

Modeling Intratumor Gene Copy Number Heterogeneity using Fluorescence in Situ Hybridization data

Charalampos E. Tsourakakis¹

Aalto University, Finland
tsourolampis@gmail.com

Abstract. ¹ Tumorigenesis is an evolutionary process which involves a significant number of genomic rearrangements typically coupled with changes in the gene copy number profiles of numerous cells. Fluorescence *in situ* hybridization (FISH) is a cytogenetic technique which allows counting copy numbers of genes in single cells. The study of cancer progression using FISH data has received considerably less attention compared to other types of cancer datasets.

In this work we focus on inferring likely tumor progression pathways using publicly available FISH data. We model the evolutionary process as a Markov chain in the positive integer cone \mathbb{Z}_+^g where g is the number of genes examined with FISH. Compared to existing work which oversimplifies reality by assuming independence of copy number changes [24,25], our model is able to capture dependencies. We model the probability distribution of a dataset with hierarchical log-linear models, a popular probabilistic model of count data. Our choice provides an attractive trade-off between parsimony and good data fit. We prove a theorem of independent interest which provides necessary and sufficient conditions for reconstructing oncogenetic trees [8]. Using this theorem we are able to capitalize on the wealth of inter-tumor phylogenetic methods. We show how to produce tumor phylogenetic trees which capture the dynamics of cancer progression. We validate our proposed method on a breast tumor dataset.

1 Introduction

Tumors are heterogeneous masses which exhibit cellular and genomic differences [13,20,21,22]. Cell-by-cell assay measurements allow us to study the phenomenon of tumor heterogeneity. Fluorescence *in situ* hybridization (FISH) is a cytogenetic technique which allows us to study *gene copy number heterogeneity* within a single tumor. It is used to count the copy number of DNA probes for specific genes or chromosomal regions. Understanding how tumor heterogeneity progresses is a major problem with significant potential impact on therapeutics. An example which illustrates the importance of heterogeneity in therapeutic resistance is

¹ Topic: Cancer Genomics. Keywords: intra-tumor heterogeneity, evolutionary dynamics, cancer phylogenetics, Markov chains, simulation, FISH

found in chronic myelogenous leukemia (CML). The presence of a subpopulation of leukemic cells of a given type significantly influences the response to therapies based on imatinib mesylate, causing the eventual relapse of the disease [20].




Fig. 1: Copy number genetic diversity within a tumor cell population after three hypothetical stages of tumorigenesis. The four observed states of cells are shown and coded with colors in the positive integer cone \mathbb{Z}_+^2 .

In this work we study the phenomenon of *gene copy number heterogeneity* within a single tumor. An illustration of the phenomenon of gene copy number heterogeneity is shown in Figure 1 which shows a hypothetical cell population of eleven somatic human cells at three stages of tumorigenesis, whose succession is indicated by arrows. We shall describe the state of a cell as a two dimensional point (g_1, g_2) in the positive integer cone \mathbb{Z}_+^2 , with the obvious meaning that the cell has g_1, g_2 copies of gene 1 and gene 2 respectively. Initially, all cells are in the healthy diploid state with respect to their gene copy numbers. At the second stage, nine cells are in state $(2,2)$ and two cells in state $(3,2)$. At the last stage, only 2 cells are in the healthy state $(2,2)$. Five, three and one cell are in states $(3,2)$, $(2,1)$ and $(1,1)$ respectively. The four observed states are shown and coded with colors in \mathbb{Z}_+^2 . Figure 1 illustrates what we observe in many tumors: existence of multiple progression states within a single tumor.

Despite the large amount of research work on modeling tumor progression using different types of tumor datasets, e.g., [12], the study of FISH datasets has received considerably less attention. Specifically, two early studies of FISH datasets were limited to either two [24,25] or three probes [19]. Pennington et al. [24,25] develop novel computational methods for analyzing FISH data. Specifically, they consider a random walk on the positive integer cone \mathbb{Z}_+^2 where at each step a coordinate $i \in \{1, 2\}$ is picked uniformly at random and is modified by $\Delta x \in \{0, 1, -1\}$ with probabilities $1 - p_{i,1} - p_{i,-1}, p_{i,1}, p_{i,-1}$, $i = 1, 2$ respectively. Given this model they optimize a likelihood-based objective over all possible trees and parameters $\{p_{i,1}, p_{i,-1}\}$. Recently, Chowdhury et al. [5] proposed a general procedure which can treat any number of probes. They reduce

the study of the progression of FISH probe cell count patterns to the rectilinear minimum spanning tree problem.

