Published in Proceedings of the Ninth Biennial Conference of the Canadian Society for Computational Studies of Intelligence (CSCSI'92), pp.221-227. Vancouver, May 1992. Morgan Kaufmann Publishers.

Reasoning about Action in First-Order Logic*

Charles Elkan
Department of Computer Science and Engineering
University of California, San Diego
La Jolla, California 92093-0114

ABSTRACT: The main contribution of this paper is a method of axiomatizing actions and their effects in standard first-order logic in a way that solves the frame problem. It has long been thought that the frame problem makes monotonic logics inadequate for reasoning about action; this paper shows how to achieve the benefits of a nonmonotonic logic using a monotonic logic. The use of first-order logic permits considerable simplicity and a frame axiom that is explicit, fixed, and intuitive. Moreover, complex worlds of actions can be axiomatized: fluents can be momentary or dynamic, the preconditions of actions need not be fully known, and the ramifications of actions can be stated unambiguously.

1 Introduction

This paper discusses reasoning about actions and their effects using the ontology of states, actions, and fluents, known as the situation calculus. The main positive contribution of the paper is a method of axiomatizing actions in standard first-order logic that solves the frame problem. It has long been thought that this problem makes monotonic logics inadequate for reasoning about action. This paper shows that the introduction of a *cancels* predicate, in addition to the more usual *holds* and *causes* predicate, and the use of bidirectional implication as a minimization operator, effectively allows nonmonotonic reasoning about actions to be performed using first-order logic.

Section 2 below introduces the situation calculus and the frame problem using the Yale shooting problem as an example. Section 3 discusses criteria for evaluating candidate solutions to the frame problem, and Section 4 then presents our solution, which satisfies the criteria advanced previously. Section 5 explains extensions to the method to accommodate transient properties of the world, to solve the (representational) qualification problem, and to solve the ramification problem: how to state the indirect effects of actions in a modular fashion. Finally, Section 6 compares the methods of this paper to related work.

*The preparation of this paper was supported by a grant from the Powell Foundation and by Grant No. IRI-9110813 from the National Science Foundation.

2 The situation calculus and the frame problem

The most complex theories of action that have been treated formally in AI research involve multiple agents with differing beliefs and agents. in some of these theories the temporal properties of actions are also considered: they take time, they may be concurrent or overlap. Substantial epistemic and temporal reasoning requires a logic with a theory of time and knowledge built into its structure [Moore, 1985; Konolige, 1986; Cohen *et al.*, 1990]. Solving the frame problem inside such logics, which are typically modal, is beyond the scope of this paper.

This work reported here is directed towards reasoning about the actions of a single agent. At this level of detail, the traditional ontology is called the situation calculus [McCarthy and Hayes, 1969]. The categories of this ontology are *states*, *objects*, *fluents*, and *actions*. Briefly, a *state* is a snapshot of the world, an *object* is something whose identity is preserved across states, a *fluent* is a relationship between objects that holds in some states, and an *action* is an event that changes one state into another.

States, objects, fluents, and actions can all be represented as terms in a first-order language. For example, if a and b are constants designating two child's building blocks, then $stack(a,b)$ can designate the action of placing a on b, where $stack(\cdot,\cdot)$ is a function symbol. Further, if s_0 is a constant designating an an initial state, then $do(s_0, stack(a,b))$ designates the state resulting from performing the action $stack(a,b)$.[1] An example of a term representing a fluent is $on(a,b)$, which presumably means that block a is on top of block b. The standard approach to encoding the situation calculus in first-order logic uses a special *holds* predicate. The arguments of *holds* are a fluent and a state, as in $holds(on(a,b), do(s_0, stack(a,b)))$.

