
Causal Expectation-Maximisation

Marco Zaffalon
IDSIA (Switzerland)
zaffalon@idsia.ch

Alessandro Antonucci
IDSIA (Switzerland)
alessandro@idsia.ch

Rafael Cabañas
IDSIA (Switzerland)
rcabanas@idsia.ch

Abstract

Structural causal models are the basic modelling unit in Pearl’s causal theory; in principle they allow us to solve counterfactuals, which are at the top rung of the ladder of causation. But they often contain latent variables that limit their application to special settings. This appears to be a consequence of the fact, proven in this paper, that causal inference is NP-hard even in models characterised by polytree-shaped graphs. To deal with such a hardness, we introduce the causal EM algorithm. Its primary aim is to reconstruct the uncertainty about the latent variables from data about categorical manifest variables. Counterfactual inference is then addressed via standard algorithms for Bayesian networks. The result is a general method to approximately compute counterfactuals, be they identifiable or not (in which case we deliver bounds). We show empirically, as well as by deriving credible intervals, that the approximation we provide becomes accurate in a fair number of EM runs. These results lead us finally to argue that there appears to be an unnoticed limitation to the trending idea that counterfactual bounds can often be computed without knowledge of the structural equations.

1 Introduction

Structural causal models (SCMs) are central to Pearl’s theory of causal inference. An SCM can be displayed graphically as in Fig. 1b, where U and V represent *exogenous* variables. By definition, they have an associated distribution; internal, *endogenous*, nodes are instead deterministic functions of their parents. These define so-called *structural equations* (SE), which assign values to endogenous nodes. Exogenous variables are often latent in SCMs, while we usually have data about the endogenous ones. In this case we talk of a *partially specified* SCM (PSCM); if we have data also for the exogenous variables, then we say that the SCM is *fully specified* (FSCM).

FSCMs permit doing counterfactual inference, which answers hypothetical questions like, what if I had done something differently in the past? This type of inference is strictly more general than computing interventions (what if I do this?), which is strictly more general than probabilistic reasoning (what if I see this?).¹ Computing probabilities of counterfactuals is therefore of utmost importance in causal inference, and one can achieve it by an FSCM.

PSCMs allow in principle to compute those probabilities too; however, most often the result is non-identifiable, which means that those probabilities can only be known to belong to an interval. Having interval-valued outcomes is often useful, in spite of their inherent imprecision. But only few methods are available to compute non-identifiable quantities, and most of them are ad hoc. Three approaches stand out as more systematic: the first, from 2020, converts the problem into one of inference in ‘credal networks’ and uses algorithms for these to yield exact counterfactual intervals [11]. The method is efficient on Markovian SCMs and can solve a class of non-Markovian ones too, albeit with increasing complexity. The second, from 2021, converts the problem into one of polynomial optimisation and, given its complexity, eventually approximates its solution via

¹Except for sets of measure zero, see [1].

Monte Carlo sampling [12]. However it does not present an extensive empirical evaluation, so that it is unclear how accurate vs. time consuming is the approximation. The third, again from 2021, is similar to the previous one [6]. The difference is that the polynomial optimisation is used to yield exact solutions in some case. Approximations are provided for the general case; extensive empirical evaluations are missing.

In this paper we present a general method to compute approximate intervals of counterfactual probabilities. We do this after observing that the computation of counterfactual probabilities is NP-hard even in polytree topologies of PSCM; remember that in FSCMs the computation is polynomial. This shows that PSCMs are particularly complex models to handle.

The solution method we propose can be understood as a specialisation of the EM algorithm to causal problems. We show that in the case of PSCMs, the EM always converges to a global optimum: that is, a distribution for the exogenous variables that generates, through the given SCM, the available empirical probabilities; we call this a *compatible* distribution. Multiple EM runs will amount to sampling the set of compatible distributions, or, stated differently, to sampling multiple FSCMs compatible with the given PSCM. Taking the extrema of the counterfactual probabilities that those FSCMs yield, we obtain an interval that is included in the original PSCM's actual interval.

Our experiments show that 20 EM runs are enough to obtain a good approximation to the actual interval. We double-check this result by analytically developing credible bounds: these indeed show that the approximate interval is close to the actual one with 95% credibility just after 20 runs.

In the special, and especially important, case where the outcome is identifiable, 6 EM runs are enough to be that confident that it is indeed identifiable (9 runs for 99% credibility). Remember that Pearl's celebrated do-calculus permits, in principle, computing the probability of the effect of interventions, when the former is identifiable. In such a case, our EM-based algorithm can be regarded as a very viable numerical extension of do-calculus to counterfactuals—and without the need to handle the algebraic manipulations of do-calculus, since it yields solutions automatically.

Finally, we discuss an issue that does not seem to have been noticed so far: that counterfactual bounds are unwarranted if obtained in absence of a pre-defined PSCM of the problem. Stated differently, in general one cannot give bounds that hold irrespectively of the actual underlying PSCM; structural equations can hardly be neglected.

Outline of the paper: Sect. 2 introduces SCMs. Sect. 3 characterises the set of compatible distributions and defines SCM-compatibility. Sect. 4 proves that computation in PSCMs is NP-hard, presents the causal EM and proves its optimality. Sect. 5 discusses the problems we meet when we neglect structural equations. Sect. 6 presents the credible intervals and the experiments. Sect. 7 concludes the paper. Proofs are gathered in Appendix A.




Figure 1: A causal diagram over two variables (a), a corresponding (Markovian) SCM over those two (endogenous) variables (b), its counterfactual graph (c), the result of a surgery to compute PN (d), and another causal diagram over three variables (e), the corresponding (still Markovian) SCM (f), its counterfactual graph (g), and the result of a surgery to compute PNS (h).

