
A LAYERED ARCHITECTURE FOR UNIVERSAL CAUSALITY *

A PREPRINT

Sridhar Mahadevan

Adobe Research and University of Massachusetts, Amherst
smahadev@adobe.com, mahadeva@umass.edu

December 18, 2022

ABSTRACT

We propose a layered hierarchical architecture called UCLA (Universal Causality Layered Architecture), which combines multiple levels of categorical abstraction for causal inference. At the top-most level, causal interventions are modeled combinatorially using a simplicial category of ordinal numbers Δ , whose objects are the ordered natural numbers $[n] = \{0, \dots, n\}$, and whose morphisms are order-preserving injections and surjections. At the second layer, causal models are defined by a graph-type category consisting of a collection of objects, such as the entities in a relational database, and morphisms between objects can be viewed as attributes relating entities. The non-random “surgical” operations on causal structures, such as edge deletion, are captured using degeneracy and face operators from the simplicial layer above. The third categorical abstract layer corresponds to the data layer in causal inference, where each causal object is mapped into a set of instances, modeled using the category of sets and morphisms are functions between sets. The fourth homotopy layer comprises of additional structure imposed on the instance layer above, such as a topological space, a measurable space or a probability space, which enables evaluating causal models on datasets. Functors map between every pair of layers in UCLA. Each functor between layers is characterized by a universal arrow, which defines an isomorphism between every pair of categorical layers. These universal arrows define universal elements and representations through the Yoneda Lemma, and in turn lead to a new category of elements based on a construction introduced by Grothendieck. Causal inference between each pair of layers is defined as a lifting problem, a commutative diagram whose objects are categories, and whose morphisms are functors that are characterized as different types of fibrations. An interesting result is that monoidal categories can be shown to be special cases of simplicial objects under a particular type of fibration called a Grothendieck opfibration. We illustrate the UCLA architecture using a range of examples, including integer-valued multisets that represent a non-graphical framework for conditional independence, and causal models based on graphs and string diagrams using symmetric monoidal categories. We define causal effect in terms of the homotopy colimit of the nerve of the category of elements.

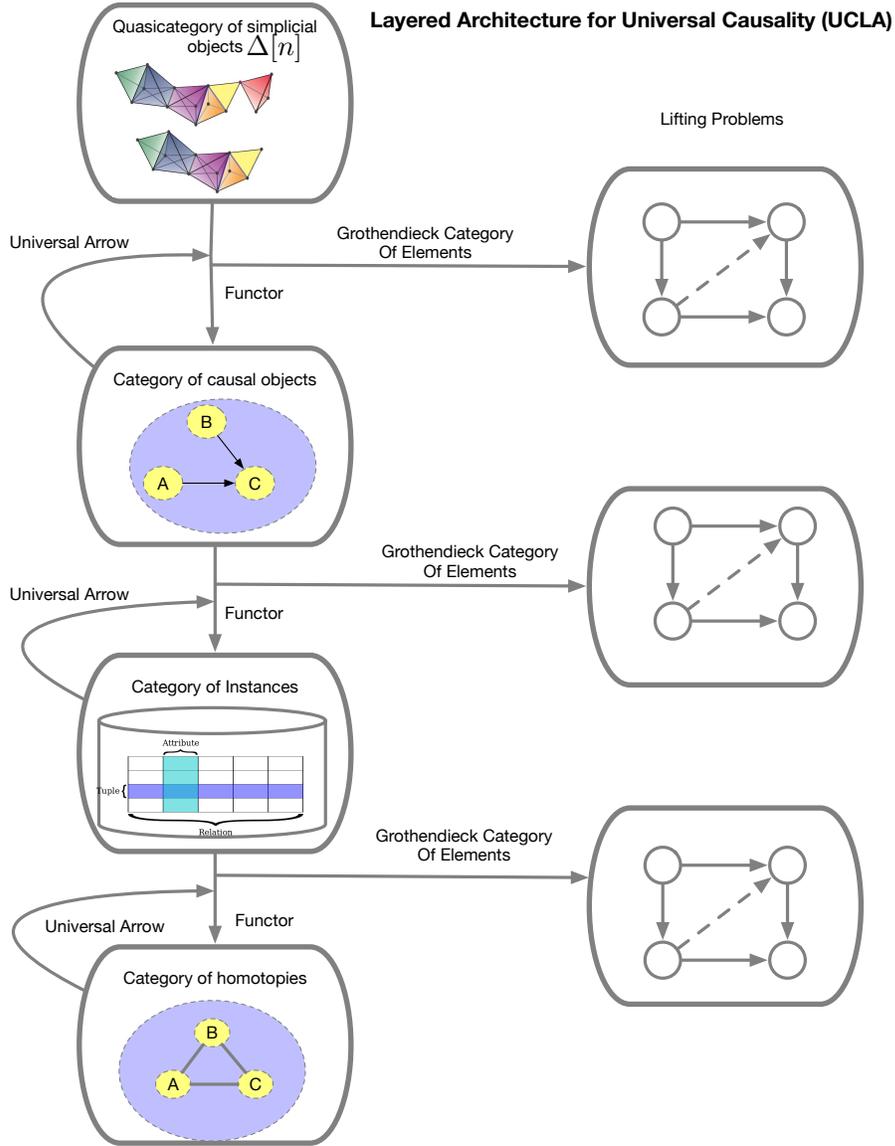
Keywords AI · Category Theory · Causal Inference · Simplicial Objects · Machine Learning · Statistics

1 Introduction

Universal Causality Fong (2012); Jacobs et al. (2018); Mahadevan (2021a, 2022a,b) is a framework for modeling causal inference using category theory, and includes a growing number of approaches, ranging from categorified representations of causal DAG models Fong (2012); Jacobs et al. (2018) to higher-order categorical representations Mahadevan (2022b), categorical representations of conditional independence Mahadevan (2022a) to universal decision models combining causal inference with other types of decision making Mahadevan (2021a). In this paper, we propose a layered architecture that defines the framework called UCLA (Universal Causality Layered Architecture). This architecture is illustrated in Figure 1. Many variants are possible, as we will discuss in the paper. As functors compose with each other, it is also possible to consider “collapsed” version of the UCLA hierarchy.

*Draft under revision. Comments welcome.

Figure 1: UCLA is a layered architecture that defines Universal Causality.



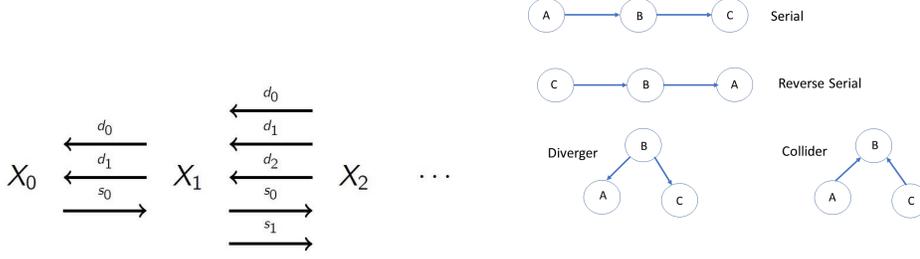
The UCLA architecture is built on the theoretical foundation of category theory MacLane (1971); MacLane and Ieke Moerdijk (1994); Riehl (2017); Joyal et al. (1996), a broad unified framework that has provided a unified mathematical language for over half a century. Previous work has investigated using category theory to model causal inference. Fong (2012) and Jacobs et al. (2018) both primarily focusing on reformulating Pearl’s DAG model using symmetric monoidal categories to represent causal models Jacobs et al. (2018). Our focus in this paper is elucidating the universal properties underlying causal inference in a representation-independent manner. We discuss the connections to Fong (2012); Jacobs et al. (2018) at length later in the paper. At the heart of the UCLA hierarchy, as Figure 1 illustrates, is the principle of *universal arrows* MacLane (1971), which we now explain with an example (see Figure 3).

The fundamental property of *universal arrows* MacLane (1971) plays a central role in category theory and in the UCLA architecture. Figure 3 explains this concept with an example, which also illustrates the connection between categories and graphs. For every (directed) graph G , there is a universal arrow from G to the “forgetful” functor U mapping the category **Cat** of all categories to **Graph**, the category of all (directed) graphs, where for any category C , its associated graph is defined by $U(C)$. To understand this functor, simply consider a directed graph $U(C)$ as a category C forgetting the rule for composition. That is, from the category C , which associates to each pair of composable

Table 1: Each layer of UCLA represents a categorical abstraction of causal inference.

Layer	Objects	Morphisms	Description
Simplicial	$[n] = \{0, 1, \dots, n\}$	$f = [m] \rightarrow [n]$	Structure manipulations
Graphical	Vertices V , Edges E	$s, t : E \rightarrow V$	Model interventions
Instances	Sets	Functions on sets $f : S \rightarrow T$	Causal Dataset
Homotopy	Topological Spaces	Continuous functions	Find homotopic equivalences

Figure 2: Simplicial set structure of a causal model, along with three canonical causal structures.



arrows f and g , the composed arrow $g \circ f$, we derive the underlying graph $U(G)$ simply by forgetting which edges correspond to elementary functions, such as f or g , and which are composites. The universal arrow from a graph G to the forgetful functor U is defined as a pair $\langle G, u : G \rightarrow U(C) \rangle$, where u is a graph homomorphism. This arrow possesses the following *universal property*: for every other pair $\langle D, v : G \rightarrow H \rangle$, where D is a category, and v is an arbitrary graph homomorphism, there is a functor $f' : C \rightarrow D$, which is an arrow in the category **Cat** of all categories, such that every graph homomorphism $\phi : G \rightarrow H$ uniquely factors through the universal graph homomorphism $u : G \rightarrow U(C)$ as the solution to the equation $\phi = U(f') \circ u$, where $U(f') : U(C) \rightarrow H$ (that is, $H = U(D)$). Namely, the dotted arrow defines a graph homomorphism $U(f')$ that makes the triangle diagram “commute”, and the associated “extension” problem of finding this new graph homomorphism $U(f')$ is solved by “lifting” the associated category arrow $f' : C \rightarrow D$. This property of universal arrows will, as we show in the paper, provide the conceptual underpinnings of the UCLA architecture, leading to the associated defining property of a universal causal representation through the Yoneda Lemma MacLane (1971).

Table 1 describes the composition of each layer. As Figure 1 illustrates, at the top layer of UCLA, we model causal inference over *simplicial objects* (May, 1992), which is a combinatorial representation of a causal model. Intuitively, causal models are made of “parts” (e.g., vertices and edges in a causal DAG Pearl (2009), objects in a symmetric monoidal category defining a causal DAG Jacobs et al. (2018), or elements of a semi-join lattice defining a conditional independence structure, such as an integer-valued multiset Studeny (2010)). Causal interventions require “surgery” of a causal model, taking parts away. We model this “surgery” process formally by the use of simplicial objects. Formally, simplicial objects May (1992) are contravariant functors $\bar{X} : \Delta^{op} \rightarrow \mathcal{C}$ from the category of ordinal numbers, whose objects are $[n] = \{0, 1, \dots, n\}, n \geq 0$, and whose arrows are non-decreasing maps $f : [m] \rightarrow [n]$, into an underlying category \mathcal{C} . When the underlying category \mathcal{C} is the category of **Sets**, we get simplicial sets. We use the simplicial layer to capture “surgery” of causal structures, including deletion of edges or more complex structures, which are represented using degeneracy and face operators.

Figure 2 illustrates the concept of causal simplicial structures. Here, X denotes a causal structure represented as a category. $X[0]$ represents the “objects” of the causal structure, defined formally as the contravariant functor $X[0] : [0] \rightarrow X$ from the simplicial category Δ to the causal category X . The arrows representing causal effects are defined as $X[1] : [1] \rightarrow X$. Note that since $[1] = \{0, 1\}$ is a category by itself, it has one (non-identity) arrow $0 \rightarrow 1$ (as well as two identity arrows). The mapping of this arrow onto X defines the “edges” of the causal model. Similarly, X_2 represents “triangles” of three objects, which can represent the canonical structures in a causal model, including serial nodes $A \rightarrow B \rightarrow C$, colliders $A \rightarrow B \leftarrow C$, and divergers $A \leftarrow B \rightarrow C$. Note that there is one edge from X_0 to X_1 , labeled by s_0 . This is a co-degeneracy operator from the simplicial layer that maps each object A into an identity edge $\mathbf{1}_A$. Similarly, there are two edges marked d_0 and d_1 from X_1 to X_0 . These are co-face operators that map an edge to its source and target vertices correspondingly. Notice also that there are three edges from X_2 to X_1 , marked d_0, d_1 , and d_2 . These are the “faces” of each triangle, which take, for example, a collider structure $A \rightarrow B \leftarrow C$ and “unpack” it in terms of its constituent edges.

Figure 3: Universal arrows play a central role in the UCLA framework. In this example, the forgetful functor U between **Cat**, the category of all categories, and **Graph**, the category of all (directed) graphs defines a universal arrow. The property asserts that every graph homomorphism problem defined by $\phi : G \rightarrow H$ uniquely factors through the universal graph homomorphism $u : G \rightarrow U(C)$ defined by the category C defining the universal arrow property. In other words, the associated *extension* problem of “completing” the triangle of graph homomorphisms in the category of **Graph** can be uniquely solved by “lifting” the associated category arrow $h : C \rightarrow D$.

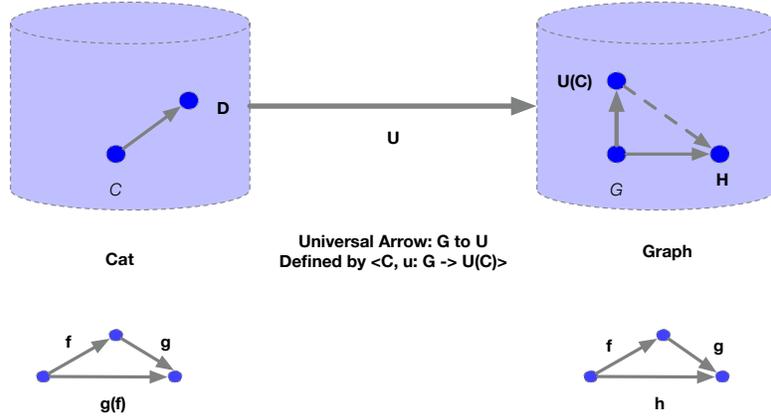
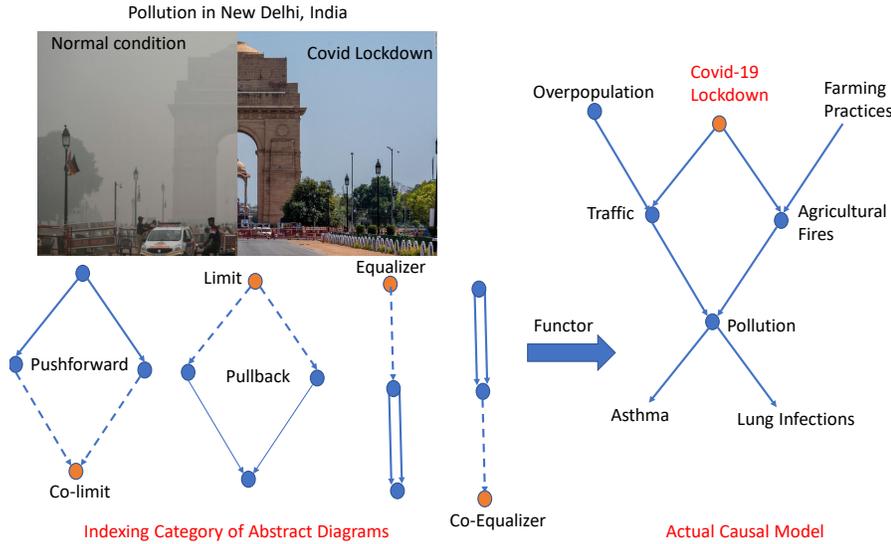


Figure 4: A causal model of climate change and Covid-19 lockdown. Universal causality interprets causal diagrams in terms of limits and co-limits of an indexing category of abstract diagrams.



As shown in Figure 4, we can interpret causal DAG-like models in category theory in terms of functors mapping from some index category of diagrams into the actual causal model. Category theory defines *universal properties*, including limits, colimits, and equalizers, which helps clarify properties of causal diagrams. We will extensively build on the property of *universal arrows* MacLane (1971) in this paper, which serves as the conceptual underpinning how causal effects can be transmitted faithfully from one layer of the UCLA hierarchy to the next.

