

# Distinguishing Cause from Effect on Categorical Data: The Uniform Channel Model

Mário A. T. Figueiredo,

MARIO.FIGUEIREDO@TECNICO.ULISBOA.PT

Catarina Oliveira

CATARINA.A.OLIVEIRA@TECNICO.ULISBOA.PT

*Instituto de Telecomunicações and LUMELIS ELLIS Unit,  
Instituto Superior Técnico, Universidade de Lisboa, Portugal*

**Editors:** Mihaela van der Schaar, Dominik Janzing and Cheng Zhang

## Abstract

Distinguishing cause from effect using observations of a pair of random variables is a core problem in causal discovery. Most approaches proposed for this task, namely *additive noise models* (ANM), are only adequate for quantitative data. We propose a criterion to address the cause-effect problem with categorical variables (living in sets with no meaningful order), inspired by seeing a conditional *probability mass function* (pmf) as a discrete memoryless channel. We select as the most likely causal direction the one which the conditional pmf is closer to a *uniform channel* (UC). The rationale is that, in a UC, as in an ANM, the conditional entropy (of the effect given the cause) is independent of the cause distribution, in agreement with the principle of *independence of cause and mechanism*. Our approach, which we call the *uniform channel model* (UCM), thus extends the ANM rationale to categorical variables. To assess how *close* a conditional pmf (estimated from data) is to a UC, we use statistical testing, supported by a closed-form estimate of a UC channel. On the theoretical front, we prove identifiability of the UCM and show its equivalence with a structural causal model with a low-cardinality exogenous variable. Finally, the proposed method compares favorably with recent state-of-the-art alternatives in experiments on synthetic, benchmark, and real data.

## 1. Introduction

Causal inference is a key problem in many areas of science and data analysis (Pearl, 2009). In principle, distinguishing statistical dependencies from causal relationships requires interventions (Pearl, 2009; Peters et al., 2017). However, intervening is often impossible (*e.g.*, analyzing past data), impractical, or unethical (*e.g.*, forcing people to smoke), which has stimulated much research aimed at inferring causal relationships (*causal discovery*) from purely observational data (Janzing, 2019; Mooij et al., 2016; Peters et al., 2017), or mixed observational-interventional data (Faria et al., 2022)

There is a vast literature on methods to learn *directed acyclic graphs* (DAGs) from data, usually by inferring *conditional independence* (CI) properties among variables (Chickering, 2002; Heckerman and Geiger, 1995; Koller and Friedman, 2010). However, without additional assumptions or criteria, those methods cannot distinguish different DAGs entailing the same CIs (in the same *Markov equivalence class* – MEC). The simplest instance of this problem involves a pair of variables ( $X, Y$ ): purely statistical methods (*e.g.*, maximum likelihood estimation) cannot recover the causal graph because  $X \rightarrow Y$  and  $Y \rightarrow X$  constitute a MEC, corresponding to the two possible factorizations of the joint distribution:  $p_{X,Y}(x, y) = p_{Y|X}(y|x) p_X(x) = p_{X|Y}(x|y) p_Y(y)$ . Although there are several methods that select particular elements of a MEC, they all include additional assumptions beyond *faithfulness*<sup>1</sup> and, for one reason or another they cannot be used to distinguish between

1. Faithfulness holds if every conditional independence in the joint probability distribution corresponds to a separation property in the graph (Koller and Friedman, 2010; Sadeghi, 2017).

$X \rightarrow Y$  and  $Y \rightarrow X$ , where  $X$  and  $Y$  are a pair of categorical variables; *e.g.*, they are only applicable to quantitative data (Park and Raskutti, 2018) or only make sense for more than two variables (Gao and Aragam, 2021).

Without interventions, choosing an element of a MEC requires additional assumptions about the underlying data-generating mechanism. For instance, *additive noise models* (ANM) (Shimizu et al., 2006; Hoyer et al., 2009; Peters et al., 2011, 2014) assume the effect is a function of the cause plus a noise term independent of the cause ( $Y = f_Y(X) + N_Y, X \perp\!\!\!\perp N_Y$ ); if the same doesn't hold in the reverse direction, the model is said to be *identifiable*. ANMs are generically identifiable (Peters et al., 2011, 2014) in the following sense: if the joint density  $p_{X,Y}$  corresponds to an ANM, the conditional density  $p_{Y|X}(\cdot|x)$  has identical *shape* for any  $x$ , simply being *shifted* by  $f_Y(x)$ , but  $p_{X|Y}(\cdot|y)$  typically depends on  $y$  in a more complicated way. Under the ANM criterion, if such a model exists in one direction but not the other, the former is selected as the causal direction.

The ANM can be seen as an instance of the principle of *independence of cause and mechanism* (ICM) (Janzing et al., 2010, 2012), according to which the cause-effect mechanism (a deterministic function followed by addition of noise, in an ANM) is *independent*<sup>2</sup> of the cause, thus of its distribution. The ICM principle has been exploited using different tools to define and assess the notion of *independence*: information geometry (Janzing et al., 2012); Kolmogorov complexity (not computable, but approximable (Li and Vitányi, 2009)) by Janzing et al. (2012) and Mian et al. (2021); stochastic complexity, via the *minimum description length* (MDL (Rissanen, 1998)) principle (Budhathoki and Vreeken, 2017a; Marx and Vreeken, 2019; Tagasovska et al., 2020) or the *minimum message length* (Wallace and Boulton, 1968) criterion, by Stegle et al. (2010).

Relatively few methods have been proposed to address the cause-effect problem with categorical variables. Peters et al. (2011) extended ANMs to the discrete case, but only for variables taking values in a set equipped with a meaningful order, in which an operation similar to addition (a shift) is defined. They consider  $\mathbb{Z}$  (or subsets thereof) for variables without cyclic structure, and  $\mathbb{Z}/n\mathbb{Z}$  with modulo- $n$  addition, for cyclic variables (*e.g.*, seasons or months of the year). However, purely categorical variables live in sets with no order, thus no meaningful notion of addition or shift, precluding the direct use of ANMs. Peters et al. (2011) consider what they call “structureless” sets, but only with a particular form of the conditional pmf, not generally applicable. Some of the MDL-based methods mentioned above can be used with categorical variables: Budhathoki and Vreeken (2017a) proposed CISC (*causal inference by stochastic complexity*); Cai et al. (2018) proposed HCR (*hidden compact representation*), based on BIC (*Bayesian information criterion*). Liu and Chan (2016) assess mechanism independence via a *distance correlation* (DC) between the cause pmf and the conditional pmf of the effect. Kocaoglu et al. (2017) select the causal direction in which the sum of the marginal entropy of the cause with that of the exogenous variable in the corresponding *structural causal model* (SCM) is minimal. Recently, Ni (2022) addressed the cause-effect problem for categorical variables by formulating the conditional distribution of the effect given the cause as ordinal regression, with optimal label permutation, and choosing the direction in which this model has the highest likelihood.

As is standard when focusing on the cause-effect problem, we assume *causal sufficiency* (*i.e.*, absence of unobserved *confounders*), no selection bias, and no feedback (Mooij et al., 2016). We propose a new approach to the cause-effect problem for categorical variables, inspired by viewing the causal mechanism as a communication channel. This view allows extending to the categorical case a key feature of ANMs: the conditional (differential) entropy of the effect given the cause

---

2. The term “independent” here does not have a probabilistic sense, but a functional sense: changing the distribution of the cause does not affect the causal mechanism, and vice-versa.

is independent<sup>1</sup> of the distribution of the cause (Janzing et al., 2012). For categorical variables (i.e., symbols, in channel terminology), a memoryless channel corresponds to the conditional *probability mass function* (pmf) of the output given the input, the *channel matrix*  $\theta^{X \rightarrow Y}$ , where  $\theta_{x,y}^{X \rightarrow Y} = p_{Y|X}(y|x) = \mathbb{P}[Y = y|X = x]$ . In a so-called *uniform channel* (UC (Hamming, 1986)), the rows of this matrix are permutations of each other<sup>3</sup>, implying (as shown below) that the conditional entropy  $H(Y|X)$  is independent of the distribution of  $X$ . Paralleling the ANM rationale, given a pair of categorical variables  $(X, Y)$ , if the conditional pmf in one direction, say of  $Y$  given  $X$ , corresponds to a UC and the same is not true in the other direction, then the causal structure is declared to be  $X \rightarrow Y$ . This criterion, which we refer to as the UCM (*uniform channel model*) is supported by an *identifiability* result proved in this paper: if a joint distribution corresponds to a UCM in one direction, in general (i.e., with probability 1 under any continuous density on the model parameters), it does not correspond to a UCM in the reverse direction.

The proposed UCM approach is further supported by the fact (proved below) that if, and only if,  $\theta^{X \rightarrow Y}$  corresponds to a UCM, is it possible to write a *structural causal model* (SCM) (Pearl, 2009) of the form  $Y = f_Y(X, U_Y)$ , where  $f_Y$  is a deterministic function and  $U_Y$  is an exogenous random variable, taking values in the same set as  $Y$  and independent of  $X$ . The importance of this independence was recently highlighted by Papineau (2022): “the probabilistic independence of exogenous terms in (...) structural equations holds the key to causal direction”.

