Causation in Flux

Alexander Bochman
Computer Science Department,
Holon Institute of Technology, Israel

Abstract

We provide a logical reformulation of the action description language C+ from (Giunchiglia et al. 2004) in the form of a dynamic causal calculus that possesses a (dynamic) nonmonotonic semantics, and describe a logical system of dynamic causal inference that constitutes a complete description of the logic that is adequate for this dynamic calculus.

Introduction

A wide range of action theories have been suggested in the AI literature, starting from direct representations in a first-order classical language (such as the situation calculus, event calculus and TAL) and ending with higher-order action description languages (Gelfond and Lifschitz 1998), as well as representations of actions in general formalisms of propositional dynamic logic (PDL) and linear temporal logic (LTL). In fact, one of the useful ways to cope with this diversity and, in particular, to provide for transferability results among different theories involves designing some further, still more general representations that subsume existing ones as special cases (see, e.g., (Thielscher 2011)).

Unlike general descriptions of temporal dynamics in various logical formalisms, action theories in AI have to deal, however, with two quite specific reasoning tasks, namely the *prediction* task (what are the results of a given sequence of actions from an initial state) and the *planning* task (what sequence of actions could lead from an initial state to a target goal state). These reasoning tasks immediately lead to a triple of famous problems, known as the frame, ramification and qualification problems (see (Shanahan 1997)). It was realized quite early that classical logic, taken by itself, encounters difficulties in resolving these problems that would yield an efficient and versatile representation of actions and changes in the world.

In recent years a dominant approach to solving these problems has been based, in one form or another, on causal reasoning. Given a set of action and causal rules describing the domain, the causal approach employs a distinction between facts that hold in a situation versus facts that are caused (explained) by other facts and the rules. The corresponding *causal closure assumption* (Reiter 2001) amounts to the requirement that all facts that hold in a situation should be either caused by other occurrent facts, or else preserve their truth-values in time (due to the accompanying *inertia principle*). A direct incorporation of such causal assertions into the language of the situation calculus has been proposed in (Lin 1995; 1996), and has been shown to provide a natural account of both the frame and ramification problems.

An elaborate formalization of the above principles of causal reasoning has been given in (Giunchiglia et al. 2004). The formalism of (Giunchiglia et al. 2004) is a multi-layered representation framework. As its top layer, it employs a (causal) action description language C+. Domain descriptions in this language are instantiated first by assigning temporal stamps i : to propositions and then by incorporating the resulting descriptions into an atemporal causal calculus. The models of the resulting causal theories are viewed then as intended models of the source, higher-level action descriptions.

In this study we will suggest a dynamic generalization of the causal calculus that will provide a direct logical description for the language $C+$. In addition, we will describe a logical system of dynamic causal inference that constitutes a concise logical framework for causal reasoning in dynamic domains.

The plan of the paper is as follows. After a brief overview of the original (atemporal) causal calculus and its use in the framework of (Giunchiglia et al. 2004), we will introduce a dynamic causal calculus as a direct reformulation of the action description language $C+$. As a second step, we will describe a logical formalism of dynamic causal inference that will be shown to constitute a precise logic for causal reasoning in action domains.

The Causal Calculus

Throughout this study, we will assume that our basic language is a classical propositional language with the usual connectives and constants $\{\wedge, \vee, \neg, \rightarrow, \mathbf{t}, \mathbf{f}\}$. $\vDash$ and Th will stand, respectively, for the classical entailment and the associated closure operator. In what follows, we will also identify propositional interpretations ('worlds') with the sets of propositional formulas that hold in them.

A *causal rule* is a rule of the form $A \Rightarrow B$, where A and B are classical propositions. We will informally interpret such rules as saying plainly *"A causes B"*.[1]

[1](Giunchiglia et al. 2004) adopted a more cautious informal

By a *causal theory* we will mean an arbitrary set of causal rules. For a set u of propositions and a causal theory Δ, we will denote by $\Delta(u)$ the set of all propositions that are caused by u in Δ, that is,

$$\Delta(u) = \{A \mid B \Rightarrow A \in \Delta, \text{ for some } B \in u\}$$

Then the nonmonotonic semantics of a causal theory can be defined as follows.