2 Structural Causal Models

We focus on models with categorical variables.² Variable V takes values from a finite set Ω_V ; v denotes its generic value. $P(V)$ is a probability mass function (PMF) over V . Notation $P(V|V') := \{P(V|v')\}_{v' \in \Omega_{V'}}$ is used for a conditional probability table, i.e., a collection of conditional PMFs over a variable indexed by the states of another one. We define an SCM M by a directed graph \mathcal{G} whose nodes are in a one-to-one correspondence with both its endogenous variables $\mathbf{X} := (X_1, \dots, X_n)$ and the exogenous ones $\mathbf{U} := (U_1, \dots, U_m)$. We focus on semi-Markovian models, i.e., \mathcal{G} is assumed acyclic. Exogenous variables correspond to the root nodes of \mathcal{G} . A PMF $P(U)$ is specified for each $U \in \mathbf{U}$. A structural equation (SE) f_X is instead provided for each $X \in \mathbf{X}$; it is a map $f_X : \Omega_{\text{pa}_X} \rightarrow \Omega_X$, where pa_X are the parents (i.e., the immediate predecessors) of X according to \mathcal{G} . If X has a single exogenous parent with states indexing all the possible deterministic relations between the endogenous parents and X , we call the specification of f_X *conservative*. To have all the states of X possibly realised, we only consider surjective SEs. We also require a joint surjectivity, meaning that any $x \in \Omega_X$ can be realised for at least a $u \in \Omega_U$. SCM M induces the following joint PMF:

$$P(\mathbf{x}, \mathbf{u}) = \prod_{X \in \mathbf{X}} P(x|\text{pa}_X) \prod_{U \in \mathbf{U}} P(u), \quad (1)$$

for each $x \in \Omega_X$ and $u \in \Omega_U$, where the values of the conditional probability tables associated with each X are degenerate, i.e., $P(x|\text{pa}_X) = \delta_{x, f_X(\text{pa}_X)}$. These are fully specified SCMs (i.e., FSCM). If the exogenous PMFs are not provided, we say instead that M is partially specified (PSCM). In that case, information about \mathbf{X} is assumed to be available, e.g., in a data set \mathcal{D} of observations.

3 Likelihood Characterisation

The factorisation properties of the joint PMF $P(\mathbf{X}, \mathbf{U})$ of an SCM M , as shown in Eq. (1), are those of a Bayesian network (BN) defined over graph \mathcal{G} . A characterisation can be provided also for the marginal graph related to the joint $P(\mathbf{X})$ via the notion of c-component [10] (see also Ex. 1 below).

Given the graph \mathcal{G} of M , its *reduction* \mathcal{G}' is obtained by removing from \mathcal{G} any arc connecting pairs of endogenous variables (e.g., see Fig. 2b). Let $\{\mathcal{G}_c\}_{c \in \mathcal{C}}$ denote the connected components of \mathcal{G}' . The *c-components* of M are the elements of the partition $\{\mathbf{X}^{(c)}\}_{c \in \mathcal{C}}$ of \mathbf{X} , where $\mathbf{X}^{(c)}$ denotes the endogenous nodes of \mathcal{G}_c for each $c \in \mathcal{C}$. This definition is consistent with the one in [10], where the more general case of possibly non-root exogenous nodes is considered. We similarly denote as $\{\mathbf{U}^{(c)}\}_{c \in \mathcal{C}}$ the partition of \mathbf{U} induced by the connected components of \mathcal{G}' . For each $c \in \mathcal{C}$, we denote as $\mathbf{Y}^{(c)}$ the union of the endogenous parents of the nodes in $\mathbf{X}^{(c)}$ and $\mathbf{X}^{(c)}$ itself. Finally, for each $X \in \mathbf{X}^{(c)}$, $\mathbf{Y}_X^{(c)}$ is obtained by removing from $\mathbf{Y}^{(c)}$ the nodes topologically following X and X itself. The c-components induce the following endogenous factorisation [10, Chp. 4, Cor. 4]:

$$P(\mathbf{x}) = \prod_{c \in \mathcal{C}} \prod_{X \in \mathbf{X}^{(c)}} P(x|\mathbf{y}_X^{(c)}), \quad (2)$$

for each $x \in \Omega_X$, with the values of x and $\mathbf{y}_X^{(c)}$ consistent with x . An example of a graph corresponding to such a factorisation is in Fig. 2c. We can also write a factorisation with respect to the c-components in Eq. (1), that is associated with the original graph (e.g., see Fig. 2a), as follows:

$$P(\mathbf{x}, \mathbf{u}) = \prod_{c \in \mathcal{C}} \left[\prod_{X \in \mathbf{X}^{(c)}} P(x|\text{pa}_X) \prod_{U \in \mathbf{U}^{(c)}} P(u) \right]. \quad (3)$$

Example 1. Fig. 2a depicts an SCM with three c-components, say $\mathcal{C} = \{1, 2, 3\}$. As shown by Fig. 2b: $\mathbf{X}^{(1)} = \{X_1\}$, $\mathbf{X}^{(2)} = \{X_2, X_4\}$, $\mathbf{X}^{(3)} = \{X_3\}$, $\mathbf{U}^{(c)} = \{U_c\}$ for each $c \in \mathcal{C}$, $\mathbf{Y}^{(1)} = \{X_1\}$, $\mathbf{Y}_{X_1}^{(1)} = \emptyset$, $\mathbf{Y}^{(2)} = \{X_1, X_2, X_3, X_4\}$, $\mathbf{Y}_{X_2}^{(2)} = \{X_1\}$, $\mathbf{Y}_{X_4}^{(2)} = \{X_1, X_2, X_3\}$, and $\mathbf{Y}^{(3)} = \{X_2, X_3\}$ and $\mathbf{Y}_{X_3}^{(3)} = \{X_2\}$. Eq. (2) gives therefore the endogenous factorisation induced by the graph in Fig. 2c.

Notice that Markovianity [9] corresponds to condition $|\mathbf{X}^{(c)}| = |\mathbf{U}^{(c)}| = 1$ for each $c \in \mathcal{C}$. Having only $|\mathbf{U}^{(c)}| = 1$ for each $c \in \mathcal{C}$ is called instead quasi-Markovianity [11].

²Exogenous variables can be assumed categorical w.l.o.g. if endogenous ones are ([6, Prop. 1]; [12, Thm. 1]).




Figure 2: A SCM (a), its reduction (b), its endogenous factorisation (c).

