
Efficient Projection onto the Perfect Phylogeny Model

Bei Jia*
jiabe@bc.edu

Surjyendu Ray
raysc@bc.edu

Sam Safavi
safavisa@bc.edu

José Bento
jose.bento@bc.edu

Boston College

Abstract

Several algorithms build on the *perfect phylogeny model* to infer evolutionary trees. This problem is particularly hard when evolutionary trees are inferred from the fraction of genomes that have mutations in different positions, across different samples. Existing algorithms might do extensive searches over the space of possible trees. At the center of these algorithms is a projection problem that assigns a fitness cost to phylogenetic trees. In order to perform a wide search over the space of the trees, it is critical to solve this projection problem fast. In this paper, we use Moreau’s decomposition for proximal operators, and a tree reduction scheme, to develop a new algorithm to compute this projection. Our algorithm terminates with an exact solution in a finite number of steps, and is extremely fast. In particular, it can search over all evolutionary trees with fewer than 11 nodes, a size relevant for several biological problems (more than 2 billion trees) in about 2 hours.

1 Introduction

The perfect phylogeny model (PPM) [1, 2] is used in biology to study evolving populations. It assumes that the same position in the genome never mutates twice, hence mutations only accumulate.

Consider a population of organisms evolving under the PPM. The evolution process can be described by a labeled rooted tree, $T = (r, \mathcal{V}, \mathcal{E})$, where r is the root, i.e., the common oldest ancestor, the nodes \mathcal{V} are the mutants, and the edges \mathcal{E} are mutations acquired between older and younger mutants. Since each position in the genome only mutates once, we can associate with each node $v \neq r$, a unique mutated position, the mutation associated to the ancestral edge of v . By convention, let us associate with the root r , a null mutation that is shared by all mutants in T . This allows us to refer to each node $v \in \mathcal{V}$ as both a mutation in a position in the genome (the mutation associated to the ancestral edge of v), and a mutant (the mutant with the fewest mutations that has a mutation v). Hence, without loss of generality, $\mathcal{V} = \{1, \dots, q\}$, $\mathcal{E} = \{2, \dots, q\}$, where q is the length of the genome, and $r = 1$ refers to both the oldest common ancestor and the null mutation shared by all.

One very important use of the PPM is to infer how mutants of a common ancestor evolve [3–8]. A common type of data used for this purpose is the frequency, with which different positions in the genome mutate across multiple samples, obtained, e.g., from whole-genome or targeted deep sequencing [9]. Consider a sample s , one of p samples, obtained at a given stage of the evolution process. This sample has many mutants, some with the same genome, some with different genomes. Let $F \in \mathbb{R}^{q \times p}$ be such that $F_{v,s}$ is the fraction of genomes in s with a mutation in position v in the genome. Let $M \in \mathbb{R}^{q \times p}$ be such that $M_{v,s}$ is the fraction of mutant v in s . By definition, the columns of M must sum to 1. Let $U \in \{0, 1\}^{q \times q}$ be such that $U_{v,v'} = 1$, if and only if mutant v is an ancestor of mutant v' , or if $v = v'$. We denote the set of all possible U matrices, M matrices and labeled rooted trees T , by \mathcal{U} , \mathcal{M} and \mathcal{T} , respectively. See Figure 1 for an illustration. The PPM implies

$$F = UM. \tag{1}$$

Our work contributes to the problem of inferring clonal evolution from mutation-frequencies: *How do we infer M and U from F ?* Note that finding U is the same as finding T (see Lemma B.2).

*Bei Jia is currently with Element AI.




Figure 1: Black lines are genomes. Red circles indicate mutations. g_i is the mutant with fewest mutations with position i mutated. Mutation 1, the mutation in the null position, $i = 1$, is shared by all mutants. g_1 is the organism before mutant evolution starts. In sample $s = 3$, $2/10$ of the mutants are type g_2 , hence $M_{2,3} = 2/10$, and $3/10$ of the mutations occur in position 7, hence $F_{7,3} = 3/10$. The tree shows the mutants’ evolution.

Although model (1) is simple, simultaneously inferring M and U from F can be hard [3]. One popular inference approach is the following optimization problem over U , M and F ,

$$\min_{U \in \mathcal{U}} \mathcal{C}(U), \quad (2)$$

$$\mathcal{C}(U) = \min_{M, F \in \mathbb{R}^{q \times p}} \|\hat{F} - F\| \text{ subject to } F = UM, M \geq 0, M^T \mathbf{1} = \mathbf{1}, \quad (3)$$

where $\|\cdot\|$ is the Frobenius norm, and $\hat{F} \in \mathbb{R}^{q \times p}$ contains the measured fractions of mutations per position in each sample, which are known and fixed. In a nutshell, we want to project our measurement \hat{F} onto the space of valid PPM models.

Problem (2) is a hard mixed integer-continuous optimization problem. To approximately solve it, we might find a finite subset $\{U_i\} \subset \mathcal{U}$, that corresponds to a “heuristically good” subset of trees, $\{T_i\} \subset \mathcal{T}$, and, for each fixed matrix U_i , solve (3), which is a convex optimization problem. We can then return T_x , where $x \in \arg \min_i \mathcal{C}(U_i)$. Fortunately, in many biological applications, e.g., [3–8], the reconstructed evolutionary tree involves a very small number of mutated positions, e.g., $q \leq 11$. In practice, a position v might be an *effective* position that is a cluster of multiple real positions in the genome. For a small q , we can compute $\mathcal{C}(U)$ for many trees, and hence approximate M , U , and get uncertainty measures for these estimates. This is important, since data is generally scarce and noisy.

Contributions: (i) we propose a new algorithm to compute $\mathcal{C}(U)$ exactly in $\mathcal{O}(q^2p)$ steps, the first non-iterative algorithm to compute $\mathcal{C}(U)$; (ii) we compare its performance against state-of-the-art iterative algorithms, and observe a much faster convergence. In particular, our algorithm scales much faster than $\mathcal{O}(q^2p)$ in practice; (iii) we implement our algorithm on a GPU, and show that it computes the cost of all (more than 2 billion) trees with ≤ 11 nodes, in ≤ 2.5 hours.

2 Related work

A problem related to ours, but somewhat different, is that of inferring a phylogenetic tree from single-cell whole-genome sequencing data. Given all the mutations in a set of mutants, the problem is to arrange the mutants in a phylogenetic tree, [10, 11]. Mathematically, this corresponds to inferring T from partial or corrupted observation of U . If the PPM is assumed, and all the mutations of all the mutants are correctly observed, this problem can be solved in linear time, e.g., [12]. In general, this problem is equivalent to finding a minimum cost Steiner tree on a hypercube, whose nodes and edges represent mutants and mutations respectively, a problem known to be hard [13].

We mention a few works on clonality inference, based on the PPM, that try to infer both U and M from \hat{F} . No previous work solves problem (2) exactly in general, even for trees of size $q \leq 11$. Using our fast projection algorithm, we can solve (2) exactly by searching over all trees, if $q \leq 11$. Ref. [3] (*Ancestree*) reduces the space of possible trees \mathcal{T} to subtrees of a heuristically constructed DAG. The authors use the element-wise 1-norm in (3) and, after introducing more variables to linearize the product UM , reduce this search to solving a MILP, which they try to solve via branch and bound. Ref. [6] (*CITUP*) searches the space of all *unlabeled* trees, and, for each unlabeled tree, tries to solve an MIQP, again using branch and bound techniques, which finds a labeling for the unlabeled tree, and simultaneously minimizes the distance $\|\hat{F} - F\|$. Refs. [5] and [14] (*PhyloSub/PhyloWGS*), use a stochastic model to sample trees that are likely to explain the data. Their model is based on [15], which generates hierarchical clusterings of objects, and from which lineage trees can be formed. A score is then computed for these trees, and the highest scoring trees are returned.

Procedure (2) can be justified as MLE if we assume the stochastic model $\hat{F} = F + \mathcal{N}(0, I\sigma^2)$, where F , U and M satisfy the PPM model, and $\mathcal{N}(0, I\sigma^2)$ represents additive, component-wise,

Gaussian measurement noise, with zero mean and covariance $I\sigma^2$. Alternative stochastic models can be assumed, e.g., as $M - U^{-1}\hat{F} = \mathcal{N}(0, I\sigma^2)$, where M is non-negative and its columns must sum to one, and $\mathcal{N}(0, I\sigma^2)$ is as described before. For this model, and for each matrix U , the cost $\mathcal{C}(U)$ is a projection of $U^{-1}\hat{F}$ onto the probability simplex $M \geq 0, M^T \mathbf{1} = \mathbf{1}$. Several fast algorithms are known for this problem, e.g., [16–20] and references therein. In a pq -dimensional space, the exact projection onto the simplex can be done in $\mathcal{O}(qp)$ steps.

Our algorithm is the first to solve (3) exactly in a finite number of steps. We can also use iterative methods to solve (3). One advantage of our algorithm is that it has no tuning parameters, and requires no effort to check for convergence for a given accuracy. Since iterative algorithms can converge very fast, we numerically compare the speed of our algorithm with different implementations of the Alternating Direction Method of Multipliers (ADMM) [21], which, if properly tuned, has the fastest convergence rate among all first-order methods [22] under some convexity assumptions, and is known to produce good solutions for several other kinds of problems, even for non-convex ones [23–29].

3 Main results

We now state our main results, and explain the ideas behind their proofs. Detailed proofs can be found in the Appendix.

Our algorithm computes $\mathcal{C}(U)$ and minimizers of (3), resp. M^* and F^* , by solving an equivalent problem. Without loss of generality, we assume that $p = 1$, since, by squaring the objective in (3), it decomposes into p independent problems. Sometimes we denote $\mathcal{C}(U)$ by $\mathcal{C}(T)$, since given U , we can specify T , and vice-versa. Let \bar{i} be the closest ancestor of i in $T = (r, \mathcal{V}, \mathcal{E})$. Let Δi be the set of all the ancestors of i in T , plus i . Let ∂i be the set of children of i in T .

Theorem 3.1 (Equivalent formulation). *Problem (3) can be solved by solving*

$$\min_{t \in \mathbb{R}} t + \mathcal{L}(t), \quad (4)$$

$$\mathcal{L}(t) = \min_{Z \in \mathbb{R}^q} \frac{1}{2} \sum_{i \in \mathcal{V}} (Z_i - Z_{\bar{i}})^2 \text{ subject to } Z_i \leq t - N_i, \forall i \in \mathcal{V}, \quad (5)$$

where $N_i = \sum_{j \in \Delta i} \hat{F}_j$, and, by convention, $Z_{\bar{i}} = 0$ for $i = r$. In particular, if t^* minimizes (4), Z^* minimizes (5) for $t = t^*$, and M^*, F^* minimize (3), then

$$M_i^* = -Z_i^* + Z_{\bar{i}}^* + \sum_{r \in \partial i} (Z_r^* - Z_{\bar{r}}^*) \text{ and } F_i^* = -Z_i^* + Z_{\bar{i}}^*, \forall i \in \mathcal{V}. \quad (6)$$

Furthermore, t^*, M^*, F^* and Z^* are unique.

Theorem 3.1 comes from a dual form of (3), which we build using Moreau’s decomposition [30].

3.1 Useful observations

Let $Z^*(t)$ be the unique minimizer of (5) for some t . The main ideas behind our algorithm depend on a few simple properties of the paths $\{Z^*(t)\}$ and $\{\mathcal{L}'(t)\}$, the derivative of $\mathcal{L}(t)$ with respect to t . Note that \mathcal{L} is also a function of N , as defined in Theorem 3.1, which depends on the input data \hat{F} .

Lemma 3.2. $\mathcal{L}(t)$ is a convex function of t and N . Furthermore, $\mathcal{L}(t)$ is continuous in t and N , and $\mathcal{L}'(t)$ is non-decreasing with t .

Lemma 3.3. $Z^*(t)$ is continuous as a function of t and N . $Z^*(t^*)$ is continuous as a function of N .

Let $\mathcal{B}(t) = \{i : Z^*(t)_i = t - N_i\}$, i.e., the set of components of the solution at the boundary of (5). Variables in \mathcal{B} are called *fixed*, and we call other variables *free*. Free (resp. fixed) nodes are nodes corresponding to free (resp. fixed) variables.

Lemma 3.4. $\mathcal{B}(t)$ is piecewise constant in t .

