
On Open-Universe Causal Reasoning

Duligur Ibeling

Computer Science Dept.
Stanford University
Stanford, CA 94305

Thomas Icard

Philosophy Dept.
Stanford University
Stanford, CA 94305

Abstract

We extend two kinds of causal models, structural equation models and simulation models, to infinite variable spaces. This enables a semantics for conditionals founded on a calculus of intervention, and axiomatization of causal reasoning for rich, expressive generative models—including those in which a causal representation exists only implicitly—in an open-universe setting. Further, we show that under suitable restrictions the two kinds of models are equivalent, perhaps surprisingly as their axiomatizations differ substantially in the general case. We give a series of complete axiomatizations in which the open-universe nature of the setting is seen to be essential.

1 Introduction

A hallmark of intelligence is the ability to reason about cause and effect. Indeed, some form of causal representation is essential for robust reasoning under uncertainty, including flexible planning, delivering useful explanations, and the capacity to transfer knowledge from one domain to another. How should causal knowledge be represented? Extant formalisms vary widely both in terms of their representational primitives and in the reasoning principles they engender, witness frameworks based on similarity-orderings (Lewis, 1973; Ginsberg, 1986), Bayesian networks, and structural equation models (Spirtes et al., 2000; Pearl, 2009), among others. While distinct and sometimes even incompatible (see, e.g., Halpern 2013; Zhang 2013 on similarity-orderings versus structural equation models, and Pearl 2009; Bottou et al. 2013 on Bayes nets versus structural equation models), each of these frameworks captures important insights about causal representation and reasoning.

A second hallmark of uncertain reasoning is the ability to deal with “open-universe” or “first-order” domains in a way that does not assume a fixed, finite set of variables. Milch et al. (2005) present an illustrative example of aircrafts producing identical blips on a radar screen. The radar blips give a noisy picture of the positions and velocities of some antecedently unknown number of aircrafts. Reasoning about evident in this setting introduces distinct challenges, chiefly to encode evidence about objects one did not know existed. Observing N blips, for example, is still consistent with any number of latent aircrafts. The field of statistical relational modeling has developed sophisticated methods for learning and reasoning in such domains (Poole, 2003; Milch et al., 2005; Carbonetto et al., 2005; Richardson and Domingos, 2006; Srivastava et al., 2014; Gogate and Domingos, 2016).

Research on these two aspects of uncertain reasoning has proceeded mostly independently. Whereas similarity-based causal representations in principle extend to first-order settings (e.g., Friedman et al. 2000), and there have been numerous extensions of directed graphical models to handle unbounded numbers of variables (e.g., Friedman et al. 1998; Pfeffer and Koller 2000), models with a perspicuous and useful causal interpretation typically assume a fixed, finite set of variables. That is, although in the course of learning one may always infer the existence of new hidden variables, no inferred causal model would typically employ an unbounded set of variables. Theoretical work on causal representation nearly always assumes this restriction (Galles and Pearl, 1998; Halpern, 2000; Spirtes et al., 2000; Zhang, 2013; Pearl, 2009).

Such a restriction is undoubtedly appropriate for many applications. But it may also be limitative. Returning to the radar blip example, for instance, we should be able to assess such claims as, “if there were more than 100 aircrafts, at least one would be missed by the radar,” or “given the observed blips, had there been five more aircrafts, at least two would have been dangerously close to one another,” and so on. These are essentially first-

order queries that depend on complex causal facts about an unknown number of entities and their properties and relations. Such causal queries ought to be answerable from a single underlying model (possibly under different observations), without having to specify a distinct set of variables for each use.

One promising approach to this problem comes from a tradition broadly within statistical relational modeling, whereby knowledge is encoded in the form of *generative probabilistic programs*. Probabilistic programs can be used as stochastic simulators and inverted to perform conditional inference given observational data, capturing an important variety of probabilistic learning (Milch et al., 2005; Goodman et al., 2008; de Raedt and Kimmig, 2015; Tran et al., 2017; Bingham et al., 2019). Encoding a generative process implicitly as a program facilitates succinct representation of complex dependencies among unbounded sets of variables.

Similarly, within research on neural networks, so called *implicit generative models*—tools such as variational autoencoders and generative adversarial networks—also represent probability distributions implicitly by means of stochastic data simulators (Mohamed and Lakshminarayanan, 2017). Though these models are typically incapable of dealing with fully open-universe domains, extensions may in principle be expressive enough to do so (see, e.g., Li et al. 2017). Related ideas involving “mental simulation” of environmental dynamics are being increasingly explored within neural approaches to reinforcement learning (Hamrick, 2019).

Probabilistic generative models are often associated with an intuitive causal interpretation, although this interpretation is not always appropriate (Peters et al., 2017). For example, there should be a reason for thinking that parameters of the model correspond appropriately to aspects of the true underlying data-generating process (Besserve et al., 2018). Our interest here is in a more abstract question: what structural features would be sufficient in order for an open-universe model even to be a candidate as a causal model?

Our aim in this paper is to establish formal (axiomatic) foundations for genuinely *open-universe causal models*. Specifically, we want to understand and assess the subjunctive conditional claims they encode. Subjunctive conditionals—“what if?” statements about what would occur under counterfactual or hypothetical conditions—are of foundational importance for causal reasoning, arguably definitional of the subject matter (Spirtes et al., 2000; Pearl, 2009; Bottou et al., 2013; Peters et al., 2017). Such conditional statements are typically formalized by appeal to a notion of *causal intervention* (Meek and Glymour, 1994; Spirtes et al., 2000; Pearl, 2009),

and axioms for conditionals play crucially into prominent algorithms for inferring counterfactual and interventional probabilities (“causal effects”) from observational data (Shpitser and Pearl, 2008; Pearl and Bareinboim, 2012; Hyttinen et al., 2015).

Open-universe models can be specified either implicitly via a (data-generating) program or process, or explicitly via a set of equations relating the variables. In general these two kinds of models—together with natural corresponding concepts of hypothetical/counterfactual intervention—are quite different, both conceptually and axiomatically (Ibeling and Icard, 2018). For instance, they handle loops and feedback in incompatible ways (cf. also Lauritzen and Richardson 2002). Nonetheless, we show that under suitable restrictions, well-motivated from a causal perspective, the two are expressively equivalent. Thus, our first contribution is to define two natural classes of simulation programs and structural equation models, respectively, and show equivalence of these two classes with respect to the conditional claims they entail. The characterizations appeal to an implicit (discrete) temporal structure, and can be seen as “open-universe data-generating processes” generalizing the idea of a “recursive” causal model in the sense of Pearl (2009), viewed either procedurally or declaratively. We offer this as a formalization of those generative models that could plausibly be interpreted as causal models.

Following this we establish a series of axiomatization results with respect to natural systems for reasoning about subjunctive conditionals. Our results build on previous axiomatic work on causal conditionals that assumed a fixed, finite set of variables (Galles and Pearl, 1998; Halpern, 1998, 2000), extending this work to the open-universe setting. Dealing with an unbounded set of variables presents unique challenges. Nevertheless, we show that an axiomatic system including quintessential causal principles (including those used in aforementioned identifiability algorithms) is sound and complete for both the procedural and the declarative classes of models, thereby substantiating the causal interpretation of both classes.

