



Inferring Tree Causal Models of Cancer Progression with Probability Raising

Loes Olde Loohuis^{1*}, Giulio Caravagna^{2,3}, Alex Graudenzi^{2,3}, Daniele Ramazzotti^{2,3}, Giancarlo Mauri², Marco Antoniotti², Bud Mishra³

1 Center for Neurobehavioral Genetics, University of California Los Angeles, Los Angeles, United States of America, **2** Dipartimento di Informatica Sistemistica e Comunicazione, Università degli Studi, Milano-Bicocca, Milano, Italy, **3** Courant Institute of Mathematical Sciences, New York University, New York, United States of America

Abstract

Existing techniques to reconstruct tree models of progression for accumulative processes, such as cancer, seek to estimate causation by combining correlation and a frequentist notion of temporal priority. In this paper, we define a novel theoretical framework called CAPRESE (CAnCER PRogression EXtraction with Single Edges) to reconstruct such models based on the notion of probabilistic causation defined by Suppes. We consider a general reconstruction setting complicated by the presence of noise in the data due to biological variation, as well as experimental or measurement errors. To improve tolerance to noise we define and use a shrinkage-like estimator. We prove the correctness of our algorithm by showing asymptotic convergence to the correct tree under mild constraints on the level of noise. Moreover, on synthetic data, we show that our approach outperforms the state-of-the-art, that it is efficient even with a relatively small number of samples and that its performance quickly converges to its asymptote as the number of samples increases. For real cancer datasets obtained with different technologies, we highlight biologically significant differences in the progressions inferred with respect to other competing techniques and we also show how to validate conjectured biological relations with progression models.

Citation: Loohuis LO, Caravagna G, Graudenzi A, Ramazzotti D, Mauri G, et al. (2014) Inferring Tree Causal Models of Cancer Progression with Probability Raising. PLoS ONE 9(10): e108358. doi:10.1371/journal.pone.0108358

Editor: Lars Kaderali, Technische Universität Dresden, Medical Faculty, Germany

Received: April 11, 2014; **Accepted:** August 27, 2014; **Published:** October 9, 2014

Copyright: © 2014 Olde Loohuis et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability: The authors confirm that all data underlying the findings are fully available without restriction. All data are included within the paper.

Funding: This work was supported by the National Science Foundation grants CCF-0836649 and CCF-0926166 and by Regione Lombardia (Italy) under the research projects RetroNet through the ASTIL [12-4-5148000-40]; U.A 053 and NEDD Project [ID14546A Rif SAL-7] Fondo Accordi Istituzionali 2009. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* Email: loldeloohuis@gmail.com

These authors contributed equally to this work.

Introduction

Cancer is a disease of evolution. Its initiation and progression are caused by dynamic somatic alterations to the genome manifested as point mutations, structural alterations, DNA methylation and histone modification changes [1].

These genomic alterations are generated by random processes, and since individual tumor cells compete for space and resources, the fittest variants are naturally selected for. For example, if through some mutations a cell acquires the ability to ignore anti-growth signals from the body, this cell may thrive and divide, and its progeny may eventually dominate some part(s) of the tumor. This *clonal expansion* can be seen as a *discrete state* of the cancer's progression, marked by the acquisition of a set of genetic events. Cancer progression can then be thought of as a sequence of these discrete steps, where the tumor acquires certain distinct properties at each state. Different progression sequences are possible, but some are more common than others, and not every order is viable [2].

In the last two decades, many specific genes and genetic mechanisms that are involved in different types of cancer have been identified (see e.g. [3,4] for an overview of common cancer

genes and [5,6] for specific genetic analyses of ovarian carcinoma and lung adenocarcinoma, respectively), and *therapies* targeting the activity of these genes are now being developed at a fast pace [2]. However, unfortunately, the *causal and temporal relations* among the genetic events driving cancer progression remain largely elusive.

The main reason for this state of affairs is that information revealed in the data is usually obtained only at one (or a few) points in time, rather than over the course of the disease. Extracting this dynamic information from the available *cross-sectional* data is challenging, and a combination of mathematical, statistical and computational techniques is needed. In recent years, several methods to extract progression models from cross-sectional data have been developed, starting from the seminal work on single-path-models by Fearon and Vogelstein [7]. In particular, different models of oncogenetic trees were developed over the years. At the core of some of these methods, e.g. [8,9], is the use of *correlation* to identify relations among genetic events. These techniques reconstruct *tree* models of progression as independent acyclic paths with branches and no confluences. Distinct models of oncogenetic trees are instead based on *maximum likelihood estimation*, e.g., [10,11,12]. More general *Markov chain* models, e.g., [13],

describe more flexible probabilistic networks, despite the computationally expensive parameter estimation. Other recent models are Conjunctive Bayesian Networks, CBNs [14,15], that extract *directed acyclic graphs*, yet imposing specific constraints on the joint occurrence of events. Finally, in a slightly different context, temporal models were reconstructed from time-course gene expression data [16,17].

In this paper we present a novel theoretical framework called CAPRESE (CAncer PRogression Extraction with Single Edges) to reconstruct cumulative progressive phenomena, such as cancer progression. We assume the original problem setting of [8], and propose a new a technique to infer *probabilistic progression trees* from cross-sectional data. Unlike maximum likelihood estimation-based techniques, our goal is the extraction of the *minimal* progression model explaining the order in which mutations occur and accumulate. The method is technology agnostic, i.e., it can be applied to dataset derived from all types of (epi-)genetic data such as deep exome sequencing, bisulfite sequencing, SNP arrays, etc., (see Results), and takes as input a set of pre-selected genetic events of which the presence or the absence of each event is recorded for each sample.

CAPRESE is based on two main ingredients: 1) rather than using *correlation* to infer progression structures, we base our technique on a notion of *probabilistic causation*, and 2), to increase robustness against noise, we adopt a *shrinkage-like estimator* to measure causation among any pair of events. More specifically, with respect to our first ingredient, we adopt the notion of (prima facie) causation proposed by Suppes in [18]. Its basic intuition is simple: event *a* causes event *b* if (i) *a* occurs *before b* and (ii) the occurrence of *a* *raises the probability* of observing *b*. This is a very basic notion of probabilistic causation that in itself does not address many of the problems associated with it (such as asymmetry, common causes, and screening off [19]), and includes *spurious* as well as *genuine* causes. However, as it turns out, this basic notion combined with a filter for independent progressions starting from the same root, is an excellent tool to guide progression extraction from cross-sectional data – one that outperforms the commonly used correlation-based methods.

Probabilistic causation was used in biomedical applications before (e.g., to find driver genes from CNV data in [20], and to extract causes from biological time series data in [21]), but, to the best of our knowledge, never to infer *progression models* in the *absence* of direct temporal information.

The extraction problem is complicated by the presence of both false positive and false negative observations (see [22] for a discussion on this issue based on the reconstruction by [8]), such as the one provided by the intrinsic variability of biological processes (e.g., *genetic heterogeneity*) and *experimental errors*. This poses a problem, because while probability raising is a very precise tool, it, by itself, is not robust enough against noise. Conditional on the amount of noise, we will rely both on probabilistic causation and on a more robust (but less precise) correlation-based metric in an optimal way. To do this we introduce our second ingredient, a *shrinkage-like estimator* to measure causation among any pair of events. The intuition behind this estimator, which is closely related to a shrinkage estimator from [23], is to find the optimal balance between probability raising on the one hand and correlation on the other, depending on the amount of noise.

We prove correctness of our algorithm by showing that with increasing sample sizes, the reconstructed tree asymptotically converges to the correct one (Theorem 3). Under mild constraints on the noise rates, this result holds for the reconstruction problem in the presence of uniform noise as well.

We also study the performance of CAPRESE in more realistic settings with limited sample sizes. Using synthetic data, we show that under these conditions, our algorithm outperforms the state-of-the-art tree reconstruction algorithm of [8] (see Results). In particular, our shrinkage-like estimator provides, on average, an increased robustness to noise which ensures it to outperform oncotrees [8]. Performance is defined in terms of *structural similarity* between the reconstructed tree and the actual tree, rather than on their induced distribution as is done, e.g., in [11]. This metric is especially appropriate for the goal of reconstructing a progression model where data-likelihood fit is secondary to “calling” the possibly minimal set of causal relations.

Also, we show that CAPRESE works well already with a relatively low number of samples and that its performance quickly converges to its asymptote as the number of samples increases. This outcome hints at the applicability of the algorithm with relatively small datasets without compromising its efficiency.

We remark that further analyses on synthetic data suggests that CAPRESE outperforms a well known bayesian probabilistic graphical model as well (i.e., *Conjunctive Bayesian Networks* [14,15]), which was originally conceived for the reconstruction of more complex topologies, e.g. DAGs, but was proven effective in reconstructing tree topologies as well [24] (see Results).

Finally, we apply our technique to alterations assessed with both Comparative Genomic Hybridization and Next Generation Sequencing techniques (see Results). In the former case, we show that the algorithm of [8] and CAPRESE highlight biologically important differences in ovarian, gastrointestinal and oral cancer, but our inferences are statistically more significant. In the latter, we validate a recently discovered relation among two key genes involved in leukemia.

Methods

Problem setting

The set-up of the reconstruction problem is as follows. Assuming that we have a set *G* of *n* mutations (*events*, in probabilistic terminology) and *s* samples, we represent a cross-sectional dataset as an *s* × *n* binary matrix in which an entry (*k,l*)=1 if the mutation *l* was observed in sample *k*, and 0 otherwise. The problem we solve in this paper is to extract a set of edges *E* yielding a progression *tree* $\mathcal{T} = (G \cup \{\blacklozenge\}, E, \blacklozenge)$ from this matrix which, we remark, only implicitly provides information of progression timing. The root of *T* is modeled using a (special) event $\blacklozenge \notin G$ such that *heterogenous progression paths* or *forests* can be reconstructed. More precisely, we aim at reconstructing a *rooted tree* that satisfies: (i) each node has at most one incoming edge, (ii) the root has no incoming edges (iii) there are no *cycles*.

