Path Dependent Structural Equation Models

Ranjani Srinivasan[1], Jaron Lee[2], Rohit Bhattacharya[2], Narges Ahmidi[2,3], Ilya Shpitser[2]

[1]Electrical and Computer Engineering, Johns Hopkins University

[2]Computer Science, Johns Hopkins University

[3]Institute of Computational Biology, Helmholtz Munich Center

Corresponding author: Ranjani Srinivasan (ranjani@jhu.edu)

November 11, 2020

Abstract

Causal analyses of longitudinal data generally assume that the qualitative causal structure relating variables remains invariant over time. In structured systems that transition between qualitatively different states in discrete time steps, such an approach is deficient on two fronts. First, time-varying variables may have *state-specific* causal relationships that need to be captured. Second, an intervention can result in state transitions downstream of the intervention different from those actually observed in the data. In other words, interventions may counterfactually alter the subsequent temporal evolution of the system. We introduce a generalization of causal graphical models, Path Dependent Structural Equation Models (PDSEMs), that can describe such systems. We show how causal inference may be performed in such models and illustrate its use in simulations and data obtained from a septoplasty surgical procedure.

1 Introduction

Many scientific questions and engineering tasks may only be approached by analyzing the behavior of a system over time. Understanding long term health risk factors (Belanger et al., 1978), trajectory tracking (Richards, 2005), speech recognition (Rabiner, 1989), game playing (Thrun, 1995), all require modeling the temporal evolution of a system. Many models for longitudinal or time series data, such as hidden Markov models or Kalman filters, are graphical models, and most may be viewed as dynamic Bayesian networks (DBNs) (Murphy, 2012). These models are used to predict the future evolution of systems, or find latent structures that best explain observations. Despite their complexity and usefulness, these models deal with fundamentally *associative* relationships. However, testing empirical hypotheses or providing decision support tools in complex domains that vary over time often requires *causal* modeling.

Causal models previously used for such tasks include graphical causal models (Pearl, 2009), marginal structural models (Robins, 1997), structural nested models (Robins, 1999), as well as models of counterfactual regret (Murphy, 2003; Chakraborty and Moodie, 2013). However, these models have generally assumed an invariant causal structure over time. For example, analysis of the impact of anti-retroviral therapy on HIV infection progression in observational studies assumed the same variables relevant for the patient health and the same causal relationships linking them at each time point in the study (Hernán et al., 2000). Changes tracked over time (such as HIV developing resistance to the current drug) are thus *quantitative*, with the underlying causal structure remaining unchanged over time. However, many systems undergo *qualitative* changes as well, where observability, relevance, and causal relationships of variables vary over time.

Consider the task of modeling surgical procedures to make informed decisions on resident surgeon training. Surgeries are often divided into discrete stages, each with an intermediate goal (Ahmidi et al., 2015). Each stage is associated with a distinct set of variables and relationships among them that may not be shared across stages. For instance, stitching together a previously made incision is a routine task requiring few tools that may be executed by a surgical robot, while reconstructing cartilage may require multiple tools, high surgical skill and manual dexterity. Another feature of surgeries is that procedures performed at a particular stage can go wrong, forcing surgeons to "double back" to correct mistakes, or deal with complications. Surgeon experience may often determine whether previous stages of the surgery are revisited. The goal of causal inference in this setting is to help assign surgeons to perform different stages of the surgery while navigating the tradeoff between the need to train resident surgeons on the one hand, and operating costs and patient safety on the other.

Addressing this tradeoff entails using retrospective data to estimate outcomes of surgery trajectories that *differ from those actually observed* due to counterfactually different choices of surgeon assignment in past stages of the surgery. Following the convention in the economics literature, we call the phenomenon where the evolution of a system changes in response to counterfactually different past choices *path dependence* (Liebowtiz and Margolis, 2002).

Other examples of path dependence include life course studies examining economic disparities in society or patient outcomes in hospitals using Electronic Health Record (EHR) data. Whether and where subjects went to elementary and high school, college, and their place of work, all represent qualitatively different stages of subjects' lives, with correspondingly

different variables and causal relationships among them. In addition, subjects might re-enroll in school, or otherwise revisit different stages in their life. Finally, life outcomes and life trajectories depend on counterfactually different life choices in the past, such as the choice to not enroll in college. Similar arguments can be made for patients being assigned to different parts of a hospital, such as the emergency room, the ward, the intensive care unit, or the operating room. Causally relevant variables and relationships between them differ drastically depending on which part of the hospital patients find themselves in. In practice, patients may revisit above parts of the hospital multiple times, often in response to various treatment decisions made. Counterfactually different decisions would result in potentially different patient trajectories.

In this paper, we introduce the *path-dependent structural equation model (PDSEM)* for causal systems that exhibit qualitative changes over time, observed or unobserved confounding, *and* path-dependence on counterfactual choices in the past. Our model can be viewed as a generalization of causal dynamic Bayesian networks, that allows complex and repeating stage transitions, and distinct causal models at each stage, or as a generalization of Markov decision processes (MDPs) that uses causal models in each state. In our formulation, each state is modeled explicitly as a (possibly hidden variable) causal model, and where observed actions are not chosen by the optimizer as in reinforcement learning problems, but instead determined by the data generating process that may contain unobserved confounders. As a result, our model combines complex and potentially looping state transitions of MDPs, and complex relationships among variables of a causal model. PDSEMs may also be viewed as a generalization of a Markov chain endowed with graphical causal model semantics, which allows handling of confounding and analysis of counterfactual state transitions.

2 Background

We first introduce necessary causal modeling ideas and previous work, before extending them to allow path-dependence.

2.1 Statistical and Causal Graphical Models

The statistical model of a directed acyclic graph (DAG) $\mathcal{G}(\mathbf{V})$ with a vertex set $\mathbf{V} \equiv \{V_1, \ldots, V_k\}$, called a *Bayesian network*, is the set of distributions that factorize with respect to the DAG as $p(\mathbf{V}) = \prod_{V_i \in \mathbf{V}} p(V_i \mid \mathrm{pa}_{\mathcal{G}}(V_i))$ where $\mathrm{pa}_{\mathcal{G}}(V_i)$ are parents of V_i in $\mathcal{G}$.

Causal models of a DAG are also sets of distributions but on *counterfactual* random variables. Each variable V_i in a causal model is determined from values of its parents $\mathrm{pa}_{\mathcal{G}}(V_i)$ and an exogenous noise variable ϵ_i via an invariant causal mechanism called a *structural equation* $f_i(\mathrm{pa}_{\mathcal{G}}(V_i), \epsilon_i)$. Causal models allow counterfactual intervention operations, denoted by the do($\mathbf{a}$) operator in (Pearl, 2009). Such operations replace each structural equation $f_i(\mathrm{pa}_{\mathcal{G}}(V_i), \epsilon_i)$ for $V_i \in \mathbf{A} \subset \mathbf{V}$ by one that sets V_i to a constant value in $\mathbf{a}$ corresponding to V_i. The joint distribution of variables in $\mathbf{Y} \equiv \mathbf{V} \setminus \mathbf{A}$ after the intervention
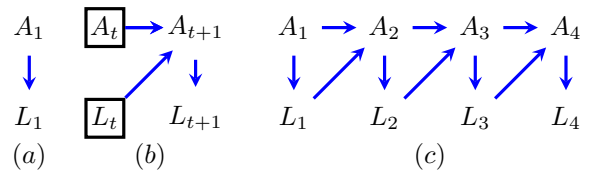

Figure 1: (a) Prior network DAG $\mathcal{G}_0$, representing the state of the dynamic Bayesian network at time $t = 0$. (b) A conditional DAG $\mathcal{G}_{t,t+1}$ representing the transitions in a dynamic Bayesian network. (c) A dynamic Bayesian network model unrolled to four time steps.

do($\mathbf{a}$) was performed is denoted by $p(\mathbf{Y} \mid \mathrm{do}(\mathbf{a}))$, equivalently written as $p(\{V_i(\mathbf{a}) : V_i \in \mathbf{Y}\})$, or $p(\mathbf{Y}(\mathbf{a}))$, where $V_i(\mathbf{a})$ is a counterfactual random variable or a potential outcome.

A popular causal model called the *non-parametric structural equation model with independent errors (NPSEM-IE)* (Pearl, 2009) assumes, aside from the structural equations for each variable being functions of their parents in the DAG $\mathcal{G}(\mathbf{V})$, that the joint distribution of all exogenous terms are marginally independent: $p(\epsilon_1, \epsilon_2, \ldots) = \prod_{V_i \in \mathbf{V}} p(\epsilon_i)$. The NPSEM-IE implies the DAG factorization of $p(\mathbf{V})$ with respect to $\mathcal{G}(\mathbf{V})$, and a truncated DAG factorization known as the *g-formula*:

$$p(\mathbf{Y}(\mathbf{a})) = \prod_{V_i \in \mathbf{Y}} p(V_i \mid \mathrm{pa}_{\mathcal{G}}(V_i))|_{\mathbf{A}=\mathbf{a}} \tag{1}$$

for every $\mathbf{A} \subseteq \mathbf{V}$, and $\mathbf{Y} = \mathbf{V} \setminus \mathbf{A}$.

2.2 Graphical Models In Discrete Time

While Bayesian networks lend themselves well to the modeling of static data, dynamic data with temporal evolution requires more sophisticated models. A generalization of the Bayesian network model for discrete time temporal systems is the *dynamic Bayesian network (DBN)* model (Murphy, 2012), which captures relationships between variables across time.

A DBN is specified by a pair of DAGs, and a corresponding pair of factorized distributions. The *prior network* $\mathcal{G}_1$ and its corresponding distribution $p(\mathbf{V}_1) = \prod_{V_i \in \mathbf{V}_1} p(V_i \mid \mathrm{pa}_{\mathcal{G}_1}(V_i))$ represent the state of the system at the first time step. The *transition network* $\mathcal{G}_{t,t+1}$ is a *conditional DAG (CDAG)* with random vertices $\mathbf{V}_{t+1}$ representing variables at time point $t+1$, and fixed vertices $\mathbf{V}_t$ representing context in the previous time point t. We will describe such conditional graphs by a shorthand "$\mathcal{G}_{t+1,t}$ on $\mathbf{V}_{t+1}$ given $\mathbf{V}_t$." In this conditional DAG no arrowheads into vertices in $\mathbf{V}_t$ are allowed. The corresponding conditional distribution $p(\mathbf{V}_{t+1} \mid \mathbf{V}_t)$ represents the way variables at point $t + 1$ depend on each other, and on variables at the prior time point t (and on no other prior variables, such as those at time point $t - 1$). This dependence leads to a *first order Markov* DBN. This distribution factorizes with respect to the CDAG $\mathcal{G}_{t,t+1}$ as follows: $p(\mathbf{V}_{t+1} \mid \mathbf{V}_t) = \prod_{V_i \in \mathbf{V}_{t+1}} p(V_i \mid \mathrm{pa}_{\mathcal{G}_{t,t+1}}(V_i))$.

The joint distribution for the DBN system over a finite number of discrete time steps T is given by the product of the prior network distribution, and the transition conditional probability distributions for a set of time steps, as follows:
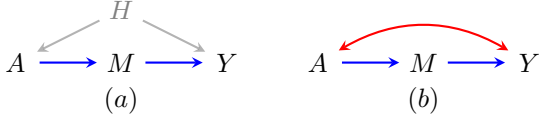
Figure 2: A hidden variable DAG and latent projection ADMG

$$\left(\prod_{V \in \mathbf{V}_1} p(V | \operatorname{pa}_{\mathcal{G}_1}(V))\right) \cdot \prod_{t=1}^{T-1} \left(\prod_{V \in \mathbf{V}_{t+1}} p(V | \operatorname{pa}_{\mathcal{G}_{t,t+1}}(V))\right) \quad (2)$$

A simple DBN is represented in Figure 1, where the prior network (1(a)) contains two variables A and L, and the transition network (1(b)) shows connections among the state variables in the prior state at time t and the subsequent state at time $t + 1$. We represent fixed vertices in a transition network via squares. Figure 1(c) shows the DBN implied by these prior and transition networks unrolled over 4 time steps. DBNs can be naturally extended to represent causal models by assuming that both prior and transition networks are causal DAGs. In other words, we assume values of every variable V_i in both the prior and the transition network is determined, via a structural equation $f_i(.)$, in terms of its observed parents $\operatorname{pa}_{\mathcal{G}_1}(V_i)$ (or $\operatorname{pa}_{\mathcal{G}_{t,t+1}}(V_i)$) and an exogenous noise term ϵ_i. If we further assume that all exogenous noise variables are marginally independent, we arrive at a DBN version of the NPSEM-IE, where in addition to the g-formula (1) holding for the prior network, the *conditional g-formula* holds for the transition network:

$$p(\mathbf{Y}_{t+1}(\mathbf{a})|\mathbf{V}_t) = \prod_{V_i \in \mathbf{Y}_{t+1}} p(V_i | \operatorname{pa}_{\mathcal{G}}(V_i))|_{\mathbf{A}=\mathbf{a}}, \quad (3)$$

for any $\mathbf{A} \subseteq \mathbf{V}_{t+1}$, and $\mathbf{Y}_{t+1} = \mathbf{V}_{t+1} \setminus \mathbf{A}$.

