


PAPER

Causal inference via string diagram surgery

A diagrammatic approach to interventions and counterfactuals

Bart Jacobs^{1*} , Aleks Kissinger², and Fabio Zanasi³

¹Department of Computing and Information Science, Radboud Universiteit Faculteit der Natuurwetenschappen Wiskunde en Informatica, Nijmegen, The Netherlands, ²Department of Computer Science, Oxford University, Oxford, UK and

³Department of Computer Science, University College London, London, UK

*Corresponding author. Email: bart@cs.ru.nl

(Received 31 January 2020; revised 1 August 2021; accepted 26 August 2021; first published online 16 November 2021)

Abstract

Extracting causal relationships from observed correlations is a growing area in probabilistic reasoning, originating with the seminal work of Pearl and others from the early 1990s. This paper develops a new, categorically oriented view based on a clear distinction between syntax (string diagrams) and semantics (stochastic matrices), connected via interpretations as structure-preserving functors. A key notion in the identification of causal effects is that of an intervention, whereby a variable is forcefully set to a particular value independent of any prior propensities. We represent the effect of such an intervention as an endo-functor which performs ‘string diagram surgery’ within the syntactic category of string diagrams. This diagram surgery in turn yields a new, interventional distribution via the interpretation functor. While in general there is no way to compute interventional distributions purely from observed data, we show that this is possible in certain special cases using a calculational tool called comb disintegration. We demonstrate the use of this technique on two well-known toy examples: one where we predict the causal effect of smoking on cancer in the presence of a confounding common cause and where we show that this technique provides simple sufficient conditions for computing interventions which apply to a wide variety of situations considered in the causal inference literature; the other one is an illustration of counterfactual reasoning where the same interventional techniques are used, but now in a ‘twinned’ set-up, with two version of the world – one factual and one counterfactual – joined together via exogenous variables that capture the uncertainties at hand.

Keywords: Probabilistic reasoning; causality; counterfactuals; string diagrams

1. Introduction

Causality is about understanding the mechanics of the world around us. This world presents itself in the form of streams of observations, in which statistical (in)dependencies can be recognised. A big question, both in science and in daily life, is: how to distinguish correlation from causation and recognise genuine causal relationships?

An important conceptual tool for distinguishing correlation from causation is the possibility of *intervention*, that is, forcing some variable to take a specific value in a way that is independent of the other variables being considered. For example, a randomised drug trial attempts to destroy any confounding ‘common cause’ explanation for correlations between drug use and recovery by randomly assigning a patient to the control or treatment group, independent of any background

factors. In an ideal setting, the observed correlations of such a trial will reflect genuine causal influence. Unfortunately, it is not always possible (or ethical) to ascertain causal effects by means of actual interventions. For instance, one is unlikely to get ethical approval to run a clinical trial on whether smoking causes cancer by randomly assigning 50% of the patients to smoke and waiting a bit to see who gets cancer. However, in certain situations, it is possible to predict the effect of such a hypothetical intervention from purely observational data, together with an assumed underlying graph structure.

In this paper, we will focus on the problem of *causal identifiability*. For this problem, we are given observational data as a joint distribution on a set of variables and we are furthermore provided with a *causal structure* associated with those variables. This structure, which typically takes the form of a directed acyclic graph or some variation thereof, tells us which variables can in principle have a causal influence on others. The problem then becomes whether we can measure how strong those causal influences are, by means of computing an *interventional* distribution. That is, can we ascertain what would have happened if a (hypothetical) intervention had occurred?

Note that this is distinct from the related problem of *causal discovery*, where a causal structure is not given from the start, but must be discovered purely from the observational data, subject to certain well-behavedness assumptions (Pearl, 2000). Causal identifiability assumes that the causal structure is already provided, either from previously doing causal discovery or by making use of some additional knowledge about the problem at hand. In particular, this means that causal identifiability is trivial when all variables are observed. However, it becomes a difficult and important problem in the presence of confounding variables (unobserved common causes) or selection bias (conditioning on common effects). In this paper, we will focus on the former.

Over the past 3 decades, a great deal of work has been done in identifying necessary and sufficient conditions for causal identifiability in the presence of confounding variables, starting with very specific special cases such as the *back-door* and *front-door* criteria (Pearl, 2000) and progressing to more general necessary and sufficient conditions for causal identifiability based on the **do**-calculus (Huang and Valtorta, 2012), or combinatoric concepts such as confounded components in semi-Markovian models (Shpitser and Pearl, 2006; Tian and Pearl, 2002).

This style of causal reasoning relies crucially on a delicate interplay between syntax and semantics, which is often not made explicit in the literature. The syntactic object of interest is the causal structure (e.g. a causal graph), which captures something about our understanding of the world, and the mechanisms which gave rise to some observed phenomena. The semantic object of interest is the data: joint and conditional probability distributions on some variables. Fixing a causal structure entails certain constraints on which probability distributions can arise, hence it is natural to see distributions satisfying those constraints as models of the syntax.

In this paper, we make this interplay precise using functorial semantics in the spirit of Lawvere Lawvere (1963) and develop basic syntactic and semantic tools for causal reasoning in this setting. We take as our starting point a functorial presentation of Bayesian networks similar to the one appearing in Fong (2012). The syntactic role is played by string diagrams, which give an intuitive way to represent morphisms of a monoidal category as boxes plugged together by wires. Given a directed acyclic graph (dag) G , we can form a free category Syn_G whose arrows are (formal) string diagrams which represent the causal structure syntactically. Structure-preserving functors from Syn_G to Stoch , the category of stochastic matrices, then correspond exactly to Bayesian networks based on the dag G .

Within this framework, we develop the notion of intervention as an operation of ‘string diagram surgery’. Intuitively, this cuts a string diagram at a certain variable, severing its link to the past. Formally, this is represented as an endofunctor on the syntactic category $\text{cut}_X: \text{Syn}_G \rightarrow \text{Syn}_G$, which propagates through a model $F: \text{Syn}_G \rightarrow \text{Stoch}$ to send observational probabilities $F(\omega)$ to interventional probabilities $F(\text{cut}_X(\omega))$.

The cut_x endofunctor gives us a diagrammatic means of computing interventional distributions given complete knowledge of F . However, more interestingly, we can sometimes compute interventionals given only partial knowledge of F , namely some observational data. We show that this can also be done via a technique we call *comb disintegration*, which is a string diagrammatic version of a technique called *c-factorisation* introduced by (Tian and Pearl, 2002). Our approach generalises disintegration, a calculational tool whereby a joint state on two variables is factored into a single-variable state and a channel, representing the marginal and conditional parts of the distribution, respectively. Disintegration has recently been formulated categorically in Clerc et al. (2017) and using string diagrams in Cho and Jacobs (2019). We take the latter as a starting point, but instead consider a factorisation of a three-variable state into a channel and a *comb*. The latter is a special kind of map which allows inputs and outputs to be interleaved. They were originally studied in the context of quantum communication protocols, seen as games (Gutoski and Watrous, 2007), but have recently been used extensively in the study of causally ordered quantum (Chiribella et al., 2008; Perinotti, 2016) and generalised (Kissinger and Uijlen, 2017) processes. While originally imagined for quantum processes, the categorical formulation given in Kissinger and Uijlen (2017) makes sense in both the classical case (Stoch) and the quantum. Much like Tian and Pearl's technique, comb factorisation allows one to characterise when the confounding parts of a causal structure are suitably isolated from each other, then exploit that isolation to perform the concrete calculation of interventional distributions.


However, unlike in the traditional formulation, the syntactic and semantic aspects of causal identifiability within our framework are connected. Namely, we can give conditions for causal identifiability in terms of factorisation of a morphism in Syn_G , whereas the actual concrete computation of the interventional distribution involves factorisation of its interpretation in Stoch. Thanks to the functorial semantics, the former immediately implies the latter.

The interventional techniques in terms of string diagrams and their interpretations can also be applied to counterfactual queries. There, we use two copies of the relevant string diagram, one for the 'actual' and one for the 'counterfactual' world. These two copies are connected via some shared states, which capture the sense in which random variables should take the same value in the real and counterfactual world. This approach, which is also described in functional form Balke and Pearl (1994); Pearl and Verma (1991), is elaborated here in terms of factorisation of stochastic matrices into a deterministic part and an exogenous state. By sharing these states between the two copies of the world, we are able to more effectively make use of our knowledge of what *did* happen in order to predict what *would have* happened had some past detail been different. One standard example from the literature is elaborated as illustration of how this works.

