

A Bayesian Approach for Inferring Local Causal Structure in Gene Regulatory Networks

Ioan Gabriel Bucur

G.BUCUR@CS.RU.NL

Tom van Bussel

TOM.VAN.BUSSEL@STUDENT.RU.NL

Tom Claassen

TOMC@CS.RU.NL

Tom Heskes

T.HESKES@SCIENCE.RU.NL

Faculty of Science, Radboud University, Postbus 9010, 6500GL Nijmegen, The Netherlands

Abstract

Gene regulatory networks play a crucial role in controlling an organism's biological processes, which is why there is significant interest in developing computational methods that are able to extract their structure from high-throughput genetic data. A typical approach consists of a series of conditional independence tests on the covariance structure meant to progressively reduce the space of possible causal models. We propose a novel efficient Bayesian method for discovering the local causal relationships among triplets of (normally distributed) variables. In our approach, we score the patterns in the covariance matrix in one go and we incorporate the available background knowledge in the form of priors over causal structures. Our method is flexible in the sense that it allows for different types of causal structures and assumptions. We apply the approach to the task of inferring gene regulatory networks by learning regulatory relationships between gene expression levels. We show that our algorithm produces stable and conservative posterior probability estimates over local causal structures that can be used to derive an honest ranking of the most meaningful regulatory relationships. We demonstrate the stability and efficacy of our method both on simulated data and on real-world data from an experiment on yeast.

Keywords: Causal discovery; Structure learning; Covariance selection; Bayesian inference; Gene regulatory networks.

1. Introduction

Gene regulatory networks (GRNs) play a crucial role in controlling an organism's biological processes, such as cell differentiation and metabolism. If we knew the structure of a GRN, we would be able to intervene in the developmental process of the organism, for instance by targeting a specific gene with drugs. In recent years, researchers have developed a number of methods for inferring regulatory relationships from data on gene expression, the process by which genetic instructions are used to synthesize gene products such as proteins. Gene regulatory relationships are inherently causal: we can manipulate the expression level of one gene (the 'cause') to regulate that of another gene (the 'effect'). Because of this, many GRN inference algorithms such as 'Trigger' (Chen et al., 2007), 'CIT' (Millstein et al., 2009) or 'CMST' (Neto et al. (2013)) are aimed at finding promising causal regulatory relationships among genes.

An efficient way to derive causal relationships from observational data, which results in clear and easily interpretable output, is to find local causal patterns in the data. The *local causal discovery* (LCD) algorithm (Cooper, 1997) makes use of a combination of observational data and background knowledge when searching for unconfounded causal relationships among triplets of variables. Trig-

ger is also designed to search for this LCD pattern in the data, using the background knowledge that genetic information is randomized at birth, before any other measurements can be made. Mani et al. (2006), on the other hand, divide the causal discovery task into identifying so-called Y structures on subsets of four variables. The Y structure is the smallest pattern containing an unconfounded causal relationship that can be learned solely from observational data in the presence of latent variables.

A key feature of Trigger is that it can estimate the probability of causal regulatory relationships, while controlling for the false discovery rate (Chen et al., 2007). The algorithm consists of a series of likelihood ratio tests for regression coefficients that are translated into statements about conditional (in)dependence, which are then used to identify the presence of the LCD pattern. Testing whether regression coefficients are significantly different from zero essentially boils down to testing whether partial correlation coefficients are significantly different from zero (Meinshausen and Bühlmann, 2006), which means that all the information needed for the tests lies in the covariance structure.

We propose a Bayesian approach for local causal discovery on triplets of (normally distributed) variables that makes use of the information in the covariance structure. With our method, we directly score patterns in the data by computing posterior probabilities over all possible three-dimensional covariance structures in one go, with the end goal of identifying plausible causal relationships. This provides a stable, efficient and elegant way of expressing the uncertainty in the underlying local causal structure, even in the presence of latent variables. Moreover, it is straightforward to incorporate background knowledge in the form of priors on causal structures. We show how we can plug in our method into an algorithm that searches for local causal patterns in a GRN and outputs a well-calibrated and reliable ranking of the most likely causal regulatory relationships.

The rest of the paper is organized as follows. In Section 2, we introduce some standard background notation and terminology. In Section 3, we describe our Bayesian approach for inferring the covariance pattern of a three-dimensional Gaussian random vector. By defining simple priors, we then derive the posterior probabilities of local causal structures given the data. In Section 4, we present the results of applying our method on simulated and real-world data. We conclude by discussing advantages and disadvantages of our approach in Section 5.

2. Background

Causal structures can be represented by directed graphs, where the nodes in the graph represent (random) variables and the edges between nodes represent causal relationships. *Maximal ancestral graphs* (MAGs) encode conditional independence information and causal relationships in the presence of latent variables and selection bias (Richardson and Spirtes, 2002). We refer to MAGs without undirected edges as *directed maximal ancestral graphs* (DMAGs). DMAGs are closed under marginalization, which means they preserve the conditional independence information in the presence of latent variables.

Two causal structures are *Markov (independence) equivalent* if they imply the same conditional independence statements. The *Markov equivalence class* of a MAG (or DMAG) is represented by a partial ancestral graph (PAG), which displays all the edge marks (arrowhead or tail) shared by all members in the class and displays circles for those marks that are not common among all members. In this work, we will consider two types of graphs: *directed acyclic graphs* and *directed maximal ancestral graphs*. However, the results presented can be applied to any causal graph structure.