Consider dividing the tree $T = (r, \mathcal{V}, \mathcal{E})$ into subtrees, each with at least one free node, using $\mathcal{B}(t)$ as separation points. See Figure 4 in Appendix A for an illustration. Each $i \in \mathcal{B}(t)$ belongs to at most $\text{degree}(i)$ different subtrees, where $\text{degree}(i)$ is the degree of node i , and each $i \in \mathcal{V} \setminus \mathcal{B}(t)$ belongs exactly to one subtree. Let T_1, \dots, T_k be the set of resulting (rooted, labeled) trees. Let $T_w = (r_w, \mathcal{V}_w, \mathcal{E}_w)$, where the root r_w is the closest node in T_w to r . We call $\{T_w\}$ the subtrees induced by $\mathcal{B}(t)$. We define $\mathcal{B}_w(t) = \mathcal{B}(t) \cap \mathcal{V}_w$, and, when it does not create ambiguity, we drop the index t in $\mathcal{B}_w(t)$. Note that different $\mathcal{B}_w(t)$ ’s might have elements in common. Also note that, by construction, if $i \in \mathcal{B}_w$, then i must be a leaf of T_w , or the root of T_w .

Definition 3.5. The (T_w, \mathcal{B}_w) -problem is the optimization problem over $|\mathcal{V}_w \setminus \mathcal{B}(t)|$ variables

$$\min_{\{Z_j : j \in \mathcal{V}_w \setminus \mathcal{B}(t)\}} (1/2) \sum_{j \in \mathcal{V}_w} (Z_j - Z_{\bar{j}})^2, \quad (7)$$

where \bar{j} is the parent of j in T_w , $Z_{\bar{j}} = 0$ if $j = r_w$, and $Z_j = Z^*(t)_j = t - N_j$ if $j \in \mathcal{B}_w(t)$.

Lemma 3.6. *Problem (5) decomposes into k independent problems. In particular, the minimizers $\{Z^*(t)_j : j \in \mathcal{V}_w \setminus \mathcal{B}(t)\}$ are determined as the solution of the (T_w, \mathcal{B}_w) -problem. If $j \in \mathcal{V}_w$, then $Z^*(t)_j = c_1 t + c_2$, where c_1 and c_2 depend on j but not on t , and $0 \leq c_1 \leq 1$.*

Lemma 3.7. *$Z^*(t)$ and $\mathcal{L}'(t)$ are piecewise linear and continuous in t . Furthermore, $Z^*(t)$ and $\mathcal{L}'(t)$ change linear segments if and only if $\mathcal{B}(t)$ changes.*

Lemma 3.8. *If $t \leq t'$, then $\mathcal{B}(t') \subseteq \mathcal{B}(t)$. In particular, $\mathcal{B}(t)$ changes at most q times with t .*

Lemma 3.9. *$Z^*(t)$ and $\mathcal{L}'(t)$ have less than $q + 1$ different linear segments.*

3.2 The Algorithm

In a nutshell, our algorithm computes the solution path $\{Z^*(t)\}_{t \in \mathbb{R}}$ and the derivative $\{\mathcal{L}'(t)\}_{t \in \mathbb{R}}$. From these paths, it finds the unique t^* , at which

$$d(t + \mathcal{L}(t))/dt = 0|_{t=t^*} \Leftrightarrow \mathcal{L}'(t^*) = -1. \quad (8)$$

It then evaluates the path $Z^*(t)$ at $t = t^*$, and uses this value, along with (6), to find M^* and F^* , the unique minimizers of (3). Finally, we compute $\mathcal{C}(T) = \|\hat{F} - F^*\|$.

We know that $\{Z^*(t)\}$ and $\{\mathcal{L}'(t)\}$ are continuous piecewise linear, with a finite number of different linear segments (Lemmas 3.7, 3.8 and 3.9). Hence, to describe $\{Z^*(t)\}$ and $\{\mathcal{L}'(t)\}$, we only need to evaluate them at the *critical values*, $t_1 > t_2 > \dots > t_k$, at which $Z^*(t)$ and $\mathcal{L}'(t)$ change linear segments. We will later use Lemma 3.7 as a criteria to find the critical values. Namely, $\{t_i\}$ are the values of t at which, as t decreases, new variables become fixed, and $\mathcal{B}(t)$ changes. Note that variables never become free once fixed, by Lemma 3.8, which also implies that $k \leq q$.

The values $\{Z^*(t_i)\}$ and $\{\mathcal{L}'(t_i)\}$ are computed sequentially as follows. If t is very large, the constraint in (5) is not active, and $Z^*(t) = \mathcal{L}(t) = \mathcal{L}'(t) = 0$. Lemma 3.7 tells us that, as we decrease t , the first critical value is the largest t for which this constraint becomes active, and at which $\mathcal{B}(t)$ changes for the first time. Hence, if $i = 1$, we have $t_i = \max_s \{N_s\}$, $Z^*(t_i) = \mathcal{L}'(t_i) = 0$, and $\mathcal{B}(t_i) = \arg \max_s \{N_s\}$. Once we have t_i , we compute the rates $Z'^*(t_i)$ and $\mathcal{L}''(t_i)$ from $\mathcal{B}(t_i)$ and T , as explained in Section 3.3. Since the paths are piecewise linear, derivatives are not defined at critical points. Hence, here, and throughout this section, these derivatives are taken from the left, i.e., $Z'^*(t_i) = \lim_{t \uparrow t_i} (Z^*(t_i) - Z^*(t))/(t_i - t)$ and $\mathcal{L}''(t_i) = \lim_{t \uparrow t_i} (\mathcal{L}'(t_i) - \mathcal{L}'(t))/(t_i - t)$.

Since $Z'^*(t)$ and $\mathcal{L}''(t)$ are constant for $t \in (t_{i+1}, t_i]$, for $t \in (t_{i+1}, t_i]$ we have

$$Z^*(t) = Z^*(t_i) + (t - t_i)Z'^*(t_i), \quad \mathcal{L}'(t) = \mathcal{L}'(t_i) + (t - t_i)\mathcal{L}''(t_i), \quad (9)$$

and the next critical value, t_{i+1} , is the largest $t < t_i$, for which new variables become fixed, and $\mathcal{B}(t)$ changes. The value t_{i+1} is found by solving for $t < t_i$ in

$$Z^*(t)_r = Z^*(t_i)_r + (t - t_i)Z'^*(t_i)_r = t - N_r, \quad (10)$$

and keeping the largest solution among all $r \notin \mathcal{B}$. Once t_{i+1} is computed, we update \mathcal{B} with the new variables that became fixed, and we obtain $Z^*(t_{i+1})$ and $\mathcal{L}'(t_{i+1})$ from (9). The process then repeats.

By Lemma 3.2, \mathcal{L}' never increases. Hence, we stop this process (a) as soon as $\mathcal{L}'(t_i) < -1$, or (b) when all the variables are in \mathcal{B} , and thus there are no more critical values to compute. If (a), let t_k be the last critical value with $\mathcal{L}'(t_k) > -1$, and if (b), let t_k be the last computed critical value. We use t_k and (9) to compute t^* , at which $\mathcal{L}'(t^*) = -1$ and also $Z^*(t^*)$. From $Z^*(t^*)$ we then compute M^* and F^* and $\mathcal{C}(U) = \|\hat{F} - F^*\|$.

The algorithm is shown compactly in Alg. 1. Its inputs are \hat{F} and T , represented, e.g., using a linked-nodes data structure. Its outputs are minimizers to (3). It makes use of a procedure *ComputeRates*, which we will explain later. This procedure terminates in $\mathcal{O}(q)$ steps and uses $\mathcal{O}(q)$ memory. Line 5 comes from solving (10) for t . In line 14, the symbols $M^*(Z^*, T)$ and $F^*(Z^*, T)$ remind us that M^* and F^* are computed from Z^* and T using (6). The correctness of Alg. 1 follows from the Lemmas in Section 3.1, and the explanation above. In particular, since there are at most $q + 1$ different linear regimes, the bound q in the for-loop does not prevent us from finding any critical value. Its time complexity is $\mathcal{O}(q^2)$, since each line completes in $\mathcal{O}(q)$ steps, and is executed at most q times.

Theorem 3.10 (Complexity). *Algorithm 1 finishes in $\mathcal{O}(q^2)$ steps, and requires $\mathcal{O}(q)$ memory.*

Theorem 3.11 (Correctness). *Algorithm 1 outputs the solution to (3).*

Algorithm 1 Projection onto the PPM (input: T and \hat{F} ; output: M^* and F^*)

```

1:  $N_i = \sum_{j \in \Delta_i} \hat{F}_j$ , for all  $i \in \mathcal{V}$  ▷ This takes  $\mathcal{O}(q)$  steps using a DFS, see proof of Theorem 3.10
2:  $i = 1, t_i = \max_r \{N_r\}, \mathcal{B}(t_i) = \arg \max_r \{N_r\}, Z^*(t_i) = \mathbf{0}, \mathcal{L}'(t_i) = \mathbf{0}$ .           ▷ Initialize
3: for  $i = 1$  to  $q$  do
4:    $(Z'^*(t_i), \mathcal{L}''(t_i)) = \text{ComputeRates}(\mathcal{B}(t_i), T)$                                ▷ Update rates of change
5:    $P = \{P_r : P_r = \frac{N_r + Z^*(t_i)_r - t_i Z'^*(t_i)_r}{1 - Z'^*(t_i)_r} \text{ if } r \notin \mathcal{B}(t_i), t_r < t_i, \text{ and } P_r = -\infty \text{ otherwise}\}$ 
6:    $t_{i+1} = \max_r P_r$                                                                  ▷ Update next critical value from (9)
7:    $\mathcal{B}(t_{i+1}) = \mathcal{B}(t_i) \cup \arg \max_r P_s$                                            ▷ Update list of fixed variables
8:    $Z^*(t_{i+1}) = Z^*(t_i) + (t_{i+1} - t_i)Z'^*(t_i)$                                ▷ Update solution path
9:    $\mathcal{L}'(t_{i+1}) = \mathcal{L}'(t_i) + (t_{i+1} - t_i)\mathcal{L}''(t_i)$                            ▷ Update objective's derivative
10:  if  $\mathcal{L}'(t_{i+1}) < -1$  then break                                               ▷ If already passed by  $t^*$ , then exit the for-loop
11: end for
12:  $t^* = t_i - \frac{1 + \mathcal{L}'(t_i)}{\mathcal{L}''(t_i)}$                                                ▷ Find solution to (8)
13:  $Z^* = Z^*(t_i) + (t^* - t_i)Z'^*(t_i)$                                            ▷ Find minimizers of (5) for  $t = t^*$ 
14: return  $M^*(Z^*, T), F^*(Z^*, T)$                                                ▷ Return solution to (3) using (6), which takes  $\mathcal{O}(q)$  steps

```

3.3 Computing the rates

We now explain how the procedure *ComputeRates* works. Recall that it takes as input the tree T and the set $\mathcal{B}(t_i)$, and it outputs the derivatives $Z'^*(t_i)$ and $\mathcal{L}''(t_i)$.

A simple calculation shows that if we compute $Z'^*(t_i)$, then computing $\mathcal{L}''(t_i)$ is easy.

Lemma 3.12. $\mathcal{L}''(t_i)$ can be computed from $Z'^*(t_i)$ in $\mathcal{O}(q)$ steps and with $\mathcal{O}(1)$ memory as

$$\mathcal{L}''(t_i) = \sum_{j \in \mathcal{V}} (Z'^*(t_i)_j - Z'^*(t_i)_{\bar{j}})^2, \quad (11)$$

where \bar{j} is the closest ancestor to j in T . We note that if $j \in \mathcal{B}(t_i)$, then, by definition, $Z'^*(t_i)_j = 1$. Assume now that $j \in \mathcal{V} \setminus \mathcal{B}(t_i)$. Lemma 3.6 implies we can find $Z'^*(t_i)_j$ by solving the $(T_w = (r_w, \mathcal{V}_w, \mathcal{E}_w), \mathcal{B}_w)$ -problem as a function of t , where w is such that $j \in \mathcal{V}_w$. In a nutshell, *ComputeRates* is a recursive procedure to solve all the (T_w, \mathcal{B}_w) -problems as an explicit function of t .

It suffices to explain how *ComputeRates* solves one particular (T_w, \mathcal{B}_w) -problem explicitly. To simplify notation, in the rest of this section, we refer to T_w and \mathcal{B}_w as T and \mathcal{B} . Recall that, by the definition of $T = T_w$ and $\mathcal{B} = \mathcal{B}_w$, if $i \in \mathcal{B}$, then i must be a leaf of T , or the root of T .

Definition 3.13. Consider a rooted tree $T = (r, \mathcal{V}, \mathcal{E})$, a set $\mathcal{B} \subseteq \mathcal{V}$, and variables $\{Z_j : j \in \mathcal{V}\}$ such that, if $j \in \mathcal{B}$, then $Z_j = \alpha_j t + \beta_j$ for some α and β . We define the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem as

$$\min_{\{Z_j : j \in \mathcal{V} \setminus \mathcal{B}\}} \frac{1}{2} \sum_{j \in \mathcal{V}} \gamma_j (Z_j - Z_{\bar{j}})^2, \quad (12)$$

where $\gamma > 0$, \bar{j} is the closest ancestor to j in T , and $Z_{\bar{j}} = 0$ if $j = r$.