To introduce our framework itself, we make use of a running example taken from Pearl's book (Pearl, 2000): identifying the causal effect of smoking on cancer with the help of an auxiliary variable (the presence of tar in the lungs). After providing some preliminaries on stochastic matrices and the functorial presentation of Bayesian networks in Sections 2 and 3, we introduce the smoking example in Section 4. In Section 5, we formalise the notion of intervention as string diagram surgery, and in Section 6 we introduce the combs and prove our main calculational result: the existence and uniqueness of comb factorisations. In Section 7, we show how to apply this theorem in computing the interventional distribution in the smoking example. In Section 8, we provide a more general version of the theorem, which captures (and slightly generalises) the conditions given in Tian and Pearl (2002). In Section 9, we focus on counterfactual reasoning, illustrating how it can also be modelled with string diagram surgery. Finally, in Section 10, we summarise our results and describe several avenues of future work.

This work is an extended version of the conference paper (Jacobs et al., 2019), which includes all the missing proofs and a completely new part on counterfactuals (Section 9).

2. Stochastic Matrices and Conditional Probabilities

Symmetric monoidal categories (SMCs) give a very general setting for studying processes which can be composed in sequence (via the usual categorical composition \circ) and in parallel (via the monoidal composition \otimes). Throughout this paper, we will use *string diagram* notation (Selinger, 2011) for depicting composition of morphisms in an SMC. In this notation, morphisms are depicted as boxes with labelled input and output wires, composition \circ as ‘plugging’ boxes together, and the monoidal product \otimes as placing boxes side by side. Identity morphisms are depicted simply as a wire and the unit I of \otimes as the empty diagram. The ‘symmetric’ part of the structure consists of symmetry morphisms, which enable us to permute inputs and outputs arbitrarily. We depict these as wire-crossings: . Morphisms whose domain is I are called *states*, and they will play a special role throughout this paper.

A monoidal category of prime interest in this paper is *Stoch*, whose objects are finite sets and morphisms $f : A \rightarrow B$ are $|B| \times |A|$ dimensional stochastic matrices. That is, they are matrices of positive numbers (including 0) whose columns each sum to 1:

$$f = \{f_i^j \in \mathbb{R}^+ \mid i \in A, j \in B\} \quad \text{with} \quad \sum_j f_i^j = 1, \text{ for all } i.$$

Note we adopt the physicists convention of writing row indices as superscripts and column indices as subscripts. Stochastic matrices are of interest for probabilistic reasoning, because they exactly capture the data of a conditional probability distribution. That is, if we take $A := \{1, \dots, m\}$ and $B := \{1, \dots, n\}$, conditional probabilities naturally arrange themselves into a stochastic matrix:

$$f_i^j := P(B = j \mid A = i) \quad \rightsquigarrow \quad f = \begin{pmatrix} P(B = 1 \mid A = 1) & \cdots & P(B = 1 \mid A = m) \\ \vdots & \ddots & \vdots \\ P(B = n \mid A = 1) & \cdots & P(B = n \mid A = m) \end{pmatrix}$$

States, that is, stochastic matrices from a trivial input $I := \{*\}$, are (non-conditional) probability distributions, represented as column vectors. There is only one stochastic matrix with trivial output: the row vector consisting only of 1’s. The latter, with notation \uparrow as on the right, will play a special role in this paper (see (1) below).

Composition of stochastic matrices is matrix multiplication. In terms of conditional probabilities, this corresponds to multiplication, followed by marginalisation over the shared variable: $\sum_B P(C \mid B)P(B \mid A)$. Identities are therefore given by identity matrices, which we will often express in terms of the Kronecker delta function δ_i^j .

The monoidal product \otimes in *Stoch* is the cartesian product on objects, and Kronecker product of matrices: $(f \otimes g)_{(i,j)}^{(k,l)} := f_i^k g_j^l$. We will typically omit parentheses and commas in the indices, writing for example h_{ij}^{kl} instead of $h_{(i,j)}^{(k,l)}$ for an arbitrary matrix entry of $h : A \otimes B \rightarrow C \otimes D$. In terms of conditional probabilities, the Kronecker product corresponds to taking product distributions. That is, if f represents the conditional probabilities $P(B \mid A)$ and g the probabilities $P(D \mid C)$, then $f \otimes g$ represents $P(B \mid A)P(D \mid C)$. *Stoch* also comes with a natural choice of ‘swap’ matrices $\sigma : A \otimes B \rightarrow B \otimes A$ given by $\sigma_{ij}^{kl} := \delta_i^l \delta_j^k$, making it into a SMC. Every object A in *Stoch* has three other pieces of structure which will play a key role in our formulation of Bayesian networks and interventions: the *copy* map, the *discarding* map and the *uniform state*:

$$\left(\begin{array}{c} \diagup \\ \bullet \\ \diagdown \end{array} \right)_i^{jk} := \delta_i^j \delta_i^k \qquad \left(\begin{array}{c} \bullet \\ \uparrow \end{array} \right)_i := 1 \qquad \left(\begin{array}{c} \downarrow \\ \bullet \end{array} \right)^i := \frac{1}{|A|} \tag{1}$$

Abstractly, this provides *Stoch* with the structure of a *CDU category*.

Definition 1. A CDU category (for *copy*, *discard*, *uniform*) is a symmetric monoidal category $(\mathcal{C}, \otimes, I)$ where each object A has a copy map $\Upsilon : A \rightarrow A \otimes A$, a discarding map $\uparrow : A \rightarrow I$, and a uniform state $\Downarrow : I \rightarrow A$ satisfying the following equations:

Equation (2) consists of three parts. The first part shows the copy map Υ as a cup with two dots on the top wire, equal to a cup with one dot on the top wire and one dot on the bottom wire. The second part shows the discard map \uparrow as a dot on a wire, equal to a vertical wire. The third part shows the uniform state map \Downarrow as a dot on a wire with a downward arrow, equal to a dashed square box.

CDU functors are symmetric monoidal functors between CDU categories preserving copy maps, discard maps and uniform states.¹

We assume that the CDU structure on I is trivial and the CDU structure on $A \otimes B$ is constructed in the obvious way from the structure on A and B . We also use the first equation in (2) to justify writing ‘copy’ maps with arbitrarily many output wires: Υ .

Similar to Bonchi et al. (2018), we can form the free CDU category $\text{FreeCDU}(A, \Sigma)$ over a pair (X, Σ) of a generating set of objects X and a generating set Σ of typed morphisms $f : u \rightarrow w$, with $u, w \in X^*$ as follows. The category $\text{FreeCDU}(A, \Sigma)$ has X^* as set of objects, and morphisms the string diagrams constructed from the elements of Σ and maps $\Upsilon : x \rightarrow x \otimes x$, $\uparrow : x \rightarrow I$ and $\Downarrow : I \rightarrow x$ for each $x \in X$, taken modulo the equations (2).

Lemma 2. *Stoch is a CDU category, with CDU structure defined as in (1).*

An important feature of *Stoch* is that $I = \{\star\}$ is the final/terminal object, with $\uparrow : B \rightarrow I$ the map provided by the universal property, for any set B . This yields equation (3) on the right, for any $f : A \rightarrow B$, justifying the name ‘discarding map’ for \uparrow .

Equation (3) shows a box labeled f with input A and output B , equal to a vertical wire with a dot on top and a label B on the right.

We conclude by recording another significant feature of *Stoch*: *disintegration* (Cho and Jacobs, 2019; Clerc et al., 2017). In probability theory, this is the mechanism of factoring a joint probability distribution $P(AB)$ as a product of the first marginal $P(A)$ and a conditional distribution $P(B|A)$. We recall from Cho and Jacobs (2019) the string diagrammatic rendition of this process. We say that a morphism $f : X \rightarrow Y$ in *Stoch* has *full support* if, as a stochastic matrix, it has no zero entries. When f is a state, it is a standard result that full support ensures uniqueness of disintegrations of f .

Proposition 3. (Disintegration). *For any state $\omega : I \rightarrow A \otimes B$ in *Stoch* with full support, there exists unique morphisms $\mathbf{a} : I \rightarrow A$, $\mathbf{b} : A \rightarrow B$ such that:*

Equation (4) shows a box labeled ω with inputs A and B , equal to a diagram where a wire from \mathbf{a} goes up and left to a dot, then right to a box labeled \mathbf{b} , and finally down to the B input of ω .