Now consider an FSCM M , from which we sample a set \mathcal{D} of i.i.d. endogenous observations. As a consequence of the factorisation in Eq. (2), the log-likelihood of \mathbf{x} given \mathcal{D} is:

$$LL[\{P(x|\mathbf{y}_X^{(c)})\}] := \sum_{\mathbf{x} \in \mathcal{D}} \sum_{c \in \mathcal{C}} \sum_{X \in \mathbf{X}^{(c)}} \log P(x|\mathbf{y}_X^{(c)}) = \sum_{c \in \mathcal{C}} \sum_{X \in \mathbf{X}^{(c)}} \sum_{x, \mathbf{y}_X^{(c)}} n(x, \mathbf{y}_X^{(c)}) \cdot \log P(x|\mathbf{y}_X^{(c)}), \quad (4)$$

where the values of x and $\mathbf{y}_X^{(c)}$ are those consistent with \mathbf{x} , $n(\cdot)$ denote the frequencies in \mathcal{D} of its argument, and indexes in the argument of LL are omitted for the sake of readability. In the rightmost-hand side of Eq. (4), we see the decomposable structure of a multinomial likelihood. Such a concave function has no local maxima and a global maximum value given by

$$LL^* := \sum_{c \in \mathcal{C}} \sum_{X \in \mathbf{X}^{(c)}} \sum_{x, \mathbf{y}_X^{(c)}} n(x, \mathbf{y}_X^{(c)}) \cdot \log \hat{P}(x|\mathbf{y}_X^{(c)}), \quad (5)$$

which is achieved if and only if the probabilities attain the relative frequencies in \mathcal{D} , i.e.,

$$\hat{P}(x|\mathbf{y}_X^{(c)}) := \frac{n(x, \mathbf{y}_X^{(c)})}{n(\mathbf{y}_X^{(c)})}, \quad (6)$$

for each x and $\mathbf{y}_X^{(c)}$, $c \in \mathcal{C}$, and $X \in \mathbf{X}^{(c)}$ [7].

Note that in the limit of infinite data, the endogenous PMF we obtain from the frequencies coincides with the marginalisation of the exogenous variables from the joint PMF in Eq. (3), i.e.,

$$\prod_{c \in \mathcal{C}} \sum_{\mathbf{u}^{(c)}} \left[\prod_{X \in \mathbf{X}_C} P(x|\text{pa}_X) \prod_{U \in \mathbf{U}_c} P(u) \right] = \prod_{c \in \mathcal{C}} \prod_{X \in \mathbf{X}^{(c)}} \hat{P}(x|\mathbf{y}_X^{(c)}). \quad (7)$$

However, in real applications we shall always make use of finite samples. And the empirical distribution learnt from a finite sample may be arbitrarily far from the limit case, so as to make Eq. (7) possibly fail. In this case, the empirical distribution is incompatible with the SCM under consideration, in the sense that the SCM simply cannot generate such a distribution in the limit. For this reason, we call Eq. (7) *M-compatibility*, or compatibility of the data with a given model.

M-compatibility is particularly important when we are only given a PSCM. In this case, we aim at reconstructing the uncertainty about \mathbf{U} from the empirical PMF $P(\mathbf{X})$. But if *M-compatibility* fails, we know that the task is hopeless, because there is no $P(\mathbf{U})$ that can eventually lead to $P(\mathbf{X})$. This may happen either because the sample is too small, or because the PSCM is a wrong model of the phenomenon under study. In either case, one should refrain from making inferences using jointly the PSCM and $P(\mathbf{X})$ as they would be logically contradicting each other.

We thus have to make sure that *M-compatibility* holds before computing inferences. To this end, let us consider the marginal log-likelihood of a given PSCM. By Eq. (3) we obtain:

$$LL[\{P(U)\}_{U \in \mathbf{U}}] := \sum_{c \in \mathcal{C}} \sum_{\mathbf{y}^{(c)}} n(\mathbf{y}^{(c)}) \log \sum_{\mathbf{u}^{(c)}} \left[\prod_{X \in \mathbf{X}^{(c)}} P(x|\text{pa}_X) \prod_{U \in \mathbf{U}^{(c)}} P(u) \right], \quad (8)$$

where by definition $\mathbf{Y}^{(c)} = \cup_{X \in \mathbf{X}^{(c)}} \{X, \mathbf{Y}_X^{(c)}\}$ and hence $n(\mathbf{y}^{(c)}) = \sum_{X \in \mathbf{X}^{(c)}} n(x, \mathbf{y}_X^{(c)})$. The following result characterises the global maxima of the marginal likelihood:

Theorem 1. *Let \mathcal{K} denote the set of quantifications for $\{P(U)\}_{U \in \mathbf{U}}$ consistent with the following constraint to be satisfied for each $c \in \mathcal{C}$ and each $\mathbf{y}^{(c)}$:*

$$\sum_{\substack{\mathbf{u}^{(c)}: f_X(\text{pa}_X)=x \\ \forall X \in \mathbf{X}^{(c)}}} \prod_{U \in \mathbf{U}^{(c)}} P(u) = \prod_{X \in \mathbf{X}^{(c)}} \hat{P}(x|\mathbf{y}_X^{(c)}), \quad (9)$$

where the values of u , x and $\mathbf{y}_X^{(c)}$ are those consistent with $\mathbf{u}^{(c)}$ and $\mathbf{y}^{(c)}$. If $\mathcal{K} \neq \emptyset$, the log-likelihood in Eq. (8) achieves its global maximum, equal to LL^* , if and only if $\{P(U)\}_{U \in \mathcal{U}} \in \mathcal{K}$. If $\mathcal{K} = \emptyset$, the marginal log-likelihood in Eq. (8) can only take values strictly lower than LL^* .

In other words, global optimality of the marginal log-likelihood is tantamount to finding an FSCM compatible with the data and the given PSCM:

Corollary 1. *The function in Eq. (8) achieving the global maximum of the marginal log-likelihood and having $\{P(U)\}_{U \in \mathcal{U}} \in \mathcal{K}$ are equivalent conditions for M -compatibility.*

This means that the marginal log-likelihood at its maximum height LL^* looks like a flat plateau. Such a plateau characterises all and only the FSCMs compatible with the given PSCM. In case there are none, the maximum will be strictly lower than LL^* , which is something that we can use as an effective test of M -incompatibility.