We refer to the solution of the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem as $\{Z_j^* : j \in \mathcal{V} \setminus \mathcal{B}\}$, which uniquely minimizes (12). Note that (12) is unconstrained and its solution, Z^* , is a linear function of t . Furthermore, the (T_w, \mathcal{B}_w) -problem is the same as the $(T_w, \mathcal{B}_w, \mathbf{1}, -N, \mathbf{1})$ -problem, which is what we actually solve.

We now state three useful lemmas that help us solve any $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem efficiently.

Lemma 3.14 (Pruning). *Consider the solution Z^* of the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem. Let $j \in \mathcal{V} \setminus \mathcal{B}$ be a leaf. Then $Z_j^* = Z_{\bar{j}}^*$. Furthermore, consider the $(\tilde{T}, \mathcal{B}, \alpha, \beta, \gamma)$ -problem, where $\tilde{T} = (\tilde{r}, \tilde{\mathcal{V}}, \tilde{\mathcal{E}})$ is equal to T with node j pruned, and let its solution be \tilde{Z}^* . We have that $Z_i^* = \tilde{Z}_i^*$, for all $i \in \tilde{\mathcal{V}}$.*

Lemma 3.15 (Star problem). *Let T be a star such that node 1 is the center node, node 2 is the root, and nodes $3, \dots, r$ are leaves. Let $\mathcal{B} = \{2, \dots, r\}$. Let $Z_1^* \in \mathbb{R}$ be the solution of the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem. Then,*

$$Z_1^* = \left(\frac{\gamma_1 \alpha_2 + \sum_{i=3}^r \gamma_r \alpha_r}{\gamma_1 + \sum_{i=3}^r \gamma_r} \right) t + \left(\frac{\gamma_1 \beta_2 + \sum_{i=3}^r \gamma_r \beta_r}{\gamma_1 + \sum_{i=3}^r \gamma_r} \right). \quad (13)$$

In particular, to find the rate at which Z_1^ changes with t , we only need to know α and γ , not β .*

Lemma 3.16 (Reduction). Consider the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem such that $j, \bar{j} \in \mathcal{V} \setminus \mathcal{B}$, and such that j has all its children $1, \dots, r \in \mathcal{B}$. Let Z^* be its solution. Consider the $(\tilde{T}, \tilde{\mathcal{B}}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ -problem, where $\tilde{T} = (\tilde{r}, \tilde{\mathcal{V}}, \tilde{\mathcal{E}})$ is equal to T with nodes $1, \dots, r$ removed, and $\tilde{\mathcal{B}} = (\mathcal{B} \setminus \{1, \dots, r\}) \cup \{j\}$. Let \tilde{Z}^* be its solution. If $(\tilde{\alpha}_i, \tilde{\beta}_i, \tilde{\gamma}_i) = (\alpha_i, \beta_i, \gamma_i)$ for all $i \in \mathcal{B} \setminus \{1, \dots, r\}$, and $\tilde{\alpha}_j, \tilde{\beta}_j$ and $\tilde{\gamma}_j$ satisfy

$$\tilde{\alpha}_j = \frac{\sum_{i=1}^r \gamma_i \alpha_i}{\sum_{i=1}^r \gamma_i}, \quad \tilde{\beta}_j = \frac{\sum_{i=1}^r \gamma_i \beta_i}{\sum_{i=1}^r \gamma_i}, \quad \tilde{\gamma}_j = \left((\gamma_j)^{-1} + \left(\sum_{i=1}^r \gamma_i \right)^{-1} \right)^{-1}, \quad (14)$$

then $Z_i^* = \tilde{Z}_i^*$ for all $i \in \mathcal{V} \setminus \{j\}$.

Lemma 3.15 and Lemma 3.16 allow us to recursively solve any $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem, and obtain for it an explicit solution of the form $Z^*(t) = c_1 t + c_2$, where c_1 and c_2 do not depend on t .

Assume that we have already repeatedly pruned T , by repeatedly invoking Lemma 3.14, such that, if i is a leaf, then $i \in \mathcal{B}$. See Figure 2-(left). First, we find some node $j \in \mathcal{V} \setminus \mathcal{B}$ such that all of its children are in \mathcal{B} . If $\bar{j} \in \mathcal{B}$, then \bar{j} must be the root, and the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem must be a star problem as in Lemma 3.15. We can use Lemma 3.15 to solve it explicitly. Alternatively, if $\bar{j} \notin \mathcal{V} \setminus \mathcal{B}$, then we invoke Lemma 3.16, and reduce the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem to a strictly smaller $(\tilde{T}, \tilde{\mathcal{B}}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ -problem, which we solve recursively. Once the $(\tilde{T}, \tilde{\mathcal{B}}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ -problem is solved, we have an explicit expression $Z_i^*(t) = c_{1i} t + c_{2i}$ for all $i \in \mathcal{V} \setminus \{j\}$, and, in particular, we have an explicit expression $Z_{\bar{j}}^*(t) = c_{1\bar{j}} t + c_{2\bar{j}}$. The only free variable of the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem to be determined is $Z_j^*(t)$. To compute $Z_j^*(t)$, we apply Lemma 3.15 to the $(\tilde{\tilde{T}}, \tilde{\tilde{\mathcal{B}}}, \tilde{\tilde{\alpha}}, \tilde{\tilde{\beta}}, \tilde{\tilde{\gamma}})$ -problem, where $\tilde{\tilde{T}}$ is a star around j , $\tilde{\tilde{\gamma}}$ are the components of γ corresponding to nodes that are neighbors of j , $\tilde{\tilde{\alpha}}$ and $\tilde{\tilde{\beta}}$ are such that $Z_i^*(t) = \tilde{\tilde{\alpha}}_i t + \tilde{\tilde{\beta}}_i$ for all i that are neighbors of j , and for which $Z_i^*(t)$ is already known, and $\tilde{\tilde{\mathcal{B}}}$ are all the neighbors of j . See Figure 2-(right).

The algorithm is compactly described in Alg. 2. It is slightly different from the description above for computational efficiency. Instead of computing $Z^*(t) = c_1 t + c_2$, we keep track only of c_1 , the rates, and we do so only for the variables in $\mathcal{V} \setminus \mathcal{B}$. The algorithm assumes that the input T has been pruned. The inputs $T, \mathcal{B}, \alpha, \beta$ and γ are passed by reference. They are modified inside the algorithm but, once *ComputeRatesRec* finishes, they keep their initial values. Throughout the execution of the algorithm, $T = (r, \mathcal{V}, \mathcal{E})$ encodes (1) a doubly-linked list where each node points to its children and its parent, which we call $T.a$, and (2) a doubly-linked list of all the nodes in $\mathcal{V} \setminus \mathcal{B}$ for which all the children are in \mathcal{B} , which we call $T.b$. In the proof of Theorem 3.17, we prove how this representation of T can be kept updated with little computational effort. The input Y , also passed by reference, starts as an uninitialized array of size q , where we will store the rates $\{Z_i^*\}$. At the end, we read Z'^* from Y .

Algorithm 2 ComputeRatesRec (input: $T = (r, \mathcal{V}, \mathcal{E}), \mathcal{B}, \alpha, \beta, \gamma, Y$)

- 1: Let j be some node in $\mathcal{V} \setminus \mathcal{B}$ whose children are in \mathcal{B} ▷ We read j from $T.b$ in $\mathcal{O}(1)$ steps
 - 2: **if** $\bar{j} \in \mathcal{B}$ **then**
 - 3: Set Y_j using (13) in Lemma 3.15 ▷ If $\bar{j} \in \mathcal{B}$, then the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem is star-shaped
 - 4: **else**
 - 5: Modify $(T, \mathcal{B}, \alpha, \beta, \gamma)$ to match $(\tilde{T}, \tilde{\mathcal{B}}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ defined by Lemma 3.16 for j in line 1
 - 6: ComputeRatesRec($T, \mathcal{B}, \alpha, \beta, \gamma, Y$) ▷ Sets $Y_i = Z_i^*$ for all $i \in \mathcal{V} \setminus \mathcal{B}$; Y_j is not yet defined
 - 7: Restore $(T, \mathcal{B}, \alpha, \beta, \gamma)$ to its original value before line 5 was executed
 - 8: Compute Y_j from (13), using for α, β, γ in (13) the values $\tilde{\tilde{\alpha}}, \tilde{\tilde{\beta}}, \tilde{\tilde{\gamma}}$, where $\tilde{\tilde{\gamma}}$ are the components of γ corresponding to nodes that are neighbors of j in T , and $\tilde{\tilde{\alpha}}$ and $\tilde{\tilde{\beta}}$ are such that $Z_i^* = \tilde{\tilde{\alpha}}_i t + \tilde{\tilde{\beta}}_i$ for all i that are neighbors of j in T , and for which Z_i^* is already known
 - 9: **end if**
-

Let q be the number of nodes of the tree T that is the input at the zeroth level of the recursion.

Theorem 3.17. Algorithm 2 correctly computes Z'^* for the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem, and it can be implemented to finish in $\mathcal{O}(q)$ steps, and to use $\mathcal{O}(q)$ memory.

The correctness of Algorithm 2 follows from Lemmas 3.14-3.16, and the explanation above. Its complexity is bounded by the total time spent on the two lines that actually compute rates during




Figure 2: Red squares represent fixed nodes, and black circles free nodes. (Left) By repeatedly invoking Lemma 3.14, we can remove nodes 2, 3, and 4 from the original problem, since their associated optimal values are equal to the optimal value for node 1. (Right) We can compute the rates for all the free nodes of a subtree recursively by applying Lemma 3.16 and Lemma 3.15. We know the linear behavior of variables associated to red squares.

the whole recursion, lines 3 and 8. All the other lines only transform the input problem into a more computable form. Lines 3 and 8 solve a star-shaped problem with at most $\text{degree}(j)$ variables, which, by inspecting (13), we know can be done in $\mathcal{O}(\text{degree}(j))$ steps. Since, j never takes the same value twice, the overall complexity is bounded by $\mathcal{O}(\sum_{j \in \mathcal{V}} \text{degree}(j)) = \mathcal{O}(|\mathcal{E}|) = \mathcal{O}(q)$. The $\mathcal{O}(q)$ bound on memory is possible because all the variables that occupy significant memory are being passed by reference, and are modified in place during the whole recursive procedure.

The following lemma shows how the recursive procedure to solve a $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem can be used to compute the rates of change of $Z^*(t)$ of a (T, \mathcal{B}) -problem. Its proof follows from the observation that the rate of change of the solution with t in (13) in Lemma 3.15 only depends on α and β , and that the reduction equations (14) in Lemma 3.16 never make α' or γ' depend on β .

Lemma 3.18 (Rates only). *Let $Z^*(t)$ be the solution of the (T, \mathcal{B}) -problem, and let $\tilde{Z}^*(t)$ be the solution of the $(T, \mathcal{B}, \mathbf{1}, 0, \mathbf{1})$ -problem. Then, $Z^*(t) = c_1 t + c_2$, and $\tilde{Z}^*(t) = c_1 t$ for some c_1 and c_2 .*

We finally present the full algorithm to compute $Z'^*(t_i)$ and $\mathcal{L}''^*(t_i)$ from T and $\mathcal{B}(t_i)$.

Algorithm 3 ComputeRates (input: T and $\mathcal{B}(t_i)$ output: $Z'^*(t_i)$ and $\mathcal{L}''(t_i)$)

- 1: $Z'^*(t_i)_j = 1$ for all $j \in \mathcal{B}(t_i)$
 - 2: **for** each (T_w, \mathcal{B}_w) -problem induced by $\mathcal{B}(t_i)$ **do**
 - 3: Set \tilde{T}_w to be T_w pruned of all leaf nodes in \mathcal{B}_w , by repeatedly evoking Lemma 3.14
 - 4: ComputeRatesRec($\tilde{T}_w, j, \mathcal{B}_w, \mathbf{1}, \mathbf{0}, \mathbf{1}, Z'^*$)
 - 5: $Z'^*(t_i)_j = \tilde{Z}'^*_j$ for all $j \in V_w \setminus \mathcal{B}$
 - 6: **end for**
 - 7: Compute $\mathcal{L}''(t_i)$ from $Z'^*(t_i)$ using Lemma 3.12
 - 8: **return** $Z'^*(t_i)$ and $\mathcal{L}''(t_i)$
-

The following theorem follows almost directly from Theorem 3.17.

Theorem 3.19. *Alg. 3 correctly computes $Z'^*(t_i)$ and $\mathcal{L}''(t_i)$ in $\mathcal{O}(q)$ steps, and uses $\mathcal{O}(q)$ memory.*

4 Reducing computation time in practice

Our numerical results are obtained for an improved version of Algorithm 1. We now explain the main idea behind this algorithm.