Functors define the correspondence between each pair of categorical abstractions in Figure 1. Thus, for example, the causal DAG model shown in layer 2 of the figure can be defined through a contravariant functor $X : \Delta^{op} \rightarrow \mathcal{C}$, where \mathcal{C} represents the category of causal models. The 0th order simplicial objects defined as $X[0] = X_0$ represent the causal variables A, B , and C . The 1-simplices X_1 define directional causal relationships represented as 0-simplices, or arrows in category \mathcal{C} , in this case, the two arrows between A and C , and between B and C . Associated with each such functor

between two layers is a universal arrow MacLane (1971), which specifies a universal property that ensures the functor faithfully represents the set of morphisms from the layer above to the layer below. Associated with the universal arrow is a construction defined by Grothendieck, which defines an associated category of elements $\int X$, whose objects are pairs (n, i) , where n is an object in the category Δ , and i is an object in the category X , such that the morphism $f : [2] \rightarrow [0]$ in Δ is contravariantly mapped by the functor to the edge between A and C in the causal model.

As Figure 1 illustrates, each pair of categorical layers is associated with a lifting problem that defines a language for expressing and answering queries between those levels. A lifting problem in category \mathcal{C} is a commutative diagram σ in \mathcal{C} in which each composable pair of arrows defines path equivalences. Gavrilovich (2017) shows a wide range of problems in topology and other areas can be represented as solutions to lifting problems. As a concrete application to causal discovery, recent work on the use of *separating sets* Kocaoglu et al. (2017) can be viewed as an instance of imposing a Kolmogorov T_0 topology on the variables in a causal model. The T_0 topology separation condition is an example of a property that can be stated in terms of a lifting problem. In simple terms, a category can be viewed as a directed graph, where certain paths can be defined as equivalent. So, in the diagram below, the path $nu \circ f = p \circ \mu$, where each arrow is labeled with a particular morphism.

$$\begin{array}{ccc} A & \xrightarrow{\mu} & X \\ \downarrow f & & \downarrow p \\ B & \xrightarrow{\nu} & Y \end{array}$$

A solution to a lifting problem in \mathcal{C} is a morphism $h : B \rightarrow X$ in \mathcal{C} satisfying $p \circ h = \nu$ and $h \circ f = \mu$ as indicated in the diagram below.

$$\begin{array}{ccc} A & \xrightarrow{\mu} & X \\ \downarrow f & \nearrow h & \downarrow p \\ B & \xrightarrow{\nu} & Y \end{array}$$

As indicated in Figure 1, each pair of layers in UCLA is accompanied by a lifting problem that is defined by the induced category of elements $\int F$, for some particular functor F . For example, the category of lifting problems between the causal model at the second layer from the top to the instance layer below it can be viewed as formalizing the process of inference in relational databases Schultz et al. (2016); Brown et al. (2019); Spivak (2010). In particular, Spivak (2013) showed that database queries in languages like SQL can be formalized in terms of lifting problems. In this particular setting, the functor used maps a relational database model defined by a category \mathcal{D} into a set of interlocking tables defined as a category of **Sets**. Thus, relational databases can be viewed as set-valued functors $F : \mathcal{D} \rightarrow \mathbf{Sets}$. In this particular case, the associated lifting problem defined by the Grothendieck category of elements enables formalizing the process of answering SQL queries. The same construction, when applied to categories other than **Sets** gives us the ability to formulate causal inference over other categories, including topological spaces.

Many properties of Grothendieck’s category of elements can be exploited (some of these are discussed in the context of relational database queries in Spivak (2013)), but for our application to causal inference, we are primarily interested in the associated class of lifting problems that define queries in a causal model. In particular, it can be shown that Grothendieck’s category of elements formally defines a type of fibration called a *Grothendieck opfibration*. More specifically, if we consider any object $c \in \mathcal{C}$ in a category \mathcal{C} , and look at its fiber $\pi_\delta^{-1}(c)$, which forms a subcategory of the category of elements I , it can be shown that the projection of every morphism in the fiber into the category \mathcal{C} is to the identity element $\mathbf{1}_c$. Richter (2020) explains how monoidal categories themselves can be defined using simplicial objects using a Grothendieck opfibration. We refer the reader to these sources for additional details. One implication of this property is that approaches to UC that build symmetric monoidal category representations of Bayesian networks Fong (2012); Jacobs et al. (2018) can be formally shown to define special cases of the UCLA hierarchy.

Finally, to explain the bottom-most layer in UCLA of homotopy categories, it is well known that many DAG models are not identifiable from observations alone Pearl (2009). Non-identifiability of DAG models, such as the serial model $A \rightarrow B \rightarrow C$, the diverging model $B \leftarrow A \rightarrow C$ and the reverse serial model $C \rightarrow B \rightarrow A$, occurs because given the parameterization of one of the models, the parameterization of the others can be derived using Bayes rule. For example, the serial model is parameterized in terms of the marginal probability $P(A)$, and the conditional probabilities $P(B|A)$ and $P(C|B)$. The diverger is parameterized using $P(B), P(A|B), P(C|B)$. Using Bayes rule, we have $P(AB) = P(A|B)P(B) = P(B|A)P(A)$, and similarly, $P(CB) = P(B|C)P(C) = P(C|B)P(B)$, we can derive the parameterization of the reverse serial model from the serial and diverger models. In contrast, for the converging model, $A \rightarrow B \leftarrow C$, its parameterization is in terms of $P(B|AC)$, which cannot be derived from the others using Bayes rule.

More broadly, a fundamental premise underlying UC is that *presheaves* – the set of morphisms $\mathbf{Hom}_{\mathcal{C}}(-, X)$ entering an object X in a category \mathcal{C} – can serve as universal representers of causal inference. Causal diagrams in UC are defined as functors, mapping from some abstract category of diagrams into a causal model, and use universal constructions, including the *pullback*, the *pushforward*, the *(co)equalizer*, and more generally, the *(co)limit* and the Kan extension. A standard construction in category theory shows that any set-valued contravariant functor can be represented as a co-limit of a diagram, which is defined through the category of elements MacLane and Ieke Moerdijk (1994). The concept of representable presheaves comes from the Yoneda Lemma, which states that any presheaf is representable in the category **Set** by a fully faithful embedding. From the Yoneda Lemma also follows straightforwardly the property that the causal influence of any object X on another object Y , denoted by $\mathbf{Hom}_{\mathcal{C}}(X, Y)$, can be defined as a natural transformation between two representable presheaves. In most applications, presheaves carry significant additional structure, which can be modeled by an enriched category over monoidal categories, including rings, topological spaces, measurable spaces, and probability spaces.

Given this high level overview of UC and the UCLA architecture, in the remainder of the paper, we describe the various layers of the UCLA architecture in more detail, and explain the significance of the underlying conceptual ideas, including the use of the universal arrows and the category of elements from the Grothendieck construction. We will illustrate these concepts using a range of examples from previous literature.

2 Categories, Functors, and Universal Arrows

As the relationship between categories and graphs is a close one, which has significant influence in causal inference, we define categories as essentially directed graphs equipped with a composition property Borceux (1994). The composition rule in category theory enables specifying which sets of paths should be considered equivalent, which manifests itself in the widespread use of commutative diagrams, as will be seen in this paper. Given a graph, we can define the “free” category associated with it where we consider all possible paths between pairs of vertices (including self-loops) as the set of morphisms between them. In the reverse direction, given a category, we can define a “forgetful” functor that extracts the underlying graph from the category, forgetting the composition rule. This process of going from one to the other embodies a fundamental underlying principle in category theory, called the *universal arrow* MacLane (1971).

Definition 1. A **graph** \mathcal{G} (sometimes referred to as a quiver) is a labeled directed multi-graph defined by a set O of *objects*, a set A of *arrows*, along with two morphisms $s : A \rightarrow O$ and $t : A \rightarrow O$ that specify the domain and co-domain of each arrow. In this graph, we define the set of composable pairs of arrows by the set

$$A \times_O A = \{\langle g, f \rangle \mid g, f \in A, s(g) = t(f)\}$$

A **category** \mathcal{C} is a graph \mathcal{G} with two additional functions: $\mathbf{id} : O \rightarrow A$, mapping each object $c \in C$ to an arrow \mathbf{id}_c and $\circ : A \times_O A \rightarrow A$, mapping each pair of composable morphisms $\langle f, g \rangle$ to their composition $g \circ f$.

It is worth emphasizing that no assumption is made here of the finiteness of a graph, either in terms of its associated objects (vertices) or arrows (edges). Indeed, it is entirely reasonable to define categories whose graphs contain an infinite number of edges. A simple example is the group \mathbb{Z} of integers under addition, which can be represented as a single object, denoted $\{\bullet\}$ and an infinite number of morphisms $f : \bullet \rightarrow \bullet$, each of which represents an integer, where composition of morphisms is defined by addition. In this example, all morphisms are invertible. In a general category with more than one object, a *groupoid* defines a category all of whose morphisms are invertible.

As our paper focuses on the use of category theory to formalize causal inference, we interpret causal changes in terms of the concept of isomorphisms in category theory. We will elaborate this definition later in the paper.

Definition 2. Two objects X and Y in a category \mathcal{C} are deemed **isomorphic**, or $X \cong Y$ if and only if there is an invertible morphism $f : X \rightarrow Y$, namely f is both *left invertible* using a morphism $g : Y \rightarrow X$ so that $g \circ f = \mathbf{id}_X$, and f is *right invertible* using a morphism h where $f \circ h = \mathbf{id}_Y$. A **causally isomorphic change** in a category is defined as a change of a causal object Y into \hat{Y} under an intervention that changes another object X into \hat{X} such that $\hat{Y} \cong Y$, that is, they are isomorphic. A **causal non-isomorphic effect** is a change that leads to a non-isomorphic change where $\hat{Y} \not\cong Y$.

In the category **Sets**, two finite sets are considered isomorphic if they have the same number of elements, as it is then trivial to define an invertible pair of morphisms between them. In the category \mathbf{Vect}_k of vector spaces over some field k , two objects (vector spaces) are isomorphic if there is a set of invertible linear transformations between them. As we will see below, the passage from a set to the “free” vector space generated by elements of the set is another manifestation of the universal arrow property.

Definition 3. A **covariant functor** $F : \mathcal{C} \rightarrow \mathcal{D}$ from category \mathcal{C} to category \mathcal{D} , and defined as the following:

- An object FX (sometimes written as $F(x)$) of the category \mathcal{D} for each object X in category \mathcal{C} .
- An arrow $F(f) : FX \rightarrow FY$ in category \mathcal{D} for every arrow $f : X \rightarrow Y$ in category \mathcal{C} .
- The preservation of identity and composition: $F id_X = id_{FX}$ and $(Ff)(Fg) = F(g \circ f)$ for any composable arrows $f : X \rightarrow Y, g : Y \rightarrow Z$.

Definition 4. A **contravariant functor** $F : \mathcal{C} \rightarrow \mathcal{D}$ from category \mathcal{C} to category \mathcal{D} is defined exactly like the covariant functor, except all the arrows are reversed. In the contravariant functor $F : \mathcal{C}^{\text{OP}} \rightarrow \mathcal{D}$, every morphism $f : X \rightarrow Y$ is assigned the reverse morphism $Ff : FY \rightarrow FX$ in category \mathcal{D} .

- For every object X in a category \mathcal{C} , there exists a covariant functor $\mathcal{C}(X, -) : \mathcal{C} \rightarrow \mathbf{Set}$ that assigns to each object Z in \mathcal{C} the set of morphisms $\mathcal{C}(X, Z)$, and to each morphism $f : Y \rightarrow Z$, the pushforward mapping $f_* : \mathcal{C}(X, Y) \rightarrow \mathcal{C}(X, Z)$.
- For every object X in a category \mathcal{C} , there exists a contravariant functor $\mathcal{C}(-, X) : \mathcal{C}^{\text{OP}} \rightarrow \mathbf{Set}$ that assigns to each object Z in \mathcal{C} the set of morphisms $\mathbf{Hom}_{\mathcal{C}}(X, Z)$, and to each morphism $f : Y \rightarrow Z$, the pullback mapping $f^* : \mathbf{Hom}_{\mathcal{C}}(Z, X) \rightarrow \mathcal{C}(Y, X)$. Note how “contravariance” implies the morphisms in the original category are reversed through the functorial mapping, whereas in covariance, the morphisms are not flipped.

2.1 Universal Arrows

Definition 5. Given a functor $S : D \rightarrow C$ between two categories, and an object c of category C , a **universal arrow** from c to S is a pair $\langle r, u \rangle$, where r is an object of D and $u : c \rightarrow Sr$ is an arrow of C , such that the following universal property holds true:

- For every pair $\langle d, f \rangle$ with d an object of D and $f : c \rightarrow Sd$ an arrow of C , there is a unique arrow $f' : r \rightarrow d$ of D with $Sf' \circ u = f$.

Example 1. Consider the example of universal arrows where \mathbf{Vect}_K denotes the category of vector spaces over some field K , and its arrows correspond to linear transformations. Let U denote the “forgetful” functor $U : \mathbf{Vect}_K \rightarrow \mathbf{Set}$ that maps each vector space V to its associated set of elements. Given any set X , we can define an associated vector space V_X choosing the elements $x \in X$ as its basis vectors, and forming all formal linear combinations $\sum_i k_i x_i$, where $k_i \in K$ and $x_i \in X$. The mapping from each set element $x \in X$ into the corresponding vector in V_X is an arrow $\lambda : X \rightarrow U(V_X)$. For any other vector space W , each function $\gamma : X \rightarrow U(W)$ can be transformed into a unique linear transformation $\delta : V_X \rightarrow W$ such that $U\delta \circ \lambda = \gamma$. Thus, here λ acts as a universal arrow from the set X to the forgetful functor U .

2.2 The Category of Fractions

The problem of defining a category with a given subclass of invertible morphisms, called the category of fractions (Gabriel et al., 1967), is another concrete illustration of the close relationships between categories and graphs. It is also useful in the context of causal inference, as for example, in defining the Markov equivalence class of directed acyclic graphs (DAGs) as a category that is localized by considering all invertible arrows as isomorphisms. Borceux (1994) has a detailed discussion of the “calculus of fractions”, namely how to define a category where a subclass of morphisms are to be treated as isomorphisms (e.g., the three models in Figure 7 are equivalent in the sense that the morphisms are invertible). The formal definition is as follows:

Definition 6. Consider a category \mathcal{C} and a class Σ of arrows of \mathcal{C} . The **category of fractions** $\mathcal{C}(\Sigma^{-1})$ is said to exist when a category $\mathcal{C}(\Sigma^{-1})$ and a functor $\phi : \mathcal{C} \rightarrow \mathcal{C}(\Sigma^{-1})$ can be found with the following properties:

1. $\forall f, \phi(f)$ is an isomorphism.
2. If \mathcal{D} is a category, and $F : \mathcal{C} \rightarrow \mathcal{D}$ is a functor such that for all morphisms $f \in \Sigma$, $F(f)$ is an isomorphism, then there exists a unique functor $G : \mathcal{C}(\Sigma^{-1}) \rightarrow \mathcal{D}$ such that $G \circ \phi = F$.

A detailed construction of the category of fractions is given in Borceux (1994), which uses the underlying directed graph skeleton associated with the category. The characterization of the Markov equivalent class of acyclic directed graphs is an example of the abstract concept of category of fractions (Andersson et al., 1997).² In our previous work Mahadevan

²Briefly, this condition states that two acyclic directed graphs are Markov equivalent if and only if they have the same skeleton and the same immoralities.

(2021b), we explored constructing homotopically invariant causal models over finite Alexandroff topological spaces, which can be seen as a special case of the UCLA framework.

2.3 Lifting Problems

The UCLA hierarchy is defined through a series of categorical abstractions of a causal model, ranging from a combinatorial model defined by a simplicial object down to a measure-theoretic or topological realization. Between each pair of layers, we can formulate a series of lifting problems Gavrilovich (2017). Lifting problems provide elegant ways to define basic notions in a wide variety of areas in mathematics. For example, the notion of injective and surjective functions, the notion of separation in topology, and many other basic constructs can be formulated as solutions to lifting problems. Database queries in relational databases can be defined using lifting problems Spivak (2013). Lifting problems define ways of decomposing structures into simpler pieces, and putting them back together again.

