

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/176429>

Please be advised that this information was generated on 2019-09-16 and may be subject to change.

Robust Causal Estimation in the Large-Sample Limit without Strict Faithfulness

Ioan Gabriel Bucur

Tom Claassen
Radboud University Nijmegen

Tom Heskes

Abstract

Causal effect estimation from observational data is an important and much studied research topic. The *instrumental variable (IV)* and *local causal discovery (LCD)* patterns are canonical examples of settings where a closed-form expression exists for the causal effect of one variable on another, given the presence of a third variable. Both rely on faithfulness to infer that the latter only influences the target effect via the cause variable. In reality, it is likely that this assumption only holds approximately and that there will be at least some form of weak interaction. This brings about the paradoxical situation that, in the large-sample limit, no predictions are made, as detecting the weak edge invalidates the setting. We introduce an alternative approach by replacing strict faithfulness with a prior that reflects the existence of many ‘weak’ (irrelevant) and ‘strong’ interactions. We obtain a posterior distribution over the target causal effect estimator which shows that, in many cases, we can still make good estimates. We demonstrate the approach in an application on a simple linear-Gaussian setting, using the MultiNest sampling algorithm, and compare it with established techniques to show our method is robust even when strict faithfulness is violated.

1 INTRODUCTION

Establishing the causal effect of one variable on another is a recurring challenge that is shared by most areas of scientific research, ranging from cell biology to economics to psychology and beyond.

Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017, Fort Lauderdale, Florida, USA. JMLR: W&CP volume 54. Copyright 2017 by the author(s).

In almost all cases the principal problem is how to account for the impact of possible confounders on the strength of the observed interaction. When experiments are possible this is easily solved, as performing a randomized trial on the cause variable nullifies all other sources of dependence, and the direct causal effect component can be read off from the resulting correlation in the data. Naturally, in many cases this is not feasible, and we need to fall back on alternative ways to handle unobserved common causes.

One such possibility is when we can establish the presence of a so-called **instrumental variable** (Bowden and Turkington, 1985) for our cause-effect pair in the data set. An instrumental variable (*IV*) is a third variable that is probabilistically dependent on the *cause*, but becomes independent of the *effect* variable after intervention on *cause*. It means all dependence between the *IV* and the *effect* variable is mediated by *cause*, and no other direct or indirect alternative path between *IV* and target *effect* exists.

When the IV setting holds in a linear-Gaussian setting, a valid causal effect estimator takes on a simple form:

$$effect\ size = \frac{Cov[IV, effect]}{Cov[IV, cause]}.$$

However, it is impossible to establish from data whether or not the IV setting applies. Sometimes we may know from background or contextual information that some particular variable is indeed an instrumental variable for the target cause-effect relation, but in general we cannot be sure.

The **local causal discovery algorithm** (Cooper, 1997) resolves this problem by checking if the *IV* and the target *effect* variable become conditionally independent given the *cause*. If true, then, assuming faithfulness, all dependence between *IV* and *effect* is indeed mediated by *cause*, at the expense of a more restricted setting to which the model applies (no confounding between *cause* and *effect*). The causal effect estimator now becomes:

$$effect\ size = \frac{Cov[cause, effect]}{Var[cause]}.$$

The LCD estimator relies on faithfulness, the assumption that any conditional independence among the variables can be read off the corresponding graph via the Markov property. However, it is possible that a direct interaction from IV on $effect$ is exactly compensated by a confounder between $cause$ and $effect$, resulting in an apparent conditional independence, even though the estimator no longer applies. As shown by (Cornia and Mooij, 2014), even for small violations this can lead to worst case arbitrarily large errors in bounds on the resulting causal estimates.

Even worse, in real-world systems the concept of *exact* conditional independences is unlikely to hold: there is bound to be at least *some* residual interaction that will start to show up as we obtain more and more data. That would imply that our methods cannot even be used on very large data sets, as the model setting no longer applies. Paradoxically, “more data hurts” then, which somehow seems very unsatisfactory.

Key Idea

Our idea of solving this situation is to not distinguish between ‘zero’ and ‘nonzero’ interactions, but between ‘irrelevant’ and ‘relevant’ interactions. Essentially, we define a prior on interaction parameters that captures the knowledge that in the real world most interactions between arbitrary variables are likely to be small, whereas interactions with a significant / measurable impact are more or less equally likely to have arbitrary (but reasonable) values. An alternative approach, based on bounds instead of prior probabilities, can be found in Silva and Evans (2016).