Definition 1. A world (= propositional interpretation) α is an *exact model* of a causal theory Δ if it is a unique model of $\Delta(\alpha)$. The set of exact models forms a *nonmonotonic semantics* of Δ.

The above semantics of causal theories coincides, in effect, with the semantics for such theories, described in (Giunchiglia et al. 2004) (see also (McCain and Turner 1997)). It can also be verified that exact models of a causal theory are precisely the worlds that satisfy the following fixed-point condition:

$$\alpha = \mathrm{Th}(\Delta(\alpha)).$$

Accordingly, exact worlds are propositional models that are not only closed with respect to the causal rules, but also such that any proposition that holds in them is caused (that is, explained) ultimately by other propositions.

Causal Inference Relations

Though the original causal calculus has been defined only semantically, (Bochman 2004) has described a logical formalism of causal inference relations, which has been shown to provide a complete formalization of logical (monotonic) reasoning in causal theories. From a logical point of view, causal inference relations are defined as sets of causal rules that are required to satisfy almost all the usual postulates of classical inference, except Reflexivity $A \Rightarrow A$. The latter feature has turned out to be essential for an adequate representation of causal reasoning.

Definition 2. A *causal inference relation* is a relation $\Rightarrow$ on the set of propositions satisfying the following conditions:

(Strengthening) If $A \vDash B$ and $B \Rightarrow C$, then $A \Rightarrow C$;

(Weakening) If $A \Rightarrow B$ and $B \vDash C$, then $A \Rightarrow C$;

(And) If $A \Rightarrow B$ and $A \Rightarrow C$, then $A \Rightarrow B \wedge C$;

(Or) If $A \Rightarrow C$ and $B \Rightarrow C$, then $A \vee B \Rightarrow C$;

(Cut) If $A \Rightarrow B$ and $A \wedge B \Rightarrow C$, then $A \Rightarrow C$;

(Truth) $\mathbf{t} \Rightarrow \mathbf{t}$;

(Falsity) $\mathbf{f} \Rightarrow \mathbf{f}$.

The rule Or permits reasoning by cases; this feature can be seen as one of the main advantages of causal reasoning as compared with, say, default logic. It indicates that the causal logic is an *objective* (extensional) logical system, a system of reasoning about the world. In this respect, it is similar to classical logic, and distinct from modal (intensional) formalisms that deal primarily with beliefs and knowledge.

Yet another important feature of causal inference stems from the validity of the following rule:

(Coherence) If $A \Rightarrow \neg A$, then $A \Rightarrow \mathbf{f}$.

The above rule says that if a proposition causes propositions that are incompatible with it, then it is causally inconsistent. This feature indicates, in effect, that the above notion of causal inference is *atemporal*. For example, the rule $p \wedge q \Rightarrow \neg q$ cannot be understood as saying that p and q jointly cause $\neg q$ (afterwards) in a temporal sense; instead, by Coherence it implies $p \wedge q \Rightarrow \mathbf{f}$, which means, in effect, that $p \wedge q$ cannot hold. Just as in classical logic, however, a representation of temporal domains in this formalism can be obtained by adding explicit temporal arguments to propositions; this is what has been actually done in the action description framework of (Giunchiglia et al. 2004).

A possible worlds semantics

A semantic interpretation of causal inference relations can be given in terms of ordinary possible worlds (or Kripke) models (W, R, V), where W is a set of possible worlds, R a binary accessibility relation on W, and V a function assigning each world a propositional interpretation. Intuitively, $R\alpha\beta$ means that α is an initial state, and β a possible output state of a causal process. A Kripke model is *quasi-reflexive* if it satisfies the condition that if $R\alpha\beta$, then $R\alpha\alpha$.

Definition 3. A rule $A \Rightarrow B$ is *valid* in a Kripke model (W, R, V) if, for any worlds α, β such that $R\alpha\beta$, if A holds in α, then B holds in β.