Paper contributions and roadmap. In this paper we achieve the following contributions:

- We introduce a novel approach to analyzing FISH datasets. The main features of our approach are its probabilistic nature which provides an attractive trade-off between parsimony and expressiveness of biological complexity and the reduction of the problem to the well-studied inter-tumor phylogeny inference problem. The former allows us to capture complex dependencies between factors while the latter opens the door to a wealth of available and established theoretical methods which exist for the inter-tumor phylogeny inference problem.
- We prove Theorem 1 which provides necessary and sufficient conditions for the unique reconstruction of an oncogenetic tree [8]. Based on the theorem’s conditions, we are able to capitalize on the wealth of inter-tumor phylogenetic methods. However, the result is of independent interest and introduces a set of interesting combinatorial questions.
- We validate our proposed method on a publicly available breast cancer dataset.

The outline of this paper is as follows: Section 2 presents our proposed methods. Section 3 performs an experimental evaluation of our methods on a breast cancer FISH dataset and an extensive biological analysis of the findings. Finally, Section 4 concludes the paper by a discussion and a brief summary.

2 Proposed Method

In Section 2.1 we model the probability distribution of FISH data with hierarchical log-linear models and show how to learn the parameters of the model for a given FISH dataset. In Section 2.2 we prove Theorem 1 which provides necessary and sufficient conditions to uniquely reconstruct an oncogenetic tree [8]. We capitalize on the theorem to harness the wealth of available methods for inter-tumor phylogenetic inference methods [1,2,3,8,9,12,11,14]. Finally in Section 2.3 we present our proposed method.

We will make the same simplifying assumptions with existing work [24,5], namely that only single gene duplication and loss events take place and that the cell population is fixed. In what follows, let $\mathcal{D} = \{x_1, \dots, x_n\}, x_i \in \mathbb{Z}_+^g$ be the input FISH dataset which measures the copy numbers of g genes in n cells taken from the same tumor.

2.1 Model and Fitting

Probabilistic Model: Let X_j be an integer-valued random variable which expresses the copies of the j -th gene with domain I_j , $j = 1, \dots, g$. We model the joint probability distribution of the random vector (X_1, \dots, X_g) as

$$\Pr[x] = \frac{1}{Z} \prod_{A \subseteq [g]} e^{\phi_A(x)} \quad (1)$$

where $x = (x_1, \dots, x_g) \in I = I_1 \times I_2 \times \dots \times I_g$ is a point of the integer positive cone \mathbb{Z}_+^g and Z is a normalizing constant, also known as the partition function, which ensures that the distribution is a proper probability distribution, i.e., $Z = \sum_{x \in I} \prod_{A \subseteq [g]} e^{\phi_A(x)}$. Each potential function ϕ_A depends only on the variables in the subset A and is parameterized by a set of weights w_A . To illustrate this, assume $g = 2$ and $I = \{0, 1\} \times \{0, 1\}$. Then,

$$\begin{aligned} \log \Pr[x] = & w_0 + w_{(1)0} \mathbb{1}\{x_1 = 0\} + w_{(1)1} \mathbb{1}\{x_1 = 1\} + w_{(2)0} \mathbb{1}\{x_2 = 0\} \\ & + w_{(2)1} \mathbb{1}\{x_2 = 1\} + w_{(12)00} \mathbb{1}\{x_1 = 0, x_2 = 0\} + w_{(12)01} \mathbb{1}\{x_1 = 0, x_2 = 1\} \\ & + w_{(12)10} \mathbb{1}\{x_1 = 1, x_2 = 0\} + w_{(12)11} \mathbb{1}\{x_1 = 1, x_2 = 1\} - \log Z, \end{aligned}$$

where w_{Ax} $A \subseteq \{1, 2\}$, $x \in \{0, 1\}^{|A|}$ are the parameters of the model. This probability distribution captures the effects of different factors through parameters w_{Ax} , $A \in \{\{1\}, \{2\}\}$, $x \in \{0, 1\}$ and pairwise interactions through parameters $w_{(12)x}$, $x \in \{00, 01, 10, 11\}$. In general, two variables X_i, X_j are defined to be directly associated if there exists at least one non-zero (or bounded away significantly from zero) parameter including the two variables. We define X_i, X_j to be indirectly associated if there exists a chain of overlapping direct associations that relate X_i, X_j .

