

Structural Causal Models Are (Solvable by) Credal Networks

Marco Zaffalon

ZAFFALON@IDSIA.CH

Alessandro Antonucci

ALESSANDRO@IDSIA.CH

Rafael Cabañas

RCABANAS@IDSIA.CH

Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland

Abstract

A structural causal model is made of endogenous (manifest) and exogenous (latent) variables. We show that endogenous observations induce linear constraints on the probabilities of the exogenous variables. This allows to *exactly* map a causal model into a *credal network*. Causal inferences, such as interventions and counterfactuals, can consequently be obtained by standard algorithms for the updating of credal nets. These natively return sharp values in the identifiable case, while intervals corresponding to the exact bounds are produced for unidentifiable queries. A characterization of the causal models that allow the map above to be compactly derived is given, along with a discussion about the scalability for general models. This contribution should be regarded as a systematic approach to represent structural causal models by credal networks and hence to systematically compute causal inferences. A number of demonstrative examples is presented to clarify our methodology. Extensive experiments show that approximate algorithms for credal networks can immediately be used to do causal inference in real-size problems.

Keywords: Structural causal models; identifiability; credal nets; interventions; counterfactuals.

1. Introduction

Since early times, dealing with causality has been—and under many respects, still is—a true challenge for scientists and philosophers (Hume, 1739). Currently causality is an emerging direction for data science (e.g., Correa and Bareinboim, 2020), with a wealth of potential applications in diverse domains such as planning (Wilkins, 2014) or NLP (Asghar, 2016), not to mention fields other than Artificial Intelligence such as Economics, Social Science or Medicine.

Pearl's structural causal models are a natural formalism for causal inference (Pearl, 2009), in particular for their appealing graphical representation. They are also very general and equivalent to the prominent alternative formalisms proposed to handle causality. However, the peculiar features of causal models may render them not always easy to access to a traditional audience, which is instead familiar with pre-existing graphical tools and related procedures.

In this paper, we focus on Pearl's non-parametric structural causal models with discrete variables, and show that they can be represented by *credal networks* (Cozman, 2000). This is a class of imprecise-probabilistic graphical models originally proposed as tools for sensitivity analysis in Bayesian networks. The representation is exact: this means that every query on the causal model can be reformulated as a query on the credal network, which can then be solved by standard algorithms for the latter. An immediate advantage of this outcome is that causal concepts naturally become more familiar concepts, such as the updating of credal networks. Another, more practical, advantage is that such an outcome allows us to systematically compute causal inference via credal nets, be it interventional queries or the more advanced counterfactual type of inference.

It is well known that causal inference is affected by ‘identification’ problems: an inference is said to be *identifiable* when it can be reduced to a precise probabilistic expression (a number representing a probability or an expectation), which can eventually be computed from data. Otherwise it is called *unidentifiable*: one that cannot be reduced to a number no matter the amount of available data. Most of the causal literature focuses on the characterization of models and queries that are identifiable, the most prominent example being the *do calculus* (Pearl, 1995). Relatively few works consider instead the far bigger arena of unidentifiable problems, which are solved by computing tight bounds on the sought probabilities or expectations (Balke and Pearl, 1997; Kang and Tian, 2006; Sachs et al., 2020). The computational results in this paper can be read as contributing in particular to this second direction: in fact, credal nets natively return intervals in the unidentifiable case that correspond to the sought bounds; these bounds automatically collapse to a precise number for the subclass of identifiable problems and queries.

The paper is organized as follows: after providing background material in Section 2, an algorithm to convert a Markovian causal model into a credal network, whose quantification is defined by the observational data, is provided in Section 3. The credal network mapping turns out to be extendible to quite a wider class of models that we call *quasi-Markovian*; the scalability of the specification is discussed in the general case in Section 4. We regard this approach as the most systematic one presented so far in the literature and we discuss its advantages by a number of examples (Section 5). Finally the applicability to real-world cases is tested by numerical simulations showing that algorithms for approximate inference in credal networks allow to compute scalable and informative inferences even for large models (Section 6). We discuss future direction in Section 7. Finally, the proofs of the theorems are in Appendix A.

2. Background Material

2.1 Bayesian and Credal Networks

Let X denote a variable taking values in a finite set Ω_X . The elements of Ω_X are ordered and notation $x^{(k)}$ and x is used, respectively, for the k -th and the generic element of Ω_X . Denote as $P(X)$ a probability mass function (PMF) over X and as $K(X)$ a *credal set* (CS) over X , which is a set of PMFs over X . Given two variables X and Y , a conditional probability table (CPT) $P(X|Y)$ is a collection of (conditional) PMFs indexed by the values of Y , i.e., $\{P(X|y)\}_{y \in \Omega_Y}$. If all PMFs in a CPT are *degenerate*, i.e., there is a state receiving probability mass one and hence all the other ones receive zero, we say that also the CPT is degenerate. A credal CPT (CCPT) $K(X|Y)$ is similarly a collection of CSs over X indexed by the values of Y . With a small abuse of terminology, we might call CPT (CCPT) also a single PMF (CS).

Consider a joint variable $\mathbf{X} := (X_1, \dots, X_n)$ and a directed acyclic graph \mathcal{G} whose nodes are in one-to-one correspondence with the variables in \mathbf{X} (whence we term a node in \mathcal{G} and its corresponding variable interchangeably). Given \mathcal{G} , a Bayesian network (BN) is a collection of CPTs $\{P(X_i|\text{Pa}(X_i))\}_{i=1}^n$, where $\text{Pa}(X_i)$ denotes the *parents* of X_i , i.e., the direct predecessors of X_i according to \mathcal{G} . A BN induces a joint PMF $P(\mathbf{X})$ that factorizes as follows: $P(\mathbf{x}) = \prod_{i=1}^n P(x_i|\text{pa}(X_i))$ for each $\mathbf{x} \in \Omega_X$, where $(x_i, \text{pa}(X_i)) \sim \mathbf{x}$, i.e., $(x_i, \text{pa}(X_i))$ are the values of $(X_i, \text{Pa}(X_i))$ consistent with \mathbf{x} . A credal network (CN) is similarly intended as a collection of CCPTs. A CN defines a CS $K(\mathbf{X})$ whose elements are PMFs factorizing as those of a BN whose CPT values are taken from the corresponding CCPTs. Computing a conditional probability for a queried variable given an evidence with respect to the joint PMF of a BN, as well as the bounds of this probability with

respect to the joint CS of a CN is NP-hard (Mauá et al., 2014). Yet, polynomial algorithms computing approximate inferences for the general case (e.g., Antonucci et al., 2015) or exact ones for classes of sub-models (e.g., Fagnoli and Zaffalon, 1998) are available for CNs, not to mention the copious tools for BNs (e.g., Koller and Friedman, 2009).