Definition 7. Let \mathcal{C} be a category. A **lifting problem** in \mathcal{C} is a commutative diagram σ in \mathcal{C} .

$$\begin{array}{ccc} A & \xrightarrow{\mu} & X \\ \downarrow f & & \downarrow p \\ B & \xrightarrow{\nu} & Y \end{array}$$

Definition 8. Let \mathcal{C} be a category. A **solution to a lifting problem** in \mathcal{C} is a morphism $h : B \rightarrow X$ in \mathcal{C} satisfying $p \circ h = \nu$ and $h \circ f = \mu$ as indicated in the diagram below.

$$\begin{array}{ccc} A & \xrightarrow{\mu} & X \\ \downarrow f & \nearrow h & \downarrow p \\ B & \xrightarrow{\nu} & Y \end{array}$$

Definition 9. Let \mathcal{C} be a category. If we are given two morphisms $f : A \rightarrow B$ and $p : X \rightarrow Y$ in \mathcal{C} , we say that f has the **left lifting property** with respect to p , or that p has the **right lifting property** with respect to f if for every pair of morphisms $\mu : A \rightarrow X$ and $\nu : B \rightarrow Y$ satisfying the equations $p \circ \mu = \nu \circ f$, the associated lifting problem indicated in the diagram below.

$$\begin{array}{ccc} A & \xrightarrow{\mu} & X \\ \downarrow f & \nearrow h & \downarrow p \\ B & \xrightarrow{\nu} & Y \end{array}$$

admits a solution given by the map $h : B \rightarrow X$ satisfying $p \circ h = \nu$ and $h \circ f = \mu$.

Example 2. Given the paradigmatic non-surjective morphism $f : \emptyset \rightarrow \{\bullet\}$, any morphism p that has the right lifting property with respect to f is a **surjective mapping**.

$$\begin{array}{ccc} \emptyset & \xrightarrow{\mu} & X \\ \downarrow f & \nearrow h & \downarrow p \\ \{\bullet\} & \xrightarrow{\nu} & Y \end{array}$$

Example 3. Given the paradigmatic non-injective morphism $f : \{\bullet, \bullet\} \rightarrow \{\bullet\}$, any morphism p that has the right lifting property with respect to f is an **injective mapping**.

$$\begin{array}{ccc} \{\bullet, \bullet\} & \xrightarrow{\mu} & X \\ \downarrow f & \nearrow h & \downarrow p \\ \{\bullet\} & \xrightarrow{\nu} & Y \end{array}$$

3 Layers 1 and 2: Simplicial Objects and Causal Interventions

We now begin our discussion of the UCLA architecture shown earlier in Figure 1, describing the top simplicial objects layer, and how it interacts with the causal category structure (layer 2). Abstractly, at the top layer, we want to define the

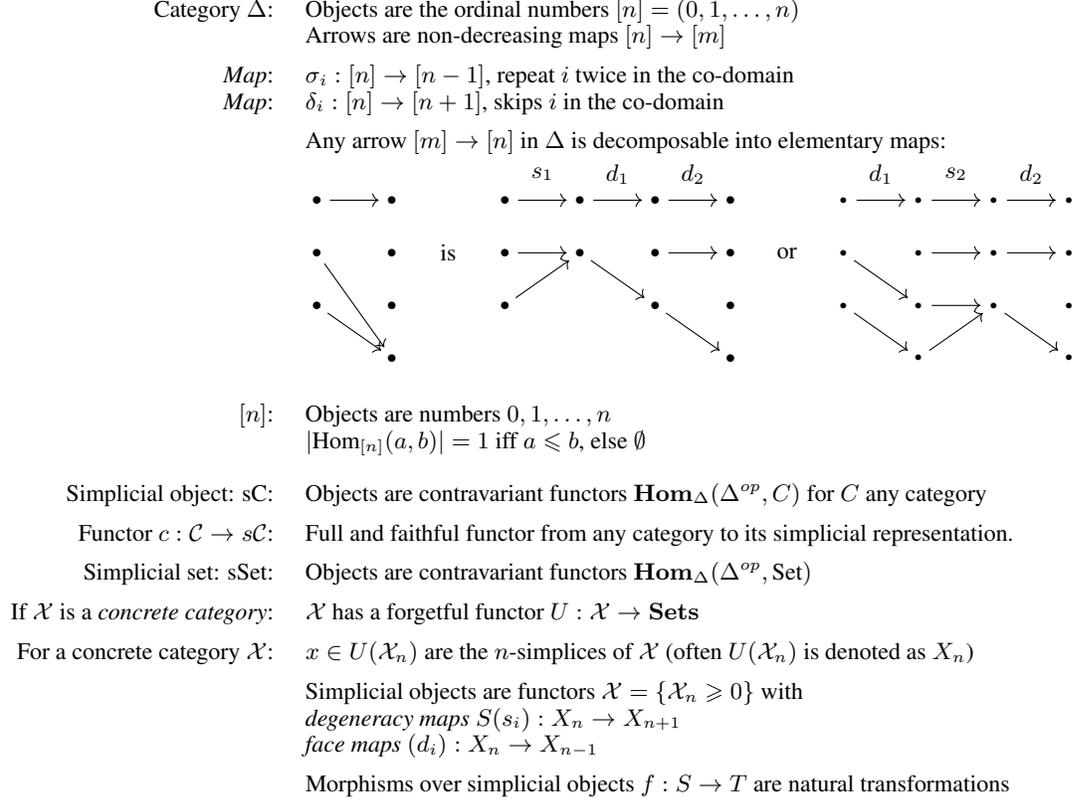


Figure 5: Simplicial objects are a combinatorial representation for manipulating causal categories.

categorical machinery to do “diagram surgery”, as in “graph surgery” Pearl (2009) or “string diagram surgery” Jacobs et al. (2018). The structure of simplicial objects and sets helps us achieve this goal in a representation-independent manner, by using a category of ordinal numbers $[n] = \{0, \dots, n\}$ with morphisms defined that respect the natural ordering, and using degeneracy and face operators to implement “surgery” operations.

Simplicial objects have long been a foundation for algebraic topology (May, 1999, 1992), and more recently in infinite category theory (Boardman and Vogt, 1973; Joyal, 2002; Lurie, 2022). The category Δ has non-empty ordinals $[n] = \{0, 1, \dots, n\}$ as objects, and order-preserving maps $[m] \rightarrow [n]$ as arrows. An important property in Δ is that any many-to-many mapping is decomposable as a composition of an injective and a surjective mapping, each of which is decomposable into a sequence of elementary injections $\delta_i : [n] \rightarrow [n+1]$, called *coface* mappings, which omits $i \in [n]$, and a sequence of elementary surjections $\sigma_i : [n] \rightarrow [n-1]$, called *co-degeneracy* mappings, which repeats $i \in [n]$. The fundamental simplex $\Delta([n])$ is the presheaf of all morphisms into $[n]$, that is, the representable functor $\Delta(-, [n])$. The Yoneda Lemma MacLane (1971) assures us that an n -simplex $x \in X_n$ can be identified with the corresponding map $\Delta([n]) \rightarrow X$. Every morphism $f : [n] \rightarrow [m]$ in Δ is functorially mapped to the map $\Delta([m]) \rightarrow \Delta([n])$ in \mathcal{S} . See Figure 5.

Any morphism in the category Δ can be defined as a sequence of *co-degeneracy* and *co-face* operators, where the co-face operator $\delta_i : [n-1] \rightarrow [n]$, $0 \leq i \leq n$ is defined as:

$$\delta_i(j) = \begin{cases} j, & \text{for } 0 \leq j \leq i-1 \\ j+1 & \text{for } i \leq j \leq n-1 \end{cases}$$

Analogously, the co-degeneracy operator $\sigma_j : [n+1] \rightarrow [n]$ is defined as

$$\sigma_j(k) = \begin{cases} j, & \text{for } 0 \leq k \leq j \\ k-1 & \text{for } j < k \leq n+1 \end{cases}$$

Note that under the contravariant mappings, co-face mappings turn into face mappings, and co-degeneracy mappings turn into degeneracy mappings. That is, for any simplicial object (or set) X_n , we have $X(\delta_i) := d_i : X_n \rightarrow X_{n-1}$, and likewise, $X(\sigma_j) := s_j : X_{n-1} \rightarrow X_n$.

The compositions of these arrows define certain well-known properties (May, 1992; Richter, 2020):

$$\begin{aligned} \delta_j \circ \delta_i &= \delta_i \circ \delta_{j-1}, \quad i < j \\ \sigma_j \circ \sigma_i &= \sigma_i \circ \sigma_{j+1}, \quad i \leq j \\ \sigma_j \circ \delta_i(j) &= \begin{cases} \sigma_i \circ \sigma_{j+1}, & \text{for } i < j \\ 1_{[n]} & \text{for } i = j, j + 1 \\ \sigma_{i-1} \circ \sigma_j, & \text{for } i > j + 1 \end{cases} \end{aligned}$$

Example 4. The ‘‘vertices’’ of a simplicial object \mathcal{C}_n are the objects in \mathcal{C} , and the ‘‘edges’’ of \mathcal{C} are its arrows $f : X \rightarrow Y$, where X and Y are objects in \mathcal{C} . Given any such arrow, the degeneracy operators $d_0 f = Y$ and $d_1 f = X$ recover the source and target of each arrow. Also, given an object X of category \mathcal{C} , we can regard the face operator $s_0 X$ as its identity morphism $1_X : X \rightarrow X$.

Example 5. Given a category \mathcal{C} , we can identify an n -simplex σ of a simplicial set \mathcal{C}_n with the sequence:

$$\sigma = C_0 \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots \xrightarrow{f_n} C_n$$

the face operator d_0 applied to σ yields the sequence

$$d_0 \sigma = C_1 \xrightarrow{f_2} C_2 \xrightarrow{f_3} \dots \xrightarrow{f_n} C_n$$

where the object C_0 is ‘‘deleted’’ along with the morphism f_0 leaving it.

Example 6. Given a category \mathcal{C} , and an n -simplex σ of the simplicial set \mathcal{C}_n , the face operator d_n applied to σ yields the sequence

$$d_n \sigma = C_0 \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots \xrightarrow{f_{n-1}} C_{n-1}$$

where the object C_n is ‘‘deleted’’ along with the morphism f_n entering it. Note this process can be used to implement ‘‘surgery’’ of a causal model, such as a causal DAG Pearl (2009), or a symmetric monoidal category Jacobs et al. (2018).

Example 7. Given a category \mathcal{C} , and an n -simplex σ of the simplicial set \mathcal{C}_n the face operator $d_i, 0 < i < n$ applied to σ yields the sequence

$$d_i \sigma = C_0 \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots C_{i-1} \xrightarrow{f_{i+1} \circ f_i} C_{i+1} \dots \xrightarrow{f_n} C_n$$

where the object C_i is ‘‘deleted’’ and the morphisms f_i is composed with morphism f_{i+1} . Note that this process can be abstractly viewed as intervening on object C_i by choosing a specific value for it (which essentially ‘‘freezes’’ the morphism f_i entering object C_i to a constant value).

Example 8. Given a category \mathcal{C} , and an n -simplex σ of the simplicial set \mathcal{C}_n , the degeneracy operator $s_i, 0 \leq i \leq n$ applied to σ yields the sequence

$$s_i \sigma = C_0 \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots C_i \xrightarrow{1_{C_i}} C_i \xrightarrow{f_{i+1}} C_{i+1} \dots \xrightarrow{f_n} C_n$$

where the object C_i is ‘‘repeated’’ by inserting its identity morphism 1_{C_i} .

Definition 10. Given a category \mathcal{C} , and an n -simplex σ of the simplicial set \mathcal{C}_n , σ is a **degenerate** simplex if some f_i in σ is an identity morphism, in which case C_i and C_{i+1} are equal.

3.1 Simplicial Subsets and Horns

We now describe more complex ways of extracting parts of causal structures using simplicial subsets and horns. These structures will play a key role in defining suitable lifting problems.

Definition 11. The **standard simplex** Δ^n is the simplicial set defined by the construction

$$([m] \in \Delta) \mapsto \mathbf{Hom}_\Delta([m], [n])$$

By convention, $\Delta^{-1} := \emptyset$. The standard 0-simplex Δ^0 maps each $[n] \in \Delta^{op}$ to the single element set $\{\bullet\}$.

Definition 12. Let S_\bullet denote a simplicial set. If for every integer $n \geq 0$, we are given a subset $T_n \subseteq S_n$, such that the face and degeneracy maps

$$d_i : S_n \rightarrow S_{n-1} \quad s_i : S_n \rightarrow S_{n+1}$$

applied to T_n result in

$$d_i : T_n \rightarrow T_{n-1} \quad s_i : T_n \rightarrow T_{n+1}$$

then the collection $\{T_n\}_{n \geq 0}$ defines a **simplicial subset** $T_\bullet \subseteq S_\bullet$.

Definition 13. The **boundary** is a simplicial set $(\partial\Delta^n) : \Delta^{op} \rightarrow \mathbf{Set}$ defined as

$$(\partial\Delta^n)([m]) = \{\alpha \in \mathbf{Hom}_\Delta([m], [n]) : \alpha \text{ is not surjective}\}$$

Note that the boundary $\partial\Delta^n$ is a simplicial subset of the standard n -simplex Δ^n .

Definition 14. The **Horn** $\Lambda_i^n : \Delta^{op} \rightarrow \mathbf{Set}$ is defined as

$$(\Lambda_i^n)([m]) = \{\alpha \in \mathbf{Hom}_\Delta([m], [n]) : [n] \not\subseteq \alpha([m]) \cup \{i\}\}$$

Intuitively, the Horn Λ_i^n can be viewed as the simplicial subset that results from removing the interior of the n -simplex Δ^n together with the face opposite its i th vertex.

3.2 Example: Causal Intervention and Horn Filling of Simplicial Complexes

Let us illustrate this abstract discussion above by instantiating it in the context of causal inference. Figure 6 instantiates the abstract discussion above in terms of an example from causal inference. We are given a simple 3 variable DAG, on which we desire to explore the causal effect of variable A on C . Using Pearl’s backdoor criterion, we can intervene on variable A by freezing its value $do(A = 1)$, for example, which will eliminate the dependence of A on B . Consider now the lifting problem where we want to know if there is a completion of this simplicial subset Λ_2^2 , which is a “outer horn”

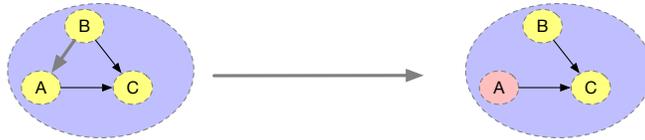
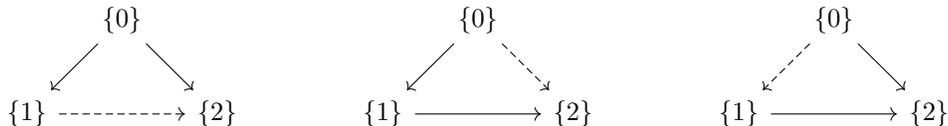


Figure 6: Causal interventions can be related to horn filling of a simplicial complex.

We can view the causal intervention problem in the more abstract setting of a class of lifting problem, shown with the following diagrams. Consider the problem of composing 1-dimensional simplices to form a 2-dimensional simplicial object. Each simplicial subset of an n -simplex induces a *horn* Λ_k^n , where $0 \leq k \leq n$. Intuitively, a horn is a subset of a simplicial object that results from removing the interior of the n -simplex and the face opposite the i th vertex. Consider the three horns defined below. The dashed arrow \dashrightarrow indicates edges of the 2-simplex Δ^2 not contained in the horns.



The inner horn Λ_1^2 is the middle diagram above, and admits an easy solution to the ‘‘horn filling’’ problem of composing the simplicial subsets. The two outer horns on either end pose a more difficult challenge. A considerable elaboration of the theoretical machinery in category theory is required to describe the various solutions proposed, which led to different ways of defining higher-order category theory (Boardman and Vogt, 1973; Joyal, 2002; Lurie, 2022). Note that in our causal intervention example above shown in Figure 6, the horn filling problem is at the rightmost (where we map $A \rightarrow \{1\}, B \rightarrow \{0\}, C \rightarrow \{2\}$). We know that this lifting problem indeed has a solution because the completion problem can be solved as the simplicial subset was created by intervention on node A . To elaborate on how we can solve such a lifting problem, we will postpone this discussion to the next section, where we show that causal interventions can be modeled as lifting problems in the category of elements (which will enable us to ‘‘bind’’ objects such as A to constants, such as $A = 1$). For the sake of completeness, we include a more detailed discussion of the horn filling problem, and its various solutions in higher-order category theory.

