

# COMPUTING THE POSTERIOR EXPECTATION OF PHYLOGENETIC TREES

PHILIPP BENNER AND MIROSLAV BAČÁK

**ABSTRACT.** Inferring phylogenetic trees from multiple sequence alignments often relies upon Markov chain Monte Carlo (MCMC) methods to generate tree samples from a posterior distribution. To give a rigorous approximation of the posterior expectation, one needs to compute the mean of the tree samples and therefore a sound definition of a mean and algorithms for its computation are necessary. To the best of our knowledge, no existing method of phylogenetic inference can handle the full set of sample trees, because such trees typically have different topologies. We develop a statistical model for the inference of phylogenetic trees based on the tree space due to Billera et al. [2001]. Since it is an Hadamard space, the mean and median are well defined, which we also motivate from a decision theoretic perspective. The actual approximation of the posterior expectation relies on some recent developments in Hadamard spaces (Bačák [2013], Miller et al. [2012]) and the fast computation of geodesics in tree space (Owen and Provan [2011]), which altogether enable to compute medians and means of trees with different topologies. We demonstrate our model on the small ribosomal subunit rRNA sequence alignment. The posterior expectations obtained on this data set are a meaningful summary of the posterior distribution and the uncertainty about the tree topology.

## 1. INTRODUCTION

Phylogenetic inference is concerned with the estimation of trees that are meant to reflect the evolutionary history of a set of species. Moreover, such point estimates are instrumental to a variety of other inferential tasks, such as the analysis of ChIP-Seq data for the prediction of regulatory elements (cf. Wasserman and Sandelin [2004]). A well motivated statistical model with a sound estimation method is therefore of utmost importance for many applications in computational genetics. A variety of such methods already exist (e.g. Huelsenbeck and Ronquist [2001], Guindon and Gascuel [2003], Drummond and Rambaut [2007], Lartillot et al. [2009]), which either search for a maximizer of the posterior or likelihood function, or rely on Markov chain Monte Carlo (MCMC) methods to generate samples from the posterior distribution. In phylogenetic inference, posterior samples are phylogenetic trees and their average is usually not well defined unless all trees have the same topology. If trees are restricted to one topology at some point of the estimation task, the computed average inevitably neglects part of the data and is not good summary of the full posterior distribution. A common approach is to construct

---

*Date:* June 11, 2013.

*Key words and phrases.* Bayesian statistics; BHV tree space; Fréchet mean; geometric median; phylogenetic trees; posterior expectation.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 267087.

a (majority rule) consensus tree from MCMC samples, for which some decision theoretic arguments have been proposed (e.g. Holder et al. [2003], Huggins et al. [2011]). However, arguments have been raised against this method (e.g. Wheeler and Pickett [2008]) and a more rigorous approach is still lacking.

A first step towards solving this issue was made by Billera et al. [2001] who introduced a space of trees, now called the *BHV tree space*, or simply tree space, where a point in this space not only identifies the tree topology, but also the edge lengths. We construct a posterior distribution on this space and show how the expectation and other posterior quantities can be computed. More specifically, since tree space is an Hadamard space, we have well-defined notions of a mean and median of probability distributions, which we motivate from decision theoretic grounds. The actual computations of the posterior mean and median rely on approximation algorithms developed by Bačák [2013], Miller et al. [2012], which in turn require additional tools, mainly the algorithm due to Owen and Provan [2011] allowing to compute geodesics between pairs of trees in polynomial time. It is important to emphasize that the construction of the BHV tree space along with the Owen-Provan algorithm provides us with a new way of measuring distances between (phylogenetic) trees, which seem to surpass the conventional metrics (e.g. the NNI distance or the Robinson-Foulds distance) at both mathematical and computational aspects.

In the present paper, we will give a full description of phylogenetic inference. After a short decision theoretic motivation (Section 2) we will outline the space of trees by Billera et al. [2001] in Section 3, on which our model is defined. To construct a distribution on this space (Section 4), we first fix a tree topology and thereby restrict the discussion to one orthant of tree space, say the  $i$ -th orthant. We construct a posterior distribution  $\mu_i$  on this orthant, which defines the probability of phylogenetic trees of this topology given a multiple sequence alignment. The posterior distribution  $\mu$  on the full tree space is then obtained by combining the single components  $\mu_i$ , i.e.  $\mu = \sum_i w_i \mu_i$ . The main obstacle of this model is the evaluation of the weights  $w_i$  since they depend on the partition function of the individual distributions  $\mu_i$ , which involves computing an intractable integral. We therefore approximate  $\mu$  with a finite combination of Dirac measures  $\pi$  representing  $K$  samples from the posterior distribution. To obtain samples from  $\mu$  a Markov chain Monte Carlo (MCMC) method is used, which we will describe in Section 5.

Even though our target reader is primarily a practitioner in computational genetics whom we provide with a detailed recipe for a rigorous approximation of posterior distributions in phylogenetic inference, we would like to point out that the presented methods stem from a fascinating mix of pure mathematics including non-Euclidean geometry, convex analysis, optimization, probability theory and combinatorics, which has recently attracted a great deal of interest among mathematicians and keeps offering challenging mathematical problems.

**Acknowledgments.** We would like to thank Pierre-Yves Bourguignon, Stephan Poppe, and Johannes Schumacher for very valuable discussions. We are extremely grateful to Ezra Miller for his comments on the manuscript, which significantly improved the exposition.

## 2. DECISION THEORETIC MOTIVATION

Before discussing the details of our model and approximations of posterior expectations, we would like to give a brief motivation. The goal of any inferential task

is to obtain predictions based on a well motivated statistical model and a set of observations. In genetics, such predictions often rely on a phylogenetic tree, which has to be estimated first. Given that we already have a posterior distribution on phylogenetic trees, we need to decide on how to obtain a point estimate. Such a tree should be a good summary of the observed data. For this, it is necessary to define a loss function  $\mathcal{L}(\hat{T}, T')$  that quantifies the error of selecting a tree  $\hat{T}$  if  $T'$  would be a better choice. To illustrate this, assume for the moment that  $\Theta$  is a real valued random variable with posterior distribution  $\mu_{\Theta | \mathbf{X}}$  conditional on some observations  $\{\mathbf{X} = \mathbf{x}\}$ . On the real line a common choice is the squared-error loss  $\mathcal{L}(\hat{\theta}, \theta) = |\hat{\theta} - \theta|^2$ . As best estimate we would take the minimum expected loss

$$\arg \min_{\hat{\theta} \in \mathbb{R}} \int_{\mathbb{R}} \mathcal{L}(\hat{\theta}, \theta) d\mu_{\Theta | \mathbf{X}}(\theta | \mathbf{x}) ,$$

and by differentiating with respect to  $\hat{\theta}$  we immediately find that the best estimate is the first moment of  $\mu_{\Theta | \mathbf{X}}$ , i.e.