When the traditional IV and LCD settings apply, this parameter prior should produce comparable results, except with a peaked distribution over the causal strength estimator instead of a point estimate. However, it should also still give very reasonable results when there is a small residual interaction present, meaning that it still works in the large-sample limit. Of course, in an ‘unlucky’ situation the resulting estimate could still be very wrong, but for an arbitrary instance this should also be extremely unlikely.

In principle, this approach does not only apply to the causal effect estimation setting considered here, but could also be extended to, e.g. causal discovery algorithms that rely on faithfulness, such as PC/FCI (Spirtes et al., 2000). As a result, Occam’s razor now takes the form of a model preference for weak interactions instead of less model parameters. We still get sparse solutions, except now the sparsity is in the number of relevant parameters. It does mean that the problem becomes more complex as we have to compute a posterior distribution. The next section de-

scribes how this approach can be applied in a linear Gaussian model context. After that we show how to handle the resulting Bayesian inference problem, followed by an experimental analysis of our approach in the IV setting.

2 MODEL DESCRIPTION

2.1 Acyclic Directed Mixed Graphs

Our goal is to propose a model that, in the linear Gaussian case, can be used to study causal discovery in the large-sample limit without having to assume strict (or ‘standard’) faithfulness. The basic idea is to pick an ordering of the variables and then allow for a fully connected model, including bi-directed edges between each of the variables to represent confounding. Such a model is clearly overspecified and, even in the limit of an infinite amount of data, will not lead to point estimates for any of the parameters. Nevertheless, we will argue that with appropriate priors that implement some preference for weak interactions over strong ones, it is still possible to infer useful probabilistic statements about causality from purely observational data.

We consider a set of variables from observational data $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$. For now, we will fix the ordering to be such that the first variable is (or can be) a parent of all other variables, the second variable a parent of all variables except the first, and so on. A different order can be implemented by relabeling the variables and, as we will argue later, we can compare or sample over different ordering following similar strategies as in Friedman and Koller (2003), Eaton and Murphy (2007).

For each pair of variables (j, i) , we introduce a latent variable $\omega_{(j,i)}$ to represent an unobserved common cause. Note that we use (i, j) and (j, i) interchangeably in our notation, e.g. $\omega_{(i,j)}$ and $\omega_{(j,i)}$ refer to the same variable. We assume all interactions to be linear with additive noise. The structural equation for variable i then reads

$$X_i = \sum_{j < i} b_{ij} X_j + \sum_{j \neq i} c_{i,(j,i)} \omega_{(j,i)} + \epsilon_i. \quad (1)$$

Here b_{ij} are the structural (path) coefficients corresponding to the direct causal effects ($X_j \rightarrow X_i$) between observed variables, while $c_{i,(j,i)}$ are the structural coefficients corresponding to the effect of the unobserved common causes, which we will refer to as the confounding coefficients. ϵ_i and $\omega_{(j,i)}$ are taken to be independent zero mean Gaussian variables with variance v_i and 1, respectively, the latter w.l.o.g. since any variance different from 1 can be compensated for by scaling the corresponding confounding coefficients.

In matrix notation, we have

$$\mathbf{X} = \mathbf{B}\mathbf{X} + \mathbf{C}\boldsymbol{\omega} + \boldsymbol{\epsilon}, \quad (2)$$

where \mathbf{B} is a $n \times n$ lower triangular matrix and \mathbf{C} is a sparse $n \times \frac{n(n-1)}{2}$ matrix, where only the entries $c_{i,(j,i)}$ and $c_{j,(j,i)}$ are not fixed at zero. The covariance matrix over the error term $\boldsymbol{\epsilon}$ is the matrix $\mathbf{V} = \text{diag}(v_1, v_2, \dots, v_n)$. For given parameters \mathbf{B} and \mathbf{C} , the zero mean normally distributed variables $\boldsymbol{\epsilon}$ and $\boldsymbol{\omega}$ induce a multivariate normal distribution over the observed variables with mean zero and covariance matrix (Bishop, 2006):

$$\boldsymbol{\Sigma} = (\mathbf{I} - \mathbf{B})^{-1}(\mathbf{V} + \mathbf{C}\mathbf{C}^T)(\mathbf{I} - \mathbf{B})^{-T}. \quad (3)$$

