous Dengue virus sequences (Lanciotti et al., 1997) under the strict clock model and a constant size coalescent prior on the tree. This data set contains 17 whole-genome sequences (each with 1485 nucleotide sites) and can be downloaded from the BEAST Github repository (<https://github.com/beast-dev/beast-mcmc/tree/master/examples/Data>). The corresponding phylogenetic posterior is informative enough to allow us to compare VBPI and the classic MCMC method on the tree topology posterior estimation. The (rooted) subplit support was gathered from 10 independent short MCMC runs of 1,000,000 iterations (sampled each 100 iterations with the first 25% discarded). In the second example, we analyzed a set of 63 RNA sequences of type 4 from the E1 region of the hepatitis C virus (HCV) genome that were sampled in 1993 in Egypt. We used a fixed substitution rate of 7.9×10^{-4} per site per year, as suggested in previous studies (Pybus et al., 2001; Fourment and Darling, 2019). We considered both the constant size coalescent prior and the Bayesian skyride prior on the tree. For the Bayesian skyride tree prior, we used a Gaussian Markov random field prior for the effective population size trajectory and a relatively uninformative gamma prior $\Gamma(0.001, 0.001)$ for the precision parameter (Minin et al., 2008). This allows us to marginalize out the precision parameter and hence remove it from the set of parameters that require variational approximations (see Section 6.2.2). We used a uniform GMRF smoothing with equal temporal smoothing coefficients $\delta_2 = \dots = \delta_{62} = 1$. Similarly, the subplit support was




Figure 8: Performance on the Dengue virus data set. **(Left)**: KL divergence for methods that use the simple split-based parameterization for the branch length distributions. **(Middle)**: KL divergence for methods that use PSP. **(Right)**: The effective population size posterior estimates. The number in parentheses specifies the number of samples used in the training objective. The error bar shows one standard deviation over 10 independent runs.

obtained from 10 independent short MCMC runs of 1,000,000 iterations (sampled each 2,000 iterations with the first 50% discarded). In both examples, the ground truth posteriors were estimated from extremely long MCMC runs of 10 billion iterations (sampled each 10,000 iterations with the first 25% discarded as burn-in). The KL divergences from trained SBN distributions to the ground truth are reported, similarly as previously done in Section 5 (see equation (27)). Moreover, we also used similar parameterization strategies, i.e., the simple split-based parameterization and the primary subsplit pair (PSP) parameterization (see Section 6.3), a similar annealing schedule $\beta_t = \min(1, 0.001 + t/50,000)$, and the same training objectives and stochastic gradient estimators as in Section 5.

7.1 Tree Topology Posterior Approximation

We first analyzed the heterochronous Dengue virus data to examine the performance of VBPI for tree topology posterior estimation on time tree models. The left and middle plots in Figure 8 show the KL divergence to the ground truth as a function of the number of parameter updates. The results for methods using the simple split-based parameterization are shown in the left plot, and the results for methods using PSP are shown in the middle plot. Unlike the case of unrooted trees, we see that increasing the number of samples K does not necessarily lead to significantly better tree topology posterior approximations even when the simple split-based parameterization was adopted. However, a better modeling of between-tree variation of the branch length distributions is still beneficial for all methods with different stochastic gradient estimators and numbers of samples K . More specifically, we see two outcomes when using a more flexible PSP parameterization: first, the gaps between different methods become considerably smaller, and second, the KL divergences of all methods were significantly improved. This suggests a moderate K and a flexible

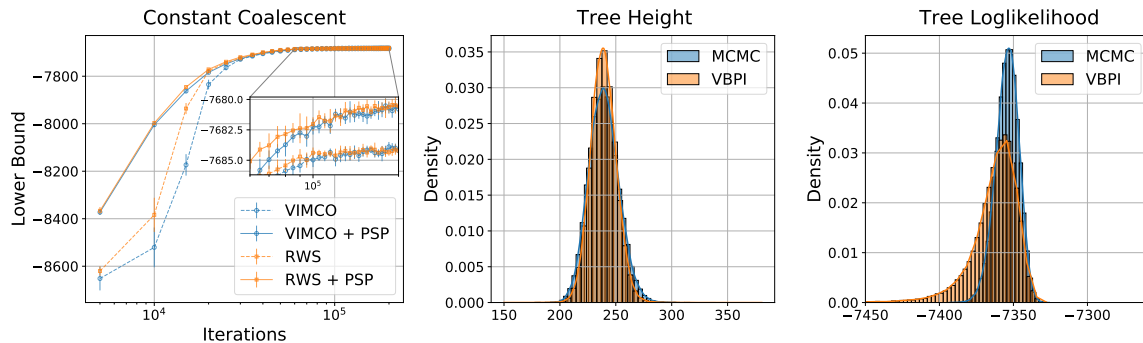


Figure 9: Performance on the HCV data set with a constant coalescent prior. **(Left)**: Lower bounds. **(Middle)**: Tree height posteriors. **(Right)**: Tree Likelihoods from posterior samples. All VBPI methods use 10 samples in the training objective. The error bar shows one standard deviation over 10 independent runs.

cross-topology branch length parameterization would be a favorable combination for VBPI on time trees. As before, we see that RWS learns faster at the beginning while VIMCO performs better in the end. To benchmark the learning efficiency of VBPI, we also compared to MCMC using BEAST 1.10.4 (Suchard et al., 2018). For that, we conducted 10 independent MCMC runs in BEAST, each for 8,000,000 iterations, and sample every 100 iterations. For each run, we computed the KL divergence to the ground truth every 200,000 iterations with the first 25% discarded as burn-in. As before, we compare 40 times the number of MCMC iterations with the number of 20-sample objective VBPI iterations to account for the additional gradient computation and multiple samples required by the training objectives. Similarly, we see that after a relatively slower starting phase, VBPI methods outperform MCMC and arrive at good approximations (especially those with PSP) with much fewer iterations.

7.2 Lower Bound and Population Dynamics Estimation

In this section, we analyzed the convergence behavior of VBPI and its approximation accuracy on parameters other than the tree topology on a set of HCV sequences using both a constant coalescent prior and the skyride model with a fixed substitution rate. The skyride model is flexible enough to capture the interesting underlying population dynamics for this data set (Minin et al., 2008). As suggested in Section 7.1, we used $K = 10$ samples in the training objective in all experiments. The posterior approximations from VBPI were obtained from 75,000 samples⁸ after training using the VIMCO gradient estimator. Figure 9 and 10 show the results for the constant coalescent prior and the skyride model respectively. In both figures, the left plots show the evidence lower bounds (ELBO) as a function of the number of parameter updates. We see that VIMCO and RWS gradient estimators

8. The same number of MCMC samples were used to form the ground truth posterior estimate.

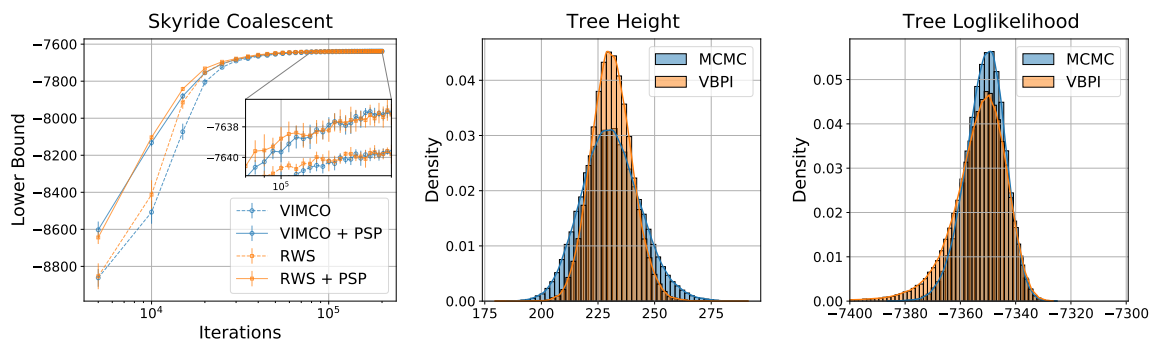


Figure 10: Performance on the HCV data set with the skyride coalescent prior. **(Left)**: Lower bounds. **(Middle)**: Tree height posteriors. **(Right)**: Tree Likelihoods from posterior samples. All VBPI methods use 10 samples in the training objective. The error bar shows one standard deviation over 10 independent runs.

perform similarly to each other with RWS being a little bit faster at the beginning. For both VIMCO and RWS, using the PSP parameterization provides much tighter lower bounds and smoother training curves, which again, demonstrates the importance of the flexibility of branch length parameterization across different topologies. The middle plots compare the posterior estimates of the tree height from MCMC and VBPI. We see that VBPI is more conservative than MCMC on tree height posterior estimation. This is partly due to the parameterization of the node heights in VBPI (see Section 6.3) where the tree height (i.e., root node height) has a strong impact on all the rest of the node heights on the tree. The right plots show the tree likelihoods of samples from MCMC and VBPI. We see that the likelihood values of trees sampled from VBPI are similar to those from MCMC, although more diffuse.