4 EM for Causal Computations

With \mathcal{K} available, one could compute the bounds of counterfactual queries by optimising the corresponding function of $\{P(U)\}_{U \in \mathcal{U}}$ over \mathcal{K} . An explicit characterisation of \mathcal{K} has been provided for the special case of quasi-Markovian models in another paper [11]; that approach faces computational issues if the model is non-Markovian, though. In fact coping with partial identifiability is demanding even with simple topologies and queries, as we show by the following:

Theorem 2. *Computing interventional queries in polytree-shaped PSCMs is NP-hard.*

This result makes it clear that we should generally consider approximate methods to compute bounds. Cor. 1 provides a good match in this respect: it tells us that sampling the global optimum points of the marginal log-likelihood corresponds to sampling from, and hence approximating, \mathcal{K} . And the crucial observation is that we can easily sample the optimum points of the marginal log-likelihood with the *expectation-maximisation* (EM) scheme [4]: in fact, exogenous variables are missing at random in \mathcal{D} , just because they are latent (missing with probability one).

In particular, given an initialisation $\{P_0(U)\}_{U \in \mathcal{U}}$, the EM algorithm consists in regarding the posterior probability $P_0(u|x)$ as a *pseudo-count* for (u, x) , for each $x \in \mathcal{D}$, $u \in \Omega_U$ and $U \in \mathcal{U}$ (E-step). A new estimate is consequently obtained as $P_1(u) := |\mathcal{D}|^{-1} \sum_{x \in \mathcal{D}} P_0(u|x)$ (M-step). This scheme, called EMCC (*EM for Causal Computation*) and detailed in Alg. 1, is iterated until convergence. Subroutine `ccomponents` (line 2) finds the c -components of M , `initialisation` (line 1) provides a random initialisation of the exogenous PMFs, line 4 corresponds to a restriction of the data set to the variables in the c -th c -component (line 4), while ε is the threshold to evaluate parameter convergence with respect to a metric δ (line 6). Note that in line 8 the d -separation properties of c -components allow to replace $P(U|x)$ with $P(U|\mathbf{y}^{(c)})$, where c is such that $U \in \mathcal{U}^{(c)}$.

Algorithm 1 EMCC: given SCM M and data set \mathcal{D} returns $\{P(U)\}_{U \in \mathcal{U}}$.

```

1:  $\{P_0(U)\}_{U \in \mathcal{U}} \leftarrow \text{initialisation}(M)$ 
2:  $\{\mathcal{U}^{(c)}, \mathbf{X}^{(c)}\}_{c \in \mathcal{C}} \leftarrow \text{ccomponents}(M)$ 
3: for  $c \in \mathcal{C}$  do
4:    $\mathcal{D}^{(c)} \leftarrow \mathcal{D} \upharpoonright \mathbf{Y}^{(c)}$ 
5:    $t \leftarrow 0$ 
6:   while  $\delta[\{P_{t+1}(U)\}_{U \in \mathcal{U}}, \{P_t(U)\}_{U \in \mathcal{U}}] > \varepsilon$  do
7:     for  $U \in \mathcal{U}^{(c)}$  do
8:        $P_{t+1}(U) \leftarrow |\mathcal{D}|^{-1} \sum_{\mathbf{y}^{(c)} \in \mathcal{D}^{(c)}} P_t(U|\mathbf{y}^{(c)})$ 
9:        $t \leftarrow t + 1$ 
10:    end for
11:  end while
12: end for

```

And since P_{t+1} gets a higher marginal log-likelihood than P_t [7, Thm. 19.3], it readily follows that

Corollary 2. *EMCC always converges returning a point of \mathcal{K} .*

Thus multiple EMCC runs, started with different seeds, will yield an approximating subset of \mathcal{K} .

As a side remark, note that each iteration of the loop in line 3 of Alg. 1 can be executed in parallel to the others because of the d-separation among c-components. Something similar can be done at the data set level when computing, by standard BN algorithms, the queries in line 8.

5 M -Compatibility and the Limits of Rung 2 of Pearl's Hierarchy

In this section we give insights on M -compatibility, as defined in Sect. 3, and on its consequences. Let us first consider a simple example to be used along the section to clarify our findings.

Example 2. Consider the example in [8, Sect. 4.1]. It refers to a study about the recovery Y of patients of gender Z possibly subject to treatment X .

A sample of 700 patients is considered (the corresponding frequencies are in Tab. 1). The authors show that for an SCM whose graph is like in Fig. 1f, given the sample, the probability of necessity and sufficiency for treatment (X) on effect (Y) is no greater than 0.01 (or PNS: $P(Y_{X=1} = 1, Y_{X=0} = 0)$ [9]). They tell this for any SCM with that graph, irrespectively of the specific SEs it uses; they only require that the empirical distribution factorises according to the graph.

Gender (Z)	Treatment (X)	Recovery (Y)	#	Gender (Z)	Treatment (X)	Recovery (Y)	#
0	0	0	2	1	0	0	107
0	0	1	114	1	0	1	13
0	1	0	41	1	1	0	109
0	1	1	313	1	1	1	1

Table 1: Data from an observational study involving three Boolean variables [8]. Positive states means female (Z), treated (X) and recovered (Y).

Let us call *conservative* an SCM whose SEs are specified in a conservative way. We focus on Markovian models because our notion of conservative SE, in Sect. 2, refers to this case only (an extension to general SMCs can be obtained relying on recent work [6, 12]). The following holds:

Theorem 3. Let M be a conservative Markovian SCM over (\mathbf{X}, U) . For any empirical $P(\mathbf{X})$, there exists at least a specification of $\{P(U)\}_{U \in \mathcal{U}}$ that guarantees M -compatibility.

In practice, by Thm. 3, if the SEs of a Markovian SCM are not available and we adopt a conservative specification, compatibility is guaranteed and, by Cor. 2, EMCC can eventually be used for inference. This is what we do in the following example, which refers to the same setup of Ex. 2.