3.3 Homotopy of Simplicial Objects

We will discuss homotopy of causal structures in more detail below, for now we briefly discuss a few notions that will be useful in this section.

Definition 15. Let C and C' be a pair of objects in a category \mathcal{C} . We say C is a **retract** of C' if there exists maps $i : C \rightarrow C'$ and $r : C' \rightarrow C$ such that $r \circ i = \text{id}_C$.

Definition 16. Let \mathcal{C} be a category. We say a morphism $f : C \rightarrow D$ is a **retract of another morphism** $f' : C \rightarrow D$ if it is a retract of f' when viewed as an object of the functor category $\mathbf{Hom}([1], \mathcal{C})$. A collection of morphisms T of \mathcal{C} is **closed under retracts** if for every pair of morphisms f, f' of \mathcal{C} , if f is a retract of f' , and f' is in T , then f is also in T .

Definition 17. Let X and Y be simplicial sets, and suppose we are given a pair of morphisms $f_0, f_1 : X \rightarrow Y$. A **homotopy** from f_0 to f_1 is a morphism $h : \Delta^1 \times X \rightarrow Y$ satisfying $f_0 = h|_{0 \times X}$ and $f_1 = h|_{1 \times X}$.

This notion of homotopy generalizes the notion of homotopy in topology, which defines why an object like a coffee cup is topologically homotopic to a doughnut (they have the same number of ‘‘holes’’).

3.4 Fibrations and Kan Complexes

Definition 18. Let $f : X \rightarrow S$ be a morphism of simplicial sets. We say f is a **Kan fibration** if, for each $n > 0$, and each $0 \leq i \leq n$, every lifting problem

$$\begin{array}{ccc} \Lambda_i^n & \xrightarrow{\sigma_0} & X \\ \downarrow & \nearrow \sigma & \downarrow f \\ \Delta^n & \xrightarrow{\bar{\sigma}} & S \end{array}$$

admits a solution. More precisely, for every map of simplicial sets $\sigma_0 : \Lambda_i^n \rightarrow X$ and every n -simplex $\bar{\sigma} : \Delta^n \rightarrow S$ extending $f \circ \sigma_0$, we can extend σ_0 to an n -simplex $\sigma : \Delta^n \rightarrow X$ satisfying $f \circ \sigma = \bar{\sigma}$.

Example 9. Given a simplicial set X , then a projection map $X \rightarrow \Delta^0$ that is a Kan fibration is called a **Kan complex**.

Example 10. Any isomorphism between simplicial sets is a Kan fibration.

Example 11. The collection of Kan fibrations is closed under retracts.

3.5 Higher-order Categories

We now formally introduce higher-order categories, building on the framework proposed in a number of formalisms (Boardman and Vogt, 1973; Joyal, 2002; Lurie, 2022).

Definition 19. (Lurie, 2022) An ∞ -category is a simplicial object S_\bullet which satisfies the following condition:

- For $0 < i < n$, every map of simplicial sets $\sigma_0 : \Lambda_i^n \rightarrow S_\bullet$ can be extended to a map $\sigma : \Delta^n \rightarrow S_i$.

This definition emerges out of a common generalization of two other conditions on a simplicial set S_i :

1. **Property K:** For $n > 0$ and $0 \leq i \leq n$, every map of simplicial sets $\sigma_0 : \Lambda_i^n \rightarrow S_\bullet$ can be extended to a map $\sigma : \Delta^n \rightarrow S_i$.
2. **Property C:** for $0 < i < n$, every map of simplicial sets $\sigma_0 : \Lambda_i^n \rightarrow S_i$ can be extended uniquely to a map $\sigma : \Delta^n \rightarrow S_i$.

Simplicial objects that satisfy property K were defined above to be Kan complexes. Simplicial objects that satisfy property C above can be identified with the nerve of a category, which yields a full and faithful embedding of a category in the category of sets. Definition 19 generalizes both of these definitions, and was called a *quasicategory* in (Joyal, 2002) and *weak Kan complexes* in (Boardman and Vogt, 1973) when \mathcal{C} is a category.

3.6 Example: Simplicial Objects over Integer-Valued Multisets

To help ground out this somewhat abstract discussion above on simplicial objects and sets, let us consider its application to two other examples. Our first example comes from a non-graphical representations of conditional independence, namely integer-valued multisets (Studený, 2010), defined as an integer-valued multiset function $u : \mathbb{Z}^{\mathcal{P}(\mathbb{Z})} \rightarrow \mathbb{Z}$ from the power set of integers, $\mathcal{P}(\mathbb{Z})$ to integers \mathbb{Z} . An imset is defined over partially ordered set (poset), defined as a distributive lattice of disjoint (or non-disjoint) subsets of variables. The bottom element is denoted \emptyset , and top element represents the complete set of variables N . A full discussion of the probabilistic representations induced by imsets is given (Studený, 2010). We will only focus on the aspects of imsets that relate to its conditional independence structure, and its topological structure as defined by the poset. A *combinatorial* imset is defined as:

$$u = \sum_{A \subset N} c_A \delta_A$$

where c_A is an integer, δ_A is the characteristic function for subset A , and A potentially ranges over all subsets of N . An *elementary* imset is defined over $(a, b \perp\!\!\!\perp A)$, where a, b are singletons, and $A \subset N \setminus \{a, b\}$. A *structural* imset is defined as one where the coefficients can be rational numbers. For a general DAG model $G = (V, E)$, an imset in standard form (Studený, 2010) is defined as

$$u_G = \delta_V - \delta_\emptyset + \sum_{i \in V} (\delta_{\mathbf{Pa}_i} - \delta_{i \cup \mathbf{Pa}_i})$$

Figure 7 shows an example imset for DAG models over three variables, defined by an integer valued function over the lattice of subsets. Each of the three DAG models shown defines exactly the same imset function. Studený (2010) gives a detailed analysis of imsets as a non-graphical representation of conditional independence.

Figure 7: An illustration of an integer-valued multiset (imset) consisting of a lattice of subsets over three elements for representing conditional independences in DAG models. All three DAG models are represented by the same imset. We can view an imset as a simplicial object (shown as the dark blue filled triangle) whose simplices map to integers.

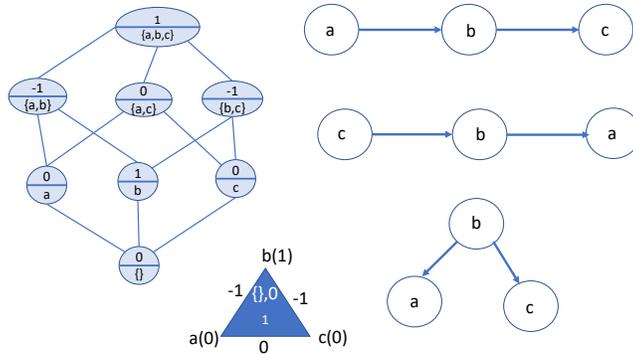
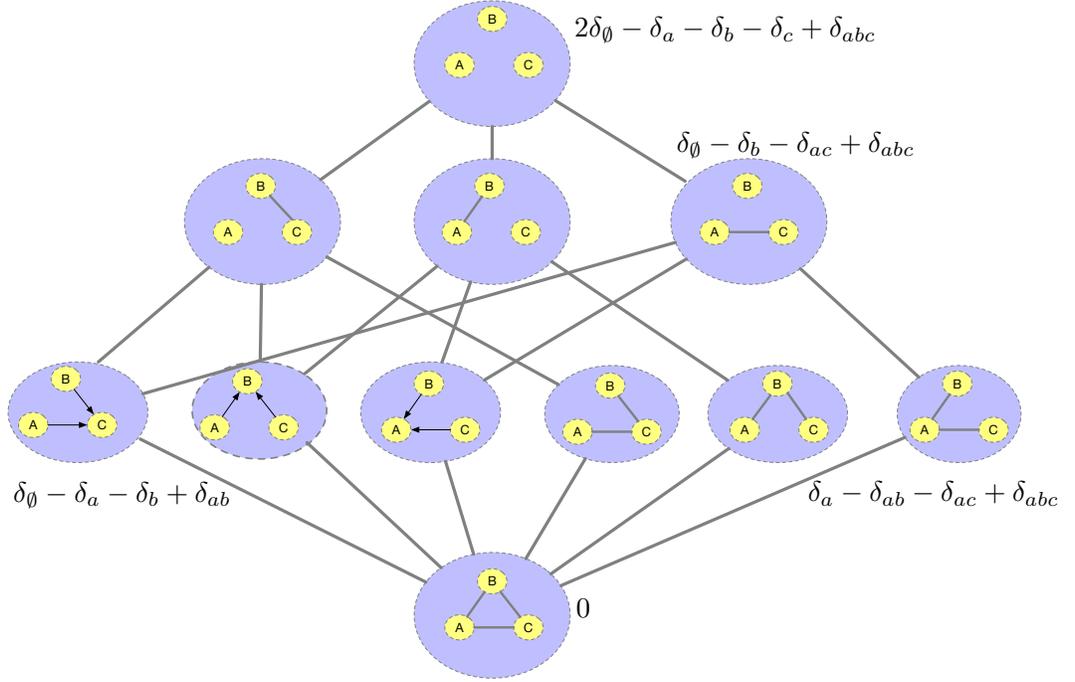


Figure 8 shows the space of possible imset representations over 3 variables, along with the accompanying causal DAG shown in its reduced “essential graph” form Andersson et al. (1997). Each vertex in this figure shows a potential DAG model, along with its imset representation. We can now see how each of these DAG models can be viewed in terms of a simplicial object using functors that maps from the category Δ into this lattice of possible DAGs over 3 variables. Note that the model at the very top can be essentially viewed as degenerate simplicial object, where all its arrows have been discarded. Similarly, the collider model $A \rightarrow C \leftarrow B$, whose integer-valued multiset representation is defined as $\delta_\emptyset - \delta_a - \delta_b + \delta_{ab}$, can be viewed as simplicial outer horn object Λ_2^2 . Similarly, the serial DAG model $A \rightarrow B \rightarrow C$ can be viewed as the simplicial inner horn object Λ_1^2 .

Figure 8: Causal discovery through the space of all models over 3 variables shown with their associated inset representations (Studený et al., 2010). Each candidate DAG defines a *causal horn*, a simplicial subobject of the complete simplex on $\Delta[2]$, and the process of causal structure discovery can be viewed in terms of the abstract horn filling problem defined above for higher-order categories. Each DAG is shown in terms of its “essential graph”, a reduced representation defined in Andersson et al. (1997), where an undirected edge is used to denote the reversibility of a directed edge.



3.7 Example: Simplicial Objects over String Diagrams

We now illustrate the above formalism of simplicial objects by illustrating how it applies to the special case where causal models are defined over symmetric monoidal categories Fong (2012); Jacobs et al. (2018). For a detailed overview of symmetric monoidal categories, we recommend the book-length treatment by Fong and Spivak (2018). We will restrict ourselves to showing how the UCLA framework provides a way to generalize the past work on string diagrams for causal inference. Symmetric monoidal categories (SMCs) are useful in modeling processes where objects can be combined together to give rise to new objects, or where objects disappear. For example, Coecke et al. (2016) propose a mathematical framework for resources based on SMCs. It is important to point out that monoidal categories can be defined as a special type of Grothendieck fibration Richter (2020), thereby formally showing that the structure of a monoidal category is a special case of our UCLA hierarchy. We will discuss Grothendieck fibrations in more detail in the next section. We leave aside the details of this construction, but refer the reader to Richter (2020) for the details.

We focus on the work of Jacobs et al. (2018), and use the example shown in Figure 9, where we have only shown the top two levels of the UCLA hierarchy. Our goal in this section is to illustrate how we can define simplicial objects over the SMC category CDU category \mathbf{Syn}_G constructed by Jacobs et al. (2018) to mimic the process of working with an actual Bayesian network DAG G . An example of such an encoding is shown in Figure 9. For the purposes of our illustration, it is not important to discuss the intricacies involved in this model, for which we refer the reader to the original paper. Our goal is to show that by encapsulating their SMC category in the UCLA framework, we can extend their approach as described below. In particular, we can solve an associated lifting problem that is defined by the functor mapping the simplicial category Δ to their SMC category. They use the category of stochastic matrices to capture the process of working with the joint distribution as shown in the figure. Instead, we show that one can use some other category, such as the category of **Sets**, or **Top** (the category of topological spaces), or indeed, the category **Meas** of measurable spaces.

Recall that Bayesian networks Pearl (1989) define a joint probability distribution

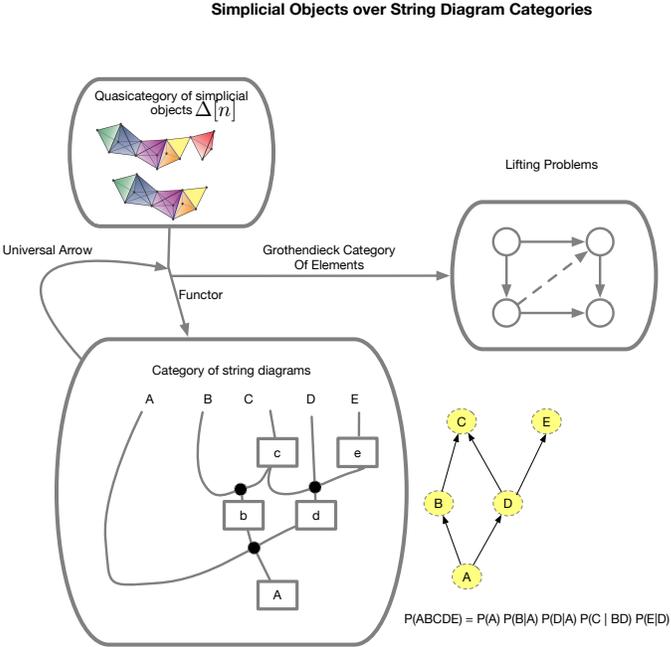


Figure 9: Causal models over symmetric monoidal categories were explored by Jacobs et al. (2018).

$$P(X_1, \dots, X_n) = \prod_{i=1}^n P(X_i | \text{Pa}(X_i)),$$

where $\text{Pa}(X_i) \subset \{X_1, \dots, X_n\} \setminus X_i$ represents a subset of variables (not including the variable itself). Jacobs et al. (2018) show Bayesian network models can be constructed using symmetric monoidal categories, where the tensor product operation is used to combine multiple variables into a “tensorized” variable that then probabilistically maps into an output variable. In particular, the monoidal category **Stoch** has as objects finite sets, and morphisms $f : A \rightarrow B$ are $|B| \times |A|$ dimensional stochastic matrices. Composition of stochastic matrices corresponds to matrix multiplication. The monoidal product \otimes in **Stoch** is the cartesian product of objects, and the Kronecker product of matrices $f \otimes g$. Jacobs et al. (2018) define three additional operations, the copy map, the discarding map, and the uniform state.

Definition 20. A CDU category (for copy, discard, and uniform) is a SMC category (\mathbf{C}, \otimes, I) , where each object A has a copy map $C_A : A \rightarrow A \otimes A$, and discarding map $D_A : A \rightarrow I$, and a uniform state map $U_A : I \rightarrow A$, satisfying a set of equations detailed in Jacobs et al. (2018). CDU functors are symmetric monoidal functors between CDU categories, preserving the CDU maps.

The key theorem we are interested in is the following from the original paper Jacobs et al. (2018):

Theorem 1. There is an isomorphism (1-1 correspondence) between Bayesian networks based on a DAG G and CDU functors $F : \mathbf{Syn}_G \rightarrow \mathbf{Stoch}$.

The significance of this theorem for the UCLA architecture is that it shows how the SMC category of CDU objects can be defined as Layer 2 of the UCLA hierarchy, whereas the category **Stoch** can be viewed as instantiating the Layer 3 of the UCLA hierarchy.

3.8 Nerve of a Category

An important concept that will play a key role in Layer 4 of the UCLA hierarchy is that of the *nerve* of a category Lurie (2022); Richter (2020). The nerve of a category \mathcal{C} enables embedding \mathcal{C} into the category of simplicial objects, which is a fully faithful embedding.

Definition 21. Let $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ be a functor from category \mathcal{C} to category \mathcal{D} . If for all arrows f the mapping $f \rightarrow \mathcal{F}f$

- injective, then the functor \mathcal{F} is defined to be **faithful**.

- surjective, then the functor \mathcal{F} is defined to be **full**.
- bijective, then the functor \mathcal{F} is defined to be **fully faithful**.