Finally, we evaluate the ability of VBPI to approximate the population dynamics of these HCV sequences (Figure 11). We see that for the simple constant coalescent model (the left plot), VBPI can provide high quality posterior approximation that matches the ground truth well. The right plot in Figure 8 shows a similar result on the Dengue virus data set. For the more challenging skyride model (the middle and the right plot), although the population trajectory estimated from VBPI still matches the ground truth, the 95% credible intervals tend to be narrower than those estimated from MCMC. This can, at least in part, be explained by the zero-forcing nature of the KL divergence used in VBPI and lack of flexibility of the mean field approximation for the population size parameters (see Section 6.3). Overall, VBPI can provide reasonably good posterior approximation of the effective population size, especially when the coalescent model is simple. We expect this problem of variance underestimation on the effective population size for complicated coalescent models to be alleviated with more flexible variational approximations and leave it for future work.




Figure 11: Posterior approximation of the effective population size parameters on the HCV data set. **(Left)**: Constant coalescent. **(Middle)**: Skyride coalescent using MCMC. **(Right)**: Skyride coalescent using VBPI. The solid lines represent the posterior medians and the shaded areas represent the 95% credible intervals.

8. Conclusion

In this work we propose VBPI as a general variational framework for Bayesian phylogenetic inference on the joint parameter space of phylogenetic models. At the core of VBPI lie subsplit Bayesian networks, an expressive probabilistic graphical model that provides distributions over the tree topology space, and an efficient structured amortization of the branch lengths over different tree topologies. We show that this provides a suitable variational family of distributions which can be trained by maximizing the multi-sample ELBO using stochastic gradient ascent. With guided exploration in the tree topology space (enabled by SBNs) and joint learning of the branch length distribution across tree topologies (via amortization), VBPI can provide competitive performance to MCMC methods especially when the posterior is diffuse. We used different branch length parameterizations and different numbers of samples in the multi-sample EBLO for training, and found that the flexibility of branch length parameterization across tree topologies has a stronger impact on the approximation accuracy and the complexity in optimization. Moreover, variational approximations provided by VBPI can be readily used for downstream tasks such as marginal likelihood estimation for model comparison via importance sampling. We report promising numerical results demonstrating the effectiveness and efficiency of VBPI on a benchmark of real data Bayesian phylogenetic inference problems.

The current VBPI framework also has many limitations. The main limitation is that it requires subsplit support estimation for variational parameterization. Subsplit support estimation is currently done with a pre-specified number of candidate tree topologies. As the exact posterior is often unknown, we may need a stopping criterion for sufficient support exploration in practice. The KL divergence, or equivalently, the ELBO, that are commonly used for variational inference is not suitable in this case due to its mode seeking behavior. As an alternative, we may conduct multiple short MCMC/bootstrap runs, and check if the newly collected subsplit supports contain significantly novel split patterns that haven't been seen before. In other lines of research, we are currently working on systematic means of

support estimation, generalizing Whidden et al. (2019) to the subsplit support context with several novel objective functions.

In addition, when the data are relatively insufficient and posteriors are diffuse, subsplit support estimation becomes challenging. However, compared to classical MCMC approaches in phylogenetics that need to traverse the enormous support of posteriors on complete trees to accurately evaluate the posterior probabilities, the subsplit-based parameterization in VBPI has a natural advantage in that it alleviates this issue by factorizing the uncertainty of complete tree topologies into local structures. This way, the SBN approximation given by VBPI can be viewed as a compact representation of the topological uncertainty of phylogenetic trees that can be learned in a timely manner. On the other hand, when there is minimal signal for structure of a region of the tree, a variational model will have a large number of parameters in that region.

Another limitation is that joint inference on model parameters could be challenging, especially when there exists an unidentifiability (e.g., rate matrix and branch length). Although the resulting posterior would still be proper when an appropriate prior is used for the related parameters, the correlation among these related parameters would be more complex. When this happens (e.g., relaxed clock models for time trees), our current variational approximation construction may be insufficient and more sophisticated design may be needed. For example, we may expand all individual variational building blocks with as many dimensions as global parameters. To reduce the number of parameters, we can attribute the global parameters to different variational building blocks when there exists conditional independence structures in the model. We can also try tying together parameters over different variational building blocks using neural networks. Another option is to provide more flexible variational families for the joint distribution of related parameters (e.g., normalizing flows which we discuss later for future work), instead of a product of conditionals.

Another potential limitation concerns optimization. Since the ELBO objective is often multimodal, the training of VBPI may end up with some local mode and hence lead to suboptimal variational approximations.

There are many extensions for future work. The first extension is to construct more flexible variational family of distributions. The current branch length variational distributions in VBPI are relatively simple and can be improved by using more flexible architectures (e.g., normalizing flows (Rezende and Mohamed, 2015)) to reduce the approximation error and amortization error. Recent work (Zhang, 2020, 2023) shows promising results in this direction where properly incorporating the inductive biases of phylogenetic models (e.g., the edges are orderless across tree topologies) seems to be the key. Similarly, we can design more flexible variational approximations for the effective population size parameters to alleviate the variance underestimation problem for time tree models.

Another potential extension is to find more efficient ways to infer the support of subsplits. Besides bootstrapping and short MCMC runs, good candidate trees can also be obtained from fast heuristic search methods (Whidden et al., 2019). Instead of fixing the support, we can also grow a dynamic support of candidate subsplits and keep adding “good” subsplits and removing “bad” subsplits in a manner analogous to maximum-likelihood phylogenetic inference. Moreover, we can also design appropriate stopping criteria as discussed above for these methods. Besides SBNs, one can also design distributions that have support over the

entire tree topology space, eliminating the necessity for support estimation (Xie and Zhang, 2023; Mimori and Hamada, 2023; Zhou et al., 2024).

When the variational approximations are not close enough to the exact posterior, pure importance sampling based marginal likelihood estimation using VBPI may become prohibitive. In this case, we can instead use these variational approximations as reference distributions for GSS approaches (Fan et al., 2011; Holder et al., 2014; Baele et al., 2016).

For the sake of simplicity, we only considered Jukes and Cantor (1969) substitution model and the strict clock model in this paper. Future methods will also consider including other interesting components of phylogenetic models, such as more complicated substitution models and the relaxed clock model, into our variational formulation.

Acknowledgments

This work supported by NIH grant AI162611. The research of Cheng Zhang was supported in part by National Natural Science Foundation of China (grant no. 12201014 and grant no. 12292983), the National Engineering Laboratory for Big Data Analysis and Applications, the Key Laboratory of Mathematics and Its Applications (LMAM) and the Key Laboratory of Mathematical Economics and Quantitative Finance (LMEQF) of Peking University. Frederick Matsen is an investigator of the Howard Hughes Medical Institute. The authors would like to thank Mathieu Fourment, Michael Karcher, Marc Suchard, and Christiaan Swanepoel for discussions. The authors are also grateful for the invaluable feedback from our anonymous reviewers and the Action Editor.

Appendix A. The PSP Parameterization




Figure 12: Branch length parameterization using primary subsplit pairs, which is the sum of parameters for a split and its neighboring subsplit pairs. Edge e represents a split (W, Z) . Parameterization for the variance is the same as for the mean.

Appendix B. Proofs of core properties of SBNs

In this section, we prove some important properties of subsplit Bayesian networks given in Section 3.

Lemma 4 There is a one-to-one mapping between rooted tree topologies on \mathcal{X} and compatible node assignments.