By a set of causal rules *determined* by a Kripke model we will mean the set of rules that is valid in it. It can be verified that such a set satisfies all the postulates of causal inference. Moreover, as for other modal formalisms, a suitable construction of a canonical semantics allows us to obtain the corresponding completeness result:

Proposition 1. *A set of causal rules forms a causal inference relation if and only if it is determined by some quasi-reflexive Kripke model.*

As a by-product, the above semantics immediately sanctions a simple modal representation of causal rules. Namely, let $\Box$ be the usual modal operator definable in a possible worlds model: $\Box A$ holds in α iff A holds in all β such that $R\alpha\beta$. Then the validity of $A \Rightarrow B$ in a possible worlds model is equivalent to validity of the formula $A \to \Box B$. Consequently, causal rules are representable by modal formulas of the latter form. As a matter of fact, this modal representation has actually been used in many approaches to formalizing causation in action theories (see, e.g., (Geffner 1990; Turner 1999; Giordano, Martelli, and Schwind 2000; Zhang and Foo 2001)).

The nonmonotonic semantics of causal inference

Causal inference relations are just a special kind of causal theory, so they also possess a nonmonotonic semantics. Moreover, due to the logical properties of causal inference, the description of this nonmonotonic semantics can be simplified as follows.

To begin with, we extend causal rules to rules having arbitrary sets of propositions as premises: given a causal inference relation $\Rightarrow$ and an arbitrary set u of propositions,

reading of such rules, namely *"If A holds, then B is caused"*.

$u \Rightarrow A$ will be taken to hold if, for some finite $a \subseteq u$, $\bigwedge a \Rightarrow A$ belongs to $\Rightarrow$. $\mathcal{C}(u)$ will denote the set of propositions caused by u, that is $\mathcal{C}(u) = \{A \mid u \Rightarrow A\}$. Then a world α is an *exact world* of a causal inference relation if and only if

$$\alpha = \mathcal{C}(\alpha).$$

Given an arbitrary causal theory Δ, we will denote by $\Rightarrow_\Delta$ the least causal inference relation that includes Δ. Now, it has been shown in (Bochman 2003) that Δ has the same nonmonotonic semantics as $\Rightarrow_\Delta$, which means that the rules of causal inference are adequate for reasoning with respect to the nonmonotonic semantics of causal theories. Moreover, it has been shown that causal inference relations constitute in this respect a maximal such logic (see (Bochman 2004) for details).

An overview of C+

Being restricted to the level of propositions, the action description language C+ is based on three kinds of propositional atoms. More precisely, propositional atoms are partitioned first into action atoms and fluent atoms, while the latter are further partitioned into simple and statically determined fluents. However, if we will ignore for a moment these distinctions, action descriptions in C+ involve only the following two kinds of rules, where A, B, C are classical propositional formulas:

- *Static laws* and *action dynamic laws*, which are expressions of the form

$$\textbf{caused } B \textbf{ if } A;$$

- *Fluent dynamic laws* - expressions of the form

$$\textbf{caused } B \textbf{ if } A \textbf{ after } C.$$

An *action description* in C+ is defined as a set of such causal laws.

Interpretations and models of action descriptions in C+ are defined, however, indirectly by transforming them into plain causal theories. To begin with, for every natural number m, an action description D is transformed into an atemporal causal theory D_m as follows. First, 'time stamps' i: for $i \in \{0, \ldots, m\}$ are inserted in front of every occurrence of every atom in propositional formulas. Then any static law or an action dynamic law is translated into the following set of causal rules, for every $i \leq m$:

$$i : A \Rightarrow i : B,$$

where $i : F$ is the result of inserting i: in front of every occurrence of every atom in a formula F. Similarly, any fluent dynamic law is translated into the following set of causal rules:

$$(i : C) \wedge (i + 1 : A) \Rightarrow i + 1 : B,$$

for every $i < m$.

Finally, in order to deal with initial states, for every simple fluent literal l, the following causal rules are added to the resulting causal theory:

$$0 : l \Rightarrow 0 : l.$$

These rules make simple literals exogenous (self-explained) in the initial state. As a result, we obtain an ordinary causal theory, and the exact models of this theory are considered to be the models of the original action description. Such models can be visualized as histories of length m of the source dynamic domain.

(Giunchiglia et al. 2004) contains also a more general semantic construction, according to which an action description D in C+ describes, in effect, a *transition model* in which states are the models of the 'minimal' (static) causal theory D_0, while transitions correspond precisely to the models of the minimal 'dynamic' causal theory D_1. It has been shown in (Giunchiglia et al. 2004, Proposition 8) that, for any $m > 0$, models of a causal theory D_m are exactly histories (paths) of length m in this transition model.