Definition 22. The **nerve** of a category \mathcal{C} is the set of composable morphisms of length n , for $n \geq 1$. Let $N_n(\mathcal{C})$ denote the set of sequences of composable morphisms of length n .

$$\{C_0 \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots \xrightarrow{f_n} C_n \mid C_i \text{ is an object in } \mathcal{C}, f_i \text{ is a morphism in } \mathcal{C}\}$$

The set of n -tuples of composable arrows in \mathcal{C} , denoted by $N_n(\mathcal{C})$, can be viewed as a functor from the simplicial object $[n]$ to \mathcal{C} . Note that any nondecreasing map $\alpha : [m] \rightarrow [n]$ determines a map of sets $N_m(\mathcal{C}) \rightarrow N_n(\mathcal{C})$. The nerve of a category \mathcal{C} is the simplicial set $N_\bullet : \Delta \rightarrow \mathbf{Set}$, which maps the ordinal number object $[n]$ to the set $N_n(\mathcal{C})$.

The importance of the nerve of a category comes from a key result Lurie (2022), showing it defines a full and faithful embedding of a category:

Theorem 2. (Lurie, 2022, Tag 002Y): The **nerve functor** $N_\bullet : \mathbf{Cat} \rightarrow \mathbf{Set}$ is fully faithful. More specifically, there is a bijection θ defined as:

$$\theta : \mathbf{Cat}(\mathcal{C}, \mathcal{C}') \rightarrow \mathbf{Set}_\Delta(N_\bullet(\mathcal{C}), N_\bullet(\mathcal{C}'))$$

Using this concept of a nerve of a category, we can now state a theorem that shows it is possible to easily embed the CDU symmetric monoidal category defined above that represents Bayesian Networks and their associated ‘‘string diagram surgery’’ operations for causal inference as a simplicial set.

Theorem 3. Define the **nerve** of the CDU symmetric monoidal category (\mathbf{C}, \otimes, I) , where each object A has a copy map $C_A : A \rightarrow A \otimes A$, and discarding map $D_A : A \rightarrow I$, and a uniform state map $U_A : I \rightarrow A$ as the set of composable morphisms of length n , for $n \geq 1$. Let $N_n(\mathcal{C})$ denote the set of sequences of composable morphisms of length n .

$$\{C_0 \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots \xrightarrow{f_n} C_n \mid C_i \text{ is an object in } \mathcal{C}, f_i \text{ is a morphism in } \mathcal{C}\}$$

The associated **nerve functor** $N_\bullet : \mathbf{Cat} \rightarrow \mathbf{Set}$ from the CDU category is fully faithful. More specifically, there is a bijection θ defined as:

$$\theta : \mathbf{Cat}(\mathcal{C}, \mathcal{C}') \rightarrow \mathbf{Set}_\Delta(N_\bullet(\mathcal{C}), N_\bullet(\mathcal{C}'))$$

This theorem is just a special case of the above theorem attesting to the full and faithful embedding of any category using its nerve, which then makes it a simplicial set. We can then use the theoretical machinery at the top layer of the UCLA architecture to manipulate causal interventions in this category using face and degeneracy operators as defined above.

Note that the functor G from a simplicial object X to a category \mathcal{C} can be lossy. For example, we can define the objects of \mathcal{C} to be the elements of X_0 , and the morphisms of \mathcal{C} as the elements $f \in X_1$, where $f : a \rightarrow b$, and $d_0 f = a$, and $d_1 f = b$, and $s_0 a, a \in X$ as defining the identity morphisms $\mathbf{1}_a$. Composition in this case can be defined as the free algebra defined over elements of X_1 , subject to the constraints given by elements of X_2 . For example, if $x \in X_2$, we can impose the requirement that $d_1 x = d_0 x \circ d_2 x$. Such a definition of the left adjoint would be quite lossy because it only preserves the structure of the simplicial object X up to the 2-simplices. The right adjoint from a category to its associated simplicial object, in contrast, constructs a full and faithful embedding of a category into a simplicial set. In particular, the nerve of a category is such a right adjoint.

As this has been a rather long section, it is time to summarize and review what we have learned. We defined the top layer of the UCLA architecture as a category over ordinal numbers Δ , and showed that the accompanying morphisms can be used to implement causal interventions using degeneracy and face operators. We illustrated using two examples – the integer-valued multisets investigated by Studeny (2010), as well as the work on string diagram surgery by Jacobs et al. (2018) – to help make concrete how the abstract machinery of simplicial objects and sets can be used to implement causal ‘‘surgery’’ operators. We also discussed how categories can be faithfully embedded into simplicial sets using the nerve of the category, the set of sequences of length n of composable morphisms. Next, we want to proceed downwards to explore the interaction between the causal layer and the layer of instances (or sets).

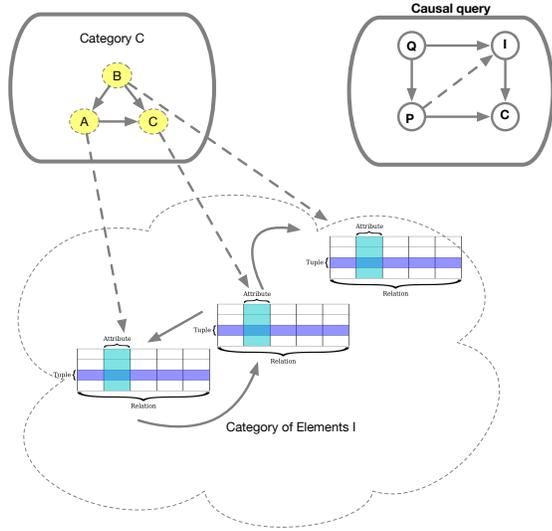


Figure 10: Causal models can be viewed as defining a relational database over instances, where the database scheme defines a category. Causal queries can then be viewed analogously to database queries, and be formulated as solutions to lifting problems Spivak (2013).

4 Layers 2 and 3 of UCLA: The Category of Elements in Causal Inference

Next, we turn to describe the second (from top) and third layers of the UCLA architecture, which pertain to the category of causal models (for example, a graph or a symmetric monoidal category), and the database of instances that support causal inferences. Figure 10 illustrates the mapping from a causal category to a dataset of instances, modeled as a category **Set**. Drawing on the close correspondences between categories and relational database schemes (see Spivak (2013) for details), we can view causal queries over data as analogous to database queries, which can then be formulated by corresponding lifting problems. In this figure, the lifting problem with respect to the category I of elements, which combines the causal model with the instances that it maps into. That is, each object in the model, e.g. a variable indicating a patient, maps into actual patients, and a variable indicating outcomes from Covid-19 exposure, maps into actual outcomes for that individual. The causal arrow from the patient variable into the exposure variable then maps into actual arrows for each patient. Causal queries of exposure to Covid-19 then become similar to database queries. In the next section, we will generalize this perspective, showing that we can map into a topological category and answer more abstract questions relating to the geometry of a dataset, or map into a category of measurable spaces to answer probabilistic queries. The structure of the lifting problem remains the same, what changes are the specifics of the underlying categories.

A central and unifying principle in UC and the UCLA architecture is that every pair of categorical layers is synchronized by a functor, along with a universal arrow. We explore the universal arrow property more deeply in this section, showing how it provides the conceptual basis behind the Yoneda Lemma, and Grothendieck’s category of elements. In the case of causal inference, universal arrows enable mimicking the effects of causal operations from one layer of the UCLA hierarchy down to the next layer. In particular, at the simplicial object layer, we can model a causal intervention in terms of face and degeneracy operators (defined below in more detail). These in turn correspond to “graph surgery” Pearl (2009) operations on causal DAGs, or in terms of “copy”, “delete” operators in “string diagram surgery” of causal models defined on symmetric monoidal categories Jacobs et al. (2018). These “surgery” operations at the next level may translate down to operations on probability distributions, measurable spaces, topological spaces, or chain complexes. This process follows a standard construction used widely in mathematics, for example group representations associate with any group G , a left \mathbf{k} -module M representation that enables modeling abstract group operations by operations on the associated modular representation. These concrete representations must satisfy the universal arrow property for them to be faithful.

A special case of the universal arrow property is that of universal element, which as we will see below plays an important role in the UCLA architecture in defining a suitably augmented category of elements, based on a construction introduced by Grothendieck.

Definition 23. If D is a category and $H : D \rightarrow \mathbf{Set}$ is a set-valued functor, a **universal element** associated with the functor H is a pair $\langle r, e \rangle$ consisting of an object $r \in D$ and an element $e \in Hr$ such that for every pair $\langle d, x \rangle$ with $x \in Hd$, there is a unique arrow $f : r \rightarrow d$ of D such that $(Hf)e = x$.

Example 12. Let E be an equivalence relation on a set S , and consider the quotient set S/E of equivalence classes, where $p : S \rightarrow S/E$ sends each element $s \in S$ into its corresponding equivalence class. The set of equivalence classes S/E has the property that any function $f : S \rightarrow X$ that respects the equivalence relation can be written as $fs = fs'$ whenever $s \sim_E s'$, that is, $f = f' \circ p$, where the unique function $f' : S/E \rightarrow X$. Thus, $\langle S/E, p \rangle$ is a universal element for the functor H .

4.1 The Category of Elements

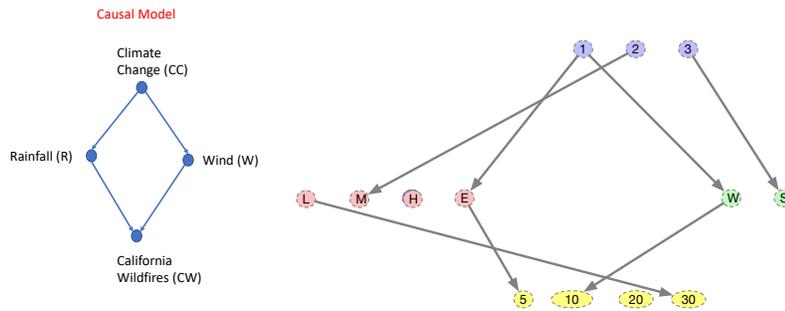
We turn next to define the category of elements, based on a construction by Grothendieck, and illustrate how it can serve as the basis for inference at each layer of the UCLA architecture. In particular, Spivak (2013) shows how the category of elements can be used to define SQL queries in a relational database.

Definition 24. Given a set-valued functor $\delta : \mathcal{C} \rightarrow \mathbf{Set}$ from some category \mathcal{C} , the induced **category of elements** associated with δ is a pair $(\int \delta, \pi_\delta)$, where $\int \delta \in \mathbf{Cat}$ is a category in the category of all categories \mathbf{Cat} , and $\pi_\delta : \int \delta \rightarrow \mathcal{C}$ is a functor that “projects” the category of elements into the corresponding original category \mathcal{C} . The objects and arrows of $\int \delta$ are defined as follows:

- $\text{Ob}(\int \delta) = \{(s, x) | x \in \text{Ob}(\delta s), x \in \delta s\}$.
- $\text{Hom}_{\int \delta}((s, x), (s', x')) = \{f : s \rightarrow s' | \delta f(x) = x'\}$

Example 13. To illustrate the category of elements construction, let us consider the toy climate change causal model shown in Figure 11. Let the category \mathcal{C} be defined by this causal DAG model, where the objects $\text{Ob}(\mathcal{C})$ are defined by the four vertices, and the arrows $\text{Hom}_{\mathcal{C}}$ are defined by the four edges in the model. The set-valued functor $\delta : \mathcal{C} \rightarrow \mathbf{Set}$ maps each object (vertex) in \mathcal{C} to a set of instances, thereby turning the causal DAG model into an associated set of tables. For example, **Climate Change** is defined as a table of values, which could be modeled as a multinomial variable taking on a set of discrete values, and for each of its values, the arrow from **Climate Change** to **Rainfall** maps each specific value of **Climate Change** to a value of **Rainfall**, thereby indicating a causal effect of climate change on the amount of rainfall in California. In the figure, **Climate Change** is mapped to three discrete levels (marked 1, 2 and 3). Rainfall amounts are discretized as well into low (marked "L"), medium (marked "M"), high (marked "H"), or extreme (marked "E"). Wind speeds are binned into two levels (marked "W" for weak, and "S" for strong). Finally, the percentage of California wildfires is binned between 5 to 30. Not all arrows that exist in the Grothendieck category of elements are shown, for clarity.

Figure 11: A toy causal DAG model of climate change to illustrate the category of elements construction.



Many properties of Grothendieck’s construction can be exploited (some of these are discussed in the context of relational database queries in Spivak (2013)), but for our application to causal inference, we are primarily interested in the associated class of lifting problems that define queries in a causal model. In particular, it can be shown that Grothendieck’s category of elements formally defines a type of fibration called a *Grothendieck opfibration*. More specifically, if we consider any object $c \in \mathcal{C}$ in a category \mathcal{C} , and look at its fiber $\pi_\delta^{-1}(c)$, which forms a subcategory of the category of elements I , it can be shown that the projection of every morphism in the fiber into the category \mathcal{C} is to the identity element 1_c . Richter (2020) explains how monoidal categories themselves can be defined using simplicial objects using a Grothendieck opfibration. We refer the reader to these sources for additional details. One consequence of this property is that approaches to UC that build symmetric monoidal category representations of Bayesian networks Fong (2012); Jacobs et al. (2018) can be formally shown to define special cases of the UCLA hierarchy.

Definition 25. If S is a collection of morphisms in category \mathcal{C} , a morphism $f : A \rightarrow B$ has the **left lifting property with respect to S** if it has the left lifting property with respect to every morphism in S . Analogously, we say a morphism $p : X \rightarrow Y$ has the **right lifting property with respect to S** if it has the right lifting property with respect to every morphism in S .

4.2 Lifting Problems in Causal Inference

We now turn to sketch some examples of the application of lifting problems for causal inference. Many problems in causal inference on graphs involve some particular graph property. To formulate it as a lifting problem, we will use the following generic template, following the initial application of lifting problems to database queries proposed by Spivak (2013).

$$\begin{array}{ccc} Q & \xrightarrow{\mu} & \int \delta \\ \downarrow f & \nearrow h & \downarrow p \\ R & \xrightarrow{\nu} & \mathcal{C} \end{array}$$

Here, Q is a generic query that we want answered, which could range from a database query, as in the original setting studied by Spivak (2013), but more interestingly, it could be a particular graph property relating to causal inference (as illustrated by the following two examples), but as we will show later, it could also be related to the combinatorial category of simplicial objects used to model causal intervention, and finally, it could also be related to questions relating to the evaluation of causal models using a measure-theoretic or probability space. By suitably modifying the base category, the lifting problem formulation can be used to encode a diverse variety of problems in causal inference. R represents a fragment of the complete causal model \mathcal{C} , and δ is the category of elements defined above. Finally, h gives all solutions to the lifting problem. Some examples will help clarify this concept.

Example 14. Consider the category of directed graphs defined by the category \mathcal{G} , where $\text{Ob}(\mathcal{G}) = \{V, E\}$, and the morphisms of \mathcal{G} are given as $\text{Hom}_{\mathcal{G}} = \{s, t\}$, where $s : E \rightarrow V$ and $t : E \rightarrow V$ define the source and terminal nodes of each vertex. Then, the category of all directed graphs is precisely defined by the category of all functors $\delta : \mathcal{G} \rightarrow \text{Set}$. Any particular graph is defined by the functor $X : \mathcal{G} \rightarrow \text{Set}$, where the function $X(s) : X(E) \rightarrow X(V)$ assigns to every edge its source vertex. For causal inference, we may want to check some property of a graph, such as the property that every vertex in X is the source of some edge. The following lifting problem ensures that every vertex has a source edge in the graph. The category of elements $\int \delta$ shown below refers to a construction introduced by Grothendieck, which will be defined in more detail later.

$$\begin{array}{ccc} V(\bullet) & \xrightarrow{\mu} & \int \delta \\ \downarrow f & \nearrow h & \downarrow p \\ \{E(\bullet) \xrightarrow{s} V(\bullet)\} & \xrightarrow{\nu} & \mathcal{G} \end{array}$$

Example 15. As another example of the application of lifting problems to causal inference, let us consider the problem of determining whether two causal DAGs, G_1 and G_2 are Markov equivalent Andersson et al. (1997). A key requirement here is that the immoralities of G_1 and G_2 must be the same, that is, if G_1 has a collider $A \rightarrow B \leftarrow C$, where there is no edge between A and C , then G_2 must also have the same collider, and none others. We can formulate the problem of finding colliders as the following lifting problem. Note that the three vertices A , B and C are bound to an actual graph instance through the category of elements $\int \delta$ (as was illustrated above), using the top right morphism μ . The bottom left morphism f binds these three vertices to some collider. The bottom right morphism ν requires this collider to exist in the causal graph \mathcal{G} with the same bindings as found by μ . The dashed morphism h finds all solutions to this lifting problem, that is, all colliders involving the vertices A , B and C .

$$\begin{array}{ccc} \{A(\bullet), B(\bullet), C(\bullet)\} & \xrightarrow{\mu} & \int \delta \\ \downarrow f & \nearrow h & \downarrow p \\ \{A(\bullet) \rightarrow B(\bullet) \leftarrow C(\bullet)\} & \xrightarrow{\nu} & \mathcal{G} \end{array}$$

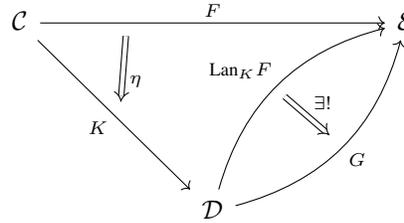
If the category of elements is defined by a functor mapping a database schema into a table of instances, then the associated lifting problem corresponds to familiar problems like SQL queries in relational databases Spivak (2013). In our application, we can use the same machinery to formulate causal inference queries by choosing the categories

appropriately. To complete the discussion, we now make the connection between universal arrows and the core notion of universal representations via the Yoneda Lemma.