A *(conditional) independence model* \mathcal{I} over a finite set of variables V is a set of triples $\langle X, Y \mid Z \rangle$, called *(conditional) independence statements*, where X, Y, Z are disjoint subsets of V and Z may

be empty (Studeny, 2006). We can induce a (probabilistic) independence model over a probability distribution $P \in \mathcal{P}$ by letting:

$$\langle A, B | C \rangle \in \mathcal{I}(P) \iff A \perp\!\!\!\perp B | C \text{ w.r.t. } P.$$

The conditional independence model induced by a multivariate Gaussian distribution is a *compositional graphoid* (Sadeghi and Lauritzen, 2014), which means that it satisfies the *graphoid axioms* and the *composition* property. Because of this, there is a one-to-one correspondence between the conditional independence models that can be induced by a multivariate Gaussian and the Markov equivalence classes of a causal graph structure.

3. Bayes Factors of Covariance Structures (BFCS)

Model	Markov Equivalence Class (PAG)	Covariance Matrix	Precision Matrix
' $X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3$ ' (Full)			
' $X_1 \perp\!\!\!\perp X_3$ ' (Acausal)			
' $X_1 \perp\!\!\!\perp X_3 X_2$ ' (Causal)			
' $(X_1, X_3) \perp\!\!\!\perp X_2$ ' (Independent)			
' $X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3$ ' (Empty)			

Figure 1: Overview of the five canonical independence patterns between three variables, depicting the equivalence between causal models, conditional independences, and covariance structures

We are interested in inferring the local covariance structure from observational data, assuming the data follows a (latent) Gaussian model. We will be working with triplets of variables. With finite data, we can never be sure about the true covariance structure underlying the data. Hence, we prefer to work with probability distributions over covariance matrices. For a general three-dimensional

covariance matrix Σ , the likelihood reads:

$$p(\mathbf{D} \mid \Sigma) = (2\pi)^{-\frac{3n}{2}} |\Sigma|^{-\frac{n}{2}} \exp \left[-\frac{1}{2} \text{tr} (\mathbf{S}\Sigma^{-1}) \right],$$

where $\mathbf{S} = \mathbf{D}^\top \mathbf{D}$ is the scatter matrix.

Under the Gaussianity assumption, there is a one-to-one correspondence between the constraints in the covariance matrix and the conditional independences among the variables. There are five specific canonical patterns to consider, which are depicted in Figure 1. The ‘full’ and ‘empty’ covariance patterns are self-explanatory. We call ‘independent’ the pattern occurring when one variable is independent of the other two. We call the pattern on the second row ‘acausal’ because X_2 cannot cause X_1 or X_3 if conditioning upon X_2 turns a conditional independence between X_1 and X_3 into a conditional dependence. We call the pattern on the third row ‘causal’ because X_2 either causes X_1 or X_3 if conditioning upon X_2 turns a conditional dependence between X_1 and X_3 into a conditional independence (Claassen and Heskes, 2011). The five patterns translate into eleven distinct covariance structures when considering all permutations of three variables. These are the only possible covariance structures on three variables, since the conditional independence model induced by a multivariate Gaussian is a compositional graphoid (Sadeghi and Lauritzen, 2014).

Our goal is to compute the posterior probability of each of the possible conditional independence models given the data. We denote by $\mathcal{J} = \{\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_{10}\}$ the set of all possible conditional independence models. The model evidence is then, for $\mathcal{M}_j \in \mathcal{J}$:

$$p(\mathbf{D} \mid \mathcal{M}_j) = \int d\Sigma p(\mathbf{D} \mid \Sigma)p(\Sigma \mid \mathcal{M}_j).$$

To facilitate computation, we derive the Bayes factors of each conditional independence model (\mathcal{M}_j) compared to a reference independence model (\mathcal{M}_0):

$$\mathcal{B}_j = \frac{p(\mathbf{D} \mid \mathcal{M}_j)}{p(\mathbf{D} \mid \mathcal{M}_0)} = \frac{\int d\Sigma p(\mathbf{D} \mid \Sigma)p(\Sigma \mid \mathcal{M}_j)}{\int d\Sigma p(\mathbf{D} \mid \Sigma)p(\Sigma \mid \mathcal{M}_0)}.$$

As we shall see in Subsection 3.2, many terms will cancel out, making the resulting ratios much simpler to compute (see for example Equation 5). Finally, we arrive at the posterior probabilities:

$$p(\mathcal{M}_j \mid \mathbf{D}) = \frac{p(\mathbf{D} \mid \mathcal{M}_j) \cdot p(\mathcal{M}_j)}{\sum_j p(\mathbf{D} \mid \mathcal{M}_j) \cdot p(\mathcal{M}_j)} = \frac{\mathcal{B}_j \cdot p(\mathcal{M}_j)}{\sum_j \mathcal{B}_j \cdot p(\mathcal{M}_j)}. \quad (1)$$