A final question is how to instantiate the UCM principle with a finite amount of data. This question parallels that of how to estimate the underlying function and noise distribution in an ANM. Naturally, with a finite dataset, we only have an estimate of the underlying distribution and the probability that this estimate corresponds exactly to a UCM in one of the two directions is vanishingly small. Although other ways to address this issue are conceivable, we resort to statistical hypothesis testing to decide in which direction, if any, the conditional pmf can be considered a UCM. A key building block of this approach is estimating a channel under the constraint that it is uniform; this is a problem that, to the best of our knowledge, had not been studied before and for which we derive a closed-form solution. We also extend the approach to the case where the rows of the channel matrix are cyclic permutations of each other (a *cyclic UCM* – CUCM), applicable when the effect variable has cyclic nature, but in this case the channel estimate has to be obtained iteratively.

The main contributions of this paper are the following:

- A new instantiation, for categorical variables, of the principle of *independence of cause and mechanism*: the *uniform channel model* (UCM) principle.
- A proof of identifiability of the UCM.
- A proof that the joint distribution of a pair of categorical random variables is entailed by an SCM in which the exogenous noise has the same cardinality as the effect variable if and only if it corresponds to a UCM.
- An instantiation of the UCM principle using statistical hypothesis testing, supported on a closed-form estimate of a UCM (which, to the best of our knowledge, is a new result, possibly of independent interest).

---

3. A uniform channel is not necessarily a *symmetric channel*, which requires additionally that all the columns are also permutations of each other (Cover and Thomas, 2006).

The paper is organized as follows. Section 2 describes the UCM and presents the corresponding identifiability theorem and equivalence to an SCM. Section 3 addresses the problem of estimating uniform and cyclic uniform channels from data. Section 4 describes how the criterion is applied to observed data. Experimental results are reported in Section 5, and Section 6 concludes the paper.

## 2. Uniform Channel Models – UCM – for Categorical Variables

This section describes the proposed causal inference principle for categorical variables, after introducing notation and reviewing the notion of uniform channel. Finally, we prove an identifiability theorem for the proposed model and show its equivalence with an SCM.

### 2.1. Categorical Variables and Uniform Channels

Let  $X \in \mathcal{X} = \{1, \dots, |\mathcal{X}|\}$  and  $Y \in \mathcal{Y} = \{1, \dots, |\mathcal{Y}|\}$  be two *categorical random variables* (although their outcomes are shown as integers, no role is played by their order). The joint pmf  $p_{X,Y}$  can be factored in two different ways,  $p_{X,Y}(x, y) = p_{Y|X}(y|x) p_X(x) = p_{X|Y}(x|y) p_Y(y)$ , corresponding to a *Markov equivalence class*. If  $X \perp\!\!\!\perp Y$ , the joint pmf factors trivially  $p_{X,Y}(x, y) = p_X(x) p_Y(y)$ . Let the vector of parameters of the first factorization be denoted as  $\theta = (\theta^X, \theta^{X \rightarrow Y})$ , *i.e.*,

$$\theta_x^X = p_X(x) = \mathbb{P}[X = x] \quad \text{and} \quad \theta_{x,y}^{X \rightarrow Y} = p_{Y|X}(y|x) = \mathbb{P}[Y = y|X = x],$$

where  $\theta^X \in \Delta_{|\mathcal{X}|-1}$ , with  $\Delta_{m-1}$  being the probability simplex in  $\mathbb{R}^m$ . The conditional probabilities are arranged in a  $|\mathcal{X}| \times |\mathcal{Y}|$  row-stochastic matrix  $\theta^{X \rightarrow Y}$ , with the  $x$ -th row denoted as  $\theta_x^{X \rightarrow Y}$ .

**Definition 1 (Discrete Memoryless Channel)** (*Cover and Thomas, 2006; Hamming, 1986*) A *discrete memoryless channel (DMC)* is a probabilistic system with a discrete input alphabet  $\mathcal{X} = \{1, \dots, |\mathcal{X}|\}$  and a discrete output alphabet  $\mathcal{Y} = \{1, \dots, |\mathcal{Y}|\}$ , specified by the conditional probabilities  $p_{Y|X}(y|x)$ , for  $x \in \mathcal{X}$  and  $y \in \mathcal{Y}$ . The adjective “memoryless” means that, given a sequence of random inputs, the corresponding outputs are conditionally independent.

**Definition 2 (Uniform channel (UC))** (*Hamming, 1986*) A UC is a DMC in which each row of the conditional probability (channel) matrix  $\theta^{X \rightarrow Y}$  is a permutation of every other row.

**Definition 3 (Cyclic uniform channel (CUC))** A CUC is a UC where each row of the channel matrix  $\theta^{X \rightarrow Y}$  is a cyclic permutation of every other row.

Let  $\mathbb{S}_{|\mathcal{Y}|}$  denote the set of all  $|\mathcal{Y}|!$  permutations of  $(1, \dots, |\mathcal{Y}|)$ . In a UC, each row  $\theta_x^{X \rightarrow Y}$  is a row-specific permutation  $\sigma_x \in \mathbb{S}_{|\mathcal{Y}|}$  of a common vector  $\gamma \in \Delta_{|\mathcal{Y}|-1}$ , *i.e.*,

$$\theta_x^{X \rightarrow Y} = (\gamma_{\sigma_x(1)}, \dots, \gamma_{\sigma_x(|\mathcal{Y}|)}) \Leftrightarrow p_{Y|X}(y|x) = \gamma_{\sigma_x(y)}. \quad (1)$$

In the case of a CUC,  $\sigma_x \in \mathbb{C}_{|\mathcal{Y}|}$ , the set of all  $|\mathcal{Y}|$  cyclic permutations of  $(1, \dots, |\mathcal{Y}|)$ .

## 2.2. Uniform Channel Model – UCM – for Categorical Variables

We propose a new principle to infer the most likely causal direction between two categorical variables by following the rationale behind ANMs. Recall that the ANM principle for real variables is as follows: if  $Y$  satisfies an ANM  $Y = f_Y(X) + N_Y$ , where  $N_Y \perp\!\!\!\perp X$  (i.e., the *exogenous* noise is independent of  $X$ ), but the same is not true in the reverse direction, then the most likely causal direction is  $X \rightarrow Y$ . In the ANM for real variables, the conditional *probability density function* (pdf) has identical *shape* for all values of  $x$ , simply being *shifted* by  $f_Y(x)$ , i.e.,  $p_{Y|X}(y|x) = p_{N_Y}(y - f_Y(x))$ , where  $p_{N_Y}$  is the pdf of the noise variable  $N_Y$ . Consequently, the *conditional differential entropy*  $h(Y|X)$  does not depend on the pfd of  $X$ , as shown next.

**Proposition 4** *If real-valued variables  $X$  and  $Y$  admit an ANM from  $X$  to  $Y$ , then the conditional differential entropy  $h(Y|X) = h(N_Y)$ , independently of the distribution of  $X$ .*

**Proof:** Using the shift-invariance (a) of differential entropy (Cover and Thomas, 2006),

$$\begin{aligned} h(Y|X) &= \mathbb{E}_{X,Y}[-\log p_{Y|X}(Y|X)] = \mathbb{E}_X[\mathbb{E}_{Y|X}[-\log p_{Y|X}(Y|X)]] \\ &= \mathbb{E}_X[\mathbb{E}_{Y|X}[-\log p_{N_Y}(Y - f_Y(X))]] \stackrel{(a)}{=} \mathbb{E}_X[h(N_Y)] = h(N_Y). \quad \blacksquare \end{aligned}$$

For categorical variables, the sets  $\mathcal{X}$  and  $\mathcal{Y}$  lack any meaningful order, thus there is no notion of *addition*, and an ANM is not directly applicable. However, the conditional entropy invariance property of ANMs can be preserved by considering the transformation group under which discrete entropy is invariant: *permutations*. Consequently, our proposed causal inference principle is:

**UCM causal inference principle for categorical variables:** given two categorical variables  $X$  and  $Y$ , if the conditional pmf  $\theta^{X \rightarrow Y}$  corresponds to a UCM, but the conditional pmf  $\theta^{Y \rightarrow X}$  does not, then we infer the causal direction to be  $X \rightarrow Y$ .

Paralleling Proposition 4, the following result is a simple consequence of the invariance of (discrete) entropy to symbol permutations.

**Proposition 5** *If  $\theta^{X \rightarrow Y}$  corresponds to a UC (each row of  $\theta^{X \rightarrow Y}$  is a permutation of a vector  $\gamma \in \Delta_{|\mathcal{Y}|-1}$ ), then the conditional entropy  $H(Y|X) = H(\gamma)$ , independently of  $p_X$ .*

**Proof:** Due to the permutation-invariance property (a) of entropy (Cover and Thomas, 2006),

$$\begin{aligned} H(Y|X) &= \mathbb{E}_{X,Y}[-\log p_{Y|X}(Y|X)] = \mathbb{E}_X[\mathbb{E}_{Y|X}[-\log p_{Y|X}(Y|X)]] \\ &= \mathbb{E}_X[\mathbb{E}_{Y|X}[-\log \gamma_{\sigma_X(Y)}]] \stackrel{(a)}{=} \mathbb{E}_X[H(\gamma)] = H(\gamma). \quad \blacksquare \end{aligned}$$

The proposed causal inference problem can be seen as an instance of the *independence of cause and mechanism* principle, with *independence* corresponding to the following property: the conditional pmf of the effect given the cause has the same collection of probability values, only their positions depend on the cause. Thus, the conditional uncertainty (entropy) of the effect, given the cause, is independent of the distribution of the cause.

A relevant fact that provides further support to the proposed principle is that the UCM can be written as an SCM (Pearl, 2009), as shown in the following proposition.