Proof For any rooted tree τ , with natural indexing, a compatible node assignment is easy to obtain by following the splitting process and obeying the compatibility requirement introduced in Section 3. Moreover, the depth of the splitting process is at most $N - 1$, where N is the size of \mathcal{X} . To show this, consider the maximum size of clades, d_i , in the i -th level of the process. The first level is for the root split, and $d_1 \leq N - 1$ since the subclades need to be proper subset of \mathcal{X} . For each $i > 1$, if $d_i > 1$, the splitting process continues and the maximum size of clades in the next level is at most $d_i - 1$, that is $d_{i+1} \leq d_i - 1$. Note that the split process ends when $d_i = 1$, its depth therefore is at most $N - 1$. The whole process is deterministic, so each rooted tree can be uniquely represented as a compatible assignment of $\mathcal{B}_{\mathcal{X}}^*$. On the other hand, given a compatible node assignment, we can simply remove all singleton clades. The remaining net corresponds to the true splitting process of a rooted tree, which then can be simply restored. ■

Proposition 6 For any consistently parameterized SBN, $\sum_{\tau} p_{\text{sbn}}(T = \tau) = 1$.

Proof Denote the set of all SBN assignments as \mathcal{A} . For any assignment $a = \{S_i = s_i^a\}_{i \geq 1}$, we have

$$p_{\text{sbn}}(a) = p(S_1 = s_1^a) \prod_{i > 1} p(S_i = s_i^a | S_{\pi_i} = s_{\pi_i}^a) > 0 \Rightarrow a \text{ is compatible.} \quad (39)$$

As a Bayesian network,

$$\sum_{a \in \mathcal{A}} p_{\text{sbm}}(a) = \sum_{s_1} p(S_1 = s_1) \prod_{i>1} \sum_{s_i} p(S_i = s_i | S_{\pi_i} = s_{\pi_i}) = 1.$$

By Lemma 4,

$$\sum_T p_{\text{sbm}}(T) = \sum_{a \sim \text{compatible}} p_{\text{sbm}}(a) = \sum_{a \in \mathcal{A}} p_{\text{sbm}}(a) = 1. \quad \blacksquare$$

Here $a \sim \text{compatible}$ denotes a sum over compatible node assignments as per Definition 3.

Appendix C. Gradients for The Multi-sample Objectives

In this section we will derive the gradients for the multi-sample objectives introduced in Section 4. The derivation presented here is a straightforward adaption of the stochastic gradient estimators introduced before (Mnih and Rezende, 2016; Kingma and Welling, 2014; Bornschein and Bengio, 2015), which we write out make explicit the slightly unusual setting of phylogenetics, which includes both a discrete tree topology and associated continuous parameters. We start with the lower bound

$$\begin{aligned} L^K(\phi, \psi) &= \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \log \left(\frac{1}{K} \sum_{j=1}^K \frac{p(\mathbf{Y} | \tau^j, \mathbf{q}^j) p(\tau^j, \mathbf{q}^j)}{Q_{\phi}(\tau^j) Q_{\psi}(\mathbf{q}^j | \tau^j)} \right) \\ &= \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \log \left(\frac{1}{K} \sum_{j=1}^K f_{\phi, \psi}(\tau^j, \mathbf{q}^j) \right). \end{aligned}$$

Using the product rule and noting that $\nabla_{\phi} \log f_{\phi, \psi}(\tau^j, \mathbf{q}^j) = -\nabla_{\phi} \log Q_{\phi}(\tau^j)$,

$$\begin{aligned} \nabla_{\phi} L^K(\phi, \psi) &= \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \nabla_{\phi} \log \left(\frac{1}{K} \sum_{j=1}^K f_{\phi, \psi}(\tau^j, \mathbf{q}^j) \right) + \\ &\quad \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \sum_{j=1}^K \frac{\nabla_{\phi} Q_{\phi}(\tau^j)}{Q_{\phi}(\tau^j)} \log \left(\frac{1}{K} \sum_{i=1}^K f_{\phi, \psi}(\tau^i, \mathbf{q}^i) \right) \\ &= \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \sum_{j=1}^K \frac{f_{\phi, \psi}(\tau^j, \mathbf{q}^j)}{\sum_{i=1}^K f_{\phi, \psi}(\tau^i, \mathbf{q}^i)} \nabla_{\phi} \log f_{\phi, \psi}(\tau^j, \mathbf{q}^j) + \\ &\quad \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \sum_{j=1}^K \log \left(\frac{1}{K} \sum_{i=1}^K f_{\phi, \psi}(\tau^i, \mathbf{q}^i) \right) \nabla_{\phi} \log Q_{\phi}(\tau^j) \\ &= \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \sum_{j=1}^K \left(\hat{L}^K(\phi, \psi) - \tilde{w}^j \right) \nabla_{\phi} \log Q_{\phi}(\tau^j). \end{aligned}$$

This gives the naive gradient of the lower bound w.r.t. ϕ .

Using the reparameterization trick, the lower bound has the form

$$\begin{aligned} L^K(\phi, \psi) &= \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \log \left(\frac{1}{K} \sum_{j=1}^K \frac{p(\mathbf{Y}|\tau^j, g_{\psi}(\epsilon^j|\tau^j))p(\tau^j, g_{\psi}(\epsilon^j|\tau^j))}{Q_{\phi}(\tau^j)Q_{\psi}(g_{\psi}(\epsilon^j|\tau^j)|\tau^j)} \right) \\ &= \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \log \left(\frac{1}{K} \sum_{j=1}^K f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j)) \right) \end{aligned}$$

Since ψ is not involved in the distribution with respect to which we take expectation,

$$\begin{aligned} \nabla_{\psi} L^K(\phi, \psi) &= \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \nabla_{\psi} \log \left(\frac{1}{K} \sum_{j=1}^K f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j)) \right) \\ &= \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \sum_{j=1}^K \frac{f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j))}{\sum_{i=1}^K f_{\phi, \psi}(\tau^i, g_{\psi}(\epsilon^i|\tau^i))} \nabla_{\psi} \log f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j)) \\ &= \mathbb{E}_{Q_{\phi, \epsilon}(\tau^{1:K}, \epsilon^{1:K})} \sum_{j=1}^K \tilde{w}^j \nabla_{\psi} \log f_{\phi, \psi}(\tau^j, g_{\psi}(\epsilon^j|\tau^j)). \end{aligned}$$

Next, we derive the gradient of the multi-sample likelihood objective used in RWS

$$\tilde{L}(\phi, \psi) = \mathbb{E}_{p(\tau, \mathbf{q}|\mathbf{Y})} \log Q_{\phi, \psi}(\tau, \mathbf{q}).$$

Again, $p(\tau, \mathbf{q}|\mathbf{Y})$ is independent of ϕ, ψ , and we have

$$\begin{aligned} \nabla_{\phi} \tilde{L}(\phi, \psi) &= \mathbb{E}_{p(\tau, \mathbf{q}|\mathbf{Y})} \nabla_{\phi} \log Q_{\phi, \psi}(\tau, \mathbf{q}) \\ &= \mathbb{E}_{Q_{\phi, \psi}(\tau, \mathbf{q})} \frac{p(\tau, \mathbf{q}|\mathbf{Y})}{Q_{\phi}(\tau)Q_{\psi}(\mathbf{q}|\tau)} \nabla_{\phi} \log Q_{\phi, \psi}(\tau, \mathbf{q}) \\ &= \frac{1}{p(\mathbf{Y})} \mathbb{E}_{Q_{\phi, \psi}(\tau, \mathbf{q})} f_{\phi, \psi}(\tau, \mathbf{q}) \nabla_{\phi} \log Q_{\phi}(\tau) \\ &\simeq \sum_{j=1}^K \frac{f_{\phi, \psi}(\tau^j, \mathbf{q}^j)}{\sum_{i=1}^K f_{\phi, \psi}(\tau^i, \mathbf{q}^i)} \nabla_{\phi} \log Q_{\phi}(\tau^j) \quad \text{with } \tau^j, \mathbf{q}^j \stackrel{\text{iid}}{\sim} Q_{\phi, \psi}(\tau, \mathbf{q}). \\ &= \sum_{j=1}^K \tilde{w}^j \nabla_{\phi} \log Q_{\phi}(\tau^j) \end{aligned}$$

The second to last step uses self-normalized importance sampling with K samples. $\nabla_{\psi} \tilde{L}(\phi, \psi)$ can be computed in a similar way.