$$\hat{\theta} = \mathbb{E}(\Theta | \mathbf{X} = \mathbf{x}) = \int_{\mathbb{R}} \theta d\mu_{\Theta | \mathbf{X}}(\theta | \mathbf{x}) .$$

Similarly we can choose  $\mathcal{L}(\hat{\theta}, \theta) = |\hat{\theta} - \theta|$  for which we obtain the median, whereas a zero-one loss results in a maximum a posteriori (MAP) estimate (cf. Schervish [1995], Robert [2001]).

Except for the zero-one loss, such functions are not well defined since trees might be of different topology, but we may take a much more direct approach. As we will outline later, the posterior distribution  $\mu$  of our model is defined on the tree space  $\mathcal{T}_n$ . In this space, all trees have  $n + 1$  leaves. By definition,  $\mathcal{T}_n$  is a geodesic metric space, which means that we have a metric  $d(T, T')$  that defines the distance between  $T$  and  $T'$  and we also have a geodesic path from  $T$  to  $T'$ , whose length is equal to  $d(T, T')$ . Actually computing the distance involves finding a geodesic path that connects the two trees, which we will discuss later. A possible choice for the loss function is for instance  $\mathcal{L}(\hat{T}, T') := d(\hat{T}, T')^2$ . We then obtain the estimate

$$\hat{T} = \arg \min_{y \in \mathcal{T}_n} \int_{\mathcal{T}_n} d(x, y)^2 d\mu(x) ,$$

which is also called the barycenter  $b(\mu)$  of the distribution  $\mu$ , or the Fréchet mean. In Euclidean spaces it coincides with the posterior expectation, which is why we define

$$\mathbb{E}T := b(\mu) ,$$

where  $T$  is a random variable on tree space with distribution  $\mu$ . For more details on probability theory in Hadamard spaces, see Sturm [2002, 2003]. We will also use

$$\text{Var}T = \min_{y \in \mathcal{T}_n} \int_{\mathcal{T}_n} d(x, y)^2 d\mu(x)$$

as a notion of variance. Similarly, we can choose  $\mathcal{L}(\hat{T}, T) := d(\hat{T}, T)$  and thereby obtain

$$\hat{T} = \arg \min_{y \in \mathcal{T}_n} \int_{\mathcal{T}_n} d(x, y) d\mu(x)$$

as estimate, which is the geometric median. The advantage of the geometric median is that it is less sensitive to long tails of the distribution, but it may not have a unique minimizer. Since the distance function on an Hadamard space is convex,

computing a point estimate of  $\mu$  reduces to finding a minimizer of a convex function. In tree space, we cannot simply differentiate the loss function and follow the gradient to find a minimizer. Appropriate algorithms for computing the mean and median are referred to in Section 5.

Of course, a valid question is whether a single  $\hat{T}$  is a good summary of the full posterior distribution  $\mu$ . In real applications part of the data will favour one tree topology, while another part clearly supports a different topology. Such seemingly contradictory data sets are very frequent in biological applications and lead to posterior distributions whose mass sits on many topologies. In this case a weighted mixture of several trees, i.e. a model average, might be a better summary of the posterior distribution. When computational time is a limiting factor, it might be too costly to use a model average. However, we will demonstrate in Section 6 that also single point estimates in tree space may allow an intuitive interpretation of multimodal posteriors. We would like to mention an alternative approach due to Nye [2011] which instead of a point estimate uses *principal component analysis* in tree space.

### 3. PHYLOGENETIC TREES AND TREE SPACE

We will now describe the construction of tree space due to L. Billera, S. Holmes, and K. Vogtmann. For the details, the interested reader is referred to the original paper Billera et al. [2001]. We first need to make precise what we mean by a (phylogenetic) tree. Given  $n \in \mathbb{N}$  with  $n \geq 3$ , a *metric  $n$ -tree* is a combinatorial tree (connected graph with no circuit) with  $n + 1$  terminal vertices called *leaves* that are labeled  $0, 1, \dots, n$ . In phylogenetics, the labels represent the species in question. The vertex connected with leaf 0 is called the *root*, since it represents a common ancestor of all species in the tree, but it will have no distinguished role in the construction of tree space.<sup>1</sup> (As a matter of fact, such trees can be considered as unrooted.) Vertices other than leaves have no labels since we consider them just as “branching points”. The edges which are adjacent to leaves are called *leaf edges*, and the remaining edges are called *inner*. We see an example of a 6-tree with three inner edges  $e_1, e_2$ , and  $e_3$  in Figure 1. All edges, both leaf and inner, have positive




FIGURE 1. An example of a 6-tree with three inner edges.

lengths. We will refer to a metric  $n$ -tree simply as a tree. The number  $n$  will be fixed and clear from the context. Later, when we consider a *set* of trees instead

<sup>1</sup>Some authors however use the term root for the leaf vertex 0 itself.

of an individual tree, it will be important that they all have the same number of leaves. For the inference of phylogenetic trees, the number of leaves is determined by the number of nucleotide sequences in the data set.

Each inner edge of a tree determines a unique partition of the set of leaves  $L$  into two disjoint and nonempty subsets  $L_1 \cup L_2 = L$  called a *split*, which we denote  $L_1|L_2$ . A split is defined as the partition of leaves that arises if we removed the inner edge under consideration. For instance, the inner edges  $e_1, e_2$ , and  $e_3$  of the tree in Figure 1 have splits  $(0, 4, 5, 6|1, 2, 3)$ ,  $(0, 1, 2, 3|4, 5, 6)$ , and  $(0, 1, 2, 3, 4|5, 6)$ , respectively. On the other hand, given a set of leaves and splits subject to certain conditions, we can uniquely construct a tree. Namely, we require any two splits  $L_1|L_2$  and  $L'_1|L'_2$  be *compatible*, that is, one of the sets

$$L_1 \cap L'_2, \quad L'_1 \cap L_2, \quad L_1 \cap L'_1, \quad L_2 \cap L'_2$$

be empty. We say that a set of inner edges  $I$  is *compatible* if for any two edges  $e, e' \in I$ , the corresponding splits are compatible. For further details, see Dress et al. [2012], Semple and Steel [2003].