Dynamic Causal Calculus

In this section we are going to provide a logical reformulation of the action description language C+ and its semantics. As a first step, we will use a more convenient, 'logic-oriented' notation, namely, we will uniformly rewrite both static and action dynamic laws of the form **caused** B **if** A as plain (static) causal rules $A \Rightarrow B$.a static law **caused** B **if** A as a familiar causal rule $A \Rightarrow B$, while a dynamic law **caused** B **if** A **after** C will be written as a rule of the form

$$C.A \Rightarrow B.$$

We will call such rules *dynamic causal rules*. Moreover, extending our previous re-interpretation of atemporal causal rules, we will assign a more 'active' informal reading to such rules, namely *"After C, A causes B"*.

Remark. As a matter of fact, the above dynamic causal rules are somewhat ambiguous from a syntactic point of view. On a most abstract level, such a rule can be viewed simply as an instantiation of a primitive ternary propositional operator. There are, however, at least two other, more articulated possibilities. Thus, a dynamic causal rule could be viewed as a plain causal rule $A \Rightarrow B$ that is conditioned by a preceding context C. In fact, this 'parsing' agrees with the informal reading of such rules, given above. Furthermore, we will see below that this understanding of dynamic causal rules as conditional static rules provides a natural justification for the postulates of the associated dynamic causal inference that will be given below. Still, a different possibility consists in viewing such rules as ordinary, binary causal rules with complex premises consisting of pairs of propositions (C, A). Again, we will see later that this reading can also be given a formal support, due to a possible translation of such rules as propositions of the form $(C \circ A) \rightarrow B$ in arrow logic.

In the version of the dynamic causal calculus that we will present in this study, we will make one further step and identify the static rules $A \Rightarrow B$ with dynamic rules of the form $\mathbf{t}.A \Rightarrow B$. This will make the above dynamic causal rules the only kind of rules of the dynamic calculus. The consequences and variations created by this identification will be discussed below.

As before, a *dynamic causal theory* will be defined as a set of dynamic causal rules. As a next step, we are going to provide a direct description of the nonmonotonic semantics of such causal theories. The guiding principle behind this nonmonotonic semantics will be a thorough enforcement of the *principle of universal causation*, according to which any state of a dynamic model should be explained (i.e., caused) by preceding states and causal rules.

Given a dynamic causal theory Δ and worlds α, β, we will denote by $\Delta(\alpha.\beta)$ the set

$$\{C \mid A.B \Rightarrow C \in \Delta \text{ for some } A \in \alpha, B \in \beta.\}$$

Definition 4. • A pair (α, β) of worlds will be called an *exact transition* with respect to a dynamic causal theory Δ if β is the unique model of $\Delta(\alpha.\beta)$, that is

$$\beta = \mathrm{Th}(\Delta(\alpha.\beta)).$$

• An *exact transition model* of a dynamic theory Δ is a set of worlds $\mathfrak{I}$ such that, for any $\beta \in \mathfrak{I}$ there is $\alpha \in \mathfrak{I}$ such that (α, β) is a an exact transition wrt Δ.

An exact transition is a transition between two states in which the resulting state is fully explained (caused), given the preceding state and the causal laws of the domain. It is important to note that if (α, β) is an exact transition, then the output world β is always closed with respect to the static rules (for our definition of the latter). Moreover, it can be easily verified that it will be a state in accordance with the definition of (Giunchiglia et al. 2004).

Remark. As a matter of fact, our definition of an exact transition almost coincides with the corresponding definition of a *causally explained transition*, given in (Giunchiglia and Lifschitz 1998) for a more restricted action description language C, a predecessor of C+. In fact, the only difference between the two definitions is that (Giunchiglia and Lifschitz 1998) required further that both the initial and resulting states of such a transition should be closed with respect to the static causal laws. On our construction, this additional requirement is accounted for, respectively, as a by-product of our definition of static causal rules on the one hand (for the resulting states), and an exact transition model on the other hand (for the initial states).

An exact transition model of a dynamic causal theory is defined above as a set of states in which every state is caused as a result of some exact transition. Consequently, any state of this model will be a state in the sense of (Giunchiglia et al. 2004), so any exact model in our sense will correspond to a model (transition system) in the sense of (Giunchiglia et al. 2004). Still, our definition of the dynamic semantics of causal theories is apparently more restrictive, since it requires that any state of the model, *including the initial one*, should be an output of some transition, whereas (Giunchiglia et al. 2004)) adopted more relaxed requirements on initial states. Roughly, it required only that such states should be closed with respect to the static laws (which are completely separated from dynamic ones) and, in addition, that any statically determined fluent literal that holds in a state should be caused by the static laws.

