Causal Theories of Action and Change

Norman McCain and Hudson Turner
Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712-1188, USA
{mccain,hudson}@cs.utexas.edu

Abstract

For many commonsense reasoning tasks associated with action domains, only a relatively simple kind of causal knowledge (previously studied by Geffner and Lin) is required. We define a mathematically simple language for expressing knowledge of this kind and describe a general approach to formalizing action domains in it. The language can be used to express ramification and qualification constraints, explicit definitions, concurrency, nondeterminism, and dynamic domains in which things change by themselves.

Introduction

It has always been clear that causal knowledge plays a central role in commonsense reasoning about actions. However, it has not always been clear what this role is, or that it cannot be played by non-causal knowledge as well. In the AI literature, this is evident in the use of state constraints for deriving the indirect effects of actions. Intuitively, a state constraint is a proposition that rules out certain states of the world as impossible but says nothing about causation. Attempts to use state constraints to infer the indirect effects of actions rely on the assumption that whatever follows from the explicitly described effects of an action and the state constraints is an indirect effect. Conceptually, this is a mistake—one which rests on a confusion of causal and non-causal grounds. Technically, it leads to unintuitive results.

In the AI literature, Judea Pearl (1988) has emphasized the importance of the distinction between causal and non-causal grounds in general default reasoning. Some of the difficulties encountered in using state constraints for determining the indirect effects of actions were recognized by Ginsberg and Smith (1988), Lifschitz (1990), and Lin and Reiter (1994). A number of authors have proposed to overcome these difficulties by replacing state constraints by representations of causal knowledge of one kind or another (Geffner 1990; Elkan 1992; Brewka & Hertzberg 1993; Baral 1995; Lin 1995; McCain & Turner 1995; Thielscher 1995).

In this paper, we define a mathematically simple language for expressing causal knowledge and describe a general approach to formalizing action domains in it.

What Kind of Causal Knowledge?

Suppose Γ is a set of true propositions about the present or past, and ϕ is a proposition about the future. Intuitively, in order to show that the prediction expressed by ϕ is true, it suffices to show that ϕ holds in every "causally possible world history" in which Γ is true. Other commonsense reasoning tasks associated with action domains—such as planning—can also be carried out on the basis of knowledge of the causally possible world histories.

Intuitively, in a causally possible world history every fact that is caused obtains. We assume in addition the *principle of universal causation*, according to which—in a causally possible world history—every fact that obtains is caused.[1] In sum, we say that a world history is causally possible if exactly the facts that obtain in it are caused in it.

It is clear that one does not need to know the causes of facts in order to determine whether a world history is causally possible. It is sufficient to know the conditions under which facts are caused. The differences between these two kinds of causal knowledge are reflected in the following sentence forms, where ϕ and ψ stand in place of sentences.

(i) *The fact that ϕ causes the fact that ψ.*
(ii) *Necessarily, if ϕ then the fact that ψ is caused.*

A sentence of form (i) describes a cause of the fact that ψ. A sentence of form (ii) describes only a condi-

[1]This rather strong philosophical commitment is rewarded by mathematical simplicity in the main definition of causal theories. Moreover, as we will see, in applications it is easily relaxed.

tion under which the fact that ψ is caused. Intuitively, it seems reasonable to suppose that (i) implies (ii), but not vice versa.[2]

Despite these differences in meaning, substituting knowledge of form (ii) for knowledge of form (i) would seem to have no effect on our determination of the causally possible world histories. We are therefore in a fortunate position; for many commonsense reasoning tasks associated with action domains, our knowledge of form (i) can be adequately represented in a simple logic of sentences of form (ii).

Causal knowledge of form (ii) was used by Geffner (1990) and, in a restricted form, by Lin (1995) in previous work on formalizing action domains. The present work owes a considerable debt to these authors.

The Language of Causal Theories

In this section, we define the language of causal theories as an extension of propositional logic.[3] We begin by considering further the motivating intuitions.

Recall that a causally possible world history is one in which exactly the facts that obtain are caused. Now assume that D is a complete description of the conditions under which facts are caused. (For example, D might be a set of sentences of form (ii).) In this case, we can say that a causally possible world history is one in which the facts that obtain are exactly those that are caused according to D. This formulation is the key to understanding the definitions that follow. It rests on two main assumptions: the principle of universal causation and the completeness of D.

