

Estimating Causal Effects from Learned Causal Networks

Anna K Raichev¹ and Jin Tian² and Alexander Ihler¹ and Rina Dechter¹
¹University of California, Irvine ²University of Iowa
araichev@uci.edu, jtian@iastate.edu, ihler@ics.uci.edu, dechter@ics.uci.edu

Abstract. The standard approach to answering an identifiable causal-effect query (e.g., $P(Y|do(X))$) when given a causal diagram and observational data is to first generate an estimand, or probabilistic expression over the observable variables, which is then evaluated using the observational data. In this paper, we propose an alternative paradigm for answering causal-effect queries over discrete observable variables. We propose to instead learn the causal Bayesian network and its confounding latent variables directly from the observational data. Then, efficient probabilistic graphical model (PGM) algorithms can be applied to the learned model to answer queries. Perhaps surprisingly, we show that this *model completion* learning approach can be more effective than estimand approaches, particularly for larger models in which the estimand expressions become computationally difficult. We illustrate our method’s potential using a benchmark collection of Bayesian networks and synthetically generated causal models.

1 Introduction

Structural Causal Models (SCMs) are a formal framework for reasoning about causal knowledge in the presence of uncertainty [23]. When the full SCM is available, it is possible to use standard probabilistic inference [6, 8] to directly answer causal queries that evaluate how forcing some variable’s assignment $X = x$ affects another variable Y , written as $P(Y|do(X = x))$. However, in practice the full causal model is rarely available; instead, only a causal diagram – a directed graph capturing the causal relationships of the underlying SCM – is assumed to be known. Causal diagrams typically include both *observable variables*, which can be measured from data, and *latent variables*, which are unobservable and for which data cannot be collected. A linchpin of causal reasoning is that many causal queries can be uniquely answered using only the observable variables [28, 24, 25, 11] and consequently estimated using observational data.

The main approach developed in the past two decades for answering causal queries under such assumptions is a two-step process which we call *estimand-based causal inference*. The first step is to determine if the causal query is *identifiable* – i.e., uniquely answerable from the model’s observational distribution – and if so, construct an expression (“estimand”) that captures the answer symbolically using probabilities over only observed variables. Then, one can use the observed data to estimate the probabilities involved in the estimand expression. Over the past few years a variety of estimand-based strategies have been developed using different statistical estimation methods [14, 3].

However, many of these approaches focus on specific, small, causal models, and do not examine the scalability of their approaches or

provide significant empirical comparisons. In larger models, estimand expressions can become large and unwieldy, making them either computationally difficult, hard to estimate accurately, or both.


In this paper, we propose a straightforward yet under-explored alternative: to learn the causal model, including latent dependencies, directly from the observed data. Although the domains and distributional forms of these latent variables are unknown, recent work [34] has shown that, if the visible variables are discrete, any SCM is equivalent to one with discrete latent variables, and gives an upper bound on their domain sizes. This allows us to apply well-known techniques for learning latent variable models, such as the Expectation-Maximization (EM) algorithm, to build a *Causal Bayesian Network (CBN)* over the observed and latent variables. We can further apply model selection techniques such as the Bayesian Information Criterion (BIC, [16]) to select appropriate domain sizes. Then, given our learned model, we can use efficient algorithms for reasoning in probabilistic graphical models (PGMs) to answer one, or even many causal queries [6, 16].

Perhaps surprisingly, we show that this approach is often significantly more effective than the estimand-based methodology. Furthermore, we provide structural conditions that help decide when each approach is likely to be more effective. The computational benefit of our approach is tied to the complexity of the causal graph: if the causal graph has low *treewidth*, then exact PGM algorithms (e.g., bucket elimination or join-tree scheme) are efficiently applicable. As an added benefit if there are multiple causal queries to perform on the same model, since the model is learned only once, the learning time can be amortized over all such queries.


Our empirical evaluation incorporates a spectrum of models, including not only the small models common in causal effect literature, but also large models based on benchmark Bayesian networks and scalable classes of synthetic models; to the best of our knowledge this is the first such extensive empirical evaluation of causal effect queries.

Contributions. This paper presents a new path for answering causal effects queries by learning a Causal Bayesian Network that is consistent with a causal graph and observational data.


- We propose to answer causal queries by directly learning the full causal Bayesian network via EM, followed by query processing using traditional PGM algorithms.
- We provide a first of its kind, empirical evaluation of algorithms for causal effect queries on varied and large synthetic and real benchmarks.
- We show empirically that our proposed learning approach gives more accurate estimates than the estimand-based alternatives.



(a) Causal diagram of a chain model with 7 observable and 3 latent variables. Dashed bidirected edges represent latent variables.



(b) The CBN for the model (a), showing latent variables explicitly.



(c) The truncated causal diagram after intervention $do(V_0)$.

Figure 1: Illustration of a causal diagram, CBN, and truncated CBN for evaluating a causal effect.

72 In settings with high-dimensional estimand expressions but low
73 treewidth causal models, our approach allows more accurate estimates
74 by retaining more information from the causal graph, and allowing
75 variance reduction without impacting computational tractability.

76 The rest of the paper is organized as follows. Section 2 provides
77 background on related work and information and definitions, Section
78 3 outlines the main idea of our approach, Section 4 provides empirical
79 evaluation, and we conclude in Section 5.

80 2 Background

81 We first provide some definitions and technical background. We use
82 capital letters X to represent variables, and lower case x for their
83 values. Bold uppercase \mathbf{X} denotes a collection of variables, with $|\mathbf{X}|$
84 its cardinality and $\mathcal{D}(\mathbf{X})$ their joint domain, while \mathbf{x} indicates an
85 assignment in that joint domain. We denote by $P(\mathbf{X})$ the joint proba-
86 bility distribution over \mathbf{X} and $P(\mathbf{X} = \mathbf{x})$ the probability of \mathbf{X} taking
87 configuration \mathbf{x} under this distribution, which we abbreviate $P(\mathbf{x})$.
88 Similarly, $P(\mathbf{Y} | \mathbf{X} = \mathbf{x})$ represents the set of conditional distributions
89 $P(\mathbf{Y} | \mathbf{X} = \mathbf{x})$, $\forall \mathbf{x}$.