The bulk of the complexity of Alg. 1 comes from line 4, i.e., computing the rates $\{Z'^*(t_i)_j\}_{j \in \mathcal{V} \setminus \mathcal{B}(t_i)}$ from $\mathcal{B}(t_i)$ and T . For a fixed $j \in \mathcal{V} \setminus \mathcal{B}(t_i)$, and by Lemma 3.6, the rate $Z'^*(t_i)_j$, depends only on one particular $(T_w = (r_w, \mathcal{V}_w, \mathcal{E}_w), \mathcal{B}_w)$ -problem induced by $\mathcal{B}(t_i)$. If exactly this same problem is induced by both $\mathcal{B}(t_i)$ and $\mathcal{B}(t_{i+1})$, which happens if the new nodes that become fixed in line 7 of round i of Algorithm 1 are not in $\mathcal{V}_w \setminus \mathcal{B}_w$, then we can save computation time in round $i + 1$, by not recomputing any rates for $j \in \mathcal{V}_w \setminus \mathcal{B}_w$, and using for $Z'^*(t_{i+1})_j$ the value $Z'^*(t_i)_j$.

Furthermore, if only a few $\{Z'^*_j\}$ change from round i to round $i + 1$, then we can also save computation time in computing \mathcal{L}'' from Z'^* by subtracting from the sum in the right hand side of equation (11) the terms that depend on the previous, now changed, rates, and adding new terms that depend on the new rates.

Finally, if the rate Z_j^* does not change, then the value of $t < t_i$ at which $Z_j^*(t)$ might intersect $t - N_j$, and become fixed, given by P_j in line 5, also does not change. (Note that this is not obvious from the formula for P_r in line 5). If not all $\{P_r\}$ change from round i to round $i + 1$, we can also save computation time in computing the maximum, and maximizers, in line 7 by storing P in a maximum binary heap, and executing lines 5 and 7 by extracting all the maximal values from the top of the heap. Each time any P_r changes, the heap needs to be updated.

5 Numerical results

Our algorithm to solve (3) exactly in a finite number of steps is of interest in itself. Still, it is interesting to compare it with other algorithms. In particular, we compare the convergence rate of our algorithm with two popular methods that solve (3) iteratively: the Alternating Direction Method of Multipliers (ADMM), and the Projected Gradient Descent (PGD) method. We apply the ADMM, and the PGD, to both the primal formulation (3), and the dual formulation (4). We implemented all the algorithms in C, and derived closed-form updates for ADMM and PG, see Appendix F. We ran all algorithms on a single core of an Intel Core i5 2.5GHz processor.

Figure 5-(left) compares different algorithms for a random Galton–Watson input tree truncated to have $q = 1000$ nodes, with the number of children of each node chosen uniformly within a fixed range, and for a random input $\hat{F} \in \mathbb{R}^q$, with entries chosen i.i.d. from a normal distribution. We observe the same behavior for all random instances that was tested. We gave ADMM and PGD an advantage by optimally tuning them for each individual problem-instance tested. In contrast, our algorithm requires no tuning, which is a clear advantage. At each iteration, the error is measured as $\max_j \{|M_j - M_j^*|\}$. Our algorithm is about $74 \times$ faster than its closest competitor (PGD-primal) for 10^{-3} accuracy. In Figure 5-(right), we show the average run time of our algorithm versus the problem size, for random inputs of the same form. The scaling of our algorithm is (almost) linear, and much faster than our $\mathcal{O}(q^2 p)$, $p = 1$, theoretical bound.




Figure 3: (Left) Time that the different algorithms take to solve our problem for trees of with 1000 nodes. (Right) Average run time of our algorithm for problems of different sizes. For each size, each point is averaged over 500 random problem instances.

Finally, we use our algorithm to exactly solve (2) by computing $\mathcal{C}(U)$ for all trees and a given input \hat{F} . Exactly solving (2) is very important for biology, since several relevant phylogenetic tree inference problems deal with trees of small sizes. We use an NVIDIA QUAD P5000 GPU to compute the cost of all possible trees with q nodes in parallel, and return the tree with the smallest cost. Basically, we assign to each GPU virtual thread a unique tree, using Prufer sequences [31], and then have each thread compute the cost for its tree. For $q = 10$, we compute the cost of all 100 million trees in about 8 minutes, and for $q = 11$, we compute the cost of all 2.5 billion trees in slightly less than 2.5 hours.

Code to solve (3) using Alg. 1, with the improvements of Section 4, can be found in [32]. More results using our algorithm can be found in Appendix G.

6 Conclusions and future work

We propose a new direct algorithm that, for a given tree, computes how close the matrix of frequency of mutations per position is to satisfying the perfect phylogeny model. Our algorithm is faster than the state-of-the-art iterative methods for the same problem, even if we optimally tune them. We use the proposed algorithm to build a GPU-based phylogenetic tree inference engine for the trees of relevant biological sizes. Unlike existing algorithms, which only heuristically search a small part of the space of possible trees, our algorithm performs a complete search over all trees relatively fast. It is an open problem to find direct algorithms that can provably solve our problem in linear time on average, or even for a worst-case input.

Acknowledgement: This work was partially funded by NIH/1U01AI124302, NSF/IIS-1741129, and a NVIDIA hardware grant.

References

- [1] Richard R Hudson. Properties of a neutral allele model with intragenic recombination. *Theoretical population biology*, 23(2):183–201, 1983.
- [2] Motoo Kimura. The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations. *Genetics*, 61(4):893–903, 1969.
- [3] Mohammed El-Kebir, Layla Oesper, Hannah Acheson-Field, and Benjamin J Raphael. Reconstruction of clonal trees and tumor composition from multi-sample sequencing data. *Bioinformatics*, 31(12):i62–i70, 2015.
- [4] Mohammed El-Kebir, Gryte Satas, Layla Oesper, and Benjamin J Raphael. Multi-state perfect phylogeny mixture deconvolution and applications to cancer sequencing. *arXiv preprint arXiv:1604.02605*, 2016.
- [5] Wei Jiao, Shankar Vembu, Amit G Deshwar, Lincoln Stein, and Quaid Morris. Inferring clonal evolution of tumors from single nucleotide somatic mutations. *BMC bioinformatics*, 15(1):35, 2014.
- [6] Salem Malikic, Andrew W McPherson, Nilgun Donmez, and Cenk S Sahinalp. Clonality inference in multiple tumor samples using phylogeny. *Bioinformatics*, 31(9):1349–1356, 2015.
- [7] Victoria Popic, Raheleh Salari, Iman Hajirasouliha, Dorna Kashef-Haghighi, Robert B West, and Serafim Batzoglou. Fast and scalable inference of multi-sample cancer lineages. *Genome biology*, 16(1):91, 2015.
- [8] Gryte Satas and Benjamin J Raphael. Tumor phylogeny inference using tree-constrained importance sampling. *Bioinformatics*, 33(14):i152–i160, 2017.
- [9] Anna Schuh, Jennifer Becq, Sean Humphray, Adrian Alexa, Adam Burns, Ruth Clifford, Stephan M Feller, Russell Grocock, Shirley Henderson, Irina Khrebtukova, et al. Monitoring chronic lymphocytic leukemia progression by whole genome sequencing reveals heterogeneous clonal evolution patterns. *Blood*, 120(20):4191–4196, 2012.
- [10] David Fernández-Baca. The perfect phylogeny problem. In *Steiner Trees in Industry*, pages 203–234. Springer, 2001.
- [11] Dan Gusfield. Efficient algorithms for inferring evolutionary trees. *Networks*, 21(1):19–28, 1991.
- [12] Zhihong Ding, Vladimir Filkov, and Dan Gusfield. A linear-time algorithm for the perfect phylogeny haplotyping (pvh) problem. *Journal of Computational Biology*, 13(2):522–553, 2006.
- [13] Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.
- [14] Amit G Deshwar, Shankar Vembu, Christina K Yung, Gun Ho Jang, Lincoln Stein, and Quaid Morris. PhyloWGS: reconstructing subclonal composition and evolution from whole-genome sequencing of tumors. *Genome biology*, 16(1):35, 2015.
- [15] Zoubin Ghahramani, Michael I Jordan, and Ryan P Adams. Tree-structured stick breaking for hierarchical data. In *Advances in neural information processing systems*, pages 19–27, 2010.
- [16] Laurent Condat. Fast projection onto the simplex and the ℓ_1 ball. *Mathematical Programming*, 158(1-2):575–585, 2016.
- [17] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l_1 -ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279. ACM, 2008.
- [18] Pinghua Gong, Kun Gai, and Changshui Zhang. Efficient euclidean projections via piecewise root finding and its application in gradient projection. *Neurocomputing*, 74(17):2754–2766, 2011.

- [19] Jun Liu and Jieping Ye. Efficient euclidean projections in linear time. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 657–664. ACM, 2009.
- [20] Christian Michelot. A finite algorithm for finding the projection of a point onto the canonical simplex of \mathbb{R}^n . *Journal of Optimization Theory and Applications*, 50(1):195–200, 1986.
- [21] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [22] Guilherme França and José Bento. An explicit rate bound for over-relaxed admm. In *Information Theory (ISIT), 2016 IEEE International Symposium on*, pages 2104–2108. IEEE, 2016.
- [23] Ning Hao, AmirReza Oghbaee, Mohammad Rostami, Nate Derbinsky, and José Bento. Testing fine-grained parallelism for the admm on a factor-graph. In *Parallel and Distributed Processing Symposium Workshops, 2016 IEEE International*, pages 835–844. IEEE, 2016.
- [24] Guilherme França and José Bento. How is distributed admm affected by network topology? *arXiv preprint arXiv:1710.00889*, 2017.
- [25] Laurence Yang, José Bento, Jean-Christophe Lachance, and Bernhard Palsson. Genome-scale estimation of cellular objectives. *arXiv preprint arXiv:1807.04245*, 2018.
- [26] Charles JM Mathy, Felix Gonda, Dan Schmidt, Nate Derbinsky, Alexander A Alemi, José Bento, Francesco M Delle Fave, and Jonathan S Yedidia. Sparta: Fast global planning of collision-avoiding robot trajectories.
- [27] Daniel Zoran, Dilip Krishnan, Jose Bento, and Bill Freeman. Shape and illumination from shading using the generic viewpoint assumption. In *Advances in Neural Information Processing Systems*, pages 226–234, 2014.
- [28] José Bento, Nate Derbinsky, Charles Mathy, and Jonathan S Yedidia. Proximal operators for multi-agent path planning. In *AAAI*, pages 3657–3663, 2015.
- [29] José Bento, Nate Derbinsky, Javier Alonso-Mora, and Jonathan S Yedidia. A message-passing algorithm for multi-agent trajectory planning. In *Advances in neural information processing systems*, pages 521–529, 2013.
- [30] Jean-Jacques Moreau. Décomposition orthogonale d’un espace hilbertien selon deux cônes mutuellement polaires. *CR Acad. Sci. Paris*, 225:238–240, 1962.
- [31] H Prufer. Neuer beweis eines satzes uber per mutationen. *Archiv derMathematik und Physik*, 27:742–744, 1918.
- [32] Github repository for the PPM projection algorithm, howpublished = <https://github.com/bentoayr/efficient-projection-onto-the-perfect-phylogeny-model>, note = Accessed: 2018-10-26.
- [33] Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- [34] Yuchao Jiang, Yu Qiu, Andy J Minn, and Nancy R Zhang. Assessing intratumor heterogeneity and tracking longitudinal and spatial clonal evolutionary history by next-generation sequencing. *Proceedings of the National Academy of Sciences*, 113(37):E5528–E5537, 2016.
- [35] Mohammed El-Kebir, Gryte Satas, Layla Oesper, and Benjamin J Raphael. Inferring the mutational history of a tumor using multi-state perfect phylogeny mixtures. *Cell systems*, 3(1):43–53, 2016.
- [36] Iman Hajirasouliha, Ahmad Mahmoody, and Benjamin J Raphael. A combinatorial approach for analyzing intra-tumor heterogeneity from high-throughput sequencing data. *Bioinformatics*, 30(12):i78–i86, 2014.

- [37] Paola Bonizzoni, Anna Paola Carrieri, Gianluca Della Vedova, Riccardo Dondi, and Teresa M Przytycka. When and how the perfect phylogeny model explains evolution. In *Discrete and Topological Models in Molecular Biology*, pages 67–83. Springer, 2014.
- [38] Ancestree data used, howpublished = https://github.com/raphael-group/ancestree/tree/master/data/simulated/cov_1000_samples_4_mut_100_clone_10_pcr_removed, note = Accessed: 2018-10-26.