Thus, a causal DBN "unrolled" to a set of time points $1, \ldots, T$ yields a standard causal DAG model with vertices $\mathbf{V}_{1:T} \equiv \mathbf{V}_1 \cup \mathbf{V}_2 \cup \ldots \cup \mathbf{V}_T$. In particular, for an intervention that sets $\mathbf{A} \subseteq \mathbf{V}_{1:T}$ to constant values $\mathbf{a}$, the interventional distribution $p(\mathbf{Y}_{1:T}(\mathbf{a}))$, where $\mathbf{Y}_{1:T} = \mathbf{V}_{1:T} \setminus \mathbf{A}$, is identified by:

$$\prod_{V \in \mathbf{V}_1 \setminus \mathbf{A}} p(V | \operatorname{pa}_{\mathcal{G}_1}(V)) \prod_{t=1}^{T-1} \prod_{V \in \mathbf{V}_{t+1} \setminus \mathbf{A}} p(V | \operatorname{pa}_{\mathcal{G}_{t,t+1}}(V)) \Bigg|_{\mathbf{A}=\mathbf{a}} \quad (4)$$

Causal DBNs have been considered in prior work. Peters et al. (2013) illustrated how structural equations can be used in the context of time series data, addressing issues of identifiability. Malinsky and Spirtes (2018, 2019); Mogensen et al. (2018) presented structure learning algorithms for causal dynamic networks and applied them to macroeconomic data.

2.3 Causal Inference With Hidden Variables

The g-formula (1) provides an elegant link between observed data and counterfactual distributions in causal models where all relevant variables are observed. Causal models that arise in practice, however, contain hidden variables. Representing such models using a DAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$ where $\mathbf{V}$ and $\mathbf{H}$ correspond to observed and hidden variables, respectively, is not very helpful, since applying (1) to $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$ results in an expression that involves unobserved variables $\mathbf{H}$.

A popular alternative is to represent a class of hidden variable DAGs $\mathcal{G}_i(\mathbf{V} \cup \mathbf{H}_i)$ by a single *acyclic directed mixed graph* ADMG $\mathcal{G}(\mathbf{V})$ that contains directed ($\rightarrow$) and bidirected ($\leftrightarrow$) edges and no directed cycles via the *latent projection* operation (Verma and Pearl, 1990) (see Section B of the Appendix). The latent projection ADMG $\mathcal{G}(\mathbf{V})$ captures relationships between observed variables $\mathbf{V}$ implied by the factorization of $p(\mathbf{V} \cup \mathbf{H})$ with respect to $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$ via the *nested Markov factorization* of $p(\mathbf{V})$ with respect to $\mathcal{G}(\mathbf{V})$ (Richardson et al., 2017).

In particular, just as identification in DAGs may be viewed in terms of a modified DAG factorization (1), identification in a hidden variable DAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$ may be viewed in terms of a modified nested factorization of $\mathcal{G}(\mathbf{V})$. The nested Markov factorization of $p(\mathbf{V})$ with respect to $\mathcal{G}(\mathbf{V})$ is defined in terms of *Markov kernels* of the form $q_{\mathbf{D}}(\mathbf{D}| \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})$, where set $\mathbf{D} \subseteq \mathbf{V}$ is *intrinsic* in $\mathcal{G}(\mathbf{V})$ (see below). Kernels $q_{\mathbf{D}}(\mathbf{D}| \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})$ are objects that resemble conditional densities $p(V_i | \operatorname{pa}_{\mathcal{G}}(V_i))$ that arise in the Markov factorization for a DAG, in the sense that they are non-negative and normalize to 1 for every value of $\operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D}$. Kernels making up the nested Markov factorization are all functionals of $p(\mathbf{V})$, however, they are not necessarily equal to conditional distributions $p(\mathbf{D} | \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})$.

A set $\mathbf{D}$ is intrinsic if it is bidirected connected and reachable. A set $\mathbf{D}$ is reachable if it is possible to find an order on variables $\langle S_1, S_2, \ldots, S_m \rangle \equiv \mathbf{V} \setminus \mathbf{D}$ such that in each subgraph $\mathcal{G}_i(\mathbf{V})$, obtained from $\mathcal{G}(\mathbf{V})$ by removing all vertices $\{S_1, \ldots, S_{i-1}\}$ and adjacent edges, there is no variable W that is a descendant of S_i in $\mathcal{G}_i(\mathbf{V})$, and simultaneously has a path to S_i consisting exclusively of bidirected edges.

The nested Markov factorization asserts that the observed margin $p(\mathbf{V})$ can be expressed as a product $\prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}(\mathbf{V}))} q_{\mathbf{D}}(\mathbf{D} | \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})$ of kernels where $\mathcal{D}(\mathcal{G}(\mathbf{V}))$ is the set of bidirected connected components, called *districts*, in $\mathcal{G}(\mathbf{V})$. Additionally, the factorization implies certain other kernels associated with reachable sets may be expressed as similar products of intrinsic kernels. Finally, the modified form of the factorization may be used to express *any* interventional distribution identified from $p(\mathbf{V})$, as follows.

Given a latent projection ADMG $\mathcal{G}(\mathbf{V})$ representing a hidden variable causal model, and any disjoint subsets $\mathbf{Y}, \mathbf{A}$ of $\mathbf{V}$, let $\mathbf{Y}^*$ be the set of ancestors of $\mathbf{Y}$ in $\mathcal{G}(\mathbf{V})$ via directed paths that do not pass through $\mathbf{A}$, and let $\mathcal{G}_{\mathbf{Y}^*}$ be the *induced subgraph* of $\mathcal{G}(\mathbf{V})$ containing only vertices in $\mathbf{Y}^*$ and edges among these vertices.

(Shpitser and Pearl, 2006; Richardson et al., 2017) showed that any interventional distribution $p(\mathbf{Y}(\mathbf{a}))$ is identified from $p(\mathbf{V})$ given $\mathcal{G}(\mathbf{V})$ if and only if every bidirected connected component in $\mathcal{G}_{\mathbf{Y}^*}$ is intrinsic. Moreover, if $p(\mathbf{Y}(\mathbf{a}))$ is identified, it is given by the following margin of the modified nested Markov factorization, made up of the appropriate kernels:

$$p(\mathbf{Y}(\mathbf{a})) = \sum_{\mathbf{Y}^* \setminus (\mathbf{Y} \cup \mathbf{A})} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})} q_{\mathbf{D}}(\mathbf{D}| \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})|_{\mathbf{A}=\mathbf{a}}. \quad (5)$$

As a simple example, consider the hidden variable DAG in Fig. 2(a). Its latent projection ADMG in Fig. 2(b), called the *front-door graph*, has intrinsic sets $\{A\}, \{M\}, \{A, Y\}, \{Y\}$, with the corresponding kernels: $q_A(A) \equiv p(A)$, $q_M(M|A) \equiv$
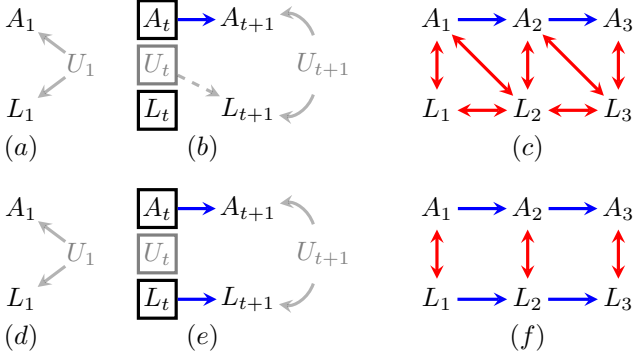
Figure 3: (a),(d) Prior network hidden variable DAGs $\mathcal{G}_0$, representing the state at time $t = 0$. (b),(e) Conditional hidden variable DAGs $\mathcal{G}_{t,t+1}$ representing the transitions in the network, with (e) leading to a first-order Markov model, and (b) leading to higher order dependences to unobserved hidden variables U_t linking multiple time points. (c),(f) Latent projection ADMGs of the unrolled hidden variable DBNs to three time steps.

$p(M|A)$, $q_{A,Y}(A,Y|M) \equiv p(Y|M,A)p(A)$, and $q_Y(Y \mid M) \equiv \sum_A p(Y|M,A)p(A)$.

By the nested Markov factorization, the observed margin $p(A, M, Y)$ factorizes as $q_{A,Y}(A, Y|M)q_M(M|A)$. In addition, certain other distributions also factorize. For example, the margin $p(A, M)$ is equal to $q_A(A)q_M(M|A)$. Further, $p(Y(a))$ is identified from $p(A, M, Y)$ and equal to $\sum_M q_Y(Y|M)q_M(M|a) = \sum_M \left(\sum_{A'} p(Y|M, A')p(A')\right) p(M|a)$, which is the *front-door formula* (Pearl, 1995). See Section B of the Appendix for a detailed exposition of the nested Markov factorization and identification theory in ADMGs.

3 Identification In Hidden Variable DBNs

Identification theory in hidden variable causal models can be extended to hidden variable causal DBN models as well. To the best of our knowledge, this extension is novel.

We start with an assumption that will allow us to view the marginal version of the DBN, defined only on observed variables, as a first-order Markov DBN.

Assumption 1 *The transition network $\mathcal{G}_{t+1,t}$ only depends on fixed variables in the previous time step t that are observed.*

If $\mathcal{G}_{t+1,t}$ depends on fixed variables that are hidden, the resulting DBN may result in observed variables in step $t + 1$ that depend on observed variables earlier than t even if observed variables in t are conditioned on.

For example, consider the DBN specified by prior and transition networks in Fig. 3 (a) and (b). Because the variable L_{t+1} depends on U_t, which is unobserved, and U_t influences L_t, "unrolling" this network, and taking the latent projection yields an ADMG shown in Fig. 3 (c), where L_3 ends up being

dependent on L_1, even after conditioning on L_2, A_2 (due to the "explaining away" phenomenon arising when a shared effect L_2 of two variables U_2 and U_1 is conditioned on). On the other hand, the DBN specified by prior and transition networks in Fig. 3 (d) and (e) does not suffer from this issue, as the transition network only depends on observed variables L_t, A_t, yielding a latent projection of the "unrolled" model shown in Fig. 3 (f), which factorizes into time step specific conditional distributions: $p(A_1, L_1)p(A_2, L_2|A_1, L_1)p(A_3, L_3|A_2, L_2)$.

In general, given a hidden variable prior network $\mathcal{G}_1$ on $\mathbf{V}_1, \mathbf{H}_1$, and transition network $\mathcal{G}_{t+1,t}$ on $\mathbf{V}_{t+1}, \mathbf{H}_{t+1}$ given $\mathbf{V}_t$, the hidden variable DBN may be represented by latent projections of the prior and transition networks: an ADMG $\mathcal{G}_1$ on $\mathbf{V}_1$, and a *conditional ADMG (CADMG)* $\mathcal{G}_{t+1,t}$ on $\mathbf{V}_{t+1}$ given $\mathbf{V}_t$, and the corresponding marginal distributions $p(\mathbf{V}_1)$ and $p(\mathbf{V}_{t+1,t}|\mathbf{V}_t)$. The "unrolled" version of the factorization of this model is: $p(\mathbf{V}_1) \prod_{t=1}^{T} p(\mathbf{V}_{t+1,t}|\mathbf{V}_t)$, where each term nested Markov factorizes with respect to either $\mathcal{G}_1$ or $\mathcal{G}_{t+1,t}$ by results in (Richardson et al., 2017). Intrinsic and reachable sets are defined in CADMGs by ignoring fixed vertices, and the nested factorization generalizes in the natural way from ADMGs from CADMGs, see Section B of the Appendix for details.