3.1 Choosing the Prior on Covariance Matrices

We consider the inverse Wishart distribution for three-dimensional covariance matrices, which is parameterized by the positive definite scale matrix Ψ and the number of degrees of freedom ν :

$$\Sigma \sim \mathcal{W}_3^{-1}(\Psi, \nu); \quad p(\Sigma) = \frac{|\Psi|^{\frac{\nu}{2}}}{2^{\frac{3\nu}{2}} \Gamma_3(\frac{\nu}{2})} |\Sigma|^{-\frac{\nu+4}{2}} \exp \left[-\frac{1}{2} \text{tr} (\Psi\Sigma^{-1}) \right].$$

The inverse Wishart is the conjugate prior on the covariance matrix of a multivariate Gaussian vector, which means the posterior is also inverse Wishart. Given the data set \mathbf{D} containing n observations and $\mathbf{S} = \mathbf{D}^\top \mathbf{D}$ the scatter matrix, the posterior reads:

$$\Sigma \mid \mathbf{D} \sim \mathcal{W}_3^{-1}(\Psi + \mathbf{S}, \nu + n). \quad (2)$$

In order to choose appropriate parameters for the inverse Wishart prior, we analyze the implied distribution in the space of correlation matrices. By transforming the covariance matrix into a correlation matrix, we end up with a so-called *projected inverse Wishart* distribution on the latter, which we denote by \mathcal{PW}^{-1} . Barnard et al. (2000) have shown that if the correlation matrix \mathbf{R} follows a projected inverse Wishart distribution with scale parameter Ψ and ν degrees of freedom, then the marginal distribution $p(R_{ij})$, $i \neq j$, for off-diagonal elements is uniform if we take Ψ to be any diagonal matrix and $\nu = p + 1$, where p is the number of variables. We are working with three variables, so we choose $\nu = 4$.

It is easy to check that for any diagonal matrix D , the projected inverse Wishart is scale invariant:

$$\mathcal{PW}^{-1}(\Psi, \nu) \equiv \mathcal{PW}^{-1}(D\Psi D, \nu).$$

From (2), it then follows that we can make the posterior distribution on the correlation matrices independent of the scale of the data by choosing the prior scale matrix $\Psi = \mathbf{0}_{3,3}$. Since that would lead to an undefined prior distribution, we can achieve the same goal by setting $\Psi = \epsilon \mathbf{I}_3$ in the limit $\epsilon \downarrow 0$, where \mathbf{I} is the identity matrix. Summarizing, we will consider the prior distribution:

$$\Sigma \sim \mathcal{W}_3^{-1}(\epsilon \mathbf{I}_3, 4), \quad \epsilon \downarrow 0.$$

3.2 Deriving the Bayes Factors

As reference model (\mathcal{M}_0), we choose the most general case in which no independences can be found in the data ($X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3$), which means that the covariance matrix is unconstrained (Figure 1, first row). We assume that the covariance matrix follows an inverse Wishart distribution

$$p(\Sigma \mid X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3) = \mathcal{W}_3^{-1}(\Sigma; \epsilon \mathbf{I}_3, \nu),$$

where we consider the limit $\epsilon \downarrow 0$ and set $\nu = 4$ (see Subsection 3.1). Using the conjugacy of the inverse Wishart prior for the covariance matrix, we immediately get the model evidence

$$p(\mathbf{D} \mid X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3) = \frac{\epsilon^{\frac{3\nu}{2}} \Gamma_3(\frac{n+\nu}{2})}{\pi^{\frac{3n}{2}} \Gamma_3(\frac{\nu}{2})} |\mathbf{S} + \epsilon \mathbf{I}_3|^{-\frac{n+\nu}{2}}, \quad (3)$$

where Γ_p is the p -variate gamma function and $\mathbf{S} = \mathbf{D}^\top \mathbf{D}$ is the scatter matrix. We first compare the evidence for the conditional independence model $X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3$ to the evidence for the reference model $X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3$ by computing the Bayes factor:

$$\mathcal{B}(X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3) = \frac{p(\mathbf{D} \mid X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3)}{p(\mathbf{D} \mid X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3)}.$$

We can implement the $X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3$ case (Figure 1, last row) by constraining Σ to be diagonal, which means we only have to consider the parameters $\Sigma_{11}, \Sigma_{22}, \Sigma_{33}$. We propose to take:

$$p(\Sigma \mid X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3) = \prod_{i=1}^3 \mathcal{W}_1^{-1}(\Sigma_{ii}; \epsilon, \nu). \quad (4)$$

The likelihood also factorizes in this case and becomes

$$p(\mathbf{D} \mid \Sigma) = \prod_{i=1}^3 \left\{ (2\pi \Sigma_{ii})^{-\frac{n}{2}} \exp \left[-\frac{1}{2} \text{tr} (S_{ii} \Sigma_{ii}^{-1}) \right] \right\},$$

yielding the model evidence

$$p(\mathbf{D} \mid X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3) = \prod_{i=1}^3 \left[\frac{\epsilon^{\frac{\nu}{2}} \Gamma_1(\frac{n+\nu}{2})}{\pi^{\frac{n}{2}} \Gamma_1(\frac{\nu}{2})} (S_{ii} + \epsilon)^{-\frac{n+\nu}{2}} \right].$$