90 **Definition 1** (Structural Causal Model). *A structural causal model*
91 *(SCM) [23] is a 4-tuple $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U}) \rangle$ where: (1) $\mathbf{U} =$
92 $\{U_1, U_2, \dots, U_k\}$ is a set of exogenous latent variables whose values
93 are affected by factors outside the model; (2) $\mathbf{V} = \{V_1, V_2, \dots, V_n\}$ is
94 a set of endogenous, observable variables of interest whose values are
95 determined by other variables in the model; (3) $\mathbf{F} = \{f_i : V_i \in \mathbf{V}\}$
96 is a set of functions f_i such that each f_i determines the value v_i of
97 V_i as a function of V_i 's causal parents $PA_i \subseteq \mathbf{U} \cup (\mathbf{V} \setminus V_i)$ so
98 that $f_i : \mathcal{D}(PA_i) \rightarrow \mathcal{D}(V_i)$ and $v_i \leftarrow f_i(pa_i)$; and (4) $P(\mathbf{U})$ is a
99 probability distribution over the latent variables. The latent variables
100 are assumed to be mutually independent, i.e., $P(\mathbf{U}) = \prod_{U \in \mathbf{U}} P(U)$.*

101 **Causal diagrams.** An SCM \mathcal{M} induces a *causal diagram* which is
102 a directed graph $\mathcal{G} = \langle \mathbf{V} \cup \mathbf{U}, E \rangle$, where each node in the graph maps
103 to a variable in the SCM, and there is an arc from every observed or
104 latent node X in the graph to V_i iff $X \in PA_i$. An SCM whose latent
105 variables connect to at most a single observable variable is referred
106 to as *Markovian*, and one whose latent variables connect to at most
107 two observable variables is called *Semi-Markovian*. It is known that
108 any SCM can be transformed into an equivalent Semi-Markovian one
109 such that answers to causal queries are preserved [27]. In the semi-
110 Markovian case it is common to use a simplified causal diagram called
111 an *Acyclic Directed Mixed Graph* (ADMG), which omits any latent
112 variables with a single child, and replaces any latent variable with two

113 children with a bidirectional dashed arc between the children, so that
114 latent variables are no longer explicitly shown (see, e.g., Figure 1).

The SCM \mathcal{M} also induces a **Causal Bayesian Network (CBN)**
 $\mathcal{B} = \langle \mathcal{G}, \mathcal{P} \rangle$ specified by \mathcal{M} 's causal diagram $\mathcal{G} = \langle \mathbf{V} \cup \mathbf{U}, E \rangle$
along with its associated conditional probability distributions $\mathcal{P} =$
 $\{P(V_i | PA_i), P(U_j)\}$. The distribution $P(\mathbf{V}, \mathbf{U})$ factors according
to the causal diagram:

$$P(\mathbf{V}, \mathbf{U}) = \prod_{V_i \in \mathbf{V}} P(V_i | PA_i) \cdot \prod_{U_j \in \mathbf{U}} P(U_j). \quad (1)$$

The *observational distribution*, $P(\mathbf{V})$, is given by

$$P(\mathbf{V}) = \sum_{\mathbf{u}} P(\mathbf{V}, \mathbf{U} = \mathbf{u}). \quad (2)$$

Causal effect and the truncation formula. An external interven-
tion, forcing variables \mathbf{X} to take on value \mathbf{x} , is modeled by replacing
the mechanism for each $X \in \mathbf{X}$ with the function $X = x$. This
is formally represented by the do-operator $do(\mathbf{X} = \mathbf{x})$. Thus the
interventional distribution of an SCM \mathcal{M} by applying $do(\mathbf{X})$ is,

$$P(\mathbf{V} \setminus \mathbf{X}, \mathbf{U} | do(\mathbf{X} = \mathbf{x})) = \prod_{V_j \notin \mathbf{X}} P(V_j | PA_j) \cdot P(\mathbf{U}) \Big|_{\mathbf{X} = \mathbf{x}} \quad (3)$$

i.e., it is obtained from Eq. (1) by truncating the factors corresponding
to the variables in \mathbf{X} and setting $\mathbf{X} = \mathbf{x}$. The effect of $do(\mathbf{X})$ on a
variable Y , denoted $P(Y | do(\mathbf{X}))$, is defined by marginalizing all the
variables other than Y .

Causal queries. While we normally have no access to the full struc-
tural causal model \mathcal{M} and cannot evaluate these expressions directly,
it is sometimes possible to evaluate the effect of an intervention given
only the observational distribution $P(\mathbf{V})$, specifically, if the answer is
unique for any full model that is consistent with the graph and $P(\mathbf{V})$
[23]. More formally:

Definition 2 (Identifiability). *Given a causal diagram \mathcal{G} , the causal*
*effect query $P(\mathbf{Y} | do(\mathbf{X}))$ is **identifiable** if any two SCMs consistent*
with \mathcal{G} that agree on the observational distribution $P(\mathbf{V})$ also agree
on $P(\mathbf{Y} | do(\mathbf{X}))$.

Identifiability makes the causal-effect query well-posed:

Definition 3 (Causal-effect query). *Given a causal diagram $\mathcal{G} =$*
 $\langle \mathbf{V} \cup \mathbf{U}, E \rangle$, data samples from an observational distribution $P(\mathbf{V})$,
and an identifiable query $P(\mathbf{Y} | do(\mathbf{X}))$, the task is to compute the
value of $P(\mathbf{Y} | do(\mathbf{X} = \mathbf{x}))$.

Estimand-based approaches. The now-standard methodology for
answering causal-effect queries is to break the task into two steps.
The first is the *identifiability step*: given a causal diagram and a
query $P(\mathbf{Y} | do(\mathbf{X}))$, determine if the query is identifiable and
if so, generate an *estimand*, or algebraic expression in terms of the
observational distribution $P(\mathbf{V})$ that answers the query. A complete
polynomial algorithm called ID has been developed for this task [27,
24]. The second step is *estimation*: use samples from the observational
distribution $P(\mathbf{V})$ to estimate the value of the estimand expression.

A number of approaches have been applied to the estimation in the
second step. A simple approach, called the *plug-in estimator*, replaces
each term in the estimand with its empirical probability value in the
observed data. More sophisticated approaches include the line of work
by Jung et al., which apply ideas from data weighting and empirical
risk minimization [12, 13] and double or debiased machine learning

[14, 15]. Evaluating the estimand’s value from a PAC perspective is analyzed in [4].

In practice, the estimand expression typically consists of multiple products or ratios of partially marginalized distributions. We illustrate with an example:

Example 1. Consider the model in Figure 1a. To evaluate the query $P(V_6 \mid do(V_0))$, the ID algorithm [27, 24] gives the expression:

$$\begin{aligned}
P(V_6 \mid do(V_0)) &= \sum_{V_1, V_2, V_3, V_4, V_5} P(V_5 \mid V_0, V_1, V_2, V_3, V_4) \\
&\times P(V_3 \mid V_0, V_1, V_2) P(V_1 \mid V_0) \sum_{v_0} P(V_6 \mid v_0, V_1, V_2, V_3, V_4, V_5) \\
&\times P(V_4 \mid v_0, V_1, V_2, V_3) P(V_2 \mid v_0, V_1) P(v_0). \quad (4)
\end{aligned}$$

Unfortunately, in large models the expression elements quickly become unwieldy; Example 1 involves estimating conditional distribution terms over the entire joint configuration space of the model. In terms of scalability, this presents several practical problems as model sizes increase. The required distributions have exponentially many configurations, suggesting they may require a significant amount of data to estimate accurately. Moreover, the required marginalizations are also over high-dimensional spaces, potentially also making them computationally intractable. These issues make it difficult to apply statistically sophisticated or machine learning-based estimators [12, 13, 14, 15] in such settings.

