

Deconvolution and phylogeny inference of structural variations in tumor genomic samples

Jesse Eaton¹, Jingyi Wang², and Russell Schwartz^{1,3}

Abstract

Phylogenetic reconstruction of tumor evolution has emerged as a crucial tool for making sense of the complexity of emerging cancer genomic data sets. Despite the growing use of phylogenetics in cancer studies, though, the field has only slowly adapted to many ways that tumor evolution differs from classic species evolution. One crucial question in that regard is how to handle inference of structural variations (SVs), which are a major mechanism of evolution in cancers but have been largely neglected in tumor phylogenetics to date, in part due to the challenges of reliably detecting and typing SVs and interpreting them phylogenetically. We present a novel method for reconstructing evolutionary trajectories of SVs from bulk whole-genome sequence data via joint deconvolution and phylogenetics, to infer clonal subpopulations and reconstruct their ancestry. We establish a novel likelihood model for joint deconvolution and phylogenetic inference on bulk SV data and formulate an associated optimization algorithm. We demonstrate the approach to be efficient and accurate for realistic scenarios of SV mutation on simulated data. Application to breast cancer genomic data from The Cancer Genome Atlas (TCGA) shows it to be practical and effective at reconstructing features of SV-driven evolution in single tumors. All code can be found at <https://github.com/jaebird123/tusv>

1 Introduction

Genomic methods have provided a wealth of information about mutational landscapes of developing cancers, but have also created a great need for sophisticated computational models to make sense of the resulting data. They have revealed extensive variation patient-to-patient (intertumor heterogeneity) as well as cell-to-cell within single patients (intratumor heterogeneity) [17] and suggested a far more complex landscape of somatic variations in cancer development than earlier mutational models [20, 11] had anticipated. Extracting meaningful biological insight from such data nonetheless remains challenging. Much effort has focused on the difficulty of identifying those variants relevant to tumorigenesis and progression, known as the drivers, from the background noise of the many more chance mutations carried along with a developing tumor despite being functional irrelevant, known as the passengers [18]. More recently, attention has shifted to understanding what one can learn even from passengers regarding how a particular tumor’s mutational spectrum [1] shapes its genome across stages of progression and how that knowledge can predict its future progression and help improve prognosis. These remain substantively unsolved problems that must be better tackled if cancer researchers are to make sense of enormous and ever-growing libraries of genetic variations in cancers.

One key advance in understanding tumor genomic data was the advent of tumor phylogenetics, i.e., the use of phylogenetic inference to reconstruct tumor progression. This field arose from the observation that cancer progression is fundamentally the evolution of clonal cell populations and thus in principle interpretable via algorithms for reconstructing evolutionary trees, i.e., phylogenetics. Tumor phylogenetics itself has greatly evolved, from its initial use in making sense of intertumor heterogeneity via oncogenetic tree models [7], through the advent of methods for interpreting variation between distinct tumor regions [15, 14], between distinct cells in single tumors [22], and ultimately to recent variants that seek to explain whole-genome evolution of numerous single-cells per tumor [24, 12, 33]. Single-cell genomic data is beginning to become available in quantity, though, most studies of non-trivial patient populations are still limited to bulk

¹Computational Biology Department, Carnegie Mellon University, Pittsburgh, PA USA

²Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA USA

³Department of Biological Sciences, Carnegie Mellon University, Pittsburgh, PA USA

sequence data, providing at best variant frequencies averaged across many single cells. Modern methods for working with such data combine phylogenetic inference with a deconvolution step, in which one infers clonal subpopulations from mixed genomic samples prior to or concurrent with inferring phylogenetic relationships between those subpopulations [27]. Numerous tumor phylogeny methods now work on this basic model of joint deconvolution and phylogenetics, with prominent examples including THETA [21], Pyclone [25], Canopy [13], PhyloWGS [6], SPRUCE [9], and CITUP [16]. See [2, 26] for recent reviews.

Despite many advances, though, key aspects of the problem of reconstructing tumor evolution from variation data remain unresolved, an important one being the interpretation of structural variations (SVs). SVs, along with the copy number aberrations (CNAs) they frequently induce, are the primary mechanism of phenotypic adaptation in developing cancers [32]. Most tumor phylogeny methods until recently focused primarily on single nucleotide variations (SNVs) (e.g., [10, 23]). SNVs are generally abundant and make for computationally simpler analyses than other marker types but omit much of the functional mutation that we often seek to understand with tumor phylogenetics. Some early methods did focus primarily on CNAs for deconvolution [29] and phylogenetics [22, 3, 28], and several tools are now available for joint inference of SNVs and CNAs (e.g., [6, 13, 9]). There is, to our knowledge, however, no method that handles phylogenetics of SVs more comprehensively. Despite their importance, SVs introduce a number of technical challenges, including difficulty of reliable detection leading to a high expected missing data rate, of reconstructing variants that by their nature are associated with copy number variant regions of the genome, and of interpreting these more complicated event types phylogenetically.

The goal of this paper is to address the lack of methods for tumor deconvolution and phylogenetics of diverse classes of SVs at nucleotide resolution. Specifically, we develop a new method for simultaneously deconvolving inferences of SVs, derived from the Weaver variant caller [30], and reconstructing the likely evolution of clonal populations via these SV events. The method relies on a novel model extending prior literature on SNV and CNA phylogenetics [9] to handle SVs. It depends on a model of joint likelihood of genomic sequence data and clonal phylogenies, which we pose and solve through a combinatorial coordinate descent inference strategy. We demonstrate, on simulated and Cancer Genome Atlas (TCGA [19]) samples that these methods are practical and effective in inferring progression of major clones from bulk whole genome sequence (WGS) data.