Appendix for “Efficient Projection onto the Perfect Phylogeny Model”

A Further illustrations




Figure 4: Four subtrees of T induced by $\mathcal{B}(t)$, represented by the red squares. The root of T_1, T_2 and T_4 is node 1. The root of T_3 is node 4. All subtrees must have nodes associated to free variables (free nodes). Any subtree is uniquely identified by any free node in it. Within each subtree, any fixed node must be the root or a leaf.

B Proof of Theorem 3.1 in Section 3

We prove Theorem 3.1, by first proving the following very similar theorem.

Theorem B.1. *Problem (3) can be solved by solving*

$$\min_t t + \mathcal{L}(t), \quad (15)$$

$$\mathcal{L}(t) = \min_{Z \in \mathbb{R}^q} \frac{1}{2} \|(U^\top)^{-1} Z\|^2 \text{ subject to } Z + N \leq t\mathbf{1}, \quad (16)$$

where $N = U^\top \hat{F}$. In particular, if t^* minimizes (15), Z^* minimizes (16) for $t = t^*$, and M^*, F^* minimize (3), then

$$M^* = -U^{-1}(U^{-1})^\top Z^*, \quad F^* = -(U^{-1})^\top Z^*. \quad (17)$$

Furthermore, t^*, M^*, F^* and Z^* are unique.

Proof of Theorem B.1. Problem (3) depends on the tree T through the matrix of the ancestors, U . To see how Theorem B.1 implies Theorem 3.1, it is convenient to make this dependency more explicit. Any tree in \mathcal{T} , can be represented through a binary matrix T , where $T_{ij} = 1$ if and only if node i is the closest ancestor of node j . Henceforth, let \mathcal{T} denote the set of all such binary matrices. We need the following lemma, which we prove later in this section of the appendix.

Lemma B.2. *Consider an evolutionary tree and its matrices $T \in \mathcal{T}$ and $U \in \mathcal{U}$. We have*

$$U = (I - T)^{-1}. \quad (18)$$

Eq. (18) implies that $((U^{-1})^\top Z)_i = (Z - T^\top Z)_i = Z_i - Z_{\bar{i}}$, and that $U^{-1}((U^{-1})^\top Z)_i = Z_i - Z_{\bar{i}} - \sum_{r \in \partial i} (Z_r - Z_{\bar{r}})$, where ∂i denotes the children of i in T , \bar{i} represents the closest ancestor of i in T . We assume by convention that $Z_{\bar{i}} = 0$ when $i = r$ is the root of T . Furthermore, the definition of U implies that $N_i = (U^\top \hat{F})_i = \sum_{j \in \Delta i} \hat{F}_j$, where Δi denotes the ancestors of j . Thus,

$$\mathcal{L}(t) = \min_{Z \in \mathbb{R}^q} \frac{1}{2} \sum_{i \in \mathcal{V}} (Z_i - Z_{\bar{i}})^2 \text{ subject to} \quad (19)$$

$$Z_i \leq t - \sum_{j \in \Delta i} \hat{F}_j, \forall i \in \mathcal{V},$$

$$M_i^* = -Z_i^* + Z_{\bar{i}}^* + \sum_{r \in \partial i} (Z_r^* - Z_{\bar{r}}^*) \text{ and} \quad (20)$$

$$F_i^* = -Z_i^* + Z_{\bar{i}}^*, \forall i \in \mathcal{V}.$$

□

Proof of Theorem 3.1. Our proof is based on Moreau's decomposition [33]. Before we proceed with the proof, let us introduce a few concepts.

Given a convex, closed and proper function $g : \mathbb{R}^q \mapsto \mathbb{R}$, we define its proximal operator by the map $G : \mathbb{R}^q \mapsto \mathbb{R}^q$ such that

$$G(n) = \arg \min_{x \in \mathbb{R}^q} g(x) + \frac{1}{2} \|x - n\|^2, \quad (21)$$

where in our case $\|\cdot\|$ is the Euclidean norm. We define the Fenchel dual of g as

$$g^*(x) = \sup_{s \in \mathbb{R}^q} \{x^\top s - g(s)\}, \quad (22)$$

and we denote the proximal operator of g^* by G^* . Note that G^* can be computed from definition (21) by replacing g by g^* .

Moreau's decomposition identity states that

$$G(n) + G^*(n) = n. \quad (23)$$

We can now start the proof. Consider the following indicator function

$$g(\tilde{M}) = \begin{cases} 0, & \text{if } (U^{-1}\tilde{M}) \geq 0 \text{ and } \mathbf{1}^\top(U^{-1}\tilde{M}) = 1, \\ +\infty, & \text{otherwise,} \end{cases} \quad (24)$$

where $\tilde{M} \in \mathbb{R}^q$, and consider its associated proximal operator G . Solving problem (3), i.e., finding a minimizer M^* , is equivalent to evaluating $U^{-1}G(\hat{F})$. Using Moreau's decomposition, we have

$$M^* = U^{-1}G(\hat{F}) = U^{-1}\hat{F} - U^{-1}G^*(\hat{F}). \quad (25)$$

We will show that $G^*(\hat{F}) = \hat{F} + (U^{-1})^\top Z^*$, where Z^* is a minimizer of (5), which proves (6) and essentially completes the proof.

To compute G^* , we first need to compute

$$g^*(Y) = \sup_M \{Y^\top \tilde{M} - g(\tilde{M})\} \quad (26)$$

$$= \max_M Y^\top \tilde{M} \quad (27)$$

$$\text{subject to } U^{-1}\tilde{M} \geq 0, \mathbf{1}^\top(U^{-1}\tilde{M}) = 1.$$

Making the change of variable $M = U^{-1}\tilde{M}$, the maximum in problem (27) can be re-written as

$$\max_M (U^\top Y)^\top M \quad (28)$$

$$\text{subject to } M \geq 0, \mathbf{1}^\top M = 1.$$

It is immediate to see that the maximum in (28) is achieved if we set all components of M equal to zero except the one corresponding to the largest component of the vector $U^\top Y$, which we should set to one. Therefore, we have

$$g^*(Y) = \max_i (U^\top Y)_i. \quad (29)$$

Now we can write

$$G^*(\hat{F}) = \arg \min_{Y \in \mathbb{R}^q} g^*(Y) + \frac{1}{2} \|Y - \hat{F}\|^2 \quad (30)$$

$$= \arg \min_{Y \in \mathbb{R}^q, t \in \mathbb{R}} t + \frac{1}{2} \|Y - \hat{F}\|^2 \quad (31)$$

$$\text{subject to } U^\top Y \leq t.$$

Making the change of variable $Z = U^\top(Y - \hat{F})$, we can write $G^*(\hat{F})$ as

$$G^*(\hat{F}) = \hat{F} + (U^{-1})^\top Z^*, \text{ where} \quad (32)$$

$$(Z^*, t^*) = \arg \min_{Z \in \mathbb{R}^q, t \in \mathbb{R}} t + \frac{1}{2} \|(U^{-1})^\top Z\|^2 \quad (33)$$

subject to $Z + U^\top \hat{F} \leq t$.

To see that M^* and F^* are unique, notice that problem (3) is a projection onto a convex set polytope, which always has a unique minimizer. Moureau's decomposition implies that $G^*(\hat{F})$ is unique, hence the minimizer Y^* of (30) is unique. Thus, $Z^* = U^\top(Y^* - \hat{F})$ and $t^* = g^*(Y^*)$ are also unique. \square

Proof of Lemma B.2. We assume that the tree has q nodes. The matrix T is such that $T_{v,v'} = 1$ if and only if v is the closet ancestor of v' . Because of this, the v th column of T^k has a one in row v' if and only if v' is an ancestor v separated by k generations. Thus, the v th column of $I + T + T^2 + \dots + T^{q-1}$, contains a one in all the rows v' such that v' is an ancestor of v , or if $v = v'$. But this is the definition of the matrix U associated to the tree T . Since no two mutants can be separated by more than $q - 1$ generations, $T^k = 0$ for all $k \geq q$. It follows that

$$U = I + T + T^2 + \dots + T^{q-1} = \sum_{i=0}^{\infty} T^i = (I - T)^{-1}.$$

\square

C Proof of useful observations in Section 3.1

Proof of Lemma 3.2. The proof follows from the following generic fact, which we prove first. Let $g(W) = \min_{Z \in \mathbb{R}^q} f(Z, W)$. If f is convex in (Z, W) , then g is convex.

Indeed, let $\alpha \geq 0$ and $\alpha' = 1 - \alpha$. We get $\alpha g(W_1) + \alpha' g(W_2) = \min_{Z_1, Z_2} \alpha f(Z_1, W_1) + \alpha' f(Z_2, W_2) \geq \min_{Z_1, Z_2} f(\alpha Z_1 + \alpha' Z_2, \alpha W_1 + \alpha' W_2) = g(\alpha W_1 + \alpha' W_2)$.

To apply this result to our problem, let $f_1(Z)$ be the objective of (19) and let $f_2(Z, t, N)$ be a function (on the extended reals) such that $f_2 = 0$ if (Z, t, N) satisfy the constraints in (19) and $+\infty$ otherwise. Now notice that $\mathcal{L}(t) = \min_Z f_1(Z) + f_2(Z, t, N)$, where $f_1 + f_2$ is convex in (Z, t, N) , since both f_1 and f_2 are convex in (Z, t, N) . Convexity implies that \mathcal{L} is continuous in N and t . It also implies that $\mathcal{L}'(t)$ is non increasing in t . \square

Proof of Lemma 3.3. Continuity of $Z^*(t)$: The objective function in (19) is convex as a function of Z and has unique minimum at $Z_i = 0, \forall i$. Hence, it is strictly convex. Due to strict convexity, if the objective takes values in a small interval, then Z must be inside some small ball.

Since we know, by the remark following Lemma 3.2, that \mathcal{L} is continuous as a function of t , if t and t' are close, then $\mathcal{L}(t)$ and $\mathcal{L}(t')$ must be close. Strict convexity then implies that $Z^*(t)$ and $Z^*(t')$ must be close. The same argument can be used to prove continuity with respect to N .

Continuity of $Z^*(t^*)$: Recall that $Z^*(t^*) = Z^*$, the solution of (3). Z^* is a continuous function of M^* , which is the solution to (3), and thus is fully determined by U and \hat{F} . Since, $\hat{F} = (U^\top)^{-1}N$, \hat{F} is a continuous function of N , and it suffices to prove that M^* is continuous in \hat{F} . Problem (3) finds the projection of \hat{F} onto a convex polytope. Let F^* be this projection. Since F^* changes continuously with \hat{F} , $M^* = U^{-1}F^*$ also changes continuously with \hat{F} . \square

Proof of Lemma 3.4. Since $Z^*(t)$ is continuous, if $Z^*(t)_i \neq t - N_i$ then $Z^*(t')_i \neq t' - N_i$ for t' in some neighborhood of t . \square

Proof of Lemma 3.6. First note that, by definition of $\mathcal{B}(t)$, we know the value of all variables in $\mathcal{B}(t)$. Hence, the unknowns in problem (19) are the variables in $\mathcal{V} \setminus \mathcal{B}(t)$, which can be partitioned into disjoint sets $\{\mathcal{V}_i \setminus \mathcal{B}(t)\}_{i=1}^k$.

Second notice that for each term in the objective (19) that involves not known variables, there is some subtree T_i that contains both of its variables. It follows that, given $\mathcal{B}(t)$, problem (19) breaks into k independent problems, the i th problem having as unknowns only the variables in $\mathcal{V}_i \setminus \mathcal{B}(t)$ and all terms in the objective where either j or \bar{j} are in $\mathcal{V}_i \setminus \mathcal{B}(t)$.

Obviously, if $j \in \mathcal{V}_w \cap \mathcal{B}(t)$, then, by definition, $Z^*(t)_j = c_1 t + c_2$, with $c_1 = 1$. To find the behavior of $Z^*(t)_j$ for $j \in \mathcal{V}_w \setminus \mathcal{B}(t)$, we need to solve 7. To solve (7), notice that the first-order optimality conditions for problem (19) imply that, if $j \in \mathcal{V} \setminus \mathcal{B}(t)$, then

$$Z_j = \frac{1}{|\partial j|} \sum_{r \in \partial j} Z_r, \quad (34)$$

where ∂j denotes the neighbors of node j . We can further write

$$\begin{aligned} Z_j &= \frac{1}{|\partial j|} \sum_{r \in \partial j \cap \mathcal{B}(t)} Z_r + \frac{1}{|\partial j|} \sum_{r \in \partial j \setminus \mathcal{B}(t)} Z_r, \\ Z_r &= \frac{1}{|\partial j|} \sum_{r \in \partial j \cap \mathcal{B}(t)} (t - N_r) + \frac{1}{|\partial j|} \sum_{r \in \partial j \setminus \mathcal{B}(t)} Z_r. \end{aligned} \quad (35)$$

It follows that $Z_j = c_1 t + c_2$, for some c_1 and c_2 that depend on T , N and \mathcal{B} . If we solve for Z_j by recursively applying (35), it is immediate to see that $c_1 \geq 0$.

