Nondeterministic Actions and the Frame Problem

Craig Boutilier
Department of Computer Science
University of British Columbia
Vancouver, BC V6T 1Z4, CANADA
Tel. (604) 822-4632
Fax. (604) 822-5485
cebly@cs.ubc.ca

Nir Friedman
Department of Computer Science
Stanford University
Stanford, CA 94305-2140, U.S.A.
Tel. (415) 725-8782
Fax. (415) 725-1449
nir@cs.stanford.edu

Abstract

We describe a logical system and methodology for the natural specification of nondeterministic actions. The logic combines elements of dynamic logic, process logic and the situation calculus and allows one to express alternative (*actual* and *possible*) paths or sequences of events. Our system permits a simple solution to the frame problem for nondeterministic actions that "completes" user-supplied theories of action. While drawing inspiration from Reiter's solution for the deterministic case, some of the main intuitions underlying this solution must be abandoned in the nondeterministic setting due to possible correlations among effects. We show our completion is unambiguous and faithful to our stated intuitions, and that in a deterministic setting our solution is equivalent to that of Reiter.

Keywords: Knowledge Representation, Reasoning about Actions, Nondeterministic actions, Frame problem.

This paper has not already been accepted by and is not currently under review for a journal or another conference. Nor will it be submitted for such during IJCAI's review period.

Nondeterministic Actions and the Frame Problem

Craig Boutilier — University of British Columbia
Nir Friedman — Stanford University

Abstract: We describe a logical system and methodology for the natural specification of nondeterministic actions. The logic combines elements of dynamic logic, process logic and the situation calculus and allows one to express alternative (*actual* and *possible*) paths or sequences of events. Our system permits a simple solution to the frame problem for nondeterministic actions that "completes" user-supplied theories of action. While drawing inspiration from Reiter's solution for the deterministic case, some of the main intuitions underlying this solution must be abandoned in the nondeterministic setting due to possible correlations among effects. We show our completion is unambiguous and faithful to our stated intuitions, and that in a deterministic setting our solution is equivalent to that of Reiter.

1 Introduction

One of the most important problems studied in AI is that of representing and reasoning about action and change. Yet, since the earliest attempts to formalize this problem, the straightforward encoding of actions and their effects has been fraught with difficulty. Roughly, given a description of the state of the world and some action, we want the ability to predict the new state after the action has been performed. Unfortunately, while our natural inclination is to specify actions in terms of those facts that change and leave unsaid those things unaffected, most logical systems are not tolerant of such implicit assumptions. The *frame problem* [14] is one of action representation: how can actions be specified in a compact and natural way; and how can a reasoning system "fill in the blanks," or treat unmentioned facts as unchanging.

A number of solutions have been proposed in the literature, including the use of nonmonotonic formalisms embodying the *default* assumption that *all* facts persist [13]. This principle of *minimization of change* allows one to infer that facts not explicitly mentioned as affected by an action are unaffected. However, as shown by Hanks and McDermott [4] anomalous behavior results when this principle is applied in the most straightforward way. Subsequent attempts to deal with these problems using nonmonotonic logics [9, 17, 1] have proven somewhat more successful through judicious application of this principle. However, all of these solutions have been shown to suffer from problems (e.g., see [7]).

1.1 Isolating the Frame Problem

Recent work seems to have adopted a clearer perspective on the problem. If the frame problem is indeed one of representation then it seems clear that a precise *representational methodology* must be tightly coupled with any reasoning mechanism one might put forth. Furthermore, one must clearly identify the nature of the problem being solved. Often minimization of change is proposed to deal with theories that have aspects of the frame problem as well as the qualification problem, the ramification problem, actions with defeasible effects and so forth. Successful resolution of these problems requires that they be isolated (conceptually) and that their solutions be studied independently. Should a single mechanism settle the score for each problem, it should be viewed as a happy coincidence; it should not be taken as an assumption from which investigations start.

This perspective has led to an increased emphasis on the representation of actions, and on solutions that can be shown to be correct with respect to restricted classes of action theories [11, 3, 6]. One particular approach that fits this mold is the model of Reiter [16], who suggests a syntactic transformation that "completes" action theories of a particular syntactic form. While isolating different aspects of reasoning about action has proven fruitful, a major question facing such a piecemeal approach is that of "scaling up": how will these solutions fare when additional features are added to the theories of action one is willing to entertain. In this paper we examine one such complicating factor, *nondeterministic actions*, and propose a solution in the spirit of Reiter's mechanism for deterministic actions.

Nondeterministic actions are actions, such as flipping a coin, that may lead to several possible outcomes. Such actions are *inherently* unpredictable[1] — all outcomes are *a priori* possible. For example, we cannot predict whether flipping a coin will result in heads or tails.