We impose the following restriction on the probabilistic model shown in equation (1): If $A \subseteq B$ and $w_A = 0$ then $w_B = 0$. This restriction reduces significantly the size of the parameter space, but allows to express complex dependencies not captured by existing work [24,25]. Since a typical FISH dataset contains detailed measurements for a handful of genes from few hundred cells the combination of these two features is crucial to avoid overfitting and obtain biological insights at the same time. Furthermore, it is worth emphasizing that in terms of biological interpretation the assumption is natural: if a set A of genes does not interact, then any superset of A maintains that property. This class of models are known as hierarchical log-linear models [4].

Learning the Parameters: Learning the parameters w of a hierarchical log-linear model is a well-studied problem, e.g., [4]. An extensive survey of learning methods can be found in [26]. Schmidt et al. [27] propose to maximize a penalized log-likelihood of the dataset \mathcal{D} where the penalty is an overlapping group l_1 -regularization term with respect to the norms of the groups. Specifically, a

spectral projected gradient method is proposed as a sub-routine for solving the following regularization problem:

$$\min_w - \sum_{i=1}^n \log \Pr[x_i|w] + \sum_{A \subseteq [g]} \lambda_A \left(\sum_{\{B|B \subseteq A\}} \|w_B\|_2^2 \right)^{1/2} \quad (2)$$

2.2 Unique Reconstruction of Oncogenetic Trees

Our main theoretical result in this section is motivated by the following natural sequence of questions:

Can we use any of the existing inter-tumor progression methods [12] on the intra-tumor progression problem? How will the resulting tree capture the evolutionary dynamics of cancer progression, i.e., how do we enforce that state (\dots, i, \dots) is reached either through $(\dots, i+1, \dots)$ or $(\dots, i-1, \dots)$ given our single gene duplication and loss event assumption?

An answer to this question is given in Section 2.3. Motivated by our intention to capitalize on inter-tumor phylogenetic methods such as [8,9], we consider a fundamental problem concerning oncogenetic trees [8]. What are the necessary and sufficient conditions to reconstruct them? Theorem 1 is likely to be of independent interest and contributes to the understanding of oncogenetic trees [8]. We briefly review the necessary definitions to state our result. Let $T = (V, E, r)$ be an oncogenetic tree, i.e., a rooted branching², on V and let $r \in V$ be the root of T . Given a finite family $\mathcal{F} = \{A_1, \dots, A_q\}$ of sets of vertices, i.e., $A_i \subseteq V(T)$ for $i = 1, \dots, q$, where each A_i is the vertex set of a rooted sub-branching of T , what are the necessary and sufficient conditions, if any, which allow us to uniquely reconstruct T ?

Theorem 1. *The necessary and sufficient conditions to uniquely reconstruct the branching T from the family \mathcal{F} are the following:*

1. *For any two distinct vertices $x, y \in V(T)$ such that $(x, y) \in E(T)$, there exists a set $A_i \in \mathcal{F}$ such that $x \in A_i$ and $y \notin A_i$.*
2. *For any two distinct vertices $x, y \in V(T)$ such that $y \not\prec x$ and $x \not\prec y$ ³ there exist sets $A_i, A_j \in \mathcal{F}$ such that $x \in A_i$, $y \notin A_i$ and $x \notin A_j$ and $y \in A_j$.*

Proof. First we prove the necessity of conditions 1,2 and then their sufficiency to reconstruct T . In the following we shall call a branching T *consistent* with the family set \mathcal{F} if all sets $A_i \in \mathcal{F}$ are vertices of rooted sub-branchings of T .

Necessity: For the sake of contradiction, assume that Condition 1 does not hold. Then, the two branchings shown in Figure 2(a) are both consistent with \mathcal{F} and therefore we cannot reconstruct T . Similarly, assume that Condition 2 does not hold. Specifically assume that for all j such that $x \in A_j$, then $y \in A_j$ too (for the symmetric case the same argument holds). Then, both branchings in Figure 2(b) are consistent with \mathcal{F} and therefore T is not reconstructable.

² Each vertex has in-degree at most one and there are no cycles.

³ We use the notation $u \prec v$ ($u \not\prec v$) to denote that u is (not) a descendant of v in T .




Fig. 2: Illustration of necessity conditions of Theorem 1, (a) condition 1 (b) condition 2.