2 Methods

2.1 Breakpoint and Structural Variant Definitions

Let $\text{chr}:\text{pos}$ denote the position and chromosome for each base pair in a reference genome. For example, 7:501 represents the base pair at position 501 on chromosome 7. We define a **breakpoint** as any base pair $c:i$ that is found nonadjacent to either base pair $c:i-1$ or base pair $c:i+1$. If base pair $c:i$ was found nonadjacent to base pair $c:i-1$ we denote the breakpoint as $[c:i[$ as the intact chromatin extends to the right while if base pair $c:i$ was found nonadjacent to base pair $c:i+1$ we denote the breakpoint as $]c:i]$. We define a **structural variant** (SV) as a pair of breakpoints found adjacent to one another in the cancer genome but at non-adjacent positions in a reference genome. We call each such pair of breakpoints a mated pair, or mates for short. For example, SV $]2:30[$, $]5:10[$ means that the segment on the reference genome on chromosome 2 at position 30 extending to the left was found next to the segment on the reference genome on chromosome 5 at position 10 in the cancer genome. This is specifically an example of a translocation SV, as the rearrangement involves different chromosomes.

To relate SVs to CNAs, we assume the reference genome is partitioned into r segments, with breakpoints positioned on the ends of segments excluding ends of chromosomes. (In practice, edges of segments are not always supported by breakpoints as mated breakpoints cannot always be supported with a sufficient number of reads). Each breakpoint is found in exactly one segment. Because of this, we can define both the number of times a mated breakpoint appears in a genome (denoted c_b for the copy number of breakpoint b) and the copy number of the segments containing each breakpoint (denoted γ_b for the copy number of the segment containing breakpoint b). A more in depth example for the appearance and copy number of breakpoints is given in Figure 1.




Figure 1: Example genomes before and after segmental deletion and duplication. Top images are the reference genome while bottom images are the genome after deletion / duplication. Each colored box represents a single base pair and base pairs between two vertical orange lines represent segments. The letters below a base pair identify the position of that base pair in the reference genome. Assume this example holds for any single chromosome labeled z . (1a) shows a deletion of segment 2 (base pairs B through C) producing structural variant $]\mathbf{z}:\mathbf{A}], [\mathbf{z}:\mathbf{D}[$. The copy number of each of the mated breakpoints ($c_{[\mathbf{z}:\mathbf{A}]}$ and $c_{[\mathbf{z}:\mathbf{D}[}$) and the copy number of each of the segments containing these breakpoints ($\gamma_{[\mathbf{z}:\mathbf{A}]}$ and $\gamma_{[\mathbf{z}:\mathbf{D}[}$) are all 1 ($c_{[\mathbf{z}:\mathbf{A}]} = c_{[\mathbf{z}:\mathbf{D}[} = \gamma_{[\mathbf{z}:\mathbf{A}]} = \gamma_{[\mathbf{z}:\mathbf{D}[} = 1$). (1b) shows a duplication of segment 2 producing structural variant $[\mathbf{z}:\mathbf{B}[$, $]\mathbf{z}:\mathbf{C}[$. The copy number of each of the mated breakpoints is 1 ($c_{[\mathbf{z}:\mathbf{B}[} = c_{[\mathbf{z}:\mathbf{C}[} = 1$) while the copy number of each of the segments containing breakpoints is 2 ($\gamma_{[\mathbf{z}:\mathbf{B}[} = \gamma_{[\mathbf{z}:\mathbf{C}[} = 2$).

2.2 Problem Statement

Our method takes as input variant calls. We currently assume these calls are of the form produced by Weaver [30], which calls SVs and CNAs from bulk genomic read data and estimates copy numbers for copy number segments and breakpoints supporting the SVs. These variant calls are used to construct an $m \times (\ell + r)$ mixed copy number matrix F , the rows of which represent tumor samples and columns of which represent mutations. The first ℓ columns correspond to breakpoints and the next r to mixed segmented copy numbers. The variant calls also provide a mapping of breakpoint positions to segments, which we code as an $\ell \times r$ binary matrix Q . We also use information mapping breakpoints to structural variants, encoded as $\ell \times \ell$ binary matrix G . From these inputs, we seek simultaneously to infer an integer copy number matrix C , which describes copy numbers across the genome regions profiled for each inferred clonal cell population; a mixture fraction matrix U , which describes how clonal populations are distributed among tumor samples; and a phylogeny T , describing ancestral relationships among the clones. We assume the number of leaves n in the phylogenetic tree containing $N = 2n - 1$ total nodes (clones) is known. More formally, given

$$\begin{aligned}
 F &\in \mathbb{R}_{\geq 0}^{m \times (\ell + r)} & f_{p,s} &\text{ is the mixed copy number of variant } s \text{ in sample } p \\
 Q &\in \{0, 1\}^{\ell \times r} & q_{b,s} &\text{ is 1 iff breakpoint } b \text{ is in segment } s \\
 G &\in \{0, 1\}^{\ell \times \ell} & g_{s,t} &\text{ is 1 iff breakpoints } s \text{ and } t \text{ are mated pairs} \\
 & & n &\in \mathbb{N} \text{ number of leaves in the phylogenetic tree} \\
 c_{max} &\in \mathbb{N}_{>2} & &\text{ maximum allowed subclonal copy number for breakpoints and segments} \\
 \lambda_1 &\in \mathbb{R}_{\geq 0} & &\text{ regularization term to weight total tree cost} \\
 \lambda_2 &\in \mathbb{R}_{\geq 0} & &\text{ regularization term to weight breakpoint consistency cost}
 \end{aligned}$$