Our structural equation model (1) is in fact a canonical DAG representation of an acyclic directed mixed graph (ADMG) over the observed variables (Richardson, 2003). In ADMGs, $b_{ij} = 0$ if there is no directed edge from i to j and $c_{i,(j,i)} = c_{j,(j,i)} = 0$ if there is no bi-directed edge between i and j . The combination of a bi-directed and a directed edge between the same variables is referred to as a bow. As indicated before, for now we will consider fully connected ADMGs with bows between all pairs of variables. The more general ADMG representation used in Bollen (1989), Brito and Pearl (2002) or Van der Zander and Liškiewicz (2016) replaces $\mathbf{C}\mathbf{C}^T$ by a matrix $\boldsymbol{\Psi}$, with $\psi_{ij} = 0$ if there is no bi-directed edge between i and j . Although the general ideas put forward in the rest of this paper would still apply, it is harder to come up with an intuitive prior distribution over such a more general matrix $\boldsymbol{\Psi}$ and obvious choices (such as an inverse Wishart distribution or graphical variants thereof) make the analysis that follows in the rest of this paper considerably more complex.

2.2 Spike-and-slab Prior

Instead of assuming strict faithfulness, we would like to implement the belief that our structural equation model may contain ‘weak’ and ‘strong’ interactions. An obvious choice is the ‘spike-and-slab’ prior introduced by George and McCulloch (1993), consisting of a mixture of two Gaussian distributions: the ‘spike’, with a small variance, and the ‘slab’, with a large variance. The strict (standard) faithfulness assumption would correspond to the special case of a spike with zero variance (see Figure 1). Having a spike with nonzero variance, we hope to be able to cope with near-conditional independencies in the large-sample limit.

We choose to first reparametrize the structural equations (1) and hence (2) by making the coefficients b_{ij} and $c_{i,(j,i)}$ dimensionless. We can achieve this by scaling them using the variance terms, i.e., through the

transformations $\tilde{b}_{ij} = \sqrt{\frac{v_j}{v_i}} b_{ij}$ and $\tilde{c}_{i,(j,i)} = \sqrt{\frac{1}{v_i}} c_{i,(j,i)}$, or, in matrix form, $\tilde{\mathbf{B}} = \mathbf{V}^{-\frac{1}{2}} \mathbf{B} \mathbf{V}^{\frac{1}{2}}$ and $\tilde{\mathbf{C}} = \mathbf{V}^{-\frac{1}{2}} \mathbf{C}$. Equation (3) then boils down to:

$$\boldsymbol{\Sigma} = \mathbf{V}^{\frac{1}{2}} (\mathbf{I} - \tilde{\mathbf{B}})^{-1} (\mathbf{I} + \tilde{\mathbf{C}} \tilde{\mathbf{C}}^T) (\mathbf{I} - \tilde{\mathbf{B}})^{-T} \mathbf{V}^{\frac{1}{2}}. \quad (4)$$

We now propose to take a spike-and-slab prior over the scaled structural coefficients \tilde{b}_{ij} :

$$p(\tilde{b}_{ij}) = w_{\text{spike}} \mathcal{N}(\tilde{b}_{ij}; 0, v_{\text{spike}}) + w_{\text{slab}} \mathcal{N}(\tilde{b}_{ij}; 0, v_{\text{slab}}). \quad (5)$$




Figure 1: ‘Traditional’ prior (left) and ‘spike-and-slab’ prior (right)

In the rest of this paper, we will be only working with the scaled confounding coefficients, so to simplify notation, we will omit the tilde and use $c_{i,(j,i)}$ to mean $\tilde{c}_{i,(j,i)}$ and so on. Furthermore, we fix the mixture weight values at $w_{\text{spike}} = w_{\text{slab}} = 0.5$, corresponding to an indifference or uniform prior (Ishwaran and Rao, 2005), we set $v_{\text{slab}} = 1$ and vary v_{spike} according to how this parameter affects the ability to handle near-conditional independencies. For the variances v_i , we take a scale-invariant log-uniform prior and, for the confounding coefficients, a zero mean Gaussian prior with unit variance (i.e. (5) with just the slab).

Of course, given prior knowledge, one can make other choices and even consider hierarchical models with prior distributions on, for example, the parameters specifying the spike-and-slab distribution, but this is beyond the scope of the current paper.

3 BAYESIAN INFERENCE

3.1 Likelihood

Our model parameters are the structural coefficients \mathbf{B} , the scaled confounding coefficients \mathbf{C} , and the variances collected in \mathbf{V} . For reasons to become clear soon, we will group the structural coefficients \mathbf{B} and the variances \mathbf{V} into the set of parameters $\boldsymbol{\Theta} = (\mathbf{B}, \mathbf{V})$.