Note that equation (3) and the CDU rules immediately imply that the unique $\mathbf{a} : I \rightarrow A$ in Proposition 3 is the marginal of ω onto A : $\begin{array}{|c|} \hline A \\ \hline \omega \\ \hline \end{array}$.


3. Bayesian Networks as String Diagrams

Bayesian networks are a widely used tool in probabilistic reasoning. They give a succinct representation of conditional (in)dependences between variables as a directed acyclic graph. Traditionally,

a Bayesian network on a set of variables A, B, C, \dots is defined as a directed acyclic graph (dag) G , an assignment of finite sets to each of the nodes $V_G := \{A, B, C, \dots\}$ of G and a joint probability distribution over those variables which factorises as $P(V_G) = \prod_{A \in V_G} P(A | \text{Pa}(A))$ where $\text{Pa}(A)$ is the set of parents of A in G . Any joint distribution that factorises this way is said to satisfy the *global Markov property* with respect to the dag G . Alternatively, a Bayesian network can be seen as a dag equipped with a set of conditional probabilities $\{P(A | \text{Pa}(A)) \mid A \in V_G\}$ which can be combined to form the joint state. Disintegration allows us to switch back and forth between joint states and conditional probabilities and thus yields the equivalence between these two perspectives.


Much like in the case of disintegration in the previous section, Bayesian networks have a categorical description as string diagrams in the category *Stoch*. This perspective has been widely adopted – with some variations – in recent years (Cabrera *et al.*, 2018; Coecke and Spekkens, 2012; Fong, 2012; Fritz, 2020; Jacobs and Zanasi, 2021, 2016, 2017). We offer here a presentation inspired by functorial semantics of algebraic theories (Lawvere, 1969).

Let us start with an example. Here is a Bayesian network in its traditional depiction as a dag with an associated joint distribution over its vertices, and as a string diagram in *Stoch*:



In the string diagram above, the stochastic matrix $\mathbf{a}: I \rightarrow A$ contains the probabilities $P(A)$, $\mathbf{b}: B \rightarrow A$ contains the conditional probabilities $P(B|A)$, $\mathbf{c}: B \otimes D \rightarrow C$ contains $P(C|BD)$, and so on. The entire diagram is then equal to a state $\omega: I \rightarrow A \otimes B \otimes C \otimes D \otimes E$ which represents $P(ABCDE)$.

Note the dag and the diagram above look similar in structure. The main difference is the use of copy maps to make each variable (even those that are not leaves of the dag, A, B and D) an output of the overall diagram. This corresponds to a variable being *observed*. We can also consider Bayesian networks with *latent* variables, which do not appear in the joint distribution due to marginalisation. Continuing the example above, making A into a latent variable yields the following depiction as a string diagram:



In general, a Bayesian network (with possible latent variables), is a string diagram in *Stoch* that (1) only has outputs and (2) consists only of copy maps and boxes which each have exactly one output.

By ‘a string diagram in *Stoch*’, we mean not only the stochastic matrix itself, but also its decomposition into components. We can formalise exactly what we mean by taking a perspective on Bayesian networks which draws inspiration from functorial semantics of algebraic

theories (Lawvere, 1969). In this perspective, which mainly elaborates on (Fong (2012) [Ch. 4]), we maintain a conceptual distinction between the purely syntactic object (the diagram) and its probabilistic interpretation.

Starting from a dag $G = (V_G, E_G)$, we construct a free CDU category Syn_G which provides the syntax of causal structures labelled by G . The objects of Syn_G are generated by the vertices of G , whereas the morphisms are generated by the following signature:

$$\Sigma_G = \left\{ \begin{array}{c} | \\ \boxed{a} \\ | \\ \dots \\ | \\ B_1 \dots B_k \end{array} \Bigg| A \in V_G \text{ with parents } B_1, \dots, B_k \in V_G \right\}$$

Then $\text{Syn}_G := \text{FreeCDU}(V_G, \Sigma_G)$. The following result establishes that models (*à la* Lawvere) of Syn_G coincide with G -based Bayesian networks.

Proposition 4. *There is a 1-1 correspondence between Bayesian networks based on the dag G and CDU functors of type $\text{Syn}_G \rightarrow \text{Stoch}$.*

Proof. In one direction, consider a Bayesian network consisting of the dag G and, for each node $A \in V_G$, an assignment of a finite set $\tau(A)$ and a conditional probability $P(A|\text{Pa}(A))$. These data yield a CDU functor $F : \text{Syn}_G \rightarrow \text{Stoch}$, defined by the following mappings:

$$F = \left\{ \begin{array}{l} A \in V_G \mapsto \tau(A) \\ \begin{array}{c} | \\ \boxed{a} \\ | \\ \dots \\ | \\ B_1 \dots B_k \end{array} \mapsto (f_{i_1 \dots i_n}^j := P(A = j | \text{Pa}(A) = (i_1, \dots, i_n))) \end{array} \right.$$

Conversely, let $F : \text{Syn}_G \rightarrow \text{Stoch}$ be a CDU functor. This defines a G -based Bayesian network by setting $\tau(A) := F(A)$ and $P(A = j | \text{Pa}(A) = (i_1, \dots, i_n)) := F(a)_{i_1 \dots i_n}^j$. It is immediate that these two mappings are inverse to each other, thus proving the statement. \square

This proposition justifies the following definition of a category BN_G of G -based Bayesian networks: objects are CDU functors $\text{Syn}_G \rightarrow \text{Stoch}$ and arrows are monoidal natural transformations between them.

4. Towards Causal Inference: the Smoking Scenario

We will motivate our approach to causal inference via a classic example, inspired by the one given in the Pearl’s book (Pearl, 2000). Imagine a dispute between a scientist and a tobacco company. The scientist claims that smoking causes cancer. As a source of evidence, the scientist cites a joint probability distribution ω over variables S for smoking and C for cancer, which disintegrates as in (5) below, with matrix $c = \begin{pmatrix} 0.9 & 0.7 \\ 0.1 & 0.3 \end{pmatrix}$. Inspecting this $c : S \rightarrow C$, the scientist notes that the probability of getting cancer for smokers (0.3) is three times as high as for non-smokers (0.1). Hence, the scientist claims that smoking has a significant causal effect on cancer.

$$\begin{array}{c} |S \quad |C \\ \boxed{\omega} \end{array} = \begin{array}{c} \begin{array}{c} |S \\ \curvearrowright \\ \bullet \\ | \\ \boxed{s} \end{array} \quad \begin{array}{c} |C \\ \boxed{c} \end{array} \end{array} \tag{5}$$


An important thing to stress here is that the scientist draws this conclusion using not only the observational data ω but also from an assumed *causal structure* which gave rise to that data,

as captured in the diagram in equation (5). That is, rather than treating diagram (5) simply as a calculation on the observational data, it can also be treated as an assumption about the actual, physical mechanism that gave rise to that data. Namely, this diagram encompasses the assumption that there is some prior propensity for people to smoke captured by $s : I \rightarrow S$, which is both observed and fed into some other process $c : S \rightarrow C$ whereby an individual's choice to smoke determines whether or not they get cancer.



The tobacco company, in turn, says that the scientist's assumptions about the provenance of this data are too strong. While they concede that *in principle* it is possible for smoking to have some influence on cancer, the scientist should allow for the possibility that there is some latent common cause (e.g. genetic conditions, stressful work environment, etc.) which leads people both to smoke and get cancer. Hence, says the tobacco company, a 'more honest' causal structure to ascribe to the data ω is (6). This structure then allows for either party to be correct. If the scientist is right, the output of $c : S \otimes H \rightarrow C$ depends mostly on its first input, that is, the causal path from smoking to cancer. If the tobacco company is right, then c depends very little on its first input, and the correlation between S and C can be explained almost entirely from the hidden common cause H .