2.2 Structural Causal Models

The background concepts in this section are reviewed from the reference book of Pearl (2009). Let us first define a *structural equation* (SE) f_X associated with variable X and based on the input variable(s) Y as a surjective function $f_X : \Omega_Y \rightarrow \Omega_X$ that determines the value of X from that of Y . Consider two sets of variables \mathbf{X} and \mathbf{U} in one-to-one correspondence. We call *endogenous* the variables in the first set and *exogenous* the others. Say that, for each $X \in \mathbf{X}$, an SE f_X associated with X is provided. Assume that the exogenous variable U corresponding to X is an input variable of f_X , possibly together with other endogenous variables, but no other exogenous ones. We call a collection $\{f_X\}_{X \in \mathbf{X}}$ of SEs of this kind a *structural causal model* M (SCM) over \mathbf{X} . The *causal diagram* \mathcal{G}_M of an SCM M is a directed graph over (\mathbf{X}, \mathbf{U}) such that the parents $\text{Pa}(X)$ of X are the input variables of the SE f_X for each $X \in \mathbf{X}$. We denote as $\underline{\text{Pa}}(X)$ the endogenous parents of X according to \mathcal{G}_M , i.e., $\underline{\text{Pa}}(X) := \text{Pa}(X) \setminus \{U\}$, for each $X \in \mathbf{X}$.

Generally speaking \mathcal{G}_M might include directed cycles. The SCM M is *semi-Markovian* if its causal diagram is acyclic. In a semi-Markovian SCM, a joint observation $\mathbf{U} = \mathbf{u}$ of the exogenous variables completely determines the state $\mathbf{X} = \mathbf{x}$ of the endogenous variables. This is achieved by obtaining the endogenous values from the SEs in a topological order for the variables in \mathbf{X} according to \mathcal{G}_M . A *probabilistic* SCM (PSCM) is a pair (M, P) such that M is a semi-Markovian SCM and P a PMF over the exogenous variables \mathbf{U} . Exactly as we obtain \mathbf{x} from \mathbf{u} in SCMs, in PSCMs we obtain an endogenous PMF $P(\mathbf{X})$ from $P(\mathbf{U})$. PSCM (M, P) is finally called *Markovian* if its exogenous variables are jointly independent, i.e., $P(\mathbf{U})$ factorizes as $P(\mathbf{u}) = \prod_{i=1}^n P(u_i)$ for each $\mathbf{u} \in \Omega_{\mathbf{U}}$ and $u_i \in \Omega_{U_i}$ with $u_i \sim \mathbf{u}$, for each $i = 1, \dots, n$. Without loss of generality, we can intend non-Markovian PSCMs as based on SCMs whose exogenous variables might have multiple endogenous children, while keeping the joint independence for $P(\mathbf{U})$.¹ In the following we compactly describe P in a PSCM by $\{P(U)\}_{U \in \mathbf{U}}$ in the place of $P(\mathbf{U})$. A Markovian PSCM is described here below, a non-Markovian one is in Example 5.

Example 1 Consider two endogenous variables (X_1, X_2) and their exogenous variables (U_1, U_2) . Let $\Omega_{X_i} := \{x_i^{(k)}\}_{k=1}^2$ for $i = 1, 2$, $\Omega_{U_1} := \{u_1^{(k)}\}_{k=1}^3$ and $\Omega_{U_2} := \{u_2^{(k)}\}_{k=1}^5$. Define an SE for X_1 given U_1 such that $f_{X_1}(u_1^{(1)}) = x_1^{(1)}$, $f_{X_1}(u_1^{(2)}) = x_1^{(2)}$, and $f_{X_1}(u_1^{(3)}) = x_1^{(2)}$. Similarly, for X_2 given U_2 and X_1 , define an SE such that $f_{X_2}(u_2^{(k)}, x_1^{(1)}) = x_2^{(1)}$ for $k = 3, 4, 5$ and $f_{X_2}(u_2^{(k)}, x_1^{(2)}) = x_2^{(1)}$ for $k = 3, 5$, while the other values of k are giving $x_2^{(2)}$. The causal diagram \mathcal{G}_M corresponding to this SCM is depicted in Figure 1.a. PSCM (M, P) based on M is obtained by also providing a PMF $P(U_1, U_2)$. Since this joint has to factorize into $P(U_1)P(U_2)$, as it follows from the graph \mathcal{G}_M , we know already that the SCM is Markovian. By finally taking both $P(U_1)$ and $P(U_2)$ uniform, we eventually get the PMF $P(X_1, X_2)$ such that $P(x_1^{(1)}, x_2^{(1)}) = \frac{1}{5}$, $P(x_1^{(1)}, x_2^{(2)}) = \frac{2}{15}$, and $P(x_1^{(2)}, x_2^{(1)}) = \frac{4}{15}$.

The following example shows that an SE f_X in an SCM M defines a degenerate CPT $P_M(X|\text{Pa}(X))$.

1. Tian and Pearl (2002a) have shown that any SCM can be mapped to one whose U variables are independent.




Figure 1: Ten causal diagrams.

Example 2 The SE f_{X_1} of the SCM M in Example 1 corresponds to the following CPT:

$$P_M(X_1|U_1) = \begin{bmatrix} u_1^{(1)} & u_1^{(2)} & u_1^{(3)} \\ 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{matrix} x_1^{(1)} \\ x_1^{(2)} \end{matrix}. \quad (1)$$

Similarly, $P_M(X_2|U_1, x_1^{(1)})$, i.e., the restriction of the CPT of X_2 for $X_1 = x_1^{(1)}$, is:

$$\begin{bmatrix} u_2^{(1)} & u_2^{(2)} & u_2^{(3)} & u_2^{(4)} & u_2^{(5)} \\ 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} x_2^{(1)} \\ x_2^{(2)} \end{matrix}. \quad (2)$$

From Example 2 and the joint independence of the exogenous variables, we have that a PSCM (M, P) defines a joint PMF $P_M(\mathbf{X}, \mathbf{U})$ such that $P_M(\mathbf{x}, \mathbf{u}) = \prod_{U \in \mathbf{U}} P(u) \cdot \prod_{X \in \mathbf{X}} P_M(x|\text{pa}(X))$, for each $\mathbf{x} \in \Omega_X$. This means that PSCMs are BNs whose endogenous degenerate CPTs are induced by the SEs of M , while the marginal PMFs over single exogenous variables are those in P .

Generally speaking SEs are non-injective and hence non-bijective maps. Yet, with a small abuse of notation, we denote as f_X^{-1} the map $\Omega_X \rightarrow 2^{\Omega_{\text{pa}_X}}$ returning the set of values of the parents of X corresponding to a particular value of X . E.g., in Example 1, $f_{X_1}^{-1}(x_1^{(2)}) = \{u_1^{(2)}, u_1^{(3)}\}$. For SEs of endogenous variables having both endogenous and exogenous parents, we perform such an inversion for a restriction of the SE obtained for given values of all the endogenous parents. E.g., the inverse of the SE for X_2 in Example 1 given that $X_1 = x_1^{(1)}$ is denoted as $f_{X_2|x_1^{(1)}}^{-1}$ and, following the table in Equation (2), writes as $f_{X_2|x_1^{(1)}}^{-1}(x_2^{(1)}) = \{u_2^{(3)}, u_2^{(4)}, u_2^{(5)}\}$. This notation for *restricted inverse* maps of an SE will be used also in the general case.