Example 3. Assume that the Markovian SCM in Fig. 1f has generated the data in Tab. 1. In particular assume M conservative. This corresponds to $|\Omega_W| = 2$, $|\Omega_U| = 4$ and $|\Omega_V| = 16$. For the SEs we have that f_Z is just the identity map, while f_X is such that:

$$f_X(Z, U = 0) = \neg Z, f_X(Z, U = 1) = 0, f_X(Z, U = 2) = Z, f_X(Z, U = 3) = 1.$$

The relations between Y and (X, Z) according to f_Y and corresponding to the sixteen states of V can be similarly listed. M -compatibility with data in Tab. 1 corresponds to the constraints in Eq. (9). For W , this trivially means $P(W = 0) = P(Z = 0) = \frac{470}{700}$ and $P(W = 1) = P(Z = 1) = \frac{230}{700}$. For U we have instead the two following independent constraints for PMF $P(U)$:

$$P(U = 1) + P(U = 2) = P(X = 0 | Z = 0) = \frac{116}{470}, \quad (10)$$

$$P(U = 0) + P(U = 1) = P(X = 0 | Z = 1) = \frac{120}{230}. \quad (11)$$

Four independent linear constraints can similarly be obtained for $P(V)$. It is a straightforward exercise to obtain a parametrisation for all $P(U)$ and $P(V)$ specifications consistent with the above constraints, and verify that the corresponding interval for PNS is the same as the one given in [8] on the basis of bounding formulae. Note that in that paper the upper bound has been rounded to the second decimal place, while a more precise estimate would be 0.015. This is the same number we get with EMCC if we compute PNS for the conservative specification of the SCM in Fig. 1f.

Thus the bounds in [8] correctly predict the right numbers we get in the above example. But we are not only assuming that the data factorise according to the graph in Fig. 1f, as the mentioned paper

does. We are also assuming that the specification is conservative (i.e., all the mechanisms are possible). What happens if we do not—considered that the latter is not an assumption required to derive the bounds. This is discussed in the next example, still based on the setup of Ex. 2.

Example 4. *Assume that, because of some expert knowledge, the state $U = 3$ is impossible, this meaning that a deterministic mechanism forcing the treatment of all the patients is considered unrealistic. Such an assumption preserves the separate surjectivity of f_X and also the joint surjectivity of the model. In other words, for each joint state of (X, Y, Z) , there is at least a joint state of (W, U, V) inducing it. Under this additional information, the negation of Eq. (10) becomes:*

$$P(U = 0) = P(X = 1|Z = 0) = \frac{354}{470}, \quad (12)$$

and this is clearly inconsistent with Eq. (11).

This means that the model obtained by dropping state $U = 3$ cannot, in the limit, generate the distribution that we see in the data: namely, the data are incompatible with such a non-conservative SCM. Note also that this shows that the empirical distribution of the data can factorise according to the graph, while at the same time, it may be incompatible with the related SCM. This may happen because the latter form of incompatibility depends on the SEs, not only on the graph.

In fact, as shown by the next example, M -incompatibility is subtle in that it may be based on an empirical probability that is inconsistent with the very SCM used to produce the data!

Example 5. *In the same setup of Ex. 4 assume that some expert tells us that, out the sixteen states of V , only three states are possible, namely those indexing the following logical relations:*

$$f_Y(X, Z, V = 0) = X \vee \neg Z, f_Y(X, Z, V = 1) = \neg X \vee \neg Z, f_Y(X, Z, V = 2) = \neg X \wedge Z.$$

Note that this preserves surjectivity. Call M' the reduced SCM that we obtain from the conservative specification (in M) by keeping only the values 0, 1, 2 both for U and V . Consider the following distributions for the exogenous variables in M' : $P(V) = [0.47, 0.439, 0.091]$, $P(U) = [0.677, 0.000, 0.323]$, and $P(W)$ taking the same empirical values of $P(Z)$ as in the case of M . Because of the previous discussion, the data in Tab. 1 are clearly M' -incompatible. And yet they can be produced by M' , because $P(X, Y, Z)$ is a positive distribution under M' . So any data can be generated—with different log-likelihoods. In particular, the ratio between the marginal log-likelihood of M' and that of M is 0.71: thus it turns out that it is not at all unlikely to produce the data in Tab. 1 with model M' . And since those data factorise according to the graph, as before, we should be allowed to apply the bounds as before, claiming that the upper bound of PNS for M' is 0.015. But the PNS for M' is 0.15, 10 times above the bound. And so the bounds in [8, Sect. 4.1] fail here.

They fail because ‘factorising according to the graph’ is not a strong enough assumption to derive bounds; we should ask for more, namely, M -compatibility with the underlying SCM. Stated differently, whenever we produce formulae, or algorithms, for computing counterfactuals, we should make sure that M -compatibility holds, otherwise the results will be unwarranted. Remember that M -compatibility can easily be verified in general through EMCC (see Cor. 1 and the paragraph that follows it); we are not aware of any other such simple and general method to verify it.

At this point, one might be tempted to escape the problem of testing M -compatibility in absence of SEs (i.e., in presence of data and the causal graph alone) by using a conservative specification of the SEs, given that the resulting model is always compatible with the data (cf. Thm. 3). Even more so that recent efforts appear to provide a conservative specification for general SCMs [6, 12]; this would seem to enable doing general counterfactual inference without SEs. Yet, conservative specifications do not seem to provide us with an easy way out, as shown next.

Let us recall that the notion of M -compatibility in Eq. (7) refers to FSCMs. For PSCMs, we should consider the set \mathcal{K} as in Thm. 1. Each element of \mathcal{K} defines an FSCM, which is M -compatible by construction. Whence M -compatibility with a PSCM corresponds to \mathcal{K} being non-empty.

Note that with Markovian (and quasi-Markovian) models, the constraints in Eq. (9) can be written separately for each $U \in \mathcal{U}$; we denote by $\mathcal{K}(U)$ the set of specifications of $P(U)$ consistent with those ‘local’ constraints. In this respect, \mathcal{K} can be understood as the set of stochastic products of elements from sets $\mathcal{K}(U)$, with $U \in \mathcal{U}$.

Given a Markovian conservative PSCM M , any other Markovian PSCM M' over the same endogenous variables and graph can be defined by dropping some values of the exogenous variables,

along with their corresponding SEs. That is, M' is defined via sets $\{\Omega'_U\}_{U \in \mathcal{U}}$ such that, for each $U \in \mathcal{U}$, $\Omega'_U \subseteq \Omega_U$ are the states of U indexing the deterministic relations of the corresponding SE of M' . As a consequence we can regard a conservative PSCM as the set of all the possible non-conservative PSCMs. Some of these may be incompatible with the data, though. So how does inference in the conservative PSCM and in the set of non-conservative PSCMs relate to each other?