Appendix D. Runtime Comparison between VBPI and MCMC

We provide approximate wall time comparison results in this section. That is, we work to understand what the wall time of phylogenetic inference using an optimized implementation of VBPI would be, including the cost of subsplit support estimation. Our current implementation is in Python, which uses automatic differentiation for likelihood gradient computation;

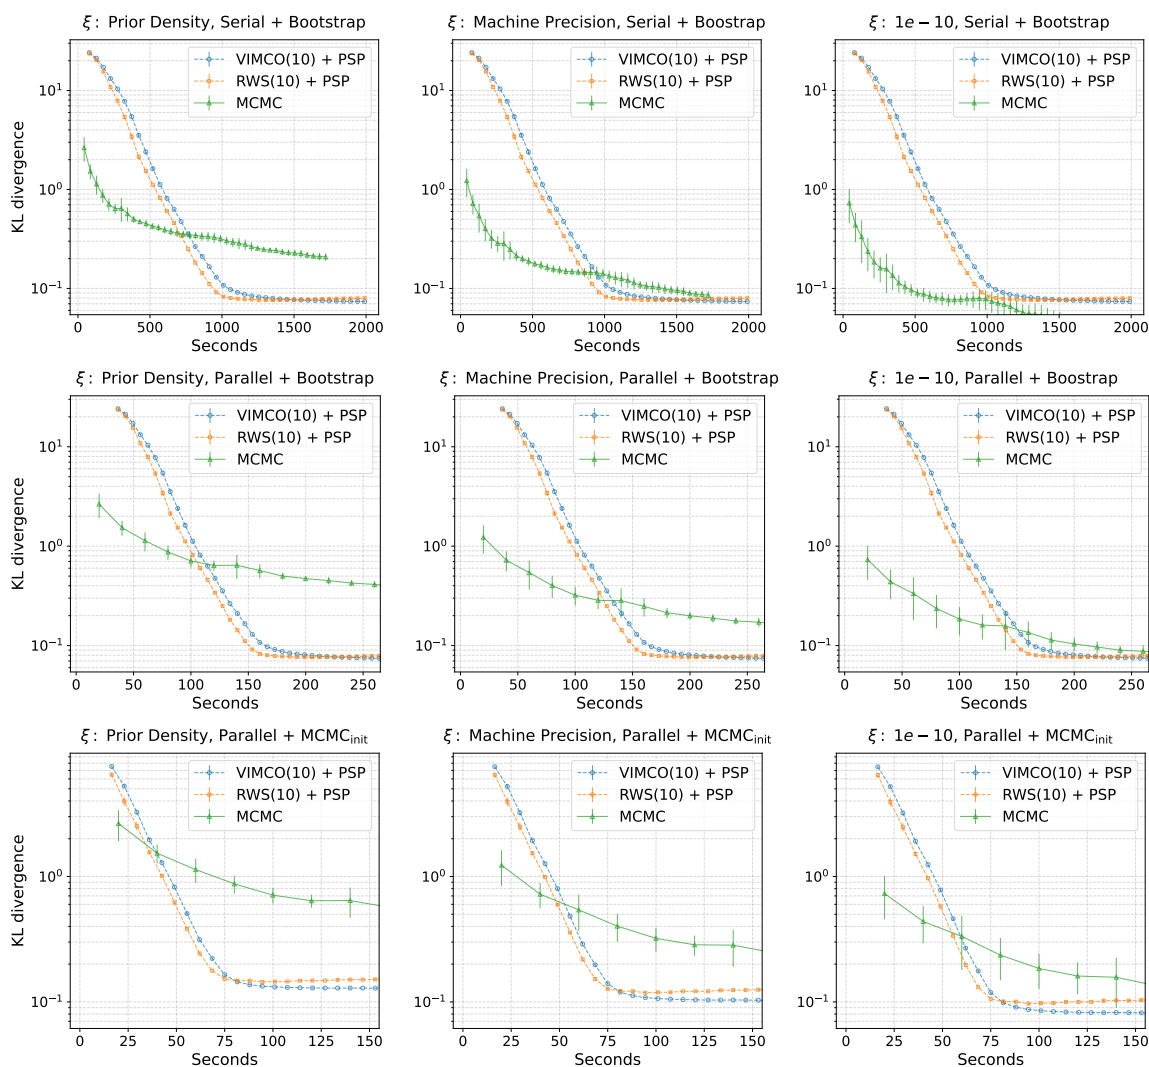


Figure 13: A runtime comparison between VBPI and MCMC on DS1. (**Top**): Serial implementation with bootstrap subsplit support. (**Middle**): Parallel implementation with bootstrap subsplit support. (**Bottom**): Parallel implementation with short run MCMC subsplit support. Each column corresponds to a different value of ξ . Error bars show one standard deviation over 10 independent runs.

this is inefficient compared to custom gradient code (Fourment et al., 2023). For example, on the DS1 example below, our python implementation takes about 14200 seconds for 200,000 iterations while the more customized `bito` library takes only around 1960 seconds for all involved phylogenetic likelihood and gradient computation. It is not straightforward to compare the number of phylogenetic likelihood computations between VBPI and MCMC as

most phylogenetic MCMC moves can leverage the caching trick for more efficient likelihood computation, without leaving a record of how many likelihood operations were performed.

Therefore we use the following strategy to estimate wall clock time, based on the assumption that the phylogenetic gradient will dominate run time. For VBPI, we record the number of gradient operations needed for a given number of iterations. For each such iteration, we lower bound the runtime by the likelihood and gradient computation for that sequence alignment using `bito` (<https://github.com/phylovi/bito>), which is a Python-interface C++ library that uses BEAGLE (Ayres et al., 2019) as the backend likelihood and gradient engine. We used $K = 10$ samples in the training objective with PSP branch length parameterization and both the VIMCO and RWS gradient estimators for training. The runtimes for VBPI also include the cost of computing a sample of tree topologies and gathering the set of subsplit supports from those trees.

For MCMC, we use MrBayes with BEAGLE backend as well. For each run, we ran MrBayes with 4 chains for 8,000,000 iterations, sampling every 100 iterations, and computed the KL divergence to the ground truth every 200,000 iterations with the first 25% discarded as burn-in. We examine the performance of VBPI and MCMC on two data sets: DS1 and the influenza data set with 100 sequences introduced in Section 5.3.

We also investigate the impact that ξ has in the KL computation. This parameter is the assigned probability for a tree in the ground truth/long MCMC run but not in the approximate distribution (short MCMC or subsplit support). We test two alternative choices for ξ in the KL computation: the prior density and $1e-10$. The prior density corresponds to treating all unsampled tree topologies equally and penalizes MCMC more for not detecting trees in the ground truth samples, especially when the number of sequences is large, while $1e-10$ is closer to the smallest ground truth posterior estimates for both data sets and hence would bias MCMC towards the ground truth. Besides serial implementations, we also compared parallel implementations of both method where VBPI can be more advantageous due to the use of multiple-sample training objectives. We used 4 processors for MPI MrBayes runs as there are 4 chains. For VBPI, we lower bound the parallel runtimes using `bito` with 10 threads as $K = 10$. The results were obtained from 10 independent runs for both methods. All experiments were done on a 2019 Mac Pro (3.2 GHz 16-Core Intel Xeon W).