We will now proceed to construct a full space of trees, denoted  $\mathcal{T}_n$ , that is, a space whose elements will be all metric  $n$ -trees. First, it is useful to realize that one can treat leaf edges and inner edges separately. Since the former can be represented in Euclidean space of dimension  $n + 1$ , the whole space  $\mathcal{T}_n$  is a product of a Euclidean space and a space that represents the inner edges. We may hence for simplicity ignore the leaf edges in the following construction.

Fix now a metric  $n$ -tree  $T$  with  $r$  inner edges of lengths  $l_1, \dots, l_r$ , where  $1 \leq r \leq n - 2$ . Clearly  $(l_1, \dots, l_r)$  lies in the open orthant  $(0, \infty)^r$ , and conversely, any point of  $(0, \infty)^r$  can be mapped to an  $n$ -tree of the same combinatorial structure as  $T$ . Note that a tree  $S$  is said to have the same combinatorial structure as  $T$  if it has the same number of inner edges as  $T$  and all its inner edges have the same splits as the inner edges of  $T$ . In other words, the trees  $S$  and  $T$  differ only by inner edge lengths.

To any point from the boundary  $\partial(0, \infty)^r$  we associate a metric  $n$ -tree obtained from  $T$  by shrinking some inner edges to zero length. Hence, each point from the closed orthant  $[0, \infty)^r$  corresponds to a metric  $n$ -tree of the same combinatorial structure as  $T$ .

Binary  $n$ -trees have the maximal possible number of inner edges, namely  $n - 2$ , which is of course equal to the dimension of the corresponding orthant. An orthant of an  $n$ -tree that is not binary appears as a face of the orthants corresponding to (at least three) binary trees. In Figure 2, we see a copy of  $[0, \infty)^2$  representing all 4-trees of a given combinatorial structure, namely, all 4-trees with two inner edges  $e_1$  and  $e_2$ , such that the split of  $e_1$  is  $(1, 2|0, 3, 4)$ , and the split of  $e_2$  is  $(1, 2, 3|0, 4)$ . If the length of  $e_1$  is zero, then the tree lies on the vertical boundary ray. If the length of  $e_2$  is zero, then the tree lies on the horizontal boundary ray. In summary, any orthant  $\mathcal{O} = [0, \infty)^r$ , where  $1 \leq r \leq n - 2$ , corresponds to a compatible set of inner edges, and conversely, any compatible set of inner edges  $I = (e_1, \dots, e_r)$  corresponds to a unique orthant  $\mathcal{O}(I)$ , which is a copy of  $[0, \infty)^r$ .

The *tree space*  $\mathcal{T}_n$  consists of  $(2n - 3)!! := (2n - 3)(2n - 5) \cdots 5 \cdot 3$  copies of the orthant  $[0, \infty)^{n-2}$  glued together along lower-dimensional faces, which correspond to non-binary trees, that is, compatible sets of inner edges of cardinality  $< n - 2$ .

We equip the tree space  $\mathcal{T}_n$  with the induced length metric. Then it becomes a geodesic metric space, that is, given a pair of trees, we have a well-defined distance




FIGURE 2. 4-trees of a given combinatorial structure.

between them and moreover they are connected by a geodesic path. One can easily observe that each geodesic consists of finitely many Euclidean line segments. An algorithm for the computation of distances and geodesics is due to Owen and Provan [2011]. The following important theorem from Billera et al. [2001] states that the tree space has nonpositive curvature.

**Theorem 3.1.** *The space  $\mathcal{T}_n$  is an Hadamard space.*

An Hadamard space is a geodesic metric space, which is complete and has non-positive curvature. Intuitively, in such spaces triangles appear “slimmer” than in Euclidean space, see Figure 3.




FIGURE 3. (a) Triangle in a space of nonpositive curvature. (b) Comparison triangle in Euclidean space.

It is impossible to isometrically embed the tree space into the Euclidean space and therefore also difficult to visualize. A piece of the tree space  $\mathcal{T}_4$  is shown in Figure 4. The geometrically oriented reader may notice that triangles in this space are deformed and squeezed inwards, that is, they are “slim” as explained above.

Since this space is not a linear space, one cannot multiply its elements by a real number and addition of two elements of tree space is also not defined. However, convex combinations of a given pair of points are meaningful. Indeed, let  $T, T' \in \mathcal{T}_n$  and  $\lambda \in [0, 1]$ , then we define a formal convex combination

$$T_\lambda := (1 - \lambda)T + \lambda T',$$




FIGURE 4. Five out of 15 orthants of  $\mathcal{T}_4$ .

which represents a unique tree  $T_\lambda \in \mathcal{T}_n$  lying on the geodesic from  $T$  to  $T'$  satisfying  $d(T, T_\lambda) = \lambda d(T, T')$ . Convex combinations are important in the algorithms for computations of medians and means.

#### 4. STATISTICAL MODEL

The posterior distribution is a conditional probability measure that depends on a multiple sequence alignment. Such an alignment is represented as a matrix, where each row is a sequence of nucleotides from one species. Within each column, observed nucleotides are assumed to have evolved according to a phylogenetic tree. The same tree is assumed for the whole alignment. A most intuitive way to describe this process is to look at it from a generative model perspective. We assume that there existed a common ancestor of all species that we are considering and the nucleotides that we observe are generated from the sequence of the common ancestor. Whenever a mutation occurs between an ancestor and its descendant, a new nucleotide is generated from the stationary distribution of the process, at which point it might happen that the same nucleotide is generated again. The stationary distribution therefore plays a crucial role. It is specific to each column of the alignment and reflects the external selective pressure that acts on each site. This is where our model differs from other commonly used methods that assume the same stationary distribution among all sites. In some improved methods it is possible to group sites into distinct classes that share a stationary distribution (e.g. Lartillot and Philippe [2004]). The nucleotide substitution model we use belongs to the family of general time reversible models with gamma prior (GTR+ $\Gamma$ , cf. Nielsen [2005]). To explicate the differences to other methods and how prior parameters should be interpreted we fully outline our model in the following.