Syntax

We begin with a language of propositional logic, whose signature is given by a nonempty set of atoms. (In application to formalizing action domains, the atoms will be taken to represent propositions about the values of fluents and the occurrences of actions at specific times.) By a *literal* we mean either A or $\neg A$, where A is an atom. We use the expressions *True* and *False* to stand for $(A \vee \neg A)$ and $(A \wedge \neg A)$ respectively, for some atom A.

[2]To see that (ii) does not imply (i), consider a domain in which there is a switch S that controls two lights, A and B. For any time t, S being closed at t causes both A and B to be on at t, and S being open at t causes both A and B to be off at t. So, intuitively, we know that: necessarily, if A is on at t then S is closed at t. Consequently, we also know that: necessarily, if A is on at t then B is caused to be on at t (which is easily rendered in form (ii)). However, A being on at t does not cause B to be on at t, so the corresponding sentence of form (i) is false.

[3]In (Lifschitz 1997), Vladimir Lifschitz reformulates the central definitions for the case of predicate logic.

By a *causal law* we mean an expression of the form

$$\phi \Rightarrow \psi \qquad (1)$$

where ϕ and ψ are formulas of the underlying propositional language. By the *antecedent* and *consequent* of (1), we mean the formulas ϕ and ψ, respectively. Note that (1) is not the material conditional $\phi \supset \psi$.

The intended reading of (1) is: *Necessarily, if ϕ then the fact that ψ is caused.* Often, but not always, we write (1) because we know something more, namely: *The fact that ϕ causes the fact that ψ.* In so doing, we substitute knowledge of form (ii) for knowledge of form (i), as previously discussed. The term "causal law" is suggested by this practice.

By a *causal theory* we mean a set of causal laws.

Semantics

We will identify an interpretation I for a propositional language with the set of literals L such that $I \models L$.

For every causal theory D and interpretation I, let

$$D^I = \{\, \psi \,:\, \text{for some } \phi,\ \phi \Rightarrow \psi \in D \text{ and } I \models \phi \,\}\,.$$

That is, D^I is the set of consequents of all causal laws in D whose antecedents are true in I. Intuitively then, D^I entails exactly the formulas that are caused to be true in I according to D.

Main definition.[4] Let D be a causal theory, and let I be an interpretation. We say I is *causally explained* according to D if I is the unique model of D^I.

We have the following easy alternative characterization. An interpretation I is causally explained according to D if and only if for every formula ϕ,

$$I \models \phi \text{ iff } D^I \models \phi\,.$$

Thus, we can say that I is causally explained according to D if and only if the formulas that are true in I are exactly the formulas caused to be true in I according to D. Intuitively then, the causally explained interpretations correspond to the causally possible world histories.

Another convenient alternative characterization is the following. An interpretation I is causally explained according to D if and only if

$$I = \{\, L \,:\, D^I \models L \,\}$$

where L is understood to stand exclusively for literals.

We say a formula ϕ is a *consequence* of a causal theory D if ϕ is true in every causally explained interpretation according to D.

[4]Herre and Wagner (1996) have independently defined *generalized logic programs* and their *strongly supported models*, which turn out to be essentially equivalent to causal theories and their causally explained interpretations.

Examples

Let D_1 be the causal theory (in the language with only the atom p) consisting of the causal law

$$p \Rightarrow p\,. \tag{2}$$

Take $I = \{p\}$. Notice that $D_1^I = \{p\}$. Since I is the unique model of D_1^I, I is causally explained according to D_1. No other interpretation is causally explained. Therefore, p is a consequence of D_1.

It is important that the language of D_1 contains only the atom p. If it contained an additional atom, no interpretation would be causally explained according to D_1.

Now, let D_2 be the causal theory obtained by adding to D_1 the causal law

$$\neg p \Rightarrow \neg p\,. \tag{3}$$

One easily checks that both $\{p\}$ and $\{\neg p\}$ are causally explained according to D_2. Thus, p is not a consequence of D_2. This shows that the consequence relation for causal theories is nonmonotonic.

As we have remarked, the definition of a causally explained interpretation reflects the principle of universal causation, by requiring that every formula true in I be caused in I according to D. Notice that by including causal laws (2) and (3) in a causal theory we effectively suspend the principle of universal causation with respect to facts about p. The ability to relax universal causation in this and similar ways is critical to the usefulness of causal theories.