Dividing by the model evidence from (3) and taking the limit $\epsilon \downarrow 0$, we obtain the Bayes factor

$$\mathcal{B}(X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3) = \frac{\Gamma_3(\frac{\nu}{2})}{\Gamma_3(\frac{n+\nu}{2})} \left[\frac{\Gamma_1(\frac{n+\nu}{2})}{\Gamma_1(\frac{\nu}{2})} \right]^3 |\mathbf{C}|^{-\frac{n+\nu}{2}} = \frac{n + \nu - 2}{\nu - 2} \frac{\Gamma(\frac{\nu-1}{2})}{\Gamma(\frac{n+\nu}{2})} \frac{\Gamma(\frac{\nu}{2})}{\Gamma(\frac{n+\nu-1}{2})} |\mathbf{C}|^{-\frac{n+\nu}{2}}, \quad (5)$$

with \mathbf{C} the sample correlation matrix and Γ the (univariate) gamma function. Due to the choice (4), the evidence for ‘ $X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3$ ’ also scales with $\epsilon^{\frac{\nu}{2}}$, so the dominant terms depending on ϵ cancel out and the Bayes factor depends only on the correlation matrix in the limit $\epsilon \downarrow 0$. The derivations for the other cases (left out due to space constraints) are similar, leading to the Bayes factors:

$$\begin{aligned} \mathcal{B}(X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3) &= f(n, \nu) g(n, \nu) |\mathbf{C}|^{\frac{n+\nu}{2}} \\ \mathcal{B}(X_3 \perp\!\!\!\perp (X_1, X_2)) &= f(n, \nu) \left(\frac{|\mathbf{C}|}{1 - C_{12}^2} \right)^{\frac{n+\nu}{2}} \\ \mathcal{B}(X_1 \perp\!\!\!\perp X_2 \mid X_3) &= g(n, \nu) \left(\frac{|\mathbf{C}|}{(1 - C_{13}^2)(1 - C_{23}^2)} \right)^{\frac{n+\nu}{2}} \\ \mathcal{B}(X_1 \perp\!\!\!\perp X_2) &= \frac{f(n, \nu)}{g(n, \nu)} (1 - C_{12}^2)^{\frac{n+\nu-1}{2}}, \end{aligned} \quad (6)$$

where $f(n, \nu) = \frac{n + \nu - 2}{\nu - 2}$ and $g(n, \nu) = \frac{\Gamma(\frac{n+\nu}{2}) \Gamma(\frac{\nu-1}{2})}{\Gamma(\frac{n+\nu-1}{2}) \Gamma(\frac{\nu}{2})} \approx \left(\frac{2n + 2\nu - 3}{2\nu - 3} \right)^{\frac{1}{2}}$.

In conclusion, for deriving the Bayes factors in (6), we only need to plug in the correlation matrix with the number of samples and compute a limited number of closed-form terms. This is then sufficient to obtain the full posterior distribution over the covariance structures, which makes the BFCS method fast and efficient.

3.3 Priors on Causal Structures

To do a full Bayesian analysis, we need to specify priors over the different conditional independence models. Assuming faithfulness, there is a one-to-one correspondence between the Markov equivalence classes of the underlying causal graph structure and the conditional independence models. By taking a uniform prior over causal graphs and denoting by $|\mathcal{M}_j|$ the number of causal graphs consistent with the independence model $\mathcal{M}_j \in \mathcal{J}$, we arrive at the prior:

$$p(\mathcal{M}_j) = \frac{|\mathcal{M}_j|}{\sum_j |\mathcal{M}_j|}, \quad \forall \mathcal{M}_j \in \mathcal{J}.$$

In Table 1 we count the number of DAGs and DMAGs (see Section 2) consistent with each covariance pattern (see Figure 1). For example, if we assume an underlying DAG structure, then $p(X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3) = \frac{1}{25}$. The addition of background knowledge (BK) reduces the number of causal graph structures corresponding to each conditional independence model. Specifically relevant for

Pattern	CI Model	Description	DAG	DAG w/ BK	DMAG	DMAG w/ BK
Full	\mathcal{M}_0	$X_1 \not\perp\!\!\!\perp X_2 \not\perp\!\!\!\perp X_3$	6	2	19	3
Acausal	\mathcal{M}_1	$X_1 \perp\!\!\!\perp X_2$	1	1	3	2
	\mathcal{M}_2	$X_2 \perp\!\!\!\perp X_3$	1	0	3	0
	\mathcal{M}_3	$X_3 \perp\!\!\!\perp X_1$	1	1	3	2
Causal	\mathcal{M}_4	$X_1 \perp\!\!\!\perp X_2 \mid X_3$	3	1	5	1
	\mathcal{M}_5	$X_2 \perp\!\!\!\perp X_3 \mid X_1$	3	1	5	1
	\mathcal{M}_6	$X_3 \perp\!\!\!\perp X_1 \mid X_2$	3	1	5	1
Independent	\mathcal{M}_7	$X_1 \perp\!\!\!\perp (X_2, X_3)$	2	2	3	3
	\mathcal{M}_8	$X_2 \perp\!\!\!\perp (X_3, X_1)$	2	1	3	1
	\mathcal{M}_9	$X_3 \perp\!\!\!\perp (X_1, X_2)$	2	1	3	1
Empty	\mathcal{M}_{10}	$X_1 \perp\!\!\!\perp X_2 \perp\!\!\!\perp X_3$	1	1	1	1
All			25	12	53	16

Table 1: Number of causal graph structures over three variables in each Markov equivalence class. In the columns marked ‘w /BK’, the background knowledge that X_1 precedes all other variables, i.e., there can be no arrowhead towards X_1 , is added when counting the number of structures.

discovering causal regulatory relationships is the background knowledge that the genetic marker precedes the expression traits, i.e., that X_1 precedes all other variables. This additional constraint leads to the counts in the columns marked ‘w/ BK’ in Table 1. Some of the covariance patterns imply acausal / causal statements (Figure 1), which is what allows us to directly translate the posterior probabilities over covariance patterns into statements over causal relationships.