It seems that for a proper assessment of the above discrepancy, we should distinguish two aspects of the difference, conceptual and practical. On the conceptual side, we believe that an ultimate reason for imposing even the above minimal restrictions on initial states stems from a broad requirement that such states should be somehow accessible in accordance with the laws of the domain. In other words, any state of a dynamic system should be consistently viewed as a result of some legitimate transition (including possible 'loops' in this state). Speaking more generally, we contend that static laws and constraints should be viewed as constraints that are effective after every legitimate transition, and vice versa, any constraint that happens to hold hold after any possible transition should be considered as a static law of the domain.

Turning to the practical side of the difference, it turns out that for a broad class of dynamic descriptions (including all the examples given in (Giunchiglia et al. 2004)), we can *guarantee* in advance that if a state satisfies the abovementioned 'static' constraints of C+, it can always be constructed as an output of some exact transition. Taking only one simple example, if we will assume that α is an arbitrary state of what has been called a simple domain in (Giunchiglia et al. 2004), then it can be 'reconstructed' as a result of an exact transition (α_1, α), where α_1 has the same fluent literals as α, and no action atom holds in α_1. This transition will be exact because simple literals in α will then be caused by inertia rules, statically determined literals will be caused by the corresponding static rules, while the action atoms will be self-explainable (exogenous). More complex cases (e.g., non-inertial fluents) may require more elaborate constructions, including, if necessary, adding some new auxiliary actions to the vocabulary. We will postpone further discussion and elaborations on this to another occasion. Still, the main conclusion that we would like to make at this point is that, for a large class of action descriptions, our semantics coincides, in effect, with the semantics of (Giunchiglia et al. 2004). However, in contrast to the action description language C+, our dynamic calculus imposes no syntactic restrictions on the occurrences of fluent or action atoms in formulas appearing in dynamic causal rules.

Dynamic Causal Inference

The framework of the dynamic causal calculus, described in the preceding section, is a typical example of a nonmonotonic formalism; conclusions that can be obtained on the basis of the exact transition semantics can change nonmonotonically if we add some further facts or causal rules to the original dynamic causal theory. Still, as with other formalisms for nonmonotonic reasoning (see (Bochman 2011)), the causal rules of the dynamic causal calculus presuppose a certain underlying logic that agrees with the above nonmonotonic semantics. Such a logic will provide us with a formal description of the associated dynamic causal inference.

By a *dynamic causal inference relation* we will mean a set of dynamic causal rules of the form $A.B \Rightarrow C$ that satisfies the conditions described below.

The first group of postulates states that a set of dynamic causal rules with a fixed first premise (D) should satisfy the postulates of an 'ordinary' causal inference:

(Strengthening) If $A \vDash B$ and $D.B \Rightarrow C$, then $D.A \Rightarrow C$;

(Weakening) If $D.A \Rightarrow B$ and $B \vDash C$, then $D.A \Rightarrow C$;

(And) If $D.A \Rightarrow B$ and $D.A \Rightarrow C$, then $D.A \Rightarrow B \wedge C$;

(Or) If $D.A \Rightarrow C$ and $D.B \Rightarrow C$, then $D.A \vee B \Rightarrow C$;

(Cut) If $D.A \Rightarrow B$ and $D.A \wedge B \Rightarrow C$, then $D.A \Rightarrow C$;

(Truth) $\mathbf{t}.\mathbf{t} \Rightarrow \mathbf{t}$;

(Falsity) $\mathbf{t}.\mathbf{f} \Rightarrow \mathbf{f}$.

In view of the above postulates, dynamic causal rules $C.A \Rightarrow B$ can be seen as ordinary, binary causal rules $A \Rightarrow B$ that are conditioned by the preceding context C.

The next two postulates describe the logical properties of this preceding context in dynamic causal rules:

(Left-Str) If $A \vDash B$ and $B.D \Rightarrow C$, then $A.D \Rightarrow C$;

(Left-Or) If $A.D \Rightarrow C$ and $B.D \Rightarrow C$, then $A \vee B.D \Rightarrow C$.