On the other hand, we can often maintain tractability by using the simple *plug-in* estimator, in which each term is estimated only on the configurations seen in the observed data. This reduces computation, since each term is non-zero on only a subset of configurations, and summations can also be performed over only the non-zero configurations. However, this approach also limits the quality of our estimates – our plug-in estimates may have high variance, especially in settings with very large probability tables, and it is difficult to apply regularization or other variance-reduction techniques without destroying the sparsity property that makes it computationally feasible.

These issues motivate us to explore the effectiveness of a *model completion* approach, in which we directly learn the Causal Bayesian Network model, including its latent variables, and then use the resulting learned model to evaluate the query. A common technique for learning latent variable models is the EM algorithm:

Expectation Maximization (EM). EM is a well-studied scheme to learn Bayesian network parameters when the graph itself is known but there are some latent variables, for which data are missing [9, 16]. EM is an iterative maximum likelihood approach that alternates between an Expectation, or E-Step, and a Maximization, or M-Step; see [19] for details. The E-step entails inference in the Bayesian network, which can be computationally hard in general; but when the graph has bounded *treewidth*¹, the E-step can be performed using join tree or bucket elimination algorithms [19, 8]. Otherwise, approximate algorithms such as belief propagation can be used [22, 20]. While EM has also been applied to structure learning [10], in our setting the structure is assumed known via the provided causal diagram.

Other related work. Motivated by similar model completion ideas, [7, 5] first convert the causal diagram to a circuit, exploiting the functional form of SCM mechanisms to yield compact circuits, then learn the circuit parameters via EM, and answer the causal queries. However, their work is restricted to binary-valued variables and, to the

best of our knowledge, no empirical evaluation is provided. Another related line of work explores neural network approaches for causal inference [30, 31]. In that work, the authors focus on *counterfactual* queries and propose learning a Neural Causal Model (NCM) whose functions are neural networks, a task which requires learning the mechanisms of the underlying SCM, which is significantly more challenging and unnecessary for our task of answering causal-effect queries. Moreover, while the authors provide a theoretical analysis, the empirical validation is only performed over very small, synthetic networks. EM has also been applied to causality in [26, 33]; in the latter, the authors assume knowledge of the functional mechanisms in the SCM, with only the distributions of the discrete latent variables unknown. In contrast, our method assumes knowledge of the graph only.

3 Learning-Based Causal Inference

The approach we propose for the causal-effect task is to first learn a full CBN $\mathcal{B} = \langle \mathcal{G}, \mathcal{P} \rangle$ given the causal diagram $\mathcal{G} = \langle \mathbf{V} \cup \mathbf{U}, E \rangle$ and samples from the observational distribution $P(\mathbf{V})$. Then, we can answer causal-effect queries based on the truncated formula Eq. (3) using probabilistic inference over the learned CBN. However, there could be many parameterizations \mathcal{P} that are consistent with the same observational distribution $P(\mathbf{V})$. Luckily, the identifiability property ensures that the problem remains well-posed: as long as the query is identifiable, any of these alternative parameterizations \mathcal{P} will generate the same answer:

Proposition 1. Assume a given SCM $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U}) \rangle$ having causal diagram \mathcal{G} and observational distribution $P(\mathbf{V})$. Any CBN $\mathcal{B} = \langle \mathcal{G}, \mathcal{P} \rangle$ inducing the same observational distribution $P(\mathbf{V})$ via Eq. (2) will agree with \mathcal{M} on any identifiable causal-effect query $P(\mathbf{Y} \mid do(\mathbf{X} = \mathbf{x}))$.

We note that if the query $P(\mathbf{Y} \mid do(\mathbf{X} = \mathbf{x}))$ is not identifiable, then CBNs $\mathcal{B} = \langle \mathcal{G}, \mathcal{P} \rangle$ with different parameters \mathcal{P} inducing the same observational distribution $P(\mathbf{V})$ may generate different answers to the query.

Consequently, our problem reduces to the well-studied task of estimating the parameters of a Bayesian network that includes latent variables from observational data, given the network structure. A widely used algorithm for this purpose is Expectation-Maximization (EM), which aims to maximize the likelihood of the data. To apply EM, however, we also require knowledge of the latent variables’ distributional forms.

Usefully, in the case of discrete visible variables V , prior work has shown that any SCM can be transformed into an equivalent SCM in which the latent variables U take on discrete, finite domains [34], and provides an upper bound on the required latent domain sizes. This allows us to assume, without loss of generality, a discrete distribution for the U_j .

However, the upper bound in [34] is conservative and often very large. Moreover, such large domain sizes are sufficient but may not be required for a given SCM. In practice, we use the discrete latent domain sizes as a complexity control mechanism, to impose an additional degree of regularity on the probability distribution over the visible variables. In order to select the appropriate amount of regularization, we take the domain sizes, $\mathbf{k} = \{k_{U_i}\}$, as hyperparameters and select their values by minimizing the model’s BIC score, which penalizes models with larger domain sizes for their increased flexibility and potential to overfit [6]:

$$BIC_{\mathcal{B}, \mathcal{D}} = -2 \cdot LL_{\mathcal{B}, \mathcal{D}} + p \cdot \log(|\mathcal{D}|) \quad (5)$$

¹ The treewidth is a graph parameter characterizing its sparseness; exact inference is exponential in the treewidth [8].

Algorithm 1: EM4CI

input : A causal diagram $\mathcal{G} = \langle U \cup V, E \rangle$, U latent and V observables; \mathcal{D} samples from $P(V)$;
Query $Q = P(Y | do(X = x))$.
output : Estimated $P(Y | do(X = x))$

// $k =$ latent domain size, $BIC_{\mathcal{B}} = BIC$ score of \mathcal{B} , \mathcal{D} ,
// $LL_{\mathcal{B}}$ is the log-likelihood of \mathcal{B} , \mathcal{D}

1. Initialize: $BIC_{\mathcal{B}} \leftarrow \infty$,
 2. If \neg identifiable(\mathcal{G}, Q), terminate.
 3. **for** $k = 2, \dots$, to upper bound, **do**
 4. $(\mathcal{B}', LL_{\mathcal{B}'}) \leftarrow \max_{LL} \{EM(\mathcal{G}, \mathcal{D}, k) | \text{for } i = 1 \text{ to } 10\}$
 5. Calculate $BIC_{\mathcal{B}'}$ from $LL_{\mathcal{B}'}$
 6. **if** $BIC_{\mathcal{B}'} \leq BIC_{\mathcal{B}}$,
 7. $\mathcal{B} \leftarrow \mathcal{B}'$, $BIC_{\mathcal{B}} \leftarrow BIC_{\mathcal{B}'}$
 8. **else**, break.
 9. **endfor**
 - 10: $\mathcal{B}_{X=x} \leftarrow$ generate truncated CBN from \mathcal{B} .
 - 11: **return** \leftarrow evaluate $P_{\mathcal{B}_{X=x}}(Y)$
-

244 where \mathcal{D} are the data samples from $P(V)$, $LL_{\mathcal{B}, \mathcal{D}}$ is the log likeli-
245 hood of the CBN model \mathcal{B} learned via EM, and p is the number of
246 free probability parameters θ in \mathcal{B} .