Figure 13 shows the KL divergence to the ground truth on DS1 as a function of approximate runtime. For each row, the left, middle and right plots show the results when ξ is set to the prior density, the machine precision and $1e-10$ respectively. The results for the serial implementations are shown in the first row. We see that the speed advantage of VBPI is reduced due to the caching trick of MCMC moves. As ξ increases, the reported KL divergences from standard MCMC runs get better and VBPI is slower than MCMC when $\xi = 1e-10$. When parallelization is enabled, as shown in the second row, the advantage of VBPI is recovered. Compared to MCMC, the performance of VBPI is not much affected by the choice of ξ which is due to the large subsplit support that provides a better coverage of the ground truth samples (however, this also makes the optimization harder). As the exact posterior here is relatively concentrated (around 3,000 unique trees in the ground truth sample), it turns out that we can reduce the size of subsplit support for faster convergence of VBPI, with only a slight decrease in the approximation accuracy. The last row shows the results when the subsplit support were obtained using 10 replicates of short MCMC runs with 100,000 iterations (sampled every 100 iterations with a 25% burn-in rate). In

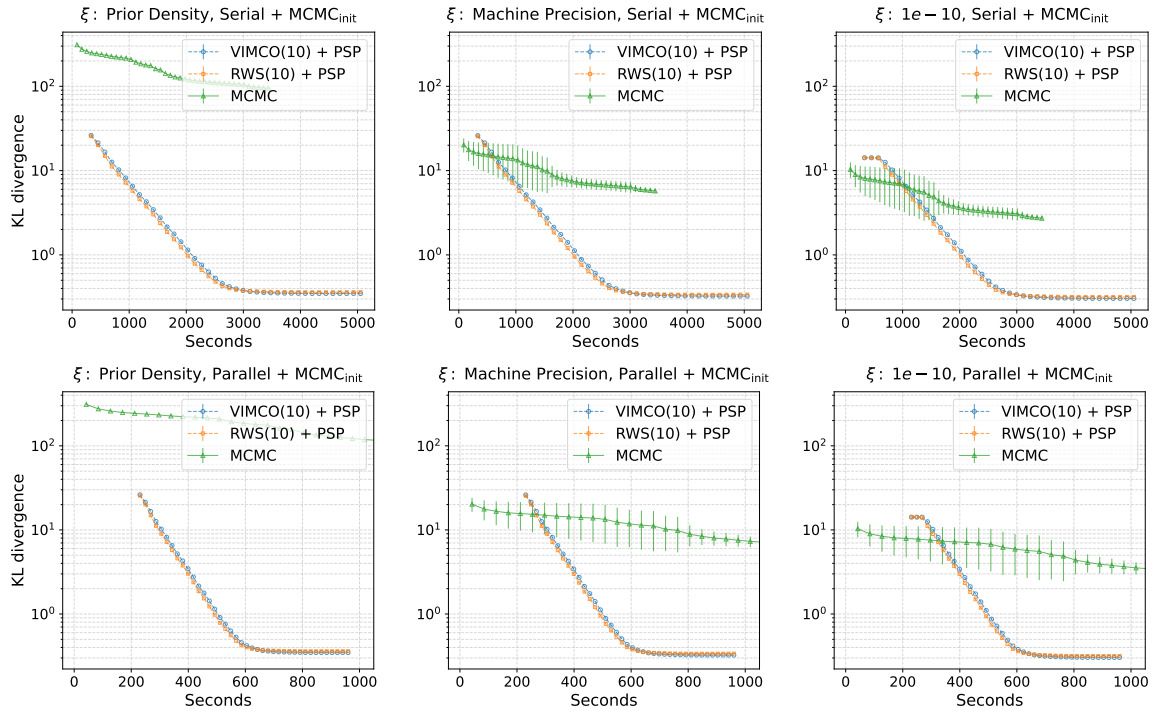


Figure 14: A runtime comparison between VBPI and MCMC on a curated influenza data set with 100 sequences. **(Top)**: Serial implementation with short run MCMC subsplit support. **(Bottom)**: Parallel implementation with short run MCMC subsplit support. Each column corresponds to a different value of ξ . Error bars show one standard deviation over 10 independent runs.

this case, we can use a faster annealing schedule where $\beta_t = \min(1, 0.001 + t/50,000)$ and the convergence time is halved compared to the previous bootstrap method, albeit the KL divergence would be a bit larger when ξ is small.

The advantage of VBPI should be more evident when the posterior is more diffuse. Figure 14 shows the runtime comparison results on a curated influenza data set with 100 sequences. Although subsplit support estimation would require some extra computation, we see that overall VBPI converges much faster than MCMC even using serial implementations. The advantage of VBPI can be further improved when parallelization is used. As the tree space is large and the posterior is quite diffuse (around 100,000 unique trees in the ground truth sample), when the prior density is used for ξ , the KL divergences of standard MCMC would be high since lots of trees in the ground truth sample would not be sampled yet (this is also a limitation of MCMC which usually requires a large number of samples for accurate posterior estimation, especially when the posterior is diffuse). When ξ increases, the KL divergences of standard MCMC runs get better but still fall behind VBPI by a large margin.

Appendix E. Heterochronous Coalescent Models

In practice, we may have different sampling times when the observed organisms are evolving rapidly. In this case, the standard coalescent model can be generalized to account for such *heterochronous* sampling (Rodrigo and Felsenstein, 1999; Felsenstein et al., 1999). Consider the ordered node times $t_{2N-1} = 0 < t_{2N-2} < \dots < t_1$, which include both the sampling times for the tip nodes and the coalescent times for the internal nodes. Let ℓ_k denote the number of lineages co-existing in the time interval (t_k, t_{k-1}) between node k and node $k-1$. Let \mathcal{C} denote the set of indices for the internal nodes. The density for heterochronous coalescent models can be obtained by modifying density (28) as

$$p(\mathbf{t}|N_e(t)) = \prod_{k \in \mathcal{C}} \frac{A_{k+1}}{N_e(t_k)} \cdot \prod_{k=2}^{2N-1} \exp\left(-\int_{t_k}^{t_{k-1}} \frac{A_k}{N_e(t)} dt\right) \quad (40)$$

where $\mathbf{t} = (t_1, \dots, t_{2N-1})$, $A_k = \binom{\ell_k}{2}$. Note that density (28) can be viewed as a special case of density (40) when $t_k = 0, \forall k \geq N$.

Appendix F. Consensus Tree Comparison



Figure 15: A comparison of majority-rule consensus trees obtained from VBPI and ground truth MCMC run on DS1. **(Left):** Ground truth MCMC. **(Right):** VBPI (10,000 sampled trees). The plot is created using the *treespace* (Jombart et al., 2017) R package.

Appendix G. Importance Sampling via Variational Approximations

In this section, we provide a detailed importance sampling procedure for marginal likelihood estimation for phylogenetic inference based on the variational approximations provided by VBPI.

Estimating Marginal Likelihood of Trees For each tree τ that is covered by the subsplit support,

$$Q_\psi(\mathbf{q}|\tau) = \prod_{e \in E(\tau)} p^{\text{Lognormal}}(q_e \mid \mu(e, \tau), \sigma(e, \tau))$$

is our variational approximation to the posterior of branch lengths on τ , where the mean and variance parameters $\mu(e, \tau), \sigma(e, \tau)$ are gathered from the structured variational parameters ψ as introduced in Section 4.3. Therefore, we can estimate the marginal likelihood of τ using importance sampling with $Q_\psi(\mathbf{q}|\tau)$ being the importance distribution as follows

$$p(\mathbf{Y}|\tau) = \mathbb{E}_{Q_\psi(\mathbf{q}|\tau)} \frac{p(\mathbf{Y}|\tau, \mathbf{q})p(\mathbf{q})}{Q_\psi(\mathbf{q}|\tau)} \simeq \frac{1}{M} \sum_{j=1}^M \frac{p(\mathbf{Y}|\tau, \mathbf{q}^j)p(\mathbf{q}^j)}{Q_\psi(\mathbf{q}^j|\tau)} \quad \text{with } \mathbf{q}^j \stackrel{\text{iid}}{\sim} Q_\psi(\mathbf{q}|\tau).$$

Estimating Model Evidence Similarly, we can estimate the marginal likelihood of the data as follows

$$p(\mathbf{Y}) = \mathbb{E}_{Q_{\phi, \psi}(\tau, \mathbf{q})} \frac{p(\mathbf{Y}|\tau, \mathbf{q})p(\tau, \mathbf{q})}{Q_\phi(\tau)Q_\psi(\mathbf{q}|\tau)} \simeq \frac{1}{K} \sum_{j=1}^K \frac{p(\mathbf{Y}|\tau^j, \mathbf{q}^j)p(\tau^j, \mathbf{q}^j)}{Q_\phi(\tau^j)Q_\psi(\mathbf{q}^j|\tau^j)} \quad \text{with } \tau^j, \mathbf{q}^j \stackrel{\text{iid}}{\sim} Q_{\phi, \psi}(\tau, \mathbf{q}).$$

In our experiments, we use $K = 1,000$. When taking a log transformation, the above Monte Carlo estimate is no longer unbiased (for the evidence $\log p(\mathbf{Y})$). Instead, it can be viewed as one sample Monte Carlo estimate of the lower bound

$$L^K(\phi, \psi) = \mathbb{E}_{Q_{\phi, \psi}(\tau^{1:K}, \mathbf{q}^{1:K})} \log \left(\frac{1}{K} \sum_{i=1}^K \frac{p(\mathbf{Y}|\tau^i, \mathbf{q}^i)p(\tau^i, \mathbf{q}^i)}{Q_\phi(\tau^i)Q_\psi(\mathbf{q}^i|\tau^i)} \right) \leq \log p(\mathbf{Y}) \quad (41)$$

whose tightness improves as the number of samples K increases. Therefore, with a sufficiently large K , we can use the lower bound estimate as a proxy for Bayesian model selection.