2.3 Interventions and Causal Effects

The basic tool for causal analysis in an SCM M is a mathematical operator called *atomic intervention* and denoted as $\text{do}(\cdot)$. Given $X \in \mathbf{X}$ and $x \in \Omega_X$, $\text{do}(X = x)$ simulates a physical action

on M forcing X to take a value $x \in \Omega_X$. Accordingly, the original SE f_X should be replaced by a constant map $X = x$. Notation M_x is used for such a modified SCM, whose causal diagram \mathcal{G}_{M_x} can be obtained by removing from \mathcal{G}_M the arcs entering X , and setting $X = x$ as an evidence. In a PSCM (M, P) , given $X, Y \in \mathbf{X}$ and $x \in \Omega_X$, we denote as $P_M(y|\text{do}(x))$ the conditional probability of $Y = y$ in the post-intervention model (M_x, P) , i.e., $P_{M_x}(y|x)$, for each $y \in \Omega_Y$. The list $\{P_M(y|\text{do}(x))\}_{x \in \Omega_X}$ is called the *causal effect* of X on $Y = y$. As interventions commute, there are no ordering issues when coping with multiple interventions. If evidence is also available, i.e., some variables have been observed, it is customary to assume that observations took place after the interventions.

3. Coping with Unidentifiability

Let us consider the problem of performing inference in a PSCM (M, P) . The exogenous variables \mathbf{U} are typically assumed to be latent, i.e., directly unobservable, while the endogenous variables \mathbf{X} are manifest. Assume for a moment that their joint PMF $\tilde{P}(\mathbf{X})$ is known. Following the discussion in Section 2.2, $\tilde{P}(\mathbf{X})$ is just the image, through the SEs of M , of the latent PMFs $\{P(U)\}_{U \in \mathbf{U}}$.

The focus of this paper is primarily on the problem of learning the (latent) latter from the (manifest) former: that is, given $\tilde{P}(\mathbf{X})$ and an SCM M over \mathbf{X} , to find out $\mathcal{K}_{M, \tilde{P}}$, namely, the collection of CSs $\{K(U)\}_{U \in \mathbf{U}}$ inducing PSCMs based on M such that:

$$\sum_{u \in \Omega_U, U \in \mathbf{U}} \left[\prod_{U \in \mathbf{U}} P(u) \cdot \prod_{X \in \mathbf{X}} P_M(x|u, \underline{\text{pa}}(X)) \right] = \tilde{P}(\mathbf{x}), \quad (3)$$

for each $\mathbf{x} \in \Omega_{\mathbf{X}}$, $(x, \underline{\text{pa}}(X)) \sim \mathbf{x}$, with $P(U) \in K(U)$ for each $U \in \mathbf{U}$. Once $\{K(U)\}_{U \in \mathbf{U}}$ have been defined, our next aim is to use them to make causal inference on a generic real-valued quantity Q that can be obtained from a PSCM (M, P) . In particular, we aim at computing the *causal bounds* of Q : $[\min_{P \in \mathcal{K}_{M, \tilde{P}}} Q(M, P), \max_{P \in \mathcal{K}_{M, \tilde{P}}} Q(M, P)]$. If this interval reduces to a point, we say that Q is *identifiable* in $\mathcal{K}_{M, \tilde{P}}$; otherwise it is called *unidentifiable*.

Note that in practice we will most likely not have the exact joint $\tilde{P}(\mathbf{X})$, but rather a sample \mathcal{D} from such a PMF; as a consequence $\tilde{P}(\mathbf{X})$ will be an approximation to the actual PFM and for this reason we shall call it ‘empirical’ in the following. We shall also assume it to be strictly positive.

3.1 Causal Bounds by Credal Networks

The collection of CSs $\mathcal{K}_{M, \tilde{P}}$ defined in the previous section provides a parametrization of PSCMs compatible with Equation (3). We call $\mathcal{K}_{M, \tilde{P}}$ the *identification* through M of \tilde{P} . It is worth noticing that, as well as a PSCM corresponds to a BN, the PSCMs induced by $\mathcal{K}_{M, \tilde{P}}$ correspond to a CN, this allowing to address causal bound computation by standard inference algorithms for CNs. An example is here below.

Example 3 Consider an SCM M made of a single endogenous variable X and a single exogenous variable U . Given an empirical PMF $\tilde{P}(X)$, consider a CN over \mathcal{G}_M such that the CCPT of X is the (degenerate) CPT $P_M(X|U)$ induced by f_X , and the CCPT of U is the CS $K(U)$ induced by the linear constraints on $P(U)$:

$$\sum_{u \in f_X^{-1}(x)} P(u) = \tilde{P}(x), \quad (4)$$

for each $x \in \Omega_X$. $K(U)$ coincides with the identification $\mathcal{K}_{M, \tilde{P}}$, i.e., every quantification of $P(U)$ consistent with Equation (4) solves Equation (3) and vice versa. To check that, let us compute the marginal probabilities over X in a PSCM (M, P) with $P \in \mathcal{K}_{M, \tilde{P}}$, i.e.,

$$P(x) = \sum_{u \in \Omega_U} P_M(x|u) \cdot P(u) = \sum_{u \in \Omega_U} \mathbb{I}[f_X(u) = x] \cdot P(u), \quad (5)$$

where the first step is by total probability theorem, the second follows from the fact that $P_M(X|U)$ is a degenerate CPT, and $\mathbb{I}[\cdot]$ is an Iverson bracket giving one if its argument is true and zero otherwise. As the rightmost-hand side of Equation (5) coincides with the left-hand side of Equation (4), we have the one-to-one mapping between the elements of $\mathcal{K}_{M, \tilde{P}}$ and those of $K(U)$. For a numerical example, consider a ternary U and a binary X whose SE coincides with f_{X_1} in Example 1. For $\tilde{P}(x^{(1)}) = \frac{1}{3}$, Equation (4) gives $P(u^{(1)}) = \frac{1}{3}$ and $P(u^{(2)}) + P(u^{(3)}) = \frac{2}{3}$, i.e., any $P(U) := [\frac{1}{3}, p, \frac{2}{3} - p]$ with $p \in [0, \frac{2}{3}]$, and no other one, is consistent with \tilde{P} .

3.2 Markovian Case

The procedure in Example 3 can be extended to any SCM under the Markovian assumption. For an endogenous variable X whose endogenous U is the unique parent, the procedure is exactly the same, and the constraints on the marginal probabilities of U are like those in Equation (4). If X has other, endogenous, parents besides U , more analysis is required to write the analogous of Equation (4). In this case, the SE of X has form $x = f_X(u, \underline{\text{pa}}(X))$. Each $\underline{\text{pa}}(X) \in \Omega_{\underline{\text{pa}}(X)}$ induces a restricted inverse of f_X and hence, for each $x \in \Omega_X$, the analogous of Equation (4) becomes:

$$\sum_{u \in f_X^{-1}[\underline{\text{pa}}(X)](x)} P(u) = \tilde{P}(x|\underline{\text{pa}}(X)), \quad (6)$$

where the conditional probabilities on the right-hand side are obtained from the empirical PMF $\tilde{P}(X)$. These are constraints on the elements of PMF $P(U)$ that can be applied separately, for each $U \in \mathbf{U}$, because of the Markovian assumption. The procedure is detailed by Algorithm 1.