Bringing out the open-universe aspect of these models, we consider augmenting the conditional language with a *causal influence* relation \rightsquigarrow , whereby $X \rightsquigarrow Y$ expresses that there is a context in which a change in X would lead to a change in Y (Woodward, 2003; Pearl, 2009). Whereas this relation is easily definable in the finite setting (Halpern, 1998), in our setting \rightsquigarrow is an essentially higher-order notion, quantifying over all possible interventions. Reasoning about this relation in any particular model—most notably in rich probabilistic programs—may in general be undecidable; however, the satisfiability problem for all the systems we consider remains NP-

complete, showing that abstract causal reasoning (e.g., as features in the do-calculus) in the open-universe setting is no more complex than in the finite setting.

Finally, as an additional application of this study, we consider the special case where all causal influence is *local*, in the sense that whenever $X \rightsquigarrow Y$, this influence is mediated by variables temporally in between X and Y . As we explain, this semantic assumption is quite natural in the open-universe setting, and axiomatically it corresponds exactly to the claim that \rightsquigarrow is a transitive relation.

The broader goal of this work is to establish foundations suitable for understanding and assessing highly expressive (open-universe) representational systems that encode causal structure, even if only implicitly. Such models feature not only in recent AI research, but also centrally in models of human cognition (see, e.g., Freer et al. 2012; Lake et al. 2017). Indeed, the open-universe nature of human reasoning under uncertainty is well-established (Kemp, 2012). We may thus expect that, however causal knowledge is represented in humans, it accommodates flexible, open-universe reasoning. At a high level, the present work brings together research on causal representation and reasoning with ideas and tools from statistical relational modeling and higher-order representation.

2 Two Classes of Models

In this section we introduce two kinds of models, one based on systems of equations (§2.1), the other based on algorithms (§2.2). They respectively emphasize declarative and procedural styles of modeling. Owing in part to this difference, the most general versions of the two can be distinguished even at a very abstract, axiomatic level (Ibeling and Icard, 2018). However, by restricting to appropriate (causally motivated) subclasses, we can demonstrate their equivalence (Thm. 1).

Throughout this section we assume a signature (χ, Σ) , specifying a countably infinite set χ of variables that take on values from a set Σ . For example, we might have a variable in χ representing “the position of the 77th aircraft,” which could take on values in some numerical range (or a special value “undefined” if there is no 77th aircraft). Even if we restrict each variable to take only finitely many possible values, the fact that χ is infinite allows encoding arbitrarily complex structures.

2.1 Structural Equation Models

The most general class of structural equations models (see, e.g., Pearl 2009) allows arbitrary equations among arbitrary sets of variables. Because our primary interest here is in (open-universe, probabilistic) generative mod-

els, we restrict attention to a special subclass of possible sets of equations. First, we assume an infinite set of variables χ . Second, similar to so called recursive models (Pearl, 2009), we will assume that variables can be given a causal (intuitively, temporal) order. Although there may well be adequate causal interpretations of nonrecursive models, e.g., as descriptions of equilibrium behavior of some underlying process (Strotz and Wold, 1960), these interpretations rely on the declarative character of equational modeling and cannot always be adequately captured by an appropriate data-generating process (Lauritzen and Richardson, 2002).

Third, unlike in dynamic Bayesian networks and related extended graphical models (Dean and Kanazawa, 1989; Friedman et al., 1998), we do not require each variable to depend on only finitely many others; a variable may depend on an unbounded number of variables preceding it in a temporal order. Fourth, we require all equations to be uniformly computable (Defn. 4).

Note, finally, that we make no use here of the distinction between exogenous and endogenous variables (though see the discussion in §2.3 below).

Definition 1 (Structural Equation Model). A *structural equation model* (SEM) is a collection of partial functions $\{f_X\}_{X \in \chi}$, with $f_X : (\chi \rightarrow \Sigma) \rightarrow \Sigma$, and a *time* map $t : \chi \rightarrow \mathbb{N}$. Each f_X is a function only of preceding variables: if $v, v' \in \text{dom}(f_X)$ and $v(X') = v'(X')$ for all X' such that $t(X') < t(X)$, then $f_X(v) = f_X(v')$.

The functions in an SEM specify *structural equations* $\{X = f_X(\cdot)\}_{X \in \chi}$ that are to be simultaneously satisfied: an SEM M has *solution* $v : \chi \rightarrow \Sigma$ if $f_X(v) = v(X)$ for all X . We write $M \models v$ in this case. Since the order t is well-founded and acyclic, v may be built up iteratively, and is thus unique. Intervention on an SEM is defined standardly (Meek and Glymour, 1994; Pearl, 2009):

Definition 2 (Intervention on an SEM). An intervention is a partial function $i : \chi \rightarrow \Sigma$. It specifies variables $\text{dom}(i) \subseteq \chi$ to be held fixed, and the values to which they are fixed. Intervention i induces a mapping of SEMs, also denoted i , as follows. Abbreviate $X \in \text{dom}(i)$ as $X \in i$. Then $i(M)$ is identical to M , but with f_X replaced by the constant function $f_X(\cdot) = i(X)$ for each $X \in i$.

Given i, i' we say i is a *restriction* of i' if $\text{dom}(i) \subseteq \text{dom}(i')$ and $i(X) = i'(X)$ for all $X \in i$. Note i' may be a total function $v : \chi \rightarrow \Sigma$. We define a natural relation of *causal influence* in SEMs as follows (after Woodward 2003; Pearl 2009).

Definition 3 (Causal influence). Let $X, Y \in \chi$ with $X \neq Y$. Then $M \models X \rightsquigarrow Y$ (read X influences Y) if there is an intervention i (the *context*) and $x_1, x_2 \in \Sigma$

so that, for $j = 1, 2$ letting i_{x_j} be the intervention fixing X to x_j and v_j be the solution $i_{x_j}(i(M)) \models v_j$, we have $v_1(Y) \neq v_2(Y)$.

If $X \rightsquigarrow Y$ then $t(X) < t(Y)$; this is easy to see again from the fact that the temporal order implies that a solution may be assembled iteratively.

For comparison with simulation models, it is helpful to identify those SEMs whose (counterfactual) solutions are computable. Similar restrictions have been explored in the literature, e.g., on approaches to causal inference based on so called algorithmic independence (Janzing and Schölkopf, 2010; Peters et al., 2017). This is the subclass in which the functions $\{f_X\}_{X \in \chi}$ are *uniformly* computable and t is also computable (both presuming an encoding of χ):

Definition 4 (Computable Structural Equation Model). Let $M = \{f_X\}_{X \in \chi}$ be an SEM with time map t , and let F be the set of computable functions $\chi \rightarrow \Sigma$. We say M is *computable* if t is computable, $\text{dom}(f_X) = F$ for all X , and $\phi : \chi \times F \rightarrow \Sigma$ defined by $\phi(X, v) = f_X(v)$ is a computable function.