References

- S. I. Amari. *Differential-Geometrical Methods in Statistics*. Springer, Berlin, Germany, 1985.
- S. I. Amari. α divergence is unique, belonging to both f -divergence and Bregman divergence classes. *IEEE Trans. Inf. Theory*, 55(11):4925–4931, 2009.
- Daniel L Ayres, Michael P Cummings, Guy Baele, Aaron E Darling, Paul O Lewis, David L Swofford, John P Huelsenbeck, Philippe Lemey, Andrew Rambaut, and Marc A Suchard. BEAGLE 3: Improved Performance, Scaling, and Usability for a High-Performance Computing Library for Statistical Phylogenetics. *Systematic Biology*, 68(6):1052–1061, 04 2019. ISSN 1063-5157. doi: 10.1093/sysbio/syz020. URL <https://doi.org/10.1093/sysbio/syz020>.

- G. Baele, P. Lemey, and M. A. Suchard. Genealogical working distributions for Bayesian model testing with phylogenetic uncertainty. *Syst. Biol.*, 65(2):250–264, 2016.
- D. Barber and W. Wiegerinck. Tractable variational structures for approximating graphical models. In *Neural Information Processing Systems*, 1999.
- C. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.
- C. Bishop, N. Lawrence, T. Jaakkola, and M. I. Jordan. Approximating posterior distributions in belief networks using mixtures. In *Neural Information Processing Systems*, 1998.
- D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. In *ICLR*, 2016.
- Sourav Chatterjee and Persi Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099 – 1135, 2018. doi: 10.1214/17-AAP1326. URL <https://doi.org/10.1214/17-AAP1326>.
- T. Dang and H. Kishino. Stochastic variational inference for Bayesian phylogenetics: a case of CAT model. *Molecular biology and evolution*, 36(4):825–833, 2019.
- Rob DeSalle and George Amato. The expansion of conservation genetics. *Nat. Rev. Genet.*, 5(9):702–712, September 2004. ISSN 1471-0056. doi: 10.1038/nrg1425. URL <http://dx.doi.org/10.1038/nrg1425>.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Conference on Learning Representations*, 2017a.
- Vu Dinh, Arman Bilge, Cheng Zhang, and Frederick A Matsen IV. Probabilistic Path Hamiltonian Monte Carlo. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1009–1018, July 2017b. URL <http://proceedings.mlr.press/v70/dinh17a.html>.
- A. J. Drummond, G. K. Nicholls, A. G. Rodrigo, and W. Solomon. Estimating mutation parameters, population history, and genealogy simultaneously from temporally spaced sequence data. *Genetics*, 161:1307–1320, 2002.
- A. J. Drummond, A. Rambaut, B. Shapiro, and O. G. Pybus. Bayesian coalescent inference of past population dynamics from molecular sequences. *Mol. Biol. Evol.*, 22:1185–1192, 2005.
- Alexei J Drummond and Marc A Suchard. Bayesian random local clocks, or one rate to rule them all. *BMC Biol.*, 8:114, August 2010. ISSN 1741-7007. doi: 10.1186/1741-7007-8-114. URL <http://dx.doi.org/10.1186/1741-7007-8-114>.

- Louis du Plessis, John T McCrone, Alexander E Zarebski, Verity Hill, Christopher Ruis, Bernardo Gutierrez, Jayna Raghvani, Jordan Ashworth, Rachel Colquhoun, Thomas R Connor, Nuno R Faria, Ben Jackson, Nicholas J Loman, Áine O’Toole, Samuel M Nicholls, Kris V Parag, Emily Scher, Tetyana I Vasylyeva, Erik M Volz, Alexander Watts, Isaac I Bogoch, Kamran Khan, COVID-19 Genomics UK (COG-UK) Consortium†, David M Aanensen, Moritz U G Kraemer, Andrew Rambaut, and Oliver G Pybus. Establishment and lineage dynamics of the SARS-CoV-2 epidemic in the UK. *Science*, January 2021. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.abf2946. URL <https://science.sciencemag.org/content/early/2021/01/07/science.abf2946>.
- Gytis Dudas, Luiz Max Carvalho, Trevor Bedford, Andrew J Tatem, Guy Baele, Nuno R Faria, Daniel J Park, Jason T Ladner, Armando Arias, Danny Asogun, Filip Bielejec, Sarah L Caddy, Matthew Cotten, Jonathan D’Ambrozio, Simon Dellicour, Antonino Di Caro, Joseph W Diclaro, Sophie Duraffour, Michael J Elmore, Lawrence S Fakoli, Ousmane Faye, Merle L Gilbert, Sahr M Gevao, Stephen Gire, Adrienne Gladden-Young, Andreas Gnirke, Augustine Goba, Donald S Grant, Bart L Haagmans, Julian A Hiscox, Umaru Jah, Jeffrey R Kugelman, Di Liu, Jia Lu, Christine M Malboeuf, Suzanne Mate, David A Matthews, Christian B Matranga, Luke W Meredith, James Qu, Joshua Quick, Suzan D Pas, My V T Phan, Georgios Pollakis, Chantal B Reusken, Mariano Sanchez-Lockhart, Stephen F Schaffner, John S Schieffelin, Rachel S Sealfon, Etienne Simon-Loriere, Saskia L Smits, Kilian Stoecker, Lucy Thorne, Ekaete Alice Tobin, Mohamed A Vandt, Simon J Watson, Kendra West, Shannon Whitmer, Michael R Wiley, Sarah M Winnicki, Shirlee Wohl, Roman Wölfel, Nathan L Yozwiak, Kristian G Andersen, Sylvia O Blyden, Fatorma Bolay, Miles W Carroll, Bernice Dahn, Boubacar Diallo, Pierre Formenty, Christophe Fraser, George F Gao, Robert F Garry, Ian Goodfellow, Stephan Günther, Christian T Happi, Edward C Holmes, Brima Kargbo, Sakoba Keita, Paul Kellam, Marion P G Koopmans, Jens H Kuhn, Nicholas J Loman, N’faly Magassouba, Dhamari Naidoo, Stuart T Nichol, Tolbert Nyenswah, Gustavo Palacios, Oliver G Pybus, Pardis C Sabeti, Amadou Sall, Ute Ströher, Isatta Wurie, Marc A Suchard, Philippe Lemey, and Andrew Rambaut. Virus genomes reveal factors that spread and sustained the ebola epidemic. *Nature*, April 2017. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature22040. URL <http://dx.doi.org/10.1038/nature22040>.
- Y. Fan, R. Wu, M.-H. Chen, L. Kuo, and P. O. Lewis. Choosing among partition models in Bayesian phylogenetics. *Mol. Biol. Evol.*, 28(1):523–532, 2011.
- J. Felsenstein. Evolutionary trees from dna sequences: A maximum likelihood approach. *Journal of Molecular Evolution*, 17:268–276, 1981.
- J. Felsenstein. Confidence limits on phylogenies: An approach using the bootstrap. *Evolution*, 39:783–791, 1985.
- Joseph Felsenstein, Mary K Kuhner, Jon Yamato, and Peter Beerli. Likelihoods on coalescents: a monte carlo sampling approach to inferring parameters from population samples of molecular data. In *Statistics in molecular biology and genetics*, volume 33, pages 163–186. Institute of Mathematical Statistics, January 1999. doi: 10.1214/lnms/1215455552.