Each progression tree subsumes a distribution of observing a subset of the mutations in a cancer sample that can be formalized as follows:

Definition 1. (Tree-induced distribution) *Let T be a tree and $\alpha : E \rightarrow [0,1]$ a labeling function denoting the independent probability of each edge, T generates a distribution where the probability of observing a sample with the set of alterations $G^* \subseteq G$ is*

$$\mathcal{P}(G^*) = \prod_{e \in E'} \alpha(e) \cdot \prod_{\substack{(u,v) \in E \\ u \in G^*, v \notin G^*}} [1 - \alpha(u,v)] \tag{1}$$

where all events in *G** are assumed to be reachable from the root \blacklozenge , and *E' ⊆ E* is the set of edges connecting the root to the events in *G**.

We would like to emphasize two properties related to the tree-induced distribution. First, the distribution subsumes that, given

any oriented edge ($a \rightarrow b$), an observed sample contains alteration b with probability $\mathcal{P}(a)\mathcal{P}(b)$, that is the probability of observing b after a . For this reason, if a causes b , the probability of observing a will be greater than the probability of observing b accordingly to the temporal priority principle which states that all causes must precede, in time, their effects [25].

Second, the input dataset is a set of samples generated, ideally, from an unknown distribution induced by an unknown tree or forest that we aim at reconstructing. However, in some cases, it could be that no tree exists whose induced distribution generates *exactly* those input data. When this happens, the set of observed samples slightly diverges from any tree-induced distribution. To model these situations a notion of *noise* can be introduced, which depends on the context in which data are gathered. Adding noise to the model complicates the reconstruction problem (see Results).

The oncotree approach. In [8] Desper *et al.* developed a method to extract progression trees, named “*oncotrees*”, from static CNV data. In [22] Szabo *et al.* extended the setting of Desper’s reconstruction problem to account for both *false positives* and *negatives* in the input data. In these oncotrees, nodes represent CNV events and edges correspond to possible progressions from one event to the next.

The reconstruction problem is exactly as described above, and each tree is rooted in the special event \blacklozenge . The choice of which edge to include in a tree is based on the estimator

$$w_{a \rightarrow b} = \log \left[\frac{\mathcal{P}(a)}{\mathcal{P}(a) + \mathcal{P}(b)} \cdot \frac{\mathcal{P}(a,b)}{\mathcal{P}(a)\mathcal{P}(b)} \right], \quad (2)$$

which assigns to each edge $a \rightarrow b$ a weight accounting for both the relative and joint frequencies of the events – thus measuring *correlation*. The estimator is evaluated after including \blacklozenge to each sample of the dataset. In this definition the rightmost term is the (symmetric) *likelihood ratio* for a and b occurring together, while the leftmost is the asymmetric *temporal priority* measured by rate of occurrence. This implicit form of timing assumes that, if a occurs *more often* than b , then it likely occurs *earlier*, thus satisfying

$$\frac{\mathcal{P}(a)}{\mathcal{P}(a) + \mathcal{P}(b)} > \frac{\mathcal{P}(b)}{\mathcal{P}(a) + \mathcal{P}(b)}.$$

An oncotree is the rooted tree whose total weight (i.e., sum of all the weights of the edges) is maximized, and can be reconstructed in $O(|G|^2)$ steps using Edmond’s algorithm [26]. By construction, the resulting graph is a proper tree rooted in \blacklozenge : each event occurs only once, *confluences* are absent, i.e., any event is caused by at most one other event. This method has been used to derive progressions for various cancer datasets e.g., [27,28,29]), and even though several methods that extend this framework exists (e.g. [9,11,15]), to the best of our knowledge, it is currently the only method that aims to solve exactly the same problem as the one investigated in this paper and thus provide a benchmark to compare against.

A probabilistic approach to causation

We briefly review the approach to probabilistic causation, on which our method is based. For an extensive discussion on this topic we refer to [19].

In his seminal work [18], Suppes proposed the following notion.

Definition 2. (Probabilistic causation, [18]). *For any two events c and e , occurring respectively at times t_c and t_e , under the mild*

assumptions that $0 < \mathcal{P}(c), \mathcal{P}(e) < 1$, the event c is a prima facie cause of the event e if it occurs before the effect and the cause raises the probability of the effect, i.e.,

$$t_c < t_e \quad \text{and} \quad \mathcal{P}(e|c) > \mathcal{P}(e|\bar{c}). \quad (3)$$

As discussed in [19] the above conditions are not, in general, sufficient to claim that event c is a cause of event e . In fact a prima facie cause is either *genuine* or *spurious*. In the latter case, the fact that the conditions hold in the observations is due either to coincidence or to the presence of a certain third *confounding factor*, related both to c and to e [18]. Genuine causes, instead, satisfy Definition 2 and are not screened-off by any confounding factor. However, they need not be direct causes. See Figure 1.

Note that we consider cross-sectional data where no information about t_c and t_e is available, so in our reconstruction setting we are restricted to consider solely the *probability raising* (PR) property, i.e., $\mathcal{P}(e|c) > \mathcal{P}(e|\bar{c})$, which makes it harder to discriminate among genuine and spurious causes. Now we review some of its properties.

Proposition 1. (Dependency). *Whenever the PR holds between two events a and b , then the events are statistically dependent in a positive sense, i.e.,*

$$\mathcal{P}(b|a) > \mathcal{P}(b|\bar{a}) \Leftrightarrow \mathcal{P}(a,b) > \mathcal{P}(a)\mathcal{P}(b). \quad (4)$$

This and the next proposition are well-known facts of the PR; their derivation as well as the proofs of all the results we present is in the File S1. Notice that the opposite implication holds as well: when the events a and b are still dependent but in a negative sense, i.e., $\mathcal{P}(a,b) < \mathcal{P}(a)\mathcal{P}(b)$, the PR does not hold, i.e., $\mathcal{P}(b|a) < \mathcal{P}(b|\bar{a})$.

We would like to use the asymmetry of the PR to determine whether a pair of events a and b satisfy a causation relation so to place a before b in the progression tree but, unfortunately, the PR satisfies the following property.

Proposition 2. (Mutual PR). $\mathcal{P}(b|a) > \mathcal{P}(b|\bar{a}) \Leftrightarrow \mathcal{P}(a|b) > \mathcal{P}(a|\bar{b})$.

That is, if a raises the probability of observing b , then b raises the probability of observing a too.

Nevertheless, in order to determine causes and effects among the genetic events, we can use our *degree of confidence* in our estimate of probability raising to decide the direction of the causation relationship between pairs of events. In other words, if a raises the probability of b *more* than the other way around, then a is a more likely cause of b than b of a . Notice that this is sound as long as each event has *at most* one cause; otherwise, *frequent late events* with more than one cause, which are rather common in biological progressive phenomena, should be treated differently. As mentioned, the PR is not symmetric, and the *direction* of probability raising depends on the relative frequencies of the events. We make this asymmetry precise in the following proposition.

Proposition 3. (Probability raising and temporal priority). *For any two events a and b such that the probability raising $\mathcal{P}(a|b) > \mathcal{P}(a|\bar{b})$ holds, we have*

$$\mathcal{P}(a) > \mathcal{P}(b) \Leftrightarrow \frac{\mathcal{P}(b|a)}{\mathcal{P}(b|\bar{a})} > \frac{\mathcal{P}(a|b)}{\mathcal{P}(a|\bar{b})}. \quad (5)$$

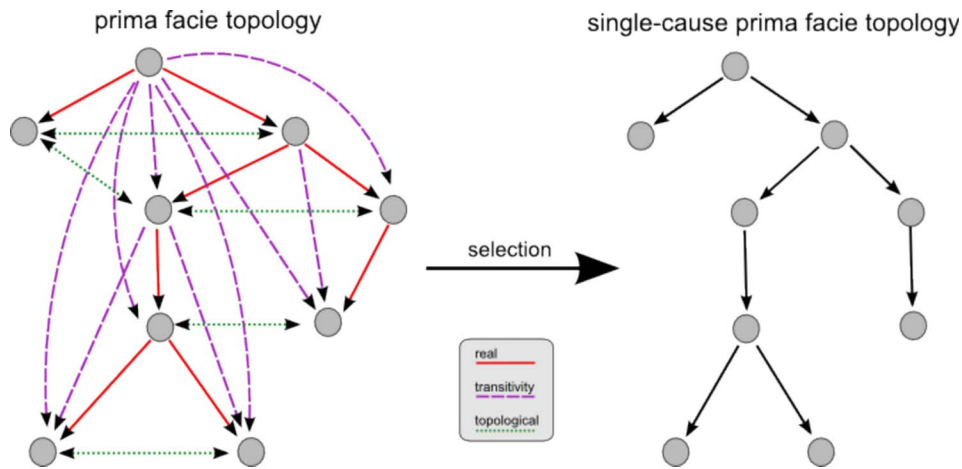


Figure 1. Prima facie topology. Example prima facie topology where all edges (a,b) represent prima facie causes, according to Definition 3: a is a probability raiser of b and it occurs more frequently. In left, we filter out spurious causes and select only the real ones among the genuine, yielding a single-cause prima facie topology.
doi:10.1371/journal.pone.0108358.g001

That is, given that the PR holds between two events, a raises the probability of b more than b raises the probability of a , if and only if a is observed more frequently than b . Notice that we use the ratio to assess the PR inequality. The proof of this proposition is technical and can be found in the File S1. From this result it follows that if we measure the timing of an event by the rate of its occurrence (that is, $\mathcal{P}(a) > \mathcal{P}(b)$ implies that a happens before b), this notion of PR subsumes the same notion of temporal priority induced by a tree. We also remark that this is also the temporal priority made explicit in the coefficients of Desper’s method. Given these results, we define the following notion of causation.