[1]The function symbol $do(\cdot,\cdot)$ is sometimes named $result(\cdot,\cdot)$. Contrary to the usual convention, the first argument of $do(\cdot,\cdot)$ is a state, and the second argument is an action. This convention allows the state resulting from a sequence of actions $a_1, a_2, \ldots a_k$ to be written perspicuously as $do(\cdots do(do(s_0, a_1), a_2), \cdots, a_k)$.

In the world of the Yale shooting problem [Hanks and McDermott, 1986][2] there are two objects, a person Fred and a gun, three fluents, *loaded*, *alive*, and *dead*, and three actions, *load*, *wait*, and *shoot*. The fluents and actions have their obvious meanings and the person and the gun are left unnamed because they are fixed.

Only the *load* and *shoot* actions have significant effects. Using the *holds* predicate these can be specified by the following axioms:

$$\begin{array}{l}\forall s \;\; holds(loaded, do(s, load)) \\ \forall s \;\; holds(loaded, s) \;\rightarrow\; holds(dead, do(s, shoot)) \\ \forall s \;\; holds(loaded, s) \;\rightarrow\; \neg holds(alive, do(s, shoot)) \\ \forall s \;\; holds(loaded, s) \;\rightarrow\; \neg holds(loaded, do(s, shoot)).\end{array}$$

It should be possible to solve inference problems of various types using an axiomatization of a world of actions and their effects. Two general classes of problems are those involving reasoning forward in time, and those requiring reasoning backwards. These can be called *projection* and *retrojection* problems respectively. More precisely, a projection problem is one where it is known what holds in an initial state, and the task is to draw conclusions about what holds after a certain sequence of actions. Retrojection problems are also loosely called explanation problems [Hanks and McDermott, 1987]. They require inferring what held before a given sequence of actions, given observations of what holds afterwards.

The original paper on the Yale shooting world discussed only a projection problem: given that Fred is alive to start with, is he dead in the state after loading the gun, then waiting, then shooting? Formally, given

$$holds(alive, s_0).$$

the task is to prove

$$holds(dead, do(do(do(s_0, load), wait), shoot)).$$

One's innocent hope is that the four axioms above about the effects of the *shoot* and *load* actions are sufficient. Unfortunately, simply stating the effects of actions is not enough. To reach the desired conclusion at least the extra axiom

$$\forall s \;\; holds(loaded, s) \;\rightarrow\; holds(loaded, do(s, wait)).$$

is needed. To answer other projection and retrojection questions, many further axioms are necessary:

$$\begin{array}{l}\forall s \;\; holds(alive, s) \;\rightarrow\; holds(alive, do(s, load)) \\ \forall s \;\; holds(alive, s) \;\rightarrow\; holds(alive, do(s, wait)) \\ \forall s \;\; holds(alive, s) \;\wedge\; \neg holds(loaded, s) \\ \qquad\qquad\qquad\qquad \rightarrow\; holds(alive, do(s, shoot))\end{array}$$

and so on.

In general, if the world consists of a different actions and e potential effects (that is, fluents), it seems that on the order of ae axioms are needed. Almost all these axioms seem redundant: they express the typical case where an action does not influence a property of the world. These axioms, of the form

$$\forall s \;\; holds(p, s) \;\wedge\; \Pi(s) \;\rightarrow\; holds(p, do(s, a)),$$

where $\Pi(s)$ is a conjunction of *holds* conditions all referring to the same state s, are called frame axioms. The frame problem is to escape the predicament of having to state so many axioms that should somehow be true by default.

3 Judging candidate solutions to the frame problem

The general issue behind the frame problem is to find a good way to state the relationships between actions and fluents in the situation calculus ontology. Papers on the frame problem generally propose a new logic and present examples of using that logic to express actions and effects. Most papers, however, do not discuss how to evaluate the candidate solutions to the frame problem that they advance. This section identifies three dimensions of adequacy along which these proposed solutions can be judged. For a more detailed and general discussion of these criteria, see [Elkan, 1991].

