

Conjunctive Bayesian Networks

Niko Beerenwinkel

`beerenw@fas.harvard.edu`

Program for Evolutionary Dynamics,
Harvard University, Cambridge, MA 02138-3758

Nicholas Eriksson*

`nke@stanford.edu`

Department of Statistics,
Stanford University, Stanford, CA 94305-4065

Bernd Sturmfels

`bernd@math.berkeley.edu`

Department of Mathematics,
University of California, Berkeley, CA 94720-3840

Abstract

Conjunctive Bayesian networks (CBNs) are graphical models that describe the accumulation of events which are constrained in the order of their occurrence. A CBN is given by a partial order on a (finite) set of events. CBNs generalize the oncogenetic tree models of Desper et al. (1999) by allowing the occurrence of an event to depend on more than one predecessor event. The present paper studies the statistical and algebraic properties of CBNs. We determine the maximum likelihood parameters and present a combinatorial solution to the model selection problem. Our method performs well on two datasets where the events are HIV mutations associated with drug resistance. Concluding with a study of the algebraic properties of CBNs, we show that CBNs are toric varieties after a coordinate transformation and that their ideals possess a quadratic Gröbner basis.

Keywords: Bayesian network, distributive lattice, Gröbner basis, maximum likelihood estimation, Möbius transform, mutagenetic tree, oncogenetic tree, sagbi basis, toric variety.

1 Introduction

The conjunctive Bayesian network (CBN) model on a finite partially ordered set (poset) was introduced in Beerenwinkel et al. (2006, §4) as well as in the form of noisy-and models in the AI literature (e.g., Pearl, 1986). Here, we give a self-contained study of the statistical and algebraic properties of this model. CBNs are specializations of Bayesian networks. They include the oncogenetic (also called mutagenetic) tree models of Desper et al. (1999) which have proven very useful in cancer research (Radmacher et al., 2001) and in the study of HIV drug resistance (Beerenwinkel et al., 2005a).

The models are motivated by the following class of problems. Consider a finite set of genetic events, for example, DNA mutations or chromosomal alterations, and assume that the genetic changes are permanent. In this situation, each individual, defined by its genotype, is completely characterized by the subset of the events that have occurred. We wish to learn the constraints on the orders in which these events have accumulated. A CBN is a probabilistic model of this process derived from a partial order on the set of events. This partial order encapsulates the dependencies between events.

For example, consider the development of drug resistance in HIV. This evolutionary process is characterized by the accumulation of resistance mutations in the viral genome. Under fixed drug pressure these mutations are virtually non-reversible, because they confer a strong selective advantage. Thus, the genetic events are fixations of specific amino acid substitutions in the virus population. In each patient, different combinations of resistance mutations will occur. We seek to determine the prevalent mutational pathways along which HIV accumulates drug resistance (cf. Beerenwinkel et al., 2005b). An order constraint might read that mutations at position 20 and 82 of the target protein must occur before we can see a mutation at position 54. This constraint appears in Figure 4(a). We will analyze such data in Section 4 and see that CBNs are an efficient, accurate tool for this problem.

For another example, the development of cancer is associated with large-scale genetic events such as the gains or losses of parts of chromosomes (Michor et al., 2004). Knowledge of the constraints on the accumulation of these genetic events helps in assessing the progression of the cancer and assigning treatments (cf. Rahnenführer et al., 2005).

A CBN consists of a set of binary random variables, called events, and a partial order on these events. While we will use the language of the theory of posets, readers can equivalently think of the partial order as a directed acyclic graph (DAG), with edges encoding the order relations. CBNs are specializations of Bayesian networks, with the difference being that in a CBN,

an event cannot occur until all of its parents have occurred. Thus the events that occur with positive probability form a *distributive lattice*. Distributive lattices are important combinatorial objects which have been studied in statistics. For example, the LCI models of (Andersson and Perlman, 1993; Andersson et al., 1995) use distributive lattices to encode conditional independence statements. Although similar in spirit, readers should beware that LCI models are not the same as CBNs.

Our original motivation for studying CBNs came from work on mutagenetic trees, introduced in Desper et al. (1999). Mutagenetic trees assume that each event depends on the occurrence of at most one other previous event. CBNs relax this assumption, allowing for an arbitrary partial order on the events. By relaxing this assumption, CBNs are able to model a larger range of biomedical problems effectively.

Even though they generalize currently used models, CBNs are still very restrictive compared to Bayesian networks in general. However, CBNs have the benefit that the maximum likelihood parameters and structure can be written down in closed form (Proposition 2 and Theorem 5). This is an uncommon phenomenon in the theory of graphical models, and should be of independent interest. In addition, the number of parameters in a CBN does not depend on the graph structure, so we do not need to use, for example, the AIC or BIC procedures.

CBNs have also been studied under the name of noisy-and models in the AI community (Meek and Heckerman, 1997; Pearl, 1986, 1988) as a model for causal inference. The basic idea is that a number of causes influence a common effect through latent intermediate variables; the noisy-and model requires all causes to have happened before the effect can. The study of these models focusses on learning the causal structure given latent variables, in contrast to our situation where we wish to learn the structure of a network on observed variables.

In this paper, we show that CBNs have desirable algebraic, statistical, and combinatorial properties. CBN models can be learned efficiently, they can be extended to take into account noise in the data, and they perform better than mutagenetic trees in our applications (cf. Figure 3). This paper is organized as follows. After formally introducing CBNs in Section 2, we compute the maximum likelihood (ML) estimator for a CBN in Section 3 and use this to give a combinatorial characterization of the CBN model of maximal likelihood. Next, in Section 4, we compare the performance of CBNs to mutagenetic trees on two data sets of HIV drug resistance mutations. Finally, in Section 5, we study algebraic properties of CBNs. These properties are surprisingly similar to other algebraic results for statistical models. This

material may ultimately become relevant for statistical inference, but may also be of independent interest to mathematicians. We determine the prime ideal of algebraic invariants of a CBN and we show that this model is a toric variety in a suitable coordinate system. Our main tool is the Möbius transform, a standard tool in working with posets which has found application in the graphical models literature, cf. Drton and Richardson (2005, §3) and Lauritzen (1996, pg. 239).

2 Conjunctive Bayesian networks

A CBN model is specified by a set \mathcal{E} of events, a partial order “ \leq ” on the events, and parameters θ_e for each event e . We will assume that there are n events, labeled as $[n] := \{1, \dots, n\}$. Therefore we write the parameters as $\theta = (\theta_1, \dots, \theta_n)$. Frequently, we will abuse notation and refer to both the model $(\mathcal{E}, \leq, \theta)$ and the poset (\mathcal{E}, \leq) as \mathcal{E} when the meaning is clear from context. A relation $e_1 < e_2$ between two events in \mathcal{E} is interpreted as requiring that event e_1 must happen before event e_2 can. The parameter θ_e is the conditional probability that the event $e \in \mathcal{E}$ has occurred given that its predecessor events have already occurred.