where $\sum_{s=1}^r q_{b,s} = 1 \forall b \in \{1, \dots, \ell\}$, $\sum_{b'=1}^{\ell} g_{b,b'} = 1 \forall b \in \{1, \dots, \ell\}$, we seek to determine

$$\begin{aligned}
 C &\in \mathbb{Z}_{\geq 0}^{N \times (\ell + r)} & c_{k,s} &\text{ is the integer copy number of segment or breakpoint } s \text{ in clone } k \\
 U &\in [0, 1]^{m \times N} & u_{p,k} &\text{ is the cell type } k \text{ that makes up sample } p \\
 E &\in \{0, 1\}^{N \times N} & e_{i,j} &\text{ is 1 iff directed edge } (v_i, v_j) \text{ exists in the inferred phylogeny } T
 \end{aligned}$$

and minimize the objective function

$$\min_{U, C} (|F - UC| + \lambda_1 R + \lambda_2 S) \quad (1)$$

where $|F - UC|$ describes the deviation between true and inferred mixed copy numbers, R is a phylogenetic cost, S is a cost capturing consistency between SVs and copy number segments, and λ_1 and λ_2 are regularization terms (constants).

2.3 Coordinate Descent Algorithm Overview

We solve for U , C , and T given F , Q , and G using coordinate descent [31]. We write two linear programs: one solving for U given F and C and the other solving for C given U and F . We then iteratively alternate between solving for U and for C while holding the other constant, either until convergence where U and C remain unchanged between iterations, or until a maximum number of iterations is reached. To avoid local minima, we run coordinate descent on multiple random initializations of U . Each row in U is independently randomly uniformly initialized so $\sum_{k=1}^N u_{p,k} = 1 \forall p \in \{1, \dots, m\}$ and samples independently distributed.

2.4 Estimating U

In solving for Eq. 1, we define the L1 distance $|F - UC|$ as

$$f_{\Delta,p,s} \geq f_{p,s} - \sum_{k=1}^N u_{p,k} \cdot c_{k,s} \quad \forall p \in \{1, \dots, m\}, s \in \{1, \dots, \ell + r\} \quad (2)$$

$$f_{\Delta,p,s} \geq -f_{p,s} + \sum_{k=1}^N u_{p,k} \cdot c_{k,s} \quad \forall p \in \{1, \dots, m\}, s \in \{1, \dots, \ell + r\} \quad (3)$$

$$|F - UC| = \sum_{p=1}^m \sum_{s=1}^{\ell+r} f_{\Delta,p,s} \quad (4)$$

Assume then that F and C are given. To ensure each element $u_{p,k} \in U$ is a percentage of cell type k in sample p and that percentages for a single sample sum to 1, we constrain $u_{p,k}$ so

$$0 \leq u_{p,k} \leq 1 \quad \forall p \in \{1, \dots, m\}, k \in \{1, \dots, N\} \quad (5)$$

$$\sum_{k=1}^N u_{p,k} = 1 \quad \forall p \in \{1, \dots, m\} \quad (6)$$

Since the regularization terms in our minimization (equation 1) do not depend on U , we can then simply find U to minimize $|F - UC|$ (equation 4) given F and C subject to constraints (2), (3), (5), and (6).

2.5 Estimate C

We then estimate C and T given F , U , Q , and G .

2.5.1 Binary Indicator Variables

Any variable x has an associated indicator variable \bar{x} defined as

$$\text{bin}(x) = \bar{x} = \begin{cases} 1 & x > 0 \\ 0 & x = 0 \end{cases} \quad (7)$$

This is used throughout the following sections. To linearly define \bar{x} , we introduce temporary variable $y_b \in \{0, 1\}$ as the bit representation of x over q bits [31]. The values of temporary variable y_b only apply to equations (8) and (9). y_b is then defined by

$$\sum_{b=0}^{\lceil \log_2 x_{max} \rceil + 1} 2^b \cdot y_b = x \quad (8)$$

and constrains \bar{x} as

$$0 \leq y_b \leq \bar{x} \leq \sum_{a=0}^{\lceil \log_2 x_{max} \rceil + 1} y_a \quad \forall b \in \{0, \dots, \lceil \log_2 x_{max} \rceil + 1\} \quad (9)$$

so \bar{x} is 0 if all bits b are 0 and 1 if any bit of x is 1. In this way, any integer variable x with a maximum value x_{max} can be represented in binary form \bar{x} . Binary indicator variables are noted with a bar on top \bar{x} or by $bin(x)$.

2.5.2 Phylogenetic Constraints

Since the individual rows of C are not independent but instead share a phylogenetic history, we create a tree structure T representing the inferred relationships between rows in C . We define a binary tree T using a $N \times N$ directed adjacency matrix E . To impose a tree structure on E , assume the first n clones are leaf nodes and clones $n + 1$ through $2n - 1 = N$ are internal nodes, with node N as the root. We constrain element $e_{i,j}$ as follows:

$$\text{root, incoming edges} \quad e_{i,N} = 0 \quad \forall i \in \{1, \dots, N\} \quad (10)$$

$$\text{non-root, incoming edges} \quad \sum_{i=1}^{N-1} e_{i,j} = 1 \quad \forall j \in \{1, \dots, N-1\} \quad (11)$$

$$\text{leaves, outgoing edges} \quad e_{i,j} = 0 \quad \forall i \in \{1, \dots, n\}, j \in \{1, \dots, N\} \quad (12)$$

$$\text{internal nodes, outgoing edges} \quad \sum_{j=1}^N e_{i,j} = 2 \quad \forall i \in \{n+1, \dots, N\} \quad (13)$$

Equations (10) and (11) ensure the root has no in-edges and all other nodes have exactly one in-edge. Equations (12) and (13) force leaves to have no out-edges and all internal nodes to have exactly two out-edges.