For a more formal description it is sufficient for the first part to develop the statistical model on the set of observations within a single column of the alignment. Let  $n + 1$  be the number of species for which we have sequences in the alignment. We introduce the random variables  $\mathbf{X} = \{X_0, X_1, \dots, X_n\}$ , where  $X_i$  takes values in an alphabet  $\mathbb{A}$  and represents the nucleotide of the  $i$ -th sequence. The alphabet

contains a character for each nucleotide and one to represent gaps in the alignment<sup>2</sup>. One of the sequences in the alignment is used for the outgroup, for which we use the random variable  $X_0$  associated with the leaf which is attached to the root of the tree. A phylogenetic tree equipped with an evolutionary model is used to relate sequences of different species. We first consider a particular phylogenetic  $n$ -tree  $T$  with  $r = n - 2$  inner edges. The leaves of the tree  $\{v_0, v_1, \dots, v_n\}$  are associated with the  $n + 1$  random variables  $\mathbf{X}$ . Inner vertices are labeled from  $n + 1$  to  $n + r - 1$  and we associate with each inner vertex  $v_k$  a random variable  $X_k$ , where  $k \in \{n + 1, \dots, n + r - 1\}$ . To discuss the evolutionary model, assume that  $v_i$  and  $v_j$  are leaves or inner vertices that are connected to the  $k$ -th inner vertex  $v_k$ . We need to define the probability of an event  $\{X_i = x_i, X_j = x_j\}$  knowing that  $\{X_k = x_k\}$ . First note that

$$X_i \perp X_j \mid X_k$$

and therefore the conditional probability of  $\{X_i = x_i, X_j = x_j\}$  given  $\{X_k = x_k\}$  factorizes. Hence, it is sufficient to specify the probability of  $\{X_i = x_i\}$  given  $\{X_k = x_k\}$ . We use the model by Felsenstein [1981], which defines a continuous-time finite Markov chain. It is given by

$$p_{X_i \mid X_k}(x_i \mid x_k) := \begin{cases} p_{M_i} p^*(x_i) + p_{\bar{M}_i} & \text{if } x_i = x_k, \\ p_{M_i} p^*(x_i) & \text{if } x_i \neq x_k, \end{cases}$$

where  $p_{M_i}$  is the probability of a mutation and  $p_{\bar{M}_i} = 1 - p_{M_i}$ . The distribution  $p^*$  is the stationary probability distribution of the process, which is common to the full tree. In this model, the case where  $x_i \neq x_k$  is simple, we have a mutation and generate the nucleotide  $x_i$  which we do with probability  $p^*(x_i)$ . More interestingly, if  $x_i$  and  $x_k$  are the same nucleotides, there is either no mutation and no nucleotide has to be generated or there is a mutation and the same nucleotide is generated again. As we will outline later, the entropy of  $p^*$  defines the level of conservation of a site. The probability of a mutation  $p_{M_i}$  depends on the evolutionary time  $t$  between species  $i$  and its ancestor  $k$ , but also on the evolutionary rate  $\alpha$ . In Felsenstein's evolutionary model, we set

$$p_{M_i} := 1 - e^{-l},$$

where  $l := \alpha t$  is defined as the length of the edge between  $v_i$  and its ancestor  $v_k$ . The model is time-reversible, which means that inference is restricted to unrooted trees. This property allows us to define our statistical model on the BHV tree space. As discussed in Section 3, our phylogenetic trees have three vertices attached to the root, which essentially makes the tree unrooted. The leaf associated with  $X_0$  can be seen to represent an outgroup. The position of the root is purely instrumental and has no importance for the computation of the likelihood (cf. Isaev [2006]). Since no observations are available for the inner vertices, it is necessary to marginalize over all corresponding random variables, such that for instance

$$p_{X_i, X_j}(x_i, x_j) = \sum_{x_k \in \mathbb{A}} p^*(x_k) p_{X_i, X_j \mid X_k}(x_i, x_j \mid x_k).$$

In this fashion we obtain the full likelihood of  $\{\mathbf{X} = \mathbf{x}\}$ .

<sup>2</sup>By including a symbol for gaps in the alphabet we explicitly state that no nucleotide is present at positions filled with a gap. If however gaps are modeled as missing data, the meaning of gaps is different, i.e. a gap indicates that any of the nucleotides is present but which one is unknown.



In this model we have two sets of unobserved parameters, namely the lengths of edges and the stationary probability distribution. We proceed by discussing the stationary distribution first, which is specific to each column of the alignment. The distribution will be integrated out in the full model, since we are only interested in the inference of phylogenetic trees. We introduce a random variable  $\Theta$  that represents the stationary distribution and obtain the conditional probability

$$p_{X|\Theta}(x|\boldsymbol{\vartheta}) = \vartheta_x = p^*(x) ,$$

of generating nucleotide  $x \in \mathbb{A}$ . We assume that  $\Theta$  is a priori Dirichlet distributed with pseudocounts  $\boldsymbol{\alpha} = (\alpha_x)_{x \in \mathbb{A}}$ . The probability of observing  $\{\mathbf{X} = \mathbf{x}\}$  becomes

$$p_{\mathbf{X}}(\mathbf{x}) = \int_{\Delta} p_{\mathbf{X}|\Theta}(\mathbf{x}|\boldsymbol{\vartheta}) f_{\Theta}(\boldsymbol{\vartheta}) d\boldsymbol{\vartheta} ,$$

where  $f_{\Theta}$  is the density function of the Dirichlet distribution. The integral is defined on the  $(|\mathbb{A}| - 1)$ -dimensional probability simplex  $\Delta$  and can be solved analytically by first expanding the polynomial of the distribution  $p_{\mathbf{X}|\Theta}$ .

It is important to select an appropriate set of parameters  $\boldsymbol{\alpha}$  for the Dirichlet distribution, as they control the expected entropy of distributions  $\boldsymbol{\vartheta}$  drawn from it. Phylogenetic trees are commonly learned on multiple sequence alignments of genes. Such genomic regions are highly conserved, which means that selective pressure causes nucleotides in a column of the alignment to be the same with high probability. To reflect this knowledge in our prior assumption, it is important that the expected entropy is low, i.e. that only the probability of one or two nucleotides is high. This can be achieved by choosing  $\alpha_x < 1$ , which puts mass on the boundaries of the probability simplex. The choice of  $\boldsymbol{\alpha}$  has a strong influence on inferred edge lengths. If we increase  $\boldsymbol{\alpha}$ , we observe that inferred branch lengths shorten to compensate for the increase in entropy of the stationary distribution. The choice of pseudocounts  $\boldsymbol{\alpha}$  therefore reflects our *a priori* assumption of how conserved we expect a genomic region to be. It is well known that within codons a heterogeneous selective pressure exists (e.g. Li et al. [1985], Yang [1996]), which can be modeled by introducing pseudocounts specific to codon positions.