Algorithm 1 Given an SCM M and a PMF $\tilde{P}(X)$, return CSs $\{K(U)\}_{U \in \mathbf{U}}$

```

1: for  $X \in \mathbf{X}$  do
2:    $U \leftarrow \text{Pa}(X) \cap \mathbf{U}$  //  $U$  as the unique exogenous parent of  $X$ 
3:    $\underline{\text{Pa}}(X) \leftarrow \text{Pa}(X) \setminus \{U\}$  // Endogenous parents of  $X$ 
4:   if  $\underline{\text{Pa}}(X) = \emptyset$  then
5:      $K(U) \leftarrow \{P'(U) : \sum_{u \in f_X^{-1}} P'(u) = \tilde{P}(x), \forall x \in \Omega_X\}$  // Eq. (4)
6:   else
7:      $K(U) \leftarrow \{P'(U) : \sum_{u \in f_X^{-1}[\underline{\text{pa}}(X)](x)} P'(u) = \tilde{P}(x|\underline{\text{pa}}(X)), \forall x \in \Omega_X, \forall \underline{\text{pa}}(X) \in \Omega_{\underline{\text{pa}}(X)}\}$  // Eq. (6)
8:   end if
9: end for

```

Overall, this allows to compute the identification through M of \tilde{P} in the Markovian case.

Theorem 1 In the Markovian case, the output $\{K(U)\}_{U \in \mathbf{U}}$ of Algorithm 1 coincides with $\mathcal{K}_{M, \tilde{P}}$.

Example 4 Consider the Markovian SCM M in Example 1 whose PMF $P(X_1, X_2)$ is used as empirical PMF $\tilde{P}(X_1, X_2)$. We obtain $\tilde{P}(x_1^{(1)}) = \frac{1}{3}$, $\tilde{P}(x_2^{(1)}|x_1^{(1)}) = \frac{3}{5}$, and $\tilde{P}(x_2^{(1)}|x_1^{(2)}) = \frac{2}{5}$. In this setup Algorithm 1 returns a CS $K(U_1)$ equal to $K(U)$ in Example 3, while Equation (6) for U_2 gives:

$$\begin{aligned} P(u_2^{(3)}) + P(u_2^{(4)}) + P(u_2^{(5)}) &= P(x_2^{(1)}|x_1^{(1)}) = 0.6, \\ P(u_2^{(3)}) + P(u_2^{(5)}) &= P(x_2^{(1)}|x_1^{(2)}) = 0.4, \\ P(u_2^{(1)}) + P(u_2^{(2)}) &= P(x_2^{(2)}|x_1^{(1)}) = 0.4, \\ P(u_2^{(1)}) + P(u_2^{(2)}) + P(u_2^{(4)}) &= P(x_2^{(2)}|x_1^{(2)}) = 0.6. \end{aligned} \quad (7)$$

This defines a CS $K(U_2)$ equivalent to any convex combination of the PMFs:

$$P_1(U_2) = \begin{bmatrix} 0.0 \\ 0.4 \\ 0.4 \\ 0.2 \\ 0.0 \end{bmatrix}, P_2(U_2) = \begin{bmatrix} 0.4 \\ 0.0 \\ 0.4 \\ 0.2 \\ 0.0 \end{bmatrix}, P_3(U_2) = \begin{bmatrix} 0.0 \\ 0.4 \\ 0.0 \\ 0.2 \\ 0.4 \end{bmatrix}, P_4(U_2) = \begin{bmatrix} 0.4 \\ 0.0 \\ 0.0 \\ 0.2 \\ 0.4 \end{bmatrix}. \quad (8)$$

As expected the uniform PMF over U_1 in Example 1 is included in $K(U_1)$ and the same happens for U_2 and $K(U_2)$.

4. Beyond Markovianity

We extend the tools of the previous section to non-Markovian models starting from an example.

Example 5 Consider an SCM M over two binary endogenous variables X_1 and X_2 whose common exogenous parent U has five states. The SE for X_1 is such that $f_{X_1}(u^{(k)}) = x_1^{(1)}$ for $k = 1, 4, 5$ and $x_1^{(2)}$ otherwise. For X_2 , we have instead $f_{X_2}(u^{(k)}, x_1^{(1)}) = x_2^{(1)}$ for $k = 1, 3$, $f_{X_2}(u^{(k)}, x_1^{(1)}) = x_2^{(2)}$ for $k = 2, 4, 5$, $f_{X_2}(u^{(k)}, x_1^{(2)}) = x_2^{(1)}$ for $k = 3$, and $f_{X_2}(u^{(k)}, x_1^{(2)}) = x_2^{(2)}$ for $k = 1, 2, 4, 5$. The causal diagram of M is the one in Figure 1.b. A (non-Markovian) PSCM based on M would be obtained by any specification of PMF $P(U)$.

In the non-Markovian case, the common exogenous parents of two or more endogenous variables are called *confounders*. Confounders express non-Markovianity also at the SCM level, being input variables common to two or more SEs. In these cases, the surjectivity we assume for single SEs is extended to the joint SE involving all the SEs with the same confounder in input.

Example 6 Let $\tilde{P}(X_1, X_2)$ denote the empirical PMF of the non-Markovian PSCM (M, P) in Example 5, whose causal diagram is in Figure 1.b. In this case, Equation (3) rewrites as:

$$\sum_{u \in \Omega_U} P(u) P_M(x_1|u) P_M(x_2|u) = \tilde{P}(x_1, x_2), \quad (9)$$

to be considered for each $x_1 \in \Omega_{X_1}$ and $x_2 \in \Omega_{X_2}$. In the sum on the left-hand side, the values of u that are not simultaneously consistent, through SEs f_{X_1} and f_{X_2} , with both x_1 and x_2 , are zero. We therefore rewrite Equation (9) as:

$$\sum_{u \in f_{X_1}^{-1}(x_1) \cap f_{X_2}^{-1}(x_2)} P(u) = \tilde{P}(x_1, x_2). \quad (10)$$

As a numerical example consider the empirical PMF $\tilde{P}(X_1, X_2)$ in Example 1. The corresponding constraints for U are: $P(u^{(1)}) = \frac{1}{5}$, $P(u^{(2)}) = \frac{2}{5}$ and $P(u^{(3)}) = \frac{4}{15}$, and $P(u^{(4)}) + P(u^{(5)}) = \frac{2}{15}$. The elements of the corresponding CS $K(U)$ are therefore $P(U) = [\frac{1}{5}, \frac{2}{5}, \frac{2}{15}, t, \frac{1}{15} - t]$ with $t \in [0, \frac{1}{15}]$.

A procedure analogous to that in Example 6 can be used for the identification of \tilde{P} for non-Markovian SCMs. As in the previous section, the key point is that the constraints on the marginal PMF of an exogenous variable $U \in \mathbf{U}$ imposed by the consistency with the empirical PMF can be specified separately from those of the other exogenous variables. We call *quasi-Markovian* a PSCM (M, P) such that each $X \in \mathbf{X}$ has only a single $U \in \mathbf{U}$ as parent.² In other words, in a quasi-Markovian model, if X is a child of U , we have $\text{Pa}(X) = (U, \underline{\text{Pa}}(X))$ where, despite the possible non-Markovianity, $\underline{\text{Pa}}(X) \subset \mathbf{X}$. This, together with the joint independence of the PMFs $\{P(U)\}_{U \in \mathbf{U}}$, induces in quasi-Markovian PSCMs the factorization:

$$P(\mathbf{x}, \mathbf{u}) = \prod_{U \in \mathbf{U}} \left[P(u) \prod_{X \in \mathbf{X}: U \in \text{Pa}(X)} P(x | \underline{\text{pa}}(X), u) \right], \quad (11)$$

for each $\mathbf{x} \in \Omega_{\mathbf{X}}$ and $\mathbf{u} \in \Omega_{\mathbf{U}}$. It is easy to see that, because of the quasi-Markovianity, in the product over \mathbf{U} in the right-hand side of Equation (11), the states of each $U \in \mathbf{U}$ appear only in the corresponding factor. This allows to define a procedure analogous to Algorithm 1 to derive the CSs $\{K(U)\}_{U \in \mathbf{U}}$ and hence obtain a CN from a quasi-Markovian model.³