2.5.3 Phylogenetic Cost

We next ensure all copy numbers are below some input maximum c_{max} and force the normal (non tumor) root node to be diploid (each segment having copy number 2) and free of structural variants (copy number of all breakpoints is 0):

$$c_{k,s} \leq c_{max} \quad \forall k \in \{1, \dots, N\}, s \in \{1, \dots, \ell + r\} \quad (14)$$

$$c_{N,b} = 0 \quad \forall b \in \{1, \dots, \ell\} \quad (15)$$

$$c_{N,s+\ell} = 2 \quad \forall s \in \{1, \dots, r\} \quad (16)$$

We next model a phylogenetic tree cost, using CNAs to estimate evolutionary distance $\rho_{i,j}$ across each tree edge $(v_i, v_j) \in E$. We approximate evolutionary distance by the L1 distance between the copy number profiles of an edge's endpoints $\sum_{s=1}^r |c_{i,s+\ell} - c_{j,s+\ell}|$. While there are more sophisticated models of copy number distance in the literature [5, 28, 4, 8], we use L1 distance as an approximation as it can be coded and computed efficiently within the ILP framework. To linearly define $\rho_{i,j}$ we use temporary variable $x_{i,j,s} \in \mathbb{N}^{N \times N \times r}$, defined as the absolute change in copy number of segment s on edge (v_i, v_j) . Here, the values of temporary variable $x_{i,j,s}$ only apply to equations (17) through (20).

$$0 \leq x_{i,j,s} \leq c_{max} \cdot e_{i,j} \quad \forall i, j \in \{1, \dots, N\}, s \in \{1, \dots, \ell + r\} \quad (17)$$

$$x_{i,j,s} \geq c_{i,s+\ell} - c_{j,s+\ell} - c_{max} \cdot (1 - e_{i,j}) \quad \forall i, j \in \{1, \dots, N\}, s \in \{1, \dots, \ell + r\} \quad (18)$$

$$x_{i,j,s} \geq -c_{i,s+\ell} + c_{j,s+\ell} - c_{max} \cdot (1 - e_{i,j}) \quad \forall i, j \in \{1, \dots, N\}, s \in \{1, \dots, \ell + r\} \quad (19)$$

Equation (17) sets the cost to zero for any pair of nodes (v_i, v_j) where v_i is not the parent of v_j , while equations (18) and (19) set the cost to be the absolute difference between copy number for of end nodes for any edge (v_i, v_j) . We then define the cost across edge (v_i, v_j) and total cost of tree as

$$\rho_{i,j} = \sum_{s=1}^r x_{i,j,s} \quad \forall i, j \in \{1, \dots, N\} \quad (20)$$

$$R = \sum_{i=1}^N \sum_{j=1}^N \rho_{i,j} \quad (21)$$

2.5.4 Perfect Phylogeny on Appearance of Breakpoints

We next impose a perfect phylogeny on breakpoints. While the perfect phylogeny assumption is problematic for other variant types, we argue that it is sufficiently unlikely for a base-resolution breakpoint to recur that it can be neglected. Note that violations of the infinite site model due to allelic loss are handled separately by treating a lost allele as having copy number zero. We therefore impose constraints to force each breakpoint to appear across exactly one edge in T and for mated breakpoints to appear together. Define $W \in \{0, 1\}^{N \times N \times \ell}$, where each element $w_{i,j,b}$ is 1 if the copy number of breakpoint b goes from 0 to a positive integer across edge (v_i, v_j) and 0 otherwise. To linearly define $w_{i,j,b}$ we define temporary variable $x_{i,j,b} \in \{0, 1, 2, 3\}$ to be

$$x_{i,j,b} = 2 + \bar{c}_{i,b} - \bar{c}_{j,b} - e_{i,j} \quad \forall i, j \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (22)$$

so $x_{i,j,b}$ is 0 iff the copy number of breakpoint b increases from 0 across edge (v_i, v_j) . The value of temporary variable $x_{i,j,b}$ only applies to equations (22) and (23). Using the binary representation $\bar{x}_{i,j,b}$ of $x_{i,j,b}$, define $w_{i,j,b}$ and ensure $w_{i,j,b}$ is 1 for a single edge in the tree.

$$w_{i,j,b} = 1 - \bar{x}_{i,j,b} \quad \forall i, j \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (23)$$

$$\sum_{i=1}^N \sum_{j=1}^N w_{i,j,b} = 1 \quad \forall b \in \{1, \dots, \ell\} \quad (24)$$

Using breakpoint mate indicator $g_{s,t} \in \{0, 1\}$, where $g_{s,t}$ is 1 iff breakpoints s and t are mates, we force breakpoint indicators to be equal for mates.

$$w_{i,j,s} - w_{i,j,t} \leq 1 - g_{s,t} \quad \forall i, j \in \{1, \dots, N\}, s, t \in \{1, \dots, \ell\} \quad (25)$$

$$-w_{i,j,s} + w_{i,j,t} \leq 1 - g_{s,t} \quad \forall i, j \in \{1, \dots, N\}, s, t \in \{1, \dots, \ell\} \quad (26)$$

Note we extend the notation of breakpoint appearance indicator $w_{i,j,b}$ to have $w_{j,b} = \sum_{i=1}^N w_{i,j,b}$ be 1 if breakpoint b appears at node v_j and 0 otherwise.