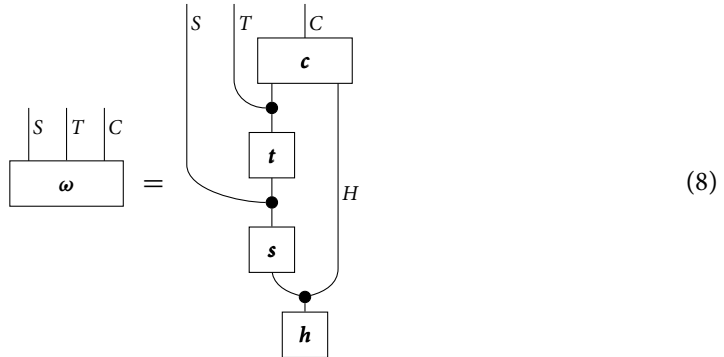
So, who is right after all? Just from the observed distribution ω , it is impossible to tell. So, the scientist proposes a clinical trial, in which patients are randomly required to smoke or not to smoke. We can model this situation by replacing s in (6) with a process that ignores its inputs and outputs the uniform state. Graphically, this looks like 'cutting' the link s between H and S :



This captures the fact that variable S is now randomised and no longer dependent on any background factors. This new distribution ω' represents the data the scientist would have obtained had they run the trial. That is, it gives the results of an *intervention* at s . If this ω' still shows a strong correlation between smoking and cancer, one can conclude that smoking indeed causes cancer even when we assume the weaker causal structure (6).

Unsurprisingly, the scientist fails to get ethical approval to run the trial, and hence has only the observational data ω to work with. Given that the scientist only knows ω (and not c and h), there is no way to compute ω' in this case. However, a key insight of statistical causal inference is that sometimes it is possible to compute interventional distributions from observational ones. Continuing the smoking example, suppose the scientist proposes the following revision to the causal structure, they posit a structure (8) that includes a third observed variable (the presence of T of tar in the lungs), which completely mediates the causal effect of smoking on cancer.

As with our simpler structure, the diagram (8) below contains some assumptions about the origins of the data ω . In particular, by omitting wires, we are asserting there is no *direct* causal link between certain variables:



The absence of an H -labelled input to t says there is no direct causal link from H to T (only mediated by S), and the absence of an S -labelled input wire into c captures that there is no direct causal link from S to C (only mediated by T). In the traditional approach to causal inference, such relationships are typically captured by a graph-theoretic property called *d-separation* (see Pearl (2000) for a standard and Jacobs and Zanasi (2017) for a categorical account) on the dag associated with the causal structure.

We can again imagine intervening at S by replacing $s : H \rightarrow S$ by $\Downarrow \circ \Uparrow$. Again, this ‘cutting’ of the diagram will result in a new interventional distribution ω' . However, unlike before, it is possible to compute this distribution from the observational distribution ω .

However, in order to do that, we first need to develop the appropriate categorical framework. In Section 5, we will model ‘cutting’ as a functor. In Section 6, we will introduce a generalisation of disintegration, which we call *comb disintegration*. These tools will enable us to compute ω' for ω , in Section 7.

5. Interventional Distributions as Diagram Surgery

The goal of this section is to define the ‘cut’ operation in (7) as an endofunctor on the category of Bayesian networks. First, we observe that such an operation exclusively concerns the string diagram part of a Bayesian network: following the functorial semantics given in Section 3, it is thus appropriate to define cut as an endofunctor on Syn_G , for a given dag G .

Definition 5. For a fixed node $A \in V_G$ in a graph G , let $\text{cut}_A : \text{Syn}_G \rightarrow \text{Syn}_G$ be the CDU functor freely obtained by the following action on the generators (V_G, Σ_G) of Syn_G :

- For each object $B \in V_G$, $\text{cut}_A(B) = B$.
- $\text{cut}_A\left(\begin{array}{c} |^A \\ \boxed{a} \\ |_{B_1 \dots B_k} \end{array}\right) = \begin{array}{c} \Downarrow^A \\ \bullet \\ |_{B_1 \dots B_k} \end{array}$ and $\text{cut}_A\left(\begin{array}{c} |^B \\ \boxed{b} \\ |_{C_1 \dots C_l} \end{array}\right) = \begin{array}{c} |^B \\ \boxed{b} \\ |_{C_1 \dots C_l} \end{array}$ for any other $\begin{array}{c} |^B \\ \boxed{b} \\ |_{C_1 \dots C_l} \end{array} \in \Sigma_G$.

Intuitively, cut_A applied to a string diagram f of Syn_G removes from f each occurrence of a box with output wire of type A .

Proposition 4 allows us to ‘transport’ the cutting operation over to Bayesian networks. Given any Bayesian network based on G , let $F : \text{Syn}_G \rightarrow \text{Stoch}$ be the corresponding CDU functor given by Proposition 4. Then, we can define its A -cutting as the Bayesian network identified by the CDU functor $F \circ \text{cut}_A$. This yields an (idempotent) endofunctor $\text{Cut}_A : \text{BN}_G \rightarrow \text{BN}_G$.

6. The Comb Factorisation

Thanks to the development of Section 5, we can understand the transition from left to right in (7) as the application of the functor Cut_S applied to the ‘Smoking’ node S . The next step is being able to actually compute the individual Stoch-morphisms appearing in (8), to give an answer to the causality question.

In order to do that, we want to work in a setting where $t: S \rightarrow T$ can be isolated and ‘extracted’ from (8). What is left behind is a stochastic matrix with a ‘hole’ where t has been extracted. To define ‘morphisms with holes’, it is convenient to pass from SMCs to compact closed categories (see e.g. Selinger (2011)). Stoch is not itself compact closed, but it embeds into $\text{Mat}(\mathbb{R}^+)$, whose morphisms are *all* matrices over non-negative numbers. $\text{Mat}(\mathbb{R}^+)$ has a (self-dual) compact closed structure; that means, for any set A there is a ‘cap’ $\cap: A \otimes A \rightarrow I$ and a ‘cup’ $\cup: I \rightarrow A \otimes A$, which satisfy the ‘yanking’ equations on the right. As matrices, caps and cups are defined by $\cap_{ij} = \cup^{ij} = \delta_i^j$. Intuitively, they amount to ‘bent’ identity wires. Another aspect of $\text{Mat}(\mathbb{R}^+)$ that is useful to recall is the following handy characterisation of the subcategory Stoch.

Lemma 6. *A morphism $f: A \rightarrow B$ in $\text{Mat}(\mathbb{R}^+)$ is a stochastic matrix (thus a morphism of Stoch) if and only if (3) holds, that is, if $\uparrow \circ f = \uparrow$.*

A suitable notion of ‘stochastic map with a hole’ is provided by a *comb*. These structures originate in the study of certain kinds of quantum channels (Chiribella *et al.*, 2008).

Definition 7. *A 2-comb in Stoch is a morphism $f: A_1 \otimes A_2 \rightarrow B_1 \otimes B_2$ satisfying, for some other morphism $f': A_1 \rightarrow B_1$,*

For this notion of 2-comb, it is important to consider the map f with an explicit description of the product \otimes on its domain and codomain.

This definition extends inductively to *n-combs*, where we require that discarding the rightmost output yields $f' \otimes \uparrow$, for some $(n - 1)$ -comb f' . However, for our purposes, restricting to 2-combs will suffice.

Remark 8. Note that Definition 7 depends on the specific decomposition of the domain and codomain into a monoidal product of subsystems. For example, any stochastic matrix is trivially a 2-comb with respect to the decomposition $f: A \otimes I \rightarrow B \otimes I$, where I is the monoidal unit. Hence, the data associated with a 2-comb is actually a pair of input/output types and a morphism: $((A_1, B_1), (A_2, B_2), f: A_1 \otimes A_2 \rightarrow B_1 \otimes B_2)$. We will suppress this extra data when it is clear from context.

The intuition behind condition (9) is that the contribution from input A_2 is only visible via output B_2 . Thus, if we discard B_2 we may as well discard A_2 . In other words, the input/output pair A_2, B_2 happen ‘after’ the pair A_1, B_1 . Hence, it is typical to depict 2-combs in the shape of a (hair) comb, with 2 ‘teeth’, as in (10) below:

(10) (11)

While combs themselves live in $\text{Stoch}, \text{Mat}(\mathbb{R}^+)$ accommodates a second-order reading of the transition \rightsquigarrow in (10), we can treat f as a map which expects as input a map $g: B_1 \rightarrow A_2$ and produces as output a map of type $A_1 \rightarrow B_2$. Plugging $g: B_1 \rightarrow A_2$ into the 2-comb can be formally defined in $\text{Mat}(\mathbb{R}^+)$ by composing f and g in the usual way, then feeding the output of g into the second input of f , using caps and cups, as in (11).