Definition 3. We state that a is a prima facie cause of b if a is a probability raiser of b , and it occurs more frequently:

$$\mathcal{P}(b|a) > \mathcal{P}(b|\bar{a}) \text{ and } \mathcal{P}(a) > \mathcal{P}(b).$$

We term *prima facie topology* a directed acyclic graph (over some events) where each edge represents a prima facie cause. When at most a single incoming edge is assigned to each event (i.e., an event has at most a *unique cause*, in the real world), we term this structure *single-cause prima facie topology*. Intuitively, this last class of topologies correspond to the trees or, more generally forests when they have disconnected components, that we aim at reconstructing.

Before moving on to introducing our algorithm let us discuss our definition of *causation*, its role in the definition of the reconstruction problem and some of its limitations. As already mentioned, it may be that for some prima facie cause c of an event e , there is a third event a prior to both, such that a causes c and ultimately c causes e . Alternatively, a may cause both c and e independently, and the causation relationship observed from c to e is merely *spurious*. In the context of the tree-reconstruction problem, namely when it is assumed that each event has at most a unique cause, the aim is to filter out the spurious edges from a general prima facie topology, so to extract a single-cause prima facie structure (see Figure 1).

Definition 3 summarizes Suppes basic notion of prima facie cause, while it is ignoring deeper discussions of causation that aim at distinguishing between actual genuine and spurious causes, e.g.

screening-off, background context, d-separation [30,31,19]. For our purposes however, the above definition is sufficient when (i) all the significant events are considered, i.e., all the genuine causes are observed as in a closed-world assumption, and (ii) we aim at extracting the *order* of progression among them (or determine that there is no apparent relation), rather than extracting causalities *per se*. Note that these assumptions are strong and might be weakened in the future (see Discussions), but are shared by us and [8].

Finally, we recall a few algebraic requirements necessary for our framework to be well-defined. First of all, the PR must be computable: every mutation a should be observed with probability strictly $0 < \mathcal{P}(a) < 1$. Moreover, we need each pair of mutations (a,b) to be *distinguishable* in terms of PR, that is, for each pair of mutations a and b , $\mathcal{P}(ab) < 1$ or $\mathcal{P}(ba) < 1$ similarly to the above condition. Any non-distinguishable pair of events can be merged as a single composite event. From now on, we will assume these conditions to be verified.

Performance measure and synthetic datasets

We made use of *synthetic data* to evaluate the performance of CAPRESE as a function of dataset size and the false positive and negative rates. Many distinct synthetic datasets were created for this purpose, as explained below. The algorithm’s performance was measured in terms of *Tree Edit Distance* (TED, [32]), i.e., the minimum-cost sequence of node edit operations (relabeling, deletion and insertion) that transforms the reconstructed trees into the ones generating the data. The choice of this measure of evaluation is motivated by the fact that we are interested in the *structure* behind the progressive phenomenon of cancer evolution and, in particular, we are interested in a measure of the genuine causes that we miss and of the spurious causes that we fail to recognize (and eliminate). Also, since topologies with similar distributions can be structurally different we choose to measure performance using structural distance rather than a distance in terms of distributions. Within the realm of ‘structural metrics’ however, we have also evaluated the performance with the *Hamming Distance* [33], another commonly used structural metric, and we obtained analogous results (not shown here).

Synthetic data generation and experimental setting. Synthetic datasets were generated by sampling from various random trees constrained to have depth $\log(|G|)$, since

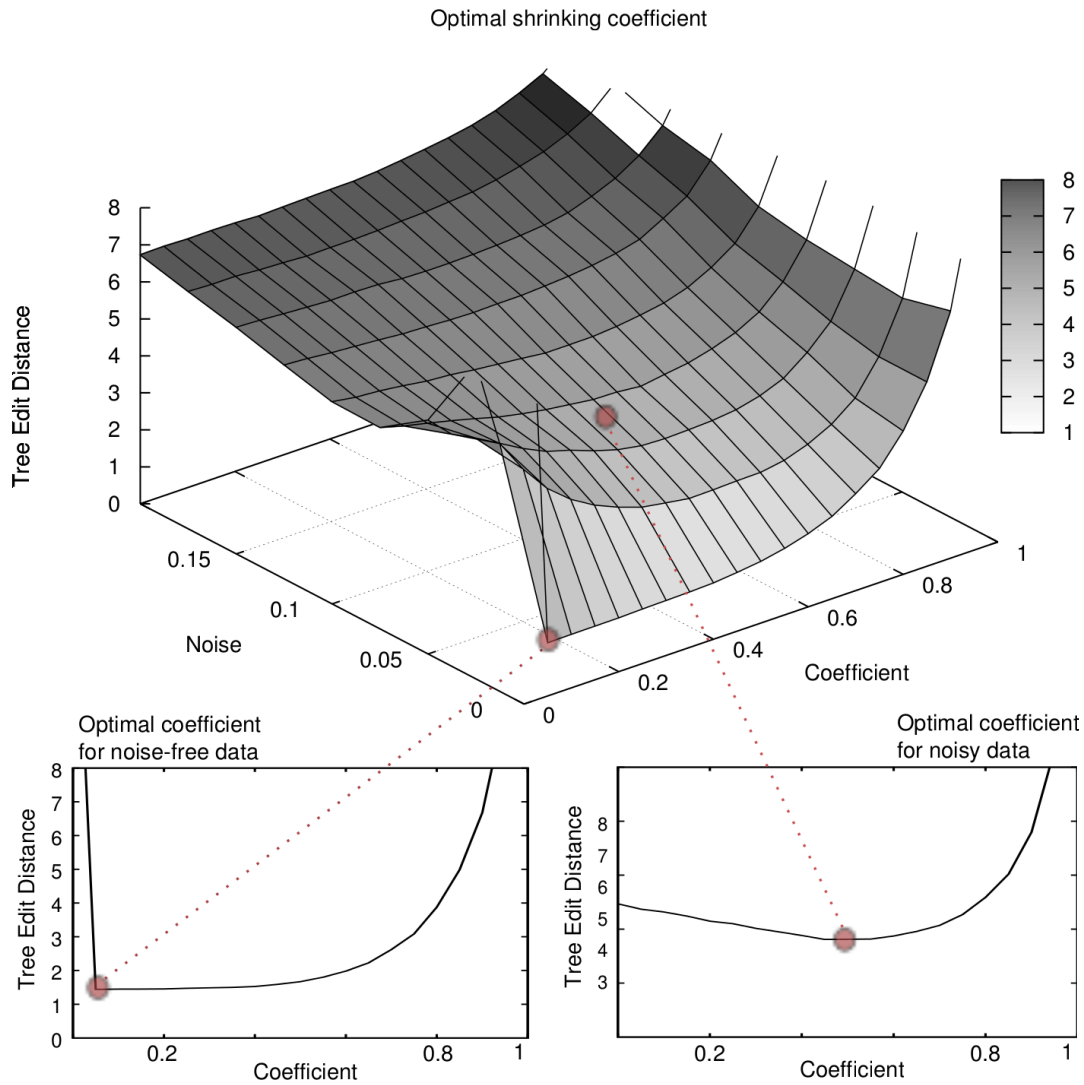


Figure 2. Optimal shrinkage-like coefficient for reconstruction performance. We show here the performance in the reconstruction of trees (TED surface) with $n = 150$ samples as a function of the shrinkage-like coefficient λ . Notice the global optimal performance for $\lambda \rightarrow 0$ when $v \rightarrow 0$ and for $\lambda \approx 1/2$ when $v > 0$. doi:10.1371/journal.pone.0108358.g002

wide branches are harder to reconstruct than straight paths, and by sampling event probabilities in $[0.05, 0.95]$ (see File S1).

Unless explicitly specified, in all the experiments we used 100 distinct random trees (or forests, accordingly to the test to perform) of 20 events each. This seems a fairly reasonable number of events and is in line with the usual size of reconstructed trees, e.g. [34,35,36,37]. The scalability of the techniques was tested against the number of samples by ranging $|G|$ from 50 to 250, with a step of 50, and by replicating 10 independent datasets for each parameters setting (see the caption of the figures for details).

We included a form of *noise* in generating the datasets, in order to account for (i) the realistic presence of *biological noise* (such as the one provided by bystander mutations, genetic heterogeneity, etc.) and (ii) *experimental errors*. A noise parameter $0 \leq v < 1$ denotes the probability that any event assumes a random value (with uniform probability), after sampling from the tree-induced distribution. Algorithmically this process generates, on average, $|G|v/2$ random entries in each sample (e.g. with $v=0.1$ we have, on average, one error per sample). We wish to assess whether these noisy samples can mislead the reconstruction process, even for low values of v . Notice that assuming a uniformly distributed noise

may appear simplistic since some events may be more robust, or easy to measure, than others. However, introducing in the data both *false positives* (at rate $\epsilon_+ = v/2$) and *negatives* (at rate $\epsilon_- = v/2$) makes the inference problem substantially harder, and was first investigated in [22].

In the Results section, we refer to datasets generated with rate $v > 0$ as noisy synthetic dataset. In the numerical experiments, v is usually discretized by 0.025, (i.e., 2.5% noise).

Results

Extracting progression trees with probability raising and a shrinkage-like estimator

The CAPRESE reconstruction method is described in Algorithm 1. The algorithm is similar to Desper's and Szabo's algorithm, the main difference being an alternative weight function based on a shrinkage-like estimator.

Algorithm 1. CAPRESE: a tree-like reconstruction with a shrinkage-like estimator.

1: consider a set of n genetic events G plus a special event \blacklozenge , added to each sample of the dataset;

2: define a $m \times n$ matrix M where each entry contains the shrinkage-like estimator $m_{i \rightarrow j} = (1 - \lambda) \frac{\mathcal{P}(j) - \mathcal{P}(j\bar{i})}{\mathcal{P}(j) + \mathcal{P}(j\bar{i})} + \lambda \frac{\mathcal{P}(i,j) - \mathcal{P}(i)\mathcal{P}(j)}{\mathcal{P}(i,j) + \mathcal{P}(i)\mathcal{P}(j)}$ according to the observed probability of the events i and j ;

3: [PR causation] define a tree $\mathcal{T} = (G \cup \{\blacklozenge\}, E, \blacklozenge)$ where $(i \rightarrow j) \in E$ for $i, j \in G$ if and only if: $m_{i \rightarrow j} > 0$ and $m_{i \rightarrow j} > m_{j \rightarrow i}$ and $\forall i' \in G, m_{i,j} > m_{i',j}$.