The state space of the CBN model is the distributive lattice $\mathcal{G} = J(\mathcal{E})$ of order ideals in \mathcal{E} . An order ideal is a subset $g \subseteq \mathcal{E}$ such that if $e_2 \in g$ and $e_1 < e_2$, then $e_1 \in g$. Readers unfamiliar with posets and their distributive lattices are referred to Beerenwinkel et al. (2006, §2) for a brief introduction. The elements of \mathcal{G} are called *genotypes*. Thus, a genotype $g \in \mathcal{G}$ is a subset of \mathcal{E} , or equivalently the binary string in which each bit indicates the occurrence of an event. This terminology presumes a well-defined ground state $0 \dots 0$ in which none of the events have yet occurred. In our examples, the unmutated virus strain or the unmutated potential cancer cell is referred to as the “wild type”. Hence, for describing mutant types, we only have to keep track of which sites differ from the wild type because of the assumption of non-reversibility of events.

We write $\min(g^c)$ for the minimal elements in the complement $g^c = \mathcal{E} \setminus \{g\}$ of a genotype g . The elements of $\min(g^c)$ are the events that have not occurred in g but could happen next. For example, in Figure 1, if $g = \{1, 2\}$ then $\min(g^c) = \{3, 4\}$. The probability of observing the genotype $g \in \mathcal{G}$ in the CBN model on the poset \mathcal{E} is defined to be

$$P_g(\theta) = \prod_{e \in g} \theta_e \cdot \prod_{e \in \min(g^c)} (1 - \theta_e).$$




Figure 1: Poset on four events, order ideals, and genotype lattice.

That is, the probability of observing g is the probability that all the events in g have happened times the probability that none of the events that depend only on g have happened.

Equivalently, the CBN model on \mathcal{E} is the directed graphical model for the binary random variables $(X_e)_{e \in \mathcal{E}}$ whose graph has edges $e \rightarrow f$ for all cover relations $e < f$ in \mathcal{E} , and whose conditional probability tables are

$$[\Pr(X_e = b \mid X_{\text{pa}(e)} = a)]_{a \in \{0,1\}^{\text{pa}(e)}, b \in \{0,1\}} = \begin{bmatrix} 1 & 0 \\ \vdots & \vdots \\ 1 & 0 \\ 1 - \theta_e & \theta_e \end{bmatrix},$$

where $\text{pa}(e)$ denotes the parents of e in the acyclic directed graph \mathcal{E} .

Example 1. Let $n = 4$ and suppose \mathcal{E} is the poset defined on four events by the cover relations $1 < 3$, $1 < 4$, $2 < 3$ and $2 < 4$. The poset \mathcal{E} has precisely seven order ideals, so the distributive lattice \mathcal{G} consists of seven genotypes. They are displayed in Figure 1. The CBN model \mathcal{E} is the family of probability distributions on \mathcal{G} which is given parametrically as follows:

$$\begin{aligned} P_\emptyset(\theta) &= (1 - \theta_1)(1 - \theta_2), & P_1(\theta) &= \theta_1(1 - \theta_2), \\ P_2(\theta) &= \theta_2(1 - \theta_1), & P_{12}(\theta) &= \theta_1\theta_2(1 - \theta_3)(1 - \theta_4), \\ P_{1234}(\theta) &= \theta_1\theta_2\theta_3\theta_4, & P_{123}(\theta) &= \theta_1\theta_2\theta_3(1 - \theta_4), \\ P_{124}(\theta) &= \theta_1\theta_2\theta_4(1 - \theta_3). \end{aligned}$$

The sum of these seven polynomials is identical to one. The parameters are the conditional probabilities $\theta_e = \Pr(X_e = 1 \mid X_{\text{pa}(e)} = (1, \dots, 1))$. \square

3 Maximum likelihood estimation

Consider any CBN model \mathcal{E} on $[n]$. The data for this model takes the form of a function $u: \mathcal{G} \rightarrow \mathbb{N}$, $g \mapsto u_g$, where u_g is the number of observations of the genotype g . Given such data $u \in \mathbb{N}^{\mathcal{G}}$, the following proposition gives an easy formula for maximum likelihood estimation of the model parameters.

Proposition 2. *For each event e in the CBN model \mathcal{E} , the ML estimator $\hat{\theta}_e$ of θ_e equals the relative frequency of the genotypes which contain e among all genotypes that contain the events which are strictly below e . In symbols,*

$$\hat{\theta}_e = \frac{\sum_{g:e \in g} u_g}{\sum_{g:\text{below}(e) \subseteq g} u_g} \quad \text{for all } e \in \mathcal{E}.$$

Proof. The log-likelihood function for the given data $u \in \mathbb{N}^{\mathcal{G}}$ equals

$$\ell_u(\theta) = \sum_{g \in \mathcal{G}} u_g \cdot \left(\sum_{e \in g} \log \theta_e + \sum_{e \in \text{min}(g^c)} \log(1 - \theta_e) \right).$$

The partial derivative of this expression with respect to a parameter θ_e is

$$\frac{\partial \ell_u}{\partial \theta_e} = \frac{A}{\theta_e} - \frac{B}{1 - \theta_e},$$

where A is the sum over all frequencies u_g of genotypes g containing e , and B is the sum over all frequencies u_g where $e \notin g$ but $\text{below}(e) \subseteq g$. Equating this partial derivative with zero, we obtain

$$\hat{\theta}_e = \frac{A}{A + B},$$

and this is precisely the formula asserted in the proposition. \square

Example 3. We illustrate Proposition 2 for the model in Example 1 and Figure 1. Since $\text{below}(1) = \emptyset$, the ML estimator for θ_1 is

$$\hat{\theta}_1 = \frac{u_1 + u_{12} + u_{123} + u_{124} + u_{1234}}{u + u_1 + u_2 + u_{12} + u_{123} + u_{124} + u_{1234}},$$

and similarly for θ_2 . For θ_3 , $\text{below}(3) = \{1, 2\}$, and hence

$$\hat{\theta}_3 = \frac{u_{123} + u_{1234}}{u_{12} + u_{123} + u_{124} + u_{1234}}.$$

The expression for the ML estimator of θ_4 is similar. \square

Remark 4. Proposition 2 shows that the ML estimator for the CBN model is a rational function in the data. In the language of Catanese et al. (2006), this says that the *ML degree* of every CBN model is equal to one.