Since the implied distribution over the observed variables is a zero-mean Gaussian with covariance matrix $\boldsymbol{\Sigma}$, the sample covariance matrix $\hat{\boldsymbol{\Sigma}}$ is a sufficient

statistic and the log-likelihood reads, up to irrelevant additive constants,

$$\begin{aligned} \log \mathcal{L}(\Theta, \mathbf{C} | \hat{\Sigma}, N) &= \\ &= -\frac{N}{2} \left\{ \text{tr} \left[\Sigma(\Theta, \mathbf{C})^{-1} \hat{\Sigma} \right] + \log \det \Sigma(\Theta, \mathbf{C}) \right\}, \end{aligned}$$

with N the number of data points and where we made explicit the dependencies of the implied covariance on the parameters.

The log-likelihood has a maximum when the implied and sample covariance matrix are identical, i.e., when $\Sigma(\Theta, \mathbf{C}) = \hat{\Sigma}$. The sample covariance matrix contains $n(n+1)/2$ independent parameters, which is exactly the number of parameters for Θ in a fully connected ADMG: $n(n-1)/2$ parameters in the lower triangular part of the matrix \mathbf{B} and n variances in \mathbf{V} . Given any set of confounding coefficients \mathbf{C} and any sample covariance matrix $\hat{\Sigma}$, we can always find a unique set of parameters $\Theta^*(\mathbf{C}, \hat{\Sigma})$ that satisfies $\Sigma(\Theta^*(\mathbf{C}, \hat{\Sigma}), \mathbf{C}) = \hat{\Sigma}$. In the Appendix we give an efficient procedure for computing $\Theta^*(\mathbf{C}, \hat{\Sigma})$ using a Cholesky decomposition.

That there is a compatible solution Θ^* for any choice of the confounding coefficients \mathbf{C} makes the problem of finding the maximum likelihood solution nonidentifiable (Brito and Pearl, 2002): there is a whole continuum of solutions, each of them leading to a different estimate of, for example, a particular causal effect size. The situation then may seem hopeless, in line with the negative result in Cornia and Mooij (2014), but it is here that our spike-and-slab prior offers a way out by preferring sensible solutions over odd ones.

3.2 Posterior Distribution

Given a particular sample covariance matrix $\hat{\Sigma}$, we are mainly interested in the posterior distribution over the structural coefficients \mathbf{B} and the variances \mathbf{V} . In the limit of an infinite amount of data, the likelihood term gets closer and closer to a delta peak around the maximum likelihood solution, eventually enforcing the constraint $\Theta = (\mathbf{B}, \mathbf{V}) = \Theta^*(\mathbf{C}, \hat{\Sigma})$. So, if we can get our hands on $p(\mathbf{C} | \hat{\Sigma})$, $p(\Theta | \hat{\Sigma})$ follows from:

$$p(\Theta | \hat{\Sigma}) = \int d\mathbf{C} \delta(\Theta - \Theta^*(\mathbf{C}, \hat{\Sigma})) p(\mathbf{C} | \hat{\Sigma}). \quad (6)$$

The posterior distribution $p(\mathbf{C} | \hat{\Sigma})$ can be computed using Laplace's method:

$$\begin{aligned} p(\mathbf{C} | \hat{\Sigma}) &= \int d\Theta p(\Theta, \mathbf{C} | \hat{\Sigma}) \\ &\propto \lim_{N \rightarrow \infty} \int d\Theta \exp[\log \mathcal{L}(\Theta, \mathbf{C} | \hat{\Sigma}, N)] p(\Theta) p(\mathbf{C}) \\ &= \det[-\mathbf{H}_{\mathbf{C}}(\Theta^*(\mathbf{C}, \hat{\Sigma}))]^{-\frac{1}{2}} p(\Theta^*(\mathbf{C}, \hat{\Sigma})) p(\mathbf{C}), \quad (7) \end{aligned}$$

where $\mathbf{H}_{\mathbf{C}}(\Theta^*(\mathbf{C}, \hat{\Sigma}))$ is the Hessian matrix of the log-likelihood per data point w.r.t. Θ evaluated at the maximum likelihood solution:

$$\mathbf{H}_{\mathbf{C}}(\Theta^*(\mathbf{C}, \hat{\Sigma})) \equiv \frac{1}{N} \left. \frac{\partial^2 \log \mathcal{L}(\Theta, \mathbf{C} | \hat{\Sigma})}{\partial \Theta \partial \Theta^T} \right|_{\Theta = \Theta^*(\mathbf{C}, \hat{\Sigma})}.$$

So, in the limit of an infinite amount of data, the posterior distribution for the confounding coefficients is proportional to its prior distribution, the priors for the matching structural coefficients and variances, and a determinant term that relates to the Jacobian for the transformation Θ^* from \mathbf{C} to the corresponding \mathbf{B} and \mathbf{V} . With relatively uninformative priors for the confounding coefficients and the variance, the prior over the structural coefficients will have the largest impact. If this prior favors 'weak' structural coefficients, the posterior will prefer confounding coefficients that map to solutions with weak structural coefficients over those that do not. Through this mechanism, the data does have an impact on the posterior distribution over the confounding and hence structural coefficients, leading to estimates of causal strength sizes that will not converge to absolute certainty in the limit of an infinite amount of data, but may still provide relevant information. We will illustrate these aspects in detail in the instrumental variable setting in the next section.