4: [Independent progressions filter] define $G_j = \{x \in G \mid \mathcal{P}(x) > \mathcal{P}(j)\}$, replace edge $(i \rightarrow j) \in E$ with edge $(\blacklozenge \rightarrow j)$ if, for all $x \in G_j$, it holds $\frac{1}{1 + \mathcal{P}(j)} > \frac{\mathcal{P}(x)}{\mathcal{P}(x) + \mathcal{P}(j)} \frac{\mathcal{P}(x,j)}{\mathcal{P}(x)\mathcal{P}(j)}$.

Definition 4. (Shrinkage-like estimator). We define the shrinkage-like estimator $m_{a \rightarrow b}$ of the confidence in the causation relationship from a to b as

$$m_{a \rightarrow b} = (1 - \lambda)\alpha_{a \rightarrow b} + \lambda\beta_{a \rightarrow b}, \tag{6}$$

where $0 \leq \lambda \leq 1$ and

$$\alpha_{a \rightarrow b} = \frac{\mathcal{P}(b|a) - \mathcal{P}(b|\bar{a})}{\mathcal{P}(b|a) + \mathcal{P}(b|\bar{a})} \quad \beta_{a \rightarrow b} = \frac{\mathcal{P}(a,b) - \mathcal{P}(a)\mathcal{P}(b)}{\mathcal{P}(a,b) + \mathcal{P}(a)\mathcal{P}(b)}. \tag{7}$$

This estimator is similar in spirit to a shrinkage estimator (see [23]) and combines a normalized version of PR, the *raw estimate* α , with a *correction factor* β (in our case a correlation-based measure of temporal distance among events), to define a proper order in the confidence of each causation relationship. Our λ is the analogous of the *shrinkage coefficient* and can have a Bayesian interpretation based on the strength of our belief that a and b are causally relevant to one another and the evidence that a raises the probability of b . In the absence of a closed form solution for the optimal value of λ , one may rely on cross-validation of simulated data. The power of shrinkage (and our shrinkage-like estimator) lies in the possibility of determining an optimal value for λ to balance the effect of the correction factor on the raw model estimate to ensure optimal performances on ill posed instances of the inference problem. A crucial difference, however, between our estimator and classical shrinkage, is that our estimator aims at improving the performance of the *overall* reconstruction process, not limited to the performance of the estimator itself as is the case in shrinkage. That is, the metric m induces an ordering to the events reflecting our confidence for their causation. Furthermore, since we make no assumption about the underlying distribution, we learn it empirically by cross-validation. In the next sections we show that the shrinkage-like estimator is an effective way to get such an ordering especially when data are noisy. In CAPRESE we use a pairwise matrix version of the estimator.

The raw estimator and the correction factor. By considering only the raw estimator α , we would include an edge $(a \rightarrow b)$ in the tree consistently in terms of (i) Definition 3 (Methods) and (ii) if a is the best probability raiser for b . When $\mathcal{P}(a) = \mathcal{P}(b)$ the events a and b are indistinguishable in terms of temporal priority, thus α is not sufficient to decide their causal relation, if any. This intrinsic ambiguity is unlikely in practice even if, in principle, it is possible. Notice that this formulation of α is a monotonic normalized version of the PR ratio.

Proposition 4. (Monotonic normalization). For any two events a and b we have

$$\mathcal{P}(a) > \mathcal{P}(b) \Leftrightarrow \frac{\mathcal{P}(b|a)}{\mathcal{P}(b|\bar{a})} > \frac{\mathcal{P}(a|b)}{\mathcal{P}(a|\bar{b})} \Leftrightarrow \alpha_{a \rightarrow b} > \alpha_{b \rightarrow a}. \tag{8}$$

This raw model estimator satisfies $-1 \leq \alpha_{a \rightarrow b}, \alpha_{b \rightarrow a} \leq 1$: when it tends to -1 the pair of events appear disjointly (i.e., they show an anti-causation pattern), when it tends to 0 no causation or anti-causation can be inferred and the two events are statistically independent and, when it tends to 1 , the causation relationship between the two events is genuine. Therefore, α provides a quantification of the degree of confidence for a PR causation relationship. In fact, for any given possible causation edge (a,b) , the term $\mathcal{P}(b|\bar{a})$ gives an estimate of the *error rate* of b , therefore the numerator of the raw model α provides an estimate of how often b is actually caused by a . The α estimator is then normalized to range between -1 and $+1$.

However, α does not provide a general criterion to disambiguate among genuine causes of a given event. We show a specific case in which α is not a sufficient estimator. Let us consider, for instance, a causal linear path: $a \rightarrow b \rightarrow c$. In this case, when evaluating the candidate parents a and b for c we have: $\alpha_{a \rightarrow c} = \alpha_{b \rightarrow c} = 1$, so a and b are genuine causes of c , though we would like to select b , instead of a . Accordingly, we can only infer that $t_a < t_c$ and $t_b < t_c$, i.e., a partial ordering, which does not help to disentangle the relation among a and b with respect to c .

In this case, the β coefficient can be used to determine which of the two genuine causes occurs closer, in time, to c (b , in the example above). In general, such a correction factor provides information on the *temporal distance* between events, in terms of statistical dependency. In other words, the higher the β coefficient, the closer two events are. Therefore, when dealing with noisy data and limited sample sizes, there are situations where, by using the α estimator alone, we could infer a wrong transitive edge to be the most likely cause even in the presence of the real cause. For this reason, we reduce the α estimator to the correction factor β , which, for each given edge (a,b) , is normalized within -1 and $(1 - \max[\mathcal{P}(a), \mathcal{P}(b)]) / (1 + \max[\mathcal{P}(a), \mathcal{P}(b)]) < +1$.

The shrinkage-like estimator m then results in the combination of the raw PR estimator α and of the β correction factor, which respects the temporal priority induced by α .

Proposition 5. (Coherence in dependency and temporal priority). The β correction factor is symmetrical and subsumes the same notion of dependency of the raw estimator α , that is

$$\mathcal{P}(a,b) > \mathcal{P}(a)\mathcal{P}(b) \Leftrightarrow \alpha_{a \rightarrow b} > 0 \Leftrightarrow \beta_{a \rightarrow b} > 0 \quad \text{and} \quad \beta_{a \rightarrow b} = \beta_{b \rightarrow a}. \tag{9}$$

The independent progressions filter. As in Desper's approach, we also add a *root* \blacklozenge with $\mathcal{P}(\blacklozenge) = 1$ in order to separate different progression paths, i.e., the different sub-trees rooted in \blacklozenge . CAPRESE initially builds a unique tree by using the estimator; typically, the most likely event will be at the top of the progression even if there may be rare cases where more than one event has no valid parent, in these cases we would already be reconstructing a forest. In the reconstructed tree, all the edges represent the most confident *prima facie* cause, although some of those could still be spurious causes. Then the correlation-like weight between any node j and \blacklozenge is computed as

$$\frac{\mathcal{P}(\blacklozenge) \mathcal{P}(\blacklozenge \rightarrow j)}{\mathcal{P}(\blacklozenge) + \mathcal{P}(j) \mathcal{P}(\blacklozenge \rightarrow j)} = \frac{1}{1 + \mathcal{P}(j)}$$

If this quantity is greater than the weight of j with each upstream connected element i , we consider the best prima facie cause of j to be a spurious cause and we substitute the edge ($i \rightarrow j$) with the edge ($\blacklozenge \rightarrow j$). Note that in this work we are ignoring deeper discussions of probabilistic causation that aim at distinguishing between actual genuine causes and spurious causes. Instead, we remove spurious causes by using a filter based on correlation because the probability raising of the omnipresent event \blacklozenge is not well defined (see Methods). In addition, we remark that the evaluation for an edge to be a genuine or a spurious cause takes into account all the given events. Because of this, if events are added or removed from the dataset, the same edge can be defined to be genuine or spurious as the set of events included in the model is varied arbitrarily. However, since we do not consider the problem of selecting the set of progression events, we assume that all and only the relevant events for the problem at hand are already known a priori and included in the model.

Finally, note that this filter is indeed implying a non-negative threshold for the shrinkage-like estimator, when an edge is valid.

Theorem 1. (Independent progressions). *Let $G^* = \{a_1, \dots, a_k\} \subset G$ a set of k prima facie causes for some $b \notin G^*$, and let $a^* = \max_{a_i \in G^*} \{m_{a_i \rightarrow b}\}$. The reconstructed tree by CAPRESE contains edge $\blacklozenge \rightarrow b$ instead of $a^* \rightarrow b$ if, for all $a_i \in G^*$*

$$\mathcal{P}(a_i, b) < \mathcal{P}(a_i) \mathcal{P}(b) \frac{1}{1 + \mathcal{P}(b)} + \frac{\mathcal{P}(b)^2}{1 + \mathcal{P}(b)} \quad (10)$$

The proof of this theorem can be found in the File S1. What this theorem suggests is that, in principle, by examining the level of statistical dependency of each pair of events, it would be possible to determine how many trees compose the reconstructed forest. Furthermore, it suggests that CAPRESE could be defined by first processing the independent progressions filter, and then using m to build the independent progression trees in the forest.

To conclude, the algorithm reconstructs a well defined tree (or, more in general, forest).

Theorem 2. (Algorithm correctness). *CAPRESE reconstructs a well defined tree \mathcal{T} without disconnected components, transitive connections and cycles.*

Additionally, asymptotically with the number of samples, the reconstructed tree is the correct one.