The next step is to formulate a prior distribution on the edge lengths where we keep the topology fixed. By this we obtain the posterior  $\mu_i$  for a single orthant  $\mathcal{O}_i$ . The same phylogenetic tree is assumed for all columns in the alignment. In fact, columns in the alignment are conditionally independent given a fixed phylogenetic tree. Since we now want to let the tree vary within one orthant of tree space, it is necessary to consider the full alignment. Let  $\mathbf{X}^{(\nu)}$ , where  $\nu = 1, 2, \dots, N$ , denote the random variables for the  $\nu$ -th column of the alignment. We also use the shorthand notation  $\bar{\mathbf{X}} = (\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(N)})$  for the full alignment. Let  $\mathbf{L} = (L_k)$  denote the random variables for the edge lengths of a tree  $T$  in orthant  $\mathcal{O}_i$ . Each  $L_k$  is a priori gamma distributed with shape parameter  $r$  and scale parameter  $\lambda$ . The likelihood of the full alignment is given by

$$p_{\bar{\mathbf{X}}|\mathbf{L}, \mathcal{O}_i}(\bar{\mathbf{x}}|\mathbf{l}) = \prod_{\nu=1}^N p_{\mathbf{X}^{(\nu)}|\mathbf{L}, \mathcal{O}_i}(\mathbf{x}^{(\nu)}|\mathbf{l}) ,$$

where the stationary distribution is integrated out, and we obtain the posterior distribution  $\mu_i$  restricted to orthant  $\mathcal{O}_i$  with density function

$$f_{\mathbf{L}|\bar{\mathbf{X}}, \mathcal{O}_i}(\mathbf{l}|\bar{\mathbf{x}}) = \frac{1}{p_{\bar{\mathbf{X}}|\mathcal{O}_i}(\bar{\mathbf{x}})} p_{\bar{\mathbf{X}}|\mathbf{L}, \mathcal{O}_i}(\bar{\mathbf{x}}|\mathbf{l}) f_{\mathbf{L}}(\mathbf{l}) .$$

The full posterior distribution of  $n$ -trees is given by

$$\mu := \sum_{i=1}^{(2n-3)!!} w_i \mu_i ,$$

where

$$w_i := \frac{p_{\bar{\mathbf{X}} | \mathcal{O}_i}(\bar{\mathbf{x}})}{\sum_j p_{\bar{\mathbf{X}} | \mathcal{O}_j}(\bar{\mathbf{x}})}$$

is the weight of the  $i$ -th component. We will denote the density function of  $\mu$  simply as  $f$ . The weight  $w_i$  depends on the normalized partition function of  $\mu_i$ , which involves computing an intractable integral. Another difficulty is that the number of orthants grows super-exponentially with the number of leaves. It is therefore necessary to approximate  $\mu$  with a Dirac mixture of posterior samples, which does not require to compute any partition functions.

## 5. APPROXIMATION OF THE POSTERIOR DISTRIBUTION

To summarize the posterior  $\mu$ , we would like to compute a point estimate

$$\hat{T} = \arg \min_{y \in \mathcal{T}_n} \int_{\mathcal{T}_n} \mathcal{L}(x, y) d\mu(x) ,$$

for an appropriate loss function  $\mathcal{L}$ , as discussed in Section 2. Unfortunately, the expected loss is difficult to compute and we therefore rely on an approximation by replacing  $\mu$  with the Dirac mixture

$$\pi := \frac{1}{K} \sum_{k=1}^K \delta_{T_k}$$

of  $K$  samples from  $\mu$ . By the ergodic theorem, we have the convergence

$$\int_{\mathcal{T}_n} \mathcal{L}(x, y) d\pi(x) = \frac{1}{K} \sum_{k=1}^K \mathcal{L}(T_k, y) \xrightarrow{a.s.} \int_{\mathcal{T}_n} \mathcal{L}(x, y) d\mu(x) ,$$

for any  $y \in \mathcal{T}_n$  as  $K \rightarrow \infty$  (cf. Robert and Casella [1999]). A set of posterior samples can be obtained with the Metropolis-Hastings algorithm (Metropolis et al. [1953], Hastings [1970]) without having to evaluate the weights  $w_i$  of the single components of  $\mu$ . The algorithm constructs a Markov chain with  $\mu$  as the stationary distribution. Let  $T_k$  be a sample from  $\mu$  with edge set  $\mathcal{E}$ . A new sample  $T_{k+1}$  is generated by the Markov chain conditional on the current sample  $T_k$ . The algorithm relies on a proposal distribution with density function  $q(\cdot | T_k)$ , which selects an edge  $e \in \mathcal{E}$  and replaces it with another edge. We thereby obtain a new tree  $T'$  that we accept as the next sample  $T_{k+1}$  with probability

$$\rho(T_k, T') = \min \left\{ 1, \frac{f(T')q(T_k | T')}{f(T_k)q(T' | T_k)} \right\} ,$$

and otherwise  $T_{k+1} = T_k$ , where  $f$  denotes the density function of  $\mu$ . Note that the normalization constant of  $\mu$  cancels in the ratio. The proposed tree  $T'$  lies in the same orthant as  $T_k$  with probability  $\tau$ . In this case, a new edge length is proposed, which is a draw from a normal distribution centered at  $|e|$ . However, with probability  $1 - \tau$  the proposed tree lies within one of the neighboring orthants,

by replacing the edge  $e$  with one of two other possible edges (see Figure 5). The transition measure of the Markov chain is given by

$$\kappa(x, dy) = \rho(x, y)q(y | x)dy + (1 - r(x))\delta_x(dy)$$

with  $r(x) = \int_{\mathcal{T}_n} \rho(x, y)q(y | x)dy$ , which satisfies the detailed balance condition and therefore has  $\mu$  as invariant distribution (cf. Robert and Casella [1999]).