To constitute a successful general solution of the problem of axiomatizing worlds of actions and their effects, a logic and a methodology for using the logic should give rise to axiomatizations that are declarative, parsimonious, and localist. These criteria seem more or less independent and exhaustive, but we shall not defend that claim. Rather, the remainder of this section is an attempt to give each of the criteria a definition that can be tested as objectively as possible.

Declarative. A declarative method of axiomatizing actions is one that uses a logic with a well-defined semantics specified independently of the rules of inference of the logic. Methods for representing actions that rely on a procedural definition of entailment, such as the STRIPS formalism and its variants [Lifschitz, 1986], do not satisfy this criterion. The main reason that a well-defined semantics is desirable is that it allows one to obtain high confidence in the correctness of a set of axioms, by working out its consequences in two independent ways, model-theoretically and proof-theoretically, and checking that these correspond.

Parsimonious. Parsimony is relatively easy to define. Suppose there are a different relevant actions, and e different relevant effects. An axiomatization that uses $O(ae)$ space is not parsimonious. If each action influences at most d properties of the world, then a parsimonious axiomatization should use at most $O(ad)$ space.[3]

Localist. Ideally, each action-effect combination should be specified by a separate axiom, or by no axiom at all if it is a special case of some general frame axiom. However, it is somewhat arbitrary what is considered a separate axiom. In classical logics, two assertions can be viewed as separate axioms if the logical connective that

[2] It is worth stressing that the contribution of this paper is not yet another "solution" to the Yale shooting problem. The contribution is a general method for axiomatizing the effects of actions. The Yale shooting problem is used as an expository example simply because it is well-known. Note also that there is no attempt here to state all claims with the maximum degree of generality.

[3] A criterion of parsimony is implicit in [Shoham, 1986].

links them implicitly is 'and'; one can imagine privileging a different connective. Moreover, nonclassical logics often have a semantics that is entirely formal, but not compositional: the meaning of a set of axioms cannot be defined in terms of independent meanings for each axiom in the set. Any type of closed world assumption, for example, destroys compositionality. Without compositionality it is not clear why separate axioms are desirable. In addition, there is no necessary correlation between separate axioms and efficient inference. Many inference procedures preprocess input sets of axioms, and represent seemingly modular sets of axioms in unmodular ways. In view of the considerations just outlined, we shall admit as "localist" any axiomatization that only has to be changed in a fixed, small number of places when a new action is taken into account.

4 The general method of axiomatizing actions

Our approach to solving the frame problem uses three basic predicates: *holds*, *causes*, and *cancels*. The arguments of *holds* are a fluent and a state, as in $holds(on(a,b), do(s_0, stack(a,b)))$. The arguments of *causes* and *cancels* are an action, a state, and a fluent, as in

$$causes(stack(a,b), s_0, on(a,b)).$$

The single, fixed, general frame axiom is

$$\forall a,s,p \quad holds(p, do(s,a)) \leftrightarrow \quad (1)$$
$$causes(a,s,p) \vee (holds(p,s) \wedge \neg cancels(a,s,p)).$$

This axiom has a common sense interpretation. It states that a fluent holds in the state resulting from an action if and only if the action "causes" the fluent, or the fluent held before the action, and the action does not "cancel" it. Note that both the *causes* and *cancels* predicates have state arguments, so whether or not an action influences a fluent can depend on the state in which the action is taken.

Using the three basic predicates, the relationships of these fluents and actions can be specified by the following axioms:[4]

$$\begin{aligned}\forall a,s,p \quad & causes(a,s,p) \leftrightarrow & (2)\\ & (a = load \wedge p = loaded) \vee \\ & (a = shoot \wedge p = dead \wedge holds(loaded, s))\\ \forall a,s,p \quad & cancels(a,s,p) \leftrightarrow & (3)\\ & (a = shoot \wedge p = alive \wedge holds(loaded, s)) \vee \\ & (a = shoot \wedge p = loaded \wedge holds(loaded, s)).\end{aligned}$$

Axioms (1)–(3) describe the Yale shooting world. They can be used to solve many different inference problems, always using the standard semantics of first-order predicate calculus. For example, in all models of sentences (1)–(3) and $holds(alive, s_0)$, it is the case that

$$holds(dead, do(do(do(s_0, load), wait), shoot))$$

is true. The shooting problem is thus solved.