Let D_3 be the causal theory (in the language with exactly the atoms p and q) consisting of (2) and (3) along with the causal law

$$\mathit{True} \Rightarrow q \equiv p\,. \tag{4}$$

Let $I = \{p, q\}$. Notice that $D_3^I = \{p,\ q \equiv p\}$. Since I is the unique model of D_3^I, I is causally explained according to D_3. Similarly, the interpretation $\{\neg p, \neg q\}$ is causally explained according to D_3. These are the only causally explained interpretations.

Notice that a causal theory that contained only (4) would have no causally explained interpretations.

Definitional Extension

We obtain a *definitional extension* D' of a causal theory D by adding a new atom A to the signature, and also adding an explicit definition of A—a causal law of the form

$$\mathit{True} \Rightarrow A \equiv \phi \tag{5}$$

where ϕ is a formula in the language of D. Causal theory D' is a *conservative extension* of D: that is, D and D' have the same consequences in the language of D. Moreover, we can replace any formula equivalent to ϕ by A anywhere in D', except in (5), without altering the set of causally explained interpretations. (These properties are easily verified.)

Example D_3 above is a definitional extension of D_2.

Literal Completion

Let D be a causal theory in which (i) the consequent of every causal law is a literal, and (ii) every literal is the consequent of finitely many causal laws. By the *literal completion* of D we mean the classical propositional theory obtained by an elaboration of the Clark completion method (Clark 1978), as follows: For each literal L in the language of D, include the formula

$$L \equiv (\phi_1 \vee \cdots \vee \phi_n)$$

where $\phi_1, \ldots, \phi_n$ are the antecedents of the causal laws with consequent L.

Proposition. *Let D be a causal theory that satisfies conditions* (i) *and* (ii) *above. The causally explained interpretations according to D are precisely the models of the literal completion of D.*

An easy proof appears in (McCain & Turner 1997).

Notice that the method of literal completion is applicable to examples D_1 and D_2.

Formalizing Actions in Causal Theories

When representing an action domain by a causal theory, it is convenient to describe the underlying propositional signature by means of three pairwise-disjoint sets: a set of *action names*, a nonempty set of *fluent names*, and a nonempty set of *time names* (corresponding to a segment of the integers). The atoms of the language are expressions of the forms a_t and f_t, where a, f, and t are action, fluent, and time names, respectively. Intuitively, a_t is true if and only if the action a occurs at time t, and f_t is true if and only if the fluent f holds at time t.

By a *fluent formula* we mean a propositional combination of fluent names. Given a fluent formula σ and a time name t, we will write σ_t to stand for the formula obtained from σ by simultaneously replacing each occurrence of each fluent name f by the atom f_t.

Suitcase Domain

We illustrate our approach by formalizing an action domain adapted from (Lin 1995) in which there is a suitcase with two latches, each of which may be in either of two positions, up or down. The suitcase is spring-loaded so that whenever both latches are in the up position the suitcase is caused to be open. We model the opening of the suitcase as a static effect (as Lin does); that is, we do not model a state of the domain in which both latches are up but the suitcase is not (yet) open.

To formalize the Suitcase domain we first choose action names and fluent names, as follows.

$Toggle(L_1)$: the action of toggling Latch 1
$Toggle(L_2)$: the action of toggling Latch 2
$Close$: the action of closing the Suitcase
$Up(L_1)$: the fluent that Latch 1 is up
$Up(L_2)$: the fluent that Latch 2 is up
$Open$: the fluent that the Suitcase is open

Next, we specify a set of time names. Here we choose time names corresponding to the nonnegative integers.

Given our choice of language, the Suitcase domain can be (partially) formalized by the following schemas, where l is a meta-variable ranging over $\{L_1, L_2\}$.[5]

$$Toggle(l)_t \wedge Up(l)_t \Rightarrow \neg Up(l)_{t+1} \tag{6}$$
$$Toggle(l)_t \wedge \neg Up(l)_t \Rightarrow Up(l)_{t+1} \tag{7}$$
$$Close_t \Rightarrow \neg Open_{t+1} \tag{8}$$
$$Up(L_1)_t \wedge Up(L_2)_t \Rightarrow Open_t \tag{9}$$

According to schemas (6) and (7), whenever a latch is toggled at a time t it is caused to be in the opposite state at time $t+1$. Schema (8) says that whenever the suitcase is closed at a time t it is caused to be not open at $t+1$. Schema (9) says that whenever both latches are up at a time t the suitcase is caused to be open also at t. Schemas (6)–(8) express "dynamic causal laws." Schema (9) expresses a "static causal law."