FIGURE 5. Possible MCMC moves to neighboring orthants. The edge  $e$  of tree (a) can be replaced by two other edges in neighboring orthants shown in trees (b) and (c). The leaves labeled from zero to three may also represent more complex subtrees.

In this paper, we will focus on approximating the mean and median of the posterior distribution. The problem of finding a point estimate therefore reduces to computing the geometric median

$$(1) \quad \Psi(\bar{T}) := \arg \min_{x \in \mathcal{T}_n} \sum_{k=1}^K d(x, T_k) ,$$

and the Fréchet mean

$$(2) \quad \Xi(\bar{T}) := \arg \min_{x \in \mathcal{T}_n} \sum_{k=1}^K d(x, T_k)^2 ,$$

of a finite set  $\bar{T} := (T_1, \dots, T_K)$  of trees from  $\mathcal{T}_n$ . Since both the median and mean are defined as minimizers of “nice” convex functions on tree space, we get the following. The median  $\Psi(\bar{T})$  always exists and it is unique unless all the trees  $T_1, \dots, T_K$  lie on a geodesic. The existence and uniqueness of  $\Xi(\bar{T})$  is a consequence of strong convexity of the minimized function. The interested reader is referred to [Jost, 1997, Theorem 3.2.1] and [Sturm, 2003, Proposition 4.4]. The proofs can be also found in [Bačák, 2013, Theorem 2.4]. Note that medians and means are well-defined on arbitrary Hadamard spaces.

We will now turn to the question of how to *compute* medians and means of a given set of trees, since the formulas (1) and (2) do not provide us with direct algorithms. It turns out that efficient approximation methods from optimization can be extended into Hadamard spaces and applied to median and mean computations. For explicit algorithms, the reader is referred to [Bačák, 2013, Section 4], who presents a random and a cyclic-order version of an approximation algorithm. Note that we consider *unweighted* medians and means here, which slightly simplifies the formulas in [Bačák, 2013, Section 4]. Interestingly, the random-order version of the algorithm for computing the mean can be alternatively justified via the law

of large numbers due to Sturm [2002], as was independently observed by Bačák [2013] and Miller et al. [2012]. For reader’s convenience, the random-order version for unweighted medians and means is outlined in Appendix A.

The approximation algorithms for computing medians and means use (at each step) the algorithm for finding a geodesic in tree space by Owen and Provan [2011]. At a crucial stage, the Owen-Provan algorithm computes a maximal flow. Our implementation [1] does this using an encoding of the max flow problem as an integer program. In contrast, our implementation [2] uses the push-relabel algorithm due to Goldberg and Tarjan [1988].

## 6. RESULTS

We tested the model on the small ribosomal subunit rRNA sequence<sup>3</sup> contained in the UCSC Hg19 multiz46way alignment. Two example computations are presented in the following. For both, a gamma prior on edge lengths with shape parameter  $r = 1$  and scale parameter  $\lambda = 0.1$  is used. This choice of parameters reflects our believe that edge lengths can be very small and allows the sampler to easily switch between orthants. On the other hand, a shape parameter of  $r > 1$  would cause the posterior to have more distinct modes. The Dirichlet prior on the stationary distribution has pseudocounts  $\alpha_x = 0.2$  for all  $x \in \mathbb{A}$ , which reflects our believe that the data set is a conserved genomic region. The unnormalized log posterior of MCMC samples is shown in Figure 6 for both examples, which indicates that the Markov chain mixes well.

In the first example we considered only five species, namely *Cavia porcellus* (cavPor3), *Dipodomys ordii* (dipOrd1), *Mus musculus* (mm9), *Ochotona princeps* (ochPri2), and *Spermophilus tridecemlineatus* (speTri1). The approximated posterior expectation was computed on the last 16000 samples and is shown in Figure 7. In this plot, edge lengths are visualized as distances in the horizontal direction only. Consider the edge

$$e_1 : (\text{ochPri2, mm9} \mid \text{dipOrd1, speTri1, cavPor3}) ,$$

which connects the subtree of ochPri2 and mm9 with the rest of the phylogenetic tree. It has a relatively small length, caused by an uncertainty about the tree topology. The density of the full posterior  $\mu$  is of course difficult to visualize, but we can have a look at a small section. For this, consider the edges

$$e_2 : (\text{ochPri2, dipOrd1} \mid \text{mm9, speTri1, cavPor3}) , \quad \text{and} \\ e_3 : (\text{ochPri2, speTri1} \mid \text{dipOrd1, mm9, cavPor3}) ,$$

which are not compatible with  $e_1$  and can replace it in the phylogenetic tree. Figure 8 shows histograms of lengths for the three edges, which can be interpreted as an estimate of a marginal posterior density. The histogram was generated by counting how often each of the edges appeared in the set of samples. We call the density marginal, because we did not consider a specific topology of the remaining tree. Hence, it does not reflect a single orthant of tree space. The estimate has positive support on all three edges. While Figure 8(a) shows only a single mode, we clearly have a bimodality in Figure 8(b). Although the edge  $e_1$  is present in the posterior expectation (Figure 7), its length is reduced due to the mass on  $|e_2|$  and

<sup>3</sup>The Hg19 coordinates of the sequence are chrM: 1686-2059.

$|e_3|$ . This correctly represents our uncertainty about the topology of the tree. For instance, an equal weight on all three edges would cause the posterior expectation to have a non-binary branching point.

In a second example we increased the number of species to 14 and used 10 Markov chains in parallel. The mean and median are shown in Figure 9. Both trees are almost identical, which shows that the posterior distribution seems to have no long tails. However, two inner edges that separate `dipOrd1`, `cavPor3`, and `speTri1`, have very small length, which indicate that also in this example there is uncertainty about the exact topology of the tree. Figure 10 shows a marginal posterior estimate for the edges

$$e_4 : (\text{cavPor3}, \text{speTri1} \mid \text{ochPri2}, \text{oryCun2}, \text{dipOrd1}, \text{mm9}, \text{rn4}, \dots) , \quad \text{and}$$

$$e_5 : (\text{dipOrd1}, \text{speTri1} \mid \text{ochPri2}, \text{oryCun2}, \text{cavPor3}, \text{mm9}, \text{rn4}, \dots) ,$$

which shows a strong bimodality. However, the interpretation of such marginal estimates is difficult because of the much richer structure of the full tree space. In both examples, edge lengths are much larger than in the phylogenetic trees published by UCSC. For instance, the leaf edge connected to `hg19` in Figure 9 has a length of 0.0628 while the same edge in the UCSC tree has only a length of 0.006591. This is a consequence of having a site specific stationary distribution.