The so-called murder mystery is a retrojection problem, introduced in [Baker, 1989], The scenario is that Fred is alive initially but not after shooting the gun and then waiting:

$$holds(alive, s_0) \wedge$$
$$\neg holds(alive, do(do(s_0, shoot), wait)).$$

The question is to discover when Fred died, and whether the gun was initially loaded. Axioms (1)–(3) and the sentence immediately above entail

$$holds(loaded, s_0) \wedge cancels(shoot, s_0, alive).$$

The mystery is solved.

A natural question is whether all axiomatizations of the type described above allow all projection problems to be solved. The following theorem states this question precisely, and affirms that the answer is yes.

Theorem: Let the theory T consist of the frame axiom (1), a *causes* axiom of the form

$$\begin{aligned}\forall a,s,p \quad & causes(a,s,p) \leftrightarrow \\ & \vdots \\ & a = \alpha \wedge p = \pi \wedge \Pi_{\alpha,\pi} \\ & \vdots\end{aligned}$$

a *cancels* axiom of the form

$$\begin{aligned}\forall a,s,p \quad & cancels(a,s,p) \leftrightarrow \\ & \vdots \\ & a = \alpha \wedge p = \pi \wedge \Omega_{\alpha,\pi} \\ & \vdots\end{aligned}$$

and an initial state axiom of the form

$$\forall p \quad holds(p, s_0) \leftrightarrow$$
$$\ldots \vee p = \pi_k \vee \ldots$$

In these axioms each subformula $\Pi_{\alpha,\pi}$ and $\Omega_{\alpha,\pi}$ is a conjunction of preconditions that may involve local existential quantification. That is, each $\Pi_{\alpha,\pi}$ and $\Omega_{\alpha,\pi}$ is of the form

$$\exists \bar{x} \;\; holds(f_1(\bar{x}), s) \wedge \cdots holds(f_k(\bar{x}), s)$$

where each f_i is a fluent and $\bar{x}$ allows fluents to be parametrized. Then T is consistent, and for all fluents p and action sequences $a_1, a_2, \ldots a_k$, either

$$T \models holds(p, do(\cdots do(do(s_0, a_1), a_2), \cdots, a_k))$$

or

$$T \models \neg holds(p, do(\cdots do(do(s_0, a_1), a_2), \cdots, a_k)).$$

Proof: Induction over the length of the action sequence $a_1, a_2, \ldots a_k$. ∎

A planning problem is a projection problem where the action sequence is unknown. Given a goal fluent p, the

[4] Almost every paper on the Yale shooting problem uses a slightly different set of fluents. The three used here are from the original circumscriptive attempt to solve the problem [Hanks and McDermott, 1986]. Our solution does not depend on this choice of fluents.

task is to find an action sequence $\alpha_1, \alpha_2, \ldots \alpha_k$ for some k such that

$$T \models holds(p, do(\cdots do(do(s_0, \alpha_1), \alpha_2), \cdots, \alpha_k)).$$

A corollary of the theorem above is that planning is semi-decidable: for any set of goal fluents, if a plan exists that achieves them simultaneously, it can be found. Planning should in fact be efficiently decidable using axioms of the form proposed here, since these axioms correspond to the completion [Clark, 1978] of logic program clauses. For the particular axioms given as an example above, the corresponding logic program is

$$\begin{aligned}
causes(A, S, P) &\rightarrow holds(P, do(S, A)) \\
holds(P, S) \wedge \neg cancels(A, S, P) &\rightarrow holds(P, do(S, A)) \\
&causes(load, S, loaded) \\
holds(loaded, S) &\rightarrow causes(shoot, S, dead) \\
holds(loaded, S) &\rightarrow cancels(shoot, S, alive) \\
holds(loaded, S) &\rightarrow cancels(shoot, S, loaded).
\end{aligned}$$