Sufficiency: Let $x \in V(T)$ and P_x be the path from the root to x , i.e., $P_x = \{r, \dots, x\}$. Also, define F_x to be the intersection of all sets in the family \mathcal{F} that contain vertex x , i.e., $F_x = \bigcap_{x \in A_i \in \mathcal{F}} A_i$. We prove that $F_x = P_x$. Since

by definition, $P_x \subseteq F_x$ we need to show that $F_x \subseteq P_x$. Assume that the latter does not hold. Then, there exists a vertex $v \in V(T)$ such that $v \notin P_x, v \in F_x$. We consider the following three cases.

- CASE 1 ($x \prec v$): Since every $A_i \in \mathcal{F}$ is the vertex set of a rooted sub-branching, $v \in P_x$ by definition.
- CASE 2 ($v \prec x$): By condition 1 and an easy inductive argument, there exists A_i such that $x \in A_i, v \notin A_i$. Therefore, $v \notin F_x$.
- CASE 3 ($x \not\prec v, v \not\prec x$): By condition 2, there exists A_i such that $x \in A_i$ and $v \notin A_i$. Therefore, in combination with the definition of F_x we obtain $v \notin F_x$.

In all three cases, we obtain a contradiction and therefore $v \in F_x \Rightarrow v \in P_x$, showing that $F_x = P_x$. Given this fact, it is easy to reconstruct the branching T . We sketch the algorithm: compute for each x the set F_x and from F_x reconstruct the ordered version of the path P_x , i.e., $(r \rightarrow v_1 \rightarrow \dots \rightarrow x)$ using sets in \mathcal{F} whose existence is guaranteed by condition 1. \square

A natural question is whether one can extend Theorem 1 to more complex classes of oncogenetic models, such as directed acyclic graphs (DAGs) [12]. The answer is negative. For instance, the oncogenetic tree with edges $1 \rightarrow 2, 1 \rightarrow 3$ and the DAG with edges $1 \rightarrow 2, 1 \rightarrow 3, 2 \rightarrow 3$ are indistinguishable. Another example is shown in Figure 3.

Oncogenetic tree reconstruction algorithms [8] and more generally established inter-tumor cancer progression methods [12] receive as input a family of sets, where each set represents the set of mutations observed in a *single* tumor. Recall, that in the case of intra-tumor cancer progression the typical input is a multiset of points in the positive integer cone where each point is the copy-gene number state of a given cell. For instance, for the tumor shown in the final step of tumorigenesis in Figure 1, the input would be $\mathcal{D} = \{(2, 2) \times 2, (3, 2) \times 5, (2, 1) \times 3, (1, 1) \times 1\}$.




Fig. 3: Extending Theorem 1 to oncogenetic directed acyclic graphs (DAGs) [12] is not possible. For instance, the oncogenetic tree and the DAG are indistinguishable in terms of the generated families of sets.

Based on the insights from the proof of Theorem 1, we convert a FISH dataset to a dataset suitable for inter-tumor cancer progression inference. Specifically, we assume we are given a FISH dataset and an algorithm f which infers an evolutionary model of cancer progression from several tumors. Notice that f could be any inter-tumor phylogenetic method, see [12].

For each cell in state $x = (x_1, \dots, x_g)$ we generate a family of sets of “mutations” as follows: for every gene $i \in [g]$ whose number of copies x_i is greater than 2 we generate $x_i - 1$ sets in order to enforce that the gene has $c + 1$ copies if at a previous stage had c copies, $c = 2, \dots, x_i - 1$. For instance if gene i has 4 copies, we generate the sets $\{gene - i - mut - 2\}$, $\{gene - i - mut - 2, gene - i - mut - 3\}$, $\{gene - i - mut - 2, gene - i - mut - 3, gene - i - mut - 4\}$ which show how the gene obtained 2 extra copies. The case $x_i < 2$ is treated in a similar way. Finally, we generate a set s_f for each cell which contains all mutations that led to state x . For instance, for the state $(0, 3)$ $s_f = \{gene - 1 - mut - 2, gene - 1 - mut - 1, gene - 1 - mut - 0, gene - 2 - mut - 2, gene - 2 - mut - 3\}$. Upon creating the dataset \mathcal{F} we use it as input to $f(\cdot)$, an existing intra-tumor phylogenetic method, see [12].

Using the conversion above, based on condition 1 of Theorem 1 the output of an inter-tumor phylogenetic method will capture the dynamic nature of the process, which will be consistent with our assumptions of single gene duplication and loss events and Ockham’s razor, e.g., the evolutionary sequence $2 \rightarrow 3 \rightarrow 4$ rather than $2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 4$.