4.3 Modeling Causal Interventions as Kan Extension

It is well known in category theory that ultimately every concept, from products and co-products, limits and co-limits, and ultimately even the Yoneda Lemma (see below), can be derived as special cases of the Kan extension (MacLane, 1971). Kan extensions intuitively are a way to approximate a functor \mathcal{F} so that its domain can be extended from a category \mathcal{C} to another category \mathcal{D} . Because it may be impossible to make commutativity work in general, Kan extensions rely on natural transformations to make the extension be the best possible approximation to \mathcal{F} along \mathcal{K} . We want to briefly show Kan extensions can be combined with the category of elements defined above to construct causal “migration functors” that map from one causal model into another. These migration functors were originally defined in the context of database migration Spivak (2013), and here we are adapting that approach to causal inference. By suitably modifying the category of elements from a set-valued functor $\delta : \mathcal{C} \rightarrow \mathbf{Set}$, to some other category, such as the category of topological spaces, namely $\delta : \mathcal{C} \rightarrow \mathbf{Top}$, we can extend the causal migration functors into solving more abstract causal inference questions. We explore the use of such constructions in the next section on Layer 4 of the UCLA hierarchy. Here, for simplicity, we restrict our focus to Kan extensions for migration functors over the category of elements defined over instances of a causal model.

Definition 26. A **left Kan extension** of a functor $F : \mathcal{C} \rightarrow \mathcal{E}$ along another functor $K : \mathcal{C} \rightarrow \mathcal{D}$, is a functor $\text{Lan}_K F : \mathcal{D} \rightarrow \mathcal{E}$ with a natural transformation $\eta : F \rightarrow \text{Lan}_K F \circ K$ such that for any other such pair $(G : \mathcal{D} \rightarrow \mathcal{E}, \gamma : F \rightarrow GK)$, γ factors uniquely through η . In other words, there is a unique natural transformation $\alpha : \text{Lan}_K F \Rightarrow G$.



A **right Kan extension** can be defined similarly. To understand the significance of Kan extensions for causal inference, we note that under a causal intervention, when a causal category S gets modified to T , evaluating the modified causal model over a database of instances can be viewed as an example of Kan extension.

First, we need to review the basic concept of adjoint functors, which will be helpful in seeing how to use Kan extensions to model causal interventions.

Definition 27. A pair of **adjoint functors** is defined as $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$, where F is considered the right adjoint, and G is considered the left adjoint,

$$\mathcal{D} \begin{array}{c} \xrightarrow{G} \\ \xleftarrow{F} \end{array} \mathcal{C}.$$

must satisfy the property that for each pair of objects C of \mathcal{C} and D of \mathcal{D} , there is a bijection of sets

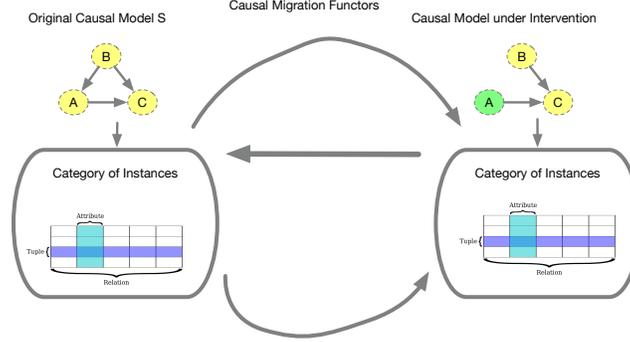
$$\phi_{C,D} : \mathbf{Hom}_{\mathcal{C}}(C, G(D)) \simeq \mathbf{Hom}_{\mathcal{D}}(F(C), D)$$

Notice the similarity of this definition to the one earlier where the universal arrow property induced a bijection of **Hom** sets that then led to universal elements, Grothendieck category of elements, and the Yoneda Lemma.

Let $\delta : S \rightarrow \mathbf{Set}$ denote the original causal model defined by the category S with respect to some dataset. Let $\epsilon : T \rightarrow \mathbf{Set}$ denote the effect of a causal intervention abstractly defined as some change in the category S to T , such as deletion of an edge, as illustrated in Figure 12. Intuitively, we can consider three cases: the *pullback* Δ_F along F , which maps the effect of a causal intervention back to the original model, the *left pushforward* Σ_F and the *right pushforward* \amalg_F , which can be seen as adjoints to the pullback Δ_F .

Following Spivak (2013), we can define three *causal migration functors* that evaluate the impact of a causal intervention with respect to a dataset of instances.

Figure 12: Kan extensions are useful in modeling the effects of a causal intervention.



1. The functor $\Delta_F : \epsilon \rightarrow \delta$ sends the functor $\epsilon : T \rightarrow \mathbf{Set}$ to the composed functor $\delta \circ F : S \rightarrow \mathbf{Set}$.
2. The functor $\Sigma_F : \delta \rightarrow \epsilon$ is the left Kan extension along F , and can be seen as the left adjoint to Δ_F .
The functor $\prod_F : \delta \rightarrow \epsilon$ is the right Kan extension along F , and can be seen as the right adjoint to Δ_F .

To understand how to implement these functors, we use the following proposition that is stated in Spivak (2013) in the context of database queries, which we are restating in the setting of causal inference.

Theorem 4. Let $F : S \rightarrow T$ be a functor. Let $\delta : S \rightarrow \mathbf{Set}$ and $\epsilon : T \rightarrow \mathbf{Set}$ be two set-valued functors, which can be viewed as two instances of a causal model defined by the category S and T . If we view T as the causal category that results from a causal intervention on S (e.g., deletion of an edge), then there is a commutative diagram linking the category of elements between S and T .

$$\begin{array}{ccc}
 \int \delta & \longrightarrow & \int \epsilon \\
 \downarrow \pi_\delta & & \downarrow \pi_\epsilon \\
 S & \xrightarrow{F} & T
 \end{array}$$

Proof: To check that the above diagram is a pullback, that is, $\int \delta \simeq S \times_T \int \delta$, or in words, the fiber product, we can check the existence of the pullback component wise by comparing the set of objects and the set of morphisms in $\int \delta$ with the respective sets in $S \times_T \int \epsilon$. •

For simplicity, we defined the migration functors above with respect to an actual dataset of instances. More generally, we can compose the set-valued functor $\delta : S \rightarrow \mathbf{Set}$ with a functor $\mathcal{T} : \mathbf{Set} \rightarrow \mathbf{Top}$ to the category of topological spaces to derive a Kan extension formulation of the definition of a causal intervention. We discuss this issue in the next section on causal homotopy.

4.4 Yoneda Lemma

The Yoneda Lemma plays a crucial role in UC because it defines the concept of a representation in category theory. We first show that associated with universal arrows is the corresponding induced isomorphisms between **Hom** sets of morphisms in categories. This universal property then leads to the Yoneda Lemma.

Theorem 5. Given any functor $S : D \rightarrow C$, the universal arrow $\langle r, u : c \rightarrow Sr \rangle$ implies a bijection exists between the **Hom** sets

$$\mathbf{Hom}_D(r, d) \simeq \mathbf{Hom}_C(c, Sd)$$

This is a well-known result whose proof can be found in MacLane (1971). The crucial point here can be illustrated with the help of a commutative diagram, showing that the bijection is defined by a natural transformation ϕ , natural in the object d in category D . By investigating a special case of this natural transformation for how it transforms the identity morphism $\mathbf{1}_r$, leads us to the Yoneda Lemma.

$$\begin{array}{ccc}
 D(r, r) & \xrightarrow{\phi_r} & C(c, Sr) \\
 \downarrow D(r, f') & & \downarrow C(c, Sf') \\
 D(r, d) & \xrightarrow{\phi_d} & C(c, Sd)
 \end{array}$$

Note here that the identity morphism $\mathbf{1}_r \in D(r, r)$ is mapped to $Sf' \circ u$ through the top and right path, and to $\phi_d(f')$ through the left and bottom path. As these paths must be equal in a commutative diagram, we get the ensuring property that a bijection between the **Hom** sets holds precisely when $\langle r, u : c \rightarrow Sr \rangle$ is a universal arrow from c to S . Note that for the case when the categories C and D are small, meaning their **Hom** collection of arrows forms a set, the induced functor $\mathbf{Hom}_C(c, S-)$ to **Set** is isomorphic to the functor $\mathbf{Hom}_D(r, -)$. This type of isomorphism defines a (universal) representation.

Lemma 1. Yoneda Lemma): If $H : D \rightarrow \mathbf{Set}$ is a set-valued functor, and r is an object in D , there is a bijection that sends each natural transformation $\alpha : \mathbf{Hom}_D(r, -) \rightarrow K$ to $\alpha_r \mathbf{1}_r$, the image of the identity morphism $\mathbf{1}_r : r \rightarrow r$.

$$y : \mathbf{Nat}(\mathbf{Hom}_D(r, -), K) \simeq Kr$$

The proof of the Yoneda Lemma follows directly from the below commutative diagram, a special case of the above diagram for universal arrows.

$$\begin{array}{ccc}
 D(r, r) & \xrightarrow{\phi_r} & Kr \\
 \downarrow D(r, f') & & \downarrow C(c, Sf') \\
 D(r, d) & \xrightarrow{\phi_d} & Kd
 \end{array}$$

4.5 The Universality of Diagrams in Causal Inference

Diagrams play a key role in defining UC and the UCLA architecture, as has already become clear from the discussion above. We briefly want to emphasize the central role played by universal constructions involving limits and colimits of diagrams, which are viewed as functors from an indexing category of diagrams to a category. To make this somewhat abstract definition concrete, let us look at some simpler examples of universal properties, including co-products and quotients (which in set theory correspond to disjoint unions). Coproducts refer to the universal property of abstracting a group of elements into a larger one.

$$\begin{array}{ccc}
 Z & \xrightarrow{p} & X \\
 \downarrow q & & \downarrow f \\
 Y & \xrightarrow{g} & X \sqcup Y \\
 & \searrow i & \downarrow r \\
 & & R
 \end{array}$$

In the commutative diagram above, the coproduct object $X \sqcup Y$ uniquely factorizes any mapping $h : X \rightarrow R$ and any mapping $i : Y \rightarrow R$, so that $h = r \circ f$, and furthermore $i = r \circ g$. Co-products are themselves special cases of the more general notion of co-limits.

4.5.1 Pullback Mappings

Figure 13 illustrates the fundamental property of a *pullback*, which along with *pushforward*, is one of the core ideas in category theory. The pullback square with the objects U, X, Y and Z implies that the composite mappings $g \circ f'$ must equal $g' \circ f$. In this example, the morphisms f and g represent a *pullback* pair, as they share a common co-domain Z . The pair of morphisms f', g' emanating from U define a *cone*, because the pullback square “commutes” appropriately. Thus, the pullback of the pair of morphisms f, g with the common co-domain Z is the pair of morphisms f', g' with common domain U . Furthermore, to satisfy the universal property, given another pair of morphisms x, y with common domain T , there must exist another morphism $k : T \rightarrow U$ that “factorizes” x, y appropriately, so that the composite morphisms $f' k = y$ and $g' k = x$. Here, T and U are referred to as *cones*, where U is the limit of the set of all cones “above” Z . If we reverse arrow directions appropriately, we get the corresponding notion of pushforward. So, in this example, the pair of morphisms f', g' that share a common domain represent a pushforward pair. As Figure 13, for

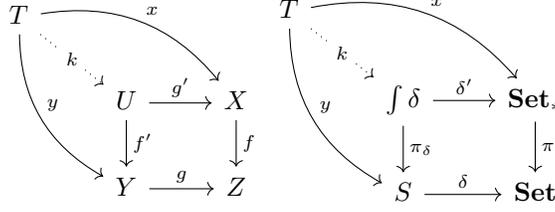


Figure 13: Left: Universal Property of pullback mappings. Right: The Grothendieck category of elements $\int \delta$ of any set-valued functor $\delta : S \rightarrow \mathbf{Set}$ can be described as a pullback in the diagram of categories.

any set-valued functor $\delta : S \rightarrow \mathbf{Sets}$, the Grothendieck category of elements $\int \delta$ can be shown to be a pullback in the diagram of categories. Here, \mathbf{Set}_* is the category of pointed sets, and π is a projection that sends a pointed set $(X, x \in X)$ to its underlying set X .

To state the next two results, we need to introduce formally the concept of limit and colimits MacLane (1971).

Definition 28. Given a functor $F : \mathcal{J} \rightarrow \mathcal{C}$ from an indexing diagram category \mathcal{J} to a category \mathcal{C} , an element A from the set of natural transformations $N(A, F)$ is called a **cone**. A **limit** of the diagram $F : \mathcal{J} \rightarrow \mathcal{C}$ is a cone η from an object $\lim F$ to the diagram satisfying the universal property that for any other cone γ from an object B to the diagram, there is a unique morphism $h : B \rightarrow \lim F$ so that $\gamma_\bullet = \eta_\bullet \circ h$ for all objects \bullet in \mathcal{J} . Dually, the **co-limit** of the diagram $F : \mathcal{J} \rightarrow \mathcal{C}$ is a cone ϵ satisfying the universal property that for any other cone γ from the diagram to the object B , there is a unique mapping $h : \text{colim} F \rightarrow B$ so that $\gamma_\bullet = h \epsilon_\bullet$ for all objects \bullet in \mathcal{J} .

A fundamental consequence of the category of elements is that every object in the functor category of presheaves, namely contravariant functors from a category into the category of sets, is the colimit of a diagram of representable objects, via the Yoneda Lemma. This is a standard result whose proof can be found in MacLane and leke Moerdijk (1994), and will be of importance below.

Theorem 6. In the functor category of presheaves $\mathbf{Set}^{C^{op}}$, every object P is the colimit of a diagram of representable objects, in a canonical way.

In UC, any causal influence of an object X upon any other object Y can be represented as a natural transformation (a morphism) between two functor objects in the presheaf category $\hat{\mathcal{C}}$. The CRP is very akin to the idea of the reproducing property in kernel methods. Reproducing Kernel Hilbert Spaces (RKHS's) transformed the study of machine learning, precisely because they are the unique subcategory in the category of all Hilbert spaces that have representers of evaluation defined by a kernel matrix $K(x, y)$ Schölkopf and Smola (2002). The reproducing property in an RKHS is defined as $\langle K(x, -), K(-, y) \rangle = K(x, y)$. An analogous but far more general reproducing property holds in the UC framework, based on the Yoneda Lemma.