2.5.5 Ancestry Condition for Non-Disappearing SVs

We next impose the two-state perfect phylogeny ancestry condition as described in [10] for the appearance of breakpoints. For any breakpoint s that appears as an ancestor to breakpoint t , the total fraction of cells with breakpoint s must be larger than the fraction with breakpoint t so long as breakpoint s never subsequently disappears. To enforce this, the fraction of cells $\phi_{p,b}$ containing breakpoint b in sample p is defined as

$$\phi_{p,b} = \sum_{k=1}^N u_{p,k} \cdot \bar{c}_{k,b} \quad \forall p \in \{1, \dots, m\}, b \in \{1, \dots, s\} \quad (27)$$

We then must define a few variables to force $\phi_{p,s} \geq \phi_{p,t}$ if breakpoint s appears before breakpoint t and is never subsequently lost. Let v_i be the i^{th} node in the phylogeny and $v_i \prec v_j$ denote that node v_i is an ancestor of v_j . We first define ancestor variables $a_{i,j} \in \{0, 1\}$ as 1 if $v_i \prec v_j$ and 0 otherwise for all $i, j \in \{1, \dots, N\}$. Linearly define $a_{i,j}$ by

$$\text{root } v_N \text{ is ancestor to all nodes} \quad a_{N,j} = 1 \quad \forall j \in \{1, \dots, N-1\} \quad (28)$$

$$\text{root } v_N \text{ has no ancestors} \quad a_{i,N} = 0 \quad \forall i \in \{1, \dots, N\} \quad (29)$$

$$\text{any parent is an ancestor} \quad a_{i,j} \geq e_{i,j} \quad \forall i, j \in \{1, \dots, N\} \quad (30)$$

$$\text{child gets parent's ancestor profile} \quad a_{g,j} \geq e_{i,j} + a_{g,i} - 1 \quad \forall i, j \in \{1, \dots, N\}, g \in \{1, \dots, i-1, i+1, \dots, N\} \quad (31)$$

$$a_{g,j} \leq 1 - e_{i,j} + a_{g,i} \quad \forall i, j \in \{1, \dots, N\}, g \in \{1, \dots, i-1, i+1, \dots, N\} \quad (32)$$

Next, define the number of descendants to node v_i with at least one copy of breakpoint b as $d_{i,b}$ for all $i \in \{1, \dots, N\}$, $b \in \{1, \dots, \ell\}$. To linearly define $d_{i,b}$, define temporary binary variables $x_{i,j,b} \in \{0, 1\}$ for equations (33) through (36) for all $i, j \in \{1, \dots, N\}$, $b \in \{1, \dots, \ell\}$ to be 1 if $a_{i,j}$ and $\bar{c}_{j,b}$ and zero otherwise.

$$x_{i,j,b} \geq a_{i,j} + \bar{c}_{j,b} - 1 \quad \forall i, j \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (33)$$

$$x_{i,j,b} \leq a_{i,j} \quad \forall i, j \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (34)$$

$$x_{i,j,b} \leq \bar{c}_{j,b} \quad \forall i, j \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (35)$$

$$d_{i,b} = \sum_{j=1}^N x_{i,j,b} \quad \forall i \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (36)$$

Define temporary binary variables $\bar{y}_{i,b} \in \{0, 1\}$ for equations (37) through (39) to be 0 to be zero if all descendants of node v_i contain at least one copy of breakpoint b and 1 otherwise.

$$y_{i,b} = \sum_{j=1}^N a_{i,j} - d_{i,b} \quad \forall i \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (37)$$

$$\bar{y}_{i,b} = \text{bin}(y_{i,b}) \quad \forall i \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (38)$$

Define temporary binary variable $\bar{z}_{i,j,s,t} \in \{0, 1\}$ for equations (39) and (40) to be 0 only if breakpoint s appears at node v_i , breakpoint t appears at node v_j , node v_i is an ancestor to node v_j , and breakpoint s never disappears.

$$z_{i,j,s,t} = 3 - w_{i,s} - w_{j,t} - a_{i,j} + \bar{y}_{i,s} \quad \forall i, j \in \{1, \dots, N\}, s, t \in \{1, \dots, \ell\} \quad (39)$$

Finally, apply the condition that the fraction of cells $\phi_{p,s}$ containing breakpoint s in sample p must be larger than the fraction of cells $\phi_{p,t}$ containing breakpoint t in sample p if breakpoint s appears in an ancestor to the node where breakpoint t appears and breakpoint s is never lost in any descendant (no descendant has copy number 0 for breakpoint s).

$$\phi_{p,s} \geq \phi_{p,t} - 1 + \sum_{i=1}^N \sum_{j=1}^N (1 - \bar{z}_{i,j,s,t}) \quad \forall p \in \{1, \dots, m\}, s, t \in \{1, \dots, \ell\} \quad (40)$$

Note that $\sum_{i=1}^N \sum_{j=1}^N (1 - \bar{z}_{i,j,s,t})$ can only take on values 0 or 1 since breakpoint appearance indicator $w_{i,s}$ and $w_{j,t}$ can only be both 1 at most once across all i, j . This means the condition $\phi_{p,s} \geq \phi_{p,t}$ only holds when breakpoint s appears before breakpoint t and never subsequently disappears. Note the ancestry condition is implied by but weaker than the sum condition described in [10], but can similarly be enforced by linear constraints.