Importantly, for generic f and g of Stoch , there is no guarantee that forming the composite (11) in $\text{Mat}(\mathbb{R}^+)$ yields a valid Stoch -morphism, that is, a morphism satisfying the finality equation (3). However, if f is a 2-comb and g is a Stoch -morphism, equation (9) enables a discarding map plugged into the output B_2 in (11) to ‘fall through’ the right side of f , which guarantees that the composed map satisfies the finality equation for discarding:


(9) (3) (3)

Remark 9. An alternative formulation of the comb composition in (11) can be made without leaving the subcategory of stochastic matrices. First, in Stoch , any stochastic matrix satisfying (9) can be *semi-localised*, that is, factored into two parts f_1, f_2 as follows:

Then, one can show that the composition:

(12)


does not depend on the particular choice of f_1, f_2 . The latter can be seen by concrete calculation, or by noting that (12), seen as a diagram in $\text{Mat}(\mathbb{R}^+)$, can be deformed into (11) by re-introducing the feedback loop on A_2 :



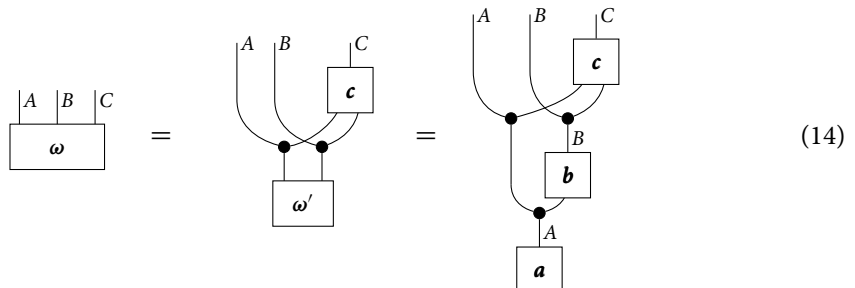
See for example Kissinger and Uijlen (2017) for more details. As it is possible to freely pass between f and f_1, f_2 , one could equivalently define a comb as an equivalence class of pairs (f_1, f_2) which compose to give f , much like those considered in the construction of completely positive maps in Coecke and Heunen (2016).

With the concept of 2-combs in hand, we can state our factorisation result.


Theorem 10. *For any state $\omega : I \rightarrow A \otimes B \otimes C$ of Stoch with full support, there exists a unique 2-comb $f : B \rightarrow A \otimes C$ and stochastic matrix $g : A \rightarrow B$ such that, in $\text{Mat}(\mathbb{R}^+)$:*



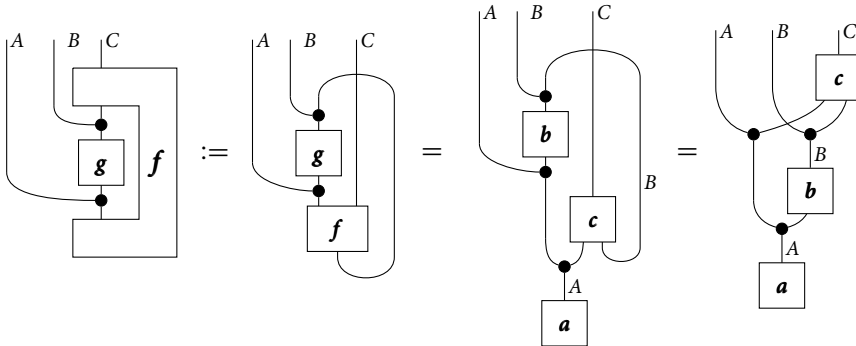
Proof. The construction of f and g mimics the construction of c -factors in Tian and Pearl (2002), using string diagrams and diagrammatic disintegration. Starting with a full-support $\omega : I \rightarrow A \otimes B \otimes C$, we apply Theorem 3 twice. First, we can disintegrate ω as $(\omega' : I \rightarrow A \otimes B, c : A \otimes B \rightarrow C)$ then further disintegrate ω' into $(a : I \rightarrow A, b : A \rightarrow B)$:



Now, we let



Then (13) holds by construction of a, b, c :



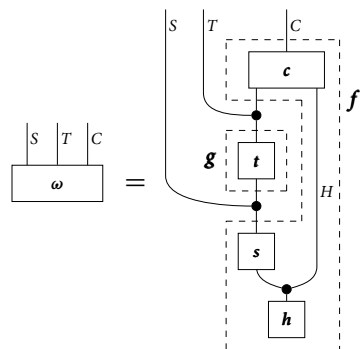
Note the last step above is just diagram deformation and the comonoid laws. The rightmost diagram above is equal to ω by (14).

For uniqueness, suppose (13) holds for some other f', g' . Then by uniqueness of disintegration, it follows that $g' = g = b$. To show that $f = f'$, we expand (13) explicitly in terms of matrices. This equation is equivalent to $\omega^{ijk} = f_j^{ik} g_i^j = (f')_j^{ik} g_i^j$. Note that if g had any zero elements, ω would not have full support, hence $g_i^j \neq 0$ and therefore $f_j^{ik} = (f')_j^{ik}$ for all i, j, k . \square


Note that Theorem 10 generalises the normal disintegration property given in Theorem 3. The latter is recovered by taking $A := I$ (or $C := I$) above.

7. Returning to the Smoking Scenario

We now return to the smoking scenario of Section 4. There, we concluded by claiming that the introduction of an intermediate variable T to the observational distribution $\omega : I \rightarrow S \otimes T \otimes C$ would enable us to calculate the interventional distribution. That is, we can calculate $\omega' = F(\text{cut}_x(\omega))$ from $\omega := F(\omega)$. Thanks to Theorem 10, we are now able to perform that calculation. We first observe that our assumed causal structure for ω fits the form of Theorem 10, where g is t and f is a 2-comb containing everything else, as in the diagram on the side.



Hence, f and g are computable from ω . If we plug them back together as in (13), we will get ω back. However, if we insert a 'cut' between f and g :




we obtain $\omega' = F(\text{cut}_X(\omega))$.

Let us now consider a concrete example. We fix interpretations for the sets $S, T,$ and C as booleans: $S = T = C = \{0, 1\}$ and let $\omega : I \rightarrow S \otimes T \otimes C$ be the stochastic matrix:

$$\omega := \begin{pmatrix} 0.5 \\ 0.1 \\ 0.01 \\ 0.02 \\ 0.1 \\ 0.05 \\ 0.02 \\ 0.2 \end{pmatrix} \begin{matrix} \leftarrow P(S = 0, T = 0, C = 0) \\ \leftarrow P(S = 0, T = 0, C = 1) \\ \leftarrow P(S = 0, T = 1, C = 0) \\ \leftarrow P(S = 0, T = 1, C = 1) \\ \leftarrow P(S = 1, T = 0, C = 0) \\ \leftarrow P(S = 1, T = 0, C = 1) \\ \leftarrow P(S = 1, T = 1, C = 0) \\ \leftarrow P(S = 1, T = 1, C = 1) \end{matrix}$$


Now, disintegrating ω :



The bottom-left element of c is $P(C = 1|S = 0)$, whereas the bottom-right is $P(C = 1|S = 1)$, so this suggests that patients are ≈ 3.5 times as likely to get cancer if they smoke (68% vs. 19%). However, comb-disintegrating ω using Theorem 10 gives $g : S \rightarrow T$ and a comb $f : T \rightarrow S \otimes C$ with the following stochastic matrices:

$$f \approx \begin{pmatrix} 0.53 & 0.21 \\ 0.11 & 0.42 \\ 0.25 & 0.03 \\ 0.12 & 0.34 \end{pmatrix} \quad g \approx \begin{pmatrix} 0.95 & 0.41 \\ 0.05 & 0.59 \end{pmatrix}$$

Recomposing these with a ‘cut’ in between, as in the left-hand side of (15), gives the interventional distribution $\omega' \approx (0.38, 0.11, 0.01, 0.02, 0.16, 0.05, 0.07, 0.22)$. Disintegrating:



From the interventional distribution, we conclude that, in a (hypothetical) clinical trial, patients are about twice as likely to get cancer if they smoke (54% vs. 25%). So, since $54 < 68$, there was

some confounding influence between S and C in our observational data, but after removing it via comb disintegration, we see there is still a significant causal link between smoking and cancer.