Theorem 3. (Asymptotic convergence). *Let $T = (G \cup \{\blacklozenge\}, E, \blacklozenge)$ be the forest to reconstruct from a set of s input samples, given as the input matrix D . If D is strictly sampled from the distribution induced by T and infinite samples are available, i.e., $s \rightarrow \infty$, CAPRESE with $\lambda \rightarrow 0$ correctly reconstructs T .*

The proof of these Theorems are also in the File S1. These theorems considered datasets where the observed and theoretical probabilities match, because of $s \rightarrow \infty$. However, data often contains false positives and negatives (i.e., data are noisy) and resistance to their effects is desirable in any inferential technique. With this in mind, we prove a corollary of the theorem analogous to a result appearing in [22].

Corollary 1. (Uniform noise). *Let the input matrix D be strictly sampled from the distribution induced by T with sample size $s \rightarrow \infty$,*

but let it be corrupted by noise levels of false positives ϵ_+ and false negatives ϵ_- . Let $p_{\min} = \min_{x \in G} \{\mathcal{P}(x)\}$, CAPRESE correctly reconstructs T for $\lambda \rightarrow 0$ whenever

$$\epsilon_+ < \sqrt{p_{\min}}(1 - \epsilon_+ - \epsilon_-)$$

and $\epsilon_+ + \epsilon_- < 1$.

Essentially, this corollary states that CAPRESE (and so the estimator m) is robust against a noise affecting all samples equally. Also, the fact that it holds for $\lambda \rightarrow 0$ is sound with the theory of shrinkage estimators for which, asymptotically, the corrector factor is not needed to regularize the ill posed problem.

Optimal shrinkage-like coefficient

Theorem 3 and Corollary 1 state that with infinite samples and mild constraints on the false positive/negative rates we get optimal results with $\lambda \rightarrow 0$. Precisely, for the uniform noise model that we applied to synthetic data (see Methods) we have $\epsilon_+ = \epsilon_- = v/2$, thus the hypothesis required by Corollary 1 is

$$v < \frac{\sqrt{p_{\min}}}{1/2 + \sqrt{p_{\min}}}$$

For $p_{\min} = 0.05$, which we set in data generation (see File S1), this inequality implies correct reconstruction for $v < 0.3$ (a 15% error rate), with infinite samples. However, we are interested in performance and the optimal value of λ in situations in which we have finite sample sizes as well. Here, we empirically estimate the optimal λ value, both in the case of trees and forests, as a function of noise and sample size. In the next section, we assess performance of our algorithm empirically.

In Figure 2, we show the variation of the performance of CAPRESE as a function of λ , for datasets with 150 samples generated from tree topologies. The optimal value, i.e., lowest Tree Edit Distance (TED, see Methods), for noise-free datasets (i.e., $v = 0$) is obtained for $\lambda \rightarrow 0$, whereas for the noisy datasets a series of U-shaped curves suggests a unique optimum value for $\lambda \rightarrow 1/2$, immediately observable for $v < 0.15$. Identical results are obtained when dealing with forests (not shown here). In addition, further experiments with n varying around the typical sample size ($n = 150$) show that the optimal λ is largely insensitive to the dataset size (see Figure 3). Thus we have limited our analysis to datasets with the typical sample size that is characteristic of data currently available.

Summarizing, Figures 2 and 3 suggest that for sample size below 250 without false positives and negatives the PR raw estimate α suffices to drive reconstruction to good results (TED is 0 with 250 samples), i.e.,

$$m_{a \rightarrow b} \stackrel{\lambda \rightarrow 0}{\approx} \alpha_{a \rightarrow b} \quad (11)$$

which is obtained by setting λ to a very small value, e.g. 10^{-2} , in order to consider at least a small contribution of the correction factor too. Conversely, when $v > 0$, the best performance is obtained by averaging the shrinkage-like effect, i.e.,

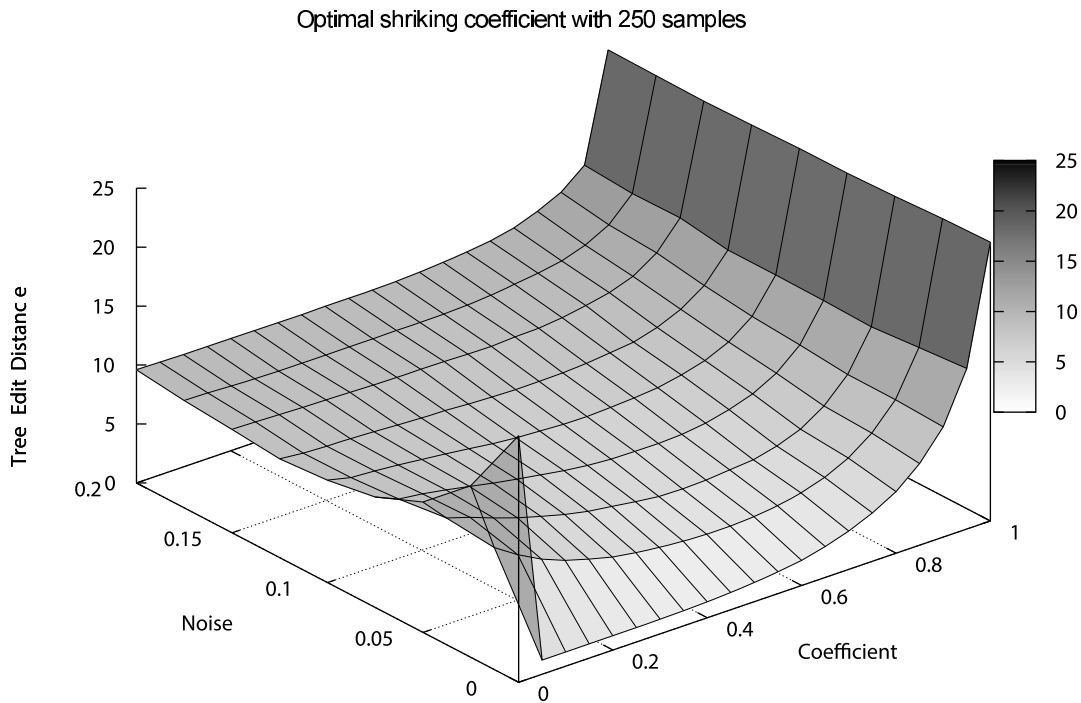


Figure 3. Optimal λ with datasets of different size. We show the analogous of Figure 2 with 50 and 250 samples. The estimation of the optimal shrinkage-like coefficient λ appears to be irrespective of the sample size. doi:10.1371/journal.pone.0108358.g003

$$m_{a \rightarrow b} \stackrel{\lambda=1/2}{=} \frac{\alpha_{a \rightarrow b}}{2} + \frac{\beta_{a \rightarrow b}}{2}. \tag{12}$$

These results suggest that, in general, a unique optimal value for the shrinkage-like coefficient can be determined, even in situations not captured by Theorem 3 and Corollary 1.

Performance of CAPRESE compared to *oncotrees*

An analogue of Theorem 3 holds for Despers’s *oncotrees* (Theorem 3.3, [8]), and an analogue of Corollary 1 holds for Szabo’s extension with uniform noise (Reconstruction Theorem 1, [22]). Thus, with infinite samples both approaches reconstruct the correct trees/forests. With finite samples and noise, however, their performance may show different patterns, as speed of convergence may vary. We investigate this issue in the current section.

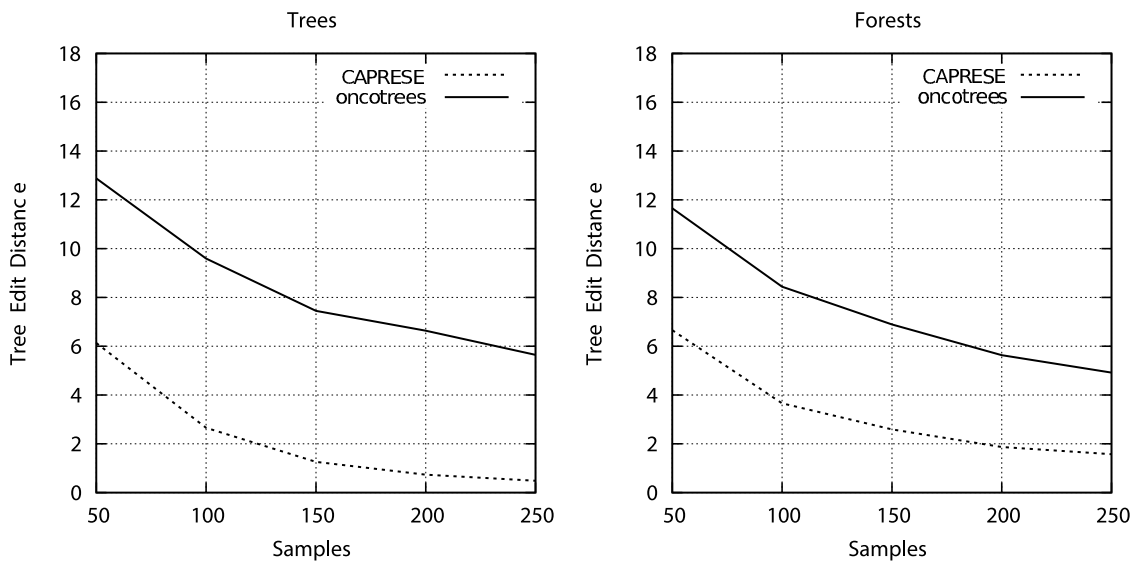


Figure 4. Comparison on noise-free synthetic data. Performance of CAPRESE (dashed line) and *oncotrees* (full line) in average TED when data are generated by random trees (left) and forests (right). In this case $\nu=0$ (no false positives/negatives) and $\lambda \rightarrow 0$ in the estimator m . doi:10.1371/journal.pone.0108358.g004