To clarify the sense in which ϕ is computable, let us associate each variable in χ with a square of an infinite Turing machine tape with alphabet Σ . We call such a tape a *variable tape*. A variable tape (no square of which is blank) stores a function $\chi \rightarrow \Sigma$. That ϕ is computable means that it is computed by a machine M with two input tapes, the second of which is a variable tape. When M is run with the encoding of some $X \in \chi$ stored on the first tape and a computable function $v : \chi \rightarrow \Sigma$ stored on the second (variable) tape, M halts outputting the value $\phi(X, v)$. M need not halt given uncomputable input v .

Let us write \mathcal{M} for the class of all computable SEMs; clearly \mathcal{M} closes under computable intervention and every SEM in \mathcal{M} has a unique solution.

Let $\mathcal{M}_{\text{local}}$ be the subclass of \mathcal{M} in which every $X \in \chi$ only depends on immediately preceding variables: if $v, v' \in \text{dom}(f_X)$ and $v(X') = v'(X')$ for all X' such that $t(X') = t(X) - 1$, then $f_X(v) = f_X(v')$. Causal influence is not transitive in SEMs generally, but it is transitive in $\mathcal{M}_{\text{local}}$ (Pearl, 2009). Furthermore, we can prove the following ‘‘interpolation’’ result:

Proposition 1. Let $M \in \mathcal{M}_{\text{local}}$ and $X, Y \in \chi$. If $M \models X \rightsquigarrow Y$ and $t(Y) > t(X) + 1$, there is a variable X' such that $M \models X \rightsquigarrow X'$ and $M \models X' \rightsquigarrow Y$.

Proof. There is a context i and x_j, v_j for $j = 1, 2$ such that $i_{x_j}(i(M)) \models v_j$ and $v_1(Y) \neq v_2(Y)$ (Defn. 3); thus $\phi(Y, v_1) \neq \phi(Y, v_2)$ (Defn. 4). Let $S_j \subset \chi$ be the set of variables whose squares the computation of $\phi(Y, v_j)$

accesses before outputting a value for Y and halting; crucially S_j is finite. Form intervention i_j as the restriction of v_j to $S_1 \cup S_2$ and let u_j be the solution $i_j(M) \models u_j$. Since $M \in \mathcal{M}_{\text{local}}$, $u_j(Y) = v_j(Y)$. Let $\{X_i\}_{2 \leq i \leq k}$ be those variables in $\text{dom}(i_1)$ for which $i_1(X_i) \neq i_2(X_i)$, and consider a finite sequence $\{i'_n\}_{1 \leq n \leq k}$ where $i'_1 = i_1$, and for $2 \leq n \leq k$, $i'_n = i'_{n-1}$ with the sole exception that $i'_n(X_n) = i_2(X_n)$. Let $\{v'_n\}_{1 \leq n \leq k}$ be the solutions $i'_n(M) \models v'_n$. Since $i'_k = i_2$, we have $v'_k(Y) = u_2(Y) = v_2(Y) \neq v_1(Y) = u_1(Y) = v'_1(Y)$; there is thus some m , $2 \leq m \leq k$ for which $v'_m(Y) \neq v'_{m-1}(Y)$. Then i'_{m-1} witnesses that $M \models X'_m \rightsquigarrow Y$ and i_1 witnesses that $M \models X \rightsquigarrow X'_m$ as $v_1(X'_m) = i_1(X'_m) \neq i_2(X'_m) = v_2(X'_m)$. \square

Taking the temporal interpretation of t seriously, Prop. 1 reflects the intuition that causal influence is always mediated through time, essentially enforcing a kind of Markov assumption.

2.2 Monotone Simulation Models

Our second class of models similarly deals with an infinite class of variables, but is intended to capture the procedural character of a simulation program, e.g., as defined in an expressive programming language. Such models have been studied in relative generality by means of (probabilistic) Turing machines (see, e.g., Freer et al. 2012 for an overview), and at this level of generality they can be shown to validate only a very weak set of conditional axioms (Ibeling and Icard, 2018), far weaker than those commonly required of causal conditionals (cf. §3.2 below). Here we restrict attention to machines in which the variables can be given a causal order, again analogously to recursive graphical models. As simulation programs are intended to capture an underlying data-generating process, this causal order can be well justified (cf. Lauritzen and Richardson 2002).

Definition 5 (Simulation model). A *simulation model* (or just a simulation) is a deterministic Turing machine with two tapes: a work tape and a variable tape (see discussion following Defn. 4). The variable tape is write-only: no variable head transitions in which a symbol of Σ is erased or rewritten are allowed.

Thanks to the write-only restriction, any variable is written at most once. Given $v : \chi \rightarrow \Sigma$, if every X is written with the value $v(X)$, v is the solution of T and we write $T \models v$. Intervention is defined via blocking rewrites (Icard, 2017; Ibeling and Icard, 2018):

Definition 6 (Intervention). Given intervention i and an oracle for i , the simulation $i(T)$ emulates T but acts as if the square for any $X \in i$ is fixed to the value $i(X)$;¹

¹Formally, suppose at some point in the execution the next

it dovetails this emulation with a procedure that writes $i(X)$ to X for all $X \in i$.

In other words, an intervention on a simulation program T is an operation on the code of T that holds some set of variables to fixed values, an operation that will typically have side-effects on other variables.

A computable intervention is one whose oracle can be effectively implemented; simulations thus close under computable intervention. The primary phenomena of interest are the behaviors of a simulation under possible interventions. This motivates an equivalence notion:

Definition 7. Simulations T, T' are *equivalent*, $T \simeq T'$, if for any computable i and $v : \chi \rightarrow \Sigma$, $i(T) \models v$ iff $i(T') \models v$. They are *weakly equivalent* if this property only holds when $\text{dom}(i) = \emptyset$.

It is immediate that interventions compose: for any interventions i, i' , there is an intervention $i' \circ i$ such that $(i' \circ i)(T) \simeq i'(i(T))$ for all T . If i, i' are computable then so is $i' \circ i$.

The mere fact that a variable has been intervened on may affect a simulation, even if it has been held fixed to its actual value (Ibeling and Icard, 2018). This behavior differs starkly from that of SEMs, so to compare the models we must consider only the class of simulations in which non-counterfactual interventions are idempotent:

Definition 8. We say T is (strongly) *functional* if for any i with $i(T) \models v$ and any restriction i' of v , $i'(i(T))$ and $i(T)$ are weakly equivalent.

We define the causal influence relation \rightsquigarrow in simulations analogously to Defn. 3. We call a simulation *monotone* if it is well-ordered with respect to causal influence:

Definition 9. A simulation T is *monotone* if there is a computable time map $t : \chi \rightarrow \mathbb{N}$ preserving \rightsquigarrow .

It follows from compositionality of interventions that if T is monotone under t , then so is $i(T)$ for any i . Here we consider only simulations that have a solution under any computable intervention. Let \mathcal{T} be the class of such simulations that are also functional and monotone.