To answer this question, let us say that M *embeds* M' if and only if M gives a chance to M' to be the actual model underlying the data, i.e., if for all $U \in \mathcal{U}$ there is at least a $P(U) \in \mathcal{K}(U)$ so that $P(U \in \Omega'_U) > 0$. The following holds:

Theorem 4. *A Markovian conservative PSCM M cannot embed incompatible models.*

As a consequence, running EMCC, or, e.g., the inference algorithms in [11] in a conservative model corresponds to *automatically discarding* the incompatible models while using only the compatible ones. The important implication here is that if the true underlying model is not compatible with the available data, the results obtained by using the conservative model will be unwarranted as an approximation to the actual one. This is in fact the reason why in Example 5 the PNS interval obtained by the conservative model does not contain the actual value of PNS.

The lesson appears to be that we cannot have guaranteed bounds without knowing the SEs of the underlying SCM. Using the conservative specification as a replacement for the actual SEs might help under the assumption that the sample is large enough so as it becomes M -compatible with the true, underlying, model. The challenge is to render such an assumption tenable.

6 Evaluating the EMCC

We start by characterising the accuracy of our procedure in terms of credible intervals.

Theorem 5. *Let $[a^*, b^*]$ denote the exact probability bounds of a causal query. Say that $\rho := \{r_i\}_{i=1}^n$ are the outputs of n EMCC iterations, while $[a, b]$ is the interval induced by ρ , i.e., $a := \min_{i=1}^n r_i$ and $b := \max_{i=1}^n r_i$. By construction $a^* \leq a \leq b \leq b^*$. The following inequality holds:*

$$P\left(a - \varepsilon L \leq a^* \leq b^* \leq b + \varepsilon L \mid \rho\right) = \frac{1 + (1 + 2\varepsilon)^{2-n} - 2(1 + \varepsilon)^{2-n}}{(1 - L^{n-2}) - (n-2)(1-L)L^{n-2}}, \quad (13)$$

where $L := (b - a)$ and $\varepsilon := \delta / (2L)$ is the relative error at each extreme of the interval obtained as a function of the absolute allowed error $\delta \in (0, L)$.

Eq. (13) can be used to evaluate the EMCC accuracy in approximating the true interval $[a^*, b^*]$. For example, if we allow for a maximum 17% of error (ε) at each extreme of the interval $[a, b]$, 20 EM runs achieve 95% credibility. By numerically taking the limit of Eq. (13) with L that tends to the machine zero, we can address the case of (candidate) identifiable queries, i.e., n EMCC runs always returning the same value. In this case 6 runs achieve 95% credibility; 9 runs lead to 99%.

We also evaluate the EMCC via empirical experiments. This will provide us with an actual measure of its performance. A benchmark of 460 SCMs is considered. These are Boolean endogenous chains of $m \in \{5, 7, 10\}$ nodes augmented with: an exogenous parent per node in the Markovian case, exogenous parents of pairs of randomly picked endogenous nodes whose distance in the chain is no more than two in the other cases. A single exogenous node per c-component is allowed in the quasi-Markovian case, two in non-quasi-Markovian models.³

We consider PNS queries involving the first and the last endogenous variable. We start from a ‘true’ FSCM M , from which a data set \mathcal{D} of 1000 instances is sampled. Afterwards the corresponding PSCM is considered. Endogenous data are processed by EMCC with random initialisation, KL divergence with double machine epsilon decides convergence (Alg. 1, line 6). PNS queries are addressed in the derived counterfactual graph and the bounds with respect to multiple EMCC runs evaluated. The baseline procedure in [11] allows one to compute the exact bounds for Markovian and quasi-Markovian models. For the non-quasi-Markovian case, exact bounds are unavailable; for the sake of evaluation, we assume they coincide with those we get by a very large number of EMCC runs. Fig. 3 depicts, as a function of the number of the EM runs, the average RMSE for the approximate bounds with respect to the baseline (i.e., $\sqrt{((a^* - a)^2 + (b^* - b)^2) / 2}$).

³The EMCC code is part of the CREDICI library [2]. The code to reproduce the models and the experiments is available at <https://github.com/IDSIA-papers/2021-NeurIPSWHY-causalEM>.




Figure 3: EMCC RMSE vs. number of runs (n).

The behaviour is consistent with Thm. 5: the EMCC gives an inner approximation that quickly ($n \leq 20$) converges to a very accurate estimate (RMSE < 0.01). For small PNS intervals (e.g., for $m = 10$ in non-quasi-Markovian models), ten runs allow to span the exact interval, again quite consistently with Thm. 5. The case $m = 10$ for quasi-Markovian models is even more extreme as the exact PNS interval is small and numerically collapses into a single sharp value. In such cases, by construction, the EMCC gives the exact value in a single run.




Figure 4: Average EMCC execution times.




Figure 5: Average interval sizes for PNS queries.

Fig. 4 depicts the average sizes of the PNS exact intervals. Fig. 5 shows the EMCC running times normalised with respect to the number of runs. The most time-consuming tasks are those involving non-quasi-Markovian models. This is expected as the complexity of our procedure is bounded by the treewidth of the reduced graph (Alg. 1, line 8), which is one for both Markovian and quasi-Markovian models. Note that quasi-Markovian models are presumably faster because of the lower number of exogenous nodes. In fact, m only reflects the number of endogenous variables, the overall number of nodes is $2m$ in the Markovian case and $1.5m$ in the other cases. The PNS query is finally addressed in the counterfactual graph, whose dimension is $3m$ in the Markovian case, and $2.5m$ in other cases.

7 Conclusions

We have presented a causal version of the EM algorithm that can be used to do counterfactual inference of unidentifiable queries. The algorithm is simple: its two components are the EM and some algorithms for BNs. It is an approximate algorithm too. This is a limitation. Yet, we have shown here that causal inference is NP-hard even in simple models, whence approximations seem necessary. Moreover, the EMCC is shown to offer a good compromise between efficiency and accuracy. In addition, EMCC gives us the opportunity to test the compatibility between a structural causal model and a sample. Checking it appears necessary to avoid drawing unwarranted conclusions. Future work should consider extending EMCC to continuous domains.