We identify the elements of \mathcal{G} with strings in $\{0, 1\}^n$. A probability distribution on \mathcal{G} is thus an element of the $(2^n - 1)$ -dimensional simplex Δ with coordinates indexed by $\{0, 1\}^n$. Write $\text{supp}(u)$ for the nonzero coordinates of u , i.e., for the genotypes that occur in the data set. We say that u *separates the events* if, for any two elements $e, f \in [n]$, there exists $g \in \text{supp}(u)$ such that $g \cap \{e, f\}$ is either $\{e\}$ or $\{f\}$. If this is not the case, then we can consider $\{e, f\}$ as a single event and replace $[n]$ by $[n - 1]$.

We call any genotype $g \subseteq [n]$ *compatible* with the model \mathcal{E} if $g \in J(\mathcal{E}) = \mathcal{G}$. This is equivalent to $P_g(\theta)$ not being the zero polynomial; see also Beerenwinkel and Drton (2005, Definition 14.2). The data u is said to be *compatible* with \mathcal{E} if all $g \in \text{supp}(u)$ are compatible with \mathcal{E} . Our next theorem is the main result of this section. It gives a combinatorial solution to the problem of model selection among CBNs. Here, any given data set $u: \{0, 1\}^n \rightarrow \mathbb{N}$ is identified with the corresponding empirical probability distribution in Δ . For such $u \in \Delta$, we can compute the ML estimator $\hat{\theta}$ for each poset \mathcal{E} on $[n]$. We define the *ML CBN model* for u to be the poset \mathcal{E} for which the log-likelihood $\ell_u(\hat{\theta})$ has the largest numerical value.

Theorem 5. *Let $u \in \Delta$ be a probability distribution which separates the events. Then there is a unique largest poset \mathcal{E}_u , such that u is compatible with \mathcal{E}_u , and the poset \mathcal{E}_u is the unique ML CBN model for u .*

Here, “largest poset” refers to the refinement relation among posets on $[n]$, i.e., $\mathcal{E} \subset \mathcal{E}'$ means that every relation $e < f$ in \mathcal{E} also holds in \mathcal{E}' . Note that this inclusion is reversed for the induced genotype lattices: $\mathcal{E} \subset \mathcal{E}'$ if and only if $\mathcal{G} = J(\mathcal{E}) \supset \mathcal{G}' = J(\mathcal{E}')$.

Proof. The probability $P_g(\theta)$ is identically zero if and only if g is not in $\mathcal{G} = J(\mathcal{E})$. This implies that the likelihood function $\prod_{g \in \text{supp}(u)} P_g(\theta)^{u_g}$ is identically zero if and only if u is not compatible with the poset \mathcal{E} . Therefore, we need to consider only posets \mathcal{E} such that u is compatible with \mathcal{E} .

We claim that there is a unique maximal poset \mathcal{E}_u to which u is compatible. Namely, \mathcal{E}_u is the set of all relations $e_1 < e_2$ such that $g \cap \{e_1, e_2\} \neq \{e_2\}$ for all $g \in \text{supp}(u)$. Then \mathcal{E}_u is an antisymmetric relation on $[n]$ because u separates the events. The relation \mathcal{E}_u is transitive because $g \cap \{e_1, e_3\} = \{e_3\}$ implies $g \cap \{e_1, e_2\} = \{e_2\}$ or $g \cap \{e_2, e_3\} = \{e_3\}$. Thus \mathcal{E}_u is a poset, and adding any relation makes u incompatible with it.

It remains to show the following statement: If $\mathcal{E}_1 \subset \mathcal{E}_2 \subseteq \mathcal{E}_u$, then \mathcal{E}_2 is more likely than \mathcal{E}_1 . It suffices to show this where \mathcal{E}_1 and \mathcal{E}_2 differ by only one relation, which without loss of generality we assume to be $1 < 2$. Thus the events 1 and 2 are incomparable in \mathcal{E}_1 but 1 must come before 2 in \mathcal{E}_2 .

Let $\mathcal{G}_1 = J(\mathcal{E}_1)$ and $\mathcal{G}_2 = J(\mathcal{E}_2)$. We begin by finding the ML parameters for the two models. Write $\hat{\theta}_e$ for the ML parameters for \mathcal{G}_1 and $\hat{\eta}_e$ for the ML parameters for \mathcal{G}_2 . According to Proposition 2, we have

$$\hat{\theta}_e = \frac{\sum_{g \in \mathcal{G}_1: e \in g} u_g}{\sum_{g \in \mathcal{G}_1: \text{below}(e) \subseteq g} u_g} \quad \text{and} \quad \hat{\eta}_e = \frac{\sum_{g \in \mathcal{G}_2: e \in g} u_g}{\sum_{g \in \mathcal{G}_2: \text{below}(e) \subseteq g} u_g}.$$

Since $\mathcal{G}_1 \supset \mathcal{G}_2 \supseteq \text{supp}(u)$, the numerators of both expressions are the same, namely, we are summing the counts u_g over all genotypes g that contain e .

We claim that $\hat{\theta}_e = \hat{\eta}_e$ except when $e = 2$. In both cases, the denominator is the sum over u_g for all genotypes g where e has either already occurred or is allowed to occur next. Since \mathcal{E}_1 and \mathcal{E}_2 differ in only one relation, $1 < 2$, the denominators are the same (and hence $\hat{\theta}_e = \hat{\eta}_e$) unless $e = 2$.

In order to further analyze the ML estimates, we set

$$\begin{aligned} V_1 &= \sum_{g: 1 \in g} u_g, & V_2 &= \sum_{g: 2 \in g} u_g, \\ N &= \sum_{g \in \mathcal{G}_1: \text{below}(2) \subseteq g} u_g, & M &= \sum_{g \in \mathcal{G}_2: \text{below}(2) \subseteq g} u_g. \end{aligned}$$

With this notation, the maximum likelihood parameters are

$$\hat{\theta}_2 = \frac{V_2}{N}, \quad \hat{\eta}_2 = \frac{V_2}{M}.$$

Notice that since event 1 always happens in the data before event 2, we have $V_2 \leq V_1$. Since \mathcal{E}_2 has more conditions than \mathcal{E}_1 , we have $M \leq N$. And since event 1 is required to happen before event 2 can in \mathcal{E}_2 , we have $V_1 \leq M$. Combining these inequalities gives us $V_2 \leq V_1 \leq M \leq N$.