Because of the mapping that implements the constraint and the corresponding Jacobian, we do not have an analytical expression for the posterior, but can numerically compute it for any value of \mathbf{C} up to a normalization constant. This then suggests a straightforward Markov Chain Monte Carlo sampling procedure, where one can take one's favorite MCMC method to draw samples for the confounding coefficients \mathbf{C} from the posterior (7) and, following (6), apply the transformation Θ^* to obtain corresponding samples for the structural coefficients \mathbf{B} and the variances \mathbf{V} .

Until now, we have only considered a fixed ordering of the variables. Using MCMC samplers that can also estimate the marginal likelihood, we can compare (as we will do in the experimental section) and eventually sample over different orderings in order to perform model selection. This brings additional challenges, in particular w.r.t. computational efficiency, which we leave for future work.

4 EMPIRICAL RESULTS

In this section, we present some empirical results that highlight the benefits of our approach. Given only the (large-sample limit) observed covariance matrix, we generate posterior samples of the assumed structural coefficients using the MultiNest (Feroz et al., 2009)

technique. The end result consists in a posterior density estimate of the structural parameters in which we are interested.

4.1 Illustrative Example

We consider the model in Figure 2, which has been studied previously in Richardson et al. (2011) and Cornia and Mooij (2014). Here X_1 is a *treatment*, presumed to be randomized, X_2 is an *exposure* subsequent to treatment assignment and X_3 is the *response*. We would like to estimate the direct causal effect of the *exposure* on the *response* in the presence of the hidden confounding variable $\omega_{(2,3)}$. For simplicity, we only consider a bow over one pair of variables, X_2 and X_3 .

The corresponding linear structural equation model is:

$$\begin{aligned} X_1 &= \epsilon_1 \\ X_2 &= b_{21}X_1 + c_{2,(2,3)}\omega_{(2,3)} + \epsilon_2 \\ X_3 &= b_{31}X_1 + b_{32}X_2 + c_{3,(2,3)}\omega_{(2,3)} + \epsilon_3 \end{aligned} \quad (8)$$




Figure 2: Illustrative model with three observed variables and one hidden confounder.

4.2 Instrumental Variable Setting

We first examine a setting where one can successfully employ the instrumental variable technique. If we assume that there is no confounding ($c_{2,(2,3)} = c_{3,(2,3)} = 0$), we can use the LCD algorithm proposed by Cooper (1997) to directly test that X_1 is an instrumental variable from observational data.

In Figure 3(a) we show the posterior density estimate of b_{32} produced by our method under these conditions in the large-sample limit. We observe a high peak around the true value $b_{32} = 1$. Moreover, roughly 85% of the posterior probability mass is concentrated in the interval $[0.9, 1.1]$. We also notice a small mode around zero, which corresponds to a small value for b_{32} , yet a large one for b_{31} . Roughly 1.86% of the probability mass is concentrated in the interval $[-0.1, 0.1]$.

The illustrated result exhibits the drawback of having introduced some uncertainty into our estimate by forgoing the strict faithfulness assumption. In this particular situation, the instrumental variable technique can provide an exact point estimate of the causal effect




Figure 3: Posterior distribution of b_{32} in the following scenarios:

- (a) instrumental variable setting without confounding ($b_{31} = 0, c_{2,(2,3)} = 0, c_{3,(2,3)} = 0$)
- (b) instrumental variable setting with confounding ($b_{31} = 0, c_{2,(2,3)} = 1, c_{3,(2,3)} = 1$)
- (c) weak departure from IV setting without confounding ($b_{31} = 0.05, c_{2,(2,3)} = 0, c_{3,(2,3)} = 0$)
- (d) weak departure from IV setting with confounding ($b_{31} = 0.05, c_{2,(2,3)} = 1, c_{3,(2,3)} = 1$)
- (e) strong departure from IV setting with confounding ($b_{31} = 1, c_{2,(2,3)} = 1, c_{3,(2,3)} = 1$)
- (f) adversarial scenario with singular coefficients ($b_{31} = 1, c_{2,(2,3)} = 1, c_{3,(2,3)} = 2$)

The other simulation parameters are the same for all scenarios: $b_{21} = 1, b_{32} = 1, v_1 = 1, v_2 = 1, v_3 = 1$.

in which we are interested, but such an ideal situation will rarely occur in practice.