The combined effect of the above pair of postulates is that the associated semantic interpretation of dynamic causal inference (described in the next section) will be again a kind of a possible world semantics, in which both the two premises and conclusion of a dynamic causal rule are evaluated with respect to worlds (complete states).

Finally, the last postulate is a formal expression of the requirement that any state that can be an input of some transition, is also an output state of at least one transition:

(Transition) If $\mathbf{t}.A \Rightarrow \mathbf{f}$, then $A.\mathbf{t} \Rightarrow \mathbf{f}$.

Recall that we have decided to identify static causal laws $A \Rightarrow B$ with dynamic causal laws of the form $\mathbf{t}.A \Rightarrow B$. Then the above postulate can be rewritten as

(Transition1) If $A \Rightarrow \mathbf{f}$, then $A.\mathbf{t} \Rightarrow \mathbf{f}$.

On this reformulation, the above postulate stipulates, in effect, that any input state of a consistent transition should be (statically) causally consistent. Combined with the other postulates, this will immediately imply that both the input and output state of a transition should be closed with respect to the valid static laws.

A possible worlds semantics

A possible worlds semantics of dynamic causal relations can be obtained by generalizing an accessibility relation on possible worlds to ternary relations.

A *causal possible world model* of dynamic causal inference is a triple (W, R, V), where W is a set of possible worlds, R a ternary accessibility relation on W, and V a function assigning each world a propositional interpretation. The accessibility relation will be required to satisfy the following two conditions:

(Quasi-reflexivity) If $R\alpha\beta\gamma$, then $R\alpha\beta\beta$.

(Transition) If $R\alpha\beta\beta$, then $R\delta\alpha\alpha$, for some $\delta \in W$.

Definition 5. A rule $A.B \Rightarrow C$ is *valid* in a model (W, R, V) if, for any worlds α, β, γ such that $R\alpha\beta\gamma$, if A holds in α and B holds in β, then C holds in γ.

Given the above definition of validity, it is easy to verify the following

Lemma 2. *The set of dynamic causal rules valid in a causal possible world model forms a dynamic causal inference relation.*

Moreover, using a suitable construction of a canonical semantics for a dynamic causal inference relation, the following completeness result can be established:

Theorem 3. *A set of dynamic causal rules forms a dynamic causal inference relation if and only if it is determined by a causal possible world model.*

Proof. (A sketch) Due to the connection between dynamic causal rules and the original, atemporal causal rules, the proof is a relatively straightforward generalization of the corresponding completeness proof for causal inference relations, given in (Bochman 2004, Theorem 7.4). More precisely, given a dynamic causal relation $\Rightarrow$, we can construct the corresponding canonical model (W, R_c) by taking W to be the set of all maximal consistent sets of propositions, and defining R_c as follows:

$$R_c\alpha\beta\gamma \equiv C(\alpha.\beta) \subseteq \beta \cap \gamma.$$

Notice that this definition directly implies quasi-reflexivity of R_c. Moreover, the use of the Transition postulate allows us to prove the transition property of R_c. Finally, it can be shown that $A.B \Rightarrow C$ holds for the source dynamic causal relation if and only if it is valid in (W, R_c). $\square$

Remark. One of the interesting consequences of the above semantic characterization of dynamic causal inference is that, similarly to a straightforward modal translation of ordinary causal rules as formulas of the form $A \to \Box B$, dynamic causal rules can be represented as formulas of *arrow logic* (see, e.g., (Venema 1997)). As a matter of fact, one of the principal motivations behind arrow logic has also consisted in providing an abstract description of dynamic (transition) models (cf. (van Benthem 1994)). Moreover, semantic interpretation of arrow logic is also based on a possible world semantics with a ternary accessibility relation, and it can be easily verified that, by the above semantic description, a dynamic causal rule $A.B \Rightarrow C$ turns out to be equivalent to a formula

$$A \circ B \to C$$

of arrow logic, where $\circ$ is a binary 'arrow conjunction' operator having the following semantic interpretation: $A \circ B$ holds in a world α if and only if there are worlds β, γ such that $R\beta\gamma\alpha$, A holds in β and B holds in γ.