If the underlying DAGs correspond to causal models, the hidden variable DBN yields identification theory where the modified nested factorization (5) is applied at every time point, just as (1) was applied at every point in a fully observed causal DBN to yield (4). That is, given a fixed set of time points $1, \ldots, T$, vertices $\mathbf{V}_{1:T} \equiv \mathbf{V}_1 \cup \mathbf{V}_2 \cup \ldots \cup \mathbf{V}_T$, and disjoint subsets $\mathbf{A}, \mathbf{Y} \subseteq \mathbf{V}_{1:T}$, we have the following result:

Lemma 1 *Under Assumption 1 , $p(\mathbf{Y}(\mathbf{a}))$ is identified from a hidden variable causal DBN model represented by latent projections $\mathcal{G}_1$ on $\mathbf{V}_1$ and $\mathcal{G}_{t+1,t}$ on $\mathbf{V}_{t+1}$ given $\mathbf{V}_t$ if and only if every bidirected connected component in $\mathcal{G}_{1\mathbf{Y}_1^*}$ (the induced subgraph of $\mathcal{G}_1$) is intrinsic in $\mathcal{G}_1$, and every bidirected component in $\mathcal{G}_{t+1,t\mathbf{Y}_i^*}$ (the induced subgraph of $\mathcal{G}_{t+1,t}$) is intrinsic in $\mathcal{G}_{t+1,t}$, where $\mathbf{Y}_1^*$ is the set of ancestors of $\mathbf{Y} \cap \mathbf{V}_1$ not through $\mathbf{A} \cap \mathbf{V}_1$ in $\mathcal{G}_1$ and for every $i \in 2, \ldots, T$, $\mathbf{Y}_i^*$ is the set of ancestors of $\mathbf{Y} \cap \mathbf{V}_i$ not through $\mathbf{A} \cap \mathbf{V}_i$ in $\mathcal{G}_{t+1,t}$. Moreover, if $p(\mathbf{Y}(\mathbf{a}))$ is identified, we have*

$$\left(\sum_{\mathbf{Y}_1^* \setminus ((\mathbf{Y} \cup \mathbf{A}) \cap \mathbf{V}_1)} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{1\mathbf{Y}_1^*})} q_{\mathbf{D}}^1(\mathbf{D}| \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})|_{\mathbf{A}=\mathbf{a}} \right) \times$$

$$\prod_{i=2}^{T} \left(\sum_{\mathbf{Y}_i^* \setminus ((\mathbf{Y} \cup \mathbf{A}) \cap \mathbf{V}_i)} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{t+1,t\mathbf{Y}_i^*})} q_{\mathbf{D}}^{t+1,t}(\mathbf{D}| \operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})|_{\mathbf{A}=\mathbf{a}}, \right)$$

where $q_{\mathbf{D}}^1$ and $q_{\mathbf{D}}^{t+1,t}$ are kernels corresponding to intrinsic sets that are districts in $\mathcal{D}(\mathcal{G}_{1\mathbf{Y}_1^*})$ and $\mathcal{D}(\mathcal{G}_{t+1,t\mathbf{Y}_1^*})$ in the nested Markov factorizations of $\mathcal{G}_1$ and $\mathcal{G}_{t+1,t}$, respectively.

We present the proof of this result and a worked example in the Section B of the Appendix.

4

4 Fully Observed PDSEMs

A crucial modeling assumption employed by causal DBNs is that both structure and parameterization remain invariant over time. Such a model is ill-suited to capture the sort of path dependence described in the introduction. We now describe our new approach to relaxing this assumption via path dependent structural equation models (PDSEMs), which are a generalization of causal DBNs that can capture path dependence.

4.1 A Simple PDSEM

To illustrate PDSEMs, we will use a simple example inspired by our surgery setting. We assume a surgery will consist of three states: s^1 ("incision"), s^2 ("modification of bone/tissue"), and s^3 ("closing the incision"). Further, each state has the following variables: A (patient status prior to any procedures in the current stage), B (experience of surgeon performing the procedure in the current stage) and C (the observed patient outcome for the stage after procedure is performed), all observed. The surgery always starts at s^1, and concludes upon reaching s^3. Procedures performed in s^2 may either succeed, leading to s^3, or fail with some probability, leading the surgeon to revisit s^1. Note that DBNs are unable to capture even this simple example, due to the fact that DBN cannot represent distinct transitions to qualitatively different states. The state transition diagram for this scenario is shown in Fig. 4 (b). By contrast, the only type of state transition diagram allowed by DBNs contains a single self-looping state.

Relationships between variables in s^1 are shown by a causal diagram in Fig. 4 (a), corresponding to a set of structural equations as described in Sec 2.1. Fig. 4 (a) serves the role played by the prior network in a causal DBN. In addition to variables A_1, B_1 and C_1, this graph contains variable S_1, representing the state to transition to, at time step 1. In general, the probability associated with this variable may depend on other variables in the current state, however in our simple model, the state s^1 transitions to s^2 with probability 1.

Transitions are specified by causal CDAGs for each possible state transition, shown in Fig. 4(c),(d) and (e) (without the dashed edge). These graphs include edges between variables in time steps t and $t + 1$, as well as state-specific relationships among variables at time $t + 1$. Note that unlike a DBN, which only had a single transition network, multiple transition networks are needed to represent multiple transitions between different states. We assume state spaces of variables associated with each state are the same across state transition and prior graphs, though variables themselves and their causal relationships may differ across graphs. For example, the state spaces of A_1, B_1, C_1 in Fig. 4(a) and A_{21}, B_{21}, C_{21} in Fig. 4(c) are the same, while the variables themselves (and the causal graphs relating them) are not. This implies values may be indexed by state, e.g. a_1 can refer without loss of generality to a value of A_1 or A_{21}. We thus can thus index conditional distributions that depend on a prior state only by the prior state itself, e.g. $p(A_{12}|A_1)$ is a shorthand for "a density over A_{12} in transition $(1, 2)$ given any value a_1 of any variable of the form A_{i1}."
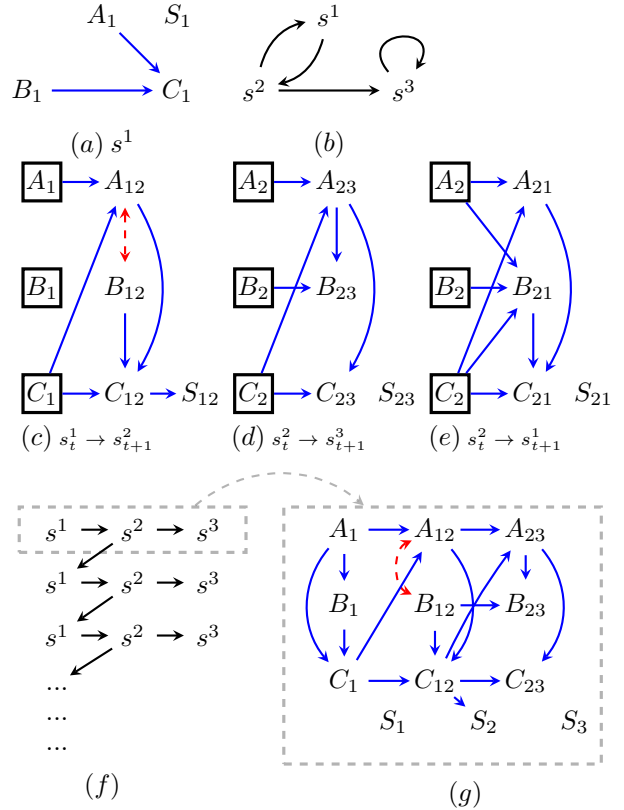

Figure 4: A simple PDSEM. (a) Causal structure of the initial state S^1. (b) The state transition diagram. (c),(d),(e) Causal diagrams representing possible transitions and subsequent states. (f) Causal relationships in a system evolving according to the state transitions: $s^1 \rightarrow s^2 \rightarrow s^3$. (g) A snapshot of a possible PDSEM trajectory represented as an unrolled DAG

Causal graphs in 4(a),(c),(d),(e), along with the state-transition diagram 4(b), completely describe the fully observed PDSEM. Complex state dynamics are captured by distinct state causal DAGs and path-dependence is simply a consequence of state-transitions that depend on variables in the current state, and not just the state itself.

The model we describe represents a randomized controlled trial where the surgeon operating during state s^2 is randomly assigned, hence B_{12} in the transition graph in Fig. 4 (c) has no parents. Otherwise, we encode standard causal relationships we expect: C in the previous state influences A, C in the next, and A in the previous state influences A in the next. Surgeon assignment B_{12} in s^2 influences assignments in subsequent stages, whether they are s^1 or s^3. The state transition at s^2 depends on the outcome C at that state. In s^3, B does not influence C, since closing the incision is a task adequately performed independent of surgeon experience.

The observed data factorization of a fully-observed PDSEM includes, in addition to factorization of the initial state causal DAG and factorizations with respect to CDAGs representing state transitions, state transition probabilities which are functions of variables in those factorizations. This factorization is not finite, but yields a well defined joint distribution p_∞ over

possible trajectories shown in 4(f). In our case, the distribution p_∞ factorizes as follows:

$$p_1 \prod_{t=1}^{\infty} (p_{12})^{\mathbb{I}(s_t^1, s_{t+1}^2)} (p_{23})^{\mathbb{I}(s_t^2, s_{t+1}^3)} (p_{21})^{\mathbb{I}(s_t^2, s_{t+1}^1)} 1^{\mathbb{I}(s_t^3)}$$

$$p_1 = p(A_1)p(B_1|A_1)p(C_1|A_1, B_1)\tilde{p}(S_1)$$
$$p_{12} = p(A_{12}|A_1, C_1)p(B_{12})p(C_{12}|B_{12}, A_{12}, C_1)p(S_{12}|C_{12})$$
$$p_{23} = p(A_{23}|A_2, C_2)p(B_{23}|B_2, A_{23})p(C_{23}|A_{23}, C_2)\tilde{p}(S_{23})$$
$$p_{21} = p(A_{21}|A_2, C_2)p(B_{21}|B_2, A_2, C_2)p(C_{21}|C_2, B_{21}, A_{21})\tilde{p}(S_{21}),$$

where s_t^i is the event "the state at time t is s^i, and all $\tilde{p}$ are deterministic by definition of our model.

PDSEMs, being causal models, allow us to reason about outcomes of interventions, for example questions such as: "what would happen if all procedures at every stage of a surgery are performed by the resident surgeon ($B = b$), possibly contrary to fact." The counterfactual joint distribution $p_\infty(b)$ corresponding to this intervention is obtained by standard structural equation replacement semantics of interventions applied to the models representing initial and transition graphs (Pearl, 2009). This distribution can be written as a product of state-specific marginal and conditional counterfactual distributions as follows:

$$p_1(b) \prod_{t=1}^{\infty} (p_{12}(b))^{\mathbb{I}(s_t^1, s_{t+1}^2)} (p_{23}(b))^{\mathbb{I}(s_t^2), s_{t+1}^3)} (p_{21}(b))^{\mathbb{I}(s_t^2, s_{t+1}^1)} 1^{\mathbb{I}(s_t^3)}$$

The distribution $p_\infty(b)$ is identified by using the g-formula for every component of the factorization of p_∞, in a generalization of (4), yielding:

$$p_0^* \prod_{t=1}^{\infty} (p_{12}^*)^{\mathbb{I}(s_t^1, s_{t+1}^2)} (p_{23}^*)^{\mathbb{I}(s_t^2, s_{t+1}^3)} (p_{21}^*)^{\mathbb{I}(s_t^2, s_{t+1}^1)} 1^{\mathbb{I}(s_t^3)}$$

$$p_1^* = p(A_1)p(C_1|A_1, b)\tilde{p}(S_1)$$
$$p_{12}^* = p(A_{12}|A_1, C_1)p(C_{12}|b, A_{12}, C_1)p(S_{12}|C_{12})$$
$$p_{23}^* = p(A_{23}|A_2, C_2)p(C_{23}|A_{23}, C_2)\tilde{p}(S_{23})$$
$$p_{21} = p(A_{21}|A_2, C_2)p(C_{21}|C_2, b, A_{21})\tilde{p}(S_{21}).$$

While the distribution $p(S_{12}|C_{12})$ governing how likely it is that s^1 or s^3 are visited from s^2 remains the same, the probability that s^1 is visited from s^2 is likely higher in $p_\infty(b)$ compared to p_∞. This is because B_{12} (counterfactually set to b) causes C_{12}, and C_{12} causes S_{12}. Thus, PDSEMs encode counterfactually changing state transition probabilities from their observed values. For example, surgeries where all stages are counterfactually performed by the less experienced resident surgeon might see more double-backs to s^1 to correct mistakes, compared to actually observed surgeries.

Having described this simple example, we turn to giving a general definition of PDSEMs.