Now that we have defined priors on the conditional independence models (Markov equivalence classes), we can derive the posterior probabilities from equations (1) and (6). The posterior probabilities could also be derived by combining the Bayesian Gaussian equivalent (BGe) score (Geiger and Heckerman, 1994) with the priors on causal structures defined in this subsection. Due to our choice of priors on covariance matrices (see Subsection 3.1), however, the Bayes factors are simpler to compute. This makes our approach more efficient when used to infer causal relationships in large regulatory networks.

To summarize, we have developed a method for computing the posterior probabilities of the covariance structures over three variables (BFCS). We will employ this procedure as part of an algorithm for discovering regulatory relationships. Similarly to LCD and Trigger, the idea is to search over triplets of variables to find potential local causal structures (see Algorithm 1).

4. Experimental Results

4.1 Consistency of Detecting Local Causal Structures

In this simulation, we assessed how well our BFCS approach is able to detect the causal structure $X_1 \rightarrow X_2 \rightarrow X_3$, which is crucial to the application of the LCD and Trigger algorithms. We considered the three generating structures depicted in Figure 2. In all three cases, the variables are mutually marginally dependent, but only in the first model $X_1 \perp\!\!\!\perp X_3 \mid X_2$ holds.

We sampled random structural parameters (interaction strengths) for each causal relationship from independent standard normal distributions. Then, given each generating model, we generated

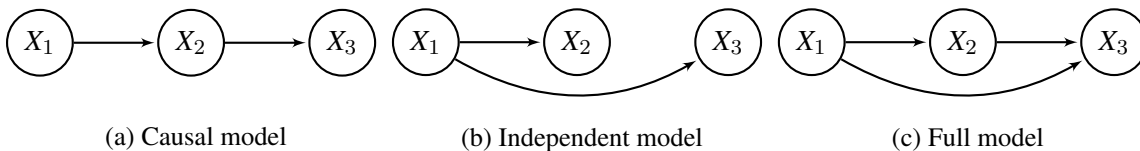


Figure 2: Generating models

data sets of different sizes (from 10^2 up to 10^6 samples). For each data set, we computed the correlation matrix, which we plugged into BFCS for computing the Bayes factors. We repeated this procedure for 1000 different parameter configurations. We assumed that X_1 precedes all other variables and we allowed for latent variables, so we did not use the knowledge that the data is causally sufficient. We considered a uniform prior over the twelve possible DMAG structures (see Table 1).

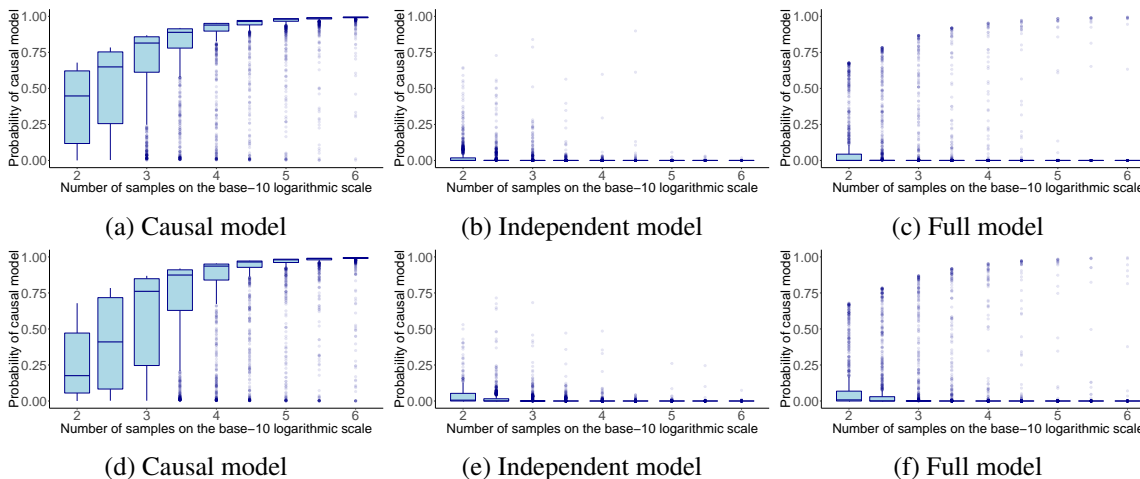


Figure 3: Box plots of the posterior probabilities of $p(X_1 \rightarrow X_2 \rightarrow X_3 \mid \mathbf{D})$ output by BFCS across the 1000 different parameter configurations for each of the generating models in Figure 2. As we increased the number of samples, the probability $p(X_1 \rightarrow X_2 \rightarrow X_3 \mid \mathbf{D})$ converged to one when the data was generated from the causal model (a) and converged to zero when it was not (b and c). **Top:** $\mathbf{x} = (X_1, X_2, X_3)$ is multivariate Gaussian; **Bottom:** X_1 is Bernoulli and $(X_2, X_3) \mid X_1$ is Gaussian.