2.3 Progression Inference

Define $\mathcal{B} = [\min_{i \in [n]} x_{i1}, \max_{i \in [n]} x_{i1}] \times \dots \times [\min_{i \in [n]} x_{ig}, \max_{i \in [n]} x_{ig}]$ to be the minimum enclosing box of \mathcal{D} , where x_{ij} is the number of copies of gene j in the i -th cell, $i \in [n]$, $j \in [g]$. Given the observed data we can calculate the empirical probability $\tilde{\pi}(s)$ of any state $s \in \mathcal{B}$ as the fraction $\frac{|\{q: q \in \mathcal{D}, q=s\}|}{n}$. The number of states in \mathcal{B} grows exponentially fast for any typical FISH dataset. We summarize parsimoniously this distribution as described in Section 2.1. Specifically,

we learn the parameters w of the hierarchical log-linear model by maximizing the overlapping l_1 penalized log-likelihood of equation (2) as described in [26]. We allow only second-order interactions between factors. It is worth mentioning that k -way interactions, $k \geq 3$ can be embedded in the model as well, see Chapter 6 [26], but we prefer not to avoid overfitting. Alternatively we can allow higher order interactions but then a penalty term for the model complexity (e.g., AIC, BIC) should be taken into account. Let π be the distribution specified by the learned parameters. We define a Metropolis-Hastings chain with stationary distribution π [17]. Initially, all n cells will be in the diploid state $(2, \dots, 2)$. Notice that all we need to compute during the execution of the chain are ratios of the form $\pi(x)/\pi(y)$, which saves us from the computational cost of computing the normalization constant Z . We simulate the chain k times in order to draw $m \geq 1$ samples from the probability distribution. Finally we use the conversion described in Section 2.2 to infer a tumor phylogeny.

There exists a subtle issue that arises in practice: there exist states of \mathcal{B} which are not observed in the dataset \mathcal{D} . We surpass this problem by adding one fictitious sample to each state $b \in \mathcal{B}$. From a Bayesian point of view this is equivalent to smoothing the data with an appropriately chosen Dirichlet prior.

To summarize, our proposed method consists of the following steps: (1) Given a FISH dataset \mathcal{D} we learn the parameters of a hierarchical log-linear model with pairwise potentials. (2) Given the learned parameters we can compute the probability distribution on Z^g . Let π be the resulting distribution. We define a Metropolis-Hastings chain with stationary distribution π . Initially cells are in the healthy diploid state $(2, \dots, 2)$. (3) Draw $m \geq 1$ samples from the probability distribution by running the Metropolis-Hastings chain simulation m times. (4) Convert the resulting FISH samples to inter-tumor phylogenetic datasets by following the procedure of Section 2.2. (5) Use an inter-tumor phylogenetic method [12] to infer a tumor phylogenetic tree.

Finally, an interesting perspective on our modeling which makes a conceptual connection to [5] is the following: upon learning the parameters of the hierarchical log-linear model, the probability distribution over \mathcal{B} assigns implicitly weights on the edges of the positive integer di-grid \mathbb{Z}^g (each undirected edge of the grid is substituted by two directed edges) according to the Metropolis-Hastings chain. Therefore, both our method and [5] assign weights to the edges of the positive integer grid. This perspective opens two natural research directions which we leave open for future research. First, instead of simulating the Markov chain, one could find an appropriate subgraph of the weighted di-grid, e.g., a maximum weighted branching rooted at the diploid state. Secondly, it is natural to ask whether there exists a natural probabilistic interpretation of the method in [5].

3 Experimental Results

Experimental Setup: In this paper, we show the results of validating our method on a breast cancer dataset from a collection of publicly available FISH datasets which can be found at <ftp://ftp.ncbi.nlm.nih.gov/pub/FISHtrees/data/>.

We used the following third-party publicly available code in our implementations: hierarchical log-linear fitting code⁴ [26], distance based oncogenetic trees [9], FISH progression trees [5] and *graphviz* for visualization purposes. Our routines are implemented in MATLAB. The number of simulations was set to $m = 10$. We experimented both with smaller values for parameter m ($m \geq 2$) and the choice of log-linear fitting method, see [6], and we found that our results are robust.

Table 2 provides a short description of the six genes that are analyzed in our trees. The breast tumor dataset consists of 187 points in \mathbb{Z}_+^6 .