**Proposition 6** *Let  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y}$  be a pair of dependent random variables such that the conditional pmf  $\theta^{X \rightarrow Y}$  corresponds to a UCM specified by  $\gamma$  and  $\sigma_1, \dots, \sigma_{|\mathcal{X}|}$ , as in (1), and the marginal  $\theta^X$  be arbitrary. Then, the joint pmf of  $X$  and  $Y$  is entailed by the following SCM:*

$$X := U_X, \quad Y := f_Y(X, U_Y), \quad (2)$$

with  $U_X \perp\!\!\!\perp U_Y$  (independent exogenous variables),  $U_X \in \mathcal{X}$  has pmf  $\theta^X$ ,  $U_Y \in \mathcal{Y}$  has pmf  $\gamma$ , and  $f_Y : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Y}$  is a function given by  $f_Y(x, u) = \tau_x(u)$ , with  $\tau_x = \sigma_x^{-1}$  (inverse<sup>4</sup> permutation of  $\sigma_x$ ). Conversely, if the conditional pmf  $\theta^{X \rightarrow Y}$  does not correspond to a UCM, it is impossible to write an SCM of the form (2), with  $U_Y \in \mathcal{Y}$ , entailing the joint pmf of  $X$  and  $Y$ .

**Proof:** Given  $x \in \mathcal{X}$ ,  $Y = f_Y(x, U)$  is a categorical random variable with conditional pmf

$$p_{Y|X}(y|x) = \mathbb{P}[Y = y|X = x] = \mathbb{P}[U_Y = \tau_x^{-1}(y)] = \gamma_{\sigma_x(y)}$$

(which coincides with (1)), where the second equality stems from the definition of  $f_Y$  and permutations being bijections. Conversely, if the conditional pmf  $\theta^{X \rightarrow Y}$  does not correspond to a UCM (neither all its rows are equal to each other, by assumption), depending on what value/category  $X$  takes, the conditional pmf of  $Y$  takes different probability values, not just a permutation of a common pmf, making it impossible to write an SCM of the form (2) with  $U_Y \in \mathcal{Y}$  independent of  $U_X$ . ■

The restriction of having  $U_Y \in \mathcal{Y}$  is what makes this result non-trivial. In fact, [Kocaoglu et al. \(2017\)](#) showed that, given any joint pmf  $p_{X,Y}$ , it is possible to write an SCM  $Y := f_Y(X, U_Y)$ , with  $U_X \perp\!\!\!\perp U_Y$ , that induces it and such that  $U_Y$  takes values in a set of cardinality  $O(|\mathcal{X}| |\mathcal{Y}|)$ .

**Remark 7** *If none, or both, of the conditional pmfs,  $\theta^{X \rightarrow Y}$  and  $\theta^{Y \rightarrow X}$ , correspond to a UCM, the proposed criterion does not select a causal direction. Of course, in practice, the conditional pmfs are estimated from a finite dataset, thus there is a very small chance that one of these estimates corresponds exactly to a UCM. In Section 4, we come back to this issue, proposing statistical tests do decide if a conditional pmf estimate can be considered to correspond to a UCM. In the following subsection, we assume that we have an infinite amount of data (equivalently, the true underlying pmf) and address the identifiability issue in this ideal condition. If this model was not identifiable in this ideal setting, it would be hard to argue that it could be useful with a finite amount of data.*

**Remark 8** *Our UCM contains as a particular case the model for “structureless” sets proposed by [Peters et al. \(2011\)](#). Their model assumes a function  $\phi : \mathcal{X} \rightarrow \mathcal{Y}$  and  $p_{Y|X}(y|x) = p$ , if  $y = \phi(x)$ , and  $p_{Y|X}(y|x) = (1 - p)/(|\mathcal{Y}| - 1)$ , if  $y \neq \phi(x)$ . This corresponds to a UC (in fact, a CUC), with*

$$\gamma = (p, (1 - p)/(|\mathcal{Y}| - 1), \dots, (1 - p)/(|\mathcal{Y}| - 1))$$

and any set of permutations such that  $\sigma_x(y) = 1$ , for  $y = \phi(x)$ . Our UC and CUC models are much more general, as they do not constrain the conditional pmf to have only two different values.

4. Given a permutation  $\sigma \in \mathbb{S}_{|\mathcal{Y}|}$ , its inverse  $\sigma^{-1} \in \mathbb{S}_{|\mathcal{Y}|}$  is such that  $\sigma^{-1}(\sigma(i)) = i$ , for any  $i = 1, \dots, |\mathcal{Y}|$ .

### 2.3. Identifiability

For the proposed criterion to be useful, it should be supported by an identifiability guarantee, *i.e.*, that the set of joint probability mass functions  $p_{X,Y}$  such that both  $p_{Y|X}$  and  $p_{X|Y}$  correspond to uniform channels should be as small as possible, ideally have zero Lebesgue measure in the space of valid parameters, thus zero probability under any continuous density (Peters et al., 2011). Before stating and proving the general identifiability result, we illustrate it for the case where both variables are binary:  $\mathcal{X} = \mathcal{Y} = \{1, 2\}$ . Let  $p_X(1) = \theta_1^X = \beta$  and let  $\theta^{X \rightarrow Y}$  correspond to a UCM (in this case, simply a *binary symmetric channel*) with error probability  $\alpha$  (Cover and Thomas, 2006):

$$\theta^{X \rightarrow Y} = \begin{bmatrix} 1 - \alpha & \alpha \\ \alpha & 1 - \alpha \end{bmatrix}.$$

Of course, a channel matrix where the two rows are equal to  $(1 - \alpha, \alpha)$  is also a UC, but in that case  $X$  and  $Y$  are independent, which is an uninteresting case. The channel in the reverse direction, *i.e.*,  $\theta^{Y \rightarrow X}$ , can be easily derived using Bayes law, yielding

$$\theta^{Y \rightarrow X} = \begin{bmatrix} \frac{(1 - \alpha)\beta}{(1 - \alpha)\beta + \alpha(1 - \beta)} & \frac{\alpha(1 - \beta)}{(1 - \alpha)\beta + \alpha(1 - \beta)} \\ \frac{\alpha\beta}{\alpha\beta + (1 - \alpha)(1 - \beta)} & \frac{(1 - \alpha)(1 - \beta)}{\alpha\beta + (1 - \alpha)(1 - \beta)} \end{bmatrix}.$$

Notice that in matrix  $\theta^{Y \rightarrow X}$ , the variable  $Y$  indexes rows and  $X$  indexes columns, so that it is row-stochastic as is standard for channel matrices. Matrix  $\theta^{Y \rightarrow X}$  represents a UC if and only if one (or both) of two conditions are satisfied: the diagonal elements are equal to each other; the elements in the first column are equal to each other (in which case,  $X \perp\!\!\!\perp Y$ ). Simple algebraic manipulation allows showing that this is equivalent to having  $(\alpha, \beta) \in \{(\alpha, \beta) \in [0, 1]^2 : \alpha = 0 \vee \alpha = 1/2 \vee \alpha = 1 \vee \beta = 0 \vee \beta = 1/2 \vee \beta = 1\}$ , which has zero Lebesgue measure. The following theorem generalizes this result for arbitrary  $|\mathcal{X}|$  and  $|\mathcal{Y}|$ .

**Theorem 9** *Let  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y}$  be two categorical random variables with a joint pmf such that the conditional  $\theta^{X \rightarrow Y}$  corresponds to a UC. Assume also that the marginals have full support<sup>5</sup>:  $p_Y(y) \neq 0$ , for any  $y \in \mathcal{Y}$ , and  $p_X(x) \neq 0$ , for any  $x \in \mathcal{X}$ . Further assume that the rows of the channel matrix  $\theta^{X \rightarrow Y}$  are not all equal to each other (*i.e.*,  $X$  and  $Y$  are not independent<sup>6</sup>). Then, the set of parameters such that the reverse channel  $\theta^{Y \rightarrow X}$  is also a UCM has zero Lebesgue measure.*

The proof, presented in Appendix A, essentially boils down to showing that the UC condition on the reverse channel  $\theta^{Y \rightarrow X}$  corresponds to the zero set of a polynomial that is not identically zero, which thus has zero Lebesgue measure, a classical result from polynomial theory (Federer, 1969).

## 3. Channel Estimation from Data

This section addresses the problem of estimating channel parameters  $\theta^{X \rightarrow Y}$  from  $N$  independent and identically distributed samples of  $(X, Y) : (x_1, y_1), \dots, (x_N, y_N)$ . This will play a key role in

5. There is no loss of generality in this assumption; if there are zeros in the marginals, we simply redefine  $\mathcal{X}$  or/and  $\mathcal{Y}$  by removing the zero-probability elements.

6. If all the rows are equal to each other, then  $Y \perp\!\!\!\perp X$ ; since independence is a symmetrical relationship, the reverse channel  $\theta^{Y \rightarrow X}$  will also have all its rows equal to each other, thus being a special case of a UC channel.

translating the causal inference criterion proposed in Section 2 to the realistic scenario where there is only access to a finite amount of data, rather than perfect knowledge of the joint pmf  $p_{X,Y}$ .