In the first experiment we generated random multivariate data from the models in Figure 2. As expected, the posterior probability $p(X_1 \rightarrow X_2 \rightarrow X_3 \mid \mathbf{D})$ converged to one (Figure 3, top row) when the true generating model was the one in Figure 2a. At the same time, $p(X_1 \rightarrow X_2 \rightarrow X_3 \mid \mathbf{D})$ converged to zero when the true generating model was the independent or full model. Note that it is easier to distinguish the causal model from the independent model than from the full model. When generating data from the full model, it is possible to generate structural parameters that are close to zero. If the direct interaction between X_1 and X_3 is close to zero, it looks as if $X_1 \perp\!\!\!\perp X_3 \mid X_2$.

In the second experiment, we considered the same generating models, but we sampled X_1 from a Bernoulli distribution to mimic a genetic variable, e.g., an allele or the parental strain in the yeast experiment (Chen et al., 2007). We sampled a different success probability for the Bernoulli variable in each repetition from a uniform distribution between 0.1 and 0.9. The random vector (X_2, X_3) was sampled from a bivariate Gaussian conditional on the value of X_1 . In the bottom row of Figure 3 we

see that when the Gaussian assumption did not hold for X_1 , we lost some power in recovering the correct model. When we increased the number of samples, this loss in power due to the violation of the Gaussian assumption became less severe and BFCS remained consistent.

4.2 Causal Discovery in Gene Regulatory Networks

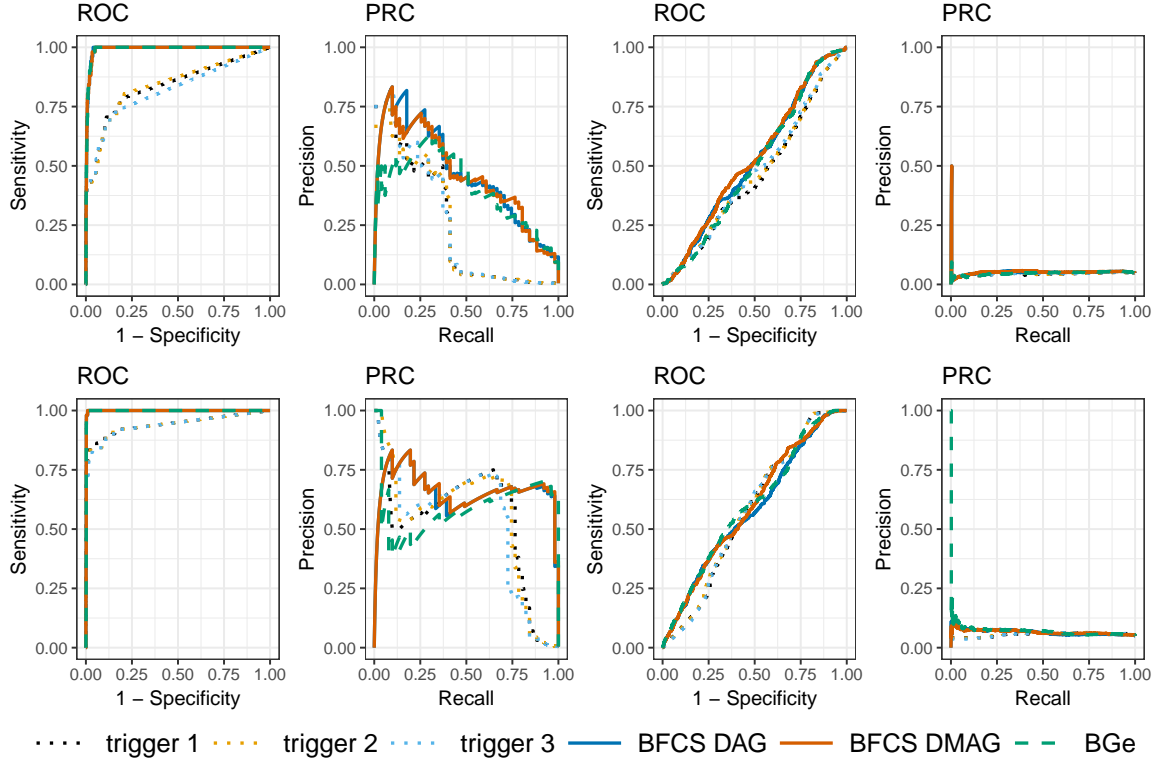


Figure 4: Comparing the performance of Trigger and BFCS in terms of the ROC and Precision-Recall (PRC) curves. We generated both a sparse GRN (**left column**) and a dense GRN (**right column**) consisting of 51 and 491 regulatory relationships, respectively. We generated 100 samples (**top row**) and then 1000 samples (**bottom row**) from both networks. We ran Trigger three times (labeled ‘trigger 1’, ‘trigger 2’, and ‘trigger 3’) on the simulated data to account for differences when sampling the null statistics. We compared Trigger against two versions of BFCS, in which we took a uniform prior over DAGs (‘BFCS DAG’) and DMAGs (‘BFCS DMAG’), respectively. For reference, we also show the performance of an equivalent method that uses the BGe score (‘BGe’).