Gene	Cytogenetic Band	Description
<i>cox-2</i>	1q25.2-q25.3	oncogene
<i>myc</i>	8q24	oncogene
<i>ccnd1</i>	11q13	oncogene
<i>cdh1</i>	16q22.1	tumor suppressor gene
<i>p53</i>	17p13.1	tumor suppressor gene
<i>znf217</i>	20q13.2	oncogene

Table 1: Genes are shown in the first column and their cytogenetic positions in the second. The third column describes whether a gene is an oncogene or a tumor suppressor gene.

Results and Analysis: Figures 4 and 5 show the cancer phylogenetic trees of a ductal carcinoma *in situ* (DCIS) obtained by our method and [5] respectively. Our tree is a distance based phylogenetic tree produced by using our reduction of intra-tumor phylogenetic inference to inter-tumor phylogenetic inference as described in Section 2.3. We observe that the tree of [5] does not explain the 27 different states that appear in the dataset. For instance the state (4, 8, 4, 4, 2, 2, 4) accounts for 0.0053% of the appearing states and is discarded by [5] but is taken into account by our method. Since there is no ground truth available to us it is hard to reach any indisputable conclusion. However, we found that our findings are strongly supported by oncogenetic literature.

The mutational events captured in our phylogenetic tree highlight putative sequential events during progression from ductal carcinoma in situ (DCIS) to invasive breast carcinoma. The first mutational events are highlighted in red. Initially one allele of *p53* and *cdh1* are lost. Concurrent loss of *cdh1* function and *p53* inactivation act synergistically in the formation, progression and metastasis of breast cancer [7]. Moreover, following the first mutational events, the next

⁴ Schmidt’s code does not scale well to more than 6-7 variables.




Fig. 4: (best viewed on screen) Cancer phylogenetic tree for a breast cancer tumor obtained by our method. Leaves of the tree are colored in the following way: the root-event leaf is in yellow; the first change in the copy number profile is in red color; euploid states are in green; states with copy number gain/loss are shown in blue/orange. The first changes are losses of one gene copy of genes *p53* and *cdh1*.

changes occur in *ccnd1*, *myc* and *znf217* which are oncogenes participating in cell cycle regulation, proliferation and cancer progression. Specifically, as shown by single invasive ductal carcinoma (IDC) cell analysis copy number loss of *cdh1* is common in DCIS. Furthermore, copy number gains of *myc* are a common feature in the transition from DCIS to IDC [13]. This is consistent with our results. The synergy between *p53* and *cdh1* appears also in the tree of Figure 5 but at the last stages of the progression. Finally, Figure 6 shows the inferred direct associations among genes.

4 Conclusion

In this work we develop a novel approach to studying FISH datasets, a type of dataset which has received considerably less attention to other types of cancer datasets. Compared to prior work we take a probabilistic approach which provides good data fit, avoids overfitting and captures complex dependencies among factors. Motivated by our intention to capitalize on inter-tumor phylogenetic methods we prove a theorem which provides necessary and sufficient conditions for reconstructing oncogenetic trees [8]. Using these conditions, we show one way to perform intra-tumor phylogenetic inference by opening the door to the wealth of established inter-tumor phylogenetic techniques [12]. We model the evolutionary dynamics as a Markov chain in the positive integer cone