Before estimating the channel parameters (conditional pmf), notice that estimating the marginal pmf  $\theta^X$  is trivial: with  $N_x$  denoting the number of samples  $(x_i, y_i)$  with  $x_i = x$ , the *maximum likelihood* (ML) estimate of  $\theta^X$  is given by

$$\hat{\theta}^X = \arg \max_{\theta \in \Delta_{|\mathcal{X}|-1}} \sum_{x \in \mathcal{X}} N_x \log \theta_x^X = \left( \frac{N_1}{N}, \dots, \frac{N_{|\mathcal{X}|}}{N} \right). \quad (3)$$

For estimating  $\theta^{X \rightarrow Y}$ , we consider the following 4 scenarios: **(1)** arbitrary channel; **(2)** UCM, with known permutations; **(3)** UCM, with *unknown* permutations; **(4)** CUCM, with *unknown* cyclic permutations. Scenarios 1 and 2 are trivial and considered only as they provide the building blocks to address scenarios 3 and 4.

### 3.1. Scenario 1: Arbitrary Channel

Let  $N_{x,y}$  be the number of samples  $(x_i, y_i)$  such that  $x_i = x$  and  $y_i = y$ . In the absence of constraints other than each row of  $\theta^{X \rightarrow Y}$  must be a valid pmf, the ML estimate is

$$\hat{\theta}^{X \rightarrow Y} = \arg \max_{\theta \in (\Delta_{|\mathcal{Y}|-1})^{|\mathcal{X}|}} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x,y} \log \theta_{x,y}. \quad (4)$$

Since both the objective function and the constraints in (4) are separable across  $x = 1, \dots, |\mathcal{X}|$ , the problem is also separable into a collection of problems, each yielding the classical ML estimates

$$\hat{\theta}_{x,y}^{X \rightarrow Y} = N_{x,y}/N_x, \quad \text{for } (x, y) \in \mathcal{X} \times \mathcal{Y}. \quad (5)$$

### 3.2. Scenario 2: UCM with Known Permutations

In a UC, each row  $\theta_x^{X \rightarrow Y}$  is as given in (1). If the permutations  $\sigma_1, \dots, \sigma_{|\mathcal{X}|}$  are known, the ML estimate of  $\gamma$  is given by

$$\hat{\gamma} = \arg \max_{\gamma \in \Delta_{|\mathcal{Y}|-1}} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x,y} \log \gamma_{\sigma_x(y)}. \quad (6)$$

This problem is not separable, as all the rows of the channel matrix share the same probability values, although with different permutations. Swapping the summation order and using the inverse permutations  $\tau_x = \sigma_x^{-1}$  to do a change of variable in the sum over  $\mathcal{Y}$ , problem (6) can be rewritten as

$$\hat{\gamma} = \arg \max_{\gamma \in \Delta_{|\mathcal{Y}|-1}} \sum_{z \in \mathcal{Y}} \log \gamma_z \sum_{x \in \mathcal{X}} N_{x, \tau_x(z)}. \quad (7)$$

Problem (7) has the same form as (3), the solution being simply

$$\hat{\gamma}_y = \frac{1}{N} \sum_{x \in \mathcal{X}} N_{x, \tau_x(y)}, \quad \text{for } y \in \mathcal{Y}. \quad (8)$$

Notice that  $N_{x, \tau_x(y)}$  is the number of samples  $(x_i, y_i)$  such that  $x_i = x$  and  $y_i = \tau_x(y)$ .



### 3.3. Scenario 3: UC with Unknown Permutations

In this case, the log-likelihood is maximized, not only w.r.t.  $\gamma$ , but also the permutations. Although, at first sight, this may look like a very hard problem, as there are  $(|\mathcal{Y}|!)^{|\mathcal{X}|}$  combinations of permutations, we show next that it can be solved very efficiently. The optimization problem in hand (formulated w.r.t. the inverse permutations, denoted as  $\tau_1, \dots, \tau_{|\mathcal{X}|} \in \mathbb{S}_{|\mathcal{Y}|}$ ) is

$$\hat{\gamma}, \hat{\tau}_1, \dots, \hat{\tau}_{|\mathcal{X}|} = \arg \max_{\substack{\gamma \in \Delta_{|\mathcal{Y}|-1} \\ \tau_1, \dots, \tau_{|\mathcal{X}|} \in \mathbb{S}_{|\mathcal{Y}|}}} \mathcal{L}(\gamma, \tau_1, \dots, \tau_{|\mathcal{X}|}), \quad (9)$$

where

$$\mathcal{L}(\gamma, \tau_1, \dots, \tau_{|\mathcal{X}|}) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x, \tau_x(y)} \log \gamma_y. \quad (10)$$

The following proposition (proved in Appendix B) provides the solution to this problem.

**Proposition 10** *A globally optimal solution to the problem specified in (9)–(10) is given as follows. For  $x \in \mathcal{X}$ ,  $\hat{\tau}_x$  is any permutation that sorts  $\{N_{x,1}, \dots, N_{x,|\mathcal{Y}|}\}$  into non-increasing order,*

$$\hat{\tau}_x \text{ is such that } N_{x, \hat{\tau}_x(1)} \geq \dots \geq N_{x, \hat{\tau}_x(|\mathcal{Y}|)}, \quad (11)$$

and, for  $y \in \mathcal{Y}$ ,

$$\hat{\gamma}_y = \frac{1}{N} \sum_{x \in \mathcal{X}} N_{x, \hat{\tau}_x(y)}. \quad (12)$$

The solution in (11)–(12) is a global, but not unique, optimum; in fact, any pmf  $\hat{\xi}$  that is a permutation of  $\hat{\gamma}$ , i.e.,  $\hat{\gamma}_y = \hat{\xi}_{\rho(y)}$ , where  $\rho \in \mathbb{S}_{|\mathcal{Y}|}$ , yields

$$\hat{\theta}_{x,y}^{X \rightarrow Y} = \hat{\gamma}_{\hat{\sigma}_x(y)} = \hat{\xi}_{\rho(\hat{\sigma}_x(y))}.$$

That is,  $\gamma$  is identifiable only up to a permutation, since any permutation of  $\hat{\gamma}$ , combined with the inverse of that permutation composed with each  $\sigma_x$ , yields the same conditional pmf estimate  $\hat{\theta}^{X \rightarrow Y}$ , thus the same maximum value of the log-likelihood. Finally, notice that the cost of computing this solution scales as  $O(|\mathcal{X}| |\mathcal{Y}| \log |\mathcal{Y}|)$ , due to the number  $|\mathcal{X}|$  of sorting operations, each of size  $|\mathcal{Y}|$ .

### 3.4. Scenario 4: CUC with Unknown Permutations

The difference between this and the previous case is that the permutations are now cyclic. Thus, the corresponding optimization problem is identical to (9)–(10), but with the constraint  $\tau_1, \dots, \tau_{|\mathcal{X}|} \in \mathbb{S}_{|\mathcal{Y}|}$  replaced with  $\tau_1, \dots, \tau_{|\mathcal{X}|} \in \mathbb{C}_{|\mathcal{Y}|}$ , where  $\mathbb{C}_{|\mathcal{Y}|}$  is the set of cyclic permutations of  $\{1, \dots, |\mathcal{Y}|\}$ .

Although the cardinality of  $\mathbb{C}_{|\mathcal{Y}|}$  is  $|\mathcal{Y}|$ , much smaller than that of  $\mathbb{S}_{|\mathcal{Y}|}$ , which is  $|\mathcal{Y}|!$ , this problem is harder than (9)–(10). Whereas the cost of the exact solution of (9)–(10) scales with  $O(|\mathcal{X}| |\mathcal{Y}| \log |\mathcal{Y}|)$ , exactly solving this problem by exhaustive search costs  $O(|\mathcal{Y}|^{|\mathcal{X}|})$ . It happens that this problem is a variant of a class of problems known as *multireference alignment*, which is known to be NP-hard (Bandeira et al., 2014). Exact solutions are thus out of the question for large problems. Here, we propose an alternating maximization approach with two steps:

- Given the current permutation estimates  $\hat{\tau}_1, \dots, \hat{\tau}_{|\mathcal{X}|}$ , update  $\hat{\gamma}$  according to (8), with  $\tau_x = \hat{\tau}_x$ .
- Given the current  $\hat{\gamma}$ , maximize w.r.t. the permutations, which is separable across  $\tau_1, \dots, \tau_{|\mathcal{X}|}$ :

$$\hat{\tau}_x = \arg \max_{\tau \in \mathbb{C}_{|\mathcal{Y}|}} \sum_{y=1}^{|\mathcal{Y}|} N_{x,\tau(y)} \log \hat{\gamma}_y, \quad \text{for } x = 1, \dots, |\mathcal{X}|. \quad (13)$$

This maximization is carried out exactly by considering all the  $|\mathcal{Y}|$  cyclic permutations.

The costs of both steps of this algorithm scale as  $O(|\mathcal{X}| |\mathcal{Y}|)$ . Convergence can be proved via the same approach that is used to prove convergence of the  $K$ -means algorithm (Selim and Ismail, 1984), since both algorithms share a common structure: alternate between exact maximization with respect to real quantities (cluster centers, in  $K$ -means,  $\gamma$  in our algorithms) and an exact combinatorial optimization (the point-to-cluster assignments in  $K$ -means, the cyclic permutations in the proposed algorithms).