- Roberto Feuda, Martin Dohrmann, Walker Pett, Hervé Philippe, Omar Rota-Stabelli, Nicolas Lartillot, Gert Wörheide, and Davide Pisani. Improved modeling of compositional heterogeneity supports sponges as sister to all other animals. *Curr. Biol.*, 27(24):3864–3870.e4, December 2017. ISSN 0960-9822, 1879-0445. doi: 10.1016/j.cub.2017.11.008. URL <http://dx.doi.org/10.1016/j.cub.2017.11.008>.
- M. Fourment and A. E. Darling. Evaluating probabilistic programming and fast variational Bayesian inference in phylogenetics. *PeerJ.*, 7:e8272, 2019. URL <https://doi.org/10.7717/peerj.8272>.
- M. Fourment and E. C. Holmes. Novel non-parametric models to estimate evolutionary rates and divergence times from heterochronous sequence data. *BMC Evol. Biol.*, 14(1):163, 2014.
- Mathieu Fourment, Christiaan J Swanepoel, Jared G Galloway, Xiang Ji, Karthik Gangavarapu, Marc A Suchard, and Frederick A Matsen IV. Automatic Differentiation is no Panacea for Phylogenetic Gradient Computation. *Genome Biology and Evolution*, 15(6):evad099, 06 2023. ISSN 1759-6653. doi: 10.1093/gbe/evad099. URL <https://doi.org/10.1093/gbe/evad099>.
- A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall, London, 2003.
- M. S. Gill, P. Lemey, N. R. Faria, A. Rambaut, B. Shapiro, and M. A. Suchard. Improving Bayesian population dynamics inference: a coalescent-based model for multiple loci. *Molecular biology and evolution*, 30(3):713–724, 2012.
- P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer Science & Business Media, 2013.
- P. Glasserman and Y. C. Ho. *Gradient estimation via Perturbation Analysis*. Springer Science & Business Media, 1991.
- P. W. Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33(10):75–84, 1990.
- P. W. Glynn and R. Szechtman. Some new perspectives on the method of control variates. In *Monte Carlo and Quasi-Monte Carlo Methods 2000*, pages 27–49. Springer, 2002.
- R. Griffiths and S. Tavaré. Sampling theory for neutral alleles in a varying environment. *Philos. Trans. R. Soc. London. B: Biol. Sci.*, 344:403–410, 1994.
- S. B. Hedges, K. D. Moberg, and L. R. Maxson. Tetrapod phylogeny inferred from 18S and 28S ribosomal RNA sequences and review of the evidence for amniote relationships. *Mol. Biol. Evol.*, 7:607–633, 1990.
- G. E. Hinton, P. Dayan, B. J. Frey, and R. M. Neal. The wake-sleep algorithm for unsupervised neural networks. *Science*, 268:1158–1161, 1995.

- S. Höhna, T. A. Heath, B. Boussau, M. J. Landis, F. Ronquist, and J. P. Huelsenbeck. Probabilistic graphical model representation in phylogenetics. *Syst. Biol.*, 63:753–771, 2014.
- Sebastian Höhna and Alexei J. Drummond. Guided tree topology proposals for Bayesian phylogenetic inference. *Syst. Biol.*, 61(1):1–11, January 2012. ISSN 1063-5157. doi: 10.1093/sysbio/syr074. URL <http://dx.doi.org/10.1093/sysbio/syr074>.
- M. T. Holder, P. O. Lewis, D. L. Swofford, and D. Bryant. Variable tree topology stepping-stone marginal likelihood estimation. In M.-H. Chen, L. Kuo, and P. O. Lewis, editors, *Bayesian phylogenetics: methods, algorithms and applications*, pages 95–111. Chapman & Hall/CRC, 2014.
- J. P. Huelsenbeck and F. Ronquist. MrBayes: Bayesian inference of phylogeny. *Bioinformatics*, 17:754–755, 2001.
- J. P. Huelsenbeck, F. Ronquist, R. Nielsen, and J. P. Bollback. Bayesian inference of phylogeny and its impact on evolutionary biology. *Science*, 294:2310–2314, 2001.
- Xiang Ji, Alexander A Fisher, Shuo Su, Jeffrey L Thorne, Barney Potter, Philippe Lemey, Guy Baele, and Marc A Suchard. Scalable bayesian divergence time estimation with ratio transformations. *Syst. Biol.*, 72(5):1136–1153, November 2023. ISSN 1063-5157. doi: 10.1093/sysbio/syad039. URL <https://academic.oup.com/sysbio/article-pdf/72/5/1136/53028357/syad039.pdf>.
- Thibaut Jombart, Michelle Kendall, Jacob Almagro-Garcia, and Caroline Colijn. treespace: Statistical exploration of landscapes of phylogenetic trees. *Molecular Ecology Resources*, 17:1385–1392, 2017. URL <https://doi.org/10.1111/1755-0998.12676>.
- M. I. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul. Introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999.
- T. H. Jukes and C. R. Cantor. Evolution of protein molecules. In H. N. Munro, editor, *Mammalian protein metabolism, III*, pages 21–132, New York, 1969. Academic Press.
- Seong-Hwan Jun, Hassan Nasif, Chris Jennings-Shaffer, David H Rich, Anna Kooperberg, Mathieu Fourment, Cheng Zhang, Marc A Suchard, and Frederick A Matsen, 4th. A topology-marginal composite likelihood via a generalized phylogenetic pruning algorithm. *Algorithms Mol. Biol.*, 18(1):10, July 2023. ISSN 1748-7188. doi: 10.1186/s13015-023-00235-1. URL <http://dx.doi.org/10.1186/s13015-023-00235-1>.
- Caleb Ki and Jonathan Terhorst. Variational phylodynamic inference using pandemic-scale data. *Mol. Biol. Evol.*, July 2022. ISSN 0737-4038, 1537-1719. doi: 10.1093/molbev/msac154. URL <http://dx.doi.org/10.1093/molbev/msac154>.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *ICLR*, 2014.

- Diederik P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pages 4743–4751, 2016.
- J. Kingman. The coalescent. *Stochastic processes and their applications*, 13(3):235–248, 1982.
- J. P. Kleijnen and R. Y. Rubinstein. Optimization and sensitivity analysis of computer simulation models by the score function method. *European Journal of Operational Research*, 88(3):413–427, 1996.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.
- S. Kullback and R. A. Leibler. On information and sufficiency. *The Ann. Math. Statist.*, 22(1):79–86, 1951.
- C. Lakner, P. van der Mark, J. P. Huelsenbeck, B. Larget, and F. Ronquist. Efficiency of Markov chain Monte Carlo tree proposals in Bayesian phylogenetics. *Syst. Biol.*, 57:86–103, 2008.
- R. S. Lanciotti, D. J. Gubler, and D. W. Trent. Molecular evolution and phylogeny of dengue-4 virus. *Journal of General Virology*, 78:2279–2286, 1997.
- Bret Larget. The estimation of tree posterior probabilities using conditional clade probability distributions. *Syst. Biol.*, 62(4):501–511, July 2013. ISSN 1063-5157. doi: 10.1093/sysbio/syt014. URL <http://dx.doi.org/10.1093/sysbio/syt014>.
- Nicolas Lartillot and Hervé Philippe. A Bayesian mixture model for across-site heterogeneities in the amino-acid replacement process. *Mol. Biol. Evol.*, 21(6):1095–1109, June 2004. ISSN 0737-4038. doi: 10.1093/molbev/msh112. URL <http://dx.doi.org/10.1093/molbev/msh112>.
- Frederick A Matsen, Aaron Gallagher, and Connor O McCoy. Minimizing the average distance to a closest leaf in a phylogenetic tree. *Syst. Biol.*, 62(6):824–836, November 2013. ISSN 1063-5157. doi: 10.1093/sysbio/syt044. URL <http://sysbio.oxfordjournals.org/content/62/6/824.abstract>.
- B. Mau, M. Newton, and B. Larget. Bayesian phylogenetic inference via Markov chain Monte Carlo methods. *Biometrics*, 55:1–12, 1999.
- Takahiro Mimori and Michiaki Hamada. Geophy: Differentiable phylogenetic inference via geometric gradients of tree topologies. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=54z8M7NTbJ>.
- B. Q. Minh, M. A. T. Nguyen, and A. von Haeseler. Ultrafast approximation for phylogenetic bootstrap. *Mol. Biol. Evol.*, 30:1188–1195, 2013.
- V. N. Minin, E. W. Bloomquist, and M. A. Suchard. Smooth skyride through a rough skyline: Bayesian coalescent-based inference of population dynamics. *Molecular biology and evolution*, 25(7):1459–1471, 2008.