## 7. CONCLUSION

We have presented a statistical model for the inference of phylogenetic trees from multiple sequence alignments. The model is formulated on tree space by Billera et al. [2001], which is an Hadamard space and therefore allows to define the mean and median of a probability distribution. The approximation of posterior quantities is complicated and we have summarized some recent developments that contributed to this work. Despite the fact that the posterior distribution will in most cases be highly nontrivial, we demonstrated on a simple example that the mean or median as a point estimate can reflect the uncertainty about the topology of the tree. Methods for phylogenetic tree inference that rely on MCMC sampling often compute a (majority rule) consensus tree. Such a tree can be justified from decision theoretic principles. However, we believe that we have proposed a more rigorous approach to solve this issue. Since our statistical model is defined on tree space, its inherent properties become part of the model, which clearly has implications on posterior estimates. Certainly, a disadvantage of MCMC approximations in phylogenetic inference is that the number of different topologies grows super-exponentially with the number of leaves. The method might thus be inappropriate for the inference of large trees as the approximation of the posterior quantities might require too many samples.

We provide two freely available implementations at [1] and [2].




FIGURE 6. Value of the log posterior density (not normalized) of MCMC samples on the small ribosomal subunit rRNA sequence a lignment.




FIGURE 7. Example 1: Fréchet mean of posterior samples with an approximate variance of 0.0248.




FIGURE 8. Example 1: Marginal posterior density estimate of three edges. The posterior expectation is shown as a vertical line.

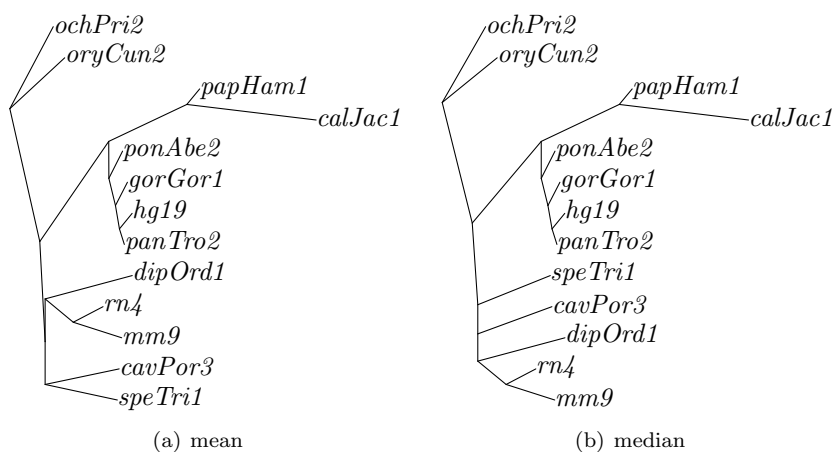


FIGURE 9. Example 2: Fréchet mean and median of posterior samples. The estimated posterior variance is 0.072.




FIGURE 10. Example 2: Marginal posterior density estimate of two edges. The posterior expectation is shown as a vertical line.

## REFERENCES

- [1]. TFBayes. URL <https://github.com/pbenner/tfbayes>.
- [2]. TrAP. URL <https://github.com/bacak/TrAP>.
- M. Bačák. Computing medians and means in Hadamard spaces. *Preprint, arXiv:1210.2145*, 2013.
- L.J. Billera, S.P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.*, 27(4):733–767, 2001. ISSN 0196-8858. doi: 10.1006/aama.2001.0759. URL <http://dx.doi.org/10.1006/aama.2001.0759>.
- A. Dress, K.T. Huber, J. Koolen, V. Moulton, and A. Spillner. *Basic phylogenetic combinatorics*. Cambridge University Press, Cambridge, 2012. ISBN 978-0-521-76832-0.
- A. Drummond and A. Rambaut. BEAST: Bayesian evolutionary analysis by sampling trees. *BMC Evolutionary Biology*, 1(214), 2007. doi: 10.1186/1471-2148-7-214. URL <http://dx.doi.org/10.1186/1471-2148-7-214>.
- J. Felsenstein. Evolutionary trees from DNA sequences: a maximum likelihood approach. *Journal of molecular evolution*, 17(6):368–376, 1981.
- A.V. Goldberg and R.E. Tarjan. A new approach to the maximum-flow problem. *J. Assoc. Comput. Mach.*, 35(4):921–940, 1988. ISSN 0004-5411. doi: 10.1145/48014.61051. URL <http://dx.doi.org/10.1145/48014.61051>.
- S. Guindon and O. Gascuel. A simple, fast and accurate algorithm to estimate large phylogenies by maximum likelihood. *Systematic Biology*, 52:696–704, 2003.
- W.K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57:97–109, 1970.
- M.T. Holder, J. Sukumaran, and P.O. Lewis. A Justification for Reporting the Majority-Rule Consensus Tree in Bayesian Phylogenetics. *Systematic Biology*, 57:814–821, 2003.
- J.P. Huelsenbeck and F. Ronquist. MRBAYES: Bayesian inference of phylogenetic trees. *Bioinformatics*, 17(8):754–755, 2001.
- P. M. Huggins, W. Li, D. Haws, T. Friedrich, J. Liu, and R. Yoshida. Bayes Estimators for Phylogenetic Reconstruction. *Systematic Biology*, 60(4):528–540, 2011. doi: 10.1093/sysbio/syr021. URL <http://sysbio.oxfordjournals.org/content/60/4/528.abstract>.
- A. Isaev. *Introduction to Mathematical Methods in Bioinformatics (Universitext)*. Springer, 2006. ISBN 3540219730.
- J. Jost. *Nonpositive curvature: geometric and analytic aspects*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 1997. ISBN 3-7643-5736-3.
- N. Lartillot and H. Philippe. A bayesian mixture model for across-site heterogeneities in the amino-acid replacement process. *Molecular Biology and Evolution*, 21(6):1095–1109, 2004. doi: 10.1093/molbev/msh112. URL <http://dx.doi.org/10.1093/molbev/msh112>.
- N. Lartillot, T. Lepage, and S. Blanquart. PhyloBayes 3: A Bayesian software package for phylogenetic reconstruction and molecular dating. *Bioinformatics*, 25(17):2286–2288, 2009.
- WH Li, CC Luo, and CI Wu. Evolution of dna sequences. *Molecular evolutionary genetics. Plenum, New York*, 1:1–94, 1985.
- N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1091, 1953.