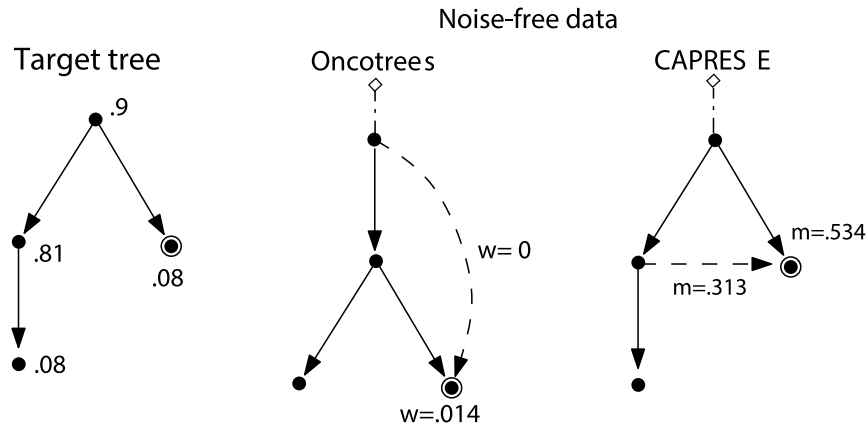


Figure 5. Example of reconstructed trees. Example of reconstruction from a dataset with 100 samples generated by the left tree (the theoretical probabilities are shown, i.e., the doubly-circled event appears in a sample with probability .08), with $\nu=0$. In the sampled dataset oncotrees mislead the cause of the doubly-circled mutation ($w=0$ for the true edge and $w=0.014$ for the wrong one). CAPRESE infers the correct cause (the values of the estimator m with $\lambda=1/2$ are shown, similar results are obtained for $\lambda \rightarrow 1$). doi:10.1371/journal.pone.0108358.g005

In Figure 4 we compare the performance of CAPRESE with oncotrees, for the case of noise-free synthetic data with the optimal shrinkage-like coefficient: $\lambda \rightarrow 0$, equation (11). Since Szabo’s algorithm is equivalent to Desper’s without false negatives and positives, we rely solely on Szabo’s implementation [22]. In Figure 5 we show an example of reconstructed tree where CAPRESE infers the correct tree while oncotrees mislead a causation relation.

In general, one can observe that the TED of CAPRESE is, on average, always bounded above by the TED of oncotrees, both in the case of trees and forests. For trees, with 50 samples the average TED of CAPRESE is around 6, whereas for Desper’s technique it is around 13. The performance of both algorithms improves as long as the number of samples is increased: CAPRESE has the best performance (i.e., TED ≈ 0) with 250 samples, while oncotrees have TED around 6. When forests are considered, the difference between the performance of the algorithms reduces

slightly, but also in this case CAPRESE clearly outperforms oncotrees.

Notice that the improvement due to the increase in the sample size seems to reach a *plateau*, and the initial TED for our estimator seems rather close to the plateau value. This empirical analysis suggests that CAPRESE has already good performances with few samples, a favorable adjoint to Theorem 3. This result has some important practical implications, particularly considering the scarcity of available biological data.

In Figure 6 we extend the comparison to *noisy* datasets. In this case, we used the optimal shrinkage-like coefficient: $\lambda \rightarrow 1/2$, equation (12). The results confirm what observed without false positives and negatives, as CAPRESE outperforms oncotrees up to $\nu=0.15$, for all the sizes of the sample sets. In the File S1 we show similar plots for the noise-free case.

We can thus draw the conclusion that our algorithm performs better with finite samples and noise, since less samples are required

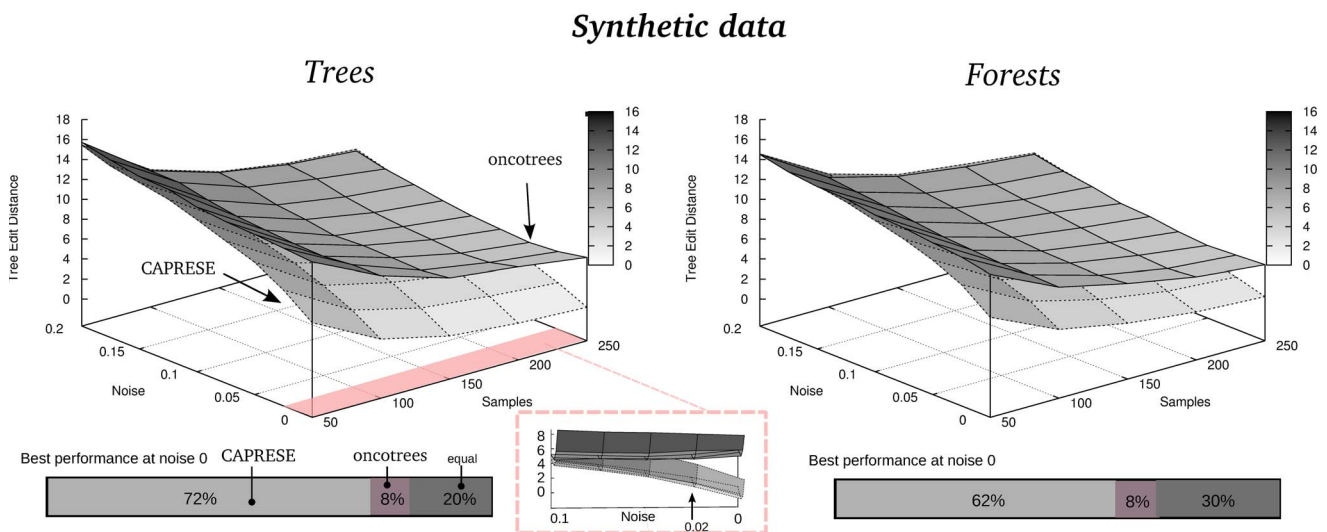


Figure 6. Reconstruction with noisy synthetic data and $\lambda=1/2$. Performance of CAPRESE and oncotrees as a function of the number of samples and noise ν . According to Figure 2 the shrinkage-like coefficient is set to $\lambda=1/2$. The magnified image shows the convergence to Desper’s performance for $\nu \approx 0.1$. The barplot represents the percentage of times the best performance is achieved at $\nu=0$. doi:10.1371/journal.pone.0108358.g006

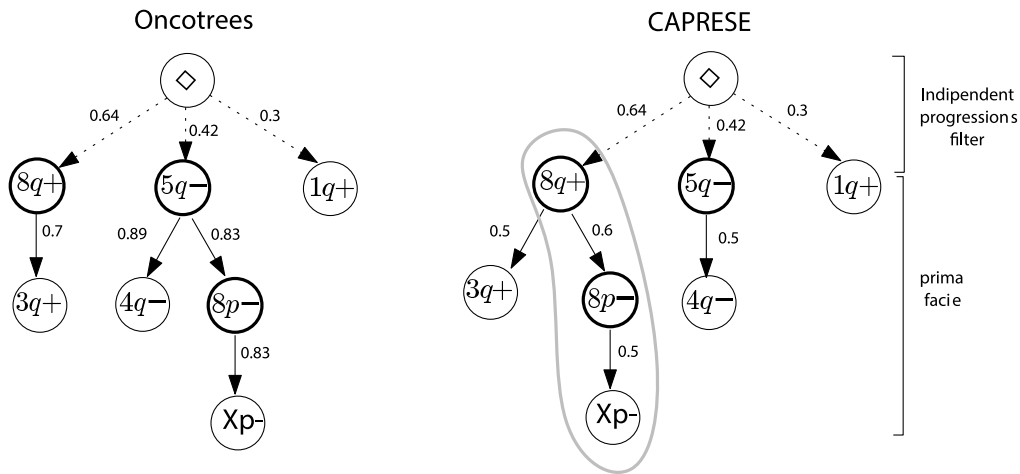


Figure 7. Reconstruction of ovarian cancer progression. Trees reconstructed by oncotrees and CAPRESE (with $\lambda \rightarrow 0$, with $\lambda = 1/2$ the same tree is reconstructed). The set of CGH events considered are gains on $8q$, $3q$ and $1q$ and losses on $5q$, $4q$, $8p$ and Xp . Events on chromosomes arms containing the key genes for ovarian cancer are in bolded circles. In the left tree all edge weights are the observed probabilities of events. In the right the full edges are the causation inferred with the PR and the weights represent the scores used by CAPRESE. Weights on dashed lines are as in the left tree.

doi:10.1371/journal.pone.0108358.g007

to get good performances and a higher resistance to false positives and negatives is shown.

Performance of CAPRESE compared to *Conjunctive Bayesian Networks*

Inspired by Desper’s seminal work, Beerenwinkel and others developed methods to estimate the constraints on the order in which mutations accumulate during cancer progression, using a probabilistic graphical model called *Conjunctive Bayesian Networks* (CBN) [14,15]. While the goal of this research was to reconstruct *direct acyclic graphs* and not trees per se, evidence presented in [24] suggests that, in the absence of noise, these models perform better than oncotrees even at reconstructing *trees*. For this reason, we performed experiments similar to the ones suggested above, comparing CAPRESE to the extension of CBN called *hidden-CBN* (h-CBN) that accounts for noisy genotype observations [15]. This method combines CBNs with a simulated annealing algorithm for structure search and a denoising of the genotypes via the maximum a posteriori estimates to compute the most likely progression. One aspect that complicates a comparison between CAPRESE and (h-)CBN is that the methods assume different models. For example, at the heart of CBN is a monotonicity assumption (i.e., an event can only occur if all its predecessors have occurred) not assumed by CAPRESE. Despite the differences between the model assumptions, we present a preliminary comparison between the methods in Appendix S3, indicating that we not only outperform oncotrees, but h-CBNs as well. In particular, this suggests that CAPRESE converges much faster than h-CBNs with respect to the sample size, also in the presence of noise.

We also analyze the rate of *false positives/negatives* reconstructed by CAPRESE when (synthetic) data are sampled from DAGs (Appendix S3). The rate of *false positives* goes to 0 as the sample size increases, implying that CAPRESE is capable of reconstructing a tree subsumed by the underlying causal DAG topology. In addition, the number of *false negatives* approaches a value proportional to the connectivity of the model from which the data was generated. This is expected, since CAPRESE will assign at most one cause to each considered event. However, it should be

noted that further experiments with samples selected from a wider array of topologies should be performed to confirm these results and compare both methods in full. While not within the scope of the current paper, these issues will be addressed in future work.