247 To optimize our CBN over both the domain sizes k and probability
248 parameters θ , we propose a practical algorithm that searches greedily
249 in the space of k while optimizing θ by EM. Larger domain sizes
250 can facilitate the learning of complex models, but require more sam-
251 ples and more time for convergence. For this reason, we prioritize
252 searching the model space starting from smaller domain sizes. At the
253 same time, having domain sizes slightly too large is unlikely to hurt
254 performance significantly, so we adopt a simple strategy of keeping
255 the latent domain sizes equal, i.e., set all $k_{U_i} = k$, and gradually
256 increase the value k .

257 Our learning strategy thus benefits from two sources of variance
258 reduction compared to the simple plug-in estimates: first, from the
259 graph structure itself, which imposes some regularity on $P(V)$, and
260 second, from preferring smaller latent domain sizes k when possible,
261 which encourages learning lower-rank distributions when supported
262 by the data.

263 Once a full CBN $\mathcal{B} = \langle \mathcal{G}, \mathcal{P} \rangle$ is learned, any causal query
264 $P(Y | do(X = x))$ can be answered based on Eq. (3). Let $\mathcal{G}_{\overline{X}}$
265 denote the graph obtained by deleting all incoming arrows to X .
266 Let $\mathcal{P}_{X=x} = \{P(V_i | PA_i), P(U)\}_{V_i \notin X} \cup \{P(X = x) = 1\}$. Let
267 $\mathcal{B}_{X=x} = \langle \mathcal{G}_{\overline{X}}, \mathcal{P}_{X=x} \rangle$ be the truncated CBN. Then we have:

Proposition 2. $P(Y | do(X = x))$ in the CBN \mathcal{B} is given by $P(Y)$
in the truncated CBN $\mathcal{B}_{X=x}$, i.e.,

$$P_{\mathcal{B}}(Y | do(X = x)) = P_{\mathcal{B}_{X=x}}(Y). \quad (6)$$

268 **The iterative learning EM4CI algorithm (Algorithm 1.)** Since
269 we use EM for learning we call our algorithm *EM for Causal Inference*
270 (*EM4CI*), described in Algorithm 1. Its input is a causal diagram \mathcal{G} ,
271 samples \mathcal{D} from $P(V)$, and a query $Q = P(Y | do(X = x))$.
272 After initialization, *Step 2* checks if the query is identifiable. This
273 can be accomplished by any variant of the well-known ID algorithm
274 [27, 24], or via the Do calculus [23]. If the query is not identifiable,
275 the algorithm terminates with failure. Otherwise *Steps 3-9* provide
276 the iterative learning scheme. Given samples \mathcal{D} from $P(V)$, we use
277 EM to learn a CBN \mathcal{B} over the graph \mathcal{G} with latent domain size k .




Figure 2: A subset of the small causal diagrams for models used in our experiments. Blue variables are intervened on and red variables are the outcome variables corresponding to the query $P(Y | do(X))$.

Table 1: Estimand Expressions for Models 1-8 in Fig. 2 & Table 3.

Model	Estimate of $P(Y do(X))$
1	$\frac{\sum_W P(X, Y, R, W) P(W)}{\sum_W P(X, R, W) P(W)}$
2	$\sum_R P(R X_1) \sum_{x_1, Z} P(Y R, x_1, X_2, Z) P(Z R, x_1) P(x_1)$
3	$P(Y)$
4	$\sum_{R, S} P(S X_1, X_2, X_3, Z) P(R X_1) \sum_{x_1, x_3, Z} P(Y R, x_1, X_2, x_3, Z) P(x_3 x_1, X_2, Z) P(x_1, Z)$
5	$\sum_{Z_1, Z_2, Z_3} P(Z_3 Z_2) P(Z_1 X, Z_2) \frac{\sum_x P(Y, Z_3 x, Z_1, Z_2) P(x, Z_2)}{\sum_x P(Z_3 x, Z_1, Z_2) P(x, Z_2)} P(Z_2)$
6	$\sum_{Z_2} P(Y X_1, X_2, Z_2) P(Z_2)$
7	$\sum_{Z_2, Z_3} P(Y X_1, X_2, Z_2, Z_3) P(Z_2, Z_3)$
8	$\sum_{R, W, Z} P(Z R, W, X) P(R W) \sum_x P(Y R, W, x, Z) P(x R, W) P(W)$

EM outputs the parameters for a network, \mathcal{B}' , and its corresponding
log-likelihood, $LL_{\mathcal{B}'}$, with which we calculate the $BIC_{\mathcal{B}'}$ score.
Since EM’s performance is known to be sensitive to initialization,
we perform ten runs of EM starting from randomly generated initial
parameters, and retain the model with the highest likelihood as the
candidate for the current latent domain size k . We then increase k ,
and repeat the process. If the BIC score decreases (step 6), we adopt
the new candidate network and the process continues with increased
 k ; otherwise the current \mathcal{B} is selected as the final learned network.
Step 10 generates the truncated CBN $\mathcal{B}_{X=x}$, and *Step 11* employs
a standard PGM inference algorithm, like join-tree decomposition
[6, 8], to the CBN $\mathcal{B}_{X=x}$ to compute $P_{X=x}(Y)$.

Complexity of EM4CI Determining identifiability is polynomial
in the graph size [27]. The complexity of EM (*Steps 3-9*) depends
on the sample size $|D|$, the variables’ domain sizes, and the num-
ber of iterations needed for convergence which is hard to predict,
but complexity-wise provide only a constant factor and can be ig-
nored. The most extensive computation in each iteration of EM is
the Expectation step, which requires probabilistic inference and can
be accomplished in time and memory exponential in the treewidth
of the graph. Otherwise, approximation algorithms such as belief
propagation may be used. In our experiments, exact inference was
always possible. Thus, each iteration of EM takes $O(|D| \cdot n l^w)$, where
 $n = |V \cup U|$ is the number of variables, l bounds the variable domain
sizes, and w is the treewidth. For T iterations we find that the com-
plexity of EM is $O(T \cdot |D| \cdot n \cdot l^w)$. *Step 11* of probabilistic inference
is $O(n \cdot l^w)$ if performing exact inference. In summary, the complexity
of algorithm *EM4CI* is $O(T \cdot |D| \cdot n \cdot l^w)$. Thus, the algorithm is most
effective for models with bounded treewidth. Note that the cost of the
EM learning process can be amortized over multiple queries.