4.2 Arbitrary PDSEMs

An arbitrary PDSEM is defined using a set of states $\mathbf{s}$, with an initial state s^1 and an absorbing state s^*, a set $\mathcal{T}$ of state index pairs of the form (i, j), where $s^i \neq s^*$ representing allowed state transitions, a DAG $\mathcal{G}_1$ on $\mathbf{V}_1$ for the initial state s^1, and for each $(i, j) \in \mathcal{T}$, a CDAG $\mathcal{G}_{ij}$ on $\mathbf{V}_{ij}$ given $\mathbf{V}_i$. Variables $S_1 \in \mathbf{V}_1, \{S_{ij} \in \mathbf{V}_{ij} : (i, j) \in \mathcal{T}\}$ determine probabilities of transitioning from state to state. Just as in a causal DBN, the DAG $\mathcal{G}_1$, and CDAGs $\mathcal{G}_{ij}$ represent structural equation

models for the initial state, and the appropriate state transitions, respectively. That is, in the initial state, each variable $V \in \mathbf{V}_1$ is determined via $f_V(\text{pa}_{\mathcal{G}}(V), \epsilon_V)$. Similarly, for each variable $V \in \mathbf{V}_{ij}$ in any state transition represented by $\mathcal{G}_{ij}$. We assume $S_1, \{S_{ij} : (i, j) \in \mathcal{T}\}$ have no outgoing edges (this is without loss of generality, as structural equations are already state-specific in a PDSEM).

In order to formulate a first order Markov PDSEM, we need the following assumption that ensures that we need not condition on any context in the past except variables in the prior state.

Assumption 2 *For every state s^j, any CDAG $\mathcal{G}_{ij}$ or DAG $\mathcal{G}_j$ will have random variables that share state spaces.*

We thus denote the values of any $\mathbf{V}_{ij}$ for any transition (i, j) into state j by $\mathbf{v}_j$ (note the lack of dependence on i). As in our example above, we index conditional densities that depend on variables in a prior state by that state only, e.g. $p(A_{12}|A_1)$.

Define $\mathbf{V} \equiv \mathbf{V}_1 \cup \left(\bigcup_{(i,j) \in \mathcal{T}} \mathbf{V}_{ij}\right)$. A PDSEM yields an observed data distribution $p_\infty(\mathbf{V})$ with the following factorization:

$$p_1(\mathbf{V}_1) \prod_{t=1}^{\infty} \left(\prod_{(i,j) \in \mathcal{T}} (p_{ij}(\mathbf{V}_{ij}|\mathbf{V}_i))^{\mathbb{I}(s_t^i, s_{t+1}^j)} \right) 1^{\mathbb{I}(s_t^*)}$$

$$p_1(\mathbf{V}_1) = \prod_{V \in \mathbf{V}_1} p(V|\text{pa}_{\mathcal{G}_1}(V)); \; p_{ij}(\mathbf{V}_{ij}|\mathbf{V}_i) = \prod_{V \in \mathbf{V}_{ij}} p(V|\text{pa}_{\mathcal{G}_{ij}}(V))$$

An intervention in a PDSEM is defined on a set of treatment variables $\mathbf{A} \equiv \bigcup_{(i,j) \in \mathcal{T}} \mathbf{A}_{ij}$ and a set to values $\mathbf{a}$ with the property that for any $(i, j), (k, j) \in \mathcal{T}$, the same values $\mathbf{a}_j$ are being set to $\mathbf{A}_{ij}$ and $\mathbf{A}_{ij}$. Define $\mathbf{Y}_{ij}$ in each transition graph $\mathcal{G}_{ij}$ to be all variables in that state not in $\mathbf{A}_{ij}$, with their corresponding values being $\mathbf{y}_j$, their union being $\mathbf{Y}$, and the values of the union being $\mathbf{y}$.

A new counterfactual distribution $p_\infty(\mathbf{Y}(\mathbf{a}))$ is obtained from the counterfactual initial state distribution $p_1(\mathbf{Y}_1(\mathbf{a}_1))$, and transition distributions $p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i))$ as:

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) \prod_{t=1}^{\infty} \left(\prod_{(i,j) \in \mathcal{T}} (p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)))^{\mathbb{I}(s_t^i, s_{t+1}^j)} \right) 1^{\mathbb{I}(s_t^*)}$$

$$(6)$$

Individual counterfactual distributions are obtained using standard structural equation replacement semantics.

Since the initial state and transitions are defined using structural equations, we obtain the following identification result, which generalizes the DBN g-formula (4) to PDSEMs.

Lemma 2 *Given a fully observed PDSEM, each factor of the distribution $p_\infty(\mathbf{Y}(\mathbf{a}))$ is identified from $p_\infty(\mathbf{V})$ as:*

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) \equiv \prod_{V \in \mathbf{Y}_1 \setminus \mathbf{A}_1} p_1(V|\text{pa}_{\mathcal{G}_1}(V))\Big|_{\mathbf{A}_1 = \mathbf{a}_1}$$

$$p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)) \equiv \prod_{V \in \mathbf{Y}_{ij} \setminus \mathbf{A}_j} p_{ij}(V|\text{pa}_{\mathcal{G}_{ij}}(V))\Big|_{\substack{\mathbf{A}_i = \mathbf{a}_i, \\ \mathbf{A}_j = \mathbf{a}_j}}, \quad (7)$$

5 PDSEMs With Hidden Variables

In extending causal inference to latent variable PDSEMs, in addition to Assumption 1 and Assumption 2, we also assume

the probabilities of any state transition trajectories are observed.

Assumption 3 *The variables S_{ij} for any $(i,j) \in \mathcal{T}$ governing state transition probabilities are observed.*

The latent variable PDSEMs then decompose into an initial state and a set of transitions such that causal inference results may be stated without loss of generality using latent projection ADMGs (and CADMGs) of appropriate DAGs and CDAGs. In addition, the fact that variables S_{ij} are observed implies we can evaluate counterfactual state transition probabilities, provided they are identified.

Formally, fix a PDSEM defined given the initial state DAG is $\mathcal{G}$ on $\mathbf{V}_1, \mathbf{H}_1$ and the set of transition CDAGs $\mathcal{G}_{ij}$ on $\mathbf{V}_{ij}, \mathbf{H}_{ij}$ given $\mathbf{V}_i$, for all $(i,j) \in \mathcal{T}$, such that $S_1 \in \mathbf{V}_1$, $S_{ij} \in \mathbf{V}_{ij}$ for every $(i,j) \in \mathcal{T}$, for every j and all $(i,j), (k,j) \in \mathcal{T}$, $\mathbf{H}_{ij} = \mathbf{H}_{kj}$ and $\mathbf{V}_{ij} = \mathbf{V}_{kj}$. We assume the variables $\mathbf{V} \equiv \{\mathbf{V}_1\} \cup \bigcup_{(i,j)\in\mathcal{T}} \mathbf{V}_{ij}$, and $\mathbf{H} \equiv \{\mathbf{H}_1\} \cup \bigcup_{(i,j)\in\mathcal{T}} \mathbf{H}_{ij}$ are observed, and hidden, respectively.

Given this definition of a latent variable PDSEM, the observed data distribution $p_\infty(\mathbf{V})$ is obtained from applying the usual transition probabilities to the margin at the initial state $p_1(\mathbf{V}_1) \equiv \sum_{\mathbf{H}_1} p_1(\mathbf{V}_1 \dot\cup \mathbf{H}_1)$, and the margins of all transition probabilities $p_{ij}(\mathbf{V}_{ij}|\mathbf{V}_i) \equiv \sum_{\mathbf{H}_{ij}} p_{ij}(\mathbf{V}_{ij} \dot\cup \mathbf{H}_{ij}|\mathbf{V}_i)$.

As before, fix a set of observed treatment variables $\mathbf{A}$, which is the union of $\{\mathbf{A}_{ij} : (i,j) \in \mathcal{T}\}$, such that the same values $\mathbf{a}_j$ are set to $\mathbf{A}_{ij}, \mathbf{A}_{kj}$ for any $(i,j), (k,j) \in \mathcal{T}$, and the set of outcomes $\mathbf{Y}_{ij} = \mathbf{V}_{ij} \setminus \mathbf{A}_{ij}$ for any $(i,j) \in \mathcal{T}$, with $\mathbf{Y}$ the union of $\{\mathbf{Y}_{ij} : (i,j) \in \mathcal{T}\}$.

Given the way the latent variable PDSEM was defined, identification theory for $p_\infty(\mathbf{Y}(\mathbf{a}))$ reduces to identification theory for $p_1(\mathbf{Y}_1(\mathbf{a}_1))$ in the latent projection ADMG $\mathcal{G}_1$ on $\mathbf{V}_1$, and $p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{V}_i(\mathbf{a}_i))$ in the latent projection CADMG $\mathcal{G}_{ij}$ on $\mathbf{V}_{ij}$ given $\mathbf{V}_i$, as follows:

Lemma 3 *Under Assumptions 1, 2 and 3, given a latent variable PDSEM represented by $\mathcal{G}_1$ and $\{\mathcal{G}_{ij} : (i,j) \in \mathcal{T}\}$, $p_\infty(\mathbf{Y}(\mathbf{a}))$ is identified from $p_\infty(\mathbf{V})$ if and only if every bidirected component in $\mathcal{G}_{1\mathbf{Y}_1}$ is intrinsic in $\mathcal{G}_1$, and every bidirected component in $\mathcal{G}_{ij\mathbf{Y}_j}$ is intrinsic in $\mathcal{G}_{ij}$ for every i and j. Moreover, if $p_\infty(\mathbf{Y}(\mathbf{a}))$ is identified, it is equal to*

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) \prod_{t=1}^\infty \left(\prod_{(i,j)\in\mathcal{T}} (p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)))^{\mathbb{I}(s_{t-1}^i, s_t^j)} \right)^{\mathbb{I}(s_{t-1}^*)} \tag{8}$$

where

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) = \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{1\mathbf{Y}_1^*})} q_\mathbf{D}^1(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_1}^s(\mathbf{D})) \Big|_{\mathbf{A}_1 = \mathbf{a}_1}, \tag{9}$$

where each kernel $q_\mathbf{D}^1(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_1}^s(\mathbf{D}))$ is in the nested Markov factorization of $p_1(\mathbf{V}_1)$ with respect to $\mathcal{G}_1$, and

$$p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)) = \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{V}_{ij}\setminus\mathbf{A}_{ij}})} q_\mathbf{D}^{ij}(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_{ij}}^s(\mathbf{D})) \Big|_{\substack{\mathbf{A}_i = \mathbf{a}_i, \\ \mathbf{A}_j = \mathbf{a}_j}}, \tag{10}$$

where each kernel $q_\mathbf{D}^{ij}(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_{ij}}^s(\mathbf{D}))$ is in the nested Markov factorization of $p_{ij}(\mathbf{V}_{ij}|\mathbf{V}_i)$ with respect to $\mathcal{G}_{ij}$.

An example of a hidden variable PDSEM and identifying functionals are given in Section B of the Appendix.

6 Experiments

6.1 Simulation of a latent variable PDSEM

We show how statistical inference may be performed in the example presented in Section 4.1 and Fig. 4, altered to include latent variables. The system has states $\{s^1, s^2, s^3\}$ and three variables in each state $\{A, B, C\}$. Additionally, s^2 has a hidden common cause of A and B. This is represented by the red (dotted) bidirected edge $A \leftrightarrow B$ in the latent projected ADMG shown in Fig. 4(c). Patient health status A, surgeon experience B, and duration of the stage of surgery C, are all continuous variables. State and transition graphs are identical to those in Fig. 4.

Two sets of parameters associated with a generative model of this kind are $p(S_{t+1} = s^j|S_t = s^i, \mathbf{V}_t)$, where $s_t^i \rightarrow s_{t+1}^j$ is a transition allowed by the model and $p(V_{t+1}^{ij} = v|S_{t+1} = s^j, S_t = s^i, \mathbf{V}_t)$, where $V^{ij} \in \{A^{ij}, B^{ij}, C^{ij}\}$, and, $s_t^i \rightarrow s_{t+1}^j$ is an allowed transition. These parameters are chosen to be reasonable for the surgery application, yielding a distribution Markov relative to appropriate graphs.

A dataset of $N = 10000$ "surgeries" was simulated, all with the same initial state. Transition probabilities were generated using a logistic regression on variables in the current state, with transitions eventually terminating at the absorbing state. Each variable V_i is generated from a set of linear structural equations with correlated errors. Using generated data, state transition probabilities were estimated using maximum likelihood. Parameters for the structural equation model were estimated using the RICF algorithm (Drton et al., 2009), implemented in the Ananke [1] package (Bhattacharya et al.).