If we introduce a confounding effect via $\omega_{(2,3)}$, we are still in the IV setting, even though we cannot directly attest this from observational data (the conditional in-

dependence $X_1 \perp\!\!\!\perp X_3|X_2$ no longer holds). The IV technique will then still provide a perfect estimate of b_{32} given the infinite sample covariance matrix.

The result of our approach in the presence of confounding is shown in Figure 3(b). We again recognize, like in Figure 3(a), a high peak around the true value $b_{32} = 1$. There is more uncertainty, however, in the estimate (only 71.1% of the probability mass falls in the interval $[0.9, 1.1]$), which is caused by the additional correlation via the hidden confounder $\omega_{(2,3)}$. A new mode around the value 1.5 starts to show, corresponding to the zero confounding solution induced by the Gaussian prior on $c_{2,(2,3)}$ and $c_{3,(2,3)}$, and by the Hessian term in the expression of the posterior distribution (see Figure 5).

4.3 ‘Weak’ Departure from the IV Setting

We now introduce a weak causal effect between the *treatment* (the IV) and the *response*: $b_{31} = 0.05$ (see Figures 3(c) and 3(d)). Given finite data, we might still conclude that we are in the IV setting, for example if we detect the conditional independence $X_1 \perp\!\!\!\perp X_3|X_2$ when applying the LCD algorithm. In that case, assuming we have no confounding, the IV estimate could still be considered reasonable provided that $b_{31} \ll b_{21}$. In the large-sample limit, however, this ‘weak’ conditional dependence will be detected, which means the IV technique can no longer be applied.

For the structural parameter values used in our simulation experiment, the partial correlation between X_1 and X_3 given X_2 is $\rho_{13.2} \approx 0.035$. We can test the null hypothesis that $\rho_{13.2} = 0$ using Fisher’s z-transformation applied to the sample correlation coefficient (Kalisch and Bühlmann, 2007). Given the partial correlation’s true value, Fisher’s test would reject the null hypothesis when the number of samples is greater than $n = 3708$. For a moderate number of samples, the LCD algorithm can be employed to estimate the causal strength from X_2 to X_3 . As we go into the large-sample limit, however, the estimated partial correlation converges to the true (nonzero) value, in which case we can no longer use the LCD algorithm for causal estimation.

The contour plot of the posterior distribution $p(\mathbf{C}|\hat{\Sigma})$ reveals the presence of a few high-density regions. In the center, we have a mode that results from combining the Hessian term $\det[-\mathbf{H}_{\mathbf{C}}(\Theta^*(\mathbf{C}, \hat{\Sigma}))]^{-\frac{1}{2}}$ with the zero-mean Gaussian prior on $c_{2,(2,3)}$ and $c_{3,(2,3)}$. The high density regions surrounding the black lines in Figure 4 are the result of putting a spike-and-slab prior on the structural coefficient b_{31} . The black contour lines are obtained by solving the linear system given by Equation (3) for $b_{31} = 0$. Similarly, the high




Figure 4: Contour plot of $\log p(\mathbf{C}|\hat{\Sigma})$ when we have a weak departure from the IV setting.

density regions surrounding the dark green lines are due to the sparsifying properties of the spike-and-slab prior on b_{32} . The dark green contour lines are obtained by solving Equation (3) for $b_{32} = 0$.




Figure 5: Contour plot of the log Hessian term $-\frac{1}{2} \log \det[-\mathbf{H}_{\mathbf{C}}(\Theta^*(\mathbf{C}, \hat{\Sigma}))]$ when we have a weak departure from the IV setting.

In Figure 3(c), we can see the posterior distribution of b_{32} when there is no confounding. We observe a

prominent mode close to the true value $b_{32} = 1$, which constitutes a reasonable estimate of the structural coefficient, even in the presence of the ‘weak’ causal effect $b_{31} = 0.05$. Roughly 80% of the probability mass is concentrated in the interval $[0.9, 1.1]$, compared to 85% in the absence of this ‘weak’ causal effect. At the same time, we rediscover the small peak around the value $b_{32} = 0$. This peak corresponds to a less probable alternative solution, induced by the spike-and-slab prior on b_{32} .