Algorithm 2 Given an SCM M and a PMF $\tilde{P}(X)$, return CSs $\{K(U)\}_{U \in \mathbf{U}}$

```

1: for  $U \in \mathbf{U}$  do
2:    $\{X_U^k\}_{k=1}^{n_U} \leftarrow \text{Sort}\{X \in \mathbf{X} : U \in \text{Pa}(X)\}$  // Children of  $U$  in topological order
3:    $\gamma \leftarrow \emptyset$ 
4:   for  $(x_U^1, \dots, x_U^{n_U}) \in \times_{k=1}^{n_U} \Omega_{X_U^k}$  do
5:     for  $(\underline{\text{pa}}(X_U^1), \dots, \underline{\text{pa}}(X_U^{n_U})) \in \times_{k=1}^{n_U} \Omega_{\underline{\text{Pa}}(X_U^k)}$  do
6:        $\Omega'_U \leftarrow \bigcap_{k=1}^{n_U} f_{X_U^k | \underline{\text{pa}}(X_U^k)}^{-1}(x_U^k)$ 
7:        $\gamma \leftarrow \gamma \cup \left\{ \sum_{u \in \Omega'_U} P(u) = \prod_{k=1}^{n_U} \tilde{P}(x_U^k | x_U^1, \dots, x_U^{k-1}, \underline{\text{pa}}(X_U^1)), \dots, \underline{\text{pa}}(X_U^k) \right\}$ 
8:     end for
9:   end for
10:   $K(U) \leftarrow \{P(U) : \gamma\}$  // CS by linear constraints on  $P(U)$ 
11: end for

```

Algorithm 2 computes the CN representation of a quasi-Markovian SCM. It is a simple exercise to check that the algorithm coincides with Algorithm 1 in the case of Markovian inputs (as $n_U = 1$ for each $U \in \mathbf{U}$). In fact, Algorithm 2 can be regarded as a (partial) extension to the non-Markovian case of Algorithm 1 as shown by the following result, analogous to Theorem 1.

Theorem 2 *In quasi-Markovian cases, the output $\{K(U)\}_{U \in \mathbf{U}}$ of Algorithm 2 coincides with $\mathcal{K}_{M, \tilde{P}}$.*

-
- 2. In principle any semi-Markovian model can be turned into quasi-Markovian, e.g., by clustering all U variables into a single one; yet this neglects the exponential blowup in the computation that follows as a consequence.
 - 3. Algorithm 2 and Theorem 2 are tightly related to the notion of *confounded component* by Tian and Pearl (2002b).

Regarding complexity, unlike Algorithm 1, where the size of the CSs constraints of the CN are the same as the SEs in the SCM, the bottleneck of Algorithm 2 are the two nested loops (lines 4–9). This roughly corresponds to a complexity exponential with respect to the maximum number of children of the exogenous variables. Setting a bound to this number (e.g., $n_U \leq 2$, meaning that confounders only act on pairs of endogenous variables), would make the procedure polynomial for quasi-Markovian models.

5. Causal Analysis by Credal Networks

In the previous sections we proved that CNs exactly represent the uncertainty about the exogenous variables of an SCM induced by an empirical endogenous PMF $\tilde{P}(X)$. Such an abstract result allows us in practice to bound causal inference. Given a (quasi-)Markovian SCM M and an empirical PMF $\tilde{P}(X)$, we first obtain its CN representation by Algorithms 1 or 2. Post-interventional queries in M are then addressed via the CN, since it preserves the separation properties of M ; the degenerate (C)CPTs of the SEs involved in the intervention are replaced by constants (*surgery*), and the queries are then obtained by standard CN algorithms. Let us start from the identifiable case, where post-intervention queries are reduced to pre-intervention ones.

Example 7 (Backdoor Identification (Pearl, 2009)) *In the non-Markovian SCM in Figure 1.d, consider the query $P(x_3|\text{do}(x_1))$. The intervention on X_1 requires the removal of the arc from U towards X_1 , and hence:*

$$P(x_3|\text{do}(x_1)) := \sum_{x_2, u, u_3} P(x_3|x_1, x_2, u_3)P(u_3)P(x_2|u)P(u) = \sum_{x_2} P(x_3|x_1, x_2)P(x_2). \quad (12)$$

In the CN obtained for such quasi-Markovian model, the task becomes the computation of the lower and upper bounds of the second term in Equation (12) when $P(U) \in K(U)$ and $P(U_2) \in K(U_2)$, where the CSs are obtained by Algorithm 2. As the CN representation is equivalent, the CN inference gives equal upper and lower bounds, which expresses query identifiability.

For identifiable tasks, our approach offers then a numeric alternative to analytical approaches such as *do calculus* (Pearl, 1995). Advantages become more evident in non-identifiable cases:

Example 8 (Unidentifiable Clinical Trials (Balke and Pearl, 1994b)) *In the SCM in Figure 1.c, the endogenous variables are binary. As f_{X_1} implements $U_1 = X_1$, we do not explicitly show U_1 . Regarding U , $|\Omega_U| = 16$ and the SEs f_{X_2} and f_{X_3} are as in the original example (not reported here for lack of space). The empirical PMF is such that $\tilde{P}(x_1^{(2)}) = 0.1$, $\tilde{P}(X_2, X_3|x_1^{(1)}) = [0.32, 0.32, 0.04, 0.32]$, and $\tilde{P}(X_2, X_3|x_1^{(2)}) = [0.02, 0.17, 0.67, 0.14]$. Algorithm 2 gives the constraints for $K(U)$ and the CN computes the causal bounds $P(x_3^{(2)}|\text{do}(x_2^{(1)})) \in [0.45, 0.46]$ and $P(x_3^{(2)}|\text{do}(x_2^{(2)})) \in [0.67, 0.68]$. The corresponding difference between the two effects is the interval $[-0.23, -0.21]$, which is narrower than $[-0.23, -0.15]$, the interval obtained by Balke and Pearl (1994b) with linear programming.*

Our procedure should be therefore regarded as a numerical counterpart of the symbolic approach of Balke and Pearl (1994b), recently extended to more general cases by Sachs et al. (2020). Apart from their restriction to binary variables only, the linear programming reduction in those approaches relaxes some of the constraints on the exogenous PMFs, thus eventually yielding

an approximate characterization of the bounds. In practice, as seen in the above example, the bounds provided by these programs will be an outer approximation of the exact ones we provide.

Our CN equivalence results can be also applied to counterfactuals, namely queries that represent more than one ‘world’ at the same time, the actual one and other, alternative worlds. Balke and Pearl (1994a) showed how to compute these queries in PSCMs by *twin nets* (more generally called ‘counterfactual graphs’). In these models, each endogenous variable has a replica, sharing the same exogenous parents and the SEs. The CSs returned by Algorithms 1 or 2 can be also used in twin nets, thus allowing us to bound such queries on the basis of empirical data.