The causal theory (6)–(9) is an incomplete description of the Suitcase domain, because it does not represent sufficient conditions for certain facts being caused—specifically, facts preserved by inertia, facts about the initial situation, and facts about which actions occur (and when). In what follows we describe standard ways of augmenting a domain description in order to provide sufficient conditions for facts of each of these kinds being caused.

Explaining Action Occurrences

Normally, in formalizing an action domain we do not describe the causes of actions. This is not because we believe that the agent's actions are not caused, or that they are "self-caused", or that the agent has free will. (We may or may not believe such things; it does not matter.) Rather, we do not describe the causes of actions because (normally at least) they are irrelevant for the purposes of deliberation and planning.

Nevertheless, a causal theory must specify conditions that are sufficient for every fact in a causally explained interpretation to be caused, including facts about the occurrences and non-occurrences of actions. Accordingly, we augment our specification of a causal theory by the following schemas, where a is a meta-variable for action names.

$$a_t \Rightarrow a_t \tag{10}$$
$$\neg a_t \Rightarrow \neg a_t \tag{11}$$

According to schema (10), the occurrence of an action a at a time t is caused whenever a occurs at t. Schema (11) is similar. By these schemas we represent, in effect, that facts about action occurrences may be *exogenous* to the causal theory.

Explaining Facts at Time 0

Whether we think of time 0 as the moment at which the suitcase came into existence, or as simply an arbitrary moment during the "life" of the suitcase, it is clear that whatever in the real world causes the latches to be either up or down at time 0, and whatever it is that causes the suitcase to be either open or closed at time 0 (except in the case when both latches are initially up) lies outside of time as it is represented in our theory. Therefore, (except in the one case mentioned) we cannot hope to write causal laws that mention the real causes of these facts. Instead, we augment our specification of a causal theory by the following schemas, where f is a meta-variable for fluent names.

$$f_0 \Rightarrow f_0 \tag{12}$$
$$\neg f_0 \Rightarrow \neg f_0 \tag{13}$$

By these schemas we represent, in effect, that facts about the initial values of fluents may be exogenous to the causal theory.

Explaining Facts by Inertia

Let $\mathbf{I}$ be a set of fluent formulas. We express that the fluents designated by the fluent formulas in $\mathbf{I}$ are inertial by writing the following schema, where σ is a meta-variable ranging over $\mathbf{I}$.

$$\sigma_t \wedge \sigma_{t+1} \Rightarrow \sigma_{t+1} \tag{14}$$

According to schema (14), whenever a fluent designated in $\mathbf{I}$ holds at two successive times, its truth at the second time is taken to be caused simply by virtue of its persistence. For the Suitcase domain, we take $\mathbf{I}$ to be the set of all fluent names and their negations.[6]

[5] Given a time name t, the expression $t+1$ stands for the name of the successor of the number named by t.

[6] Thus, the inertia laws for the Suitcase domain can also be represented by the schemas $f_t \wedge f_{t+1} \Rightarrow f_{t+1}$ and $\neg f_t \wedge \neg f_{t+1} \Rightarrow \neg f_{t+1}$, where f is a meta-variable for fluent

Suitcase Domain (continued)

Schemas (6)–(14) express the complete causal theory D for the Suitcase domain.[7] Schemas (6)–(9) are domain dependent. We call the remaining schemas (10)–(14) the *standard schemas*. Intuitively, the role of the standard schemas is to exempt specific classes of facts from the principle of universal causation.

Let I be the interpretation characterized below.

$\bullet\ \neg Toggle(L_1)_0$	$\bullet\ \neg Toggle(L_1)_1$	$\bullet\ \neg Toggle(L_1)_2$	$\cdots$
$\bullet\ Toggle(L_2)_0$	$\bullet\ \neg Toggle(L_2)_1$	$\bullet\ \neg Toggle(L_2)_2$	$\cdots$
$\bullet\ \neg Close_0$	$\bullet\ \neg Close_1$	$\bullet\ \neg Close_2$	$\cdots$
$\bullet\ Up(L_1)_0$	$\bullet\ Up(L_1)_1$	$\bullet\ Up(L_1)_2$	$\cdots$
$\bullet\ \neg Up(L_2)_0$	$Up(L_2)_1$	$\bullet\ Up(L_2)_2$	$\cdots$
$\bullet\ \neg Open_0$	$Open_1$	$\bullet\ Open_2$	$\cdots$