To see that $c_1 \leq 1$, we study how Z_j , defined by (35), depends on t algebraically. To do so, we treat t as a variable. The study of this algebraic dependency in the proof should not be confused with t being fixed in the statement of the theorem.

Define $\rho = |\partial i \cup \mathcal{B}(t)| / |\partial j|$, and notice that

$$\max_j \{Z_j\} \leq \rho t + (1 - \rho) \max_j \{Z_j\} + C, \quad (36)$$

in which C is some constant. Recursively applying the above inequality we get

$$\max_j \{Z_j\} \leq t + C', \quad (37)$$

in which C' is some constant. This shows that no Z_j can grow with t faster than $1 \times t$ and hence $c_1 \leq 1$. \square

Proof of Lemma 3.7. Lemma 3.6 implies that, for any j , $Z_j^*(t)$ depends linearly on t . The particular linear dependency, depends on $\mathcal{B}(t)$, which is piecewise constant by Lemma 3.4. Therefore, $Z_j^*(t)$ is a continuous piecewise linear function of t . This in turn implies that $\mathcal{L}'(t)$ is a continuous piecewise linear function of t , since it is the derivative of the continuous piecewise quadratic $\mathcal{L}(t) = (1/2) \sum_{i \in \mathcal{V}} (Z^*(t)_i - Z^*(t)_{\bar{i}})^2$. Finally, since the particular linear dependency of Z^* , depends on $\mathcal{B}(t)$, it follows that $Z^*(t)$ and $\mathcal{L}'(t)$ change linear segment if and only if $\mathcal{B}(t)$ changes. \square

Proof of Lemma 3.8. Let us assume that there exists $t < t'$ for which $\mathcal{B}(t) \subset \mathcal{B}(t')$. We can assume without loss of generality that t is sufficiently close to t' such that $\mathcal{B}(s)$ is constant for $s \in [t, t']$. Let j be such that $j \in \mathcal{B}(t')$ but $j \notin \mathcal{B}(t)$. This means that $Z_j^*(s) < s - N_j$ for all $s \in [t, t']$ and that $Z_j^*(t') = t' - N_j$. Since by Lemma 3.6, $Z_j^*(s) = c_1 s + c_2$, for some constants c_1 and c_2 , the only way that $Z_j^*(s)$ can intersect $s - N_j$ at $s = t'$ is for $c_1 > 1$, which is a contradiction.

If $\mathcal{B}(t)$ decreases as t increase, and given that the largest that $\mathcal{B}(t)$ can be is $\{1, \dots, q\}$, it follows that $\mathcal{B}(t)$ can only take $q + 1$ different configurations. One configuration per size of $\mathcal{B}(t)$, from q to 0. \square

Proof of Lemma 3.9. Lemma 3.8 implies that $\mathcal{B}(t)$ changes at most $q + 1$ times. Lemma 3.7 then implies that $Z^*(t)$ and $\mathcal{L}'(t)$ have less than $q + 1$ different linear segments. \square

D Proofs of the properties of the algorithm in Section 3.2

Proof of Theorem 3.10. Run-time: Recall that $Z^*, Z'^* \in \mathbb{R}^q$ and that $\mathcal{L}' \in \mathbb{R}$. Line 1 is done in $\mathcal{O}(q)$ steps by doing a DFS on T . Here, we assume that T is represented as a linked list. Specifically,

starting from the root, we keep a variable x where we accumulate the values of \hat{F}_j visited from the root to the current node being explored in T as we move down the tree. As we move up the tree, we subtract values of the nodes \hat{F}_j from x . Then, at each node i visited by the DFS, we can read from x the value N_i . Line 2 takes $\mathcal{O}(q)$ steps to finish. The procedure `ComputeRates` takes $\mathcal{O}(q)$ steps to finish, which we prove in Theorem 3.19. All of the other lines inside the for-loop are manipulations that take at most $\mathcal{O}(q)$ steps. Lines 13 and 12 take $\mathcal{O}(q)$ steps. From (6), the complexity to compute F^* is $\mathcal{O}(q)$, and the complexity to compute M^* is $\mathcal{O}(\sum_{i \in \mathcal{V}} |\partial i|) = \mathcal{O}(|\mathcal{E}|) = \mathcal{O}(q)$.

Memory: The DFS in line 1 only requires $\mathcal{O}(q)$ memory. Throughout the algorithm, we only need to keep the two most recent values of t_i , $\mathcal{B}(t_i)$, $Z^*(t_i)$, $Z'^*(t_i)$, $\mathcal{L}'(t_i)$ and $\mathcal{L}''(t_i)$. This takes $\mathcal{O}(q)$ memory. The procedure `ComputeRates` takes $\mathcal{O}(q)$ memory, which we prove in Theorem 3.19. \square

Proof of Theorem 3.11. The proof of Theorem 3.11 amounts to checking that, at every step of Algorithm 1, the quantities computed, e.g., the paths $\{Z^*(t)\}$ and $\{\mathcal{L}'(t)\}$, are correct.

Lemmas 3.7 and 3.9 prove that $Z^*(t)$ and $\mathcal{L}'(t)$ are piecewise linear and continuous with at most q changes in linear segment. Hence, the paths $\{Z^*(t)\}$ and $\{\mathcal{L}'(t)\}$ are fully specified by their value at $\{t_i\}_{i=1}^k$, and $k \leq q$.

Lemma 3.7 proves that these critical values are determined as the instants, at which $\mathcal{B}(t)$ changes. Furthermore, Lemma 3.8 proves that, as t decreases, variables are only added to $\mathcal{B}(t)$. Hence, to find $\{t_i\}$ and $\{\mathcal{B}(t_i)\}$, we only need to find the times and components at which, as t decreases, $Z^*(t)_r$ goes from $Z^*(t)_r < t - N_r$ to $Z^*(t)_r = t - N_r$. Also, since \mathcal{B} can have at most q variables, the for-loop in line 3 being bounded to the range $1-q$, does not prevent the algorithm from finding any critical value.

Theorem 3.19 tells us that we can compute $Z'^*(t_i)$ from $\mathcal{B}(t_i)$ and T . Since we have already proved that the path $\{Z^*(t)\}$ is piecewise linear and continuous, we can compute t_{i+1} , and the variables that become fixed, by solving (10) for t for each $r \notin \mathcal{B}(t_i)$, and choosing for t_{i+1} the largest such t , and choosing for the new fixed variables, i.e., $\mathcal{B}(t_{i+1}) - \mathcal{B}(t_i)$, the components r for which the solution of (10) is t_{i+1} .

Since we have already proved that $Z^*(t)$ and $\mathcal{L}'(t)$ are piecewise linear and constant, we can compute $Z^*(t_{i+1})$ and $\mathcal{L}'(t_{i+1})$ from $Z^*(t_i)$, $\mathcal{L}'(t_i)$, $Z'^*(t_i)$ and $\mathcal{L}''(t_i)$ using (9).

Lemma 3.2 proves that $\mathcal{L}'(t)$ decreases with t , and Theorem 3.1 proves that t^* is unique. Hence, as t decreases, there is a single t at which $\mathcal{L}'(t)$ goes from > -1 to < -1 . Since we have already proved that we correctly, and sequentially, compute $\mathcal{L}'(t_i)$, $\mathcal{L}''(t_i)$, and that $\mathcal{L}'(t)$ is piecewise linear and constant, we can stop computing critical values whenever we can determine that $\mathcal{L}'(t) = \mathcal{L}'(t_k) + (t - t_k)\mathcal{L}''(t_k)$ will cross the value -1 , where t_k is the latest computed critical value. This is the case when $\mathcal{L}'(t_k) > -1$ and $\mathcal{L}'(t_{k+1}) < -1$, or when $\mathcal{L}'(t_k) > -1$ and t_k is the last possible critical value, which happens when $|\mathcal{B}(t_i)| = q$. From this last critical value, t_k , we can then find t^* and Z^* by solving $-1 = \mathcal{L}'(t_k) + (t^* - t_k)\mathcal{L}''(t_k)$ and $Z^* = Z^*(t_k) + (t^* - t_k)Z'^*(t_k)$. Finally, once we have Z^* , we can use (6) in Theorem 3.1 to find M^* and F^* . \square

E Proofs for computing the rates in Section 3.3

Proof of Lemma 3.12. Let $t \in (t_{i+1}, t_i)$. We have,

$$\begin{aligned} \mathcal{L}'(t) &= \frac{d}{dt} \frac{1}{2} \sum_{j \in \mathcal{V}} (Z^*(t)_j - Z^*(t)_{\bar{j}})^2 = \\ &= \sum_{i \in \mathcal{V}} (Z^*(t)_j - Z^*(t)_{\bar{j}})(Z'^*(t)_j - Z'^*(t)_{\bar{j}}). \end{aligned} \quad (38)$$

Taking another derivative, and recalling that $Z''^*(t) = 0$ for $t \in (t_{i+1}, t_i)$, we get

$$\mathcal{L}''(t) = \sum_{j \in \mathcal{V}} (Z'^*(t)_j - Z'^*(t)_{\bar{j}})^2, \quad (39)$$

and the lemma follows by taking the limit $t \uparrow t_i$. \square

Proof of Lemma 3.14. The $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem is unconstrained and convex, hence we can solve it by taking derivatives of the objective with respect to the free variables, and setting them to zero. Let us call the objective function $F(Z)$. If $j \in \mathcal{V} \setminus \mathcal{B}$ is a leaf, then $\frac{dF}{dZ_j} = 0$ implies that $Z_j^* = Z_j^*$. We now prove the second part of the lemma. Let $\tilde{F}(Z)$ be the objective of the modified problem. Clearly, $\frac{dF}{dZ_i} = \frac{d\tilde{F}}{dZ_i}$ for all $i \in \tilde{T} \setminus \tilde{j}$. Let C be the children of \tilde{j} in T and \tilde{C} be the children of \tilde{j} in \tilde{T} . We have $\tilde{C} = C \setminus j$. Furthermore, $\frac{d\tilde{F}}{dZ_j} = 0$ is equivalent to $\gamma_j(Z_j - Z_{\tilde{j}}) + \sum_{s \in \tilde{C}} \gamma_s(Z_j - Z_s) = 0$, and $\frac{dF}{dZ_j} = 0$ is equivalent to $\gamma_j(Z_j - Z_{\tilde{j}}) + \sum_{s \in C} \gamma_s(Z_j - Z_s) = 0$. However, we have already proved that the optimal solution for the original problem has $Z_j^* = Z_j^*$. Hence, this condition can be replaced in $\frac{dF}{dZ_j}$, which becomes $\gamma_j(Z_j - Z_{\tilde{j}}) + \sum_{s \in \tilde{C}} \gamma_s(Z_j - Z_s) = 0$. Therefore, the two problems have the same optimality conditions, which implies that $Z_i^* = \tilde{Z}_i^*$, for all $i \in \tilde{\mathcal{V}}$. \square

Proof of Lemma 3.15. The proof follows directly from the first order optimality conditions, a linear equation that we solve for Z_1^* . \square

Proof of Lemma 3.16. The first order optimality conditions for both problems are a system of linear equations, one equation per free node in each problem. All the equations associated to the ancestral nodes of j are the same for both problems. The equation associated to variable j in the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem is

$$\gamma_j(Z_{\tilde{j}} - Z_j) + \sum_{i=1}^r \gamma_i(Z_i - Z_j) = 0, \quad (40)$$

which implies that

$$Z_j = \frac{\gamma_j Z_{\tilde{j}} + \sum_{i=1}^r \gamma_i Z_i}{\gamma_j + \sum_{i=1}^r \gamma_i}. \quad (41)$$

The equation associated to the variable \tilde{j} in the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem is

$$F(Z, \alpha, \beta, \gamma) + \gamma_j(Z_j - Z_{\tilde{j}}) = 0, \quad (42)$$

where $F(Z)$ is a linear function of Z determined by the tree structure and parameters associated to the ancestral edges and nodes of \tilde{j} . The equation associated to the variable \tilde{j} in the $(\tilde{T}, \tilde{\mathcal{B}}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ -problem is

$$F(\tilde{Z}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}) + \tilde{\gamma}_j(\tilde{\alpha}_j t + \tilde{\beta}_j - \tilde{Z}_{\tilde{j}}) = 0, \quad (43)$$

for the same function F as in (42). Note that the components of $\tilde{\alpha}$, $\tilde{\beta}$ and $\tilde{\gamma}$ associated to the ancestral edges and nodes of \tilde{j} are the same as in α , β and γ . Hence, $F(\tilde{Z}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}) = F(\tilde{Z}, \alpha, \beta, \gamma)$.