4 Empirical Evaluation

We perform an extensive empirical evaluation of EM4CI and compare
against estimand-based approaches such as the brute-force empirical
plug-in method and a state-of-the-art scheme called weighted empiri-
cal risk minimization (WERM) [13]. All experiments were run on a
2.66 GHz processor with 8 GB of memory.

4.1 Benchmarks.

The inputs to a causal inference algorithm are a triplet: 1) a causal dia-
gram (of an underlying SCM), 2) data from the model’s observational
distribution, and 3) a query $Q = P(Y | do(X))$. We use two sources




Figure 3: Example larger synthetic models.

for our benchmarks: synthetically generated scalable model classes, and real domains from the academic literature of various fields with causal interpretations [1].

Synthetic networks. We generate each triplet benchmark instance by first choosing a causal diagram and a query. Then, picking domain sizes of observed and latent variables, we generate the conditional probability tables (CPTs), one per variable and its parents, to yield the full CBN. Given the full model we generate samples from its observational distribution by forward sampling [6] over all variables, then discard the values of the latent variables. We compute the exact answer to the target query on the full CBN via an exact algorithm (e.g., join-tree [6, 8]).

The causal diagrams were selected in two ways. First, we take a set of 8 small diagrams from the causality literature [23, 14]; see Figure 2 and the Appendix. Second, we examine three scalable classes of graphs whose treewidths vary but can be controlled. These are: chain networks (treewidth 2; Figure 3c), diamond networks (treewidth 5; Figure 3a), and cone-cloud networks (treewidth $O(\sqrt{n})$, Figure 3b), abbreviated CH, D, CC respectively.

In the small diagrams (e.g., Figure 2), we set the domain size of the observed variables to $d = 2$, and of the latent variables to $k = 10$. For the chain, diamond, and cone-cloud models, we use $(d, k) = (4, 10)$. The parameters of each CPT were generated by sampling from a Dirichlet distribution [6] with parameters $\alpha = [\alpha_1, \dots, \alpha_n]$. We chose $\alpha_i \in [0, 1]$ uniformly at random in order to generate conditional probabilities near the edges of the simplex (i.e., far from uniform).

Use-case benchmarks. We also test on four networks created for real-world domains. The “Alarm” network was developed for on-line monitoring of patients in intensive care units [2]; the “A” network, from the UAI literature, is synthetic but known to be daunting to exact algorithms [17]; the “Barley” network was built for a decision support system for growing barley without pesticides [18]; and the “Win95” model is an expert system for printer troubleshooting in Windows 95. Since no variables in these networks are specified as latent, we chose source vertices with at least two children to be latent confounding variables. We generated data by forward sampling and then discarding the latent variables’ values.

Queries. We hand selected identifiable queries, aiming for non-trivial ones and prioritizing those that appear to be complex. The queries and their corresponding estimands are given in the Appendix. We selected several identifiable queries per model to better evaluate

Table 2: Algorithms and Packages

Algorithm	Use	Software Package	Programming Language
ID Algorithm	Plug-In & EM4CI	<i>causaleffect</i>	R
Join-tree Decomposition	EM4CI inference	SMILE	C++
EM Algorithm	EM4CI learning	SMILE	C++
Sparse Plug-In Evaluation	Plug-In		Python
WERM			R

Table 3: Results of EM4CI & Plug-In on $P(Y | do(\mathbf{X}))$ ($d, k) = (2, 10)$

Model	k_{lrm}	100 Samples				1,000 Samples				
		EM4CI mad	EM4CI time(s)	Plug-In mad	Plug-In time(s)	EM4CI mad	EM4CI time(s)	Plug-In mad	Plug-In time(s)	
1	2	0.0037	0.4759	0.0104	1.9	2	0.0032	3.1	0.0025	2.3
2	2	0.1832	1.8643	0.1436	2.3	2	0.0490	8.4	0.0867	2.0
3	2	0.1288	0.9288	0.0569	1.1	2	0.0040	3.6	0.0039	0.7
4	2	0.1819	1.8169	0.1469	2.3	2	0.1438	12.0	0.0704	2.1
5	2	0.4910	1.6539	0.5000	2.0	2	0.0044	17.3	0.0058	2.2
6	2	0.2663	0.3004	0.3930	2.0	2	0.1263	0.5	0.1319	2.1
7	2	0.2520	0.7757	0.2509	1.9	2	0.0891	7.1	0.0238	2.0
8	2	0.1372	0.6348	0.1579	2.0	2	0.2340	4.7	0.1303	1.9

the methods, as well as to emphasize that the learning time can be amortized over multiple queries.

4.2 Algorithms and performance measures.

EM4CI In our EM4CI algorithm we used the EM implementation of SMILE: Structural Modeling, Inference, and Learning Engine package [1], written in C++. In SMILE, inference is carried out via join trees [22, 19], and used both in the E step of the EM algorithm [19] and for answering queries over the learned model.

Estimand-based algorithms: Plug-In and WERM. We used the ID algorithm in the causal effect package [29] to compute the estimands. Subsequently, the plug-in method evaluates the estimand directly from the observational data, by plugging in the empirical conditional probabilities. Our plug-in implementation uses sparse table representations in Python, to ensure that the representation size of each estimated probability term is linear in the data size, rather than exponential in the number of variables. We also compare against a state-of-the-art scheme called WERM [13] using the author-provided implementation in R. A summary of the algorithms and implementations is given in Table 2.

Measures of performance. We report the results of EM4CI along the two phases of the algorithm; the *learning phase* (steps 3-9) done using the EM algorithm and the *inference phase* (steps 10-11) done using the join tree algorithm. For the learning process, we report the selected latent domain sizes and the total time at termination. For the inference phase we report the *mean absolute deviation (mad)* between the true answer and the estimated answer. The measure *mad* for a query $P(Y | do(\mathbf{X}))$ is computed by averaging the absolute error over all single-value queries over all instantiations of the intervened and queried variables, $P(Y = y | do(\mathbf{X} = \mathbf{x}))$ for $\mathbf{x}, y \in \mathcal{D}(\mathbf{X}) \times \mathcal{D}(Y)$. We also report average inference time over the same set of query instantiations. In our experiments we increased k by twos ($k = 2, 4, 6, 8, 10 \dots$) in the EM4CI algorithm.