We used the PDSEM to consider the causal impact of surgeon experience (measured by total operating time in their career) on average surgery length. This outcome is easy to measure, and is known to serve as an informative proxy for other measures of surgery quality, such as follow-up assessments of quality of life (Rambachan et al., 2013; Jackson et al., 2011). We assessed this causal question by generating two sets of sampled surgery trajectories where, in each stage of the surgery, the surgeon was intervened to have higher (vs. lower) career operating time by one unit. These trajectories may be viewed as a Monte Carlo sampling scheme for evaluating the functional given by (13), (14) and (15). The comparison of these two sets of trajectories may be viewed as a generalization of the *average causal effect (ACE)* from classical longitudinal causal models to PDSEMs.

The results are shown in Fig. 5. Surgeries performed by experienced surgeons are shorter ($\mu = 5.79$, $\mathbf{q}_{0.05} = 3$, $\mathbf{q}_{0.95} = 13$) than those performed by trainees ($\mu = 7.02$, $\mathbf{q}_{0.05} = 3$, $\mathbf{q}_{0.95} = 17$) where $\mathbf{q}_p$ denotes the p^{th} quantile. Surgeries performed by trainees have higher variance.

6.2 Data Application of the PDSEM

We now illustrate how a PDSEM may be applied to analyze data obtained from a surgery. Our dataset consists of 236 septoplasty procedures conducted at our institution's research hospital. A total of 57343 timestamped records were collected,
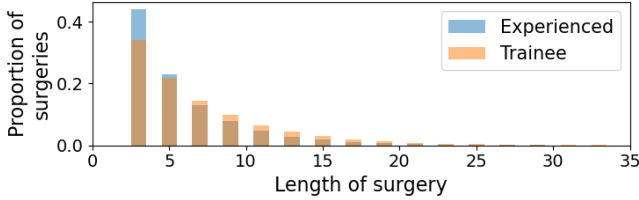
Figure 5: Histograms of the number of transitions in a surgery under two different interventions: when a more experienced surgeon performs the entire procedure, and when a less experienced trainee performs the entire procedure.

including tool and personnel activity. Surgeries consist of six distinct phases: s^1 (opening of the septum), s^2 (raising septal flaps), s^3 (removal of deviated septal cartilage and bone), s^4 (reconstruction), s^5 (closing of the incision), and s^6 (other activity). An artificial absorbing state s^{end} represents the end of procedures. Procedures are often led by an attending, with a surgeon trainee assisting. Of the surgeries, 42.79% of them were performed fully by the leading attending; the others by a team. Also, attending surgeons perform for 64.98% of all operating time and trainees the rest. Twelve different surgical tools were tracked for use. Each phase of the surgery requires different techniques and tools. The progression of the surgery is not monotonic – surgeons commonly revisit earlier stages. The state transition diagram representing allowed state transitions is presented in Fig 7. We chose to discretize all variables into two categories. Model parameters were estimated by maximum likelihood. More details about the data and model can be found in Section D of the Appendix.
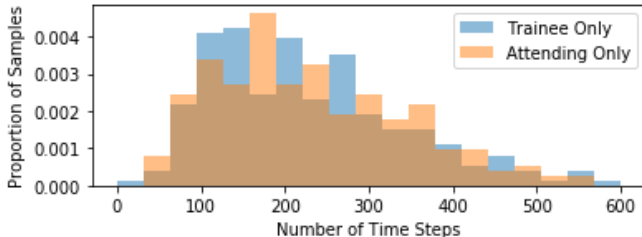

Figure 6: Histograms of hypothetical surgeries performed only by a junior trainee surgeon (blue) versus hypothetical surgeries performed only by a senior attending surgeon (orange). Surgeries performed by the attending are slightly longer ($\mu = 244.3.91, \sigma = 139.9$) than those of the trainee ($\mu = 233.5, \sigma = 125.9$).
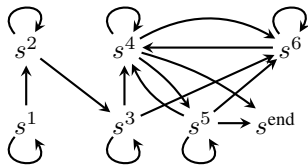

Figure 7: The state transition diagram for the surgery data application.

As before, we considered the causal impact of surgeon experience on average length of surgery, evaluated by consider-

ing counterfactual trajectories and comparing to trajectories observed in the data. Estimation of $p(s_t|s_{t-1}, \mathbf{v}_{t-1})$ at all levels of $s_{t-1}, \mathbf{v}_{t-1}$ is not always possible due to finite sample limitations. To address this, we apply additive smoothing to $p(s_t|s_{t-1}, \mathbf{v}_{t-1})$, based on the empirical distribution $p(s_t|s_{t-1})$. Goodness of fit is illustrated in Fig 10 of the Appendix and results are presented in Fig 6. We have made considerable assumptions in modeling our PDSEM and have closely matched the generative model to the empirical distribution (Fig 10). We observe that the causal effect of surgeon skill on surgery length, given our learned parameters, is close to zero. This indicates that policies that govern the trade-off between the need to train surgeons, and overall surgery quality (as quantified by our outcome) are effective at our institution.

Generalizing statistical inference in PDSEMs with hidden variables to likelihoods based on parameterizations of the nested Markov model (Richardson et al., 2017) presents a number of open problems, especially for variables that are not binary or not multivariate normal. We discuss these issues in Section D of the Appendix.

7 Conclusions

In this paper, we have introduced the Path Dependent Structural Equation Model (PDSEM) for longitudinal data which unifies complex state structure from DBNs and complex state transition dynamics from MDPs. It can also be seen as a graphical model generalizing the dynamics of a Markov chain with state-specific dynamics. We have described counterfactuals associated with these causal models that can alter the subsequent temporal evolution of the system, identification theory for such counterfactuals in terms of the observed data distribution, and described estimation. We showed the utility of the model in clinical settings using simulations as well as real data from a septoplasty surgical procedure. Developing novel methods for efficient Monte Carlo sampling based statistical inference for hidden variable versions of PDSEMs based on the nested Markov model is a promising area of future work.

Acknowledgements

References

Narges Ahmidi, Piyush Poddar, Jonathan D. Jones, Swaroop S. Vedula, Lisa Ishii, Gregory D. Hager, and Masaru Ishii. Automated objective surgical skill assessment in the operating room from unstructured tool motion in septoplasty. *International Journal of Computer Assisted Radiology and Surgery*, 10(6):981–991, 2015.

Charlene F. Belanger, Charles H. Hennekens, Bernard Rosner, and Frank E. Speizer. The nurses' health study. *The American Journal of Nursing*, 78(6):1039–1040, June 1978. ISSN 0002-936X.

Rohit Bhattacharya, Jaron J. R. Lee, Razieh Nabi, and Ilya Shpitser. *Ananke*: A python package for causal inference with graphical models. URL https://ananke.readthedocs.io/en/latest/index.html.

Rohit Bhattacharya, Razieh Nabi, and Ilya Shpitser. Semiparametric inference for causal effects in graphical models with hidden variables. *arXiv preprint arXiv:2003.12659*, 2020.

Bibhas Chakraborty and Erica E. Moodie. *Statistical methods for dynamic treatment regimes: reinforcement learning, causal inference, and personalized medicine*. New York: Springer-Verlag, 2013.

Mathias Drton. Likelihood ratio tests and singularities. *The Annals of Statistics*, 37(2):979–1012, 2009.

Mathias Drton, Michael Eichler, and Thomas S. Richardson. Computing maximum likelihood estimates in recursive linear models with correlated errors. *Journal of Machine Learning Research*, 10(10), 2009.

Robin J. Evans and Thomas S. Richardson. Markovian acyclic directed mixed graphs for discrete data. *Annals of Statistics*, pages 1–30, 2014.

Robin J. Evans and Thomas S. Richardson. Smooth, identifiable supermodels of discrete DAG models with latent variables. *Bernoulli*, 2018. (to appear).

Nicholas Fettman, Thomas Sanford, and Raj Sindwani. Surgical management of the deviated septum: techniques in septoplasty. *Otolaryngologic Clinics of North America*, 42 (2):241–252, 2009.

Miguel A. Hernán, Babette Brumback, and James M. Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*, pages 561–570, 2000.

Timothy D. Jackson, Jeffrey J. Wannares, Todd R. Lancaster, David W. Rattner, and Matthew M. Hutter. Does speed matter? The impact of operative time on outcome in laparoscopic surgery. *Surgical Endoscopy*, 25(7):2288–2295, 2011.

Steffan L. Lauritzen. *Graphical Models*. Oxford, U.K.: Clarendon, 1996.

Stan Liebowtiz and Stephen Margolis. Path dependence. *Encyclopedia of Law and Economics*, 2002.

Daniel Malinsky and Peter Spirtes. Causal structure learning from multivariate time series in settings with unmeasured confounding. In *Proceedings of 2018 ACM SIGKDD Workshop on Causal Discovery*, pages 23–47, August 2018.

Daniel Malinsky and Peter Spirtes. Learning the structure of a nonstationary vector autoregression. *Proceedings of Machine Learning Research*, 89:2986–2994, April 2019. ISSN 2640-3498.

Søren W. Mogensen, Daniel Malinsky, and Niels R. Hansen. Causal learning for partially observed stochastic dynamical systems. In *UAI*, pages 350–360, 2018.

Kevin P. Murphy. *Machine learning: A probabilistic perspective*. MIT Press, September 2012. ISBN 978-0-262-30432-0.

Susan A. Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–709, 1995. URL citeseer.ist.psu.edu/55450.html.

Judea Pearl. *Causality: models, reasoning, and inference*. Cambridge University Press, 2 edition, 2009. ISBN 978-0521895606.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Causal inference on time series using restricted structural equation models. *Advances in Neural Information Processing Systems 26*, pages 154–162, 2013.

Lawrence R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

Aksharananda Rambachan, Lauren M. Mioton, Sujata Saha, Neil Fine, and John Y. S. Kim. The impact of surgical duration on plastic surgery outcomes. *European Journal of Plastic Surgery*, 36(11):707–714, 2013.

Mark A. Richards. *Fundamentals of radar signal processing*. Tata McGraw-Hill Education, 2005.

Thomas S. Richardson, James M. Robins, and Ilya Shpitser. Nested Markov properties for acyclic directed mixed graphs. In *Twenty Eighth Conference on Uncertainty in Artificial Intelligence (UAI-12)*, 2012.

Thomas S. Richardson, Robin J. Evans, James M. Robins, and Ilya Shpitser. Nested Markov properties for acyclic directed mixed graphs, 2017. Working paper.

James M. Robins. Marginal structural models. In *American Statistical Association, Section on Bayesian Statistical Science*, pages 1–10, 1997.

James M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical Models in Epidemiology: The Environment and Clinical Trials*. NY: Springer-Verlag, 1999.

Eli Sherman and Ilya Shpitser. Identification and estimation of causal effects from dependent data. In *Advances in neural information processing systems*, pages 9424–9435, 2018.

Ilya Shpitser and Judea Pearl. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*. AAAI Press, Palo Alto, 2006.

Ilya Shpitser, Thomas S. Richardson, and James M. Robins. An efficient algorithm for computing interventional distributions in latent variable causal models. In *Uncertainty in Artificial Intelligence*, volume 27. AUAI Press, 2011.

Ilya Shpitser, Robin J. Evans, and Thomas S. Richardson. Acyclic linear sems obey the nested markov property. In *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.

Bobby A. Tajudeen and David W. Kennedy. Thirty years of endoscopic sinus surgery: What have we learned? *World Journal of Otorhinolaryngology-Head and Neck Surgery*, 3 (2):115–121, 2017.

Sebastian Thrun. Learning to play the game of chess. In *Advances in neural information processing systems*, pages 1069–1076, 1995.

Thomas S. Verma and Judea Pearl. Equivalence and synthesis of causal models. Technical Report R-150, Department of Computer Science, University of California, Los Angeles, 1990.

Appendix

A   Graph preliminaries

Let capital letters X denote random variables, and let lower case letters x values of X. Sets of random variables are denoted $\mathbf{V}$, and sets of values $\mathbf{v}$. For a subset $\mathbf{A} \subseteq \mathbf{V}$, $\mathbf{v_A}$ denotes the subset of values in $\mathbf{v}$ of variables in $\mathbf{A}$. Domains of X and $\mathbf{X}$ are denoted by $\mathfrak{X}_X$ and $\mathfrak{X}_{\mathbf{X}}$, respectively.

Standard genealogic relations on graphs are as follows: parents, children, descendants, siblings and ancestors of X in a graph $\mathcal{G}$ are denoted by $\mathrm{pa}_{\mathcal{G}}(X), \mathrm{ch}_{\mathcal{G}}(X), \mathrm{de}_{\mathcal{G}}(X), \mathrm{si}_{\mathcal{G}}(X), \mathrm{an}_{\mathcal{G}}(X)$, respectively (Lauritzen, 1996). These relations are defined disjunctively for sets, e.g. $\mathrm{pa}_{\mathcal{G}}(\mathbf{X}) \equiv \bigcup_{X \in \mathbf{X}} \mathrm{pa}_{\mathcal{G}}(X)$. By convention, for any X, $\mathrm{an}_{\mathcal{G}}(X) \cap \mathrm{de}_{\mathcal{G}}(X) \cap \mathrm{dis}_{\mathcal{G}}(X) = \{X\}$.