Our analysis will involve the ratios of the ML parameters

$$\frac{\hat{\theta}_2}{\hat{\eta}_2} = \frac{M}{N}, \quad \frac{1 - \hat{\theta}_2}{1 - \hat{\eta}_2} = \frac{M}{N} \frac{N - V_2}{M - V_2}. \quad (1)$$

For $i = 1, 2$, the likelihood function for the given distribution u equals

$$L_u(\theta; \mathcal{G}_i) = \prod_g \left(\prod_{e \in g} \theta_e^{u_g} \right) \cdot \left(\prod_{e \in \min(g^c)} (1 - \theta_e)^{u_g} \right).$$

Substitute $\theta_e = \widehat{\theta}_e$ for $i = 1$ and $\theta_e = \widehat{\eta}_e$ for $i = 2$. Our assertion states that

$$L_u(\widehat{\theta}; \mathcal{G}_1) \leq L_u(\widehat{\eta}; \mathcal{G}_2).$$

To prove this, we consider the ratio $L_u(\widehat{\theta}; \mathcal{G}_1)/L_u(\widehat{\eta}; \mathcal{G}_2)$, written as a product over $g \in \text{supp}(u)$, and we examine the four possibilities for $g \cap \{1, 2\}$:

Case 1: $g = 00^*$. Here event 2 can happen in \mathcal{E}_1 , but it can't yet in \mathcal{E}_2 , since it needs event 1 to happen first. This contributes a factor $(1 - \widehat{\theta}_2)^{u_g}$ to the product over g in $L_u(\widehat{\theta}; \mathcal{G}_1)/L_u(\widehat{\eta}; \mathcal{G}_2)$. Since event 2 has not happened yet, there are no factors $\widehat{\theta}_2/\widehat{\eta}_2$ in this product, so everything else cancels.

Case 2: $g = 01^*$. This case cannot happen by compatibility.

Case 3: $g = 10^*$. Event 2 has not happened in either case, so all the terms in the product over $e \in g$ cancel. The same set of events can happen in both \mathcal{G}_1 and \mathcal{G}_2 , so everything in the product over $e \in \min(g^c)$ cancels except the factor with $e = 2$, which occurs both in $L_u(\widehat{\theta}; \mathcal{G}_1)$ and in $L_u(\widehat{\eta}; \mathcal{G}_2)$.

Case 4: $g = 11^*$. This case is similar to case 3, except that event 2 has now happened in both cases.

The result of this analysis is the identity

$$\frac{L_u(\widehat{\theta}; \mathcal{G}_1)}{L_u(\widehat{\eta}; \mathcal{G}_2)} = \prod_{g=00^*} (1 - \widehat{\theta}_2)^{u_g} \prod_{g=10^*} \left(\frac{1 - \widehat{\theta}_2}{1 - \widehat{\eta}_2} \right)^{u_g} \prod_{g=11^*} \left(\frac{\widehat{\theta}_2}{\widehat{\eta}_2} \right)^{u_g}. \quad (2)$$

Notice that

$$\sum_{g=10^*} u_g + \sum_{g=11^*} u_g = V_1 \quad \text{and} \quad \sum_{g=11^*} u_g = V_2.$$

Therefore, $\sum_{g=00^*} u_g = 1 - V_1$. Substituting (1) into (2), we get

$$\begin{aligned} \frac{L_u(\widehat{\theta}; \mathcal{G}_1)}{L_u(\widehat{\eta}; \mathcal{G}_2)} &= \left(\frac{N - V_2}{N} \right)^{1-V_1} \left(\frac{M(N - V_2)}{N(M - V_2)} \right)^{V_1-V_2} \left(\frac{M}{N} \right)^{V_2} \\ &= \frac{M^{V_1} (N - V_2)^{1-V_2}}{N (M - V_2)^{V_1-V_2}}. \end{aligned} \quad (3)$$

The following lemma shows that (3) is less than or equal to one for all $0 \leq V_2 \leq V_1 \leq M \leq N \leq 1$. This completes the proof of Theorem 5. \square

Lemma 6. *If x, y, a, b are real numbers with $0 \leq a \leq b \leq x \leq y \leq 1$ then*

$$\frac{x^b (y - a)^{1-a}}{y (x - a)^{b-a}} \leq 1 \quad (4)$$

Proof. We fix a and b and regard the left hand side of (4) as a function $f_{a,b}(x, y)$ of x and y . The two partial derivatives of this function satisfy

$$\frac{\partial f_{a,b}}{\partial x} = \frac{a(x-b)}{x(x-a)} \cdot f_{a,b}(x, y) \quad \text{and} \quad \frac{\partial f_{a,b}}{\partial y} = \frac{a(1-y)}{y(y-a)} \cdot f_{a,b}(x, y).$$

Both expressions are positive on the triangle $\{(x, y) \in \mathbb{R}^2 : \max(a, b) \leq x \leq y \leq 1\}$, hence $f_{a,b}(x, y)$ is bounded above by $f_{a,b}(1, 1) = (1-a)^{1-b} \leq 1$. \square

We summarize the results of this section in the following algorithm.

Algorithm 7. (*Model selection and parameter estimation for CBN models*)

INPUT: A probability distribution $u \in \Delta$ on the set of genotypes $\{0, 1\}^n$.

OUTPUT: The ML CBN model \mathcal{E}_u and the ML parameters $\hat{\theta}$.

STEP 1: Check whether u separates the n events. If not, group non-distinguished events together, thus decrementing n , and replace u by the probability distribution which is induced on the smaller set of genotypes.

STEP 2: Define the poset \mathcal{E}_u on $[n]$ as follows: For any two events $e, f \in [n]$ we set $e < f$ in \mathcal{E}_u if and only if $g \cap \{e, f\} \neq \{f\}$ for all $g \in \text{supp}(u)$.

STEP 3: For each event $e \in [n]$, compute $\hat{\theta}_e$ by the formula in Proposition 2.

STEP 4: Output the poset \mathcal{E}_u and the vector $\hat{\theta} \in [0, 1]^n$.

4 Application to HIV genetic data

The use of Algorithm 7 to obtain the ML CBN model is complicated by the presence of noise in real world data sets. Any relation $e < f$ between two events e and f will be estimated to be part of the poset \mathcal{E} only if no genotype containing f but not e has been observed. Thus, the algorithm will miss relations $e < f$ that have strong but imperfect support. The problem of noisy data has been analyzed in earlier work on mutagenetic tree models. It can be addressed by explicit error models in a ML framework as described in Beerenwinkel and Drton (2005, §14.2) and Beerenwinkel and Drton (2007, §3.3). Also, Szabo and Boucher (2002) have incorporated an error model directly into the reconstruction algorithm of Desper et al. (1999).

We propose the following method for constructing a range of CBN models as the error tolerance ϵ varies. Let \mathcal{E}_ϵ be the poset on $[n]$ which consists of all relations $e < f$ which are violated by at most a fraction ϵ of the data. Thus, for $\epsilon = 0$, we recover \mathcal{E}_u . Generally, some observations $g \in \text{supp}(u)$ will be incompatible with the model \mathcal{E}_ϵ . These samples are removed prior to ML estimation of the model parameters θ . In order to account for both the compatible and incompatible data, we use a simple mixture model.