Correspondences

Recall that a dynamic causal theory is an arbitrary set of dynamic causal rules. For any dynamic causal theory Δ there exists a least dynamic causal inference relation that includes Δ. We will denote it by $\Rightarrow_\Delta$. Clearly, $\Rightarrow_\Delta$ is the set of all dynamic causal rules that can be derived from Δ using the postulates of dynamic causal inference.

As before, we will extend the notation of dynamic causal rules to sets of propositional formulas in premises: for sets u, v of propositional formulas, $u.v \Rightarrow A$ will be taken to hold if $\bigwedge a.\bigwedge b \Rightarrow A$, for some finite $a \subseteq u, b \subseteq v$. In addition, $\mathcal{C}(u.v)$ will denote the set of propositions $\{A \mid u.v \Rightarrow A\}$.

Due to the logical properties of a dynamic causal inference relation, the definition of an exact transition can now be simplified, namely a pair of worlds (α, β) will be an *exact transition* with respect to a dynamic causal inference relation if and only if

$$\beta = \mathcal{C}(\alpha.\beta).$$

Then the following key result of this study shows, in effect, that the logic of causal dynamic inference is adequate for reasoning with respect to the exact semantics of dynamic causal theories, since it preserves the latter.

Theorem 4. *The exact transition models of a dynamic causal theory Δ coincide with the exact transition models of $\Rightarrow_\Delta$.*

Proof sketch. It can be shown that if C_Δ is the provability operator of $\Rightarrow_\Delta$, then, for any worlds α, β, if $C_\Delta(\alpha.\beta)$ is a consistent set, then it coincides with $\mathrm{Th}(\Delta(\alpha.\beta))$. Consequently, $\alpha = C_\Delta(\alpha.\beta)$ iff $\alpha = \mathrm{Th}(\Delta(\alpha.\beta))$, and therefore exact transitions of Δ will coincide with exact transitions of $\Rightarrow_\Delta$. Hence the result. □

Moreover, it can be shown that the logic of causal dynamic inference constitutes a maximal logic that is adequate for reasoning with exact causal models.

Definition 6. Two dynamic causal theories Δ and Γ will be said to be *strongly equivalent* if, for any set Φ of causal rules, $\Delta \cup \Phi$ has the same exact models as $\Gamma \cup \Phi$.

Strongly equivalent theories are 'equivalent forever', that is, they are interchangeable in any larger causal theory without changing the nonmonotonic semantics. Consequently, strong equivalence can be seen as an equivalence with respect to the background monotonic logic of causal theories. And the next result shows that this logic is precisely the logic of dynamic causal inference.

Theorem 5. *Dynamic causal theories Δ and Γ are strongly equivalent if and only if $\Rightarrow_\Delta = \Rightarrow_\Gamma$.*

The above result states that dynamic causal theories are strongly equivalent if and only if each of them can be obtained from the other using the postulates of dynamic causal inference. Again, the proof of this result can be established as a certain 'dynamic' generalization of the corresponding proof for the original causal calculus that has been given in (Bochman 2004).

Summary and Perspectives

The primary objective of this study consisted in showing that causal reasoning in dynamic action domains can be given a direct and concise logical representation. The action description language C+ has turned out to be especially suitable for this purpose, first of all because it has been formulated explicitly as a causal language and, on the other hand, because of the wealth of representation capabilities of this language as a working action theory in AI that have been demonstrated in (Giunchiglia et al. 2004). Taken jointly with this practical support, the results of the present study strongly indicate that a theory of dynamic causal inference can be viewed as a self-subsistent logical theory that provides an adequate and comprehensive representation framework for reasoning in dynamic domains. The study creates also obvious incentives for broader questions about the role and scope of causation in commonsense reasoning, as well as in knowledge representation for AI.

Causation has always been one of the most discussed concepts in the philosophy of science. It is intimately related to practically all notions that are essential both for a commonsense and scientific view of the world, such as laws, counterfactuals, explanation and abduction. On the other hand, causation and related notions have shown to be extremely elusive and problematic concepts. Efforts of many philosophers and logicians in the past have been focused on a formal, logical explication of these notions, but the task has turned out to be surprisingly difficult. Furthermore, starting with David Hume, an influential line of philosophical thought has argued, in effect, that causation should be expelled from the language of Science and Logic.