We will also define the set of *strict parents* as follows: $\mathrm{pa^s}_{\mathcal{G}}(\mathbf{X}) = \mathrm{pa}_{\mathcal{G}}(\mathbf{X}) \setminus \mathbf{X}$. Given any vertex V in an ADMG $\mathcal{G}$, define the *ordered Markov blanket* of V as $\mathrm{mb}_{\mathcal{G}}(V) \equiv (\mathrm{dis}_{\mathcal{G}}(V) \cup \mathrm{pa}_{\mathcal{G}}(\mathrm{dis}_{\mathcal{G}}(V))) \setminus V$. Given a graph $\mathcal{G}$ with vertex set $\mathbf{V}$, and $\mathbf{S} \subseteq \mathbf{V}$, define the *induced subgraph* $\mathcal{G}_{\mathbf{S}}$ to be a graph containing the vertex set $\mathbf{S}$ and all edges in $\mathcal{G}$ among elements in $\mathbf{S}$.

In the subsequent discussion, we will denote an ADMG $\mathcal{G}$ on $\mathbf{V}$ by notation $\mathcal{G}(\mathbf{V})$, and a CADMG $\mathcal{G}$ on $\mathbf{V}$ given $\mathbf{W}$ by notation $\mathcal{G}(\mathbf{V}, \mathbf{W})$.

B   The Nested Markov Factorization

It is recommended that the reader look up notation for graphs in Section A of the Appendix to follow this section.

B.1   Why do we need an alternative factorization?

A hidden variable CDAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H}, \mathbf{W})$ may be used to define a factorization on distributions $p(\mathbf{V}|\mathbf{W})$ in terms of the CDAG as: $p(\mathbf{V}|\mathbf{W}) = \sum_{\mathbf{H}} \prod_{V \in \mathbf{V} \cup \mathbf{H}} p(V|\mathrm{pa}_{\mathcal{G}}(V))$. However, inferences may be sensitive to assumptions made about the state spaces for the unobserved variables and the latent variable model may contain singularities at which asymptotics are irregular (Drton, 2009). Additionally, such a model does not form a tractable search space: an arbitrary number of hidden variables and associated structures may be incorporated that are consistent with observed data distributions.

Alternatively, a factorization of the marginal distribution $p(\mathbf{V}|\mathbf{W})$ can be defined directly on the latent projection CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$. This *nested Markov factorization*, described in (Richardson et al., 2017) completely avoids modeling hidden variables, and leads to a regular likelihood in special cases (Evans and Richardson, 2018). It captures all equality constraints a hidden variable CDAG factorization imposes on the observed margin $p(\mathbf{V}|\mathbf{W})$ (Shpitser et al., 2018). In addition, $p(Y(a)|\mathbf{W})$ (an interventional distribution given a fixed context $\mathbf{W}$) identified in a hidden variable causal model represented by $\mathcal{G}(\mathbf{V} \cup \mathbf{H}, \mathbf{W})$ is always equal to a modified version of a nested factorization (Richardson et al., 2017) associated with $\mathcal{G}(\mathbf{V}, \mathbf{W})$, described here.

B.2   The nested Markov factorization

The nested Markov factorization of $p(\mathbf{V}|\mathbf{W})$ with respect to a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$ links *kernels*, mappings derived from $p(\mathbf{V}|\mathbf{W})$ and CADMGs derived from $\mathcal{G}(\mathbf{V}, \mathbf{W})$ via a *fixing* operation.

Kernel: A kernel $q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$ is a mapping from values in $\mathbf{W}$ to normalized densities over $\mathbf{V}$ (Lauritzen, 1996). A conditional distribution is a familiar example of a kernel, in that $\sum_{\mathbf{v} \in \mathbf{V}} q_{\mathbf{V}}(\mathbf{v}|\mathbf{w}) = 1$. Conditioning and marginalization are defined in kernels in the usual way: For $\mathbf{A} \subseteq \mathbf{V}$, $q_{\mathbf{V}}(\mathbf{A}|\mathbf{W}) \equiv \sum_{\mathbf{V} \setminus \mathbf{A}} q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$ and $q_{\mathbf{V}}(\mathbf{V} \setminus \mathbf{A}|\mathbf{A} \cup \mathbf{W}) \equiv \frac{q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})}{q_{\mathbf{V}}(\mathbf{A}|\mathbf{W})}$.

Fixability and the fixing operator: A variable $V \in \mathbf{V}$ in a CADMG $\mathcal{G}$ is fixable if $\mathrm{de}_{\mathcal{G}}(V) \cap \mathrm{des}_{\mathcal{G}}(V) = \emptyset$. In other words, V is fixable if paths $V \leftrightarrow ... \leftrightarrow B$ and $V \to ... \to B$ do not both exist in $\mathcal{G}$ for any $B \in \mathbf{V} \setminus \{V\}$.

We define a fixing operator $\phi_V(\mathcal{G})$ for graphs, and a fixing operator $\phi_V(q; \mathcal{G})$ for kernels. Given a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, with a fixable $V \in \mathbf{V}$, $\phi_V(\mathcal{G}(\mathbf{V}, \mathbf{W}))$ yields a new CADMG $\mathcal{G}(\mathbf{V} \setminus \{V\}, \mathbf{W} \cup \{V\})$ obtained from $\mathcal{G}(\mathbf{V}, \mathbf{W})$ by moving V from $\mathbf{V}$ to $\mathbf{W}$, and removing all edges with arrowheads into V. Given a kernel $q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$, and a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, the operator $\phi_V(q_{\mathbf{V}}(\mathbf{V}|\mathbf{W}), \mathcal{G}(\mathbf{V}, \mathbf{W}))$ yields a new kernel:

$$q_{\mathbf{V} \setminus \{V\}}(\mathbf{V} \setminus \{V\}|\mathbf{W} \cup \{V\}) \equiv \frac{q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})}{q_{\mathbf{V}}(V|\mathrm{mb}_{\mathcal{G}}(V))}$$

Fixing sequences: A sequence $\langle V_1, \ldots, V_k \rangle$ is said to be *valid* in $\mathcal{G}(\mathbf{V}, \mathbf{W})$ if V_1 fixable in $\mathcal{G}(\mathbf{V}, \mathbf{W})$, V_2 is fixable in $\phi_{V_1}(\mathcal{G}(\mathbf{V}, \mathbf{W}))$, and so on. If any two sequences σ_1, σ_2 for the same set $\mathbf{S} \subseteq \mathbf{V}$ are fixable in $\mathcal{G}$, they lead to the same CADMG.

The graph fixing operator can be extended to a set $\mathbf{S}$: $\phi_{\mathbf{S}}(\mathcal{G})$. This operator is defined as applying the vertex fixing operation in any valid sequence σ for set $\mathbf{S}$.

Given a sequence $\sigma_{\mathbf{S}}$, define $\eta(\sigma_{\mathbf{S}})$ to be the first element in $\sigma_{\mathbf{S}}$, and $\tau(\sigma_{\mathbf{S}})$ to be the subsequence of $\sigma_{\mathbf{S}}$ containing all elements but the first. Given a sequence $\sigma_{\mathbf{S}}$ on elements in $\mathbf{S}$ valid in $\mathcal{G}(\mathbf{V}, \mathbf{W})$, the kernel fixing operator $\phi_{\sigma_{\mathbf{S}}}(q_{\mathbf{V}}(\mathbf{V}|\mathbf{W}), \mathcal{G}(\mathbf{V}, \mathbf{W}))$ is defined to be equal to $q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$ if $\sigma_{\mathbf{S}}$ is the empty sequence, and $\phi_{\tau(\sigma_{\mathbf{S}})}(\phi_{\eta(\sigma_{\mathbf{S}})}(q_{\mathbf{V}}(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W})), \phi_{\eta(\sigma_{\mathbf{S}})}(\mathcal{G}(\mathbf{V}, \mathbf{W})))$ otherwise.

Reachability: Given a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, a set $\mathbf{R} \subseteq \mathbf{V}$ is called *reachable* if there exists a sequence for $\mathbf{V} \setminus \mathbf{R}$ valid in $\mathcal{G}(\mathbf{V}, \mathbf{W})$. In other words, if $\mathbf{S}$ is fixable in $\mathcal{G}$, $\mathbf{V} \setminus \mathbf{S}$ is reachable.

Intrinsic sets: A set $\mathbf{R}$ reachable in $\mathcal{G}(\mathbf{V}, \mathbf{W})$ is *intrinsic* in $\mathcal{G}(\mathbf{V}, \mathbf{W})$ if $\phi_{\mathbf{V} \setminus \mathbf{R}}(\mathcal{G})$ contains a single district, $\mathbf{R}$ itself. The set of intrinsic sets in a CADMG $\mathcal{G}$ is denoted by $\mathcal{I}(\mathcal{G})$.

Nested Markov factorization: A distribution $p(\mathbf{V}|\mathbf{W})$ is said to obey the *nested Markov factorization* with respect to the CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$ if there exists a set of kernels of the form $\{q_{\mathbf{S}}(\mathbf{S}| \mathrm{pa}_{\mathcal{G}}(\mathbf{S})) : \mathbf{S} \in \mathcal{I}(\mathcal{G})\}\}$ such that for every valid sequence $\sigma_{\mathbf{R}}$ for a reachable set $\mathbf{R}$ in $\mathcal{G}$, we have:

$$\phi_{\sigma_{\mathbf{R}}}(p(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W})) = \prod_{\mathbf{D} \in \mathcal{D}(\phi_{\mathbf{R}}(\mathcal{G}(\mathbf{V}, \mathbf{W})))} q_{\mathbf{D}}(\mathbf{D}| \mathrm{pa}^{\mathrm{s}}_{\mathcal{G}}(\mathbf{D}))$$

If a distribution obeys this factorization, then for any reachable $\mathbf{R}$, any two valid sequences on $\mathbf{R}$ applied to $p(\mathbf{V}|\mathbf{W})$ yield the same kernel $q_{\mathbf{R}}(\mathbf{R}|\mathbf{V} \setminus \mathbf{R})$. Hence, kernel fixing may be defined on sets, just as graph fixing. In this case, for every $\mathbf{D} \in \mathcal{I}(\mathcal{G})$, $q_{\mathbf{D}}(\mathbf{D}| \mathrm{pa}^{\mathrm{s}}_{\mathcal{G}}(\mathbf{D})) \equiv \phi_{\mathbf{V} \setminus \mathbf{D}}(p(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W}))$.

The *district factorization* or *Tian factorization* of $p(\mathbf{V}|\mathbf{W})$ results from the nested factorization:

$$p(\mathbf{V}|\mathbf{W}) = \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}(\mathbf{V}, \mathbf{W}))} q_{\mathbf{D}}(\mathbf{D}| \mathrm{pa}^{\mathrm{s}}_{\mathcal{G}}(\mathbf{D}))$$

$$= \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}(\mathbf{V}, \mathbf{W}))} \left(\prod_{D \in \mathbf{D}} p(D \mid \mathrm{pre}_{\prec}(D)) \right),$$

where $\mathrm{pre}_{\prec}(D)$ is the set of predecessors of D according to a topological total ordering $\prec$. Each factor $\prod_{D \in \mathbf{D}} p(D \mid \mathrm{pre}_{\prec}(D))$ is only a function of $\mathbf{D} \cup \mathrm{pa}_{\mathcal{G}}(\mathbf{D})$ under the nested factorization.

An important result in (Richardson et al., 2017) states that if $p(\mathbf{V} \cup \mathbf{H}|\mathbf{W})$ obeys the factorization for a CDAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H}, \mathbf{W})$, then $p(\mathbf{V}|\mathbf{W})$ obeys the nested factorization for the latent projection CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$.