Write $\mathcal{G}_\epsilon = J(\mathcal{E}_\epsilon)$ for the genotype space of the model \mathcal{E}_ϵ . We assume that the incompatible genotypes $g \notin \mathcal{G}_\epsilon$ are generated with uniform probability $1/(2^n - |\mathcal{G}_\epsilon|)$. Our mixture model \mathcal{E}'_ϵ is given parametrically by the event probabilities θ_ϵ and a mixture parameter λ as

$$P'_g(\theta, \lambda) = \begin{cases} \lambda P_g(\theta) & \text{if } g \in \mathcal{G}_\epsilon \\ (1 - \lambda) (2^n - |\mathcal{G}_\epsilon|)^{-1} & \text{if } g \notin \mathcal{G}_\epsilon, \end{cases}$$

for each observation $g \in \{0, 1\}^n$. This expression gives an explicit trade-off between a large number of compatible samples and a good model fit.

Since the mixing distributions of the model \mathcal{E}'_ϵ have disjoint support, the log-likelihood function of the data $u: \{0, 1\}^n \rightarrow \mathbb{N}$ decomposes as follows:

$$\begin{aligned} \ell'_u(\theta, \lambda) &= \sum_{g \in \mathcal{G}_\epsilon} u_g [\log \lambda + \log P_g(\theta)] \\ &\quad + \sum_{g \notin \mathcal{G}_\epsilon} u_g [\log(1 - \lambda) - \log(2^n - |\mathcal{G}_\epsilon|)]. \end{aligned} \quad (5)$$

Proposition 8. *The ML estimators $\hat{\theta}_\epsilon$ of θ_ϵ under the model \mathcal{E}'_ϵ are given by Proposition 2. The ML estimator $\hat{\lambda}$ of λ under the model \mathcal{E}'_ϵ is given by the fraction of the data u which is compatible with the model \mathcal{E}_ϵ . That is,*

$$\hat{\lambda} = \frac{\sum_{g \in \mathcal{G}_\epsilon} u_g}{\sum_g u_g}.$$

Proof. The partial derivatives of (5) with respect to θ_ϵ are the same as they were in Proposition 2. Next, if we solve

$$0 = \frac{\partial \ell'}{\partial \lambda} = \sum_{g \in \mathcal{G}_\epsilon} \frac{u_g}{\lambda} - \sum_{g \notin \mathcal{G}_\epsilon} \frac{u_g}{1 - \lambda},$$

we get the above formula for $\hat{\lambda}$. □

We now apply these methods to mutation data from HIV that was obtained from patients under antiretroviral therapy. The set \mathcal{E} of genetic events consists of seven amino acid alterations in the HIV genome that confer drug resistance. Specifically, as an unordered set,

$$\mathcal{E} = \{\text{K20R, M36I, M46I, I54V, A71V, V82A, I84V}\},$$

where, for example, K20R indicates the amino acid mutation from lysine (K) to arginine (R) at position 20 of the HIV protease. We consider two datasets

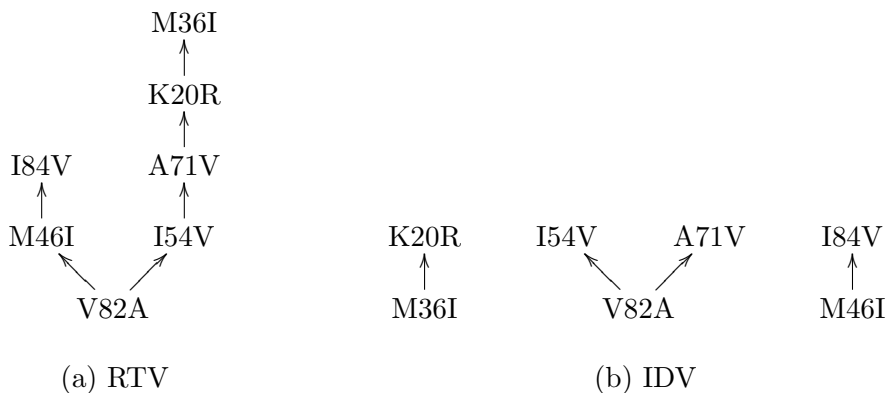


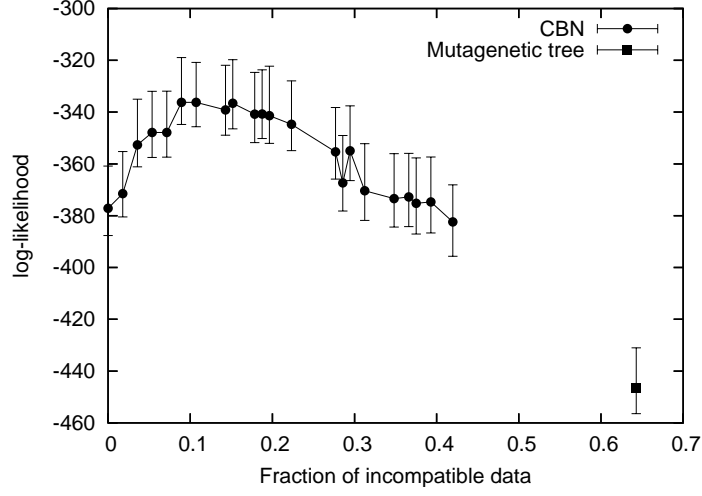
Figure 2: Posets corresponding to the mutagenetic trees that were found in Beerenwinkel et al. (2006, Figure 3) for (a) ritonavir (RTV) and (b) indinavir (IDV).

from the Stanford HIV Drug Resistance Database (Rhee et al., 2003), which consist of 112 and 691 observed genotypes under therapy with the protease inhibitors ritonavir (RTV) and indinavir (IDV) respectively.

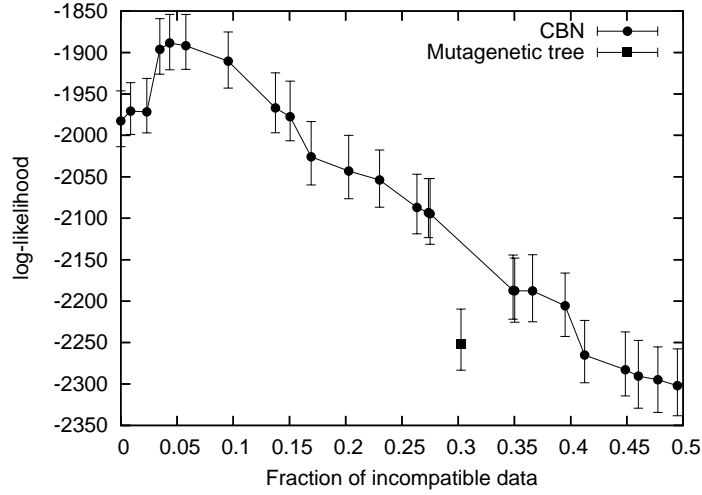
Previous studies identified correlations and preferred pathways among the resistance mutations (Condra et al., 1996; Molla et al., 1996). In particular, in Beerenwinkel et al. (2006) we used mutagenetic trees to infer the underlying dependency structure. The posets are displayed in Figure 2.