A Proofs

Proof of Theorem 1. Let us first consider the case $\mathcal{K} \neq \emptyset$ and take $\{P'(U)\}_{U \in \mathcal{U}} \in \mathcal{K}$. By definition, Eq. (9) for SCM M rewrites as:

$$\sum_{\mathbf{u}^{(c)}} \left[\prod_{X \in \mathbf{X}^{(c)}} P(x|\text{pa}_X) \prod_{U \in \mathcal{U}^{(c)}} P'(u) \right] = \prod_{X \in \mathbf{X}^{(c)}} \hat{P}(x|\mathbf{y}_X^{(c)}). \quad (14)$$

Putting this in Eq. (8) we get:

$$LL = \sum_{c \in \mathcal{C}} \sum_{\mathbf{y}^{(c)}} n(\mathbf{y}^{(c)}) \sum_{X \in \mathbf{X}^{(c)}} \log \hat{P}(x|\mathbf{y}_X^{(c)}), \quad (15)$$

and this is Eq. (4) when the probabilities are those in Eq. (6). This proves the sufficient condition. To prove the necessary condition, consider again Eq. (14), which is equivalent to Eq. (9). Assume, ad absurdum, that there is a $\{P'(U)\}_{U \in \mathcal{U}}$ where the log-likelihood attains its global maximum, but $\{P'(U)\}_{U \in \mathcal{U}} \notin \mathcal{K}$. Thus, Eq. (14) should be violated for at least a value of $\mathbf{y}^{(c)}$ and c . Yet, putting Eq. (14) in Eq. (4) produces a value smaller than the global maximum, as this is achieved if and only if the values in Eq. (6) are used. The same discussion applies to any $\{P'(U)\}_{U \in \mathcal{U}}$ when $\mathcal{K} = \emptyset$. \square

Proof of Theorem 2. The proof follows from the analogous result for credal networks in [3, Thm. 1]. The polytree used in that proof has deterministic conditional probability tables for non-root nodes. We can therefore understand that model as an SCM. Its endogenous nodes form a chain, each node has a single exogenous parent apart from the first one in the chain that has a second exogenous parent. The set-valued quantification of the exogenous PMFs corresponds to a partially specified SCM. As a query, the authors consider the upper bound of a marginal query in the last node of the chain. The task amounts to the identification of the upper bound of the causal effect on the last node of the chain given an intervention in an additional endogenous parent of the first node of the chain. An algorithm to bound interventional queries in SCMs would therefore solve inference in polytree-shaped credal networks. This contradicts the result of the authors. \square

Proof of Theorem 3. The thesis is equivalent to the fact that the constraints in Eq. (9) can be satisfied for at least a $P(\mathbf{U})$. As M is Markovian, for each $c \in \mathcal{C}$, $|\mathcal{U}_c| = |\mathbf{X}_c| = 1$ and we denote as X and U the unique elements of \mathbf{X}_c and \mathcal{U}_c , while \mathbf{Y} is the joint variable corresponding to the endogenous parents of X . Eq. (9) rewrites as the following linear constraint:

$$\sum_{u \in \Omega_U} P(x|\mathbf{y}, u) \cdot P(u) = \sum_{u \in \Omega_U: f(\mathbf{y}, u) = x} P(u) = P(x|\mathbf{y}), \quad (16)$$

to be satisfied for each $x \in \Omega_X$ and $\mathbf{y} \in \Omega_{\mathbf{Y}}$. The representability result provided in [5, Thm. 1] can be used to prove that the linear constraints in Eq. (16) can be satisfied. In their proof the authors consider a setup analogous to the current one but, instead of U , a continuous $U' \in [0, 1]$ with a uniform density $P(U')$ is considered. Say that $\Omega_X := \{x_1, \dots, x_q\}$. For each $\mathbf{y} \in \Omega_{\mathbf{Y}}$, we define the vector $H_{\mathbf{y}} := \{h_{\mathbf{y}}^{(i)}\}_{i=0}^q$ such that $h_{\mathbf{y}}^{(0)} := 0$ and $h_{\mathbf{y}}^{(i)} := \sum_{j=1}^i P(x_j|\mathbf{y})$ for each $i = 1, \dots, q$. Note that $h_{\mathbf{y}}^{(i)} \leq h_{\mathbf{y}}^{(i+1)}$ for each $i = 1, \dots, q-1$ and $h_{\mathbf{y}}^{(q)} = 1$. The authors show that Eq. (16) is satisfied if the SE $X = f'_X(\mathbf{Y}, U')$ is such that:

$$f'_X(\mathbf{y}, u') := \left\{ x_i \in \Omega_X : u' \in [h_{\mathbf{y}}^{(i)}, h_{\mathbf{y}}^{(i+1)}] \right\}, \quad (17)$$

for each $u' \in [0, 1]$ and $\mathbf{y} \in \Omega_{\mathbf{Y}}$. A partition of $[0, 1]$ is obtained by removing the left endpoints from the intervals in Eq. (17) apart from the first one. Thus, for a given $\mathbf{y} \in \Omega_{\mathbf{Y}}$, Eq. (17) can be regarded as a discrete SE, mapping to X the values of a discretisation of U' based on the partition of $[0, 1]$ induced by $H_{\mathbf{y}}$.

A least common partition is obtained from a set of partitions by simply putting together and sorting the endpoints of the intervals of all the partitions. Take the discretisation of U' induced by such a least common partition when considering all the partitions of $[0, 1]$ induced by each $\mathbf{y} \in \Omega_{\mathbf{Y}}$. In practice, the SE $X = f'_X(\mathbf{Y}, U')$ can be equivalently described by $X = \hat{f}_X(\mathbf{Y}, \hat{U})$, where \hat{U} is a discrete variable whose states are in correspondence with the above considered set of discretisation intervals for U . The uniform density $P(U')$ is consequently mapped to a PMF $P(\hat{U})$ such that $P(\hat{U} = \hat{u})$ is equal to the integral of $P(U')$ on the interval associated with \hat{u} and hence it is equal to its width. Finally, observe that, for each $\hat{u} \in \Omega_{\hat{U}}$, \hat{f}_X defines a deterministic relation between \mathbf{Y} and \mathbf{X} and

this should correspond to a state of U in M , as U is enumerating all these possible relations because of the conservative specification. This defines a map $\mu: \Omega_{\hat{U}} \rightarrow \Omega_U$ with $\mu(\hat{u}) := \{u \in \Omega_U : f_X(u, \mathbf{Y}) = \hat{f}(\hat{u}, \mathbf{Y})\}$. For $P(U)$ we have:

$$P(u) = \sum_{\hat{u} \in \Omega_{\hat{U}}: \gamma(\hat{u})=u} P(\hat{u}), \quad (18)$$

and $P(u) = 0$ for the states of u that are not in the domain of μ . This is the exogenous quantification that proves the thesis. \square