2.3 Relation to Existing Work

Thm. 1 below shows that the class \mathcal{M} of computable SEMs can essentially be seen as declarative versions of

variable tape transition of T is $\delta(q, s) = (q', s', d)$. Before simulating this transition $i(T)$ calls the oracle on X . If $X \notin i$ then the simulation proceeds unimpeded: T transitions to state q' , s' is written to square X , and the head moves in direction d . If $X \in i$ then instead, the simulation does the action $(q'', i(X), d')$ where $\delta(q, i(X)) = (q'', s'', d')$.

the class \mathcal{T} of generative programs. Both classes of models naturally accommodate probability. For SEMs, we simply identify a subset of the variables as *exogenous* and associate those variables with an appropriate probability distribution, which in turn induces (conditional and interventional) distributions on the remaining (“endogenous”) variables (Pearl, 2009). For simulation programs, we simply add an additional read-only *random bit tape* whose distribution induces random behavior in the machine, which in turn implicitly defines (conditional and interventional) distributions on the variables χ (cf. Freer et al. 2012). The resulting models encompass a wide range of familiar and powerful formalisms.

For instance, \mathcal{M} clearly incorporates all computable (Defn. 4) and recursive (in the sense of Pearl 2009) SEMs. Likewise, \mathcal{T} certainly includes Bayesian networks as well as dynamic Bayesian networks (Dean and Kanazawa, 1989), in addition to standard dynamic programming algorithms and feedforward neural networks. But crucially, \mathcal{T} also encompasses programs in which variables may depend on infinitely many others, such as those defined in common probabilistic programming languages (Milch et al., 2005; Goodman et al., 2008; de Raedt and Kimmig, 2015; Tran et al., 2017; Bingham et al., 2019). This includes, e.g., the BLOG program formalizing the radar blip example, among many others.

As in previous work (e.g., Galles and Pearl 1998; Halpern 1998, 2000; Zhang 2013, etc.), the axiomatizations below in §3 factor out probabilities (in effect, holding all probabilistic choices fixed) and focus only on the deterministic core, as defined above in §2.1, §2.2.

2.4 The Equivalence of Monotone Simulations and Computable SEMs

Analogously to Defn. 7, we can assess equivalence between between any two causal models, whether simulations or SEMs. It turns out that the classes \mathcal{T} and \mathcal{M} are equivalent in the following sense.

Theorem 1. For any $M \in \mathcal{M}$ there is a $T \in \mathcal{T}$ such that $M \simeq T$ and vice versa.

Proof. Let $M \in \mathcal{M}$. Consider $T \in \mathcal{T}$ with this pseudocode: T consists of an infinite loop calling the following subroutine $\text{CALC}(X)$ on every $X \in \chi$. The subroutine $\text{CALC}(X)$ returns immediately if the variable tape square for X isn’t blank. Otherwise it emulates a computation of $\phi(X, \cdot)$ (Defn. 4) and writes the result to the square for X . The *emulated* input variable tape (which holds the second argument of ϕ) behaves as follows on access to the square for some variable Y . If the square for Y in the *original* output variable tape (which T writes to) is not blank, then the emulated tape holds the value from

the original tape. If it is blank, then the simulation computes $t(Y)$ and $t(X)$; if $t(Y) < t(X)$ then $\text{CALC}(Y)$ is recursively called so that the Y -square is no longer blank and one can proceed as above. Otherwise, the emulated tape contains some arbitrary value in its Y -square.

Note T is functional since the square for any X can be written only by a call to $\text{CALC}(X)$. Next let us show that $M \simeq \text{T}$, which immediately implies that T is monotone. Let i be an intervention, and suppose $i(M) \models v$. We show by induction that $i(\text{T}) \models v$. The base case: a variable X with $t(X) = t_0 = \min_{X' \in \mathcal{X}} t(X')$. If $X \notin i$ then $v(X) = f_X(v') = \phi(X, v')$ for any input v' ; thus $\text{CALC}(X)$ writes $v(X)$. $X \in i$ is trivial. Now consider an X with $t(X) > t_0$. By induction and construction of the emulator, the value output by $\text{CALC}(X)$ is $f_X(v')$ at some v' such that $v'(Y) = v(Y)$ for all Y such that $t(Y) < t(X)$. By Defn. 1, $f_X(v') = f_X(v) = v(X)$ so we are done. An analogous induction shows that if $i(\text{T}) \models v$ then $i(M) \models v$.

Now let $\text{T} \in \mathcal{T}$. Define $M \in \mathcal{M}$ by giving the following pseudocode to compute $\phi(X, v)$: emulate the run of $i'(\text{T})$, where i' is the restriction of v to all variables but X , until reaching a write to the emulated X -square, and halt outputting this value. This is computable in the sense of Defn. 4 since the oracle (Defn. 6) for i' on X' can simply check that $X' \neq X$ and then look on the input variable tape to find the required value $v(X')$. If $i(\text{T}) \models v$, when computing any $\phi(X, v)$ the i' above is a restriction of v , so $i(M) \models v$ by functionality (Defn. 8). Consider Lem. 1 below for the converse.

Lemma 1. Let $\text{T} \in \mathcal{T}$, i be an intervention, and $t \in \mathbb{N}$. Suppose i' is the restriction of i to those X' such that $t(X') < t$, with $i(\text{T}) \models v$ and $i'(\text{T}) \models v'$. Then $v(X) = v'(X)$ for all $X \notin i$ with $t(X) \leq t$.

Proof. Such an X is written to at some point in executing both $i(\text{T}), i'(\text{T})$ so we have finite restrictions i_∞, i'_∞ of i, i' such that $i_\infty(\text{T}), i'_\infty(\text{T})$ write the same value to X as do $i(\text{T}), i'(\text{T})$ resp. Transforming i'_∞ into i_∞ stepwise (cf. proof of Prop. 1) by adding one more variable Y to the domain in each step, if $v(X) \neq v'(X)$ then some step yields a Y such that $\text{T} \models Y \rightsquigarrow X$ but $t(Y) \geq t(X)$. \square

Suppose $i(M) \models v$; we show $i(\text{T}) \models v$ by induction. The base case is again a variable X with $t(X) = t_0$. If $X \notin i$, apply Lem. 1 choosing i', t_0 as i, t ; the restriction is empty, so the value output for X by $i'(\text{T})$ agrees with that output by T , which in turn agrees with that output by $i(\text{T})$ since $t(X) = t_0$ and $X \notin i$. For the inductive case with $t(X) = t > t_0$, choosing i', t in Lem. 1 and applying functionality and composition gives that $i(\text{T})$ writes $v(X)$ to its X -square. \square

3 The Conditional Logic of Open-Universe Causal Models

Thm. 1 establishes a strong equivalence between simulation models and (computable) SEMs, showing that they encode exactly the same underlying conditional theories. In this section we offer an axiomatic characterization of these conditional theories.

We introduce two formal languages of subjunctive conditionals interpreted in either \mathcal{M} or \mathcal{T} under their respective semantics of intervention. We then find sound and complete axiomatizations of their validities. For reasons of simplicity and elegance, but not necessity, we assume a binary alphabet $\Sigma = \{0, 1\}$ in this section.