- T. Minka. Divergence measures and message passing. Technical Report Technial Report TR-2005-173, Microsoft Research, 2005.
- A. Mnih and K. Gregor. Neural variational inference and learning in belief networks. In *Proceedings of The 31th International Conference on Machine Learning*, pages 1791–1799, 2014.
- Andriy Mnih and Danilo Rezende. Variational inference for monte carlo objectives. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1791–1799, 2016.
- Radford Neal. MCMC using hamiltonian dynamics. In S Brooks, A Gelman, G Jones, and XL Meng, editors, *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC Handbooks of Modern Statistical Methods. Taylor & Francis, 2011. ISBN 9781420079425. URL <http://books.google.com/books?id=qfRsAIKZ4rIC>.
- D. J. Nott, S. L. Tan, M. Villani, and R. Kohn. Regression density estimation with variational methods and stochastic approximation. *Journal of Computational and Graphical Statistics*, 21(3):797–820, 2012.
- J. W. Paisley, D. M. Blei, and M. I. Jordan. Variational Bayesian inference with stochastic search. In *Proceedings of the 29th International Conference on Machine Learning ICML*, 2012.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- G. C. Pflug. *Optimization of Stochastic Models: The Interface between Simulation and Optimization*. Springer Science & Business Media, 1996.
- O. G. Pybus, A. Rambaut, and P. H. Harvey. An integrated framework for the inference of viral population history from reconstructed genealogies. *Genetics*, 155(3):1429–1437, 2000.
- O. G. Pybus, M. A. Charleston, S. Gupta, A. Rambaut, E. C. Holmes, and P. H. Harvey. The epidemic behaviour of the hepatitis c virus. *Science*, 292(5525):2323–2325, 2001. doi: 10.1126/science/1058321.

- Tom Rainforth, Adam R. Kosioreck, Tuan Anh Le, Chris J. Maddison, Maximilian Igl, Frank Wood, and Yee Whye Teh. Tighter variational bounds are not necessarily better. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- R. Ranganath, S. Gerrish, and D. M. Blei. Black box variational inference. In *AISTATS*, pages 814–822, 2014.
- D. Rezende and S. Mohamed. Variational inference with normalizing flow. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 1530–1538, 2015.
- D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, 2014.
- H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22:400–407, 1951.
- A. G. Rodrigo and J. Felsenstein. Coalescent approaches to hiv population genetics. In K. Crandall, editor, *The Evolution of HIV*, pages 233–272. Johns Hopkins Univ. Press, Baltimore, 1999.
- F. Ronquist, M. Teslenko, P. van der Mark, D. L. Ayres, A. Darling, S. Höhna, B. Larget, L. Liu, M. A. Suchard, and J. P. Huelsenbeck. MrBayes 3.2: efficient Bayesian phylogenetic inference and model choice across a large model space. *Syst. Biol.*, 61:539–542, 2012.
- R. Y. Rubinstein. Sensitivity analysis of discrete event systems by the push out method. *Annals of Operations Research*, 39(1):229–250, 1992.
- R. Y. Rubinstein and J. Kreimer. About one monte carlo method for solving linear equations. *Mathematics and Computers in Simulation*, 25(4):321–334, 1983.
- P. Sagulenko, V. Puller, and R. A. Neher. Treetime: maximum-likelihood phylodynamic analysis. *Virus Evol.*, 4, 2018.
- I. Sason and S. Verdú. f -divergence inequalities. *IEEE Trans. Inf. Theory*, 62(11):5973–6006, 2016.
- L. Saul and M. I. Jordan. Exploiting tractable substructures in intractable networks. In *Neural Information Processing Systems*, 1996.
- Eric E. Schadt, Janet S. Sinsheimer, and Kenneth Lange. Computational advances in maximum likelihood methods for molecular phylogeny. *Genome Res.*, 8:222–233, 1998. doi: 10.1101/gr.8.3.222.
- K. Strimmer and V. Moulton. Likelihood analysis of phylogenetic networks using directed graphical models. *Molecular biology and evolution*, 17:875–881, 2000.
- M. A. Suchard, P. Lemey, G. Baele, D. L. Ayres, A. J. Drummond, and A. Rambaut. Bayesian phylogenetic and phylodynamic data integration using beast 1.10. *Virus Evolution*, 4(1), 2018. doi: 10.1093/ve/vey016.

- David L. Swofford. Paup*: Phylogenetic analysis using parsimony (and other methods) version 4.0 beta, 2001.
- M. Titsias and M. Lázaro-Gredilla. Doubly stochastic variational bayes for non-conjugate inference. In *International Conference on Machine Learning*, 2014.
- M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- C. Whidden, B. C. Claywell, T. Fisher, A. F. Magee, M. Fourment, and F. A. Matsen IV. Systematic exploration of the high likelihood set of phylogenetic tree topologies. *Syst. Biol.*, 69(2):280–293, 2019.
- Chris Whidden and Frederick A Matsen IV. Quantifying MCMC exploration of phylogenetic tree space. *Syst. Biol.*, 64(3):472–491, May 2015. ISSN 1063-5157, 1076-836X. doi: 10.1093/sysbio/syv006. URL <http://dx.doi.org/10.1093/sysbio/syv006>.
- R. J. Williams. Simple statistical gradient-following algorithm for connectionist reinforcement learning. *Machine Learning*, 8:5–32, 1992.
- Alan S. Willsky. Multiresolution markov models for signal and image processing. *Proceedings of the IEEE*, 90(8):1396–1458, 2002.
- Tianyu Xie and Cheng Zhang. ARTree: A deep autoregressive model for phylogenetic inference. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=SoLebIqHgZ>.
- W. Xie, P. O. Lewis, Y. Fan, L. Kuo, and M.-H. Chen. Improving marginal likelihood estimation for Bayesian phylogenetic model selection. *Syst. Biol.*, 60:150–160, 2011.
- Z. Yang. Paml 4: phylogenetic analysis by maximum likelihood. *Molecular Biology and Evolution*, 24(8):1586–1591, 2007. doi: 10.1093/molbev/msm088.
- Z. Yang and B. Rannala. Bayesian phylogenetic inference using DNA sequences: a Markov chain Monte Carlo method. *Mol. Biol. Evol.*, 14:717–724, 1997.
- Cheng Zhang. Improved variational Bayesian phylogenetic inference with normalizing flows. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18760–18771. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/d96409bf894217686ba124d7356686c9-Paper.pdf>.
- Cheng Zhang. Learnable topological features for phylogenetic inference via graph neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=hVVUY7p64WL>.
- Cheng Zhang and Frederick A Matsen IV. Generalizing tree probability estimation via Bayesian networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1444–1453. Curran Associates, Inc., 2018. URL <http://papers.nips.cc/paper/7418-generalizing-tree-probability-estimation-via-bayesian-networks.pdf>.

- Cheng Zhang and Frederick A Matsen IV. Variational Bayesian phylogenetic inference. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SJVmjR9FX>.
- Yun Zhang, Brian D Aebermann, Tavis K Anderson, David F Burke, Gwenaelle Dauphin, Zhiping Gu, Sherry He, Sanjeev Kumar, Christopher N Larsen, Alexandra J Lee, Xiaomei Li, Catherine Macken, Colin Mahaffey, Brett E Pickett, Brian Reardon, Thomas Smith, Lucy Stewart, Christian Suloway, Guangyu Sun, Lei Tong, Amy L Vincent, Bryan Walters, Sam Zaremba, Hongtao Zhao, Liwei Zhou, Christian Zmasek, Edward B Klem, and Richard H Scheuermann. Influenza research database: An integrated bioinformatics resource for influenza virus research. *Nucleic Acids Res.*, 45(D1):D466–D474, January 2017. ISSN 0305-1048, 1362-4962. doi: 10.1093/nar/gkw857. URL <http://dx.doi.org/10.1093/nar/gkw857>.
- Ming Yang Zhou, Zichao Yan, Elliot Layne, Nikolay Malkin, Dinghuai Zhang, Moksh Jain, Mathieu Blanchette, and Yoshua Bengio. PhyloGFN: Phylogenetic inference with generative flow networks. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=hB7S1fEmze>.