For each dataset, we built posets \mathcal{E}_ϵ for various values of ϵ . For each estimated poset, we report two numbers: the log-likelihood of the data given the mixture model \mathcal{E}'_ϵ and the mixture parameter $\hat{\lambda}$ (i.e., the fraction of the data which was explained by the model \mathcal{E}_ϵ). We also calculated these numbers for the mutagenetic trees (Figure 2). These results are shown in Figure 3. Software for building the posets \mathcal{E}_ϵ and computing the likelihood is available at <http://bio.math.berkeley.edu/CBN/>.

The two posets that maximize ℓ'_u for RTV and IDV, respectively, are displayed in Figure 4. Notice that almost all CBNs \mathcal{E}_ϵ constructed performed better than the mutagenetic trees. In order to estimate the significance of this difference, we repeated the log-likelihood calculation for each poset using 1000 bootstrap samples from the original data. The difference in log-likelihood between these optimal posets and the mutagenetic tree induced posets is large enough that their distributions derived from the bootstrap analysis never overlapped. Thus the difference between the optimal CBN models and the mutagenetic trees is found to be highly significant.



(a) RTV



(b) IDV

Figure 3: Log-likelihood ℓ'_u for the CBN models \mathcal{E}_ϵ (filled circles) for various choices of the error tolerance ϵ as a function of the fraction of incompatible genotypes $g \notin \mathcal{G}_\epsilon$. The filled squares correspond to the trees shown in Figure 2. Quartile bars have been derived from 1000 bootstrap samples. Subfigures correspond to (a) ritonavir (RTV) and (b) indinavir (IDV).



Figure 4: Conjunctive Bayesian networks for (a) ritonavir (RTV) and (b) indinavir (IDV) that maximize the likelihood ℓ'_u . The posets represent the models corresponding to the maxima of the graphs shown in Figure 3.

Comparing the optimal CBNs (Figure 4) to the mutagenetic trees (Figure 2) suggests that the mutagenetic trees may induce too many relations and may be handicapped by the requirement that the output is a tree. The posets for RTV share two relations ($V82A < M46I$ and $V82A < I54V$), while those for IDV share none. The RTV poset (Figure 4a) includes the conjunction that both mutations K20R and V82A must occur before I54V, which can not be represented in a mutagenetic tree. By contrast, the IDV poset (Figure 4b) could be represented by a mutagenetic tree, but this tree has not been found by the tree building procedure of Desper et al. (1999). Although the posets and trees do not share many relations, they display a similar structure in that the development of ritonavir resistance is a much more ordered process than for indinavir (see also Beerenwinkel et al., 2006).

This comparison suggests several advantages of CBNs. First, they provide better model fits than the posets derived from the mutagenetic tree models. Second, they rely on a ML method both for parameter estimation and for model selection. This stands in contrast to the algorithm of Desper et al. (1999), which is not an ML procedure. Finally, the perturbed CBNs \mathcal{E}_ϵ can cover a wide range of fractions of unexplained samples, providing a “parametric” picture of the relations present in the data.

5 Algebraic study of the CBN model

In this final section we study CBNs from the perspective of algebraic statistics. Following Pachter and Sturmfels (2005), we regard a CBN as an algebraic variety in a space of dimension $|\mathcal{G}|$. The objective is to compute the prime ideal of all polynomials which vanish on this variety. These poly-

mials are the *algebraic invariants* of the CBN model.

Example 9. For the model with four events and seven genotypes in Example 1, the algebraic invariants are generated by the three polynomials

$$p_{123} \cdot p_{124} - p_{12} \cdot p_{1234}, \quad p_1 \cdot p_2 - p_\emptyset \cdot p_{12} - p_\emptyset \cdot p_{123} - p_\emptyset \cdot p_{124} - p_\emptyset \cdot p_{1234},$$

and $p_\emptyset + p_1 + p_2 + p_{12} + p_{123} + p_{124} + p_{1234} - 1.$

Indeed, these three expressions vanish identically if we replace p_g by $P_g(\theta)$, for each genotype g . Here “are generated” means that every other polynomial with this property is a linear combination of the three polynomials. \square

The main theorem in this section exhibits an explicit Gröbner basis for the algebraic invariants of any CBN model. This Gröbner basis consists of a set of quadratic polynomials together with the trivial invariant $\sum_{g \in \mathcal{G}} p_g - 1$, just like in Example 9. For the special case where \mathcal{E} is a forest, this result was proven in Beerenwinkel and Drton (2005, Theorem 14.11). Other widely used statistical models have Gröbner bases of the same form, for example, decomposable Markov random fields (Geiger et al., 2006) and Jukes-Cantor models in phylogenetics (Sturmfels and Sullivant, 2005). Among Markov random fields, having a Gröbner basis of quadrics is equivalent to having ML degree one (Geiger et al., 2006, Theorem 4.4). This suggests a possible relationship between Theorem 5 and Theorem 10 below. The analogy to Jukes-Cantor models is noteworthy. These models become toric varieties after a linear change of coordinates, known as the Fourier transform or Hadamard conjugation. The same property will be shown in Corollary 11 for the CBN models, but now the role of the Fourier transform is played by Möbius inversion on the distributive lattice \mathcal{G} .

To state our algebraic results, we regard the probabilities p_g , for each genotype g in $\mathcal{G} = J(\mathcal{E})$, as unknowns. These generate the polynomial ring

$$\mathbb{R}[\mathcal{G}] = \mathbb{R}[p_g : g \in \mathcal{G}].$$

In this ring, we consider the prime ideal $I_{\mathcal{E}}$ consisting of all polynomials that vanish on the family of probability distributions defined by the CBN model \mathcal{E} . Equivalently, $I_{\mathcal{E}}$ is the kernel of the ring map $\mathbb{R}[\mathcal{G}] \rightarrow \mathbb{R}[\mathcal{E}]$, $p_g \mapsto P_g(\theta)$, where $\mathbb{R}[\mathcal{E}]$ is the polynomial ring generated by the parameters θ_e , $e \in \mathcal{E}$.

We fix a linear extension of the reverse inclusion order on \mathcal{G} , where $g = \emptyset$ is the largest element and $g = \mathcal{E}$ is the smallest element. We define \prec to be the degree reverse lexicographic monomial ordering on $\mathbb{R}[\mathcal{G}]$ induced by the variable ordering given by the fixed linear extension.

Theorem 10. *The reduced Gröbner basis of the ideal $I_{\mathcal{E}}$ with respect to the monomial ordering \prec consists of the trivial linear invariant $\sum_{g \in \mathcal{G}} p_g - 1$, with leading term p_{\emptyset} , together with one homogeneous quadratic polynomial*

$$\underline{p_g \cdot p_h} - p_{g \cup h} \cdot p_{g \cap h} + \prec\text{-lower terms} \quad (6)$$

for each incomparable pair of genotypes $\{g, h\}$ in the distributive lattice \mathcal{G} .