4.3 Results

Results on small synthetic models. Results on models 1-8 (Figure 2) are presented in Table 3. Since these models are quite small we report the total time (learning plus inference) for EM4CI. We compare against the *Plug-In* estimand method (see estimand expressions in Table 1), which is guaranteed to converge to the exact answer. Therefore, we expect *Plug-In* to produce fairly accurate results on these small models if given enough samples. We observe that the accuracy of both methods are similar at both 100 and 1,000 samples, with EM4CI being more accurate on some cases, and *Plug-In* on others. EM4CI

Table 4: Results of absolute error on Query $P(Y = 1|do(\mathbf{X} = 1))$ on Models 1, 8, & 3' by WERM and EM4CI. $k_{l_{rn}}$ is the learned domain sizes of latent variables.

Model	1,000 Samples					10,000 Samples				
	WERM error	WERM time(s)	EM4CI error	EM4CI time(s)	$k_{l_{rn}}$	WERM error	WERM time(s)	EM4CI error	EM4CI time(s)	$k_{l_{rn}}$
1	0.0071	18.7	0.0059	8.8	2	0.0031	32.6	0.0046	63.5	2
8	0.1082	25.8	0.1566	7.6	2	0.11	47.7	0.0001	81.4	2
3'	0.027	27.2	0.0004	3.5	2	0.001	44.1	0.0009	53.1	2

Table 5: Results for EM4CI and Plug-In on $P(Y|do(\mathbf{X}))$ (d, k) = (4, 10)

(a) 1,000 samples

Model	Query	$k_{l_{rn}}$	EM4CI		Plug-In		
			mad	Learn-time(s)	mad	time(s)	
5-CH	$P(V_4 do(V_6))$	4	0.0902	3.5	0.0001	0.1509	2.3
9-CH	$P(V_8 do(V_6))$	4	0.1204	11.5	0.0002	0.1516	2.4
25-CH	$P(V_2 do(V_6))$	2	0.0070	77.7	0.0003	0.0959	6.1
49-CH	$P(V_4 do(V_6))$	4	0.0005	161.2	0.0007	0.0319	17.8
99-CH	$P(V_8 do(V_6))$	6	0.0093	413.4	0.0023	0.0611	88.1
9-D	$P(V_8 do(V_6))$	2	0.0719	24.6	0.0002	0.1832	3.4
17-D	$P(V_{16} do(V_6))$	6	0.0542	202.3	0.0006	0.0700	4.5
65-D	$P(V_4 do(V_6))$	4	0.0074	432.4	0.0012	0.1716	232.5
6-CC	$P(V_6 do(V_6))$	4	0.0088	23.5	0.0001	0.0156	2.3
15-CC	$P(V_6 do(V_{14}))$	4	0.0147	60.8	0.0001	0.0659	4.5
45-CC	$P(V_6 do(V_{14}, V_{36}, V_{44}))$	6	0.0097	199.2	2.7429	0.1509	18.6

(b) 10,000 samples

Model	Query	$k_{l_{rn}}$	EM4CI		Plug-In		
			mad	Learn-time(s)	mad	time(s)	
5-CH	$P(V_4 do(V_6))$	4	0.0508	17.3	0.0001	0.0537	2.5
9-CH	$P(V_8 do(V_6))$	4	0.0236	150.0	0.0002	0.1074	3.1
25-CH	$P(V_2 do(V_6))$	6	0.0068	697.1	0.0005	0.0714	26.4
49-CH	$P(V_4 do(V_6))$	10	0.0017	2412.6	0.0036	0.0160	133.7
99-CH	$P(V_8 do(V_6))$	6	0.0028	3887.9	0.0022	0.0433	850.6
9-D	$P(V_8 do(V_6))$	4	0.0611	390.7	0.0002	0.1481	3.0
17-D	$P(V_{16} do(V_6))$	6	0.0360	1849.6	0.0007	0.0582	8.4
65-D	$P(V_4 do(V_6))$	4	0.0022	4787.2	0.0013	0.1376	2258.5
6-CC	$P(V_6 do(V_6))$	6	0.0138	116.9	0.0003	0.0136	2.7
15-CC	$P(V_6 do(V_{14}))$	4	0.0022	489.5	0.0043	0.0321	10.9
45-CC	$P(V_6 do(V_{14}, V_{36}, V_{44}))$	6	0.0026	1833.7	2.757	0.1561	105.8

had better time performance for 100 samples, but for 1,000 samples the *Plug-In* was faster since learning time of EM4CI was longer.

WERM comparison. The results for Models 1, 8, and 3' comparing WERM [13] to EM4CI are given in Table 4. We use domain size $d = 2$, and 1,000 and 10,000 samples. Again, EM4CI produced more accurate results in several instances, though the disparities are smaller than with the *Plug-In* method. EM4CI was faster than WERM with 1,000 samples but slower with 10,000 samples.

Unfortunately, the available code for WERM is specific to these models, so we were unable to compare against it more generally. This also highlights a general lack of available estimand-based implementations applicable to general settings.

Results on large synthetic models. In Tables 5 and 6 we show results on larger models of chains, diamonds, and cone-cloud graphs. The first two tables compare EM4CI to the *plug-in* method for a single query over all the models. Specifically, Table 5a presents time and accuracy results for 1,000 samples. We see that EM4CI was consistently more accurate, and in many cases significantly better (e.g., in 45-cone-cloud). We see the same trend for 10,000 samples in Table 5b. The superior accuracy of learning for model completion compared with the *Plug-In* is visualized well in Figure 4, which shows the accuracy trends for each class as a function of model sizes, for both 1,000 and 10,000 samples. We can see that generally, EM4CI using only 1,000 samples is even more accurate than *Plug-In* with 10,000 samples.

Again, we expect this improvement is due to the variance reduction of the estimation process. By its nature, model completion exploits more information from the causal graph than is apparent in the estimand expression alone, implicitly capturing the known conditional independence structure in its estimates of $P(\mathbf{V})$. We can also include

² model 3' is the same as Model 3 in Figure 2, but with edge $Y \rightarrow Z$ reversed to match the hard-coded model in WERM.