In this experiment, we simulated a transcriptional regulatory network meant to emulate the yeast data set analyzed in Chen et al. (2007). We randomly generated data for 100 genetic markers, where each marker is an independent Bernoulli variable with ‘success’ probability uniformly sampled between 0.1 and 0.5. We then generated transcript level expression data from the structural equation model:

$$\mathbf{t} := \mathbf{B}\mathbf{t} + \mathbf{l} + \varepsilon,$$

where \mathbf{B} is a lower triangular matrix, $\mathbf{t} = (T_1, T_2, \dots, T_{100})^\top$ is the random vector of the expression trait data, $\mathbf{l} = (L_1, L_2, \dots, L_{100})^\top$ is the random vector of the genetic markers, and $\varepsilon \sim$

$\mathcal{N}(\mathbf{0}_{100}, \mathbf{I}_{100})$ is added noise. The true causal relationships are in the directed graph structure defined by \mathbf{B} . Each expression trait T_i was causally linked to the genetic marker $L_i, \forall i \in \{1, 2, \dots, 100\}$.

For each pair of expression traits (T_i, T_j) and for every genetic marker L_k , we derived the posterior probability of $L_k \rightarrow T_i \rightarrow T_j$ using BFCS. We then took the maximum of these probabilities over k , which is a lower bound of the probability that there exists L_k such that the triple (L_k, T_i, T_j) has the causal structure $L_k \rightarrow T_i \rightarrow T_j$. Similarly to Trigger, we reported this value as a conservative estimate for the probability of $T_i \rightarrow T_j$. For reference, we also evaluated the performance of an equivalent method that uses the BGe score instead of the Bayes factors computed in BFCS.

In Figure 4, we compare the performance of Trigger and BFCS by showing the ROC and Precision-Recall (PRC) curves. Both Trigger and BFCS performed much better when the underlying network structure was sparse. For sparse networks, BFCS shows a significant improvement in the AUC measure for both curves. For dense networks, Trigger and BFCS have fairly similar ROC curves, but BFCS shows a noticeable improvement in the precision-recall curve. More specifically, the higher precision for the first causal relationships that are recalled shows that BFCS is better at ranking the top regulatory relationships. Furthermore, the probabilities output by BFCS are better calibrated than those by Trigger (see Table 2). This is because Trigger is overconfident in its predictions, which explains the large spread in probabilities compared to BFCS.

Samples	Sparse GRN				Dense GRN			
	trigger	BFCS DAG	BFCS DMAG	BGe	trigger	BFCS DAG	BFCS DMAG	BGe
100	0.037	0.022	0.015	0.009	0.481	0.201	0.148	0.231
1000	0.028	0.011	0.008	0.006	0.343	0.376	0.310	0.445

Table 2: Comparing the calibration of Trigger and BFCS using the Brier score (lower is better).

4.3 Comparing Results from an Experiment on Yeast

Algorithm 1 Running BFCS on the yeast data set

- 1: *Input*: Yeast data set consisting of 3244 markers and 6216 gene expression measurements
 - 2: **for all** expression traits T_i **do**
 - 3: **for all** expression traits $T_j, j \neq i$ **do**
 - 4: **for all** genetic markers L_k **do**
 - 5: Compute the Bayes factors for the triplet (L_k, T_i, T_j)
 - 6: Derive the posterior probability of the structure $L_k \rightarrow T_i \rightarrow T_j$ given the data
 - 7: **end for**
 - 8: Save $\max_k p(L_k \rightarrow T_i \rightarrow T_j)$ as the probability of gene i regulating gene j
 - 9: **end for**
 - 10: **end for**
 - 11: *Output*: Matrix of regulation probabilities
-

Chen et al. (2007) showcased the Trigger algorithm by applying it to an experiment on yeast, in which two distinct strains were crossed to produce 112 independent recombinant segregant lines. Genome-wide genotyping and expression profiling were performed on each segregant line. We computed the probabilities over the triples in the yeast data set using `trigger` (Chen et al., 2017) and BFCS (Algorithm 1). For BFCS, we used DMAGs to allow for the possibility of latent variables.

Rank	Gene	Chen et al.	trigger	BFCS	Rank	Gene	Chen et al.	trigger	BFCS
1	MDM35	0.973	0.999	0.678	1	FMP39	0.176	0.401	0.691
2	CBP6	0.968	0.997	0.683	2	DIA4	0.493	0.987	0.691
3	QRI5	0.960	0.985	0.678	3	MRP4	0.099	0.260	0.691
4	RSM18	0.959	0.984	0.672	4	MNP1	0.473	0.999	0.691
5	RSM7	0.953	0.977	0.684	5	MRPS18	0.527	0.974	0.690
6	MRPL11	0.924	0.999	0.670	6	MTG2	0.000	0.000	0.690
7	MRPL25	0.887	0.908	0.675	7	YNL184C	0.299	0.768	0.690
8	DLD2	0.871	0.896	0.660	8	YPL073C	0.535	0.993	0.690
9	YPR126C	0.860	0.904	0.634	9	MBA1	0.290	0.591	0.690
10	MSS116	0.849	0.997	0.659	10	ACN9	0.578	0.927	0.690

(a) Genes regulated by NAM9, sorted by ‘Chen et al.’. (b) Genes regulated by NAM9, sorted by ‘BFCS’.