- E. Miller, M. Owen, and S. Provan. Averaging metric phylogenetic trees. *Preprint, arXiv:1211.7046v1*, 2012.
- R. Nielsen, editor. *Statistical methods in molecular evolution*. Springer Science+Business Media, Inc., 2005.
- T. M. W. Nye. Principal components analysis in the space of phylogenetic trees. *Ann. Statist.*, 39(5):2716–2739, 2011. ISSN 0090-5364. doi: 10.1214/11-AOS915. URL <http://dx.doi.org/10.1214/11-AOS915>.
- M. Owen and S. Provan. A fast algorithm for computing geodesic distances in tree space. *IEEE/ACM Trans. Computational Biology and Bioinformatics*, 8:2–13, 2011.
- C.P. Robert. *The Bayesian choice*. Springer Texts in Statistics. Springer-Verlag, New York, second edition, 2001. ISBN 0-387-95231-4. From decision-theoretic foundations to computational implementation, Translated and revised from the French original by the author.
- C.P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer-Verlag, 1 edition, 1999. ISBN 038798707X.
- M.J. Schervish. *Theory of Statistics*. Springer Series in Statistics. Springer, 1995. ISBN 9780387945460.
- C. Semple and M. Steel. *Phylogenetics*, volume 24 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2003. ISBN 0-19-850942-1.
- K.-T. Sturm. Nonlinear martingale theory for processes with values in metric spaces of nonpositive curvature. *Ann. Probab.*, 30(3):1195–1222, 2002. ISSN 0091-1798. doi: 10.1214/aop/1029867125. URL <http://dx.doi.org/10.1214/aop/1029867125>.
- K.-T. Sturm. Probability measures on metric spaces of nonpositive curvature. In *Heat kernels and analysis on manifolds, graphs, and metric spaces (Paris, 2002)*, volume 338 of *Contemp. Math.*, pages 357–390. Amer. Math. Soc., Providence, RI, 2003.
- W.W. Wasserman and A. Sandelin. Applied bioinformatics for the identification of regulatory elements. *Nature Review Genetics*, 5(4):276–287, 2004.
- W.C. Wheeler and K.M. Pickett. Topology-Bayes versus Clade-Bayes in Phylogenetic Analysis. *Molecular Biology and Evolution*, 25(2):447–453, 2008. doi: 10.1093/molbev/msm274.
- Z. Yang. Among-site rate variation and its impact on phylogenetic analyses. *Trends in Ecology and Evolution*, 11(9):367–372, 1996.

#### APPENDIX A. ALGORITHMS FOR COMPUTING MEDIANS AND MEANS

The approximation algorithms for computing medians and means were introduced by Bačák [2013] and we refer the interested reader therein for the proofs of convergence and further details. The algorithms rely upon a well-known optimization technique called the proximal point method. Interestingly, the algorithm for computing the mean can be alternatively justified via the law of large numbers due to Sturm [2002], as was independently observed by Bačák [2013] and Miller et al. [2012].

**A.1. Algorithms for computing medians.** Let us first describe the algorithm for computing a median of a given set  $T_1, \dots, T_K \in \mathcal{T}_n$ .

We set  $x_0 := T_1$  and suppose that at the  $i$ -th iteration we have an approximation  $x_i \in \mathcal{T}_n$  of  $\Psi(\bar{T})$ . To find  $x_{i+1}$ , a tree  $T_k$  is selected from our set of trees  $T_1, \dots, T_K$  at random and we define  $x_{i+1}$  as a point on the geodesic between  $x_i$  and  $T_k$ . (In other words  $x_{i+1}$  is a convex combination of  $x_i$  and  $T_k$ .) The position of  $x_{i+1}$  on this geodesic is determined by a parameter  $t_i \in [0, 1]$ , which is computed at each iteration. By this procedure, we obtain a sequence of trees  $x_1, x_2, \dots$  which is known converge to a median of  $T_1, \dots, T_K$ .

**Algorithm A.1** (Computing median, random order version). *Let  $x_0 := T_1$ . At each step  $i \in \mathbb{N}_0$ , choose randomly  $r_i \in \{1, \dots, K\}$  according to the uniform distribution and put*

$$(3) \quad x_{i+1} := (1 - t_i)x_i + t_i T_{r_i},$$

with  $t_i$  defined by

$$t_i := \min \left\{ 1, \frac{1}{(i+1)d(T_{r_i}, x_i)} \right\},$$

for each  $i \in \mathbb{N}_0$ .

It is important to insist on the uniform distribution on the set  $\{1, \dots, K\}$ , that is, no tree of  $T_1, \dots, T_K$  be privileged. Only then we obtain a sequence of trees  $x_1, x_2, \dots$  which converges to a median of  $T_1, \dots, T_K$ .

**A.2. Algorithms for computing means.** Computing the mean is similar to the computation of the median. As a matter of fact it only differs in the coefficients determining the position of  $x_{i+1}$  on the geodesic from  $x_i$  to  $T_k$ . Again, let  $T_1, \dots, T_K \in \mathcal{T}_n$  be a finite set of trees from  $\mathcal{T}_n$ . The following approximation algorithms generate a sequence of trees  $x_1, x_2, \dots$  from  $\mathcal{T}_n$  which converges to  $\Xi(\bar{T})$ . At each iteration a tree  $T_k$  is selected at random and we obtain the following algorithm.

**Algorithm A.2** (Computing mean, random order version). *Let  $x_0 := T_1$  and at each step  $i \in \mathbb{N}_0$ , choose randomly  $r_i \in \{1, \dots, K\}$  according to the uniform distribution and put*

$$x_{i+1} := \frac{1}{i+1}x_i + \frac{i}{i+1}T_{r_i}.$$

The above algorithms have also their deterministic counterparts, where we choose the trees from the input set in a cyclic order instead of randomly; see Bačák [2013]. Even though both random and cyclic versions converge to the same value, there is no theorem on which one converges faster. Our computational studies however suggest that the random version is better.

MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES, INSELSTR. 22, 04103 LEIPZIG, GERMANY

*E-mail address:* philipp.benner@mis.mpg.de

*E-mail address:* bacak@mis.mpg.de