In recent years, however, we witness a revival of interest in the concept of causation, accompanied with new, more practical, insights about its role in our reasoning. Most prominent in this respect is Pearl's theory of causal reasoning (Pearl 2000) and its applications in statistics, economics, cognitive and social sciences.

An important alternative source of the new understanding of causation and its role in our reasoning comes from Artificial Intelligence, especially from theories of action and change. Though existing causal theories in this area do not always agree with a commonsense view of causation, they have provided a working concept of causation that has turned out to be crucial for singling out intended models of commonsense action descriptions. Moreover, they have made especially vivid the fact that causal reasoning, that is, asking why and seeking explanations, is germane to our reasoning about the world. These theories have also made evident that, though causal reasoning includes an important logical part, it is *not reducible* to a plain logical derivation in some ingenious causal logic. Instead, causal reasoning should be viewed as an important case of general nonmonotonic (assumption-based) reasoning. Accordingly, the tools and formalisms of nonmonotonic reasoning should hopefully provide us with a more adequate understanding of the concept of causation itself.

Acknowledgments. I would like to thank Vladimir Lifschitz for instructive comments and suggestions. He has pointed, in particular, to the similarity between our definition of an exact transition and the corresponding definition of a causally explained transition, given in (Giunchiglia and Lifschitz 1998). I am also grateful to Joohyung Lee for discussing the essential aspects of the syntax and semantics of the language C+.

References

Bochman, A. 2003. A logic for causal reasoning. In *IJCAI-03, Proceedings of the Eighteenth International Joint Con-*

ference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003, 141–146. Acapulco: Morgan Kaufmann.

Bochman, A. 2004. A causal approach to nonmonotonic reasoning. *Artificial Intelligence* 160:105–143.

Bochman, A. 2011. Logic in nonmonotonic reasoning. In Brewka, G.; Marek, V. W.; and Truszczynski, M., eds., *Nonmonotonic Reasoning. Essays Celebrating its 30th Anniversary*. College Publ. 25–61.

Geffner, H. 1990. Causal theories for nonmonotonic reasoning. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 524–530. AAAI Press.

Gelfond, M., and Lifschitz, V. 1998. Action languages. *Electronic Transactions on Artificial Intelligence* 3:195–210.

Giordano, L.; Martelli, A.; and Schwind, C. 2000. Ramification and causality in a modal action logic. *J. Log. Comput.* 10(5):625–662.

Giunchiglia, E., and Lifschitz, V. 1998. An action language based on causal explanation: preliminary report. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 623–630. AAAI Press.

Giunchiglia, E.; Lee, J.; Lifschitz, V.; McCain, N.; and Turner, H. 2004. Nonmonotonic causal theories. *Artificial Intelligence* 153:49–104.

Lin, F. 1995. Embracing causality in specifying the inderect effect of actions. In *Proc. Int. Joint Conf. on Artificial Intelligence, IJCAI-95*, 1985–1991. Montreal: Morgan Kaufmann.

Lin, F. 1996. Embracing causality in specifying the indeterminate effects of actions. In *Proceedings AAAI-96*, 670–676.

McCain, N., and Turner, H. 1997. Causal theories of action and change. In *Proceedings AAAI-97*, 460–465.

Pearl, J. 2000. *Causality*. Cambridge UP.

Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamic Systems*. MIT Press.

Shanahan, M. P. 1997. *Solving the Frame Problem*. The MIT Press.

Thielscher, M. 2011. A unifying action calculus. *Artificial Intelligence* 175(1):120 – 141.

Turner, H. 1999. A logic of universal causation. *Artificial Intelligence* 113:87–123.

van Benthem, J. 1994. A note on dynamic arrow logic. In van Eijck, J., and Visser, A., eds., *Logic and Information Flow*. Cambridge, MA, USA: MIT Press. 15–29.

Venema, Y. 1997. A crash course in arrow logic. In Marx, M.; Pólos, L.; and Masuch, M., eds., *Arrow Logic and Multi-Modal Logic*. Stanford, CA, USA: Center for the Study of Language and Information. 3–34.

Zhang, D., and Foo, N. 2001. EPDL: A logic for causal reasoning. In *Proc Int. Joint Conf. on Artificial Intelligence, IJCAI-01*, 131–136. Seattle: Morgan Kaufmann.