Theorem 7. Causal Reproducing Property: All causal influences between any two objects X and Y can be derived from its presheaf functor objects, namely

$$\mathbf{Hom}_{\mathcal{C}}(X, Y) \simeq \mathbf{Nat}(\mathbf{Hom}_{\mathcal{C}}(-, X), \mathbf{Hom}_{\mathcal{C}}(-, Y))$$

Proof: The proof of this theorem is a direct consequence of the Yoneda Lemma, which states that for every presheaf functor object F in $\hat{\mathcal{C}}$ of a category \mathcal{C} , $\mathbf{Nat}(\mathbf{Hom}_{\mathcal{C}}(-, X), F) \simeq FX$. That is, elements of the set FX are in 1 – 1 bijections with natural transformations from the presheaf $\mathbf{Hom}_{\mathcal{C}}(-, X)$ to F . For the special case where the functor object $F = \mathbf{Hom}_{\mathcal{C}}(-, Y)$, we get the result immediately that $\mathbf{Hom}_{\mathcal{C}}(X, Y) \simeq \mathbf{Nat}(\mathbf{Hom}_{\mathcal{C}}(-, X), \mathbf{Hom}_{\mathcal{C}}(-, Y))$. •

The significance of the Causal Reproducing Property is that presheaves act as “representers” of causal information, precisely analogous to how kernel matrices act as representers in an RKHS.

5 Layer 4 of UCLA: Evaluating Causal Effects using Homotopy Colimits

Finally, we turn to discuss the role of the causal homotopy layer. Before delving into the details of this category, we pause to review what is meant by “causal effect”. Previous work in causal inference has adopted a wide variety of definitions of causality. For example, in the work on potential outcomes Imbens and Rubin (2015), a causal effect is defined through the average treatment effect (ATE) estimator, where the mean of the outcomes $E(Y_1)$ under treatment is compared with the mean of the outcomes of the outcomes under control $E(Y_0)$. In the work on DAG-based causal

inference (Pearl, 2009), causal effects are defined as a change in the distribution of values of the treated variable $P(Y|X = 1)$ compared to the untreated variable $P(Y|X = 0)$. Janzing et al. (2013) propose an information-theoretic approach to quantifying causal influence, where the causal effect is defined in terms of the KL-divergence between the observational $P(Y|X = x)$ and interventional distributions $P(Y|do(X = x))$. Work on reproducing kernel Hilbert space embeddings of causal models, such as mean counterfactual embeddings Muandet et al. (2021), propose evaluating the ATE estimator in the induced RKHS, so that in the original data space, the estimator is defined with respect to a reproducing kernel.

Since our approach to causality is based on category theory, we prefer to derive an estimator that is entirely based on the structure of categories, namely objects and morphisms. In particular, we defined the nerve of the category above as a full and faithful embedding of a category as a simplicial object. Associated with the nerve is an important topological invariant of a category called its *classifying space* Richter (2020), which provides a way of defining algebraic invariants with a causal model.

5.1 Intuition Underlying Our Definition of Causal Effects

Before proceeding to give rigorous definitions, it is useful to build up some intuition. Evaluating causal effects in terms of probabilistic models, or statistical models, or information-theoretic measures is so widespread that it may be difficult to imagine other ways of thinking about this problem. Let us consider what it means to evaluate the average treatment effect in a clinical trial, e.g., patients being tested for a Covid-19 vaccine, or indeed the counterfactual evaluation of the efficacy of a Covid-19 vaccine on patients in a hospital suffering from the effects of Covid-19. Consider the causal pathways that occur for patients who did take the vaccine, vs. the pathways for patients who did not take the vaccine. For the latter group, those who ended up in the hospital will end up suffering potential long-term effects from the Covid-19 infection, incur other medical complications such as shortness of breath, and end up requiring long-term treatment. Patients who were vaccinated and did not require hospitalizations would not in general require some medical intervention. In the population at large, of course, there are bound to be patients who despite being vaccinated end up in the hospital and end up requiring possibly long-term care, and similarly, patients who were never vaccinated still do not end up being infected by Covid-19 at a level serious enough to require hospitalization.

However, given the widespread published statistics showing patients admitted to hospitals for serious Covid-19 infections are preponderantly those who were not vaccinated, it is clear there is sufficient statistical regularity to discriminate between the two groups using an ATE estimator (either for testing a drug, or counterfactually, to determine whether they indeed would have required hospitalizations had they been vaccinated).

Using our formulation in terms of the category of elements, we can discriminate between two sets of causal pathways through the space, one requiring hospitalizations and long-term care, and the other who do not. The statistical regularity that discriminates between vaccinated and non-vaccinated patients also manifests itself in a change in the underlying density of trajectories between the two cases. The topological model we propose below can uncover the difference in these densities, and provides a non-statistical approach to evaluating causal effects.

For example, the method UMAP (uniform manifold approximation and projection) McInnes et al. (2018) is a powerful dimensionality reduction and data visualization method that is derived using the framework of category theory. UMAP can embed data in a lower-dimensional space in a way that separates objects from one class from objects in another class. UMAP relies on simplicial objects in terms of its design (actually, it uses an extended notion of “fuzzy” simplicial sets, where each simplicial set is mapped to a locale $[0, 1]$, the unit interval). Thus, by using such a category-theoretic or topologically aware method, one can discriminate between the classes of objects that were intervened on from those who were not.

5.2 Singular Homology

Before we describe our approach to defining causal effect more rigorously, we need to first define how to construct the topological realization of a category. We will then explore the properties of the classifying space of a category, and show how it can be useful in quantifying causal effects. We will also clarify how it relates to determining equivalences among causal models, namely homotopical invariance, and also how it sheds light on causal identification.

First, we need to define more concretely the topological n -simplex that provides a concrete way to attach a topology to a simplicial object. Our definitions below build on those given in (Lurie, 2022). For each integer n , define the topological space $|\Delta_n|$ realized by the object Δ_n as

$$|\Delta_n| = \{t_0, t_1, \dots, t_n \in \mathbb{R}^{n+1} : t_0 + t_1 + \dots + t_n = 1\}$$

This is the familiar n -dimensional simplex over n variables. For any causal model, its classifying space $|\mathcal{N}_\bullet(\mathcal{C})|$ defines a topological space. We can now define the *singular n -simplex* as a continuous mapping $\sigma : |\Delta_n| \rightarrow |\mathcal{N}_\bullet(\mathcal{C})|$. Every singular n -simplex σ induces a collection of $n - 1$ -dimensional simplices called *faces*, denoted as

$$d_i\sigma(t_0, \dots, t_{n-1}) = (t_0, t_1, \dots, t_{i-1}, 0, t_i, \dots, t_{n-1})$$

Note that as discussed above, a causal intervention on a variable in a DAG can be modeled as applying one of these degeneracy operators d_i . The above definition shows that every such intervention has an effect on the topology associated with the causal model. Define the set of all morphisms $\text{Sing}_n(X) = \mathbf{Hom}_{\mathbf{TOP}}(\Delta_n, |\mathcal{N}_\bullet(\mathcal{C})|)$ as the set of singular n -simplices of $|\mathcal{N}_\bullet(\mathcal{C})|$.

Definition 29. For any topological space defined by a causal model $|\mathcal{N}_\bullet(\mathcal{C})|$, the **singular homology groups** $H_*(|\mathcal{N}_\bullet(\mathcal{C})|; \mathbf{Z})$ are defined as the homology groups of a chain complex

$$\dots \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_2(|\mathcal{N}_\bullet(\mathcal{C})|)) \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_1(|\mathcal{N}_\bullet(\mathcal{C})|)) \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_0(|\mathcal{N}_\bullet(\mathcal{C})|))$$

where $\mathbf{Z}(\text{Sing}_n(|\mathcal{N}_\bullet(\mathcal{C})|))$ denotes the free Abelian group generated by the set $\text{Sing}_n(|\mathcal{N}_\bullet(\mathcal{C})|)$ and the differential ∂ is defined on the generators by the formula

$$\partial(\sigma) = \sum_{i=0}^n (-1)^i d_i\sigma$$

Intuitively, a chain complex builds a sequence of vector spaces that can be used to construct an algebraic invariant of a PSR from its classifying space by choosing the left \mathbf{k} module \mathbf{Z} to be a vector space. Each differential ∂ then becomes a linear transformation whose representation is constructed by modeling its effect on the basis elements in each $\mathbf{Z}(\text{Sing}_n(X))$.

Example 16. Let us illustrate the singular homology groups defined by an integer-valued multiset Studeny (2010) used to model conditional independence. Imsets over a DAG of three variables $N = \{a, b, c\}$ shown previously as Figure 7 can be viewed as a finite discrete topological space. For this topological space X , the singular homology groups $H_*(X; \mathbf{Z})$ are defined as the homology groups of a chain complex

$$\mathbf{Z}(\text{Sing}_3(X)) \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_2(X)) \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_2(X)) \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_1(X)) \xrightarrow{\partial} \mathbf{Z}(\text{Sing}_0(X))$$

where $\mathbf{Z}(\text{Sing}_i(X))$ denotes the free Abelian group generated by the set $\text{Sing}_i(X)$ and the differential ∂ is defined on the generators by the formula

$$\partial(\sigma) = \sum_{i=0}^4 (-1)^i d_i\sigma$$

The set $\text{Sing}_n(X)$ is the set of all morphisms $\mathbf{Hom}_{\mathbf{TOP}}(|\Delta_n|, X)$. For an imset over the three variables $N = \{a, b, c\}$, we can define the singular n -simplex σ as:

$$\sigma : |\Delta^4| \rightarrow X \text{ where } |\Delta^n| = \{t_0, t_1, t_2, t_3 \in [0, 1]^4 : t_0 + t_1 + t_2 + t_3 = 1\}$$

The n -simplex σ has a collection of faces denoted as $d_0\sigma, d_1\sigma, d_2\sigma$ and $d_3\sigma$. If we pick the k -left module \mathbf{Z} as the vector space over real numbers \mathbb{R} , then the above chain complex represents a sequence of vector spaces that can be used to construct an algebraic invariant of a topological space defined by the integer-valued multiset. Each differential ∂ then becomes a linear transformation whose representation is constructed by modeling its effect on the basis elements in each $\mathbf{Z}(\text{Sing}_n(X))$. An alternate approach to constructing a chain homology for an integer-valued multiset is to use Möbius inversion to define the chain complex in terms of the nerve of a category (see our recent work on categoroids (Mahadevan, 2022a) for details).

5.3 Classifying Spaces and Homotopy Colimits

Building on the intuition proposed above, we now introduce a formal way to define causal effects in our framework, which relies on the construction of a topological space associated with the nerve of a category. As we saw above, the nerve of a category is a full and faithful embedding of a category as a simplicial object.

Definition 30. The **classifying space** of a causal model defined as a category \mathcal{C} is the topological space associated with the nerve of the category $|N_\bullet \mathcal{C}|$.

We now want to bring in the set-valued functor mapping each causal category \mathcal{C} to the actual experiment used, e.g., in a clinical trial Imbens and Rubin (2015), to evaluate average treatment effect or quantify the effect of a **do** calculus intervention. We can then compute the topological space prior to intervention, and subsequent to intervention, and compare the two topological spaces in terms of their algebraic invariants (e.g., the chain complex, as described below).

In general, we may want to evaluate a causal model not only with respect to the data used in a causal experiment, but also with respect to some underlying topological space or some measurable space. We can extend the above definition straightforwardly to these cases using an appropriate functor $\mathcal{T} : \mathbf{Set} \rightarrow \mathbf{Top}$, or alternatively $\mathcal{M} : \mathbf{Set} \rightarrow \mathbf{Meas}$. These augmented constructions can then be defined with respect to a more general notion called the *homotopy colimit* Richter (2020) of a causal model.

Definition 31. The **homotopy colimit** $\text{hocolim}_{\mathcal{T} \circ \delta}$ of a causal model associated with a category \mathcal{C} , along with its associated category of elements associated with a set-valued functor $\delta : \mathcal{C} \rightarrow \mathbf{Set}$, and a topological functor $\mathcal{T} : \mathbf{Set} \rightarrow \mathbf{Top}$ is isomorphic to topological space associated with the nerve of the category of elements, that is $\text{hocolim}_{\mathcal{T} \circ \delta} \simeq |N_\bullet (f \delta)|$.

To understand the classifying space $|N_\bullet \mathcal{C}|$ of a causal model defined as a category \mathcal{C} , let us go over some simple examples to gain some insight.

Example 17. For any set X , which can be defined as a discrete category \mathcal{C}_X with no non-trivial morphisms, the classifying space $|N_\bullet \mathcal{C}_X|$ is just the discrete topology over X (where the open sets are all possible subsets of X).

Example 18. If we take a causal model defined as a partially ordered set $[n]$, with its usual order-preserving morphisms, then the nerve of $[n]$ is isomorphic to the representable functor $\delta(-, [n])$, as shown by the Yoneda Lemma, and in that case, the classifying space is just the topological space Δ_n defined above.

Example 19. The classifying space $|N_\bullet \mathcal{C}_{CDU}|$ associated with CDU symmetric monoidal category encoding of a causal Bayesian DAG (see Figure 9) is defined using the monoidal category (\mathbf{C}, \otimes, I) , where each object A has a copy map $C_A : A \rightarrow A \otimes A$, and discarding map $D_A : A \rightarrow I$, and a uniform state map $U_A : I \rightarrow A$, is defined as the topological realization of its nerve. As before, the nerve $N_n(\mathcal{C})$ of the CDU category is defined as the set of sequences of composable morphisms of length n .

$$\{C_o \xrightarrow{f_1} C_1 \xrightarrow{f_2} \dots \xrightarrow{f_n} C_n \mid C_i \text{ is an object in } \mathcal{C}, f_i \text{ is a morphism in } \mathcal{C}\}$$

Note that the CDU category was associated with a CDU functor $F : \mathbf{Syn}_G \rightarrow \mathbf{Stoch}$ to the category of stochastic matrices. We can now define the homotopy colimit $\text{hocolim}_{\mathcal{F}}$ of the CDU causal model associated with the CDU category \mathcal{C} , along with its associated category of elements associated with a set-valued functor $\delta : \mathcal{C} \rightarrow \mathbf{Set}$, and a topological functor $\mathcal{F} : \mathbf{Set} \rightarrow \mathbf{Stoch}$ is isomorphic to topological space associated with the nerve of the category of elements over the composed functor, that is $\text{hocolim}_{\mathcal{F} \circ \delta}$.

5.4 Defining Causal Effect

Finally, we turn to defining causal effect using the notion of classifying space and homotopy colimits, as defined above. Space does not permit a complete discussion of this topic, but the basic idea is that once a causal model is defined as a topological space, there are a large number of ways of comparing two topological spaces from analyzing their chain complexes, or using a topological data analysis method such as UMAP McInnes et al. (2018).

In conventional approaches, such as potential outcomes, one defines a random variable Y for outcomes (e.g., of a clinical trial), and then compares the difference of the outcomes for patients that were administered a drug vs. patients who were administered a placebo, which is formally written as $E(Y(1) - Y(0))$, where the expectation is over the population of all patients, and requires counterfactual estimation of quantities that are not in the data (e.g., patients who took the drug have no known outcomes for their outcomes under the placebo, and vice versa). In our UCLA framework, the causal effect between the two corresponding treatment and no-treatment cases amounts to a comparison of the classifying spaces. We formally define causal effect as follows.

Definition 32. Let the classifying space under ‘‘treatment’’ be defined as the topological space $|N_\bullet \mathcal{C}_1|$ associated with the nerve of category \mathcal{C}_1 under some intervention, which may result in a topological deformation of the model (e.g., deletion of an edge). Similarly, the classifying space under ‘‘no treatment’’ be defined as the $|N_\bullet \mathcal{C}_0|$ under a no-treatment setting, with no intervention. A **causally non-isomorphic effect** exists between categories \mathcal{C}_1 and \mathcal{C}_0 , or $\mathcal{C}_1 \not\cong \mathcal{C}_0$ if and only if there is no invertible morphism $f : |N_\bullet \mathcal{C}_1| \rightarrow |N_\bullet \mathcal{C}_0|$ between the ‘‘treatment’’ and ‘‘no-treatment’’ topological spaces, namely f must be both *left invertible* and *right invertible*.

There is an equivalent notion of causal effect using the homotopy colimit definition proposed above, which defines the nerve functor using the category of elements. This version is particularly useful in the context of evaluating a causal model over a dataset.

Definition 33. Let the homotopy colimit $\text{hocolim}_1 = |N_\bullet(\int \delta_1)|$ be the topological space associated with a causal category \mathcal{C}_1 under the “treatment” condition be defined with respect to an associated category of elements defined by a set-valued functor $\delta_1 : \mathcal{C} \rightarrow \mathbf{Set}$ over a dataset of “treated” variables, and corresponding “no-treatment” $\text{hocolim}_0 = |N_\bullet(\int \delta_0)|$ be the topological space of a causal model associated with a category \mathcal{C}_0 be defined over an associated category of elements defined by a set-valued functor $\delta_0 : \mathcal{C} \rightarrow \mathbf{Set}$ over a dataset of “placebo” variables. A **causally non-isomorphic effect** exists between categories \mathcal{C}_1 and \mathcal{C}_0 , or $\mathcal{C}_1 \not\cong \mathcal{C}_0$ if and only if there is no invertible morphism $f : |N_\bullet(\int \delta_1)| \rightarrow |N_\bullet(\int \delta_0)|$ between the “treatment” and “no-treatment” homotopy colimit topological spaces, namely f must be both *left invertible* and *right invertible*.