Table 3: The column ‘Chen et al.’ shows the original results of the Trigger algorithm as reported in Chen et al. (2007). The ‘trigger’ column contains the probabilities we obtained when running the algorithm from the Bioconductor `trigger` package (Chen et al., 2017) on the entire yeast data set with default parameters. The column ‘BFCS’ contains the output of running Algorithm 1 on the yeast data set, for which we took a uniform prior over DMAGs.

In Table 3a, we report the top ten genes purported to be regulated by the putative regulator NAM9, sorted according to the probability estimates reported in Chen et al. (2007). We see that BFCS also assigns relatively high, albeit better calibrated (Table 2) and much more conservative, probabilities to the most significant regulatory relationships found by Trigger. In Table 3b, we see that some relationships ranked significant by BFCS are assigned a very small probability by Trigger. The regulatory relationship $\text{NAM9} \rightarrow \text{MTG2}$, of which both genes are part of the mitochondrial ribosome assembly, is ranked sixth by BFCS, but is assigned zero probability by Trigger. This is because, in the Trigger algorithm, the genetic marker L_i exhibiting the strongest linkage with T_i (NAM9 in this case) is preselected and then only the probability of $L_i \rightarrow T_i \rightarrow T_j$ is estimated. With BFCS, on the other hand, we estimated the probability of this structure for all genetic markers.

5. Discussion

We have introduced a novel Bayesian approach for inferring gene regulatory networks that uses the information in the local covariance structure over triplets of variables to make statements about the presence of causal relationships. One key advantage of our method is that we consider all possible causal structures at once, whereas other methods only look at and test for a subset of structures. Because we focus on discovering local causal structures, our method is simple, fast, and inherently parallel, which makes it applicable to very large data sets. Furthermore, the probability estimates produced by BFCS constitute a measure of reliability in the inferred causal relations. We have demonstrated the effectiveness of our approach by comparing it against the Trigger algorithm, a state-of-the-art procedure for inferring causal regulatory relationships. Other methods for inferring gene regulatory networks such as ‘CIT’ (Millstein et al., 2009) or ‘CMST’ (Neto et al., 2013) output p -values instead of probability estimates, which is why they are not directly comparable to BFCS.

In this paper, we have proposed a simple uniform prior on two types of causal graph structures, namely DAGs and DMAGs. However, our method allows for more informative causal priors to be incorporated, taking into consideration properties such as the sparsity of the networks. Moreover, our approach is structure-agnostic, by which we mean we can consider different causal graph struc-

tures incorporating various data-generating assumptions. The tricky part is then to come up with an appropriate prior on the set of causal graph structures from which we assume the data is generated.

Acknowledgments

This research has been partially financed by the Netherlands Organisation for Scientific Research (NWO), under project 617.001.451. We would like to thank the anonymous reviewers for their thoughtful comments and efforts towards improving our manuscript.

References

- J. Barnard, R. McCulloch, and X.-L. Meng. Modeling Covariance Matrices in Terms of Standard Deviations and Correlations, with Application to Shrinkage. *Stat. Sin.*, 10(4):1281–1311, 2000.
- L. S. Chen, F. Emmert-Streib, and J. D. Storey. Harnessing naturally randomized transcription to infer regulatory relationships among genes. *Genome Biology*, 8:R219, Oct. 2007.
- L. S. Chen, D. P. Sangurdekar, and J. D. Storey. trigger, 2017. URL <https://doi.org/doi:10.18129/B9.bioc.trigger>.
- T. Claassen and T. Heskes. A Logical Characterization of Constraint-based Causal Discovery. UAI'11, pages 135–144, Arlington, Virginia, United States, 2011. AUAI Press.
- G. F. Cooper. A Simple Constraint-Based Algorithm for Efficiently Mining Observational Databases for Causal Relationships. *Data Mining and Knowledge Discovery*, 1(2):203–224, June 1997.
- D. Geiger and D. Heckerman. Learning Gaussian Networks. UAI'94, pages 235–243, San Francisco, CA, USA, 1994. Morgan Kaufmann Publishers Inc. ISBN 978-1-55860-332-5.
- S. Mani, P. Spirtes, and G. F. Cooper. A Theoretical Study of Y Structures for Causal Discovery. UAI'06, pages 314–323, Arlington, Virginia, United States, 2006. AUAI Press.
- N. Meinshausen and P. Bühlmann. High-Dimensional Graphs and Variable Selection with the Lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006.
- J. Millstein, B. Zhang, J. Zhu, and E. E. Schadt. Disentangling molecular relationships with a causal inference test. *BMC Genetics*, 10:23, May 2009.
- E. C. Neto, A. T. Broman, M. P. Keller, A. D. Attie, B. Zhang, J. Zhu, and B. S. Yandell. Modeling Causality for Pairs of Phenotypes in System Genetics. *Genetics*, 193(3):1003–1013, Mar. 2013.
- T. Richardson and P. Spirtes. Ancestral Graph Markov Models. *The Annals of Statistics*, 30(4):962–1030, 2002. ISSN 0090-5364.
- K. Sadeghi and S. Lauritzen. Markov properties for mixed graphs. *Bernoulli*, 20(2):676–696, May 2014. ISSN 1350-7265. doi: 10.3150/12-BEJ502.
- M. Studeny. *Probabilistic Conditional Independence Structures*. Springer Science & Business Media, June 2006. ISBN 978-1-84628-083-2. Google-Books-ID: NJ4iwCMoznIC.