Case studies

In the next subsections we apply CAPRESE to real cancer data obtained by *Comparative Genomic Hybridization* (CGH) and *Next Generation Sequencing* (NGS). This shows the potential application of reconstruction techniques to various types of mutational profiles and various cancers.

Performance on cancer CGH datasets

Encouraged by the results in previous sections, we test our reconstruction approach on a real *ovarian cancer* dataset made available within the oncotree package [8]. The data was collected through the public platform SKY/M-FISH [38], used to allow investigators to share molecular cytogenetic data. The data was obtained by using the CGH technique on samples from *papillary serous cystadenocarcinoma* of the ovary. This technique uses fluorescent staining to detect CNV data at the resolution of chromosome arms. While the recent emergence of NGS approaches make the dataset itself rather outdated, the underlying principles remain the same and the dataset provides a valid test-case for our approach. The seven most commonly occurring events are selected from the 87 samples, and the set of events are the following gains and losses on chromosomes arms $G = \{8q+, 3q+, 1q+, 5q-, 4q-, 8p-, Xp-\}$ (e.g., $4q-$ denotes a deletion of the q arm of the 4^{th} chromosome).

In Figure 7 we compare the trees reconstructed by the two approaches. Our technique differs from Desper’s by predicting the causal sequence of alterations

$$8q+ \rightarrow 8p- \rightarrow Xp-,$$

when used either $\lambda \rightarrow 0$ or $\lambda = 1/2$. Notice that among the samples in the dataset some are not generated by the distribution induced by the recovered tree, thus comparing the reconstruction for both λ s is necessary.

Table 1. Estimated confidence for ovarian progression.

Oncotrees (overall confidence 8.3%)							
→	8q+	3q+	5q-	4q-	8p-	1q+	Xp-
◆	.99	.06	.51	.22	.004	.8	.06
8q+	0	.92	.08	0.16	.04	.02	.007
3q+	.002	0	.04	0	0	.09	.04
5q-	.001	.002	0	.52	.39	.009	.16
4q-	0	0	.27	0	.14	.05	.11
8p-	0	0	.07	.08	0	.004	.59
1q+	0	0	0	.004	0	0	0
Xp-	0	0	.003	.003	.04	.01	0

CAPRESE (overall confidence 8.6%)							
→	8q+	3q+	5q-	4q-	8p-	1q+	Xp-
◆	.99	.06	.51	.22	.004	.8	.06
8q+	0	.92	.06	.16	.62	.01	.008
3q+	.002	0	.03	.002	0	.09	.04
5q-	.001	.002	0	.5	.26	.009	.17
4q-	0	0	.29	0	.09	.05	.12
8p-	0	0	.07	.08	0	.004	.59
1q+	0	0	0	.004	0	0	0
Xp-	0	.001	.003	.004	.01	.01	0

Frequency of edge occurrences in the non-parametric bootstrap test, for the trees shown in Figure 7. Bold entries are the edges recovered by the algorithms. doi:10.1371/journal.pone.0108358.t001

At this point, we do not have a biological interpretation for this result. However, we do know that common cancer genes reside in these regions, e.g. the tumor suppressor gene PDGFR on 5q and the oncogene MYC on 8q, and loss of heterozygosity on the short arm of chromosome 8 is quite common (see, e.g., <http://www.genome.jp/kegg/>). Recently, evidence has been reported that 8p contains many cooperating cancer genes [39].

In order to assign a confidence level to these inferences we applied both parametric and non-parametric *bootstrapping methods* to our results. Essentially, these tests consist of using the reconstructed trees (in the parametric case), or the probability observed in the dataset (in the non-parametric case) to generate new synthetic datasets, and then reconstructs again the progressions (see, e.g., [40] for an overview of these methods and [41] for the use of bootstrap for evaluating the confidence of oncogenic




Figure 8. Validating the SETBP1 - ASXL1 relation in atypical Chronic Myeloid Leukemia. Progression models where ASXL1 indel and nonsense point are merged (left) and separate (right) suggest that a missense point mutation hitting SETBP1 can cause a non-sense point mutation in ASXL1, that the observed ASXL1 mutations might be independent and that indel ASXL1 is an early event with high confidence. doi:10.1371/journal.pone.0108358.g008

trees.). The confidence is given by the number of times the trees in Figure 7 are reconstructed from the generated data. A similar approach can be used to estimate the confidence of every edge separately. For oncotrees the *exact tree* is obtained 83 times out of 1000 non-parametric resamples, so its estimated confidence is 8.3%. For CAPRESE the confidence is 8.6%. In the parametric case with false positive and false negative error rates of 0.21 and 0.027, following [22], the confidence of oncotrees is 17% while the confidence of our method is much higher: 32%. When error rates are forced to 0 the confidence of oncotrees raises to 86.6% and 90.9% respectively.

For the non-parametric case, edges confidence is shown in Table 1. Most notably, our algorithm reconstructs the inference $8q+ \rightarrow 8p-$ with high confidence (confidence 62%, and 26% for $5q- \rightarrow 8p-$), while the confidence of the edge $8q+ \rightarrow 5q-$ is only 39%, almost the same as $8p- \rightarrow 8q+$ (confidence 40%). The confidences are similar with either $\lambda \rightarrow 0$ or $\lambda = 1/2$.

Analysis of other CGH datasets. We report the differences between the reconstructed trees also based on datasets of gastrointestinal and oral cancer ([35,37] respectively). In the case of gastrointestinal stromal cancer, among the 13 CGH events considered in [35] (gains on $5p$, $5q$ and $8q$, losses on $14q$, $1p$, $15q$, $13q$, $21q$, $22q$, $9p$, $9q$, $10q$ and $6q$), oncotrees identify the path progression

$$1p- \rightarrow 15q- \rightarrow 13q- \rightarrow 21q-$$

while CAPRESE reconstructs the branch

$$1p- \rightarrow 15q- 1p- \rightarrow 13q- \rightarrow 21q-.$$

In the case of oral cancer, among the 12 CGH events considered in [37] (gains on $8q$, $9q$, $11q$, $20q$, $17p$, $7p$, $5p$, $20p$ and $18p$, losses on $3p$, $8p$ and $18q$), the reconstructed trees differ since oncotrees identifies the path

$$8q+ \rightarrow 20q+ \rightarrow 20p+$$

while our algorithm reconstructs the path

$$3p- \rightarrow 7p+ \rightarrow 20q+ \rightarrow 20p+.$$

These examples show that CAPRESE provides important differences in the reconstruction compared to oncotrees.

Performance on cancer NGS datasets

In this section we show the application of reconstruction techniques to the validation of a specific relation among recurrent mutations involved in *atypical Chronic Myeloid Leukemia* (ACML).

In [42] Piazza *et al.* used high-throughput *exome sequencing technology* to identify somatically acquired mutations in 64 ACML patients, and found a previously unidentified recurring *missense point mutation* hitting SETBP1. By re-sequencing SETBP1 in samples with ACML and other common human cancers, they found that around 25% of the ACML patients tested positive for SETBP1, while most of the other types of tumors were negative. Assessing the possible relationship between SETBP1 variants and the mutations in many driver ACML oncogenes such as (e.g., ASXL1, TET2, KRAS, etc.) no significant association or mutual exclusion with SETBP1 was found but for ASXL1, which was

frequently mutated together with SETBP1, hinting at a potential relation among the events. In particular, ASXL1 was presenting either a *non-sense point* or a *indel* type of somatic mutation.

Hence, we reconstructed ACML progression models from the datasets provided in [42], with the goal of assessing a *potential causal dependency* between mutated SETBP1 and ASXL1. A more extensive analysis is postponed, as we only seek to clearly illustrate the functionalities of the algorithm here.

As a first case (Figure 8, left), we treated the ASXL1 missense point and indel mutations as indistinguishable, and we merged the two events in the dataset. Afterwards, we separated the two types of mutations for ASXL1 (Figure 8, right).

In particular, it is interesting to notice that, when the ASXL1 mutations are considered equivalent, the inference suggests that the mutations belong to two independent progression paths (i.e., the independent progression filters “breaks” every potential causal relation among the events). Conversely, when the mutations are kept separate, the progression model suggests that: (i) the missense point mutation hitting SETBP1 can cause a non-sense point mutation in ASXL1 and (ii) the observed ASXL1 mutations seems to be independent. Concerning edges confidence, as before assessed via non-parametric bootstrap, it is worth noting that the confidence in the indel ASXL1 mutation being an early event raises consistently in the latter case.

All in all, it seems that a progression model allows to test the significance of the association firstly observed in [42] and also refines the knowledge by suggesting a specific causal and temporal relations among events. With this in mind, ad-hoc sequencing experiments might be set up to assess these predictions, eventually providing a strong evidence that could be used to, e.g., synthesize a progression-specific ACML-effective drug.

Discussion

In this work we have introduced a novel framework for the reconstruction of the causal topologies underlying cumulative progressive phenomena, based on the *probability raising* notion of causation. Besides such a probabilistic notion, we also introduced the use of a *shrinkage-like estimator* to efficiently unravel ambiguous causal relations, often present when data are noisy. As a first step towards the definition of our new framework, we have here presented an effective novel technique called CAPRESE (CAncer PRogression Extraction with Single Edges) for the reconstruction of tree or, more generally, forest models of progression which combines probabilistic causation and a shrinkage-like estimation.

We prove correctness of CAPRESE by showing asymptotic convergence to the correct tree. Under mild constraints on the noise rates, this result holds for the reconstruction problem in the presence of uniform noise as well. Moreover, we also compare our technique to the standard tree reconstruction algorithm based on correlation (i.e., Oncotrees) and to a more general bayesian probabilistic graphical model (i.e., Conjunctive Bayesian Networks), and show that CAPRESE outperforms the state-of-the-art on synthetic data, also exhibiting a noteworthy efficiency with relatively small datasets. Furthermore, we tested our technique on ovarian, gastrointestinal and oral cancer CGH data and NGS leukemia data. The CGH analysis suggested that our approach can infer, with high confidence, novel causal relationships which would remain unpredictable in a correlation-based attack. The NGS analysis allowed validating a causal and temporal relation among key mutations in atypical chronic myeloid leukemia.