3.1 Syntax and Semantics

We admit the variables of \mathcal{X} as atoms and let $\mathcal{L}_{\text{prop}}$ be the language of propositional formulas over these atoms. Let $\mathcal{L}_{\text{int}} \subset \mathcal{L}_{\text{prop}}$ be the language of purely conjunctive, ordered formulas of literals formed from distinct atoms. To be precise, let us call ℓ an X -literal if $\ell \in \{X, \neg X\}$. Then \mathcal{L}_{int} consists of formulas of the form $\ell_1 \wedge \dots \wedge \ell_n$, where each ℓ_i is an X_i -literal for some distinct $X_1, \dots, X_n \in \mathcal{X}$. Each $\alpha \in \mathcal{L}_{\text{int}}$ specifies a finite intervention by giving fixed values for a fixed list of variables, and we also write α for this intervention. \mathcal{L}_{int} includes the empty intervention \top with $\text{dom}(\top) = \emptyset$.

Let $\mathcal{L}_{\text{cond}}$ be the language of formulas of the form $[\alpha]\beta$ for $\alpha \in \mathcal{L}_{\text{int}}, \beta \in \mathcal{L}_{\text{prop}}$. We call such a formula a *subjunctive conditional*, and call α the *antecedent* and β the *consequent*. The overall conditional language \mathcal{L} is the language of propositional formulas over atoms in $\mathcal{L}_{\text{cond}}$. For $\alpha, \beta \in \mathcal{L}$, $\alpha \rightarrow \beta$ abbreviates $\neg\alpha \vee \beta$, and $\alpha \leftrightarrow \beta$ means $(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)$.

To give the semantics of \mathcal{L} , we first define a satisfaction relation between solutions $v : \mathcal{X} \rightarrow \Sigma$ and formulas of $\mathcal{L}_{\text{prop}}$. For $X \in \mathcal{X}$, write $v \models X$ iff $v(X) = 1$. For arbitrary $\varphi \in \mathcal{L}_{\text{prop}}$, $v \models \varphi$ is defined familiarly by recursion. Now we may define satisfaction for $\mathcal{L}_{\text{cond}}$. Let M be a model in either \mathcal{M} or \mathcal{T} . Write $M \models [\alpha]\beta$ if $\alpha(M)$ has solution v and $v \models \beta$. Finally, satisfaction for the entire language \mathcal{L} is defined by recursion.

We also consider an augmented language \mathcal{L}^+ , defined by extending \mathcal{L} to include additional atoms $X \rightsquigarrow Y$ for $X, Y \in \mathcal{X}$. This is interpreted as causal influence: $X \rightsquigarrow Y$ holds in model M if $M \models X \rightsquigarrow Y$ (recall Defn. 3).

Thm. 1 implies that any $\varphi \in \mathcal{L}^+$ is satisfiable in \mathcal{T} iff satisfiable in \mathcal{M} . Hence any semantic property is oblivious to the class in which to interpret it and we omit this.

A first observation is that our language and interpreta-

tions do not enjoy the property of *compactness* (the property that every unsatisfiable set of formulas has a finite unsatisfiable subset), and our axiomatizations are hence only complete relative to finite sets of assumptions:

Proposition 2. The languages \mathcal{L} , \mathcal{L}^+ are not compact.

Proof. Enumerate the variables as $\chi = \{X_i\}_{i \in \mathbb{N}}$. Consider the set $\Omega = \{[X_{i+1}]X_i \wedge [\neg X_{i+1}]\neg X_i : i \in \mathbb{N}\}$. Every finite subset of Ω is easily seen to be satisfiable. But each formula implies that $X_{i+1} \rightsquigarrow X_i$, contradicting that t is well-founded; thus Ω is unsatisfiable. \square

3.2 Axiomatizations

We now give an axiomatization **AX** of the validities of \mathcal{L} . Start with the base system $\mathbf{AX}_{\text{det}}^\downarrow$ from Ibeling and Icard (2018), obtained by the axioms and rules:

- PC. Propositional calculus
- RW. From $\beta \rightarrow \beta'$ infer $[\alpha]\beta \rightarrow [\alpha]\beta'$
- R. $[\alpha]\alpha$
- K. $[\alpha](\beta \rightarrow \gamma) \rightarrow ([\alpha]\beta \rightarrow [\alpha]\gamma)$
- F/D. $[\alpha]\neg\beta \leftrightarrow \neg[\alpha]\beta$

The causal axiom of *effectiveness* (Galles and Pearl, 1998; Halpern, 2000) is encoded by **RW** and **R**, while **F/D** expresses the fact that solutions always exist uniquely. System $\mathbf{AX}_{\text{det}}^\downarrow$, without **F/D**—a system significantly weaker than any existing logic of causal conditionals in the literature—has been shown sound and complete with respect to the general class of simulation models, while **F/D** can be added to axiomatize those programs that are deterministic and/or always halt (Ibeling and Icard, 2018). In the present setting we are assuming determinism (but allowing a natural incorporation of probability; recall §2.3).

In order to capture the logic of our restricted classes of SEMs and simulation programs, we add two further axioms to obtain the system $\mathbf{AX} = \mathbf{AX}_{\text{det}}^\downarrow + \mathbf{C} + \mathbf{Rec}$. Axiom **C** is known variously as *cautious monotonicity* (Kraus et al., 1990) or *composition* (Galles and Pearl, 1998; Pearl, 2009), and schema **Rec** is known as *recursiveness* (Halpern, 1998, 2000).

- C. $([\alpha]\beta \wedge [\alpha]\gamma) \rightarrow [\alpha \wedge \beta]\gamma$
- Rec. $([\alpha_1 \wedge X_1]\ell_2 \wedge [\alpha_1 \wedge \neg X_1]\neg\ell_2 \wedge \dots \wedge [\alpha_{k-1} \wedge X_{k-1}]\ell_k \wedge [\alpha_{k-1} \wedge \neg X_{k-1}]\neg\ell_k) \rightarrow \neg([\alpha_k \wedge X_k]\ell_1 \wedge [\alpha_k \wedge \neg X_k]\neg\ell_1)$

In **Rec**, each ℓ_i is an X_i -literal and $X_1 \neq X_k$. Note that the axiom commonly known as *reversibility* is easily derivable in **AX** (Galles and Pearl, 1998).

For \mathcal{L}^+ , consider the additional axioms

- Wit. $([\alpha \wedge X_1]\ell_2 \wedge [\alpha \wedge \neg X_1]\neg\ell_2) \rightarrow X_1 \rightsquigarrow X_2$
- Rec⁺. $(X_1 \rightsquigarrow X_2 \wedge \dots \wedge X_{k-1} \rightsquigarrow X_k) \rightarrow X_k \not\rightsquigarrow X_1$
- Trans. $(X \rightsquigarrow Y \wedge Y \rightsquigarrow Z) \rightarrow X \rightsquigarrow Z$

where in **Wit**, ℓ_2 is an X_2 -literal and $X_1 \neq X_2$. Let $\mathbf{AX}^+ = \mathbf{AX} + \mathbf{Wit} + \mathbf{Rec}^+$.