Example 9 (What-If At the Party (Balke and Pearl, 1994a)) *In the Markovian SCM of Figure 1.f, the binary endogenous variables (X_1, X_2, X_3, X_4) have exogenous counterparts with cardinalities: $|\Omega_{U_1}| = 2$, $|\Omega_{U_2}| = |\Omega_{U_3}| = 4$, and $|\Omega_{U_4}| = 3$. We refer to the original paper for the specification of the SEs and $\{P(U_k)\}_{k=1}^4$. For the corresponding PSCM, in that paper, the twin net is built as in Figure 1.g and used to compute the counterfactual $P(x_4^{(2)} | x_3^{(1)}) = 0.79$, modelling the hypothetical effect on X_4 of an intervention on X_3 forcing a state different from the observed one. We address the same task without exploiting $\{P(U_k)\}_{k=1}^4$, using only the empirical PMF. Algorithm 1 can compute CSs $\{K(U_k)\}_{k=1}^4$ in the original model, and use them to make the twin net as a CN. In spite of the looser information, the above counterfactual in the twin CN takes exactly the same sharp value.*

Remark 1 (Counterfactuals without structural equations) *The same can be done even if no explicit information about the SEs is available (this follows straightforwardly from ideas by, e.g., Balke and Pearl 1994b, Section 3.1). Consider a SE $x = f_X(x', u)$. The number of degenerate CPTs $P(X|X')$ modelling a deterministic relation between X' and X is $k := |\Omega_X|^{|\Omega_{X'}|}$. Accordingly, we set $\Omega_U := k$ and let $P(X|X', u^{(k)})$ correspond to the k -th degenerate CPT $P(X|X')$. Modelling ignorance about SE f_{X_4} in Example 9 requires for instance $|\Omega_{U_4}| := |\Omega_{X_4}|^{|\Omega_{X_2}| \cdot |\Omega_{X_3}|} = 16$. By letting U_4 enumerate all the functional relations between (X_2, X_3) and X_4 and keeping the same empirical PMF, we still obtain the same counterfactual value. To make the query unidentifiable, we add a 0.01 cut-off to the degenerate values of \tilde{P} ; this induces the bounds $P(x_4^{(2)} | x_3^{(1)}) \in [0.75, 0.85]$.*

Many other tools for causal modelling and analysis can easily be embedded in our CN formalism. This is for instance the case of *measurement bias* considered by Pearl (2010). A likelihood $\{P(z|U)\}_{z \in \Omega_Z}$ modelling a noisy observation of an exogenous variable U might be available under the assumption of conditional independence between Z and the other variables given U . This can be modelled by setting Z as a binary child of U (e.g., see Figure 1.e) in the CN, and using the likelihood for the specification of its CPT exactly as in Pearl’s *virtual evidence* method.

Expert judgements about U (e.g., a comparative statement such as $P(u^{(i)}) > P(u^{(j)})$) can be trivially embedded in our model as additional constraints in line 10 of Algorithm 2. Finally note also that, exactly as we perform standard surgery in the CN by exploiting the fact that the CN maintains the same separation properties of the original model, also other state-of-the-art techniques such as the *non-atomic interventions* proposed by Correa and Bareinboim (2020) could be supported, thus providing a numerical alternative to their σ -calculus.

6. Numerical Tests

For validation, we use quasi-Markovian SCMs of increasing size and different topologies (of the endogenous restriction of the causal diagram): *trees* (Figure 1.h), *polytrees* (Figure 1.i), *multi-*

ply connected (Figure 1.j). We call *length* of the model $l := |\mathbf{X}|$ the number of endogenous variables in M . We assume *stationarity*, i.e., the same SEs are associated with the same variables with different indexes. A random $P(\mathbf{U})$ is sampled, and the corresponding empirical PMF $\hat{P}(\mathbf{X})$ obtained. The latter is the input of Algorithm 2. Endogenous and exogenous variables have resp. two and six states, while SEs are randomly generated. As unidentifiable queries, we consider $P(X_{l/2} | \text{do}(x_1^{(0)}), x_l^{(0)})$ for trees, $P(X_{l/4} | \text{do}(x_1^{(0)}), x_{l-1}^{(0)})$ for polytrees, and $P(X_{l/4} | \text{do}(x_1^{(0)}), x_l^{(0)})$ for multiply connected models. Descriptors reported here are averages of 100 iterations.

A credal version of *variable elimination* (CVE) is used to compute exact inferences in small networks. Figure 2 (right) depicts the average size of the causal bounds computed with this method. Interestingly, relatively small (< 0.1), and hence informative, interval sizes are obtained. For larger models we use the *ApproxLP* algorithm for approximate inference in general CNs (Antonucci et al., 2015). Figure 2 (three plots on the left) shows the average execution times of the two methods for different values of l . CVE cannot handle large models due to time (a timeout of five minutes is set) and space limits. For the sake of readability, the x-axis of these plots ends at $l = 20$, but ApproxLP allows to query larger models: with a one-minute timeout, the length of the largest model that can be computed is $l = 141$ for trees, $l = 92$ for polytrees and $l = 38$ for multiply connected models. We also compared the exact intervals returned by CVE with the approximation obtained with ApproxLP. Notably, the average RMSE is low, being 0.61% for trees, 0.03% for polytrees and 0.026% for multiply connected models. The tests have been performed by means of a Java library implementing all the techniques discussed in the paper.⁴




Figure 2: Execution times (left) and bounds size (right) for experiments.

7. Conclusions

The present work proposes a fully automatic, numeric, approach to causal inference that is alternative to analytical avenues such as do- or σ -calculus and their specializations, and that is natively capable to solve unidentifiable problems too. It shows that we can take data and an incompletely specified SCM, and turn them, exactly, into a credal network. Standard algorithms for credal nets will then deliver us causal inference up to the top level of *Pearl's causal hierarchy* (Bareinboim et al., to appear): counterfactuals. This holds true even under the extreme condition where the SCM is provided as a bare causal graph, with no information at all about the structural equations—including the cardinality of the U variables (see Remark 1). Our algorithms to efficiently convert SCMs into credal nets apply to a wide class of models, which we call quasi-Markovian. It seems possible that these ideas can be extended to more general models and also to the continuous case. The relation with credal networks is, however, already general. In this sense, an SCM is (solvable by) a credal network.