2.5.6 Structural Variant and Segment Consistency

Since each breakpoint belongs to exactly one segment, we define the copy number of each segment containing a breakpoint b and constrain it so a breakpoint's copy number never exceeds that of its containing segment:

$$c_{k,b} \leq \gamma_{k,b} = \sum_{s=1}^r q_{b,s} \cdot c_{k,s+\ell} \quad \forall k \in \{1, \dots, N\}, b \in \{1, \dots, \ell\} \quad (41)$$

where input $q_{b,s} \in \{0, 1\}$ is 1 if segment s contains breakpoint b . $\sum_{s=1}^r q_{b,s} = 1$ as each breakpoint belongs to a single segment. We similarly define $\psi_{p,b}$ directly from the input to be the mixed copy number of the segment containing breakpoint b .

$$\psi_{p,b} = \sum_{s=1}^r q_{b,s} \cdot f_{p,s+\ell} \quad \forall p \in \{1, \dots, m\}, b \in \{1, \dots, \ell\}$$

$$\pi_{p,b} = \frac{f_{p,b}}{\psi_{p,b}} \quad \forall p \in \{1, \dots, m\}, b \in \{1, \dots, \ell\}$$

Intuitively, the ratio $\pi_{p,b}$ of the mixed copy number of a breakpoint to the mixed copy number of the segment containing that breakpoint should be maintained in the integer output as this preserves the difference in mutation types (duplication, deletion). To penalize for discrepancies between the inferred ratio of breakpoint and its segment copy number given $\pi_{p,b}$, we incorporate the following quantity into our objective function:

$$\left| \pi_{p,b} - \frac{\sum_{k=1}^N (u_{p,k} \cdot c_{k,b})}{\sum_{k=1}^N (u_{p,k} \cdot \gamma_{k,b})} \right|$$

To convert this from a ratio to units of copy numbers, we rearrange the expression and define S for the final term in the objective function (1) to be

$$z_{p,b} \geq \pi_{p,b} \cdot \sum_{k=1}^N (u_{p,k} \cdot \gamma_{k,b}) - \sum_{k=1}^N (u_{p,k} \cdot c_{k,b}) \quad \forall p \in \{1, \dots, m\}, b \in \{1, \dots, \ell\} \quad (42)$$

$$z_{p,b} \geq -\pi_{p,b} \cdot \sum_{k=1}^N (u_{p,k} \cdot \gamma_{k,b}) + \sum_{k=1}^N (u_{p,k} \cdot c_{k,b}) \quad \forall p \in \{1, \dots, m\}, b \in \{1, \dots, \ell\} \quad (43)$$

$$S = \sum_{p=1}^m \sum_{b=1}^{\ell} z_{p,b} \quad (44)$$

In this way, increased emphasis is placed on the relationship between segments and breakpoints. The solution for C and T is found by minimizing (1) subject to constraints (2) through (44).

3 Results

3.1 Simulated Data

To validate accuracy of the method on data of known ground truth, we assess accuracy in inference of copy number profiles across clones. For each such test, we generate a copy number matrix C_{tru} containing breakpoints and segments, mix this matrix with a mixture fraction matrix U_{tru} to get the mixed copy number matrix ($C_{tru} \times U_{tru} \rightarrow F$), run our deconvolution algorithm, and compare the inferred copy number matrix C_{inf} with the original true copy number matrix C_{tru} . We score our result as the L1 distance ($|C_{tru} - C_{inf}|$) between copy number matrices after a maximum matching between copy number profiles (for clones).

To generate C_{tru} , we simulated mutation data varying the expected number of mutations l , number of samples m , and number cell types n . For each triplet (l, m, n) , five synthetic patients were generated. Reported scores are averaged across those 5 patients. For each run of the simulation, we generated a binary tree T with n leaves and a random topology. Mutations were assigned so that the expected numbers of mutations across all edges in each tree are equal. We start with a genomic profile for the root (assumed to be a normal diploid cell containing no structural variants) and progressively added a Poisson-distributed number of mutations across each edge down to the leaves. Initially, the root node contains three pairs of homologous chromosomes of the same lengths as human chromosomes 1-3. To generate mutations, a central location is uniformly chosen across all chromosomes, then a mutation size is sampled from an exponential distribution, with expectation equal to the mean structural variant size found across 59 TCGA samples. The mutation type is uniformly randomly selected to be either a tandem duplication, deletion, or inversion. From the generated tree, we obtain a copy number matrix C_{tru} . We then create a cell type mixture matrix U_{tru} by uniformly randomly assigning cell type fractions such that the fraction of all cell types in each sample sums to 1. U_{tru} and C_{tru} are subsequently multiplied to generate mixed copy number matrix F .

Since there is no method for validating how accurate the choice of regularization terms λ_1 and λ_2 are on real data, we define empirical values for these terms based on each sample and show they perform well on simulated data. We choose regularization terms λ_1 and λ_2 empirically from the data to be $\lambda_1 = \frac{\ell+r}{\ell} \cdot \frac{m}{N}$ and $\lambda_2 = \frac{\ell+r}{\ell}$. This allows the maximum error in the $|F - UC|$ term in the minimization, which is $m \cdot (\ell+r) \cdot c_{max}$, to equal the maximum errors in $\lambda_1 R$ and $\lambda_2 S$ terms, which are $\ell \cdot N \cdot c_{max}$ and $\ell \cdot m \cdot c_{max}$ respectively. To

show these empirical definitions do as well as iteratively choosing the hyperparameters, we test on simulated data generated for $n = 3$ leaves, $m = 3$ samples, and $l = 50$ mutations as this produces approximately 100 breakpoints, a value comparable to the average number of breakpoints found in real, TCGA samples. Figure 2 shows that empirically selecting hyperparameters λ_1 and λ_2 (solid green curve) performs reasonably well as compared to other hyper parameter values (dotted blue curve). Both outperform the algorithm when excluding the regularization terms (dashed red curve) indicating the usefulness of including phylogenetic cost and breakpoint-segment consistency into the model.