We are now in a position to present our three axiomatization results.

Theorem 2. **AX** is sound and complete with respect to the validities of \mathcal{L} .

Proof. By Thm. 1 it suffices to carry out the proof for \mathcal{M} only. Soundness is straightforward. As for completeness, we show that any $\varphi \in \mathcal{L}$ consistent with **AX** is satisfiable. Let $\chi_\varphi \subset \chi$ be the set of variables that appear as atoms in φ , let $\mathcal{L}_\varphi \subset \mathcal{L}$ be the fragment of \mathcal{L} using only atoms from χ_φ , and extend φ to a maximal consistent set $\Gamma \subset \mathcal{L}_\varphi$. We construct an SEM M satisfying all of Γ , and, in particular, φ .

Let us first give t . Define an irreflexive relation \prec on χ_φ by $X_1 \prec X_2$ if any instance of the schema $[\alpha_1 \wedge X_1]\ell_2 \wedge [\alpha_1 \wedge \neg X_1]\neg\ell_2$ is in Γ , where ℓ_2 is an X_2 -literal. By consistency with **Rec** there is a total order consistent with \prec . Assign the least element to $t = 0$, and iteratively remove least elements to obtain a t injective on χ_φ .

Now we may give the structural equations. Note that for every $X \in \chi_\varphi$ and $\alpha \in \mathcal{L}_{\text{int}} \cap \mathcal{L}_\varphi$ there is exactly one formula of the form $[\alpha]\ell$, where ℓ is an X -literal, in Γ : the forward direction of **F/D** shows there is at most one, and the backward direction shows there is one. Thus form a vector $v_\alpha : \chi \rightarrow \{0, 1\}$ by $v_\alpha(X) = 0, 1$ if $[\alpha]\neg X, [\alpha]X$ are in Γ respectively, and $v_\alpha(X) = 0$ if $X \notin \chi_\varphi$. Suppose α includes exactly the variables $X' \in \chi_\varphi$ for which $t(X') < t(X)$. Let α, \dots represent an extension of α to all of χ ; then define f_X on any such extension as $f_X(\alpha, \dots) = v_\alpha(X)$. For all $X \notin \chi_\varphi$, let $f_X(\cdot) = 0$ and $t(X) = 0$.

It remains to prove that M satisfies every formula of Γ , for which it suffices to show that $M \models [\alpha]X$ iff $v_\alpha(X) = 1$. To this effect we prove that $v(X) = v_\alpha(X)$, where $\alpha(M) \models v$, by induction on the time $t(X)$. If $X \in \alpha$ then by **RW**, **R** we have $v(X) = v_\alpha(X) = \alpha(X)$; thus suppose $X \notin \alpha$. Consider an α' that includes the value $v(X') = v_\alpha(X')$, where the equality holds inductively, for each $X' \in \chi_\varphi$ such that $t(X') < t(X)$. Thus by **K** we have that $[\alpha']\alpha' \in \Gamma$. Suppose that v, v_α differ at X , e.g. $v(X) = 0$ but $v_\alpha(X) = 1$. Then $[\alpha]X \in \Gamma$ so by **C**, $[\alpha \wedge \alpha']X \in \Gamma$. By **C** and **Rec** we must have $[\alpha']X \in \Gamma$ also. But this term specifies f_X , and hence this contradicts that $v(X) = 0$. \square

Theorem 3. AX^+ is sound and complete with respect to the validities of \mathcal{L}^+ .

Proof. Soundness is again straightforward; we show a consistent $\varphi \in \mathcal{L}^+$ is satisfiable. Extend to a maximal consistent set Γ using only χ_φ -atoms. Define a relation \prec^+ by $X \prec^+ Y$ if $X \rightsquigarrow Y \in \Gamma$. By consistency with Rec^+ we have that there is a total order consistent with \prec^+ . Now, reproduce the construction from the proof of Thm. 2 on Γ . By consistency with Wit , \prec is consistent with \prec^+ , so we may extend them to the same total order, from which we obtain t . Let M be the model thus constructed.

Now, obtain a new model M' by extending M as follows. For every pair X, Y , introduce a new variable $Z_{X,Y}$ with $f_{Z_{X,Y}}(\cdot) = 0$ and $t(Z_{X,Y}) = 0$. It is possible to add such variables since χ is infinite. Now, modify the structural functions for each $Y \in \chi_\varphi$ as follows. Labelling the variables X' for which $X' \prec^+ Y$ as X_1, \dots, X_k :

$$f_Y(v) = \begin{cases} f_Y(v), & v(Z_{X',Y'}) = 0 \text{ for all } X', Y' \\ X_1, & v(Z_{X_1,Y}) = 1 \text{ and } v(Z_{X',Y'}) = 0 \\ & \text{whenever } (X', Y') \neq (X_1, Y) \\ \dots & \\ X_k, & v(Z_{X_k,Y}) = 1 \text{ and } v(Z_{X',Y'}) = 0 \\ & \text{whenever } (X', Y') \neq (X_k, Y) \\ 0, & \text{otherwise} \end{cases}$$

This modification is admissible because \prec^+ is consistent with t . We claim that M' satisfies all of Γ . Any $\psi \in \Gamma \cap \mathcal{L}$ is clearly satisfied: ψ is satisfied by M and in M' we have $v(Z_{X,Y}) = 0$ for all X, Y , so the structural equations are unmodified from those of M . Thus, suppose that an extended atom $X \rightsquigarrow Y \in \Gamma$. Consider an intervention α that holds $Z_{X,Y}$ to 1, and holds all other $Z_{X',Y'}$ to 0. Then under α , we have the modified structural equation $f_Y(v) = v(Y) = v(X)$, so $M' \models X \rightsquigarrow Y$.

Now suppose $X \not\rightsquigarrow Y \in \Gamma$. Suppose toward a contradiction that $M' \models X \rightsquigarrow Y$. Then there is an intervention α under which toggling X changes Y . If α does not set any $Z_{X',Y'}$ -variables, then the structural equations are the same in M, M' and we have that $M \models [\alpha \wedge X] \ell_Y \wedge [\alpha \wedge \neg X] \neg \ell_Y$, where ℓ_Y is a Y -literal, contradicting consistency with Wit . If α sets any $Z_{X',Y'}$ to 1 for $Y' \neq Y$, then the value of Y is fixed to 0. The only case remaining is when only $\alpha(Z_{X',Y}) = 1$ for some X' . If $X' = X$ then there is a contradiction since the modification was only made if $X \rightsquigarrow Y \in \Gamma$. Thus $X' \neq X$. The structural equation for Y then becomes $f_Y(v) = v(Y) = v(X')$; further, X' is fixed to 0 since $v(Z_{X',Y}) = 1$, so this is impossible. \square

Theorem 4. $AX^+ + \text{Trans}$ is sound and complete with respect to the validities of \mathcal{L}^+ over $\mathcal{M}_{\text{local}}$.