Note this conclusion depends totally on the particular observational data that we picked. For a different interpretation of ω in Stoch, one might conclude that there is *no* causal connection, or even that smoking *decreases* the chance of getting cancer. Interestingly, all three cases can arise even when a naïve analysis of the data shows a strong direct correlation between S and C . To see and/or experiment with these cases, we have provided the Python code² used to perform these calculations. See also Nielsen (2018) for a paedagogical overview of this example (using traditional Bayesian network language) with some sample calculations.

8. Generic Single Interventions

While we applied the comb decomposition to a particular example, this technique applies essentially unmodified to many examples where we intervene at a single variable (called X below) within an arbitrary causal structure.

Theorem 11. *Let G be a dag with a fixed node X that has corresponding generator $x: Y_1 \otimes \dots \otimes Y_n \rightarrow X$ in Syn_G . Then, let ω be a morphism in Syn_G of the following form:*

(16)

for some morphisms f_1, f_2 and g in Syn_G not containing x as a subdiagram. Then the interventional distribution $\omega' := F(\text{cut}_x(\omega))$ is computable from any observational distribution $\omega = F(\omega)$ with full support.

Proof. The proof is very close to the example in the previous section. Interpreting ω into Stoch, we get a diagram of stochastic maps, which we can comb-disintegrate, then recompose with $\nabla \circ \uparrow$ to produce the interventional distribution:

(3)

The right-hand side above is then $F(\text{cut}_x(\omega))$. □

This is general enough to cover several well-known sufficient conditions from the causality literature, including single-variable versions of the so-called *front-door* and *back-door* criteria,

as well as the sufficient condition based on confounding paths given by Tian and Pearl (2002). As the latter subsumes the other two, we will say a few words about the relationship between the Pearl/Tian condition about confounding paths and Theorem 11. In Tian and Pearl (2002), the authors focus on *semi-Markovian* models, where the only latent variables have exactly two observed children and no parents. Suppose we write $A \leftrightarrow B$ if two observed variables are connected by a latent common cause, then one can characterise *confounding paths* as the transitive closure of \leftrightarrow . They go on to show that the interventional distribution corresponding to cutting X is computable whenever there are no confounding paths connecting X to one of its children.

We can compare this to the form of expression ω in equation (16). First, note this factorisation implies that all boxes which take X as an input must occur as sub-diagrams of g . Hence, any ‘confounding path’ connecting X to its children would yield at least one (un-copied) wire from f_1 to g , hence it cannot be factored as (16). Conversely, if there are no confounding paths from X to its children, then we can place the boxes involved in any other confounding path either entirely inside of g or entirely outside of g and obtain factorisation (16). Hence, restricting to semi-Markovian models, the non-confounding-path condition from Tian and Pearl (2002) is equivalent to ours. However, Theorem 11 is slightly more general: its formulation does not rely on the causal structure ω being semi-Markovian.

9. Counterfactuals

While interventional distributions are a powerful tool for extracting causal information from a probabilistic model, there are certain cases where genuine causal influences remain hidden. In those cases, we can extend string diagram surgery techniques to do *counterfactual reasoning*. This enables us to reason about alternatives for events that have already occurred. For example, we can consider the likelihood of statements like:

‘Had Sally been able to play, the team would have won the game’.

Implicit in this statement is the assumption that Sally was, in fact, *not* able to play. We can make this assumption explicit as follows:

‘Given Sally was not able to play, the team would have won, had she been able to play’.

We therefore end up in the seemingly paradoxical situation of needing to condition on a real-world event (Sally not being able to play) and considering the outcome of a contradictory event (Sally being able to play).

In order to evaluate the likelihood of such a statement, we need to be able to compare two worlds: the real world, containing events that have already happened, with a hypothetical world which is the same in every way except for a single intervention, namely making Sally able to play.

To see how we could compute probabilities for counterfactual statements using string diagram surgery, we will work out a concrete example of Balke and Pearl (1994). We will do this by giving a diagrammatic version of the ‘twin model’ technique used there.

Consider three people, Ann, Bob and Carl, who may or may not go to a party. The likelihood that Ann goes to the party is 60%. Bob is very likely (90%) to go to the party if Ann is there, but he will almost surely not go otherwise (97%). For Carl, it is the opposite: he is very likely to go (95%) if Ann does not go, and otherwise he will probably not go (90%). If Bob and Carl both go to the party, they are very likely to have a scuffle (95%). If only one of them is going, it is certain there will be no scuffle (100%), whereas if both are not going there is still a small chance (5%) of them getting in a scuffle somewhere else.

We can express this situation with the following string diagram, where the wires labelled A , B and C represent whether Ann, Bob or Carl go to the party, and S whether there is a scuffle.³




The probabilities given before are then modelled by the following stochastic matrices:

$$a := \begin{pmatrix} 0.40 \\ 0.60 \end{pmatrix} \quad f := \begin{pmatrix} 0.90 & 0.03 \\ 0.10 & 0.97 \end{pmatrix} \quad g := \begin{pmatrix} 0.10 & 0.95 \\ 0.90 & 0.05 \end{pmatrix} \quad h := \begin{pmatrix} 0.95 & 1 & 1 & 0.05 \\ 0.05 & 0 & 0 & 0.95 \end{pmatrix}$$

An example of a counterfactual statement is the following:

‘Given Bob did not go to the party, there would have been a scuffle, had he gone’. (18)

Since we know all the stochastic matrices already, we can compute the (naïve) interventional distribution to predict how likely a scuffle is if we intervene and *make* Bob go to the party:



Disintegrating on *B* yields

$$\begin{array}{c} S \\ | \\ \boxed{h} \\ | \\ \boxed{g} \\ | \\ \boxed{a} \end{array} \Big|_B \approx \begin{pmatrix} 0.97 & 0.58 \\ 0.03 & 0.42 \end{pmatrix}$$

The bottom-right corner of this matrix corresponds to $P(S = 1|B = 1)$, and as we can see, the probability of a scuffle, given the intervention setting $B = 1$, is 42%.

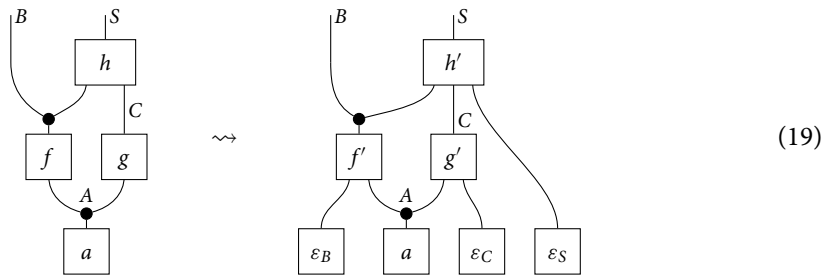
However, we have ignored an important piece of information: Bob actually did *not* go to the party. Hence, to assign a more accurate probability to statement (18), we should take this into account. As we said before, to do this requires us to consider two copies of the world, one which actually happened (where Bob did not go to the party) and one where we introduce a hypothetical intervention (making Bob go to the party). Crucially, everything else should be kept the same between the two copies of the world.

So, what does it mean to ‘keep everything else the same’? For example, if we look at the stochastic map f relating Ann’s attendance to Bob’s, the randomness in f can be interpreted as the presence of other, possibly unknown variables affecting Bob’s attendance (e.g. the weather or Bob not feeling well). We can isolate those variables by factoring f into a deterministic part f' and an *exogenous state* ϵ_B . For example, one such factorisation is

$$\begin{array}{c} \boxed{f} \end{array} = \begin{array}{c} \boxed{f'} \\ \boxed{\varepsilon_B} \end{array} \quad \text{where} \quad f' := \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \end{pmatrix} \quad \varepsilon_B := \begin{pmatrix} 0.027 \\ 0.873 \\ 0.003 \\ 0.097 \end{pmatrix}$$

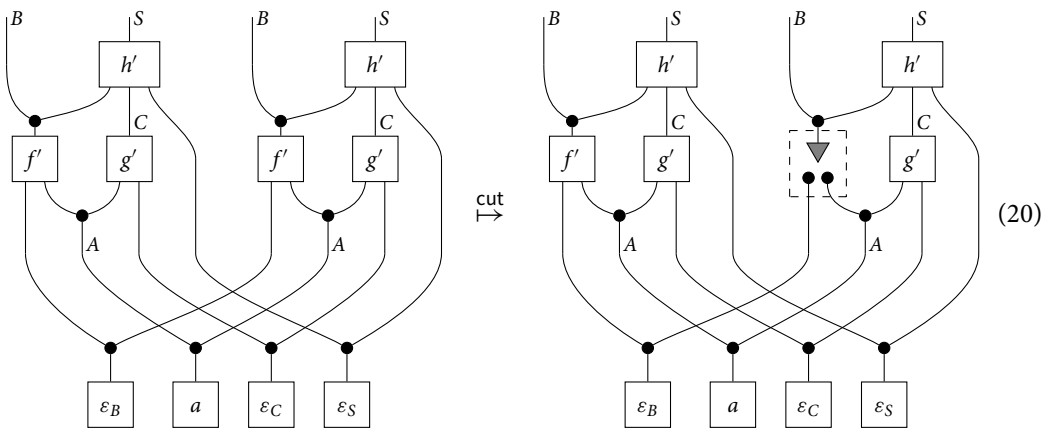
The deterministic part is a function, which can be expressed as a stochastic matrix whose entries are all only 0 or 1. It is well known that such a factorisation for stochastic matrices is always possible. For example, it is suggested by Fritz (2020) as a candidate axiom for ‘well-behaved’ probabilistic categories (of which Stoch, which Fritz calls FinStoch, is an example).