#### 4. Applying the UCM Principle from Data

Applying the proposed causal inference principle amounts to performing hypothesis testing concerning the UC nature of the conditional pmf estimates  $\hat{\theta}^{X \rightarrow Y}$  and  $\hat{\theta}^{Y \rightarrow X}$ . This is closely related to classical tests for two-way contingency tables (Agresti, 2013; Read and Cressie, 1988). Given a table of counts  $N_{x,y}$ , let the *null hypothesis*  $H_0$  be that these counts can be explained by a UCM in the  $X \rightarrow Y$  direction. To test this hypothesis, consider the corresponding maximum log-likelihood (noting that  $p_{X,Y}(x, y) = \mathbb{P}[X = x, Y = y] = \theta_x^X \gamma_{\sigma_x(y)}$ , for a UCM),

$$\mathcal{L}_{H_0} = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x,y} \log(\hat{\theta}_x^X \hat{\gamma}_{\hat{\sigma}_x(y)}) = \sum_{x \in \mathcal{X}} N_x \log \hat{\theta}_x^X + \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x,y} \log(\hat{\gamma}_{\hat{\sigma}_x(y)}), \quad (14)$$

where  $\hat{\theta}^X$ ,  $\hat{\sigma}_1, \dots, \hat{\sigma}_{|\mathcal{X}|}$ , and  $\hat{\gamma}$  are the ML estimates obtained as shown in Section 3.

The alternative hypothesis is that the channel is arbitrary, with maximum log-likelihood

$$\mathcal{L}_{\bar{H}_0} = \sum_{x \in \mathcal{X}} N_x \log \hat{\theta}_x^X + \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x,y} \log(N_{x,y}/N_x), \quad (15)$$

since the ML estimates of the conditional pmf parameters are as given in (5), and the ML estimate of the marginal  $\theta^X$  is the same, regardless of the channel being uniform or not. These models are nested: a UCM is a particular case of the set of all valid channels, thus it is always true that  $\mathcal{L}_{H_0} \leq \mathcal{L}_{\bar{H}_0}$ .

The *likelihood-ratio statistic* (LRS), denoted  $G^2$ , is then given by

$$G_{X \rightarrow Y}^2 = 2(\mathcal{L}_{\bar{H}_0} - \mathcal{L}_{H_0}) = 2 \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x,y} \log\left(\frac{N_{x,y}}{\hat{\gamma}_{\hat{\sigma}_x(y)} N_x}\right); \quad (16)$$

notice that  $\hat{\gamma}_{\hat{\sigma}_x(y)} N_x$  is the expected value of  $N_{x,y}$  under the null hypothesis. This is the LRS in the  $X \rightarrow Y$  direction, which we indicate with the subscript  $X \rightarrow Y$ . The LRS in the reverse direction, denoted  $G_{Y \rightarrow X}^2$ , is computed in the same way, after swapping the roles of  $X$  and  $Y$ .

It is well known that  $G^2$  is asymptotically  $\chi^2$ -distributed with  $\text{df} = (|\mathcal{X}| - 1)(|\mathcal{Y}| - 1)$  *degrees of freedom*, yielding the p-value

$$p = \mathbb{P}[\chi_{\text{df}}^2 \geq G^2] = 1 - \mathbb{P}[\chi_{\text{df}}^2 < G^2],$$

where  $\mathbb{P}[\chi_{df}^2 < G^2]$  is the cumulative distribution function of a  $\chi_{df}^2$  distribution. If  $p$  is less than some significance level (*i.e.*, the test statistic  $G^2$  is too large), the null hypothesis is rejected.

Let  $p^{X \rightarrow Y}$  and  $p^{Y \rightarrow X}$  be the  $p$ -values of the LRS for testing the uniformity of the channels in both directions, and let  $\alpha$  be a significance level for the test (Agresti, 2013; Read and Cressie, 1988), *i.e.*, the null hypothesis is rejected if the  $p$ -value is less than  $\alpha$ . Having a statistical test of whether an estimated conditional pmf corresponds to a UCM, we adopt a procedure similar to the one proposed by Peters et al. (2011).

- If  $p^{X \rightarrow Y} \geq \alpha$  and  $p^{Y \rightarrow X} < \alpha$ , declare  $X \rightarrow Y$ .
- If  $p^{X \rightarrow Y} < \alpha$  and  $p^{Y \rightarrow X} \geq \alpha$ , declare  $Y \rightarrow X$ .
- If  $p^{X \rightarrow Y} < \alpha$  and  $p^{Y \rightarrow X} < \alpha$ , declare "undecided: wrong model".
- If  $p^{X \rightarrow Y} \geq \alpha$  and  $p^{Y \rightarrow X} \geq \alpha$ , declare "undecided: both directions possible".

The fourth case (*i.e.*, the hypotheses that the channel is uniform in both directions cannot be rejected) is asymptotically improbable, unless  $X$  and  $Y$  are independent, due to the identifiability guarantee. Alternatively, to force the method to make a decision between the two causal directions, one may simply decide for  $X \rightarrow Y$ , if  $p^{X \rightarrow Y} > p^{Y \rightarrow X}$ , and for  $Y \rightarrow X$ , otherwise.

## 5. Results

We compare the proposed approach, on synthetic, benchmark, and real data, with two state-of-the-art methods for categorical variables, for which code is publicly available: DC (Liu and Chan, 2016) ([eda.mmci.uni-saarland.de/prj/cisc/](http://eda.mmci.uni-saarland.de/prj/cisc/), with  $\epsilon = 0$ ) and HCR (Cai et al., 2018) (a Python version of the  $R$  code available at [cran.r-project.org/web/packages/HCR/index.html](http://cran.r-project.org/web/packages/HCR/index.html)). The code for all the experiments will be made available upon acceptance of the manuscript.

In the ML estimates underlying our approach (namely (5) and (8)), to avoid the problem of zero or vanishing probabilities, we use a small amount ( $10^{-3}$ ) of additive (a.k.a. Dirichlet) smoothing.

### 5.1. Identifying the UCM Direction

The first set of experiments is a sanity check, assessing the ability of the proposed criterion to identify the UCM direction, using synthetic data, with different sample sizes  $N$  and different sizes of the support sets,  $|\mathcal{X}|$ , and  $|\mathcal{Y}|$ . For each pair  $(|\mathcal{X}|, |\mathcal{Y}|)$  and each  $N$ , we generate 100 independent datasets using randomly generated UCMs in the  $X \rightarrow Y$  direction and the results reported for each  $N$  are the corresponding averages. The decision rule is simply to choose  $X \rightarrow Y$ , if  $p^{X \rightarrow Y} \geq p^{Y \rightarrow X}$  (equivalently,  $G_{X \rightarrow Y}^2 \leq G_{Y \rightarrow X}^2$ ), and  $Y \rightarrow X$  (which is wrong), otherwise. The results in Fig. 1 show that the accuracy achieves high values, close to 100%, for  $N > 500 \sim 1000$ , without a clear effect of the sizes of the support sets or difference between the non-cyclic and the cyclic cases.

### 5.2. Benchmark Data

We use the 112 pairs in the *cause-effect pairs* benchmark set (Guyon et al., 2013) where both variables are categorical and that have as ground truth that either  $X \rightarrow Y$  or  $Y \rightarrow X$ . We set  $\alpha = 0.05$  and compare UCM with the two methods mentioned above: DC (Liu and Chan, 2016) and HCR (Cai et al., 2018). Decisions of "undecided" are counted as wrong. The average accuracies of UCM, DC,

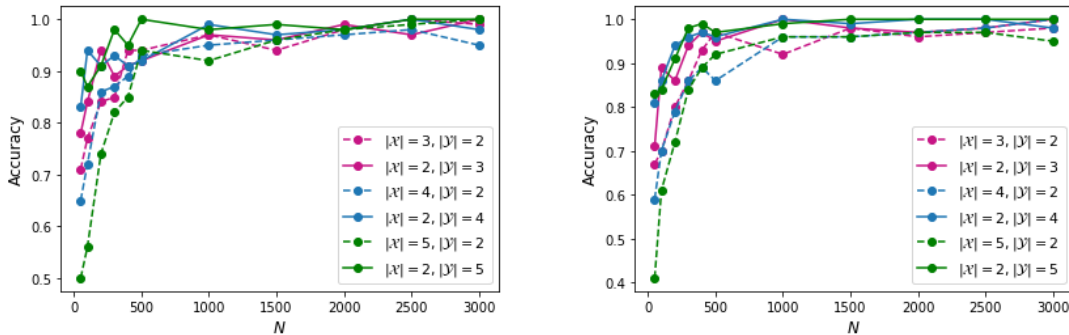


Figure 1: Accuracy in selecting the UCM direction, using the proposed criterion, for different sample and support sizes. Left plot: general UCMs; right plot: cyclic UCMs (CUCM).

and HCR reported in Table 1 shows that UCM outperforms both HCR and DC on this dataset. Notice that a random decision would yield accuracy equal to  $1/3$ .

Table 1: Average accuracy results on the 112 pairs of the benchmark dataset.

UCM	DC	HCR
0.61	0.41	0.47

### 5.3. Real Data

Finally, we evaluate the UCM method (again with  $\alpha = 0.05$ ) on real data from the *UCI Machine Learning Repository* (Bache and Lichman, 2013). We use pairs of variables from the following datasets: Adult, Pittsburgh Bridges, Accute Inflammation, Temperature, and Horse Colic. The datasets and selected pairs, as well as the criteria used to decide what is the ground truth causal direction, are described in Appendix C. We include only pairs for which a test of independence, at significance level 0.05 (Agresti, 2013), rejects the null hypothesis of independence. Furthermore, we include only pairs for which at least one of the three tested methods chooses one of the causal directions. Table 2 shows that the UCM and HCR approaches found the “correct” causal direction in 5 out of 9 pairs, and DC in 4 pairs. UCM returned only correct decisions or abstained from deciding. This small number of experiments does not allow for reaching any strong conclusions but suggests that UCM performs on par, arguably somewhat better, with DC and HCR.