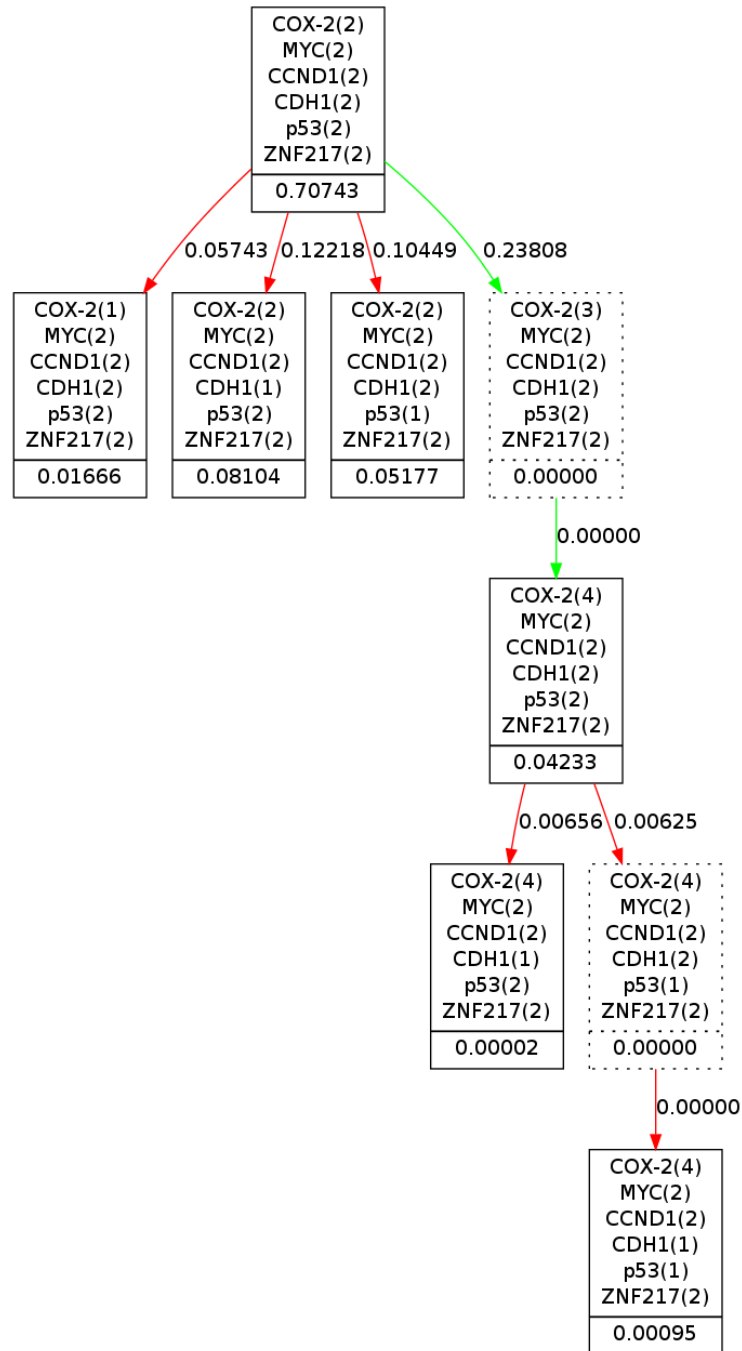


Fig. 5: Cancer phylogenetic tree for a breast cancer tumor obtained by [5]. Nodes with dotted borders represent Steiner nodes, i.e., states that do not appear in the dataset. Green and red edges model gene gain and loss respectively. The weight value on each edge does not have the semantics of probability, but it is the rectilinear distance between the two connected states. See [5] for further details. The weight on each node describes the fraction of cells in the FISH dataset with the particular copy number profile.

\mathbb{Z}_+^g where g is the number of genes examined with FISH. Finally, we validate our approach to a breast cancer FISH dataset.




Fig. 6: Direct associations among genes inferred from fitting a hierarchical log-linear model to a ductal carcinoma *in situ* FISH dataset.

Our work leaves numerous problems open for future research. Improved models need to be developed that remove the simplifying assumption of a fixed cell population and take the clonal evolution model into account [23], namely cancer is initiated once multiple mutations occur in a random single cell which gives birth to the uncontrolled proliferation of cancerous cells. Secondly, clustering patients and finding consensus FISH progression trees per cluster is another interesting problem. Furthermore, we plan to experiment with (a) other choices of inter-tumor phylogenetic methods and (b) fitting approaches that allow higher order interactions but will also account for the increased complexity of the resulting model. Finally, using features from our inferred trees as features for classification is an interesting question, see [5].

5 Acknowledgments

The author would like to thank Prof. Russell Schwartz, Dr. Stanley Shackney and Maria Tsiarli for helpful discussions on intratumor heterogeneity and the anonymous reviewers for their constructive feedback.

References

1. Beerenwinkel, N., Eriksson, N., Strumfels, B.: *Conjunctive bayesian networks*. Bernoulli, Vol. 13, pp. 893909 (2007)
2. Beerenwinkel, N., Rahnenführer, J., Däumer, M., Hoffmann, D., Kaiser, R., Selbig, J., Lengauer, T.: *Learning multiple evolutionary pathways from cross-sectional data*. Journal of Computational Biology, Vol. 12, pp. 584598 (2005)