Proof of Theorem 4. Assume, ad absurdum, that M embeds an incompatible M' . For each $U \in \mathcal{U}$, let $\{\mathcal{K}'(U)\}_{U \in \mathcal{U}}$ denote the sets of compatible specifications of M' . By Cor. 1, the incompatibility of M' implies that $\mathcal{K}'(U)$ should be empty for at least a $U \in \mathcal{U}$. By the definition of embedding, for each $U \in \mathcal{U}$, we should have at least a $P(U) \in \mathcal{K}(U)$ such that $P(U \in \Omega'_U) > 0$.

Obtaining M' from M can be regarded as the result of conditioning the exogenous PMFs on the events $U \in \Omega'_U$, for each $U \in \mathcal{U}$. As $P(U \in \Omega'_U) > 0$, such conditioning is well defined for P . Let $P'(U)$ denote the resulting PMF for M' .

Consider the compatibility constraints of M as in Eq. (9) involving PMF $P(U)$. As $P(U) \in \mathcal{K}(U)$, these constraints should be satisfied by $P(U)$. With Markovian models, the constraints are linear. Dividing by $P(U \in \Omega'_U)$ on both sides gives the analogous constraints for $P'(U)$. But this means $P'(U) \in \mathcal{K}'(U)$ and hence $\mathcal{K}'(U) \neq \emptyset$ (for each $U \in \mathcal{U}$), which is a contradiction. \square

Proof of Theorem 5. Consider the probability on the left-hand side of Eq. (13). The corresponding joint probability rewrites as:

$$P\left(\Delta_a \leq \frac{\delta}{2}, \Delta_b \leq \frac{\delta}{2}, \rho\right) = \gamma \int_0^{\frac{\delta}{2}} \int_0^{\frac{\delta}{2}} P(\rho | \Delta_a = x, \Delta_b = y) dx dy, \quad (19)$$

where a uniform prior density is considered, i.e., $P(\Delta_a = x, \Delta_b = y) := \gamma dx dy$ with γ normalisation constant, $\Delta_a := (a - a^*)$ and $\Delta_b := (b^* - b)$.

To obtain $P(\rho | \Delta_a = x, \Delta_b = y)$, let us first notice that, for each $r_i \in \rho$:

$$P(r_i \in [a, b] | \Delta_a = x, \Delta_b = y) = \frac{b - a}{b - a + x + y}. \quad (20)$$

Eq. (20) is obtained by assuming the output r_i of a single EMCC run to be uniformly distributed over the interval $[a^*, b^*]$. Under i.i.d. assumptions we get $P(\rho | \Delta_a = x, \Delta_b = y) = \left(\frac{b-a}{b-a+x+y}\right)^n$ from Eq. (20). Eq. (19) rewrites therefore as:

$$P\left(\rho, \Delta_a \leq \frac{\delta}{2}, \Delta_b \leq \frac{\delta}{2}\right) = \gamma \int_0^{\frac{\delta}{2}} \int_0^{\frac{\delta}{2}} \left(\frac{b-a}{b-a+x+y}\right)^n dx dy. \quad (21)$$

We can similarly obtain the marginal for ρ by integrating Δ_a and Δ_b , i.e.,

$$P(\rho) = \gamma \int_0^{a+(1-b)} \left[\int_0^{a+(1-b)-y} \left(\frac{b-a}{b-a+x+y}\right)^n dx \right] dy. \quad (22)$$

The integrals in Eqs. (21) and (22) are trivial and their ratio gives the right-hand side of Eq. (13). \square

References

- [1] E. Bareinboim, J.D. Correa, D. Ibeling, and T. Icard. On Pearl’s hierarchy and the foundations of causal inference. *ACM Special Volume in Honor of Judea Pearl*, 2(3), 2020.
- [2] R. Cabañas, A. Antonucci, D. Huber, and M. Zaffalon. CREDICI: a Java library for causal inference by credal networks. In *Proceedings of the 10th International Conference on Probabilistic Graphical Models*, Proceedings of Machine Learning Research, pages 597–600. PMLR, 2020.
- [3] J. C. F. da Rocha and F. G. Cozman. Inference with separately specified sets of probabilities in credal networks. In *Proceedings of the Eighteenth Annual Conference on Uncertainty in Artificial Intelligence*, pages 430–437. Morgan Kaufmann, 2002.
- [4] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39(1):1–22, 1977.
- [5] M. J. Druzdzel and H. A. Simon. Causality in Bayesian belief networks. In *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence*, pages 3–11. Morgan Kaufmann, 1993.
- [6] G. Duarte, Finkelsteinm N., D. Knox, J. Mummolo, and I. Shpitser. An automated approach to causal inference in discrete settings. *arXiv:2109.13471*, 2021.
- [7] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT, 2009.
- [8] S. Mueller, A. Li, and J. Pearl. Causes of effects: Learning individual responses from population data. *arXiv:2104.13730*, 2021.
- [9] J. Pearl. *Causality*. Cambridge University Press, 2009.
- [10] J. Tian. *Studies in Causal Reasoning and Learning*. PhD thesis, UCLA, 2002.
- [11] M. Zaffalon, A. Antonucci, and R. Cabañas. Structural causal models are (solvable by) credal networks. In *Proceedings of the tenth International Conference on Probabilistic Graphical Models*, Proceedings of Machine Learning Research, pages 581–592. PMLR, 2020.
- [12] J. Zhang, J. Tian, and E. Bareinboim. Partial identification of counterfactual distributions. Technical Report R-78, Causal Artificial Intelligence Lab, Columbia University, Jun 2021.