Proof. We start our proof of Theorem 10 by recalling that the sum of the polynomials $P_g(\theta)$ equals one. If we take the subsum of all polynomials $P_g(\theta)$ where g runs over all genotypes containing some fixed genotype $h \in \mathcal{G}$ then this is essentially the same sum with \mathcal{E} replaced by $\mathcal{E} \setminus h$, and we conclude:

$$\sum_{g: h \subseteq g} P_g(\theta) = \prod_{e \in h} \theta_e.$$

This expression represents the probability that each event in h has happened. The identity suggests that we perform the following linear change of coordinates in the polynomial ring $\mathbb{R}[\mathcal{G}]$:

$$q_h := \sum_{g: h \subseteq g} p_g \quad (\text{for all } h \in \mathcal{G}). \quad (7)$$

Thus, in the new coordinates q_h , the CBN model is precisely the toric variety associated with the distributive lattice \mathcal{G} . A well-known theorem of Hibi (1987) states that the ideal of this toric variety is generated by the binomials

$$\underline{q_g \cdot q_h} - q_{g \cup h} \cdot q_{g \cap h}, \quad (8)$$

where $\{g, h\}$ runs over all incomparable pairs of elements of \mathcal{G} . Moreover, these binomials form a Gröbner basis with the underlined terms as the leading terms. Thus $I_{\mathcal{E}}$ is generated by the quadrics (8) together with the relation $q_{\emptyset} - 1$, which is obtained from (7) under the assumption that the probabilities p_g sum to one. Now, if we rewrite (8) in terms of the original coordinates p_g then we get quadrics of the form (6).

We claim that the quadrics (8) in the original coordinates p_g form a Gröbner basis for $I_{\mathcal{E}}$. This Gröbner basis will be minimal but not reduced. We shall verify the Gröbner basis property by using the theory of sagbi bases (or canonical bases), as described by Sturmfels (1996, §11).

Let \prec denote a negative degree monomial ordering on the polynomial ring of parameters, $\mathbb{R}[\mathcal{E}] = \mathbb{R}[\theta_e : e \in \mathcal{E}]$. Thus \prec is a local monomial ordering in which $1 = \theta^{\emptyset}$ is the largest monomial, and monomials of higher

total degree are \prec -smaller than monomials of lower total degree. See Greuel and Pfister (2002) for an introduction to local monomial orderings.

We shall prove that the coordinate polynomials $P_g(\theta)$ of the CBN model form a sagbi basis for the local ordering \prec , i.e., the \prec -leading monomials

$$\text{in}_{\prec}(P_g(\theta)) = \prod_{e \in g} \theta_e \quad (9)$$

generate the algebra of all \prec -leading monomials of polynomials in the image of our ring map $\mathbb{R}[\mathcal{G}] \rightarrow \mathbb{R}[\mathcal{E}]$, $p_g \mapsto P_g(\theta)$. Let $J_{\mathcal{E}}$ be the prime ideal in $\mathbb{R}[\mathcal{G}]$ consisting of all algebraic relations on the initial monomials (9). By Hibi's result, $J_{\mathcal{E}}$ is generated by the binomials $p_g \cdot p_h - p_{g \cup h} \cdot p_{g \cap h}$, and these binomials form the reduced Gröbner basis of $J_{\mathcal{E}}$ with respect to \prec .

Let $w \in \mathbb{R}^{\mathcal{E}}$ be a weight vector which represents the local ordering \prec for the coordinate polynomials $P_g(\theta)$, and let $\mathcal{A}^T w$ be the induced weight vector in $\mathbb{R}^{\mathcal{G}}$. By Sturmfels (1996, Lemma 11.2) we have

$$\text{in}_{\mathcal{A}^T w}(I_{\mathcal{E}}) \subseteq J_{\mathcal{E}}. \quad (10)$$

Importantly, $p_g \cdot p_h - p_{g \cup h} \cdot p_{g \cap h}$ is the initial form of (8) with respect to $\mathcal{A}^T w$, so the reverse inclusion holds as well. Thus, equality holds in (10), and the desired sagbi basis property holds by Sturmfels (1996, Theorem 11.4).

By Sturmfels (1996, Corollary 11.6(a)) we conclude that the quadratic model invariants (8) form a Gröbner basis of $I_{\mathcal{E}}$ with respect to \prec . This Gröbner basis is minimal and it can be transformed into the reduced Gröbner basis by autoreduction. This completes the proof of Theorem 10. \square

A few remarks are in order. The linear transformation between the p -coordinates and the q -coordinates on the polynomial ring $\mathbb{R}[\mathcal{G}]$, given in (7), is precisely the *Möbius inversion* on the genotype lattice \mathcal{G} . Equivalently, the coefficients (+1, -1, or 0) in the monomials in the expanded model coordinates $P_g(\theta)$ are precisely the values of the Möbius function on \mathcal{G} .

Example 9 illustrates Theorem 10 for the model in Example 1. The three model invariants form a Gröbner basis with leading terms $p_{123} \cdot p_{124}$, $p_1 \cdot p_2$, and p_{\emptyset} . Möbius inversion on the genotype lattice pictured in Figure 1 gives

$$\begin{aligned} p_{\emptyset} &= q_{\emptyset} - q_1 - q_2 + q_{12}, & p_1 &= q_1 - q_{12}, & p_2 &= q_2 - q_{12}, \\ p_{12} &= q_{12} - q_{123} - q_{124} + q_{1234}, & p_{1234} &= q_{1234}, \\ p_{123} &= q_{123} - q_{1234}, & p_{124} &= q_{124} - q_{1234}. \end{aligned}$$

If we perform these substitutions in the reduced Gröbner basis listed in Example 9, then the three given model invariants simplify to

$$q_{\emptyset} - 1, \quad q_1 \cdot q_2 - q_{\emptyset} \cdot q_{12}, \quad q_{123} \cdot q_{124} - q_{1234} \cdot q_{12}.$$

Corollary 11. *The Möbius inversion (7) on the distributive lattice $\mathcal{G} = J(\mathcal{E})$ is a linear change of coordinates which identifies the CBN model \mathcal{E} with the toric variety of the distributive lattice \mathcal{G} defined by Hibi.*

We close with the remark that the sagbi basis property of the coordinate polynomials of the CBN model, which was established in the course of proving Theorem 10, can be used to express any polynomial in the coordinate subalgebra rapidly in terms of the generators $P_g(\theta)$. This process, which is known as *subduction* (Sturmfels, 1996, Algorithm 11.1) generalizes the classical procedure of expressing any symmetric polynomial in terms of elementary symmetric functions, and it may be of interest to statisticians.