Figure 2: Score on simulated data for varying hyperparameters λ_1 and λ_2 . Score is the L1 distance between the true C_{tru} and inferred C_{inf} copy number matrices after maximum matching (lower means better performance). Dotted blue line is score when $\lambda_2 = 0$ is held constant and λ_1 varies from 0.01, 0.05, 0.25, 1.25, 6.25 across the x-axis. Solid green line shows the score when hyperparameters are automatically chosen based on the number of SVs and CNAs in the data set while dashed red line shows score for when only a single term in the optimization expression is used.

3.2 TCGA Data

We next apply the methods to a selection of TCGA breast cancer (BRCA) samples [19], restricting analysis to a subset of 59 samples for which WGS data was available. Of these, 31 ran successfully within our prescribed run time limit, while 28 with the highest SV counts timed out before completion or required more memory than was available to us (we have access to 128Gb of RAM). Since there is no known ground truth for these samples, we cannot assess their individual accuracy. Nonetheless, they provide some basis for analysis of trends across samples. Space does not permit us to display all observed trees, so for purposes of illustration we classify them into seven observed topologies (A-G), shown in Fig. 3, with frequencies of occurrence shown in Fig. 4. None of the inferred trees are purely linear, consistent with a model of significant subclonal heterogeneity rather than a simple sequential model of clonal progression. Quantitation by several measures of heterogeneity, as shown in Fig. 5, likewise suggests a wide diversity among samples. The data is suggestive of a possible clustering into distinct low-diversity and high-diversity subclusters, but with substantial overlap between clusters.




Figure 3: Tree topologies observed across 31 TCGA BRCA samples, grouped into seven categories (A–G).




Figure 4: Histogram of occurrences of tree topologies across 31 TCGA BRCA samples.




Figure 5: Violin plots quantifying heterogeneity of the 31 TCGA trees. (a) Number of unique clones. (b) Diversity, defined in terms of the clonal frequency vector $u_{p,k}$ as $1 - \sum_{k=1}^N u_{p,k}$. (c) Number of breakpoints.

4 Discussion

We have developed a new method for automated joint deconvolution and phylogeny inference of tumor genomic data designed to address the important unsolved problem of describing progression via SVs. We specifically learn a model encompassing CNAs and SVs of major clones, mixture fractions of these clones across samples, and phylogenetic tree relating the clones. We pose the model inference problem to balance the likelihood of sequence read data with respect to copy numbers and observed breakpoints against the evolutionary cost of the phylogenetic tree. We solve the resulting model via a coordinate descent algorithm posed as a pair of MILPs. We demonstrate that the method can accurately and efficiently reconstruct clonal populations and phylogenetic histories from both simulated tumor data. Application to WGS data from the TCGA shows the method to be effective on real data supportive of a range of tree topologies and complexities.

This work provides a proof-of-concept demonstration of the feasibility of more comprehensively modeling the important role of SVs in tumor evolution, but also suggests a number of avenues for future work. Our methods currently rely on a sometimes costly and potentially suboptimal model fitting algorithm, and further algorithmic advances might plausibly lead both to greater efficiency and improved solution quality.

In particular, there are currently practical limits on the total SNV counts the method can handle without excessive run time and memory usage. While most of the TCGA breast cancers considered fell within those limits, a significant minority did not. Furthermore, our work focuses only on the subproblem of handling SVs (and associated CNAs), and will likely benefit from incorporating other variant types, most notably SNVs but also potentially expression, methylation, or other markers of cell state. Furthermore, biotechnology for data generation is continuing to advance, with growing numbers of computational methods taking advantage of single cell sequence or long read technologies that can provide direct single-cell readouts of SNV or CNA data. SV detection is problematic for all current single-cell technologies, and we can anticipate value in combining single-cell methods with bulk deconvolution methods such as ours for SVs. Finally, the present work has focused only on the development of the new technology and its validation. The ultimate value of the work will lie in bringing SV-aware phylogenetics to diverse patient cohorts, to begin to develop a comprehensive understanding of the landscape of SV variation in tumor progression and its implications for patient prognosis and treatment.

Acknowledgments: We thank Ashok Rajaraman and Jian Ma for helpful discussions and assistance with Weaver. Portions of this work have been funded by the U.S. National Institutes of Health via award R21CA216452 and Pennsylvania Dept. of Health Grant GBMF4554 4100070287. The Pennsylvania Department of Health specifically disclaims responsibility for any analyses, interpretations or conclusions. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. Specifically, it used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC). The results published here are in whole or part based upon data generated by The Cancer Genome Atlas managed by the NCI and NHGRI. Information about TCGA can be found at <http://cancergenome.nih.gov>.