Table 6: EM4CI & Plug-In on the 45 Cone Cloud $|\mathbf{V}| = 45$; $|\mathbf{U}| = 16$; $d = 4$; $k = 10$; treewidth ≈ 11

(a) Plug-In

Query	1,000 Samples		10,000 Samples	
	mad	time(s)	mad	time(s)
$P(V_0 do(V_{14}, V_{36}, V_{44}))$	0.1510	18.6	0.1561	105.8
$P(V_0 do(V_{12}, V_{29}, V_{34}))$	0.1744	19.0	0.1043	100.98
$P(V_0 do(V_{10}, V_{14}, V_{40}))$	0.1906	18.4	0.1201	131.4
$P(V_4 do(V_{15}, V_{20}, V_{40}))$	0.2080	26.2	0.1282	121.6

(b) EM4CI

Learning:	1,000 Samples		10,000 Samples	
	time = 199.2s	$k_{l_{rn}} = 6$	time = 1833.7s	$k_{l_{rn}} = 6$
Inference:				
Query	mad	time(s)	mad	time(s)
$P(V_0 do(V_{14}, V_{36}, V_{44}))$	0.0097	2.7682	0.0026	2.7566
$P(V_0 do(V_{12}, V_{29}, V_{34}))$	0.0118	18.9881	0.0024	18.9896
$P(V_0 do(V_{10}, V_{14}, V_{40}))$	0.0097	0.0590	0.0029	0.0577
$P(V_4 do(V_{15}, V_{20}, V_{40}))$	0.0082	0.0175	0.0013	0.0164

simple and easy-to-impose complexity control, in the form of the latent domain sizes, to further reduce variance. In contrast, it is difficult to impose meaningful regularity or variance reduction on the *plug-in* estimates of $P(\mathbf{V})$ without creating significant computational burdens. Thus our results suggest that whenever the causal graph's treewidth is bounded and the estimand expression has large scope functions, we should prefer using model completion. Other settings may require more study to determine the best approach.

Focusing on time, we see that the time of EM4CI is significantly more costly than *Plug-In*, with EM learning being the most time consuming part. Interestingly, while time grows with model size for both schemes, the inference time component remains efficient, likely due to the low treewidth of some of the models (e.g., the chains and the diamonds). In the 45-cone case, inference time is impacted more by model size, since its treewidth increases with the square root of the number of variables. Generally, for a single query, we find *Plug-In* has better time performance, and its time increases at a slower pace. In both methods, time also increases with sample size, e.g., when moving from 1,000 samples table to 10,000 samples.

Results on real-world data sets. Table 7, compares EM4CI against *Plug-In* on a single query for all 4 networks. We observe the same pattern of performance as in the synthetic networks: EM4CI provides far more accurate results on all these models but with some time cost, attributed to learning the models. Evaluation for multiple queries for EM4CI and *Plug-In* are given in Table 8 for the **A network**. We find the same latent domain size ($k_{l_{rn}} = 4$) is selected for both sample sizes. As before, learning time grows with sample size. Accuracy results are excellent and improve with increased sample sizes as

Table 7: Results on EM4CI & Plug-In on Real Networks

(a) 1,000 Samples

Model	Query	$k_{l_{rn}}$	mad	EM4CI		Plug-In	
				Learn-time(s)	inf-time(s)	mad	time(s)
A	$P(V_{51} do(V_{10}))$	4	0.0139	71	0.0012	0.0584	8
Alarm	$P(HRB P(Shunt))$	2	0.0076	15.7	0.0002	0.0190	2.9
Barley	$P(\text{protein} do(\text{expYield}))$	14	0.0066	199	0.9038	0.0687	9.9
Win95	$P(\text{Output} do(NnPsGrphc))$	2	0.0113	109.1	0.0002	0.2674	1.5

(b) 10,000 samples

Model	Query	$k_{l_{rn}}$	mad	EM4CI		Plug-In	
				Learn-time(s)	inf-time(s)	mad	time(s)
A	$P(V_{51} do(V_{10}))$	4	0.0083	540.6	0.0012	0.0114	55.7
Alarm	$P(HRB P(Shunt))$	4	0.0043	181.3	0.0002	0.0075	5.7
Barley	$P(\text{protein} do(\text{expYield}))$	10	0.0031	819.5	0.8882	0.0655	63.4
Win95	$P(\text{Output} do(NnPsGrphc))$	4	0.0008	1080.8	0.0002	0.2894	2.0




Figure 4: Comparing the accuracy of EM4CI and Plug-In. While both methods improve with more samples (solid to dashed lines), the error (*mad*) of EM4CI is smaller, even when compared to Plug-In with more samples.

well. Note that increased sample size improves accuracy but does not impact inference time, although more data may allow selection of a more complex model (larger latent domain size), which may also improve inference accuracy but may require more time. Results for **Alarm**, **Win95**, and **Barley** networks are similar (see Appendix).

Answering multiple queries. Tables 6b and 8b highlight how learning time of EM4CI can be amortized effectively over multiple queries. The table presents EM4CI’s learning and inference time separately on the multiple queries. In contrast, the *Plug-In* method requires estimation of each new query from scratch, even if on the same model; see Table 6a and 8a. Thus, for multiple queries, EM4CI may take far less time per query, while providing superior quality estimates. (See additional results on multiple queries in the Appendix.)

Summary. Our experiments illustrate the strength of the direct-learning approach as a viable alternative for answering causal effect queries. We saw that model-completion by learning implemented in our EM4CI yields highly accurate estimates of causal effect queries on a variety of benchmarks, both synthetic networks and real life networks, on both small and large models. Moreover EM4CI shows clear superiority to the Plug-In estimands approach. In terms of time efficiency however, EM4CI was consistently slower due to learning time overhead. Yet, when answering multiple queries is desired the time overhead can be amortized over multiple queries.

Table 8: Plug-In & EM4CI results on the **A Network** $|V| = 46$; $|U| = 8$; $d = 2$; $k = 2$ treewidth ≈ 16

(a) Plug-In				
Query	1,000 Samples		10,000 Samples	
	mad	time(s)	mad	time(s)
$P(V_{51} \text{do}(V_{10}))$	0.0584	8.0	0.0114	55.7
$P(V_{51} \text{do}(V_{14}))$	0.0319	8.3	0.0056	51.3
$P(V_{51} \text{do}(V_{41}))$	0.0255	13.9	0.0092	48.3
$P(V_{51} \text{do}(V_{45}))$	0.0496	9.8	0.0206	49.1

(b) EM4CI				
Learning	1,000 Samples		10,000 Samples	
	time = 71(s)	$k_{lrm} = 4$	time(s) = 541	$k_{lrm} = 4$
Inference:				
Query	mad	time(s)	mad	time(s)
$P(V_{51} \text{do}(V_{10}))$	0.0139	0.0012	0.0083	0.0012
$P(V_{51} \text{do}(V_{14}))$	0.0143	0.0047	0.0086	0.0046
$P(V_{51} \text{do}(V_{41}))$	0.0147	0.0042	0.0079	0.0041
$P(V_{51} \text{do}(V_{45}))$	0.0140	0.0031	0.0082	0.0030

5 Conclusion

In causal inference, the estimand-based approach has become standard: generating a potentially complex expression in terms of the visible distribution, and then estimating the required probabilities from observed data and evaluating the expression. While mathematically correct, this approach tends to ignore the difficulties inherent in estimating the complex, conditional probabilities required, and in computationally evaluating the resulting marginalized expression. These difficulties become increasingly apparent as model size increases.