4. See github.com/idsia/credici and github.com/IDSIA-papers/2020-PGM-structural.

References

- A. Antonucci, C. P. de Campos, D. Huber, and M. Zaffalon. Approximate credal network updating by linear programming with applications to decision making. *International Journal of Approximate Reasoning*, 58:25–38, 2015.
- N. Asghar. Automatic extraction of causal relations from natural language texts: a comprehensive survey. *CoRR*, abs/1605.07895, 2016. URL <http://arxiv.org/abs/1605.07895>.
- A. Balke and J. Pearl. Probabilistic evaluation of counterfactual queries. In B. Hayes-Roth and R. E. Korf, editors, *Proceedings of AAAI/IAAI 1994*, pages 230–237. AAAI Press/The MIT Press, 1994a.
- A. Balke and J. Pearl. Counterfactual probabilities: Computational methods, bounds and applications. In R. L. de Mántaras and D. Poole, editors, *Proceedings of UAI '94*, pages 46–54. Morgan Kaufmann, 1994b.
- A. Balke and J. Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.
- E. Bareinboim, J. D. Correa, D. Ibeling, and T. Icard. On Pearl’s hierarchy and the foundations of causal inference. In *ACM Special Volume in Honor of Judea Pearl (provisional title)*. ACM, to appear.
- J. D. Correa and E. Bareinboim. A calculus for stochastic interventions: causal effect identification and surrogate experiments. In *Proceedings of AAAI 2020*, pages 10093–10100. AAAI Press, 2020.
- F. G. Cozman. Credal networks. *Artificial intelligence*, 120(2):199–233, 2000.
- E. Fagioli and M. Zaffalon. 2U: An exact interval propagation algorithm for polytrees with binary variables. *Artificial Intelligence*, 106(1):77–107, 1998.
- D. Hume. *A Treatise of Human Nature*. Oxford University Press, 1739.
- C. Kang and J. Tian. Inequality constraints in causal models with hidden variables. In *Proceedings of UAI '06*. AUAI Press, 2006.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT, 2009.
- D. D. Mauá, C. P. De Campos, A. Benavoli, and A. Antonucci. Probabilistic inference in credal networks: new complexity results. *Journal of Artificial Intelligence Research*, 50:603–637, 2014.
- J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- J. Pearl. *Causality*. Cambridge University Press, 2009.
- J. Pearl. On measurement bias in causal inference. In P. Grünwald and P. Spirtes, editors, *Proceedings of UAI 2010*, pages 425–432. AUAI Press, 2010.
- M. C. Sachs, E. E. Gabriel, and A. Sölander. Symbolic computation of tight causal bounds. *CoRR*, abs/2003.10702, 2020. URL <http://arxiv.org/abs/2003.10702>.
- R. D. Shachter. Evaluating influence diagrams. *Operations research*, 34(6):871–882, 1986.
- J. Tian and J. Pearl. On the testable implications of causal models with hidden variables. In A. Darwiche and N. Friedman, editors, *Proceedings of UAI '02*, pages 519–527. Morgan Kaufmann, 2002a.
- J. Tian and J. Pearl. A general identification condition for causal effects. In R. Dechter, M. J. Kearns, and R. S. Sutton, editors, *Proceedings of AAAI/IAAI 2002*, pages 567–573. AAAI Press/The MIT Press, 2002b.
- D. E. Wilkins. *Practical Planning: Extending the Classical AI Planning Paradigm*. Elsevier, 2014.

Appendix A. Proofs

Proof [of Theorem 1] Let us first note that in a Markovian PSCM (M, P) , for each $X \in \mathbf{X}$, $x \in \Omega_X$ and $\underline{\text{pa}}(X) \in \Omega_{\underline{\text{pa}}(X)}$, we have:

$$P(x|\underline{\text{pa}}(X)) = \sum_{u \in \Omega_U} P(x|u, \underline{\text{pa}}(X)) \cdot P(u|\underline{\text{pa}}(X)) = \sum_{u \in \Omega_U} P(x|u, \underline{\text{pa}}(X)) \cdot P(u), \quad (13)$$

where $U \in \mathbf{U}$ is the, unique because of Markovianity, exogenous parent of X and, as usual, $\underline{\text{pa}}(X)$ are the other, endogenous, parents of X . The first derivation in Equation (13) follows from total probability theorem, while the second is because of the d -separation between U , which is a root node of \mathcal{G}_M , and the other parents of its child X (Koller and Friedman, 2009).

To prove the theorem, first check the inclusion $\mathcal{K}_{M, \tilde{P}} \subseteq \{K(U)\}_{U \in \mathbf{U}}$. Take $\{P(U)\}_{U \in \mathbf{U}} \in \mathcal{K}_{M, \tilde{P}}$. The corresponding PSCM, based on M , satisfies Equation (13) with $P(x|\underline{\text{pa}}(X)) = \tilde{P}(x|\underline{\text{pa}}(X))$ because of the definition of $\mathcal{K}_{M, \tilde{P}}$. Thus:

$$\sum_{u \in \Omega_U} P(x|u, \underline{\text{pa}}(X)) \cdot P(u) = \tilde{P}(x|\underline{\text{pa}}(X)). \quad (14)$$

Conditional probabilities $P(x|u, \underline{\text{pa}}(X))$ in Equation (14) are from a (degenerate) CPT based on the SE $f_X \in M$. We consequently rewrite the equation as:

$$\sum_{u \in \Omega_U} \mathbb{1}_{f_X(u, \underline{\text{pa}}(X)) = x} \cdot P(u) = \tilde{P}(x|\underline{\text{pa}}(X)). \quad (15)$$

Equation (15) corresponds to Equation (4) (i.e., line 5 of Algorithm 1) when $\underline{\text{pa}}(X)$ is empty and Equation (6) (i.e., line 7 of Algorithm 1) otherwise. This proves $P(U) \in K(U)$ for each $U \in \mathbf{U}$ and hence $\mathcal{K}_{M, \tilde{P}} \subseteq \{K(U)\}_{U \in \mathbf{U}}$. Vice versa, to prove $\{K(U)\}_{U \in \mathbf{U}} \subseteq \mathcal{K}_{M, \tilde{P}}$ let us take a PMF $P(U) \in K(U)$ for each $U \in \mathbf{U}$. This induces a (Markovian) PSCM based on M that should satisfy Equation (13), but also Equation (14) because of lines 5 and 7 of Algorithm 1. But this proves the consistency of P with \tilde{P} , hence $\{P(U)\}_{U \in \mathbf{U}} \subseteq \mathcal{K}_{M, \tilde{P}}$ and finally the thesis. \blacksquare

Lemma 3 *In a quasi-Markovian PSCM, for each $U \in \mathbf{U}$, let $\{X_U^k\}_{k=1}^{n_U}$ denote the children of U , with the index k sorting them according to a topological order. For each $x_U^k \in \Omega_{X_U^k}$ and $\underline{\text{pa}}(X_U^k) \in \Omega_{\underline{\text{pa}}(X_U^k)}$, with $k = 1, \dots, n_U$, we have:*

$$\sum_{u \in \Omega_U} P(u) \prod_{k=1}^{n_U} P(x_U^k | \underline{\text{pa}}(X_U^k), u) = \prod_{k=1}^{n_U} P(x_U^k | x_U^1, \dots, x_U^{k-1}, \underline{\text{pa}}(X_U^1), \dots, \underline{\text{pa}}(X_U^k)). \quad (16)$$

Proof Equation (11) gives a factorization of the joint PMF of the quasi-Markovian PSCM corresponding to that of a BN based on the directed acyclic graph \mathcal{G}_M . Arc reversal (Shachter, 1986) allows to express the joint PMF of a BN as that of a second BN based on a different directed graph where the orientation of one or more arcs is changed. In our case we want to reverse the arcs from U to X_U^k , for each $k = 1, \dots, n_U$. To preserve the joint PMF after the reversal of $U \rightarrow X_U^k$ we should: (i) add the parents of U to the parents of X_U^k , if not already present, and, (ii) add the parents of X_U^k to the parents of U , again, if not already present. Although in principle these operation might add cycles to the directed graph, as stated in the proof of Theorem 3 of Shachter (1986), this is not the case if we follow a topological order like the one associated to the index k .