The estimate produced by our method remains robust to the presence of confounding, even with the ‘weak’ causal effect b_{31} . The probability mass in the interval $[0.9, 1.1]$ is 68.8%, compared to 71.1% in the previous subsection, as can be seen in Figure 3(d). Again, we have a mode around $b_{32} = 1.5$, corresponding to the solution when the confounding coefficients are zero.

4.4 ‘Strong’ Departure from the IV Setting

We now consider a scenario where the causal effect from X_1 to X_3 is ‘strong’ and where we also have confounding. These conditions lead to more uncertainty in the estimate of b_{32} , which is reflected by the large spread of the posterior distribution (see Figure 3(e)). In this scenario, the IV and LCD estimators will both give the wrong result: $\hat{b}_{32}^{IV} = \hat{b}_{32}^{LCD} = 2$. As far as our method is concerned, the output includes a significant proportion of probability mass around the correct solution, $b_{32} = 1$ (roughly 7.2% in the interval $[0.9, 1.1]$), even though sparser solutions are still preferred, as indicated by the modes at $b_{32} = 1.5$ and $b_{32} = 2$. We highlight this as an advantage of our approach, which provides richer information about the data generating process than a simple point estimate.

4.5 Adversarial Scenario




Figure 6: Ground truth model (left; $v_1 = 1, v_2 = 1, v_3 = 1$) and equivalent preferred sparser model (right; $v_1 = 1, v_2 = 2, v_3 = 3$)

Finally, we consider a degenerate case where both traditional methods and our own run into trouble (see Figure 6). Even though the model on the left side of the figure is the ground truth, methods promoting model sparsity will go for the equivalent model on the right side. The singular choice of structural coefficients in the ground truth model produces an apparent ‘zero’

partial correlation, $\rho_{13.2} = 0$, which implies the conditional independence $X_1 \perp\!\!\!\perp X_3 | X_2$. If one were to use the IV or LCD estimator, one would obtain the same incorrect result: $\hat{b}_{32}^{IV} = \hat{b}_{32}^{LCD} = 2$. In Figure 3(f), the highest mode is also at $b_{32} = 2$. Despite this clear preference, there is still some mass around the correct solution (roughly 0.91% in the interval $[0.9, 1.1]$).

4.6 Model Selection

So far, we have assumed a fixed temporal ordering of our variables. Under the assumptions of our simulation model ($\text{An}(X_1) \cap \{X_2, X_3\} = \emptyset$), two temporal orderings are possible: $1 \rightarrow 2 \rightarrow 3$ and $1 \rightarrow 3 \rightarrow 2$. We can use MultiNest to compute the log-evidence for each of the two possible models. Assuming that $1 \rightarrow 2 \rightarrow 3$ is the correct temporal ordering of our data generating process, we obtain the evidence ratio $\frac{p(1 \rightarrow 2 \rightarrow 3)}{p(1 \rightarrow 3 \rightarrow 2)} = 3.358$ for the simulation scenario described in Subsection 4.3. In Figure 7 we see how this ratio varies with the width of the spike.




Figure 7: Evidence ratio in favor of $1 \rightarrow 2 \rightarrow 3$ versus the alternative $1 \rightarrow 3 \rightarrow 2$.

For a low precision (high variance) spike, the preference for sparse models disappears and the ordering does not matter. When the true interaction can be explained by a spike which is sufficiently different from a slab, the ordering leading to the sparser model is preferred to the ordering leading to the denser model. For an extremely high precision (low variance, tending towards the faithfulness assumption), the obtained results suggest that both models again become equally likely, but this may well be due to numerical issues: MultiNest appears to have increasing difficulty to properly sample from the areas corresponding to spikes that become more and more narrow (see the contours in Figure 4).

4.7 Robustness

We can see the effect of changing the variance of the spike in Figure 8. For a wide interval of values, from spike variance 10^{-3} to 10^{-7} , the posterior output qual-




Figure 8: Posterior distribution of b_{32} when we have a weak departure from the IV setting for various spike variances (the slab variance is fixed at 1)

itatively expresses the same result: it strongly implies that the causal effect of interest is significant, as indicated by the high peak close to one, but at the same time it does not exclude the unlikely possibility that this effect is irrelevant, as indicated by the small peak around zero. This shows that the method we propose is robust with respect to the width of the spike. When the spike variance is close to that of the slab, the peak at zero vanishes and the distribution becomes more spread, reflecting additional uncertainty in the estimation of the causal effect.