An alternative path, relatively less explored, is to leverage the causal model structure more directly via learning. By performing model completion – i.e., learning a Causal Belief Network, including its latent confounding variables, from the observed data – we exploit additional information about the co-dependence structure in the distribution, and can more easily apply complexity control and variance reduction strategies to the parameter estimation process. Then, once an approximate full model is available, traditional computationally efficient PGM algorithms can be applied to answer the query, either exactly or approximately. In settings where multiple queries are desired, their cost can be amortized by performing the learning process once.

Our algorithm EM4CI uses the well-known EM algorithm to learn the model and its latent variable distributions, then uses tree-decomposition algorithms to answer the queries.

We carried out an extensive empirical evaluation, the first of its scale in the causal community, over a collection incorporating both synthetic networks and real world distributions. We evaluated and compared EM4CI’s performance to the plug-in estimand approach, in terms of both accuracy and time.

Our results appear conclusive: we show clearly that the model completion using EM4CI yields consistently superior results compared to the estimand plug-in. This benefit does come with a cost in time from the learning phase, which grows more quickly with networks size and number of samples compared to the plug-in approach. However as we show, learning time can be amortized over multiple queries on the learned models, making a collection of queries significantly more efficient.

Our EM4CI approach relies in part on the efficiency of inference; when the treewidth of the graph is bounded, both learning and inference using EM4CI are likely to be effective. Additionally, since the structure of the graph is better exploited by model completion, we reduce the variance of our estimators, resulting in better performance for a given dataset size. Although in some cases, the estimand approach may generate a simple and easy-to-estimate expression, in larger models the expression is often complex. In such settings, the estimand expression loses structural information, forcing us to estimate high-dimensional probabilities and making it difficult to apply

528 variance reduction strategies without creating a computationally diffi-
529 cult evaluation problem. Overall this suggests that model completion
530 should be considered a competitive strategy for causal estimation.

References

- [1] BayesFusion LLC. *QGeNIe Modeler USER MANUAL*, 2022.
- [2] I. Beinlich, H. Suermondt, R. Chavez, and G. Cooper. The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks. In *Proceedings of the Second European Conference on Artificial Intelligence in Medical Care*, pages 247–256. Springer-Verlag, Berlin, 1989.
- [3] R. Bhattacharya, R. Nabi, and I. Shpitser. Semiparametric inference for causal effects in graphical models with hidden variables. *Journal of Machine Learning Research*, 23:1–76, 2022.
- [4] A. Bhattacharyya, S. Gayen, S. Kandasamy, V. Raval, and V. N. Variyam. Efficient interventional distribution learning in the pac framework. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pages 7531–7549, 2022.
- [5] Y. Chen and A. Darwiche. On the definition and computation of causal treewidth. In J. Cussens and K. Zhang, editors, *Uncertainty in Artificial Intelligence, UAI, 2022*, 2022.
- [6] A. Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, 2009.
- [7] A. Darwiche. Causal inference using tractable circuits. *CoRR*, 2022.
- [8] R. Dechter. *Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2013.
- [9] 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, Series B.*, 39:1–38, 1977.
- [10] N. Friedman. The Bayesian structural EM algorithm. In *Uncertainty in Artificial Intelligence (UAI)*, pages 129–138, 1998.
- [11] Y. Huang and M. Valtorta. On the completeness of an identifiability algorithm for semi-Markovian models. *Annals of Mathematics and Artificial Intelligence*, 54(4):363–408, 2008.
- [12] Y. Jung, J. Tian, and E. Bareinboim. Estimating causal effects using weighting-based estimators. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 10186–10193, 2020.
- [13] Y. Jung, J. Tian, and E. Bareinboim. Learning causal effects via weighted empirical risk minimization. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33*, 2020.
- [14] Y. Jung, J. Tian, and E. Bareinboim. Estimating identifiable causal effects through double machine learning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*, pages 12113–12122, 2021.
- [15] Y. Jung, J. Tian, and E. Bareinboim. Estimating identifiable causal effects on markov equivalence class through double machine learning. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, 2021.
- [16] N. Koller, Daphne; Friedman. *Probabilistic Graphical Models*. The MIT Press, 2009.
- [17] A. V. Kozlov and J. P. Singh. Parallel implementations of probabilistic inference. *IEEE Computer*, pages 33–40, 1996.
- [18] K. Kristensen and I. Rasmussen. The use of a Bayesian network in the design of a decision support system for growing malting barley without use of pesticides. *Computers and Electronics in Agriculture*, 33:197–217, 2002.
- [19] S. L. Lauritzen. The EM algorithm for graphical association models with missing data. *Computational Statistics & Data Analysis*, pages 191–201, 1995.
- [20] R. Mateescu, K. Kask, V. Gogate, and R. Dechter. Join-graph propagation algorithms. *J. Artif. Intell. Res.*, 37:279–328, 2010.
- [21] Microsoft, Inc. win95: an expert system for trouble shooting, 1995.
- [22] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann, 1989.
- [23] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.
- [24] I. Shpitser and J. Pearl. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence*, page 1219, 2006.
- [25] I. Shpitser and J. Pearl. Identification of conditional interventional distributions. In *Uncertainty in Artificial Intelligence*, pages 442–449, 2006.
- [26] I. Shpitser, T. S. Richardson, and J. M. Robins. An efficient algorithm for computing interventional distributions in latent variable causal models. *CoRR*, abs/1202.3763, 2012.
- [27] J. Tian. *Studies in Causal reasoning and Learning*. PhD thesis, University of California, Los Angeles, 2002.
- [28] 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 the Eighteenth National Conference on Artificial Intelligence (AAAI)*, pages 567–573, 2002.
- [29] S. Tikka. *Package 'causaleffect'*. CRAN, 2022. URL <https://github.com/santikka/causaleffect/>. Version 1.3.15, October 12, 2022.
- [30] K. Xia, K. Lee, Y. Bengio, and E. Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. In M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, editors, *NeurIPS 34*, pages 10823–10836, 2021.
- [31] K. Xia, Y. Pan, and E. Bareinboim. Neural causal models for counterfactual identification and estimation. *CoRR*, abs/2210.00035, 2022.
- [32] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Generalized belief propagation. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *NIPS*, 2000.
- [33] M. Zaffalon, A. Antonucci, and R. Cabañas. EM based bounding of unidentifiable queries in structural causal models. *Why-21 workshop at NeurIPS*, 2021.
- [34] J. Zhang, J. Tian, and E. Bareinboim. Partial counterfactual identification from observational and experimental data. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato, editors, *International Conference on Machine Learning, ICML 2022*, volume 162 of *Proceedings of Machine Learning Research*, pages 26548–26558. PMLR, 2022.