Proof. Soundness follows since causal influence in $\mathcal{M}_{\text{local}}$ is transitive, as we remarked before. To show completeness, as before, extend a given consistent $\varphi \in \mathcal{L}^+$ to a maximal consistent set Γ , and define \prec^+ by $X \prec^+ Y$ if $X \rightsquigarrow Y \in \Gamma$. By consistency with Rec^+ , Trans , \prec^+ is a strict partial order. Again reproduce the construction from the proof of Thm. 2; by consistency with Wit , \prec is consistent with \prec^+ and we extend them to the same total order, obtaining t and M .

First convert M to a model $M' \in \mathcal{M}_{\text{local}}$. In M , the structural function for a variable Y may be written as $f_Y(X_1, \dots, X_k)$, where $X_1, \dots, X_k \in \chi_\varphi$ are precisely the variables preceding Y . To form M' , for each X_i that does not satisfy $t(X_i) = t(Y) - 1$, add a chain of variables $W'_{X_i, t(X_i)+1}, \dots, W'_{X_i, t(Y)-1}$ such that $t(W'_{X_i, t}) = t$ and with structural equations $v(W'_{X_i, t(X_i)+1}) = v(X_i)$, $v(W'_{X_i, t}) = v(W'_{X_i, t-1})$ for all $t(X_i) + 1 < t < t(Y)$, and change the structural function for Y to $f_Y(W'_{X_i, t(Y)-1}, \dots, W'_{X_i, t(Y)-1}, \dots, W'_{X_k, t(Y)-1})$. It is clear that $M' \models \psi$ iff $M \models \psi$ for any $\psi \in \mathcal{L}^+$.

Now modify M' to obtain a model $M'' \in \mathcal{M}_{\text{local}}$ that we will show to satisfy all of Γ . To form M'' , for every pair X, Y such that $X \prec^+ Y$, add a chain $W_{X,Y, t(X)+1}, \dots, W_{X,Y, t(Y)-1}$ and a single variable $Z_{X,Y}$. These have the times $t(W_{X,Y, t'}) = t'$, and $t(Z_{X,Y}) = t(Y) - 1$; and the structural equations $v(W_{X,Y, t(X)+1}) = v(X)$, $v(W_{X,Y, t'}) = v(W_{X,Y, t'-1})$ for $t(X) < t' < t(Y)$, and $v(Z_{X,Y}) = 0$. Modify the structural equation for Y as follows, labelling the variables X' for which $X' \prec^+ Y$ as X_1, \dots, X_k :

$$f_Y(v) = \begin{cases} f_Y(v), & v(Z_{X',Y}) = 0 \text{ for all } X' \\ v(W_{X_1, Y, t(Y)-1}), & v(Z_{X_1, Y}) = 1 \text{ and} \\ & v(Z_{X', Y}) = 0, \text{ if } X' \neq X_1 \\ \dots & \\ v(W_{X_k, Y, t(Y)-1}), & v(Z_{X_k, Y}) = 1 \text{ and} \\ & v(Z_{X', Y}) = 0, \text{ if } X' \neq X_k \\ 0, & \text{otherwise} \end{cases}$$

The modification is admissible because all variables involved are assigned to $t(Y) - 1$. We claim that M'' satisfies all of Γ . Again, this holds for any \mathcal{L} -literal, so suppose $X \rightsquigarrow Y \in \Gamma$. Under an intervention that holds the $Z_{X,Y}$ to 1, we have that $v(Y) = v(W_{X,Y, t(Y)-1}) = \dots = v(W_{X,Y, t(X)+1}) = v(X)$ so that $M'' \models X \rightsquigarrow Y$.

Now suppose that $M'' \models X \rightsquigarrow Y$. The intervention witnessing this fixes either none, or one of the $Z_{X_i, Y}$ to

1. If it fixes none, then taking the values of all the χ_φ variables into an intervention gives a witness that $M \models X \rightsquigarrow Y$, so by Wit, $X \rightsquigarrow Y \in \Gamma$. Thus suppose it fixes just one $Z_{X_i, Y}$ to 1. This gives that $X_i \rightsquigarrow Y \in \Gamma$, with $t(X) \leq t(X_i) < t(Y)$. Now $M'' \models X \rightsquigarrow X_i$; we thus repeat the argument obtaining a sequence $X \prec^+ X'_1 \prec^+ \dots \prec^+ X'_k \prec^+ Y$. Since \prec^+ is transitive (from Trans) we finally obtain $X \rightsquigarrow Y \in \Gamma$. \square

Proposition 3. The satisfiability problem (determining whether a formula of $\mathcal{L}/\mathcal{L}^+$ has a model in the relevant class) in all three cases (Thms. 2, 3, 4) is NP-complete.

Proof Sketch. In each case the proof of completeness produces a small model, usable as a certificate: $O(n), O(n^2), O(n^3)$ variables have to be considered in the first, second, and third cases respectively. \square

3.3 Relation to Existing Work

Thm. 2 shows that the class of open-universe generative programs defined in §2.2 satisfies a natural set of axioms. In fact, the system AX encompasses all of the principles about counterfactual conditionals used in the complete identification algorithm of Shpitser and Pearl (2008). This, together with Thm. 1, lends at least some credence to the idea that these “merely implicit” causal models can be understood on a par with more familiar explicit causal representations such as SEMs.

The previous literature on causal conditionals has explored classes of axioms and models that are quite different from those considered here. For instance, in addition to studying the (finite) recursive structural equation models, Halpern (1998, 2000) axiomatizes the class of SEMs that have a unique solution (but may not be recursive), as well as the class of all SEMs built of arbitrary equations, which may in general lack solutions. Similarly, Zhang (2013) considers classes of SEMs with desirable *sets* of solutions. Some of the central axioms for these classes make reference to all variables of the signature and thus cannot be translated into the open-universe setting. More fundamentally, it is not evident which of these models have an adequate procedural interpretation.

An analogous generalization in our setting might be to consider simulation models that “crash” under certain interventions and fail to have a solution. If we allow ϕ (Defn. 4) to be only partial computable on $\chi \times F$, the construction from Thm. 1 shows that every partial SEM has an equivalent simulation model in this wider class; but simple counterexamples show the reverse direction does not hold. Logically, only the forward direction of axiom F/D remains sound, and we leave the question of axiomatizing this wider class for future work.

The relation \rightsquigarrow has been used as a defined symbol in previous work (e.g., Halpern 2000). Because the relevant definition again relies on reference to all variables in the signature, this is not possible in the open-universe setting. Thms. 3 and 4 are thus new in the present work.

4 Conclusion

We have identified two equivalent classes of models—one declarative, one procedural—formalizing the notion of an *open-universe causal model*. Both classes validate an intuitive and familiar set of principles about subjunctive conditionals and the relation of causal influence. This highlights an important class of *implicit* generative models that can plausibly be treated as genuine causal models, on a par with (an infinitary generalization of) computable, recursive) structural equation models. More detailed work is of course needed to identify concrete cases in which components of learned generative models support legitimate causal counterfactuals (see, e.g., Besserve et al. 2018 for progress on this question).