3. Beerenwinkel, N., Sullivant, S.: *Markov models for accumulating mutations*. Biometrika, Vol. 96, pp. 663-676 (2009)
4. Bishop, Fienberg, S., Holland, P.: *Discrete Multivariate Analysis*. MIT Press (1975)
5. Chowdhury S.A. et al.: *FISHtrees: Modeling tumor progression from fluorescence in situ hybridization (FISH) data from many single cells of solid tumors and their metastases*. ISMB 2013
6. Dellaportas, P., Forster, J.: *Markov Chain Monte Carlo Model Determination for Hierarchical and Graphical Log-linear Models*. Biometrika (1996)
7. Derksen, PW., Liu, X., Saridin, F. et al.: *Somatic inactivation of E-cadherin and p53 in mice leads to metastatic lobular mammary carcinoma through induction of anoikis resistance and angiogenesis*, Cancer Cell, Vo.10(5), pp.437-449 (2006)
8. Desper, R., et al.: *Inferring tree models for oncogenesis from comparative genome hybridization data*. Journal of Computational Biology, Vol. 6(1), pp. 37-51 (1999)
9. Desper, R., et al.: *Distance-based reconstruction of tree models for oncogenesis*. J. Comput. Biol., Vol. 7(6), pp. 789-803 (2000)
10. Gasco, M., Shami, S., Crook, T.: *The p53 pathway in breast cancer* Breast Cancer Res., Vol.54 (2), pp.70-76 (2002)
11. Gerstung, M., Baudis, M., Moch, H., Beerenwinkel, N.: *Quantifying cancer progression with conjunctive bayesian networks*. Bioinformatics, Vol. 25, pp. 2809-2815 (2009)
12. Hainke, K., Rahnenführer, J., Fried, R.: *Disease progression models: A review and comparison*. Dortmund University, Technical Report, Available at <http://tinyurl.com/ceyr9wx>
13. Heselmeyer-Haddad, K., et. al.: *Single-cell genetic analysis of ductal carcinoma in situ and invasive breast cancer reveals enormous tumor heterogeneity yet conserved genomic imbalances and gain of MYC during progression*. Journal of American Pathology, 181(5), pp. 1807-1822 (2012)
14. Heydebreck, A., Gunawan, B., Füzesi, L.: *Maximum likelihood estimation of oncogenetic tree models*. Biostatistics, Vol. 5(4), pp. 545-556 (2004)
15. Krig, SR., Miller, JK., Fritze, S., et al.: *ZNF217, a candidate breast cancer oncogene amplified at 20q13, regulates expression of the ErbB3 receptor tyrosine kinase in breast cancer cells.*, Oncogene, Vol.29(40), pp.5500-5510 (2010)
16. Letouzé, E., Allory, E., Bollet, M. et al.: *Analysis of the copy number profiles of several tumor samples from the same patient reveals the successive steps in tumorigenesis*. Genome Biology Vol. 11:R76 (2010)
17. Levin, D., Peres, Y., Wilmer, E.: *Markov Chains and Mixing Times*. American Mathematical Society (2008)
18. Lin, SY., Xia, W., Wang, JC. et al.: *β -catenin, a novel prognostic marker for breast cancer: its roles in cyclin D1 expression and cancer progression*. Proc Natl Acad Sci., Vol.97(8) pp.4262-4266 (2000)
19. Martins, F. et al.: *Evolutionary pathways in BRCA1-associated breast tumors*, Cancer discovery, Vol. 2(6), pp. 503-511 (2012)
20. Marusyk, A. and Polyak, K.: *Tumor heterogeneity: causes and consequences*. Biochimica et Biophysica Acta (BBA)-Reviews on Cancer, 1805(1), pp. 105-117 (2010)
21. Navin, N., Krasnitz, A., Rodgers, L. et al.: *Inferring tumor progression from genomic heterogeneity*. Genome Res., Vol. 20, pp. 68-80 (2010)
22. Navin, N. et al.: *Tumour evolution inferred by single-cell sequencing*. Nature, Vol. 472, pp. 90-94 (2011)
23. Nowell, P.C.: *The clonal evolution of tumor cell populations*. Science, Vol. 194, pp. 23-28 (1976)

24. Pennington, G., Smith, C., Shackney, S., Schwartz, R.: *Reconstructing tumor phylogenies from heterogeneous single-cell data*. Journal of Bioinformatics and Computational Biology, Vol. 5(02A), pp. 407-427 (2007)
25. Pennington, G., Shackney, S., Schwartz, R.: *Cancer Phylogenetics from Single-Cell Assays*. Technical Report CMU-CS-06-103, available at <http://tinyurl.com/bvjlgch>
26. Schmidt, M.: *Graphical Model Structure Learning with l_1 -Regularization*. Ph.D. Thesis, University of British Columbia (2010)
27. Schmidt, M., Murphy, K.: *Convex Structure Learning in Log-Linear Models: Beyond Pairwise Potentials*. AISTATS (2010)