We also compute the bounds derived from the Witness Protection Program (WPP) algorithm for the linear case (see Section 6.1 in Silva and Evans (2016)) by taking $\epsilon_c = b_{32} \sqrt{\frac{\Sigma_{22}}{\Sigma_{33}}}$ as the constraint on the ‘weak’ interaction. The bounds are plotted in Figure 8. When the spike variance takes a reasonable value (in the interval $[10^{-7}, 10^{-3}]$), the posterior distribution produced by our method has most of its probability mass within the WPP bounds. Additionally, the posterior reveals interesting details, such as unlikely alternative solutions, which are not captured by the bounds.

5 DISCUSSION

We have shown a new approach to causal inference based on a spike-and-slab prior that promotes small parameter interactions. The prior derives from the notion that, in the real-world, many interactions are likely to be very small, but not exactly zero. In that setting, standard notions that promote model sparsity based on the *number* of (nonzero) parameters no longer suffice, since all models have an equal number of parameters. What differs in our approach is that model sparsity now becomes a minimization of strong interaction parameters relative to weak parameters. In our experiments, the resulting method shows meaningful and consistent results: it provides very informative posterior distributions on target causal effect estimations. It shows when we can be confident about a certain nonzero effect size, or when we do not have enough information to decide on a specific value.

Above all, the method is no longer susceptible to the standard faithfulness violations that plague many existing causal inference algorithms. This should make the method much more robust in many practical applications. In this paper we have only shown how to implement the method for a simple linear Gaussian model, but the principle should apply equally to larger graphs and more complex model forms. As such, we consider it a very promising start and we are now looking to incorporate this idea into causal discovery algorithms like PC/FCI in order to make them less susceptible to errors from violations of faithfulness.

Acknowledgements

We would like to thank Dr. Ricardo Silva for providing the code we used to compute the WPP bounds. This research has been partially financed by the Netherlands Organisation for Scientific Research (NWO) under project 617.001.451. TC was supported by NWO grant 612.001.202 (MoCoCaDi), and EU-FP7 grant n.603016 (MATRICS).

References

- Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006. ISBN 0387310738.
- K. A. Bollen. *Structural Equations with Latent Variables*. John Wiley & Sons, New York, 1989.
- Roger J. Bowden and Darrell A. Turkington. *Instrumental Variables*. Cambridge University Press, Cambridge, 001 1985. ISBN 9781139052085.
- Carlos Brito and Judea Pearl. A new identification condition for recursive models with correlated errors. *Structural Equation Modeling*, 9(4):459–474, 2002.
- Gregory F. Cooper. A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Mining and Knowledge Discovery*, 1(2):203–224, 1997.
- Nicholas Cornia and Joris M. Mooij. Type-II errors of independence tests can lead to arbitrarily large errors in estimated causal effects: An illustrative example. In Joris M. Mooij, Dominik Janzing, Jonas Peters, Tom Claassen, and Antti Hyttinen, editors, *UAI 2014 Workshop Causal Inference: Learning and Prediction*, number 1274 in CEUR Workshop Proceedings, pages 35–42, Aachen, 2014.
- Daniel Eaton and Kevin Murphy. Bayesian structure learning using dynamic programming and MCMC. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, UAI’07, pages 101–108, Arlington, Virginia, United States, 2007. AUAI Press. ISBN 0-9749039-3-0. URL <http://dl.acm.org/citation.cfm?id=3020488.3020501>.
- F. Feroz, M. P. Hobson, and M. Bridges. MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Monthly Notices of the Royal Astronomical Society*, 398(4):1601–1614, 2009.
- N. Friedman and D. Koller. Being Bayesian about network structure: A Bayesian approach to structure discovery in Bayesian networks. *Machine Learning*, 50:95–126, 2003.
- Edward I. George and Robert E. McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993.
- Hemant Ishwaran and J. Sunil Rao. Spike and slab variable selection: frequentist and Bayesian strategies. *Annals of Statistics*, pages 730–773, 2005.
- Markus Kalisch and Peter Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8(Mar):613–636, 2007.
- Thomas Richardson. Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1):145–157, 2003.
- Thomas S. Richardson, Robin J. Evans, and James M. Robins. Transparent parameterizations of models for potential outcomes. *Bayesian Statistics*, 9:569–610, 2011.
- Ricardo Silva and Robin Evans. Causal inference through a witness protection program. *Journal of Machine Learning Research*, 17(56):1–53, 2016. URL <http://jmlr.org/papers/v17/15-130.html>.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. The MIT Press, Cambridge, Massachusetts, 2nd edition, 2000.
- Benito Van der Zander and Maciej Liśkiewicz. On searching for generalized instrumental variables. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 1214–1222, 2016.