From an axiomatic perspective, it would be desirable to extend the present treatment to the full probabilistic setting, since, as remarked in §2.3, both classes of models can be augmented with a natural probabilistic source. Axioms for probabilistic formal systems are well studied (e.g., Fagin et al. 1990). In the direction of a fully formalized do-calculus (Pearl, 2009), one would like to embed an axiom system like AX into an appropriate probability calculus, and combine these with a logic of *direct* causal influence (a direct and probabilistic version of \rightsquigarrow), so that the do-calculus rules could be expressed and studied in a precise formal system. While the basic definitions would be clear for SEMs (and this could of course already be investigated for finite SEMs), the extension to simulation programs is less clear. Existing identifiability algorithms require specific assumptions about exogenous noise variables, e.g., that each has only two endogenous children (Shpitser and Pearl, 2008). Some work would need to be done to ensure that probabilistic programs (or Turing machines) satisfy analogous restrictions.

Finally, other extensions to the languages considered here would also be natural to investigate. We studied one higher-order relation, namely \rightsquigarrow , which involves quantification over an infinite domain (the space of interventions). One of the advantages of open-universe models is precisely that they enable reasoning beyond the propositional level. Thus, e.g., systems for reasoning about causal and counterfactual statements involving explicit quantification (as in the examples from §1) are easily motivated, and ought to be understood.

References

- Besserve, M., Shajarisales, N., Schölkopf, B., and Janzing, D. (2018). Group invariance principles for causal generative models. In *Proceedings of the Twenty-first International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Bingham, E., Chen, J. P., Jankowiak, M., Obermeter, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P., Horsfall, P., and Goodman, N. D. (2019). Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research*, 28:1–6.
- Bottou, L., Peters, J., Quiñero-Candela, J., Charles, D. X., Chikering, D. M., Portugaly, E., Ray, D., Simard, P., and Snelson, E. (2013). Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14:3207–3260.
- Carbonetto, P., Kisiński, J., de Freitas, N., and Poole, D. (2005). Nonparametric Bayesian logic. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI)*.
- de Raedt, L. and Kimmig, A. (2015). Probabilistic (logic) programming concepts. *Machine Learning*, 100(1):5–47.
- Dean, T. and Kanazawa, K. (1989). A model for reasoning about persistence and causation. *Computational Intelligence*, 5(2):142–150.
- Fagin, R., Halpern, J. Y., and Megiddo, N. (1990). A logic for reasoning about probabilities. *Information and Computation*, 87:78–128.
- Freer, C. E., Roy, D. M., and Tenenbaum, J. B. (2012). Towards common-sense reasoning via conditional simulation: Legacies of Turing in artificial intelligence. In Downey, R., editor, *Turing’s Legacy*. ASL Lecture Notes in Logic.
- Friedman, N., Halpern, J. Y., and Koller, D. (2000). First-order conditional logic for default reasoning revisited. *ACM Transactions on Computational Logic*, 1(2):175–207.
- Friedman, N., Murphy, K., and Russell, S. (1998). Learning the structure of dynamic probabilistic networks. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Galles, D. and Pearl, J. (1998). An axiomatic characterization of causal counterfactuals. *Foundations of Science*, 3(1):151–182.
- Ginsberg, M. L. (1986). Counterfactuals. *Artificial Intelligence*, 30:35–79.
- Gogate, V. and Domingos, P. (2016). Probabilistic theorem proving. *Communications of the ACM*, 59(7):107–115.
- Goodman, N. D., Mansinghka, V. K., Roy, D., Bonawitz, K., and Tenenbaum, J. B. (2008). Church: a language for generative models. In *Proceedings of the Twenty-fourth Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Halpern, J. Y. (1998). Axiomatizing causal reasoning. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Halpern, J. Y. (2000). Axiomatizing causal reasoning. *Journal of AI Research*, 12:317–337.
- Halpern, J. Y. (2013). From causal models to counterfactual structures. *Review of Symbolic Logic*, 6(2):305–322.
- Hamrick, J. B. (2019). Analogies of mental simulation and imagination in deep learning. *Current Opinions in Behavioral Sciences*, 29:8–16.
- Hytinen, A., Eberhardt, F., and Järvisalo, M. (2015). Do-calculus when the true graph is unknown. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Ibeling, D. and Icard, T. (2018). On the conditional logic of simulation models. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*.
- Icard, T. F. (2017). From programs to causal models. In Cremers, A., van Gessel, T., and Roelofsen, F., editors, *Proceedings of the 21st Amsterdam Colloquium*, pages 35–44.
- Janzing, D. and Schölkopf, B. (2010). Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*, 56(10):5168–5194.
- Kemp, C. (2012). Exploring the conceptual universe. *Psychological Review*, 119(4):685–722.
- Kraus, S., Lehmann, D., and Magidor, M. (1990). Non-monotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(2):167–207.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
- Lauritzen, S. L. and Richardson, T. S. (2002). Chain graphical models and their causal interpretations. *Journal of the Royal Statistical Society, B*, 64(3):321–361.
- Lewis, D. (1973). *Counterfactuals*. Harvard University Press.

- Li, J., Xu, K., Chaudhuri, S., Yumer, E., Zhang, H., and Guibas, L. (2017). GRASS: Generative recursive autoencoders for shape structures. *ACM Transactions on Graphics*, 36(4):1–14.
- Meek, C. and Glymour, C. (1994). Conditioning and intervening. *The British Journal for the Philosophy of Science*, 45:1001–1021.
- Milch, B., Marthi, B., Russell, S., Sontag, D., Ong, D. L., and Kolobov, A. (2005). BLOG: Probabilistic models with unknown objects. In *Proc. 19th IJCAI*, pages 1352–1359.
- Mohamed, S. and Lakshminarayanan, B. (2017). Learning in implicit generative models. arXiv:1610.03483v4.
- Pearl, J. (2009). *Causality*. CUP.
- Pearl, J. and Bareinboim, E. (2012). External validity: From do-calculus to transportability across populations. *Statistical Science*, 29(4):579–595.
- Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press.
- Pfeffer, A. and Koller, D. (2000). Semantics and inference for recursive probability models. In *Proc. 7th AAAI*, pages 538–544.
- Poole, D. (2003). First-order probabilistic inference. In *Proc. 18th IJCAI*.
- Richardson, M. and Domingos, P. (2006). Markov logic networks. *Machine Learning*, 62(1-2):107–136.
- Shpitser, I. and Pearl, J. (2008). Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research*, 9:1941–1979.
- Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press.
- Srivastava, S., Russell, S., Ruan, P., and Cheng, X. (2014). First-order open-universe POMDPs. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Strotz, R. H. and Wold, H. O. A. (1960). Recursive vs. nonrecursive systems: an attempt at synthesis. *Econometrica*, 28(2):417–427.
- Tran, D., Hoffman, M. D., Saurous, R. A., Brevdo, E., Murphy, K., and Blei, D. M. (2017). Deep probabilistic programming. In *International Conference on Learning Representations (ICLR)*.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford university press.
- Zhang, J. (2013). A Lewisian logic of causal counterfactuals. *Minds and Machines*, 23:77–93.