## 6. Conclusions

We introduced the *uniform channel model* (UCM) to address the cause-effect problem with categorical variables. The proposed approach is based on viewing conditional distributions as communication channels. The UCM can be seen as an ANM-type instantiation of the principle of *independence of cause and mechanism*, preserving a key feature of ANM for quantitative data: the conditional entropy (uncertainty) of the effect given the cause is independent of the cause. The core results of

Table 2: Results on real data. Wrong decisions are shown in red; UWM stands for "undecided: wrong model". Month is a cyclic variable, thus a CUC was used in the  $Y \rightarrow X$  direction.

Dataset	$X$	$Y$	UCM	DC	HCR
Adult	Occupation	Income	UWM	$X \rightarrow Y$	$X \rightarrow Y$
Adult	Work Class	Income	UWM	$X \rightarrow Y$	$X \rightarrow Y$
Acute Inflammation	Inflam. of urinary bladder	Lumbar pain	$Y \rightarrow X$	Inconcl.	Inconcl.
Acute Inflammation	Inflam. of urinary bladder	Nausea	$Y \rightarrow X$	Inconcl.	Inconcl.
Acute Inflammation	Inflam. of urinary bladder	Burning urethra	$Y \rightarrow X$	Inconcl.	Inconcl.
Pittsburgh Bridges	Material	Lanes	$X \rightarrow Y$	$Y \rightarrow X$	$X \rightarrow Y$
Pittsburgh Bridges	Purpose	Type	UWM	$Y \rightarrow X$	$X \rightarrow Y$
Temperature	Month	Temperature	$X \rightarrow Y$	$X \rightarrow Y$	$Y \rightarrow X$
Horse Colic	Abdomen Status	Surgical Lesion	UWM	$X \rightarrow Y$	$X \rightarrow Y$

this paper are a proof of the identifiability of the UCM and a proof of its equivalence to a structural causal model with an exogenous variable of fixed cardinality.

To instantiate the approach on finite data, we used classical statistical tests to decide in which of the two directions (if any) the conditional distribution is close enough to correspond to a UCM. The experimental results confirmed the adequacy of the proposed method. By comparing our method with two other recent methods (DC, by [Liu and Chan \(2016\)](#), and HCR, by [Cai et al. \(2018\)](#)), we found that UCM outperforms those other methods on benchmark datasets and performs on par with those methods on real data.

As future work, we will aim to extend the proposed method to handle more than two variables. For example, closeness to a uniform channel of the conditional distribution of each variable given its parents can be used in a score-based method. Another direction of research will look at cases where one variable is categorical and the other is continuous. In fact, we envision a generalization that subsumes both ANM and the proposed UCM as follows. In the correct causal direction, the conditional distributions of the effect should be equivariant under some transformation group that is relevant to its domain: shifts, in the ANM case, permutations, for categorical variables, cyclic permutations for variables with cyclic structure. If this equivariance holds in the causal direction but not in the reverse one, we have identifiability.

## Acknowledgments

We thank Afonso Bandeira, André Gomes, Francisco Andrade, João Xavier, and José Mourão, for fruitful conversations and relevant suggestions. We also thank the anonymous reviewers for some important suggestions.

## References

- A. Agresti. *Categorical Data Analysis*. Wiley, 2013. 3rd Edition.
- K. Bache and M. Lichman. UCI Machine Learning Repository, 2013. <http://archive.ics.uci.edu/ml/index.php>.
- A. Bandeira, M. Charikar, A. Singer, and A. Zhu. Multireference alignment using semidefinite programming. In *5th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 459–470, 2014.
- M. Best and N. Chakravarti. Active set algorithms for isotonic regression: An unifying framework. *Mathematical Programming*, 47:425–439, 1990.
- K. Budhathoki and J. Vreeken. MDL for causal inference on discrete data. In *IEEE International Conference on Data Mining*, pages 751–756, 2017a.
- K. Budhathoki and J. Vreeken. Causal inference by stochastic complexity, 2017b. arXiv:1702.06776.
- K. Budhathoki and J. Vreeken. Accurate causal inference on discrete data. In *IEEE International Conference on Data Mining*, pages 881–886, 2018.
- R. Cai, J. Qiao, K. Zhang, Z. Zhang, and Z. Hao. Causal discovery from discrete data using hidden compact representation. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Neural Information Processing Systems*, 2018.
- D. Chickering. Learning equivalence classes of Bayesian-network structures. *Journal of Machine Learning Research*, 2:445–498, 2002.
- T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2006.
- G. Faria, A. Martins, and M. Figueiredo. Differentiable causal discovery under latent interventions. In B. Schölkopf, C. Uhler, and K. Zhang, editors, *First Conference on Causal Learning and Reasoning (CLEaR)*, volume 177 of *Proc. of Machine Learning Research*, pages 253–274, 2022.
- H. Federer. *Geometric Measure Theory*. Springer-Verlag, 1969.
- M. Gao and B. Aragam. Efficient Bayesian network structure learning via local markov boundary search. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, pages 4301–4313, 2021.
- I. Guyon, A. Statnikov, and M. Henaff. Causality challenge #3: Cause-effect pairs, 2013. <http://www.causality.inf.ethz.ch/cause-effect.php>.
- R. Hamming. *Coding and Information Theory*. Prentice-Hall, 1986.
- G. Hardy, J. Littlewood, and G. Pólya. *Inequalities*. Cambridge University Press, 1952.
- D. Heckerman and D. Geiger. Learning Bayesian networks: A unification for discrete and Gaussian domains. In *Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 274–284, 1995.

- P. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Processing Information Systems*, page 689–696, 2009.
- D. Janzing. The cause-effect problem: Motivation, ideas, and popular misconceptions. In I. Guyon, A. Statnikov, and B. Batu, editors, *Cause Effect Pairs in Machine Learning*, pages 3–26. Springer, 2019.
- D. Janzing, P. Hoyer, and B. Schölkopf. Telling cause from effect based on high-dimensional observations. In *International Conference on Machine Learning*, pages 479–486, 2010.
- D. Janzing, J. Mooij, K. Zhang, J. Lemeire, J. Zscheischler, P. Daniušis, B. Steudel, and B. Schölkopf. Information-geometric approach to inferring causal directions. *Artificial Intelligence*, 182-183: 1–31, 2012.
- M. Kocaoglu, A. Dimakis, S. Vishwanath, and B. Hassibi. Entropic Causal Inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2010.
- S. Lang. *Algebra*. Springer–Verlag, 2002. 3rd Edition.
- M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, 2009.
- F. Liu and L. Chan. Causal inference on discrete data via estimating distance correlations. *Neural Computation*, 28(5):801–814, 2016.
- A. Marx and J. Vreeken. Identifiability of cause and effect using regularized regression. In *ACM International Conference on Knowledge Discovery and Data Mining*, pages 852–861, 2019.
- O. Mian, A. Marx, and J. Vreeken. Discovering fully oriented causal networks. *AAAI Conference on Artificial Intelligence*, 35(10):8975–8982, 2021.
- J. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf. Distinguishing cause from effect using observational data: Methods and benchmarks. *Journal of Machine Learning Research*, 17: 1–102, 2016.
- Y. Ni. Bivariate causal discovery for categorical data via classification with optimal label permutation. In A. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- D. Papineau. The statistical nature of causation. *The Monist*, 105(2):247–275, 2022.
- G. Park and G. Raskutti. Learning quadratic variance function (QVF) DAG models via overdispersion scoring (ODS). *Journal of Machine Learning Research*, 18(224):1–44, 2018.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2009.
- J. Peters, D. Janzing, and B. Schölkopf. Causal inference on discrete data using additive noise models. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 33(12):2436–2450, 2011.

- J. Peters, J. Mooij, D. Janzing, and B. Schölkopf. Causal discovery with continuous additive noise models. In *Journal of Machine Learning Research*, volume 15, pages 2009–2053, 2014.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference*. MIT Press, 2017.
- T. Read and N. Cressie. *Goodness-of-Fit Statistics for Discrete Multivariate Data*. Springer, 1988.
- J. Rissanen. Stochastic complexity in statistical inquiry theory. *World Scientific*, 1998.
- Kayvan Sadeghi. Faithfulness of probability distributions and graphs. *Journal of Machine Learning Research*, 18:1–29, 2017.
- S. Selim and M. Ismail. K-means-type algorithms: A generalized convergence theorem and characterization of local optimality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(1):81–87, 1984.
- S. Shimizu, P. Hoyer, A. Hyvarinen, and A. Kerminen. A linear non-Gaussian acyclic model for causal discovery. In *Journal of Machine Learning Research*, volume 7, pages 2003–2030, 2006.
- O. Stegle, D. Janzing, K. Zhang, J. Mooij, and B. Schölkopf. Probabilistic latent variable models for distinguishing cause and effect. In *Neural Information Processing Systems*, 2010.
- N. Tagasovska, V. Chavez-Demoulin, and T. Vatter. Distinguishing cause from effect using quantiles: Bivariate quantile causal discovery. In *International Conference on Machine Learning*, volume PMLR 119, pages 9311–9323, 2020.
- C. Wallace and D. Boulton. An information measure for classification. *Computer Journal*, 11: 185–194, 1968.