Interpretation I specifies, for all actions a and times t, whether or not a occurs at t, and, for all fluents f and times t, whether or not f holds at t. (The ellipses indicate that after time 2 no action occurs and no fluent changes its value.) It is not difficult to see that I is causally explained according to D. The bullets indicate the literals at times 0, 1 and 2 that appear in D^I due to the standard schemas. The two literals not marked by bullets appear in D^I due to the schemas (7) and (9). The remaining literals (represented by the ellipses) appear in D^I due to the standard schemas (10), (11), and (14). The atoms of the form $Open_t$ ($t > 1$) also appear in D^I due to schema (9). Since D^I contains no other formulas, we have shown that $I = D^I$. Therefore, I is the unique model of D^I.

The following formula is a consequence of D.

$$Up(L_1)_0 \wedge Up(L_2)_0 \wedge Close_0 \supset Toggle(L_1)_0 \vee Toggle(L_2)_0$$

In general, when both latches are up, it is impossible to perform *only* the action of closing the suitcase; one must also concurrently toggle at least one of the latches. If this seems unintuitive, recall that we have chosen to model the suitcase being open as a static effect of the latches being up, so there is no time in any causally possible world history at which both latches are up and the suitcase is closed.

Additional Expressive Possibilities

We have demonstrated that causal theories can be used to represent some standard "complications" of action theories, such as indirect effects of actions, implied action preconditions and concurrent actions. Next we briefly describe a few of the additional expressive possibilities of our approach.

Ramification and Qualification Constraints

Ramification and qualification constraints, in the sense of Lin and Reiter (1994), are formalized by schemas of the forms

$$True \Rightarrow \sigma_t$$

$$\neg\sigma_t \Rightarrow False$$

respectively, where σ (the constraint) is a fluent formula. A similar result appears in (McCain & Turner 1995).

Nondeterministic Actions

The semantics of causal theories rests on the principle of universal causation, according to which every fact is caused. Intuitively, in the case of a nondeterministic action, there is no cause for one of its possible effects rather than another. We have already seen, however—in schemas (10)–(14)—that there are ways of effectively exempting facts from the principle of universal causation. We can use laws of a similar form to describe nondeterministic actions. For instance, coin tossing can be described (in part) as follows.

$$Toss_t \wedge Heads_{t+1} \Rightarrow Heads_{t+1} \qquad (15)$$

$$Toss_t \wedge \neg Heads_{t+1} \Rightarrow \neg Heads_{t+1} \qquad (16)$$

Intuitively, according to schemas (15) and (16), for every time t, $Toss_t$ renders $Heads_{t+1}$ exogenous.

Actions with Delayed Effects, Things that Change by Themselves, and Events

Because we refer explicitly to time points in our action descriptions, we may, if we wish, describe actions with delayed effects. We may also model things that change by themselves. This we can do simply by writing causal laws that relate fluents at different times, without mentioning any actions. Alternatively, we may explicitly introduce events, which, like actions, can be conceived to be the causes of change. Previously, we explained why we are normally uninterested in modeling the causes of actions, and instead wish to view them as exogenous to our causal theories. In the case of events, this may or may not be so.

names. In other cases, there may be "fluent literals" that do not designate inertial fluents and inertial fluents that are not designated by fluent literals. Hence, the greater generality of (14) is significant. As an example, consider a coin tossing domain, represented in a language with two fluent names, *Heads* and *Tails*. Suppose we wish to model three possible orientations of the coin: lying heads, lying tails, and standing on edge. Intuitively, for each of these orientations, the corresponding fluent is inertial. Therefore, the set **I** should contain the following formulas (or equivalent ones): *Heads*, *Tails*, and $\neg Heads \wedge \neg Tails$. Notice that the last formula, which holds exactly when the coin is standing on edge, is not equivalent to a literal. Notice also that, intuitively, the literal $\neg Heads$ does not designate an inertial fluent, since it holds when the coin is in either of two orientations. The same is true of $\neg Tails$ and $\neg(\neg Heads \wedge \neg Tails)$.

[7] We note in passing that one can do query answering and satisfiability planning (Kautz & Selman 1992) in the Suitcase domain by using the literal completion of D.