References

- [1] L. B. Alexandrov and M. R. Stratton. “Mutational signatures: the patterns of somatic mutations hidden in cancer genomes”. In: *Curr. Opin. Genet. Devel.* 24 (2014), pp. 52–60.
- [2] N. Beerenwinkel et al. “Cancer evolution: mathematical models and computational inference”. In: *Syst. Biol.* 64.1 (2015), e1–e25.
- [3] S. A. Chowdhury et al. “Phylogenetic analysis of multiprobe fluorescence in situ hybridization data from tumor cell populations”. In: *Bioinformatics* 29.13 (2013), pp. i189–i198.
- [4] S.A. Chowdhury et al. “Inferring models of multiscale copy number evolution for single-tumor phylogenetics”. In: *Bioinformatics* 31.12 (2015), pp. i258–i267.
- [5] Salim Akhter Chowdhury et al. “Algorithms to model single gene, single chromosome, and whole genome copy number changes jointly in tumor phylogenetics”. In: *PLoS computational biology* 10.7 (2014), e1003740.
- [6] A. G. Deshwar et al. “PhyloWGS: Reconstructing subclonal composition and evolution from whole-genome sequencing of tumors”. In: *Genome Biol.* 16 (2015), p. 35.
- [7] R. Desper et al. “Inferring tree models of oncogenesis from comparative genomic hybridization data”. In: *J. Comput. Biol.* 6.1 (1999), pp. 37–51.
- [8] M. El-Kebir et al. “Complexity and algorithms for copy-number evolution problems”. In: *Algorithms for Molecular Biology* 12.1 (2017), p. 13.
- [9] M. El-Kebir et al. “Inferring the mutational history of a tumor using multi-state perfect phylogeny mixtures”. In: *Cell Systems* 3.1 (2016), pp. 43–53.
- [10] M. El-Kebir et al. “Reconstruction of clonal trees and tumor composition from multi-sample sequencing data”. In: *Bioinformatics* (2015).
- [11] E. Fearon and B. Vogelstein. “A genetic model for colorectal tumorigenesis”. In: *Cell* 61.5 (1990), pp. 759–767.

- [12] K. Jahn, J. Kuipers, and N. Beerenwinkel. “Tree inference for single-cell data”. In: *Genome Biol.* 17 (2016), p. 96.
- [13] Y. Jiang et al. “Assessing intratumor heterogeneity and tracking longitudinal and spatial clonal evolutionary history by next-generation sequencing”. In: *Proc. Natl. Acad. Sci. USA* (2016), p. 201522203.
- [14] L. Khaliq et al. “The clonal evolution of metastases from primary serous epithelial ovarian cancers”. In: *Int. J. Cancer* 124.7 (2009), pp. 1579–1586.
- [15] C. C. Maley et al. “Genetic clonal diversity predicts progression to esophageal adenocarcinoma”. In: *Nat. Genet.* 38.4 (2006), pp. 468–473.
- [16] S. Malikic et al. “Clonality inference in multiple tumor samples using phylogeny”. In: *Bioinformatics* 31.9 (2015), pp. 1349–1356.
- [17] A. Marusyk and K. Polyak. “Tumor heterogeneity: causes and consequences”. In: *Biochim. Biophys. Acta (BBA)-Rev. Cancer* 1805.1 (2010), pp. 105–117.
- [18] N. McGranahan et al. “Clonal status of actionable driver events and the timing of mutational processes in cancer evolution”. In: *Sci. Transl. Med.* 7.283 (2015), 283ra54.
- [19] Cancer Genome Atlas Network et al. “Comprehensive molecular portraits of human breast tumours”. In: *Nature* 490.7418 (2012), pp. 61–70.
- [20] P. C. Nowell. “The clonal evolution of tumor cell populations”. In: *Science* 194.4260 (1976), pp. 23–28.
- [21] L. Oesper, G. Satas, and B.J. Raphael. “Quantifying tumor heterogeneity in whole-genome and whole-exome sequencing data.” In: *PubMed* (2014). THetA 2.
- [22] G. Pennington et al. “Reconstructing tumor phylogenies from heterogeneous single-cell data”. In: *J. Bioinform. Comput. Biol.* 5.2a (2007), pp. 407–427.
- [23] V. Popic et al. “Fast and scalable inference of multi-sample cancer lineages”. In: *Genome Biol.* 16 (2015), p. 91.
- [24] E. M. Ross and F. Markowitz. “OncoNEM: inferring tumor evolution from single-cell sequencing data”. In: *Genome Biol.* 17.1 (2016), p. 69.
- [25] A. Roth et al. “PyClone: statistical inference of clonal population structure in cancer”. In: *Nat. Meth.* 11.4 (2014), pp. 396–398.
- [26] R. Schwartz and A. A. Schäffer. “The evolution of tumour phylogenetics: principles and practice”. In: *Nature Reviews Genetics* 18.4 (2017), pp. 213–229.
- [27] R. Schwartz and S. E. Shackney. “Applying unmixing to gene expression data for tumor phylogeny inference”. In: *BMC Bioinformatics* 11 (2010), p. 42.
- [28] R. F. Schwarz et al. “Phylogenetic quantification of intra-tumour heterogeneity”. In: *PLoS Comput. Biol.* 10.4 (2014), e1003535.
- [29] D. Tolliver et al. “Robust unmixing of tumor states in array comparative genomic hybridization data”. In: *Bioinformatics* 26.12 (2010), pp. i106–i114.
- [30] Li Y et al. “Allele-Specific Quantification of Structural Variations in Cancer Genomes.” In: *Cell* 3 (2016).
- [31] S. Zaccarian et al. “The Copy-Number Tree Mixture Deconvolution Problem and Applications to Multi-sample Bulk Sequencing Tumor Data”. In: *Springer Link* (2017).
- [32] T. I. Zack et al. “Pan-cancer patterns of somatic copy number alteration”. In: *Nature Genetics* 45.10 (2013), pp. 1134–1140.
- [33] H. Zafar et al. “SiFit: inferring tumor trees from single-cell sequencing data under finite-sites models”. In: *Genome Biology* 18.1 (2017), p. 178.