One of the strong points of CAPRESE is that it can be applied to genomic data of any kind, even heterogeneous, and at any

resolution, as shown. In fact, it simply requires a set of samples in which the presence or the absence of some alterations supposed to be involved in a causal cumulative process have been assessed. Notice also that the results of our technique can be used not only to describe the *progression* of the process, but also to *classify* different progression types. In the case of cancer, for instance, this genome-level classifier could be used to group patients according to the position of the detected individual genomic alterations in the progression model (e.g., at a specific point of the tree) and, consequently, to set up a genome-specific *therapy design* aimed, for instance, at blocking or slowing certain progression paths instead of others, as was studied in [43].

Several future research directions are possible. First, more complex models of progression, e.g. directed acyclic graphs, could be inferred with probability raising and compared to the standard approaches of [14,15,44], as we explained in the introduction. These models, rather than trees, could explain the common phenomenon of *preferential progression paths* in the target process via, e.g., *confluence* among events. In the case of cancer, for instance, these models would be certainly more suitable than trees to describe the accumulation of mutations.

Second, the shrinkage-like estimator itself could be improved by introducing, for instance, different correction factors. In addition, an analytical formulation of the optimal shrinkage-like coefficient could be investigated by starting from the hypotheses which apply to our problem setting, along the lines of [45].

References

1. Hanahan D, Weinberg RA (2011) Hallmarks of cancer: The next generation. *Cell* 144: 646–674.
2. Luo J, Solimini NL, Elledge SJ (2009) Principles of cancer therapy: Oncogene and non-oncogene addiction. *Cell* 136: 823–837.
3. Vogelstein B, Kinzler K (2004) Cancer genes and the pathways they control. *Nature medicine* 10: 789–799.
4. Frank SA (2007) *Dynamics of Cancer*. Princeton University Press.
5. Bell D, Berchuck A, Birrer M, Chien J, Cramer D, et al. (2011) Integrated genomic analyses of ovarian carcinoma. *Nature*.
6. Imielinski M, et al. (2012) Mapping the hallmarks of lung adenocarcinoma with massively parallel sequencing. *Cell* 150: 1107–1120.
7. Vogelstein B, Fearon ER, Hamilton SR, Kern SE, Preisinger AC, et al. (1988) Genetic alterations during colorectal-tumor development. *New England Journal of Medicine* 319: 525–532.
8. Desper R, Jiang F, Kallioniemi O, Moch H, Papadimitriou C, et al. (1999) Inferring tree models for oncogenesis from comparative genome hybridization data. *Journal of Computational Biology* 6: 37–51.
9. Desper R, Jiang F, Kallioniemi O, Moch H, Papadimitriou C, et al. (2000) Distance-based reconstruction of tree models for oncogenesis. *Journal of Computational Biology* 7: 789–803.
10. von Heydebreck A, Gunawan B, Füzesi L (2004) Maximum likelihood estimation of oncogenetic tree models. *Biostatistics* 5: 545–56.
11. Beerwinkel N, Rahnenführer J, Däumer M, Hoffmann D, Kaiser R, et al. (2005) Learning multiple evolutionary pathways from cross-sectional data. *Journal of Computational Biology* 12: 584–598.
12. Beerwinkel N, Rahnenführer J, Kaiser R, Hoffmann D, Selbig J, et al. (2005) Mtreemix: a software package for learning and using mixture models of mutagenetic trees. *Bioinformatics* 21: 2106–2107.
13. Hjeltn M, Höglund M, Lagergren J (2006) New probabilistic network models and algorithms for oncogenesis. *Journal of Computational Biology* 13: 853–865.
14. Beerwinkel N, Eriksson N, Sturmfels B (2007) Conjunctive bayesian networks. *Bernoulli*: 893–909.
15. Gerstung M, Baudis M, Moch H, Beerwinkel N (2009) Quantifying cancer progression with conjunctive bayesian networks. *Bioinformatics* 25: 2809–2815.
16. Gupta A, Bar-Joseph Z (2008) Extracting dynamics from static cancer expression data. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 5: 172–182.
17. Ramakrishnan N, Tadeipalli S, Watson LT, Helm RF, Antonioti M, et al. (2010) Reverse engineering dynamic temporal models of biological processes and their relationships. *PNAS* 107: 12511–12516.
18. Suppes P (1970) *A probabilistic theory of causality*. North Holland Publishing Company.
19. Hitchcock C (2012) Probabilistic causation. In: Zalta E, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2012 edition.

Third, advanced statistical techniques such as *bootstrapping* [40] could be used to account for more sophisticated models of noise within data, so as to decipher complex causal dependencies.

Finally, a further development of the framework could involve the introduction of *timed data*, in order to extend our techniques to settings where a temporal information on the samples is available.

Software availability

The implementation of CAPRESE is part of the TRanslational ONCOlogy (TRONCO) R package and is available for download at standard R repositories.

Supporting Information

File S1 Contains Proofs, Synthetic Data Generation, and Further Results.

(PDF)

Acknowledgments

We are grateful for the many excellent comments we received from anonymous reviewers.

Author Contributions

Conceived and designed the experiments: LO GC AG DR GM MA BM. Performed the experiments: LO GC AG DR. Analyzed the data: LO GC AG DR. Wrote the paper: LO GC AG DR MA BM GM. Software implementation: DR.

20. Ionita I, Daruwala R, Mishra B (2006) Mapping Tumor-Suppressor genes with multipoint statistics from Copy-Number-Variation data. *American Journal of Human Genetics* 79: 13–22.
21. Kleinberg S (2012) *Causality, Probability, and Time*. Cambridge University Press.
22. Szabo A, Boucher K (2002) Estimating an oncogenetic tree when false negatives and positives are present. *Mathematical biosciences* 176: 219–236.
23. Efron B (2013) *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press.
24. Hainke K, Rahnenführer J, Fried R (2012) Cumulative disease progression models for cross-sectional data: A review and comparison. *Biometrical Journal* 54: 617–640.
25. Reichenbach H (1956) *The Direction of Time*. University of California Press.
26. Edmonds J (1967) Optimum branchings. *Journal of Research of the National Bureau of Standards B* 71: 233–240.
27. Kainu T, et al. (2000) Somatic deletions in hereditary breast cancers implicate 13q21 as a putative novel breast cancer susceptibility locus. *Proceedings of the National Academy of Sciences* 97: 9603–9608.
28. Huang Q, Yu G, McCormick S, Mo J, Datta B, et al. (2002) Genetic differences detected by comparative genomic hybridization in head and neck squamous cell carcinomas from different tumor sites: construction of oncogenetic trees for tumor progression. *Genes, Chromosomes and Cancer* 34: 224–233.
29. Radmacher M, Simon R, Desper R, Taetle R, Schäffer A, et al. (2001) Graph models of oncogenesis with an application to melanoma. *Journal of theoretical biology* 212: 535–548.
30. Cartwright N (1979) Causal laws and effective strategies. *Noûs* 13: 419–437.
31. Pearl J (2000) *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
32. Zhang K, Shasha D (1989) Simple fast algorithms for the editing distance between trees and related problems. *SIAM journal on computing* 18: 1245–1262.
33. Hamming RW (1950) Error-detecting and error-correcting codes. *Bell System Technical Journal* 29: 147–160.
34. Samuelson E, Karlsson S, Partheen K, Nilsson S, Szpirer C, et al. (2012) Bac cgh-array identified specific small-scale genomic imbalances in diploid dmba-induced rat mammary tumors. *BMC cancer* 12: 352.
35. Gunawan B, et al. (2007) An oncogenetic tree model in gastrointestinal stromal tumours (gists) identifies different pathways of cytogenetic evolution with prognostic implications. *The Journal of pathology* 211: 463–470.
36. Longerich T, Mueller M, Breuhahn K, Schirmacher P, Benner A, et al. (2012) Oncogenetic tree modeling of human hepatocarcinogenesis. *International Journal of Cancer* 130: 575–583.
37. Pathare S, Schäffer A, Beerwinkel N, Mahimkar M (2009) Construction of oncogenetic tree models reveals multiple pathways of oral cancer progression. *International journal of cancer* 124: 2864–2871.

38. Knutsen T, Gobu V, Knaus R, Padilla-Nash H, Augustus M, et al. (2005) The interactive online sky/m-fish & cgh database and the entrez cancer chromosomes search database: Linkage of chromosomal aberrations with the genome sequence. *Genes, Chromosomes and Cancer* 44: 52–64.
39. Xue W, et al. (2012) A cluster of cooperating tumor-suppressor gene candidates in chromosomal deletions. *Proceedings of the National Academy of Sciences* 109: 8212–8217.
40. Efron B (1982) *The jackknife, the bootstrap and other resampling plans*, volume 38. SIAM.
41. Szabo A, Boucher KM (2008) *Oncogenetic trees. Handbook of Cancer Models with Applications to Cancer Screening, Cancer Treatment and Risk Assessment* World Scientific.
42. Piazza R, Valletta S, Winkelmann N, Redaelli S, Spinelli R, et al. (2013) Recurrent setbp1 mutations in atypical chronic myeloid leukemia. *Nature Genetics* 45: 18–24.
43. Olde Loohuis L, Witzel A, Mishra B (2014) Cancer hybrid automata: Model, beliefs & therapy. *Information and Computation* 236: 68–86.
44. Gerstung M, Eriksson N, Lin J, Vogelstein B, Beerenwinkel N (2011) The temporal order of genetic and pathway alterations in tumorigenesis. *PLoS one* 6: e27136.
45. Efron B, Morris C (1973) Stein's estimation rule and its competitors—an empirical bayes approach. *Journal of the American Statistical Association* 68: 117–130.