Concluding Remarks

We have defined a simple nonmonotonic formalism for expressing causal theories and a novel approach to formalizing action domains. The main assumptions underlying the formalism—universal causation and completeness—are already present in the causal framework of (McCain & Turner 1995).[8] In this earlier framework, causal knowledge is represented by inference rules. An inference rule ϕ/ψ can be used to represent the knowledge that: *Necessarily, if the fact that* ϕ *is caused, then the fact that* ψ *is caused.* Although this is a natural form of causal knowledge, in some ways inference rules are less expressive than causal laws. For example, using only inference rules, one cannot explicitly represent inertia and the exogenous role of facts.[9]

The form of knowledge represented in causal theories has been used previously by both Geffner and Lin.[10] The language of causal theories is simpler than the formalism of Geffner and in some ways (e.g., in allowing causal laws with non-literal consequents) more expressive than the formalism of Lin.[11] As a basis for formalizing action domains, the language of causal theories is noteworthy, in our opinion, for its combination of mathematical simplicity and expressiveness.

Acknowledgments

We are grateful to the participants of the Mini-Workshop on Causality (held at UT Austin, May 5–6, 1996) for helpful discussions. Our deepest thanks to our advisor, Vladimir Lifschitz. Special thanks also to Enrico Giunchiglia. This work was partially supported by National Science Foundation grant #IRI-9306751.

References

Baral, C. 1995. Reasoning about actions: non-deterministic effects, constraints, and qualification. In *Proc. of IJCAI-95*, 2017–2023.

Brewka, G., and Hertzberg, J. 1993. How to do things with worlds: On formalizing actions and plans. *Journal of Logic and Computation* 3(5).

[8]The relationship with (McCain & Turner 1995) is investigated in (McCain & Turner 1997).

[9]As demonstrated in (Turner 1997), these ideas *can* be represented in default logic, where the "causal laws" of (McCain & Turner 1995) correspond to justification-free default rules. It is interesting to note also that causal theories correspond in a simple way to prerequisite-free default theories (see (McCain & Turner 1997)).

[10]The relationship between causal theories and (Lin 1995) is also investigated in (McCain & Turner 1997) and in (Giunchiglia & Lifschitz 1997).

[11]Non-literal consequents are needed, for instance, to capture traditional ramification constraints. Recall that they are also used when introducing explicit definitions.

Clark, K. 1978. Negation as failure. In Gallaire, H., and Minker, J., eds., *Logic and Data Bases*. New York: Plenum Press. 293–322.

Elkan, C. 1992. Reasoning about action in first-order logic. In *Proc. of the 1992 Canadian Conf. on Artificial Intelligence.*

Geffner, H. 1990. Causal theories of nonmonotonic reasoning. In *Proc. AAAI-90*, 524–530.

Ginsberg, M. L., and Smith, D. E. 1988. Reasoning about action II: the qualification problem. *Artificial Intelligence* 35(3):311–342.

Giunchiglia, E., and Lifschitz, V. 1997. An action language based on causal logic. Unpublished manuscript.

Herre, H., and Wagner, G. 1996. Stable generated models of generalized logic programs (talk). Dagstuhl Seminar 9627, Disjunctive Logic Programming & Databases: Nonmonotonic Aspects, July 1–5.

Kautz, H., and Selman, B. 1992. Planning as satisfiability. In *Proceedings of the 10th European Conference on Artificial Intelligence.*

Lifschitz, V. 1990. Frames in the space of situations. *Artificial Intelligence* 46:365–376.

Lifschitz, V. 1997. A logic for causal reasoning. Unpublished manuscript.

Lin, F., and Reiter, R. 1994. State constraints revisited. *Journal of Logic and Computation, Special Issue on Actions and Processes* 4(5):655–678.

Lin, F. 1995. Embracing causality in specifying the indirect effects of actions. In *Proc. of IJCAI-95*, 1985–1991.

McCain, N., and Turner, H. 1995. A causal theory of ramifications and qualifications. In *Proc. of IJCAI-95*, 1978–1984.

McCain, N., and Turner, H. 1997. On relating causal theories to other formalisms. Unpublished manuscript.

Pearl, J. 1988. Embracing causality in default reasoning. *Artificial Intelligence* 35:259–271.

Thielscher, M. 1995. Computing ramifications by postprocessing. In *Proc. of IJCAI-95*, 1994–2000.

Turner, H. 1997. Representing actions in logic programs and default theories: A situation calculus approach. *Journal of Logic Programming* 31(1–3):245–298.