## Appendix A. Proof of Theorem 9

**Proof:** Let us denote  $\theta^X = \beta \in \Delta_{|\mathcal{X}|-1}$  and recall that under the UC assumption, matrix  $\theta^{X \rightarrow Y}$  has the form

$$\theta^{X \rightarrow Y} = \begin{bmatrix} \gamma_{\sigma_1(1)} & \gamma_{\sigma_1(2)} & \cdots & \gamma_{\sigma_1(|\mathcal{Y}|)} \\ \gamma_{\sigma_2(1)} & \gamma_{\sigma_2(2)} & \cdots & \gamma_{\sigma_2(|\mathcal{Y}|)} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{\sigma_{|\mathcal{X}|}(1)} & \gamma_{\sigma_{|\mathcal{X}|}(2)} & \cdots & \gamma_{\sigma_{|\mathcal{X}|}(|\mathcal{Y}|)} \end{bmatrix},$$

where  $\sigma_1, \dots, \sigma_{|\mathcal{X}|} \in \mathbb{S}_{|\mathcal{Y}|}$  are permutations and  $\gamma = (\gamma_1, \dots, \gamma_{|\mathcal{Y}|}) \in \Delta_{|\mathcal{Y}|-1}$ . The assumption that the rows of this matrix are not all equal to each other precludes the two following condition from holding:  $\sigma_1 = \sigma_2 = \dots = \sigma_{|\mathcal{X}|}$  and  $\gamma = (1, 1, \dots, 1)/|\mathcal{Y}|$ .

Using Bayes' law, it is trivial to obtain the reverse channel, the elements of which are given by

$$\theta_{y,x}^{Y \rightarrow X} = p_{X|Y}(x|y) = \frac{p_{Y|X}(y|x) p_X(x)}{p_Y(y)} = \frac{\gamma_{\sigma_x(y)} \beta_x}{\sum_{x' \in \mathcal{X}} \gamma_{\sigma_{x'}(y)} \beta_{x'}} = \frac{a_{y,x}}{A_y}, \quad (17)$$

where  $a_{y,x} = \beta_x \gamma_{\sigma_x(y)}$  and  $A_y = p_Y(y) \neq 0$  (by assumption). As in the binary example, using variables  $Y$  and  $X$  to index rows and columns, respectively,  $\theta^{Y \rightarrow X}$  is a row-stochastic matrix:

$$\theta^{Y \rightarrow X} = \begin{bmatrix} a_{1,1}/A_1 & \cdots & a_{1,|\mathcal{X}|}/A_1 \\ \vdots & \ddots & \vdots \\ a_{|\mathcal{Y}|,1}/A_{|\mathcal{Y}|} & \cdots & a_{|\mathcal{Y}|,|\mathcal{X}|}/A_{|\mathcal{Y}|} \end{bmatrix}.$$

For  $\theta^{Y \rightarrow X}$  to correspond to a UC, its rows must be permutations of each other, which is equivalent to all being permutations of one of them, say the first, without loss of generality. We exclude the case where these permutations are all equal to identity, since that would correspond to all rows of  $\theta^{Y \rightarrow X}$  being equal to each other, *i.e.*,  $X \perp\!\!\!\perp Y$ , which is excluded in the conditions of the theorem. The condition that all the rows are permutations of the first one can be written formally as

$$\exists (\rho_2, \dots, \rho_{|\mathcal{Y}|}) \in \mathbb{L} : \forall y \in \mathcal{Y} \setminus \{1\}, \forall x \in \mathcal{X}, a_{1,x}/A_1 = a_{y,\rho_y(x)}/A_y, \quad (18)$$

where  $\mathbb{L} = (\mathbb{S}_{|\mathcal{X}|})^{|\mathcal{Y}|-1} \setminus \mathbb{I}$ , with  $\mathbb{I} = \{\rho_2, \dots, \rho_{|\mathcal{Y}|} : \rho_2 = \dots = \rho_{|\mathcal{Y}|} = \iota\}$ , and  $\iota$  is the identity permutation. In words,  $\mathbb{L}$  is the set of all  $(|\mathcal{Y}|-1)$ -tuples of permutations of  $|\mathcal{X}|$  elements, except for the one in which all permutations are identity.

The equality  $a_{1,x}/A_1 = a_{y,\rho_y(x)}/A_y$  is equivalent to  $(a_{1,x} A_y - a_{y,\rho_y(x)} A_1)^2 = 0$ , thus the following equivalence holds:

$$(\forall y \in \mathcal{Y} \setminus \{1\}, \forall x \in \mathcal{X}, a_{1,x}/A_1 = a_{y,\rho_y(x)}/A_y) \Leftrightarrow Q_\rho(\theta) = 0,$$

with

$$Q_\rho(\theta) = \sum_{y \in \mathcal{Y} \setminus \{1\}} \sum_{x \in \mathcal{X}} (a_{1,x} A_y - a_{y,\rho_y(x)} A_1)^2, \quad (19)$$

where we have written the model parameters compactly as  $\theta = (\beta, \gamma) \in \Delta_{|\mathcal{X}|-1} \times \Delta_{|\mathcal{Y}|-1} \subset \mathbb{R}^{|\mathcal{X}|} \times \mathbb{R}^{|\mathcal{Y}|}$ , and denoted  $\rho = (\rho_2, \dots, \rho_{|\mathcal{Y}|}) \in \mathbb{L}$ . A key observation is that  $Q_\rho(\theta)$  is a polynomial in the elements of  $\theta$ , since the  $a_{y,x}$  and the  $A_y$  are themselves polynomials (either products of two elements or sums of products of pairs of elements) as is clear in (17).

Finally, the existential quantifier in (18) can be re-written using a product, *i.e.*,

$$(\exists \boldsymbol{\rho} \in \mathbb{L} : Q_{\boldsymbol{\rho}}(\boldsymbol{\theta}) = 0) \Leftrightarrow R(\boldsymbol{\theta}) = 0, \quad \text{where } R(\boldsymbol{\theta}) = \prod_{\boldsymbol{\rho} \in \mathbb{L}} Q_{\boldsymbol{\rho}}(\boldsymbol{\theta}).$$

Since  $R(\boldsymbol{\theta})$  a product of polynomials, it is itself a polynomial. Consequently, we have shown that the UC condition in (18) corresponds to having  $\boldsymbol{\theta}$  as a root of a polynomial.

The rest of the proof relies on a classical result about polynomials (Federer, 1969): let  $S : \mathbb{R}^n \rightarrow \mathbb{R}$  be a polynomial that is not identically zero; then, the set  $S^{-1}(0) = \{\mathbf{u} \in \mathbb{R}^n : S(\mathbf{u}) = 0\}$  has zero Lebesgue measure in  $\mathbb{R}^n$ . All that is left to show then is that  $R(\boldsymbol{\theta})$  is not identically zero. For this purpose, we can ignore the valid parameter space  $\Delta_{|\mathcal{X}|-1} \times \Delta_{|\mathcal{Y}|-1}$ , because if  $R^{-1}(0)$  has zero Lebesgue measure in  $\mathbb{R}^{|\mathcal{X}|} \times \mathbb{R}^{|\mathcal{Y}|}$ , so does the intersection  $R^{-1}(0) \cap (\Delta_{|\mathcal{X}|-1} \times \Delta_{|\mathcal{Y}|-1})$ . We can also ignore the condition  $\boldsymbol{\gamma} \neq (1, \dots, 1)/|\mathcal{Y}|$ , since this is a single point, thus a set of zero measure.

A sufficient and necessary condition for  $R(\boldsymbol{\theta})$  not to be identically zero is that none of its factors  $Q_{\boldsymbol{\rho}}(\boldsymbol{\theta})$  is identically zero<sup>7</sup>. To show that no  $Q_{\boldsymbol{\rho}}(\boldsymbol{\theta})$  is identically zero, let us write it explicitly, using the definitions of  $a_{y,x}$  and  $A_y$  in (17):

$$Q_{\boldsymbol{\rho}}(\boldsymbol{\theta}) = \sum_{y \in \mathcal{Y} \setminus \{1\}} \sum_{x \in \mathcal{X}} \left( \beta_x \gamma_{\sigma_x(1)} \sum_{x' \in \mathcal{X}} \beta_{x'} \gamma_{\sigma_{x'}(y)} - \beta_{\rho_y(x)} \gamma_{\sigma_{\rho_y(x)}(y)} \sum_{x' \in \mathcal{X}} \beta_{x'} \gamma_{\sigma_{x'}(1)} \right)^2. \quad (20)$$

Since  $Q_{\boldsymbol{\rho}}(\boldsymbol{\theta})$  is a sum of non-negative terms, to show that it is not identically zero, it suffices to show that one of the terms in the sum is strictly positive for some choice of  $\boldsymbol{\theta}$ . The condition  $\boldsymbol{\rho} = (\rho_2, \dots, \rho_{|\mathcal{Y}|}) \in \mathbb{L}$  means that at least one of the permutations  $\rho_2, \dots, \rho_{|\mathcal{Y}|}$  is not the identity, which implies that there is at least one pair  $(x, y)$  such that  $\rho_y(x) \neq x$ . Let  $y$  and  $x$  be one such pair. Choosing  $\boldsymbol{\gamma} = (1, \dots, 1)$  and  $\boldsymbol{\beta} \in \Delta_{|\mathcal{X}|-1}$  such that all components are different from each other ( $i \neq j \Rightarrow \beta_i \neq \beta_j$ ), we have (noticing that  $\sum_{x' \in \mathcal{X}} \beta_{x'} = 1$ )

$$Q_{\boldsymbol{\rho}}(\boldsymbol{\theta}) = \left( \beta_x \sum_{x' \in \mathcal{X}} \beta_{x'} - \beta_{\rho_y(x)} \sum_{x' \in \mathcal{X}} \beta_{x'} \right)^2 + \sum_{y' \neq y} \sum_{x' \neq x} (\dots)^2 \quad (21)$$

$$= (\beta_x - \beta_{\rho_y(x)})^2 + \text{non-negative terms} > 0. \quad (22)$$