B.3   Identification

Not every interventional distribution $p(\mathbf{Y}(\mathbf{a}))$ is identified in a hidden variable causal model. However, *every* $p(\mathbf{Y}(\mathbf{a})|\mathbf{W})$ identified from $p(\mathbf{V}|\mathbf{W})$ can be expressed as a modified nested factorization as follows:

$$p(\mathbf{Y}(\mathbf{a})|\mathbf{W}) = \sum_{\mathbf{Y}^* \setminus \mathbf{Y}} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})} p(\mathbf{D}| \mathrm{do}(\mathrm{pa}^{\mathrm{s}}_{\mathcal{G}}(\mathbf{D})))|_{\mathbf{A}=\mathbf{a}}$$

$$= \sum_{\mathbf{Y}^* \setminus \mathbf{Y}} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})} \phi_{\mathbf{V} \setminus \mathbf{D}}(p(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W}))|_{\mathbf{A}=\mathbf{a}}$$

where $\mathbf{Y}^* \equiv \mathrm{an}_{\mathcal{G}(\mathbf{V}(\mathbf{a}), \mathbf{W})}(\mathbf{Y}) \setminus \mathbf{a}$. That is, $p(\mathbf{Y}(\mathbf{a})|\mathbf{W})$ is only identified if it can be expressed as a factorization, where every piece corresponds to a kernel associated with a set intrinsic in $\mathcal{G}(\mathbf{V}, \mathbf{W})$. Moreover, no piece in this factorization contains elements of $\mathbf{A}$ as random variables.

B.4   Example of the nested factorization of a hidden variable PDSEM

A hidden variable PDSEM can be unrolled into a latent-projected ADMG if the model obeys restrictions given in Section 5. For instance, Fig. 8 in this Appendix shows an example where the first two states of the system involve hidden variables. In particular, the system at s^2 is the *front-door-graph* previously encountered in Section 2. Transition graphs are in Fig. 8(c)-(e).

The nested factorization for the initial graph in Fig. 8 (a) has intrinsic sets

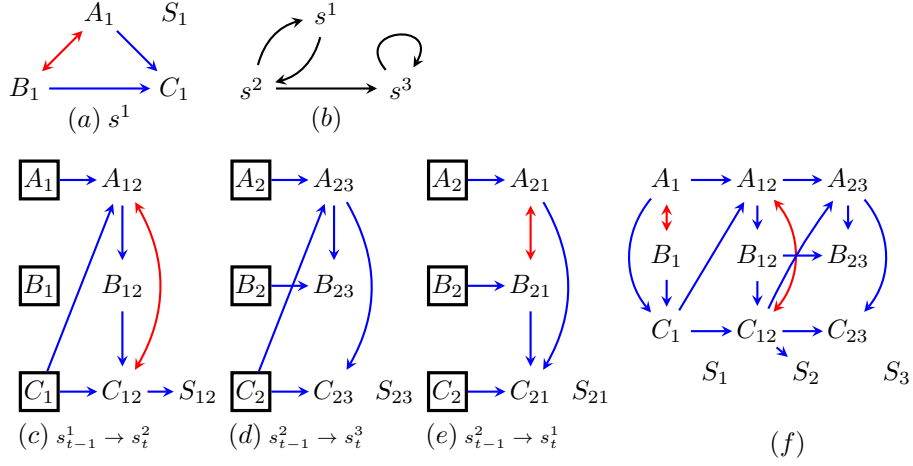$$(a) : \{A_1\}, \{B_1\}, \{C_1\}, \{A_1, B_1\}, \{S_1\}$$

Figure 8: A hidden variable PDSEM. (a) Causal structure of the initial state S^1. (b) The state transition diagram. (c),(d),(e) Latent projected causal diagrams representing possible transitions and subsequent states. (f) A snapshot of a possible PDSEM trajectory represented as an unrolled ADMG

with corresponding kernels

$$(a): q_{A_1}(A_1) \equiv p(A_1); q_{B_1}(B_1) = p(B_1); q_{C_1}(C_1|A_1, B_1) \equiv p(C_1|A_1, B_1); q_{A_1, B_1}(A_1, B_1) \equiv p(A_1, B_1); q_{S_1}(S_1) \equiv p(S_1).$$
(11)

Similarly, the nested factorizations for the transition graphs in Fig. 8 (c),(d),(e) have intrinsic sets:

$$(c): \{A_{12}\}, \{B_{12}\}, \{C_{12}\}, \{A_{12}, C_{12}\}, \{S_{12}\}$$
$$(d): \{A_{23}\}, \{B_{23}\}, \{C_{23}\}, \{S_{23}\}$$
$$(e): \{A_{21}\}, \{B_{21}\}, \{A_{21}, B_{21}\}, \{C_{21}\}, \{S_{21}\},$$

with corresponding kernels

$$(c): q_{A_{12}}(A_{12}|C_1) \equiv p(A_{12}|C_1); q_{B_{12}}(B_{12}|A_{12})) \equiv p(B_{12}|A_{12}); q_{C_{12}}(C_{12}|C_1, B_{12}) \equiv \sum_{A_{12}} p(C_{12}|B_{12}, A_{12}, C_1)p(A_{12}|C_1);$$

$$q_{A_{12}, C_{12}}(A_{12}, C_{12}|B_{12}, C_1) \equiv p(C_{12}|B_{12}, A_{12}, C_1)p(A_{12}|C_1); q_{S_{12}}(S_{12}|C_{12}) \equiv p(S_{12}|C_{12}).$$
$$(d): q_{A_{23}}(A_{23}|A_2, C_2) \equiv p(A_{23}|A_2, C_2); q_{B_{23}}(B_{23}|B_2, A_{23}) \equiv p(B_{23}|B_2, A_{23}); q_{C_{23}}(C_{23}|C_2, A_{23}) \equiv p(C_{23}|C_2, A_{23});$$
$$q_{S_{23}}(S_{23}) \equiv p(S_{23}).$$
$$(e): q_{A_{21}}(A_{21}|A_2) \equiv p(A_{21}|A_2); q_{B_{21}}(B_{21}|B_2) \equiv p(B_{21}|B_2); q_{A_{21}, B_{21}}(A_{21}, B_{21}) \equiv p(A_{21}, B_{21});$$
$$q_{C_{21}}(C_{21}|C_2, B_{21}, A_{21}) \equiv p(C_{21}|C_2, B_{21}, A_{21}); q_{S_{21}}(S_{21}) \equiv p(S_{21}).$$
(12)

Applying the Nested Markov factorization on the trajectory in 8 (f), we obtain the following factorization:

$$p(A_1, B_1, C_1) \cdot p(A_{12}, B_{12}, C_{12}|A_1, B_1, C_1) \cdot p(A_{23}, B_{23}, C_{23}|A_{12}, B_{12}, C_{12})$$
$$= \underbrace{\{q_{A_1, B_1}(A_1, B_1)q_{C_1}(C_1 \mid A_1, B_1)\}}_{(a)} \cdot \underbrace{\{q_{A_{12}, C_{12}}(A_{12}, C_{12}|B_{12}, A_1, C_1)q_{B_{12}}(B_{12}|A_{12})\}}_{(c)} \cdot$$
$$\underbrace{\{q_{A_{23}}(A_{23}|A_{12}, C_{12}) \cdot q_{A_{23}}(B_{23}|B_{12}, A_{23}) \cdot q_{C_{23}}(C_{23}|C_{12}, A_{23})\}}_{(d)},$$

where the kernels are given in (11) and (12) above.

C   Proofs

Lemma 1 *Under Assumption 1, $p(\mathbf{Y}(\mathbf{a}))$ is identified from a hidden variable causal DBN model represented by latent projections $\mathcal{G}_1$ on $\mathbf{V}_1$ and $\mathcal{G}_{t+1,t}$ on $\mathbf{V}_{t+1}$ given $\mathbf{V}_t$ if and only if every bidirected connected component in $\mathcal{G}_{1\mathbf{Y}_1^*}$ (the induced subgraph of $\mathcal{G}_1$) is intrinsic in $\mathcal{G}_1$, and every bidirected component in $\mathcal{G}_{t+1,t\mathbf{Y}_i^*}$ (the induced subgraph of $\mathcal{G}_{t+1,t}$) is intrinsic in $\mathcal{G}_{t+1,t}$, where*

$\mathbf{Y}_1^*$ is the set of ancestors of $\mathbf{Y} \cap \mathbf{V}_1$ not through $\mathbf{A} \cap \mathbf{V}_1$ in $\mathcal{G}_1$, and for every $i \in 2, \ldots, T$, $\mathbf{Y}_i^*$ is the set of ancestors of $\mathbf{Y} \cap \mathbf{V}_i$ not through $\mathbf{A} \cap \mathbf{V}_i$ in $\mathcal{G}_{t+1,t}$. Moreover, if $p(\mathbf{Y}(\mathbf{a}))$ is identified, we have

$$\left(\sum_{\mathbf{Y}_1^* \setminus ((\mathbf{Y} \cup \mathbf{A}) \cap \mathbf{V}_1)} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{1\mathbf{Y}_1^*})} q_{\mathbf{D}}^1(\mathbf{D}|\operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})|_{\mathbf{A}=\mathbf{a}} \right) \times$$

$$\prod_{i=2}^T \left(\sum_{\mathbf{Y}_i^* \setminus ((\mathbf{Y} \cup \mathbf{A}) \cap \mathbf{V}_i)} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{t+1,t\mathbf{Y}_i^*})} q_{\mathbf{D}}^{t+1,t}(\mathbf{D}|\operatorname{pa}_{\mathcal{G}}(\mathbf{D}) \setminus \mathbf{D})|_{\mathbf{A}=\mathbf{a}}, \right)$$

where $q_{\mathbf{D}}^1$ and $q_{\mathbf{D}}^{t+1,t}$ are kernels corresponding to intrinsic sets representing elements of $\mathcal{D}(\mathcal{G}_{1\mathbf{Y}_1^*})$ and $\mathcal{D}(\mathcal{G}_{t+1,t\mathbf{Y}_1^*})$ in the nested Markov factorizations of $\mathcal{G}_1$ and $\mathcal{G}_{t+1,t}$, respectively.

Proof: We want to obtain $p(\mathbf{Y}(\mathbf{a}))$ from the observed joint $p(\mathbf{V}_{1:T})$. Using identification result 5 on the unrolled ADMG gives $\sum_{Y^* \setminus Y} p(\mathbf{Y}^*(a)) = \sum_{\mathbf{Y}^* \setminus \mathbf{Y}} \prod_{D \in \mathcal{D}(\mathcal{G}_{\text{unrolled}Y^*})} p(\mathbf{D}(\operatorname{pa}(\mathbf{D}) \setminus \mathbf{D}))|_{\mathbf{A}=\mathbf{a}}$. Assumption 1 ensures that no district $\mathbf{D}$ spans time points, and parents $\operatorname{pa}(\mathbf{D})$ at time $t+1$ lie either at t or $t+1$. This allows us to write $\sum_{\mathbf{Y}^* \setminus \mathbf{Y}} p(\mathbf{Y}^*(\mathbf{a})) = \sum_{\mathbf{Y}^* \setminus \mathbf{Y}} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{1Y^*})} p(\mathbf{D}(\operatorname{pa}(\mathbf{D}) \setminus \mathbf{D}))|_{\mathbf{A}=\mathbf{a}} \times \prod_{t=1}^{T-1} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{t+1,tY^*})} p(\mathbf{D}(\operatorname{pa}(\mathbf{D}) \setminus \mathbf{D}))|_{\mathbf{A}=\mathbf{a}}$. Applying the identification results in Richardson et al. (2012) to the prior network ADMG $\mathcal{G}_1$ and extensions of these results in Sherman and Shpitser (2018) to the transition network CADMGs $\mathcal{G}_{t+1,t}$, these counterfactual conditionals can be replaced by given modified nested factorizations, provided every appropriate bidirected connected set in the prior or transition graph is intrinsic in that graph.   □

Lemma 2 *Given a fully observed PDSEM, each factor of the distribution $p_\infty(\mathbf{Y}(\mathbf{a}))$ is identified from $p_\infty(\mathbf{V})$ as:*

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) \equiv \prod_{V \in \mathbf{Y}_1 \setminus \mathbf{A}_1} p_1(V|\operatorname{pa}_{\mathcal{G}_1}(V))\Big|_{\mathbf{A}_1=\mathbf{a}_1}$$

$$p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)) \equiv \prod_{V \in \mathbf{Y}_{ij} \setminus \mathbf{A}_j} p_{ij}(V|\operatorname{pa}_{\mathcal{G}_{ij}}(V))\Big|_{\substack{\mathbf{A}_i=\mathbf{a}_i, \\ \mathbf{A}_j=\mathbf{a}_j}}$$

Proof: This follows from the factorization of $p_\infty(\mathbf{V}(\mathbf{a}))$ into elements of the form $p_1(\mathbf{Y}_1(\mathbf{a}_1))$, and $p_{ij}(\mathbf{Y}_j(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a_i}))$, the fact that $\mathcal{G}_1, \{\mathcal{G}_{ij} : (i,j) \in \mathcal{T}\}$ define causal models under standard structural equation semantics, and equation 1 .   □