Decomposing f , g and h into function and exogenous variables, we are able to *exogonize* the diagram, that is, obtain a new diagram where all of the randomness occurs in boxes with no inputs:

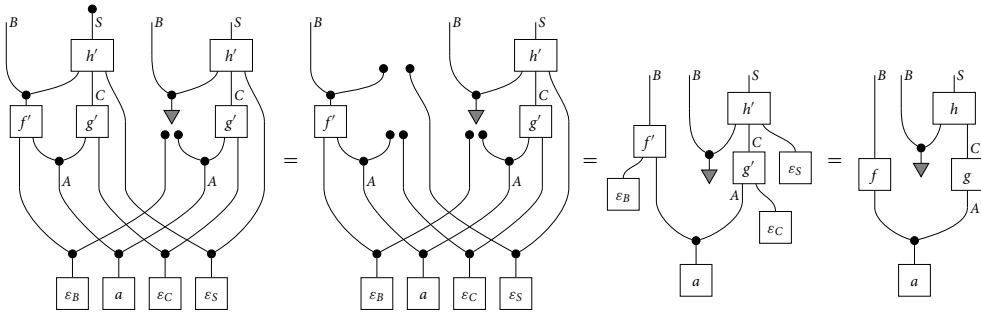


Note that, since a does not have any inputs, we can already treat it as an exogenous state.


In order to answer the counterfactual query, we transform our diagram into a ‘twin’ diagram, which contains two copies of all the variables: the first copy represents the ‘factual world’ in which Bob did not go to the party, whereas the second copy represents a counterfactual world, in which Bob went to the party. In order to keep everything else but Bob’s attendance the same, the two copies share *the same* exogenous states. We then intervene on the counterfactual copy of B by cutting out the rightmost f' :



To evaluate the probability of the counterfactual statement (18), we first marginalise out the left copy of S and do some simplification:



The result is a state $\omega : I \rightarrow B \otimes B \otimes S$. Disintegrating over the two copies of B , we obtain a stochastic map from ‘real’ B and ‘counterfactual’ B to S :




We can condition on the ‘real’ $B = 0$ by plugging the state $e_0 := (1, 0)$ (meaning ‘Bob did not go to the party’) into the first wire of k , to obtain

$$\begin{array}{c} S \\ | \\ \boxed{k} \\ | \\ \boxed{e_0} \quad B \end{array} \approx \begin{pmatrix} 0.99 & 0.21 \\ 0.01 & 0.79 \end{pmatrix} \tag{21}$$

Again looking at the bottom-right corner, we see the hypothetical probability of a scuffle has gone up to 79%, matching the conclusion from Balke and Pearl (1994).

This is because knowing that Bob actually did not attend the party enables us to draw certain conclusions about what happened at the actual party. Namely, if Bob was not there, it is less likely that Ann was there, which in turn means it was more likely that Carl was. Hence, if we want to think about what would have happened at that *same* party, had Bob gone, he most likely would have seen Carl, and there most likely would have been some fisticuffs.

Crucially, the exogenous states are *shared* between the two copies of the diagram, rather than split into two independent copies. This is what enables us to imagine that the real and counterfactual halves are talking about the exact same event, rather than a separate event that happens to be modelled by the same stochastic processes. We can see in that case that the shared a box is what enables information to flow between the two worlds:



That is, when we condition on the counterfactual copy of B in equation (21), this has an effect on the value of S for all values of the ‘real’ copy of B .

One other thing to note from this example is that it does not actually depend on the decompositions of stochastic maps into functions and exogenous variables. This is not the case in general, and furthermore such decompositions are not unique. Hence counterfactual reasoning is only well defined when either (i) the decomposition into functions and exogenous states is given in advance, or (ii) the resulting counterfactual query does not depend on the choice of decomposition.

In Balke and Pearl (1994), the authors use this fact to state that causal Bayesian networks alone are not adequate for counterfactual reasoning and emphasise instead the utility of structural equational models for this task. However, in certain situations, like the party example, Bayesian networks do suffice.

10. Conclusion and Future Work

This paper takes a fresh, systematic look at the problem of causal identifiability. By clearly distinguishing syntax (string diagram surgery and identification of comb shapes) and semantics (comb-disintegration of joint states), we obtain a clear methodology for computing interventional distributions, and hence causal effects, from observational data. Furthermore, we show that diagram surgery can be used to model counterfactual queries.

A natural next step is moving beyond single-variable interventions to the general case, that is, situations where we allow interventions on multiple variables which may have some arbitrary causal relationships connecting them. This would mean extending the comb factorisation Theorem 10 from a 2-comb and a channel to arbitrary n -combs. This seems to be straightforward, via an inductive extension of the proof in Section 8. A more substantial direction of future work will be the strengthening of Theorem 11 from sufficient conditions for causal identifiability to a full characterisation. Indeed, the related condition based on confounding paths from Tian and Pearl (2002) is a necessary and sufficient condition for computing the interventional distribution on a single variable. Hence, it will be interesting to formalise this necessity proof (and more general versions, e.g. Huang and Valtorta (2008)) within our framework and investigate, for example, the extent to which it holds beyond the semi-Markovian case.

Throughout this work, we have relied crucially on the fact that observational data have full support. This assumption seems to often be made without comment in the literature on causal identifiability. For example, Tian and Pearl make this assumption implicitly in Tian and Pearl (2002) by dividing by conditional probabilities at will. However, it is a rather strong assumption, and in particular rules out deterministic relationships between variables. There are some tricks one can do to cope with this problem (see e.g. the discussion of determinism in the context of causal discovery in Mooij and Magliacane (2020)), so it would be interesting to see if they can be adapted to our setting.

While we focus exclusively on the case of taking models in *Stoch* in this paper, the techniques we gave are posed at an abstract level in terms of composition and factorisation. Hence, we are optimistic about their prospects to generalise to other probabilistic (e.g. infinite discrete and continuous variables) and quantum settings. In the latter case, this could provide insights into the emerging field of *quantum causal structures* (Costa and Shrapnel, 2016; Henson *et al.*, 2014; Leifer and Spekkens, 2013; Pienaar and Brukner, 2015; Ried *et al.*, 2015), which attempts in part to replay some of the results coming from statistical causal reasoning, but where quantum processes play a role analogous to stochastic ones. A key difficulty in applying our framework to a category of quantum processes, rather than *Stoch*, is the unavailability of canonical (basis-independent) ‘copy’ morphisms due to the quantum no-cloning theorem Wootters and Zurek (1982). However, a recent proposal for the formulation of ‘quantum common causes’ Allen *et al.* (2017) suggests a

(partially defined) analogue to the role played by ‘copy’ in our formulation constructed via multiplication of certain commuting Choi matrices. Hence, it may yet be possible to import results from classical causal reasoning into the quantum case just by changing the category of models.

Acknowledgements. Fabio Zanasi is supported by EPSRC grant EP/V002376/1.

Notes

1 CDU-categories are closely related to *gs-monoidal* categories, structured adopted in the categorical description of resource-sensitive algebraic theories (Corradini and Gadducci, 1999). The difference is that CDU categories also include uniform states

▼ and the corresponding equation. As we will see, this extra structure is needed in order to account for causal intervention. A CDU category is also very much like a Markov category defined in Fritz (2020), but Markov categories do not have these uniform states either; additionally, in Markov categories the discard equation (3) holds for all maps.