Projection and planning problems can be reduced to query-answering against this pure PROLOG program. The state of the art in PROLOG implementation is about nine RISC cycles per logical inference [Mills, 1989]. One could in principle obtain a specialized reasoner running at a comparable speed for any world of actions and fluents axiomatized following our guidelines.

5 Solving further problems in axiomatizing actions

Incomplete knowledge. The qualification problem is that any formalization of a fragment of common sense knowledge is always incomplete. However many preconditions have been taken into account in the formal statement of a rule, further preconditions exist whose truth influences the truth of the consequent of the rule, so the formal rule should have even more antecedents. Given that one can never be sure that all relevant preconditions for an action have been mentioned, the issue from a knowledge representation perspective is to axiomatize actions and their effects in such a way that when new preconditions are discovered, they can be stated easily [McCarthy, 1977]. For example in the Yale shooting world, one should be able to add the condition that the gun cannot be loaded if it is locked away.

In our framework, the preconditions of actions can be stated in a modular way using a new predicate *prevented*. Here is an example:

$$\begin{aligned}
\forall a, s, p \quad & causes(a, s, p) \leftrightarrow \neg prevented(a, s) \wedge \\
& \qquad \vdots \\
& \quad (a = load \wedge p = loaded) \vee \\
& \qquad \vdots \\
\forall a, s \quad & prevented(a, s) \leftrightarrow \\
& \qquad \vdots \\
& \quad (a = load \wedge holds(lockedaway, s)) \\
& \qquad \vdots
\end{aligned}$$

When a new precondition for an action is discovered, the use of the *prevented* predicate allows it to be taken into account by changing an axiomatization in just one place.

Using a *prevented* predicate in the way just described allows an axiomatization that is discovered to be incomplete to be updated in a straightforward way. A second issue of incomplete knowledge is how to assert in an axiomatization that the preconditions of an action are only partially stated. Intuitively, the $\leftrightarrow$ connective in the axiom defining *prevented* says that *prevented* is fully defined: there are no a and s such that the truth or falsity of $prevented(a, s)$ is unknown. Conditionally replacing the $\leftrightarrow$ connective with a unidirectional $\leftarrow$ connective permits the truth of $prevented(a, s)$ to be partially unspecified. The following axioms say that $prevented(a, s)$ is true if $prevents(a, s)$ is true, and if $determined(a, s)$ is true, then $prevented(a, s)$ is true only if $prevents(a, s)$ is true. That is, *prevented* is fully specified for those a and s for which *determined* is true, and partially specified for all other a and s.

$$\begin{aligned}
\forall a, s \quad & prevented(a, s) \leftarrow prevents(a, s) \\
\forall a, s \quad & (prevented(a, s) \rightarrow prevents(a, s)) \\
& \qquad\qquad \leftarrow determined(a, s)) \\
\forall a, s \quad & prevents(a, s) \leftrightarrow \\
& \qquad \vdots \\
& (a = load \wedge holds(lockedaway, s)) \\
& \qquad \vdots \\
\forall a, s \quad & determined(a, s) \leftrightarrow \\
& \qquad \vdots \\
& \qquad a = load \\
& \qquad \vdots
\end{aligned}$$

Momentary and dynamic fluents. Once they start to hold, most properties of the world tend to continue to hold. The frame axiom (1) formalizes this observation. However some properties of the world tend to change by themselves. Fluents of this nature are called *momentary* [Lifschitz and Rabinov, 1989], and a flexible logic for reasoning about action must allow them to be decoupled from any frame axioms. The following modified axiom achieves this in our approach:

$$\begin{aligned}
\forall a, s, p \quad & holds(p, do(s, a)) \leftrightarrow \\
& \qquad causes(a, s, p) \vee \\
& (holds(p, s) \wedge \neg momentary(p) \wedge \neg cancels(a, s, p)).
\end{aligned}$$