In conclusion, since none of the  $Q_{\boldsymbol{\rho}}(\boldsymbol{\theta})$  polynomials is identically zero,  $R(\boldsymbol{\theta})$  is also not identically zero, consequently its zero set has zero Lebesgue measure.  $\blacksquare$

## Appendix B. Proof of Proposition 10

**Proof:** Noticing that the permutations that map  $\boldsymbol{\gamma}$  to each row of  $\boldsymbol{\theta}^{Y|X}$  are arbitrary, there is no loss of generality in assuming  $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_{|\mathcal{Y}|}$ , *i.e.*,  $\boldsymbol{\gamma} \in \mathcal{K}_{|\mathcal{Y}|}$ , the so-called *monotone cone* (Best and Chakravarti, 1990). Furthermore, it is more convenient to formulate the problem w.r.t. the inverse permutations, denoted as  $\tau_1, \dots, \tau_{|\mathcal{X}|} \in \mathbb{S}_{|\mathcal{Y}|}$ . The problem can thus be written as

$$\hat{\boldsymbol{\gamma}}, \hat{\tau}_1, \dots, \hat{\tau}_{|\mathcal{X}|} = \arg \max_{\substack{\boldsymbol{\gamma} \in (\Delta_{|\mathcal{Y}|-1} \cap \mathcal{K}_{|\mathcal{Y}|}) \\ \tau_1, \dots, \tau_{|\mathcal{X}|} \in \mathbb{S}_{|\mathcal{Y}|}} \mathcal{L}(\boldsymbol{\gamma}, \tau_1, \dots, \tau_{|\mathcal{X}|}), \quad (23)$$

7. Recall that a product of two polynomials with real coefficients is identically zero only if at least one of the factors is identically zero. This is a classical result from abstract algebra, which in the language thereof is stated as follows: the ring of all polynomials in  $n$  variables with real coefficients is an *integral domain* or *entire ring*, that is, it does not have divisors of zero (Lang, 2002). The result generalizes trivially, by induction, to products of more than two polynomials.

where

$$\mathcal{L}(\gamma, \tau_1, \dots, \tau_{|\mathcal{X}|}) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} N_{x, \tau_x(y)} \log \gamma_y. \quad (24)$$

The assumption  $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_{|\mathcal{Y}|}$  makes the maximization w.r.t.  $\tau_1, \dots, \tau_{|\mathcal{X}|}$  independent of the particular values of  $\gamma$  as well as separable into a collection of  $|\mathcal{X}|$  independent maximizations,

$$\hat{\tau}_x = \arg \max_{\tau_x \in \mathbb{S}_{|\mathcal{Y}|}} \sum_{y \in \mathcal{Y}} N_{x, \tau_x(y)} \log \gamma_y, \quad (25)$$

for  $x \in \mathcal{X}$ . Solving (25) is a simple application of the *rearrangement inequality*<sup>8</sup> (Hardy et al., 1952). Since  $\log \gamma_1 \geq \log \gamma_2 \geq \dots \geq \log \gamma_{|\mathcal{Y}|}$ , the solution  $\hat{\tau}_x$  is any permutation that also sorts  $\{N_{x,1}, \dots, N_{x,|\mathcal{Y}|}\}$  into non-increasing order:

$$\hat{\tau}_x \text{ is such that } N_{x, \hat{\tau}_x(1)} \geq \dots \geq N_{x, \hat{\tau}_x(|\mathcal{Y}|)}. \quad (26)$$

If all the elements of  $\{N_{x,1}, \dots, N_{x,|\mathcal{Y}|}\}$  are different, the optimal permutation is unique; otherwise, there are several optimal permutations, all achieving the same maximum. The cost of finding  $\hat{\tau}_1, \dots, \hat{\tau}_{|\mathcal{X}|}$  is  $O(|\mathcal{X}| |\mathcal{Y}| \log |\mathcal{Y}|)$ , since it requires  $|\mathcal{X}|$  sorting operations, each with  $|\mathcal{Y}|$  elements.

Plugging  $\hat{\tau}_1, \dots, \hat{\tau}_{|\mathcal{X}|}$  back into (9)–(10) and swapping the summation order, yields

$$\hat{\gamma} = \arg \max_{\gamma \in (\Delta_{|\mathcal{Y}|-1} \cap \mathcal{K}_{|\mathcal{Y}|})} \sum_{y \in \mathcal{Y}} \log \gamma_y \sum_{x \in \mathcal{X}} N_{x, \hat{\tau}_x(y)}. \quad (27)$$

This problem is the same as (7), with  $\hat{\tau}_x$  in the place of  $\tau_x$  and with the additional constraint  $\gamma \in \mathcal{K}_{|\mathcal{Y}|}$ . Temporarily ignoring this constraint leads to (see (8)),

$$\hat{\gamma}_y = \frac{1}{N} \sum_{x \in \mathcal{X}} N_{x, \hat{\tau}_x(y)}, \quad \text{for } y \in \mathcal{Y}. \quad (28)$$

The fact that  $N_{x, \hat{\tau}_x(1)} \geq N_{x, \hat{\tau}_x(2)} \geq \dots \geq N_{x, \hat{\tau}_x(|\mathcal{Y}|)}$ , for any  $x \in \mathcal{X}$ , implies that  $\hat{\gamma} \in \mathcal{K}_{|\mathcal{Y}|}$ , without having to include this constraint. Consequently, problem (9)–(10) has a global solution given by (28) and (26).  $\blacksquare$

### Appendix C. Detailed Description of the Datasets Used in Section 5.3

**Adult** - This dataset consists of 48832 records from the census database of the US in 1994. We consider the following pairs: (*occupation, income*) and (*work class, income*). The variable *occupation* takes values in the set  $\{\text{admin, armed-force, blue-collar, white-collar, service, sales, professional, other-occupation}\}$ . The variable *work class* takes categories in  $\{\text{private, self-employed, public servant, unemployed}\}$ . Finally, *income* is a binary variable taking value in  $\{> 50, \leq 50\}$ . Following

8. Given any two non-decreasing sequences of  $n$  real numbers,  $x_1 \leq \dots \leq x_n$  and  $y_1 \leq \dots \leq y_n$ ,

$$\forall \sigma \in \mathbb{S}_n, \quad \sum_{i=1}^n x_{\sigma(i)} y_i \leq \sum_{i=1}^n x_i y_i.$$

Budhathoki and Vreeken (2017b), we assume that the ground truth is *occupation*  $\rightarrow$  *income* and *work class*  $\rightarrow$  *income*.

**Pittsburgh Bridges** - This dataset contains records about 108 bridges and some of their characteristics. We consider the pairs (*purpose, type*) and (*material, lanes*). The variable *purpose* takes values in {*Walk, Aqueduct, RR, Highway*}; variable *type* takes values in {*Wood, Suspen, Simple-T, Arch, Cantilev, CONT-T*}; variable *material* takes values in {*Steel, Iron, Wood*}; variable *lanes* takes values in {1, 2, 4, 6}. Following Cai et al. (2018), the ground truths assumed are *material*  $\rightarrow$  *lanes* and *purpose*  $\rightarrow$  *type*.

**Acute Inflammations** - This dataset consists of 120 patients and whether each patient is experiencing a specific symptom, the temperature, and whether he/she suffers from acute inflammations of the urinary bladder and/or acute nephritis. We consider the binary variables *ocurrence of nausea* ( $Y_1$ ), *lumbar pain* ( $Y_2$ ), and *burning of urethra* ( $Y_3$ ) (naturally, taking value 1 if the patient has that symptom, and 0 otherwise). Moreover,  $X$  represent the diagnosis *inflammation of urinary bladder*. Following Peters et al. (2011), the goal is to model the diagnosis process, thus we expect  $Y_j \rightarrow X$ , for  $j = 1, \dots, 3$ . Notice that the variables  $X_i$  only correspond to the diagnosis, not necessarily the truth, otherwise they would be considered the cause, rather than the effect.

The  $\chi^2$  test did not reject the null hypothesis of independence (at a significance level of 5%) for the following pairs of variables in this dataset: (*Inflammation of urinary bladder, Occurrence of nausea*); (*Inflammation of urinary bladder, Burning of urethra*); (*Nephritis, Micturition pains*). Although, intuitively, we would expected a causal relation between those, the statistical evidence is not strong enough in favour of their mutual dependency (according to the  $\chi^2$  test), therefore, these pairs were discarded from the experiments.

**Temperature** - This dataset consists of 9162 daily values of temperature measured in Furtwangen (Germany), with the variable *day of the year* taking integer values from 1 to 365 (or 366 for leap years) and *temperature* in  $^{\circ}C$ . Here, we aggregate days associated with each month and take *month*  $\rightarrow$  *temperature* as the ground truth, as Peters et al. (2011). Notice that *month* assumes a cyclic structure.

**Horse Colic** - This dataset contains 368 medical records of horses. We study the causal relationship between the variable *abdomen status*, which takes the values in {*Normal, Other, Firm feces in the large intestine, Distended small intestine, Distended large intestine*}, and the binary variable *surgical lesion*, indicating whether the lesion was surgical or not. As Budhathoki and Vreeken (2018), we regard *abdomen status*  $\rightarrow$  *surgical lesion* as ground truth.