We can define an equivalent “do-calculus” like version of the causal effect definitions above for the case when a causal model defined as a graph structure is manipulated by an intervention that deletes an edge, or does some more sophisticated type of “category” surgery.

6 Contributions and Future Work

In this paper, we explored a layered hierarchical architecture for universal causality, based on the framework of category theory. Among the unique aspects of our formulation are the following novel contributions:

1. **Simplicial objects for causal “surgery”:** We introduced the framework of simplicial objects as a generic way to implement causal “surgery” on a model, which defines contravariant functors from Δ , the category of ordinal numbers, into a causal model. Simplicial objects provides an elegant and general way of extracting parts of a compositional structure, and its associated lifting problems define when a partial fragment of a causal model can be “put back” together into a complete model. Such “extension” problems, which were first explained at the beginning of the paper in connection with universal arrows, formalizes the notion of causal identifiability n in our framework.
2. **Universal Arrow:** We identified the process of “implementing” causal interventions from one categorical layer to the next with the definition of a universal arrow from objects in the co-domain category to the functor mapping the domain category to the co-domain category. It is the property of universal arrows that permits mapping causal interventions from the simplicial layer down to the causal category layer, and from the causal category layer into the category of elements.
3. **Lifting Problem:** Associated with each pair of layers of the UCLA hierarchy is a lifting problem over a suitable category of elements, from simplicial category of elements, to a category of elements over a dataset, to a category of elements over a topological space. In general, the Grothendieck category of elements is a way to embed each object in a category into the category of all categories **Cat**. This construction has many elegant properties, which deserves further exploration in a subsequent paper.
4. **Homotopy colimits and Classifying Spaces:** We defined causal effect in terms of the classifying space associated with the nerve of a causal category, and with the homotopy colimit of the nerve of the category of elements. These structures have been extensively explored in the study of homotopy in category theory Richter (2020), and there are many advanced techniques that can be brought to bear on this problem.

6.1 Future Work

There are many directions for future work, and we briefly describe a few of them below.

1. **Simplicial Causal Information Fields:** Although we described one particular instantiation of our UCLA framework, there are many other ways to define the various categorical layers that we could not include for reasons of space. In particular, we can easily adapt our approach to capture the work on causal information fields Heymann et al. (2020), where causal DAG models are formalized using Witsenhausen’s measure-theoretic formalization of causality Witsenhausen (1975). In particular, Heymann et al. (2021) show that Pearl’s do-calculus can be generalized using causal information fields to include causal models with feedback and other enhancements. In this setting, instead of conditional probability tables defining each variable of a model, a measure-theoretic information field is used that is defined as a measurable function from the product σ -algebras of its parents in the model. Witsenhausen defined causality as a condition on a given information field model of variables with a defined information field structure, whereby a given sequence of

variables could be determined such that each variable’s information field was uniquely computable given the values of the variables that preceded it in the ordering. Using the simplicial objects framework, we can define Witsenhausen’s causality condition as a contravariant functor from the ordinal numbers $[n]$ that defines the partial ordering he used in his theoretical characterization of causality. In addition, we can define a modified category of elements that combines the causal category model with the category of measurable spaces, and reformulate his framework (and that of causal information fields) in our approach.

2. **Causal Discovery from Conditional Independence Oracles:** We did not discuss the issue of causal discovery at length in this paper, but many approaches in the literature rely on conditional independence oracles Kocaoglu et al. (2017). In a previous paper, we defined categoroids Mahadevan (2022a), a category-theoretic formulation of universal conditional independence, which generalizes well-known axiomatizations of conditional independence, such as *separoids* Dawid (2010), *graphoids* Pearl (1989) and *imsets* Studeny (2010). Categoroids are formally defined as a *join* of two categories Lurie (2022), combining a subcategory specifying a causal model, and another subcategory specifying the conditional independence structure of the model. It is possible to combine categoroids with the UCLA hierarchy, as it is known that simplicial objects can be defined over the join of two categories Richter (2020). The problem of causal discovery can then be rigorously formulated as a lifting problem as well, using a simplicial extension of categoroids, where the conditional independence oracle is defined as a solution to a lifting problem (which is able to answer questions of the form $(x \perp\!\!\!\perp y | z)$ (is x conditionally independent of y given z ?)).
3. **Grothendieck Topology:** It is possible to define an abstract Grothendieck topology \mathcal{J} for any category, which leads to the concept of a *site* MacLane and Ieke Moerdijk (1994). In simple terms, for any object c in \mathcal{C} , a *sieve* S is a family of morphisms, all with co-domain c such that

$$f \in S \rightarrow f \circ g \in S$$

for any g where the composition is defined. A Grothendieck topology \mathcal{J} on category \mathcal{C} then defines a sieve $J(c)$ for each object c such that the following properties hold: (i) the maximal sieve $t_c = \{f | \text{cod}(f) = c\}$ is in $J(c)$. There is an additional stability condition and a transitive closure condition. An interesting problem for future work is to define causal inference over sheaves of a site, using the concept of Grothendieck topologies. Any causal intervention that, for instance, deletes an edge, would cause a change in the structure of sieves.

4. **Limitations of causal DAG models:** One reason for pursuing this line of research is to expand the scope of causal models to beyond well-studied paradigms, such as Bayesian networks and causal DAGs. Studeny (2010) astutely observes that even in a relatively toy problem of 4 variables, there are as many as 18,000 conditional independence structures, but that DAG models can barely capture more than a few hundred of them. This fact suggests that DAG models are severely limited, and that there is much to be gained by exploring more powerful formalisms for causal inference. In particular, Studeny (2010) himself proposed integer-valued sets (imsets) as a much more powerful formalism for representing conditional independence structures. The application of imsets to causal inference has not been well-studied, and the approach of constructing categories over imsets, which we explored in a previous paper Mahadevan (2022a), may be an interesting direction for future research.
5. **Causal confounding:** A notable omission of our paper is the lack of any discussion of the issue of confounding, one of the basic challenges facing any causal experiment in the real world. In the literature, many approaches to confounding have been studied, from models of missing data in statistics Imbens and Rubin (2015) to using marginalized DAG and hyperedge DAG models Evans (2018); Forré and Mooij (2017). Usually, the assumption made is that while confounders are not directly observable, their impact on the underlying model is limited, and their effects can be ascertained under various assumptions. To treat confounding properly in a genuine category-theoretic manner is a problem that is outside the scope of this paper, and we have chosen to leave a detailed study of this important topic for future work.

7 Summary

In this paper, we proposed a framework called Universal Causality (UC) for causal inference using the tools of category theory. Specifically, we described a layered hierarchical architecture called UCLA (Universal Causality Layered Architecture), where causal inference is modeled at multiple levels of categorical abstraction. At the top-most level, causal inference is modeled using a simplicial quasi-category of ordinal numbers Δ , whose objects are the ordered natural numbers $[n] = \{0, \dots, n\}$, and whose morphisms are order-preserving injections and surjections. Causal “surgery” is then modeled as the action of a contravariant functor from the category Δ into a causal model. At the second layer, causal models are defined by a category consisting of a collection of objects, such as the entities in a relational

database, and morphisms between objects can be viewed as attributes relating entities. The third categorical abstract layer corresponds to the data layer in causal inference, where each causal object is mapped into a set of instances, modeled using the category of sets and morphisms are functions between sets. The fourth layer comprises of additional structure imposed on the instance layer above, such as a topological space, a measurable space or a probability space, or more generally, a locale. Between every pair of layers in UCLA are functors that map objects and morphisms from the domain category to the co-domain category. Each functor between layers is characterized by a universal arrow, which defines an isomorphism between every pair of categorical layers. These universal arrows define universal elements and representations through the Yoneda Lemma, and in turn lead to a new category of elements based on a construction introduced by Grothendieck. Causal inference between each pair of layers is defined as a lifting problem, a commutative diagram whose objects are categories, and whose morphisms are functors that are characterized as different types of fibrations. We defined causal effect in the UCLA framework using the notion of homotopy colimits associated with the nerve of a category. We illustrate the UCLA architecture using a diverse set of examples.

References

- Brendan Fong. Causal theories: A categorical perspective on Bayesian networks, 2012.
- Bart Jacobs, Aleks Kissinger, and Fabio Zanasi. Causal inference by string diagram surgery, 2018. URL <https://arxiv.org/abs/1811.08338>.
- Sridhar Mahadevan. Universal decision models. *CoRR*, abs/2110.15431, 2021a. URL <https://arxiv.org/abs/2110.15431>.
- Sridhar Mahadevan. Categoroids: Universal conditional independence, 2022a. URL <https://arxiv.org/abs/2208.11077>.
- Sridhar Mahadevan. Unifying causal inference and reinforcement learning using higher-order category theory, 2022b. URL <https://arxiv.org/abs/2209.06262>.
- Saunders MacLane. *Categories for the Working Mathematician*. Springer-Verlag, New York, 1971. Graduate Texts in Mathematics, Vol. 5.
- Saunders MacLane and Ieke Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer, 1994.
- E. Riehl. *Category Theory in Context*. Aurora: Dover Modern Math Originals. Dover Publications, 2017. ISBN 9780486820804. URL <https://books.google.com/books?id=6B9MDgAAQBAJ>.
- André Joyal, Mogens Nielsen, and Glynn Winskel. Bisimulation from open maps. *Information and Computation*, 127(2):164–185, 1996. ISSN 0890-5401. doi:<https://doi.org/10.1006/inco.1996.0057>. URL <https://www.sciencedirect.com/science/article/pii/S0890540196900577>.
- J.P. May. *Simplicial Objects in Algebraic Topology*. Chicago Lectures in Mathematics. University of Chicago Press, 1992. ISBN 9780226511818. URL <https://books.google.com/books?id=QGjwV0gyQnIC>.
- Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, USA, 2nd edition, 2009. ISBN 052189560X.
- M. Studeny. *Probabilistic Conditional Independence Structures*. Information Science and Statistics. Springer London, 2010. ISBN 9781849969482. URL <https://books.google.com.gi/books?id=bGFRcgAACAAJ>.
- Misha Gavrilovich. The unreasonable power of the lifting property in elementary mathematics, 2017. URL <https://arxiv.org/abs/1707.06615>.
- Murat Kocaoglu, Karthikeyan Shanmugam, and Elias Bareinboim. Experimental design for learning causal graphs with latent variables. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 7018–7028, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/291d43c696d8c3704cdbe0a72ade5f6c-Abstract.html>.
- Patrick Schultz, David I. Spivak, Christina Vasilakopoulou, and Ryan Wisnesky. Algebraic databases, 2016. URL <https://arxiv.org/abs/1602.03501>.
- Kristopher Brown, David I. Spivak, and Ryan Wisnesky. Categorical data integration for computational science, 2019. URL <https://arxiv.org/abs/1903.10579>.
- David I. Spivak. Functorial data migration, 2010. URL <https://arxiv.org/abs/1009.1166>.
- David I. Spivak. Database queries and constraints via lifting problems. *Mathematical Structures in Computer Science*, 24(6), oct 2013. doi:10.1017/s0960129513000479. URL <https://doi.org/10.1017%2Fs0960129513000479>.
- Francis Borceux. *Handbook of Categorical Algebra*, volume 1 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1994. doi:10.1017/CBO9780511525858.
- P. Gabriel, P. Gabriel, and M. Zisman. *Calculus of Fractions and Homotopy Theory*. Calculus of Fractions and Homotopy Theory. Springer-Verlag, 1967. ISBN 9780387037776. URL <https://books.google.com/books?id=UEQZAQAAIAAJ>.
- Steen A. Andersson, David Madigan, and Michael D. Perlman. A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25(2):505 – 541, 1997. doi:10.1214/aos/1031833662. URL <https://doi.org/10.1214/aos/1031833662>.
- Sridhar Mahadevan. Causal homotopy, 2021b. URL <https://arxiv.org/abs/2112.01847>.
- J.P. May. *A Concise Course in Algebraic Topology*. Chicago Lectures in Mathematics. University of Chicago Press, 1999. ISBN 9780226511832. URL <https://books.google.com/books?id=g8SG03R1bpgC>.

- M. Boardman and Rainer Vogt. *Homotopy invariant algebraic structures on topological spaces*. Springer, Berlin, 1973.
- A. Joyal. Quasi-categories and kan complexes. *Journal of Pure and Applied Algebra*, 175(1):207–222, 2002. ISSN 0022-4049. doi:[https://doi.org/10.1016/S0022-4049\(02\)00135-4](https://doi.org/10.1016/S0022-4049(02)00135-4). URL <https://www.sciencedirect.com/science/article/pii/S0022404902001354>. Special Volume celebrating the 70th birthday of Professor Max Kelly.
- Jacob Lurie. Kerodon. <https://kerodon.net>, 2022.
- B. Richter. *From Categories to Homotopy Theory*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2020. ISBN 9781108479622. URL <https://books.google.com/books?id=pnzUDwAAQBAJ>.
- Milan Studený, Jíří Vomlel, and Raymond Hemmecke. A geometric view on learning bayesian network structures. *International Journal of Approximate Reasoning*, 51:573–586, 06 2010. doi:10.1016/j.ijar.2010.01.014.
- Brendan Fong and David I Spivak. *Seven Sketches in Compositionality: An Invitation to Applied Category Theory*. Cambridge University Press, 2018. URL <http://arxiv.org/abs/1803.05316>. cite arxiv:1803.05316Comment: 341+xii pages.
- Bob Coecke, Tobias Fritz, and Robert W. Spekkens. A mathematical theory of resources. *Information and Computation*, 250:59–86, oct 2016. doi:10.1016/j.ic.2016.02.008. URL <https://doi.org/10.1016%2Fj.ic.2016.02.008>.
- Judea Pearl. *Probabilistic reasoning in intelligent systems - networks of plausible inference*. Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann, 1989.
- B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002.
- Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, USA, 2015. ISBN 0521885884.
- Dominik Janzing, David Balduzzi, Moritz Grosse-Wentrup, and Bernhard Schölkopf. Quantifying causal influences. *The Annals of Statistics*, 41(5):2324 – 2358, 2013. doi:10.1214/13-AOS1145. URL <https://doi.org/10.1214/13-AOS1145>.
- Krikamol Muandet, Motonobu Kanagawa, Sorawit Saengkyongam, and Sanparith Marukatat. Counterfactual mean embeddings. *Journal of Machine Learning Research*, 22(162):1–71, 2021. URL <http://jmlr.org/papers/v22/20-185.html>.
- Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction, 2018. URL <http://arxiv.org/abs/1802.03426>. cite arxiv:1802.03426Comment: Reference implementation available at <http://github.com/lmcinnes/umap>.
- Benjamin Heymann, Michel de Lara, and Jean-Philippe Chancelier. Causal information with information fields. In *Neural Information Processing Systems Workshop on Causal Discovery and Causality-inspired Machine Learning*, 2020.
- H. S. Witsenhausen. The intrinsic model for discrete stochastic control: Some open problems. In A. Bensoussan and J. L. Lions, editors, *Control Theory, Numerical Methods and Computer Systems Modelling*, pages 322–335, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg. ISBN 978-3-642-46317-4.
- Benjamin Heymann, Michel de Lara, and Jean-Philippe Chancelier. Causal inference theory with information dependency models, 2021. URL <https://arxiv.org/abs/2108.03099>.
- A. Philip Dawid. Beware of the dag! In Isabelle Guyon, Dominik Janzing, and Bernhard Schölkopf, editors, *Causality: Objectives and Assessment (NIPS 2008 Workshop)*, Whistler, Canada, December 12, 2008, volume 6 of *JMLR Proceedings*, pages 59–86. JMLR.org, 2010. URL <http://proceedings.mlr.press/v6/dawid10a.html>.
- Robin J. Evans. Margins of discrete Bayesian networks. *The Annals of Statistics*, 46(6A):2623 – 2656, 2018. doi:10.1214/17-AOS1631. URL <https://doi.org/10.1214/17-AOS1631>.
- Patrick Forré and Joris M. Mooij. Markov properties for graphical models with cycles and latent variables, 2017.