This axiom is identical to (1), except a fluent p holding in a state s may cease to hold because it is momentary as well as because it is canceled. For an example of how to use this axiom, consider the loud noise presumably caused by a *shoot* action. This fluent can be named *bang* and axiomatized by adding a description of what induces it to the *causes* axiom:

$$\begin{aligned}
\forall a, s, p \quad & causes(a, s, p) \leftrightarrow \\
& \qquad \vdots
\end{aligned}$$

$$(a = shoot \wedge p = bang \wedge holds(loaded, s)) \vee$$
$$\vdots$$

and stating that it is momentary:

$$\forall p \;\; momentary(p) \;\leftrightarrow\; p = bang.$$

A dynamic fluent is a property of the world whose nature is to change spontaneously [Lifschitz and Rabinov, 1989]. For example, the level of water in a swimming pool that is being filled changes even after unrelated actions. Dynamic properties of the world are difficult to axiomatize in any logic that uses the situation calculus ontology, because they are typically time-dependent, and time is deliberately excluded from the situation calculus, in a tradeoff of tractability for expressiveness [Levesque, 1988]. However if one assumes that all actions take the same amount of time, for example, then dynamic fluents can be accommodated. One can write

$$\begin{aligned}
\forall a, s, p \quad & causes(a, s, p) \;\leftrightarrow \\
& \vdots \\
& \exists k \; p = level(pool, k) \;\wedge \\
& \quad holds(level(pool, k-1), s) \wedge holds(filling(pool), s) \\
& \vdots \\
\forall a, s, p \quad & cancels(a, s, p) \;\leftrightarrow \\
& \vdots \\
& \exists k \; p = level(pool, k) \;\wedge \\
& \quad holds(level(pool, k), s) \wedge holds(filling(pool), s) \\
& \vdots
\end{aligned}$$

The ramification problem. This problem is to axiomatize actions in such a way that their indirect effects can be stated independently. For example, suppose that the world of interest is a living room with two air vents. If both vents are blocked the room will be stuffy. One does not want to state this repeatedly as a potential effect of every action that may block a vent. The standard solution in the literature [Ginsberg and Smith, 1988; Baker and Ginsberg, 1989] to the ramification problem is to use so-called domain constraints that state relationships between fluents without mentioning actions at all. Following this approach one would write

$$\begin{gathered} holds(stuffy, s) \;\equiv \\ holds(blocked(vent_1), s) \wedge holds(blocked(vent_2), s). \end{gathered} \tag{4}$$

Axioms of this sort are acausal: they are silent as to what events in the world have the power actually to induce other events. Unfortunately they often lead to unwanted ambiguity. Suppose that

$$holds(blocked(vent_1), s_0) \;\wedge\; \neg holds(blocked(vent_2), s_0)$$

is true. This entails $\neg holds(stuffy, s_0)$. What can one conclude is true after an action a given the fact $causes(a, s_0, blocked(vent_2))$? All that can be deduced using (4) is

$$cancels(a, s_0, blocked(vent_1)) \vee causes(a, s_0, stuffy).$$

Intuitively, there should be no ambiguity: insufficient ventilation cannot normally cause a vent to become unblocked.. We want to conclude

$$causes(a, s_0, stuffy).$$

This ambiguity problem is discussed in [Ginsberg and Smith, 1988], but no satisfactory solution is presented. The ambiguity can be resolved in our approach, because we have a theory of causation that, however limited, is explicit. We can state that if certain fluents are caused or canceled, then other fluents are also caused or canceled, given perhaps that various preconditions hold. In the present example, we can write

$$\begin{aligned}
\forall a, s, p \quad & causes(a, s, p) \;\leftrightarrow \\
& \vdots \\
& (p = stuffy \wedge holds(blocked(vent_1), s) \\
& \qquad \wedge\; causes(a, s, blocked(vent_2))) \vee \\
& (p = stuffy \wedge holds(blocked(vent_2), s) \\
& \qquad \wedge\; causes(a, s, blocked(vent_1))) \vee \\
& \vdots
\end{aligned}$$

This axiom captures reality better: it entails $causes(a, s_0, stuffy)$ as wanted.