Acknowledgements

Part of this work was done while Niko Beerenwinkel was at UC Berkeley and supported by the DFG (BE 3217/1-1). Nicholas Eriksson was supported by the NSF (DMS-0603448) and Bernd Sturmfels was partially supported by the NSF (DMS-0456960) and by DARPA (HR0011-05-1-0057).

References

- S. Andersson and M. Perlman. Lattice models for conditional independence in a multivariate normal distribution. *The Annals of Statistics*, 21(3): 1318–1358, 1993.
- S. Andersson, D. Madigan, M. Perlman, and C. Triggs. On the relation between conditional independence models determined by finite distributive lattices and by directed acyclic graphs. *J. Statist. Planning Inference*, 48:25–46, 1995.
- N. Beerenwinkel and M. Drton. Mutagenetic tree models. In L. Pachter and B. Sturmfels, editors, *Algebraic Statistics for Computational Biology*, chapter 14, pages 278–290. Cambridge University Press, Cambridge, UK, 2005.
- N. Beerenwinkel and M. Drton. A mutagenetic tree hidden Markov model for longitudinal clonal HIV sequence data. *Biostatistics*, 8(1):53–71, Jan 2007. URL <http://dx.doi.org/10.1093/biostatistics/kxj033>.
- N. Beerenwinkel, M. Däumer, T. Sing, J. Rahnenführer, T. Lengauer, J. Selbig, D. Hoffmann, and R. Kaiser. Estimating HIV evolutionary path-

- ways and the genetic barrier to drug resistance. *J Infect Dis*, 191(11): 1953–1960, Jun 2005a. URL <http://dx.doi.org/10.1086/430005>.
- N. Beerenwinkel, J. Rahnenführer, M. Däumer, D. Hoffmann, R. Kaiser, J. Selbig, and T. Lengauer. Learning multiple evolutionary pathways from cross-sectional data. *J Comput Biol*, 12(6):584–598, 2005b. URL <http://dx.doi.org/10.1089/cmb.2005.12.584>. RECOMB 2004.
- N. Beerenwinkel, N. Eriksson, and B. Sturmfels. Evolution on distributive lattices. *J Theor Biol*, 242(2):409–420, Sep 2006. URL <http://dx.doi.org/10.1016/j.jtbi.2006.03.013>.
- F. Catanese, S. Hoşten, A. Khetan, and B. Sturmfels. The maximum likelihood degree. *Am J Math*, 128:671–697, 2006.
- J. H. Condra, D. J. Holder, W. A. Schleif, O. M. Blahy, R. M. Danovich, L. J. Gabryelski, D. J. Graham, D. Laird, J. C. Quintero, A. Rhodes, H. L. Robbins, E. Roth, M. Shivaprakash, T. Yang, J. A. Chodakewitz, P. J. Deutsch, R. Y. Leavitt, F. E. Massari, J. W. Mellors, K. E. Squires, R. T. Steigbigel, H. Teppler, and E. A. Emini. Genetic correlates of in vivo viral resistance to indinavir, a human immunodeficiency virus type 1 protease inhibitor. *J Virol*, 70(12):8270–8276, Dec 1996.
- R. Desper, F. Jiang, O. P. Kallioniemi, H. Moch, C. H. Papadimitriou, and A. A. Schäffer. Inferring tree models for oncogenesis from comparative genome hybridization data. *J Comput Biol*, 6(1):37–51, 1999.
- M. Drton and T. S. Richardson. Binary models for marginal independence. Technical Report 474, Department of Statistics, University of Washington, 2005. URL <http://www.stat.washington.edu/www/research/reports/2005/tr474.pdf>.
- D. Geiger, C. Meek, and B. Sturmfels. On the toric algebra of graphical models. *Ann Statist*, 34(3):1463–1492, 2006.
- G.-M. Greuel and G. Pfister. *A Singular Introduction to Commutative Algebra*. Springer, 2002.
- T. Hibi. Distributive lattices, affine semigroup rings and algebras with straightening laws. *Adv Stud Pure Math*, 11:93–109, 1987.
- S. Lauritzen. *Graphical Models*. Clarendon Press, 1996.
- C. Meek and D. Heckerman. Structure and parameter learning for causal independence and causal interaction models. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, pages 366–375, 1997.

- F. Michor, Y. Iwasa, and M. A. Nowak. Dynamics of cancer progression. *Nat Rev Cancer*, 4(3):197–205, Mar 2004. URL <http://dx.doi.org/10.1038/nrc1295>.
- A. Molla, M. Korneyeva, Q. Gao, S. Vasavanonda, P. J. Schipper, H. M. Mo, M. Markowitz, T. Chernyavskiy, P. Niu, N. Lyons, A. Hsu, G. R. Granneman, D. D. Ho, C. A. Boucher, J. M. Leonard, D. W. Norbeck, and D. J. Kempf. Ordered accumulation of mutations in HIV protease confers resistance to ritonavir. *Nat Med*, 2(7):760–766, Jul 1996.
- L. Pachter and B. Sturmfels, editors. *Algebraic Statistics for Computational Biology*. Cambridge University Press, Cambridge, UK, 2005. doi: 10.2277/0521857007.
- J. Pearl. Fusion, propagation, and structuring in belief networks. *Artif. Intell.*, 29(3):241–288, 1986. ISSN 0004-3702.
- J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kauffman Publishers, 1988.
- M. D. Radmacher, R. Simon, R. Desper, R. Taetle, A. A. Schäffer, and M. A. Nelson. Graph models of oncogenesis with an application to melanoma. *J Theor Biol*, 212(4):535–548, Oct 2001. doi: 10.1006/jtbi.2001.2395. URL <http://dx.doi.org/10.1006/jtbi.2001.2395>.
- J. Rahnenführer, N. Beerenwinkel, W. A. Schulz, C. Hartmann, A. von Deimling, B. Wullich, and T. Lengauer. Estimating cancer survival and clinical outcome based on genetic tumor progression scores. *Bioinformatics*, 21(10):2438–2446, May 2005. doi: 10.1093/bioinformatics/bti312. URL <http://dx.doi.org/10.1093/bioinformatics/bti312>.
- S.-Y. Rhee, M. J. Gonzales, R. Kantor, B. J. Betts, J. Ravela, and R. W. Shafer. Human immunodeficiency virus reverse transcriptase and protease sequence database. *Nucleic Acids Res*, 31(1):298–303, Jan 2003.
- B. Sturmfels. *Gröbner Bases and Convex Polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, 1996.
- B. Sturmfels and S. Sullivant. Toric ideals of phylogenetic invariants. *J Comput Biol*, 12(4):457–481, May 2005. doi: 10.1089/cmb.2005.12.457. URL <http://dx.doi.org/10.1089/cmb.2005.12.457>.
- A. Szabo and K. Boucher. Estimating an oncogenetic tree when false negatives and positives are present. *Math Biosci*, 176(2):219–236, Apr 2002.