Lemma 3 *Under Assumptions 1, 2 and 3, given a latent variable PDSEM represented by $\mathcal{G}_1$ and $\{\mathcal{G}_{ij} : (i,j) \in \mathcal{T}\}$, $p_\infty(\mathbf{Y}(\mathbf{a}))$ is identified from $p_\infty(\mathbf{V})$ if and only if every bidirected component in $\mathcal{G}_{1\mathbf{Y}_1}$ is intrinsic in $\mathcal{G}_1$, and every bidirected component in $\mathcal{G}_{ij\mathbf{Y}_j}$ is intrinsic in $\mathcal{G}_{ij}$ for every i and j. Moreover, if $p_\infty(\mathbf{Y}(\mathbf{a}))$ is identified, it is equal to*

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) \prod_{t=1}^\infty \left(\prod_{(i,j) \in \mathcal{T}} (p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)))^{\mathbb{I}(s_{t-1}^i, s_t^j)} \right) \mathbf{1}^{\mathbb{I}(s_{t-1}^*)} \tag{13}$$

where

$$p_1(\mathbf{Y}_1(\mathbf{a}_1)) = \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{1\mathbf{Y}_1^*})} q_{\mathbf{D}}^1(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_1}^s(\mathbf{D}))\Big|_{\mathbf{A}_1=\mathbf{a}_1}, \tag{14}$$

where each kernel $q_{\mathbf{D}}^1(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_1}^s(\mathbf{D}))$ is in the nested Markov factorization of $p_1(\mathbf{V}_1)$ with respect to $\mathcal{G}_1$, and

$$p_{ij}(\mathbf{Y}_{ij}(\mathbf{a}_j)|\mathbf{Y}_i(\mathbf{a}_i)) = \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{V}_{ij} \setminus \mathbf{A}_{ij}})} q_{\mathbf{D}}^{ij}(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_{ij}}^s(\mathbf{D}))\Big|_{\substack{\mathbf{A}_i=\mathbf{a}_i, \\ \mathbf{A}_j=\mathbf{a}_j}}, \tag{15}$$

where each kernel $q_{\mathbf{D}}^{ij}(\mathbf{D}|\operatorname{pa}_{\mathcal{G}_{ij}}^s(\mathbf{D}))$ is in the nested Markov factorization of $p_{ij}(\mathbf{V}_{ij}|\mathbf{V}_i)$ with respect to $\mathcal{G}_{ij}$.

Proof: Assumption 3 implies all state transitions are known, and thus allows us to proceed by induction on any sequence of state transitions with positive probability after t steps.

Unrolling the prior network, and appropriate transition networks for such a sequence yields an ADMG representing the observed data distribution had that transition taken place, with Assumption 1 implying that districts in this ADMG do not span multiple time steps. This immediately implies the conclusion by the same argument used in the proof of Lemma 1.

In fact, this argument works for any transition sequence of any size.   □

D   The Septoplasty Surgical Procedure, and its PDSEM Model

Septoplasty is a surgical procedure performed on the nasal cartilage, called the septum, to relieve nasal obstruction (Tajudeen and Kennedy, 2017). A deviated or deformed septum is the most common cause of such an obstruction. Apart from nasal obstruction, a significantly deviated nasal septum has also been implicated in epistaxis, sinusitis, obstructive sleep apnea, and headaches which can act as diagnosis factors. The procedure involves cartilage resection, modification or a graft. The outcome

of septoplasty is typically a score/index constructed from a questionnaire investigating quality of life measures and perceived nasal obstruction levels, like Nasal Obstruction Septoplasty Effectiveness (NOSE) and the Fairley Nasal Questionnaire (FNQ) (Fettman et al., 2009).

For instructional and evaluation purposes, surgeries are often divided into discrete steps or "stages", each with its own intermediate goal (Ahmidi et al., 2015). Our data from the septoplasty procedure was manually annotated by clinical experts and divided into the following states:

- s_1: opening of the septum,

- s_2: raising septal flaps,

- s_3: removal of deviated septal cartilage and bone,

- s_4: reconstruction,

- s_5: closing of the incision,

- s_6: activity not otherwise included in the above 5 phases,

- s_{end}: end of surgery state (which contains no variables).

The variables in our data are the following: $\mathbf{V}$ = { K: knife, G: gorney scissors, C_1: cottle, D_1: short needle driver, D_2: long needle driver, O: other tools, C_2: suction cannula, M: main surgeon exists, S: suction exists, A_1: main surgeon is an attending, A_2: suction done by attending, T: duration of that phase is greater than 10 seconds }

- $\mathbf{V}_{s_1} = \{K, O, C_2, M, S, A_1, A_2, T\}$,

- $\mathbf{V}_{s_2} = \{K, C_1, O, C_2, M, S, A_1, A_2, T\}$,

- $\mathbf{V}_{s_3} = \{K, C_1, D_1, D_2, O, C_2, M, S, A_1, A_2, T\}$,

- $\mathbf{V}_{s_4} = \{K, C_1, G, O, C_2, M, S, A_1, A_2, T\}$,

- $\mathbf{V}_{s_5} = \{D_1, D_2, O, C_2, M, S, A_1, A_2, T\}$,

- $\mathbf{V}_{s_6} = \{K, C_1, O, C_2, M, S, A_1, A_2, T\}$,

To determine the allowed state transitions , we retained observed data state transitions where at least 5 such transitions occurred. The permitted state transitions $s_i \rightarrow s_j$ are summarized in Figure 7 in the main paper – note that transitions other than those depicted have probability $p(s_j \mid s_i, \mathbf{v}_{s_i}) = 0$ for all $\mathbf{v}_{s_i}$. To determine the state transition distributions $p(s_j \mid s_i, \mathbf{v}_{s_i})$, we restricted the set $\mathbf{v}_{s_i}$ for all i to be $\{A_1, M, A_2, S\}$ to increase tractability of estimation, and estimated this discrete conditional distribution via a conditional probability table. The prior distribution on the initial state was set to $p(s_1) = 1$.

State DAGs were determined based on clinician recommendation and have been reproduced in Figure 9 for reference. These immediately lead to prior variable distributions $p(\mathbf{V}_{s_i})$ for each state s_i.

Transition graphs from $s_i \rightarrow s_j$ are constructed using a simple rule: the for any variable v in any state s_j, the parents $\mathrm{pa}(v_{s_j})$ consists of the variable with the same name in the previous state s_i if it exists, and all parents in the state DAG for s_j point indicated by state DAGs. For example, in the transition $s_1 \rightarrow s_1$ moving from time step $t - 1 \rightarrow t$, variable K at time step t has parents A_1 at t, as given in Figure 9(a), as well as K from time step $t - 1$. However, in the transition $s_1 \rightarrow s_2$ moving from time step $t - 1$ to time step t, variable C_1 has parent A_1 in time step t, as given in Figure 9(b), but no parents from the previous time step $t - 1$ since C_1 does not exist in s_1. Based on this rule, probability distributions $p(v_{s_j} \mid \mathrm{pa}(v_{s_j}))$ are estimated using conditional probability tables.

Goodness of fit of our model with respect to the original data distribution is shown in Figure 10. Trajectories simulated by our model are able to capture the distribution of surgery duration originally seen in the data, quite well.

Implications for Statistical Inference For Latent Variable PDSEMs: If a PDSEM is fully observed, causal inference may be performed by obtaining maximum likelihood estimates $\hat{\eta}$ of all parameters, and evaluating the g-formula functionals using Monte-Carlo sampling using learned distributions of the form $p(V \mid \mathrm{pa}_\mathcal{G}(V); \eta_V)$. This method is computationally efficient as long as the initial DAG and transition CDAGs in a PDSEM are sufficiently sparse. Indeed, our data application was based on this approach.

However, an analogous approach is not straightforward for nested Markov parameterizations of the marginal PDSEM representing a PDSEM with hidden variables. In our simulations, we use a specific generative model for our continuous variables, i.e, the linear Gaussian Structural Equation model. Another choice based on work in (Evans and Richardson, 2014) is the Möbius parameterization for binary variables. However, this is ill-suited for drawing samples. Instead, existing approaches to sampling
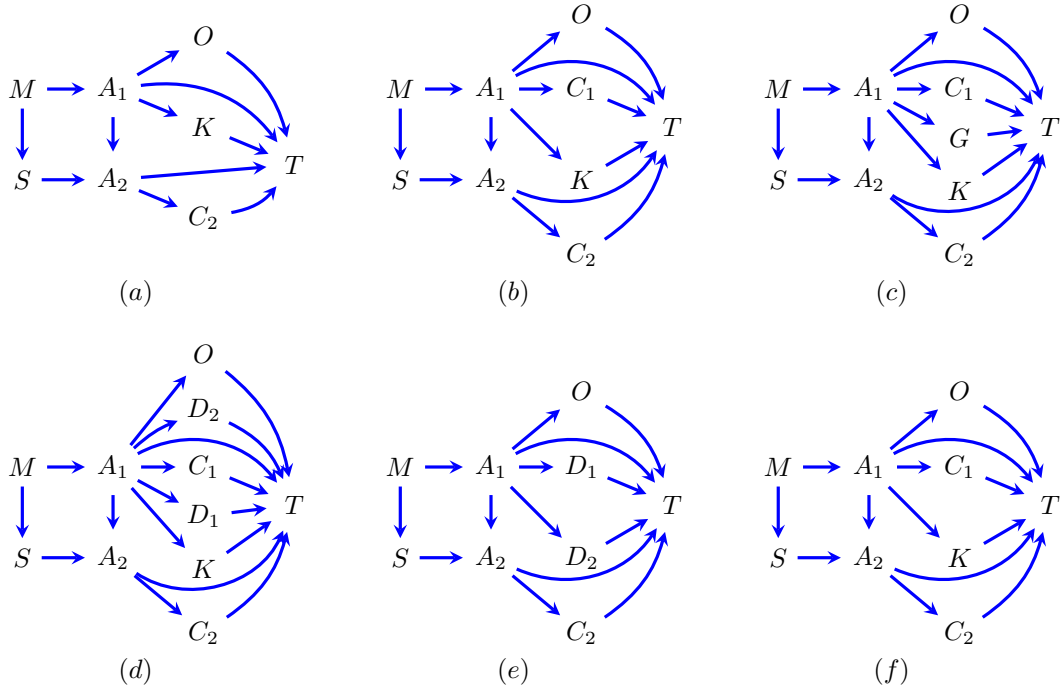
Figure 9: (a)-(f) State DAGs corresponding to states s^1: Opening of the septum, s^2: Raising septal flaps, s^3: Removal of deviated septal cartilage and bone, s^4: Reconstruction, s^5: Closing of incision and s^6: Activity not otherwise included in the above phases.
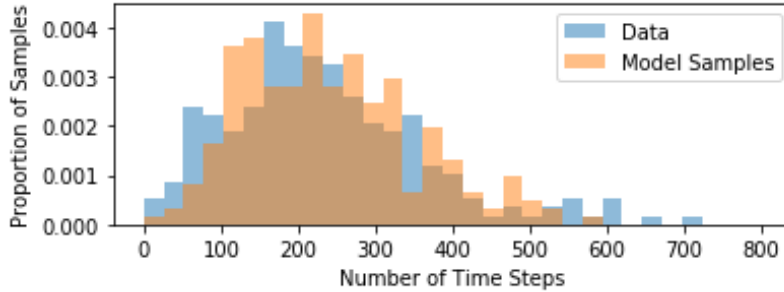

Figure 10: Histograms of observed surgery (blue) versus simulated surgeries from the estimated model (orange).

from a nested Markov discrete likelihood involve first converting the likelihood expressed in terms of the Möbius parameters to one expressed as a the joint distribution $p(\mathbf{V})$ (from which it is easy to generate samples for a discrete sample space of $\mathbf{V}$). Importantly, such a conversion leads to an intractable object that requires storage and running time exponential in $|\mathbf{V}|$. This holds *even if* the underlying model dimension of the nested Markov model is small. The situation is radically different from that of DAG models, where a small model dimension directly leads to a computationally efficient sampling scheme. For settings beyond Gaussian and discrete data, statistical inference strategies are significantly more complicated and have been discussed in Bhattacharya et al. (2020).

While there exist promising approaches, based on the nested Markov generalization of the variable elimination algorithm (Shpitser et al., 2011), in general the problem remains open.

E   Computation Details

The septoplasty data application presented in Section 6 was computed on a Lenovo X1 Carbon with an Intel i7 1.8 GHz processor and 16 GB of RAM. Computation for each scenario (generating from the model without interventions, attending performing the whole surgery, and trainee performing the whole surgery) took between 1.5 to 2 hours each.