The real issue is that if the vents are blocked, that brings about stuffiness *causally*, whereas in the other direction, stuffiness only suggests that the vents are blocked *evidentially*. Causal and evidential implications have different logical behaviours [Pearl, 1988]. It does not work to use the same $\rightarrow$ connective to encode both, whether the connective is understood inside a monotonic or a nonmonotonic logic.

One could still object that stating causal relationships explicitly is less parsimonious. However this difference is unimportant: computationally an undirected axiom such as (6) has to be indexed to make it usable whenever an instance of any of its left- or right-hand-side conjuncts becomes true or false. If anything, inference is more efficient with a causal description of indirect effects.

6 Discussion

The simplicity of the ideas for reasoning about action introduced in this paper should not be misconstrued. Basing a solution to the frame problem on a weak, but explicit, theory of causality has been suggested before, notably in [Lifschitz, 1987]. The general idea is to specify what actions induce what effects, and to use the minimization machinery of a nonmonotonic logic to say that actions induce no other effects. There are two principal new ideas here: the introduction of a *cancels* predicate dual to the *causes* predicate, and the observation that bidirectional implication is a powerful minimization operator that can be used to obtain nonmonotonic effects.

Given any *fixed* information about a world of actions and their effects, all valid conclusions about that world can be obtained by monotonic inference using first-order logic. If *new* information about an action or an effect must be taken into account, then the truth of some conclusions may change. New information can therefore not be taken into account just by adding new axioms to the prior axiomatization. However, with an axiomatization constructed according to the method laid out above, new information can be taken into account by modifying the prior axiomatization in a straightforward, precisely specified, computationally easy way. Some nonmonotonic logics appear to allow new information to be taken into account just by adding new axioms, but in those logics, after new axioms are added, the overall minimization called for in the semantics of the logic must be performed again. This process is analogous to inserting new disjuncts in the right hand side of an $\leftrightarrow$ sentence, which is how new facts are included in the axiomatizations advocated in this paper.

An alternative approach to axiomatizing the effects of actions in first-order logic has recently been published by Reiter [Reiter, 1991], building on work by Pednault [Pednault, 1989] and Schubert [Schubert, 1989]. The methods of this paper are considerably simpler, for two main reasons. First, Reiter uses multiple frame axioms, each one involving quantification over actions and states. A single frame axiom is sufficient above, because it involves quantification over fluents in addition. Second, no distinction is made in this paper between the preconditions under which an action is possible and the preconditions which must hold for it to have a particular effect. The absence of this distinction makes the theorem of Section 4 much simpler than the corresponding theorem in [Reiter, 1991].

Some further contrasts between the techniques of this paper and alternative approaches to formalizing the effects of actions are worth mentioning. The TWEAK calculus of [Chapman, 1987] is inherently incapable of solving the ramification problem. Approaches based on a separate logical theory or model for each state of the world [Ginsberg and Smith, 1988; Rao and Foo, 1989], with actions corresponding to database updates, do not have clear frame axioms. Approaches based on circumscription [Lifschitz and Rabinov, 1989; Baker and Ginsberg, 1989] do not have practical proof procedures.

Acknowledgement: Many of the ideas here first arose in discussions with Ray Reiter.

References

[Baker and Ginsberg, 1989] Andrew B. Baker and Matthew L. Ginsberg. Temporal projection and explanation. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 906–911, August 1989.