To describe the changes induced by these operations to \mathcal{G}_M , let us set \mathcal{G}^0 equal to \mathcal{G}_M , but without the arcs we want to remove and, for each $k = 1, \dots, n_U$, denote as \mathcal{G}^k the causal diagram after the k -th reversal, performed together with the above parent augmentations. We similarly denote as $\text{Pa}_k(Z)$ the parents of a generic variable Z in \mathcal{G}^k . E.g., by our definition of \mathcal{G}_0 , $\text{Pa}_0(X_U^k) = \text{Pa}(X_U^k)$, for each $k = 1, \dots, n_U$. It is easy to check that, after the first reversal: $\text{Pa}_1(U) := X_U^1 \cup \text{Pa}_0(X_U^1)$, and as $\text{Pa}_0(U) = \emptyset$ (i.e., U is originally a root), $\text{Pa}_1(X_U^1) := \text{Pa}_0(X_U^1)$. All the other sets of parents are unchanged. After k reversals we have instead:

$$\text{Pa}_k(U) := \bigcup_{i=1}^k \left[X_U^i \cup \text{Pa}_0(X_U^i) \right], \quad (17)$$

$$\text{Pa}_k(X_U^k) := \left[\bigcup_{i=1}^{k-1} X_U^i \right] \cup \left[\bigcup_{i=1}^k \text{Pa}_0(X_U^i) \right]. \quad (18)$$

The validity of these two equations for each $k = 1, \dots, n_U$ can be proved by induction. First note that they are satisfied for $k = 1$. After that, assume Equations (17) and (18) valid for $k = j - 1$ and prove that these also hold for $k = j$. After the j -th reversal, the parents of U can be expressed as:

$$\text{Pa}_j(U) = \text{Pa}_{j-1}(U) \cup X_U^j \cup \text{Pa}_{j-1}(X_U^j) = \bigcup_{i=1}^{j-1} \left[X_U^i \cup \text{Pa}_0(X_U^i) \right] \cup X_U^j \cup \text{Pa}_0(X_U^j), \quad (19)$$

where in the first derivation we added to the parents of U , as they are before the reversal, X_U^j and its parents (again before the reversal). The second step follows from the validity of the Equation (17) for $k = j - 1$ and from the general fact $\text{Pa}_l(X_U^j) = \text{Pa}_0(X_U^j)$ for $l < j$ (i.e., the parents of X_U^j remain unchanged before the j -th reversal). It is easy to check that the rightmost-hand side of Equation (19) corresponds to the right-hand side of Equation (17) for $k = j$. Similarly, for X_U^j :

$$\text{Pa}_j(X_U^j) = \text{Pa}_{j-1}(X_U^j) \cup \text{Pa}_{j-1}(U) = \text{Pa}_0(X_U^j) \cup \bigcup_{i=1}^{j-1} \left[X_U^i \cup \text{Pa}_0(X_U^i) \right], \quad (20)$$

where in the first derivation we added to the parents of X_U^j , as they are before the reversal, the parents of U (again before the reversal). The second step follows from the above discussed fact that the parents of X_U^j are unchanged before the j -th reversal and the validity of Equation (18) for $k = j - 1$. It is easy to check that the rightmost-hand side of Equation (20) corresponds to the right-hand side of Equation (18) for $k = j$. Overall, we have Equations (17) and (18) valid for each $k = 1, \dots, n_U$. An example of these reversals is in Figure 3.

Note that after the n_U reversals, U is a barren node as: (i) all the n_U original children of U are now its parents; (ii) for each $k = 1, \dots, n_U$, U does not belong to the new parents of X_U^k in Equation (18). After the last reversal, the joint PMF associated with the quasi-Markovian PSCM is still factorizing as a BN, but based on a different directed acyclic graph. The new CPTs are those associated with U and its children, while all the other ones remain unchanged. The parents of U are now as in Equation (17), to be considered for $k = n_U$. For each $k = 1, \dots, n_U$, the parents of X_U^k are provided instead by Equation (18) (remember that after the k -th reversal, the parents of X_U^k will not be further modified). As the new “reversed” joint PMF with the original one in Equation (11) coincide, we can simplify from both sides all the probabilities associated to the CPTs that are




Figure 3: Arc reversals in a quasi-Markovian SCM (exogenous nodes are in gray).

not associated with U and its children and obtain:

$$P(u) \prod_{k=1}^{n_U} P(x_U^k | \underline{\text{pa}}(X_U^k), u) = \quad (21)$$

$$P(u | x_U^1, \dots, x_U, \underline{\text{pa}}(X_U^1), \dots, \underline{\text{pa}}(X_U^k)) \prod_{k=1}^{n_U} P(x_U^k | x_U^1, \dots, x_U^{k-1}, \underline{\text{pa}}(X_U^1), \dots, \underline{\text{pa}}(X_U^k)).$$

Finally, Equation (16) follows from Equation (21) by taking the sum over U on both sides and noticing that, as U is a barren node, it only appears in its own CPT. \blacksquare

Proof [of Theorem 2] The proof is analogous to that of Theorem 1, with Equation (16) from Lemma 3 playing the role of Equation (13).

To prove the theorem, first check the inclusion $\mathcal{K}_{M, \tilde{P}} \subseteq \{K(U)\}_{U \in \mathcal{U}}$. Take $\{P(U)\}_{U \in \mathcal{U}} \in \mathcal{K}_{M, \tilde{P}}$. The corresponding quasi-Markovian PSCM, based on M , satisfies Equation (16) with the product of conditional probabilities in the right-hand side as in the empirical case because of the definition of $\mathcal{K}_{M, \tilde{P}}$. Thus:

$$\sum_{u \in \Omega_U} P(u) \prod_{k=1}^{n_U} P(x_U^k | \underline{\text{pa}}(X_U^k), u) = \prod_{k=1}^{n_U} \tilde{P}(x_U^k | x_U^1, \dots, x_U^{k-1}, \underline{\text{pa}}(X_U^1), \dots, \underline{\text{pa}}(X_U^k)). \quad (22)$$

The conditional probabilities in the left-hand side of Equation (22) are from (degenerate) CPTs based on the SEs of the children of U . We consequently rewrite the equation as:

$$\sum_{u \in \Omega_U} P(u) \prod_{k=1}^{n_U} \mathbb{I}[f_{X_U^k}(\underline{\text{pa}}(X_U^k), u) = x_U^k] = \prod_{k=1}^{n_U} \tilde{P}(x_U^k | x_U^1, \dots, x_U^{k-1}, \underline{\text{pa}}(X_U^1), \dots, \underline{\text{pa}}(X_U^k)). \quad (23)$$

Equation (23) corresponds to the linear constraint specification in line 7 of Algorithm 2. This proves $P(U) \in K(U)$ for each $U \in \mathbf{U}$ and hence $\mathcal{K}_{M, \tilde{P}} \subseteq \{K(U)\}_{U \in \mathbf{U}}$.

Vice versa, to prove $\{K(U)\}_{U \in \mathbf{U}} \subseteq \mathcal{K}_{M, \tilde{P}}$ let us take a PMF $P(U) \in K(U)$ for each $U \in \mathbf{U}$. This induces a quasi-Markovian PSCM based on M that should satisfy Equation (16), but also Equation (22) because of line 7 of Algorithm 2. But this proves the consistency of P with \tilde{P} , hence $\{P(U)\}_{U \in \mathbf{U}} \subseteq \mathcal{K}_{M, \tilde{P}}$ and finally the thesis. ■