2 <https://gist.github.com/akissinger/aeec1751792a208253bda491ead587b6>

3 Note in (17), we treat A and C as latent variables and choose to observe B and S . This is because B and S are involved in the counterfactual query (18) under consideration.

References

- Allen, J.-M. A., Barrett, J., Horsman, D. C., Lee, C. M. and Spekkens, R. W. (2017). Quantum common causes and quantum causal models. *Physical Review X* 7 031021.
- Balke, A. and Pearl, J. (1994). Probabilistic evaluation of counterfactual queries. In: Hayes-Roth, B. and Korf, R. (eds.) *Proceedings of the 12th National Conference on Artificial Intelligence*. AAAI Press/The MIT Press, 230–237.
- Bonchi, F., Sobociński, P. and Zanasi, F. (2018). Deconstructing Lawvere with distributive laws. *Journal of Logical Algebra Methods Programming* 95 128–146.
- Cabrera, B., Heindel, T., Heckel, R. and König, B. (2018). Updating probabilistic knowledge on condition/event nets using Bayesian networks. In: *29th International Conference on Concurrency Theory, CONCUR 2018, September 4–7, 2018, Beijing, China*, 27:1–27:17.
- Chiribella, G., D’Ariano, G. M. and Perinotti, P. (2008). Quantum circuit architecture. *Physical Review Letters* 101 060401.
- Cho, K. and Jacobs, B. (2019). Disintegration and Bayesian inversion via string diagrams. *Mathematical Structures in Computer Science* 29 (7) 938–971.
- Clerc, F., Danos, V., Dahlqvist, F. and Garnier, I. (2017). Pointless learning. In: Esparza, J. and Murawski, A. (eds.) *Foundations of Software Science and Computation Structures*, vol. 10203. Lecture Notes Computer Science. Berlin: Springer, 355–369.
- Coecke, B. and Heunen, C. (2016). Pictures of complete positivity in arbitrary dimension. *Information and Computation* 250 50–58.
- Coecke, B. and Spekkens, R. W. (2012). Picturing classical and quantum Bayesian inference. *Synthese* 186 (3) 651–696.
- Corradini, A. and Gadducci, F. (1999). An algebraic presentation of term graphs, via GS-monoidal categories. *Applied Categorical Structures* 7 (4) 299–331.
- Costa, F. and Shrapnel, S. (2016). Quantum causal modelling. *New Journal of Physics* 18 (6) 063032.
- Fong, B. (2012). Causal theories: A categorical perspective on Bayesian networks. Master’s thesis, Univ. of Oxford. see arxiv.org/abs/1301.6201.
- Fritz, T. (2020). A synthetic approach to Markov kernels, conditional independence, and theorems on sufficient statistics. *Advances in Mathematics* 370 107239.
- Gutoski, G. and Watrous, J. (2007). Toward a general theory of quantum games. In: *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*. ACM, 565–574.
- Henson, J., Lal, R. and Pusey, M. F. (2014). Theory-independent limits on correlations from generalized Bayesian networks. *New Journal of Physics* 16 (11), 113043.
- Huang, Y. and Valorta, M. (2008). On the completeness of an identifiability algorithm for semi-Markovian models. *Annals of Mathematics and Artificial Intelligence* 54 (4) 363–408.
- Huang, Y. and Valorta, M. (2012). Pearl’s calculus of intervention is complete. *CoRR*, abs/1206.6831.
- Jacobs, B. and Zanasi, F. (2021). The logical essentials of Bayesian reasoning. In: Barthe, G., Katoen, J.-P. and Silva, A. (eds.) *Foundations of Probabilistic Programming*. Cambridge Univ. Press, 295–331.
- Jacobs, B., Kissinger, A. and Zanasi, F. (2019). Causal inference by string diagram surgery. In: *Foundations of Software Science and Computation Structures - 22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings*, 313–329.

- Jacobs, B. and Zanasi, F. (2016). A predicate/state transformer semantics for Bayesian learning. *Electr. Notes Theor. Comput. Sci.* **325**, 185–200.
- Jacobs, B. and Zanasi, F. (2017). A formal semantics of influence in Bayesian reasoning. In: Larsen, K., Bodlaender, H. and Raskin, J.-F. (eds.) *Mathematics Foundation of Computer Science*, vol. 83. *LIPICs*. Schloss Dagstuhl, 21:1–21:14.
- Kissinger, A. and Uijlen, S. (2017). A categorical semantics for causal structure. In: *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20–23, 2017*, 1–12.
- Lawvere, F. W. (1963). Functorial semantics of algebraic theories. *Proceedings of the National Academy of Sciences of the United States of America* **50** (5) 869.
- Lawvere, F. W. (1969). Ordinal sums and equational doctrines. In: Eckmann, B. (ed.) *Seminar on Triples and Categorical Homology Theory*, vol. 80. *Lecture Notes in Mathematics*. Springer-Verlag, 141–155.
- Leifer, M. S. and Spekkens, R. W. (2013). Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Physical Review A* **88** 052130.
- Mooij, J. M., Magliacane, S. and Claassen, T. (2020). Joint causal inference from multiple contexts. *Journal of Machine Learning Research* **21** (99) 1–108.
- Nielsen, M. (2018). If correlation doesn't imply causation, then what does? Available at <http://www.michaelnielsen.org/ddi/if-correlation-doesnt-imply-causation-then-what-does>, accessed: 2018-11-15.
- Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge University Press.
- Pearl, J. and Verma, T. S. (1991). *A Theory of Inferred Causation*. San Mateo, CA: Morgan Kaufmann.
- Perinotti, P. (2016). Causal structures and the classification of higher order quantum computations. Available at https://secure-web.cisco.com/1jxvUSMjzVpPjt982gd5jZAxp3vC8KIP36joNcSngFcAbsMxhZ9jrv7BeH_DFrAHXTqtSLstLvifcfBXseDkJFuodtvmZVOD_BFH8itpo9pVoYdlcCe4obhzyPK5Um_mS0lk-o7xokvRA_Zbd0Z3LZteW-Cvzr0VIUHMEx1u5tVm3hcMVNKK5hwG0Zk7w5_iMdbl9hOgPAIdNQSqWvrqoGtTyWwFhKEfn_F00cV3bnk-ACNMbMKLjLHbiSkyeJ4TF3u7tYTIISFP9hIVWnd7e3ImpQi2-R6Ew0es66Eza-23H5napxga3Lwd_PD0VeW0aXOHck-wpFVPxWIXWBHI165qCgILZ_ikOcMckX4ieDNz17BjbqPexDDLiCaQkt9hiZD-DYCjWRHCdiOweBfw/https%3A%2F%2Flink.springer.com%2Fchapter%2F10.1007%2F978-3-319-68655-4_7
- Pienaar, J. and Brukner, Č. (2015). A graph-separation theorem for quantum causal models. *New Journal of Physics* **17** (7) 073020.
- Ried, K., Agnew, M., Vermeyden, L., Janzing, D., Spekkens, R. W. and Resch, K. J. (2015). A quantum advantage for inferring causal structure. *Nature Physics* **11** 1745–2473.
- Selinger, P. (2011). A survey of graphical languages for monoidal categories. In: Coecke, B. (ed.) *New Structures for Physics*, vol. 813. *Lecture Notes in Physics*. Springer, 289–355.
- Shpitser, I. and Pearl, J. (2006). Identification of joint interventional distributions in recursive semi-Markovian causal models. In: *Proceedings of the 21th National Conference on Artificial Intelligence*, Menlo Park, CA; Cambridge, MA; London: AAAI Press; MIT Press; 1999, 1219–1226.
- Tian, J. and Pearl, J. (2002). A general identification condition for causal effects. In: *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence, July 28–August 1, 2002, Edmonton, Alberta, Canada*, 567–573.
- Wootters, W. K. and Zurek, W. H. (1982). A single quantum cannot be cloned. *Nature* **299** (5886) 802–803.

Cite this article: Jacobs B, Kissinger A and Zanasi F (2021). Causal inference via string diagram surgery: A diagrammatic approach to interventions and counterfactuals. *Mathematical Structures in Computer Science* **31**, 553–574. <https://doi.org/10.1017/S096012952100027X>