[Baker, 1989] Andrew B. Baker. A simple solution to the Yale shooting problem. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, pages 11–20, 1989.

[Chapman, 1987] David Chapman. Planning for conjunctive goals. *Artificial Intelligence*, 32:333–377, 1987.

[Clark, 1978] Kenneth. L. Clark. Negation as failure. In Hervé Gallaire and Jack Minker, editors, *Logic and Databases*, pages 293–322. Plenum Press, New York, 1978.

[Cohen *et al.*, 1990] Philip R. Cohen, Hector J. Levesque, and Joseph Nunes. On acting together: Joint intentions for intelligent agents. In *Proceedings of the National Conference on Artificial Intelligence*, Boston, Massachusetts, August 1990.

[Elkan, 1991] Charles Elkan. Formalizing causation in first-order logic: Lessons from an example. In *Working Notes of the AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning*, Stanford University, March 1991.

[Ginsberg and Smith, 1988] Matthew L. Ginsberg and David E. Smith. Reasoning about action I: A possible worlds approach. *Artificial Intelligence*, 35(2):165–195, June 1988.

[Hanks and McDermott, 1986] Steve Hanks and Drew McDermott. Default reasoning, nonmonotonic logics, and the frame problem. In *Proceedings of the National Conference on Artificial Intelligence*, pages 328–333, August 1986.

[Hanks and McDermott, 1987] Steve Hanks and Drew McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33(3):379–412, November 1987.

[Konolige, 1986] Kurt Konolige. *A Deduction Model of Belief.* Pitman, 1986.

[Levesque, 1988] Hector J. Levesque. Logic and the complexity of reasoning. *The Journal of Philosophical Logic*, 17:355–389, 1988.

[Lifschitz and Rabinov, 1989] Vladimir Lifschitz and Arkady Rabinov. Things that change by themselves. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 864–867, August 1989.

[Lifschitz, 1986] Vladimir Lifschitz. On the semantics of STRIPS. In *Proceedings of the Workshop on Reasoning about Actions and Plans*, pages 1–9. Morgan Kaufmann Publishers, Inc., 1986.

[Lifschitz, 1987] Vladimir Lifschitz. Formal theories of action. In *Proceedings of the Workshop on the Frame Problem in Artificial Intelligence*, Lawrence, Kansas, 1987.

[McCarthy and Hayes, 1969] John McCarthy and Patrick J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, 1969.

[McCarthy, 1977] John McCarthy. Epistemological problems of artificial intelligence. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1038–1044, 1977.

[Mills, 1989] Jonathan W. Mills. A pipelined architecture for logic programming with a complex but single-cycle instruction set. In *Proceedings of the IEEE First International Tools for AI Workshop*, September 1989.

[Moore, 1985] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25(1):75–94, 1985.

[Pearl, 1988] Judea Pearl. Embracing causality in default reasoning. *Artificial Intelligence*, 35(2):259–271, June 1988.

[Pednault, 1989] Edwin P. D. Pednault. ADL: Exploringthe middle ground between STRIPS and the situation calculus. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, pages 324–332. Morgan Kaufmann Publishers, Inc., 1989.

[Rao and Foo, 1989] Anand S. Rao and Norman Y. Foo. Minimal change and maximal coherence: A basis for belief revision and reasoning about actions. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 966–971, August 1989.

[Reiter, 1991] Raymond Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, 1991.

[Schubert, 1989] Lenhart K. Schubert. Monotonic solution of the frame problem in the situation calculus: An efficient method for worlds with fully specified actions. In *Knowledge Representation and Defeasible Reasoning*, pages 23–67. Kluwer Academic Publishers, 1989.

[Shoham, 1986] Yoav Shoham. What is the frame problem? In *Proceedings of the Workshop on Reasoning about Actions and Plans*, pages 83–98. Morgan Kaufmann Publishers, Inc., 1986.