By replacing (41) into (42), one can easily check the following. Equations (42) and (43), as linear equations on Z and \tilde{Z} respectively, have the same coefficients if (14) holds. Hence, if (14) holds, the solution to the linear system associated to the optimality conditions in both problem gives the same optimal value for all variables ancestral to \tilde{j} and including \tilde{j} . \square

Proof of Theorem 3.17. Although T changes during the execution of the algorithm, in the proof we let $T = (r, \mathcal{V}, \mathcal{E})$ be the tree, passed to the algorithm at the zeroth level of the recursion. Recall that $|\mathcal{V}| = q$ and $\mathcal{E} = q - 1$.

Correctness: The correctness of the algorithm follows directly from Lemmas 3.14, 3.15, and 3.16 and the explanation following these lemmas.

Run-time: It is convenient to think of the complexity of the algorithm by assuming that it is running on a machine with a single instruction pointer that jumps from line to line in Algorithm 2. With this in mind, for example, the recursive call in line 6 simply makes the instruction pointer jump from line 6 to line 1. The run-time of the algorithm is bounded by the sum of the time spent in each line in Algorithm 2, throughout its entire execution. Each basic step costs one unit of time. Each node in \mathcal{V} is only chosen as j at most once, throughout the entire execution of the algorithm. Hence, line 1 is executed at most q times, and thus any line is executed at most q times, at most once for each possible choice for j .

Assuming that we have $T.b$ updated, j in line 1 can be executed in $\mathcal{O}(1)$ time, by reading the first element of the linked list $T.b$. Lines 2 and 6 also take $\mathcal{O}(1)$ time. Here, we are thinking of the cost of line 6 as simply the cost to make the instruction pointer jump from line 1 to line 6, not the cost to fully completing the call to *ComputeRatesRec* on the modified problem. The modification made to the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem by lines 5 and 7, is related to the addition, or removal, of at most $\text{degree}(j)$ nodes, where $\text{degree}(j)$ is the degree of j in T . Hence, they can be executed in $\mathcal{O}(\text{degree}(j))$ steps. Finally, lines 3 and 8 require solving a star-shaped problem with $\mathcal{O}(\text{degree}(j))$ variables, and thus take $\mathcal{O}(\text{degree}(j))$, which can be observed by inspecting (13).

Therefore, the run-time of the algorithm is bounded by $\mathcal{O}(\sum_j \text{degree}(j)) = \mathcal{O}(q)$.

To see that it is not expensive to keep T updated, notice that, if T changes, then either $T.b$ loses j (line 5) or has j reinserted (line 7), both of which can be done in $\mathcal{O}(1)$ steps. Hence, we can keep $T.b$ updated with only $\mathcal{O}(1)$ effort each time we run line 5 and line 7. Throughout the execution of the algorithm, the tree T either shrinks by losing nodes that are children of the same parent (line 5), or T grows by regaining nodes that are all siblings (line 7). Hence, the linked list $T.a$ can be kept updated with only $\mathcal{O}(1)$ effort each time we run line 5 and line 7. Across the whole execution of the algorithm, $T.a$ and $T.b$ can be kept updated with $\mathcal{O}(\sum_j \text{degree}(j)) = \mathcal{O}(q)$ effort.

Memory: All the variables with a size that depend on q are passed by reference in each call of *ComputeRatesRec*, namely, $Y, T, \mathcal{B}, \alpha, \beta$ and γ . Hence, we only need to allocate memory for them once, at the zeroth level of the recursion. All these variables take $\mathcal{O}(q)$ memory to store. \square

Proof of Lemma 3.18. From Definition 3.13, we know that the $(T, \mathcal{B}, \mathbf{1}, -N, \mathbf{1})$ -problem and the (T, \mathcal{B}) -problem are the same. Hence, it is enough to prove that the solutions of (i) any $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem and of (ii) the $(T, \mathcal{B}, \alpha, 0, \gamma)$ -problem change at the same rate as a function of t .

We have already seen that the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem can be solved by recursively invoking Lemma 3.16 until we arrive at problems that are small enough to be solved via Lemma 3.15.

We now make two observations. First, while recursing, Lemma 3.16 always transform a $(\tilde{T}, \tilde{\mathcal{B}}, \tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ -problem into a smaller problem $(\tilde{\tilde{T}}, \tilde{\tilde{\mathcal{B}}}, \tilde{\tilde{\alpha}}, \tilde{\tilde{\beta}}, \tilde{\tilde{\gamma}})$ -problem where, by (14), $\tilde{\tilde{\gamma}}$ and $\tilde{\tilde{\alpha}}$ only depend on $\tilde{\alpha}$ and $\tilde{\gamma}$ but not on $\tilde{\beta}$.

Second, while recursing, and each time Lemma 3.15 is invoked to compute an explicit value for some component of the solution via solving some star-shaped $(\tilde{\tilde{T}}, \tilde{\tilde{\mathcal{B}}}, \tilde{\tilde{\alpha}}, \tilde{\tilde{\beta}}, \tilde{\tilde{\gamma}})$ -problem, the rate of change of this component with t , is a function of $\tilde{\tilde{\alpha}}$ and $\tilde{\tilde{\gamma}}$ only. We can see this from (13).

Hence, the rate of change with t of the solution of the $(T, \mathcal{B}, \alpha, \beta, \gamma)$ -problem does not depend on β . So we can assume $\beta = 0$. \square

Proof of Theorem 3.19. Correctness: The correctness of Algorithm 3 follows from the correctness of Algorithm 2.

Run-time and memory: We can prune each T_w in $\mathcal{O}(|T_w|)$ steps and $\mathcal{O}(1)$ memory using DFS. In particular, once we reach a leaf of T_w that is free, i.e., not in \mathcal{B}_w , and as DFS travels back up the tree, we can prune from T_w all the nodes that are free. By Theorem 3.17, the number of steps and memory needed to completely finish line 4 is $\mathcal{O}(|T_w|)$. The same is true to complete line 5. Hence, the number of steps and memory required to execute the for-loop is $\mathcal{O}(\sum_w |T_w|) = \mathcal{O}(|T|) = \mathcal{O}(q)$. Finally, by Theorem 3.12, \mathcal{L}'' can be computed from Z'^* in $\mathcal{O}(q)$ steps using $\mathcal{O}(1)$ memory. \square

F Details of the ADMM and the PGD algorithms in Section 5

Here we explain the details of our implementations of the Alternating Direction Method of Multipliers (ADMM) and the Projected Gradient Descent (PGD) methods, applied to our problem.

F.1 ADMM

F.1.1 ADMM for the primal problem

We start by putting our initial optimization problem (3) into the following equivalent form:

$$\min_{M \in \mathbb{R}^q} \{f(M) = \frac{1}{2}\|F - UM\|^2\} + g(M), \quad (44)$$

where $g(M)$ is the indicator function imposing the constraints on M :

$$g(M) := \begin{cases} 0, & M \geq 0, M^\top \mathbf{1} = \mathbf{1}, \\ +\infty, & \text{otherwise.} \end{cases} \quad (45)$$

In this formulation, our target function is a sum of two terms. We now proceed with the standard ADMM procedure, utilizing the splitting f, g . Our ADMM scheme iterates on the following variables $M, M_1, M_2, u_1, u_2, \in \mathbb{R}^q$. M_1 and M_2 are primal variables, M is a consensus variable, and u_1 and u_2 are dual variables. It has tuning parameters $\alpha, \rho \in \mathbb{R}$.

First, we evaluate the proximal map associated with the first term

$$M_1 \leftarrow \arg \min_{S \in \mathbb{R}^q} \frac{1}{2}\|F - US\|^2 + \frac{\rho}{2}\|S - M + u_1\|^2, \quad (46)$$

where S is a dummy variable. This map can be evaluated in closed form,

$$M_1 = (\rho I + U^\top U)^{-1}(\rho M - \rho u_1 + U^\top F). \quad (47)$$

Second, we evaluate the proximal map associated with the second term

$$M_2 \leftarrow \arg \min_{S \in \mathbb{R}^q} g(S) + \frac{\rho}{2}\|S - M + u_2\|^2, \quad (48)$$

where S is again a dummy variable. This map is precisely the projection onto the simplex, which has been extensively studied in the literature; there are many fast algorithms that solve this problem exactly. We implemented the algorithm proposed in [16].

Lastly, we perform the rest of the standard ADMM updates:

$$\begin{aligned} M &\leftarrow \frac{1}{2}(M_1 + u_1 + M_2 + u_2), \\ u_1 &\leftarrow u_1 + \alpha(M_1 - M), \\ u_2 &\leftarrow u_2 + \alpha(M_2 - M). \end{aligned} \quad (49)$$

We repeat the above steps until a satisfactory precision is reached, and read off the final solution from the variable M .

F.1.2 ADMM for the dual problem

We now apply ADMM to the dual problem (4). We start by incorporating the constraints into the target function to rewrite (4) as

$$\min_{Z, t} \{f(t) = t\} + \{h(Z) = \frac{1}{2}\|(U^\top)^{-1}Z\|^2\} + g(t, Z), \quad (50)$$

where

$$g(t, Z) := \begin{cases} 0, & t\mathbf{1} - Z \geq N, \\ +\infty, & \text{otherwise,} \end{cases} \quad (51)$$

is the indicator function imposing the constraints on t, Z . ADMM now splits the problem into three parts, each associated to one of the functions f, g and h .

Our ADMM scheme will iterate on the following variables $Z, X_Z, X_{gZ}, u_Z, u_{gZ} \in \mathbb{R}^q$, and $t, X_t, X_{gt}, u_t, u_{gt} \in \mathbb{R}$. The variables X_Z, X_{gZ}, X_t, X_{gt} are primal variables, t, Z are consensus variables, and u_Z, u_{gZ}, u_t, u_{gt} are dual variables. It has tuning parameters $\alpha, \rho \in \mathbb{R}$.

First, we evaluate the proximal map for the first term

$$X_Z \leftarrow \arg \min_{S \in \mathbb{R}^q} \frac{1}{2} \|(U^\top)^{-1} S\|^2 + \frac{\rho}{2} \|S - Z + u_Z\|^2, \quad (52)$$

where S is a dummy variable. This map can be evaluated using an closed form formula:

$$X_Z = (\rho I + U^{-1}(U^{-1})^\top)^{-1} \rho(Z - u_Z). \quad (53)$$

Next, we evaluate the proximal map for the second term

$$X_t \leftarrow \arg \min_{S \in \mathbb{R}} S + \frac{\rho}{2} (S - t + u_t)^2, \quad (54)$$

where S is a dummy variable. Again, this can be solved straightforwardly:

$$X_t = \frac{\rho t - \rho u_t - 1}{\rho}. \quad (55)$$

We then evaluate the proximal map for the third term, which involves the constraints

$$\begin{aligned} (X_{gZ}, X_{gt}) \leftarrow \arg \min_{S \in \mathbb{R}^q, S_t \in \mathbb{R}} g(S, S_t) + \frac{\rho}{2} \|(S, S_t) \\ - (Z - u_{gZ}, t - u_{gt})\|^2, \end{aligned} \quad (56)$$

where S, S_t are dummy variables. This problem is a projection onto the polyhedron defined by the constraints, $t\mathbf{1} - Z \geq N$, in \mathbb{R}^{q+1} . We developed an algorithm that solves this problem exactly in $\mathcal{O}(q \log q)$ steps. This is discussed in Section F.3.

What is left to be done is the following part of the ADMM:

$$\begin{aligned} Z &\leftarrow \frac{1}{2}(X_Z + u_Z + X_{gZ} + u_{gZ}), \\ u_Z &\leftarrow u_Z + \alpha(X_Z - Z), \\ u_{gZ} &\leftarrow u_{gZ} + \alpha(X_{gZ} - Z), \\ t &\leftarrow \frac{1}{2}(X_t + u_t + X_{gt} + u_{gt}), \\ u_t &\leftarrow u_t + \alpha(X_t - t), \\ u_{gt} &\leftarrow u_{gt} + \alpha(X_{gt} - t). \end{aligned} \quad (57)$$

We repeat the above steps until a satisfactory precision is reached, and read off the final solution from the variables t and Z .

F.2 PGD

F.2.1 PGD for the primal problem

Implementing PGD is rather straightforward. For the initial problem (3), we simply do the following update:

$$M \leftarrow \text{Proj-onto-Simplex}(M + \alpha U^\top (F - UM)), \quad (58)$$

where $\text{Proj-onto-Simplex}()$ refers to projection onto the simplex, for which we implemented the algorithm proposed in [16]. $\alpha \in \mathbb{R}$ is the step size, a tuning parameter. We perform this update repeatedly until a satisfactory precision is reached.

F.2.2 PGD for the dual problem

For the dual problem (4), the updates we need are

$$\begin{aligned} Z &\leftarrow Z - \alpha U^{-1}(U^{-1})^\top Z, \\ t &\leftarrow t - \alpha, \\ (Z, t) &\leftarrow \text{Proj-onto-Polyhedron}((Z, t)), \end{aligned} \quad (59)$$

where $\text{Proj-onto-Polyhedron}()$ refers to projection onto the polyhedron defined by $t\mathbf{1} - Z \geq N$ in \mathbb{R}^{q+1} , while $\alpha \in \mathbb{R}$ is the step size. This is explicitly explained in F.3. Again, we perform these updates repeatedly until a satisfactory precision is reached, and tune the parameters to achieve the best possible performance.

E3 Projection onto the polyhedron $t\mathbf{1} - Z \geq N$

We would like to solve the following optimization problem:

$$\arg \min_{Z \in \mathbb{R}^q, t \in \mathbb{R}} \frac{1}{2} \|(Z, t) - (A, B)\|^2, \quad (60)$$

$$\text{subject to } t\mathbf{1} - Z \geq N, \quad (61)$$

which is the problem of projection onto the polyhedron $t\mathbf{1} - Z \geq N$ in \mathbb{R}^{q+1} . The Lagrangian of this optimization problem is

$$\mathcal{L} = \frac{1}{2} \|(Z, t) - (A, B)\|^2 + \lambda^\top (Z + N - t\mathbf{1}), \quad (62)$$

where $\lambda \in \mathbb{R}^q$ is the Lagrange multiplier. We solve problem (60) by solving the dual problem $\max_{\lambda \geq 0} \min_{Z, t} \mathcal{L}$.

We first solve the minimization over variables Z and t . It is straightforward to find the closed form solutions:

$$Z^* = A - \lambda, \quad t^* = B + \mathbf{1}^\top \lambda. \quad (63)$$

Using these expressions, we can rewrite the Lagrangian as

$$\mathcal{L} = -\frac{1}{2} \lambda^\top (I + \mathbf{1}\mathbf{1}^\top) \lambda + R^\top \lambda, \quad (64)$$

where $R = A + N - B\mathbf{1}$.

Now our goal becomes solving the following optimization problem:

$$\arg \min \frac{1}{2} \lambda^\top (I + \mathbf{1}\mathbf{1}^\top) \lambda - R^\top \lambda, \quad (65)$$

$$\text{subject to } \lambda \geq 0. \quad (66)$$

The KKT conditions for (65) are

$$\lambda_i + \mathbf{1}^\top \lambda - R_i - s_i = 0, \quad \lambda_i \geq 0, \quad s_i \geq 0, \quad \lambda_i s_i = 0, \quad i = 1, \dots, q, \quad (67)$$

where s_i are Lagrange multipliers associated with the constraint $\lambda \geq 0$.

We proceed with sorting the vector R first, and maintain a map $f : \{1, 2, \dots, q\} \rightarrow \{1, 2, \dots, q\}$ that maps the sorted indices back to the unsorted indices of R . Let us call the sorted R by \tilde{R} . Then, from the above KKT conditions, it is straightforward to derive the following expression for λ_i :

$$\lambda_i = \begin{cases} \tilde{R}_i - \mathbf{1}^\top \lambda, & i \geq \tau, \\ 0, & i < \tau, \end{cases} \quad i = 1, 2, \dots, q \quad (68)$$

where

$$\tau = \min\{i \mid \tilde{R}_i - \mathbf{1}^\top \lambda \geq 0\}. \quad (69)$$

Then it follows that

$$\mathbf{1}^\top \lambda = \sum_{i=\tau}^q (\tilde{R}_i - \mathbf{1}^\top \lambda) = \frac{1}{2 + q - \tau} \sum_{i=\tau}^q \tilde{R}_i, \quad (70)$$

and hence we have that

$$c(\tau) := \tilde{R}_\tau - \mathbf{1}^\top \lambda = \tilde{R}_\tau - \frac{1}{2 + q - \tau} \sum_{j=\tau}^q \tilde{R}_j. \quad (71)$$

According to (69), to find τ , we only need to find the smallest value of i that makes $c(i)$ non negative. That is, $\tau = \min\{i \mid c(i) \geq 0\}$.

Therefore, by sorting the components of R from small to large, and checking $c(i)$ for each component, from large i to small i , we can obtain the desired index τ . Combining equations (68) and (70) with τ , we find a solution that equals λ^* , the solution to problem (65), apart from a permutation of its components. We then use our index map f to undo the sorting of the components introduced by sorting R .

Finally, by plugging λ^* back into equation (63), we obtain the desired solution to our problem (60). The whole projection procedure can be done in $\mathcal{O}(q \log q)$, the slowest step being the sorting of R .

G More results using our algorithm

In this section, we use our fast projection algorithm to infer phylogenetic trees from frequency of mutation data.

The idea is simple. We scan all possible trees, and, for each tree T , we project \hat{F} into a PPM for this T using our fast projection algorithm. This gives us a projected F and M such that $F = UM$, the columns of M are in the probability simplex, and $\|\hat{F} - F\|$ is small. Then, we return the tree whose projection yields the smallest $\|\hat{F} - F\|$. Since all of these projections can be done in parallel, we assign the projection for different subsets of the set of all possible trees to different GPU cores. Since we are performing an exhaustive search over all possible trees, we can only infer small trees. As such, when dealing with real-size data, similar to several existing tools, we first cluster the rows of \hat{F} , and produce an “effective” \hat{F} with a small number of rows. We infer a tree on this reduced input. Each node in our tree is thus associated with multiple mutated positions in the genome, and multiple mutants, depending on the clustering. We cluster the rows of \hat{F} using k -means, just like in [6]. We decide on the numbers of clusters, and hence tree size, based on the same BIC procedure as in [6]. It is possible that other pre-clustering, and tree-size-selection strategies, yield better results. We call the resulting tool EXACT.

We note that it is not our goal to show that the PPM is adequate to extract phylogenetic trees from data. This adequacy, and its limits, are well documented in well-cited biology papers. Indeed, several papers provide open-source tools based on the PPM, and show their tools’ good performance on data containing the frequencies of mutation per position in different samples, \hat{F} in our paper. A few tools are PhyloSub [5], AncesTree [3], CITUP [6], PhyloWGS [14], Canopy [34], SPRUCE [35], rec-BTP [36], and LICHeE [7]. These papers also discuss the limitations of the PPM regarding inferring evolutionary trees, and others propose extensions to the PPM to capture more complex phenomena, see e.g., [37].

It is important to further distinguish the focus of our paper from the focus of the papers cited in the paragraph above. In this paper, we start from the fact that the PPM is already being used to infer trees from \hat{F} , and with substantiated success. However, all of the existing methods are heuristics, leaving room for improvement. We identify one subproblem that, if solved very fast, allows us to do exact PPM-based tree inference for problems of relevant biological sizes. It is this subproblem, a projection problem in Eq. (3), that is our focus. We introduce the first non-iterative algorithm to solve this projection problem, and show that it is $74\times$ faster than different optimally-tuned iterative methods. We are also the first to show that a full-exact-enumeration approach to inferring U and M from \hat{F} is possible, in our case, using a GPU and our algorithm to compute and compare the cost of all the possible trees that might explain the data \hat{F} . EXACT often outperforms the above tools, none of which does exact inference. Our paper is not about EXACT, whose development challenges and significance for biology go beyond solving our projection problem, and which is the focus of our future work.

Despite this difference in purpose, in this section we compare the performance of inferring trees from a full exact search over the space of all possible PPM models with the performance of a few existing algorithms. In Figure 5, we compare EXACT, PhyloWGS, CITUP and AncesTree on recovering the correct ancestry relations on biological datasets also used by [3]. A total of 30 different datasets [38], i.e., \hat{F} , were tested. We use the default parameters in all of the algorithms tested.




Figure 5: Comparison of different phylogenetic tree inference algorithms.

In each test, and for every pair of mutations i and j , we use the tree output by each tool to determine if (a) i is an ancestor of j or if j is an ancestor of i , (b) if i and j are in the same node, (c) if either i or j are missing in the tree, or, otherwise, (d) if i and j are incomparable. We give these four possible ancestral relations, the following names: *ancestral*, *clustered*, *missing*, and *incomparable*. A random guess correctly identifies 25% of the ancestral categories, on average. If the fraction of misidentified relations is 0, the output tree equals the ground-truth tree. All methods do better than random guesses.

For example, in Figure 6, according to EXACT, mutation 63, at the root, is an ancestor of mutation 57, at node 3. However, according to the ground truth, in Figure 6, they belong to the same node. So, as far as comparing 63 with 57 goes, EXACT makes a mistake. As another example, according to EXACT, mutations 91 and 55 are incomparable, while according to the ground truth, 91 is a descendent of 55. Hence, as far as comparing 91 with 55 goes, EXACT makes another mistake. The fraction of errors, per ancestral relation error type, that each of these tools makes is: EXACT = {23%, 10%, 0%, 13%}; PhyloWGS = {3%, 2%, 0%, 1%}; AncesTree = {54%, 16%, 95%, 25%}; CITUP = {27%, 13%, 0%, 21%}.

In our experiments, EXACT performs, on average, better than the other three methods. PhyloWGS performs close to EXACT, however, it has a much longer run time. Although AncesTree does fairly well in terms of accuracy, we observe that it often returns trees with the same topology, a star-shaped tree. The other methods, produce trees whose topology seems to be more strongly linked to the input data. Finally, AncesTree’s inferred tree does not cover all of the existing mutations. This behaviour is expected, as, by construction, AncesTree tries to find the largest tree that can be explained with the PPM. See Figure 6, and Figure 7, for an example of the output produced by different algorithms, and the corresponding ground truth.

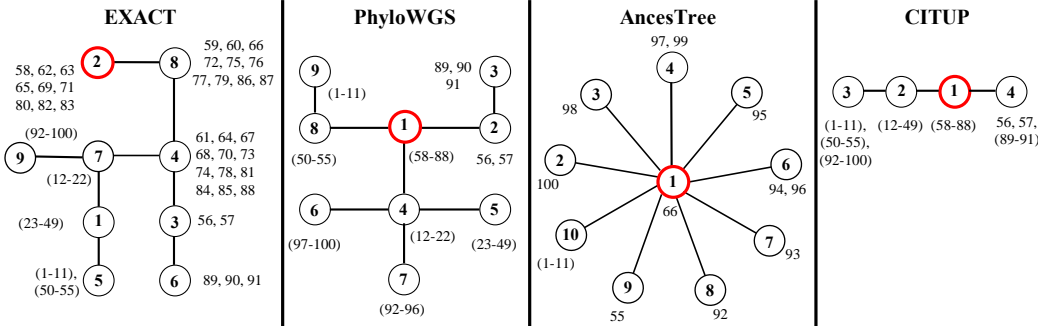


Figure 6: Tree reconstructed by different algorithms for the first file in the folder [38]. AncesTree often outputs star-shaped trees. The small numbers listed next to each node represent mutations. Mutations indexed by the same number in different trees are the same real mutation. The root of each tree is circled in thick red. Nodes are labeled by numbers, and these labels are assigned automatically by each tool. Labels of different trees are incomparable.

We end this section by discussing a few extra properties that distinguished an approach like EXACT from the existing tools. Because our algorithm’s speed allows a complete enumeration of all of the trees, EXACT has two unique properties. First, EXACT can exactly solve

$$\min_{U \in \mathcal{U}} \mathcal{J}(\mathcal{C}(U)) + \mathcal{Q}(U), \quad (72)$$

where U encodes ancestral relations, $\mathcal{C}(U)$ is the fitness cost as defined in our paper, \mathcal{J} is an arbitrary, fast-to-compute, 1D scaling function, and $\mathcal{Q}(U)$ is an arbitrary, fast-to-compute, tree-topology penalty function. No other tool has this flexibility. Second, EXACT can find the k trees with the smallest objective value in (72). A few existing tools can output multiple trees, but only when these all have the same “heuristically-optimal” objective value. This feature is very important because, given that the input data is noisy, and the number of samples is often small, it allows, e.g., one to give a confidence score for the ancestry relations in the output tree. Furthermore, experiments show that the ground-truth tree can often be found among these k best trees. Hence, using other biological principles, the ground-truth tree can often be identified from this set. Outputting just “heuristically-optimal” trees prevents this finding.




Figure 7: Ground truth tree for the input file that generated Figure 6. The small numbers listed next to each node represent mutations. Mutations indexed by the same number in different trees are the same real mutation. The root of each tree is circled in thick red. Nodes are labeled by numbers, and these labels are assigned automatically by each tool. Labels of different trees are incomparable.