Reiter's approach is expressed in the framework of the situation calculus (SC), which is restricted to dealing with deterministic actions: an action a applied at situation s has only one possible outcome, namely *result*(s, a). Furthermore, the main intuitions underlying Reiter's approach are not directly applicable in nondeterministic contexts. More precisely, Reiter's method identifies, for each atomic proposition, the conditions under which it can change, and then asserts that these are the *only* conditions under which it can change. The most straightforward generalization to a nondeterministic setting suggests that we examine, for each proposition, the conditions under which it *might* change (i.e., *Heads* might change after *flip*) and then state that these are the only such conditions. Unfortunately, such a method fails due to the presence of *correlations* among action effects. For example, the action *force* (forcing a door) might have three possible outcomes when the door is closed: nothing happens (the door remains closed), the door opens without triggering the alarm, and the door opens and triggers the alarm. If we examine the possible outcomes, we note that *Open* can be either true or false after *force*, and similarly *Alarm*

[1] This is, of course, relative to the level of detail one is willing to model.

can be either true or false. This separation of effects seems to suggest that there are four possible outcomes after *force*. However, one of these, $Alarm \wedge \neg Open$, should not be possible due to the correlation between *Alarm* and *Open*.

The contribution of this paper is twofold. First, we espouse a general methodology for reasoning about action that allows one to semantically isolate different problems in action representation and solve them (in particular, defeasible effects, information-producing actions, qualifications and ramifications). Second, we provide a solution to the frame problem in the presence of nondeterministic actions, one that, because of the methodology adopted, seems especially intuitive. An important aspect of our treatment is the insistence on explicit representation of possible outcomes of actions. This allows us to distinguish nondeterministic actions from *indefinite* (or indeterminate) actions. We elaborate on both of these issues in the concluding section.

1.2 Outline

In this paper we describe a logical system and natural methodology for reasoning about actions, and a particular application of this methodology for dealing with nondeterministic actions. We approach the problem is several stages: we first propose a specific semantic interpretation of actions; we then present a language and logic for describing the relevant aspects of this semantic model; we must next identify a particular methodology for specification of action theories, and determine the role of a user's input; finally we provide a solution to the frame problem based on the prescribed semantics, language and user specification.

We first describe the semantic models used to interpret actions and action sequences. These models are essentially *transition systems*, familiar from the study of dynamical systems [12], stochastic processes [15], and dynamic logic [2]. Actions are mappings between states of the system, nondeterministic actions leading to several resulting states. Also of interest are *paths* or *trajectories* through the state space that describe possible evolutions of the world.[2]

Second, we present a language and a logic MPL that allows one to reason about nondeterministic actions — with several *possible* outcomes — and the properties of specific trajectories — describing an *actual* evolution of the system. This includes a description of the actual states along a path and the actions that occur. In addition, from any state one can express properties of *alternative*, unrealized trajectories rooted at that state — both those that result from different action choices and those that result from different (from the actual) outcomes of the actual action choices. The language of MPL thus combines aspects of SC, dynamic logic (PDL) and process logics [5] and bears a direct relationship to our semantic model.

Third, we must propose a specific methodology for, and examine the role of, a user's specification of actions. We introduce a language $\mathcal{A}^{\mathrm{ND}}$ for reasoning about

[2]We focus on discrete, nondeterministic transition systems; however, the intuitions underlying our methodology should be applicable to stochastic and continuous-time systems.

nondeterministic actions. (This language is somewhat reminiscent of the language $\mathcal{A}$ of Gelfond and Lifschitz [3]; we defer comparisons to Section 4.) The language $\mathcal{A}^{\mathrm{ND}}$ admits two types of statements: *action descriptions*, that describe the effects of actions; and *path descriptions*, that describe properties of the actual system trajectory, or course of events. We call a collection of action descriptions an *action theory*. Such a theory describes the possible effects of actions in different states of the system, and leads us to adopt a particular model of action that captures the intended system dynamics, or "physics" of the world in question. The chief desideratum for the specification of action theories is is the ability to only specify *changes* induced by an action, leaving "non-changes" implicit. The key feature of our proposal is our insistence that all outcomes be listed as explicit possibilities. Action specification in this form is quite natural, since only possible changes need to be specified. For example, the possible effects of flipping a coin can be specified as follows:

$$\mathit{flip}\ \textbf{causes}\ \mathit{Heads} \mid\mid \neg\mathit{Heads}\ \textbf{when}\ \mathit{HaveCoin} \tag{1}$$

$$\mathit{flip}\ \textbf{necessarily causes}\ \neg\mathit{HaveCoin}\ \textbf{when}\ \mathit{HaveCoin} \tag{2}$$

These actions descriptions state that if an agent has a coin and flips it, it might come up heads *and* it might come up tails (i.e., not heads) and that the agent will not have the coin after the toss. We take the *intent* of such a theory to describe a system in which a coin flip has exactly one of two possible outcomes when the agent has a coin, one where *Heads* is true and one where *Heads* is false. Furthermore, we intend that each of these possibilities completely characterizes the changes associated with the outcome in question: if $\neg$*Heads* results, we know only *Heads* and *HaveCoin* become false (other unmentioned propositions persist, or are unaffected by the flip). This second assumption is the usual assumption of persistence.

We note that such action descriptions are similar in intent to statements of PDL. Assertions (1) and (2) capture constraints imposed by the PDL sentence

$$\mathit{HaveCoin} \supset \langle \mathit{flip} \rangle\, \mathit{Heads} \wedge \langle \mathit{flip} \rangle\, \neg\mathit{Heads} \wedge [\mathit{flip}]\, \neg\mathit{HaveCoin}$$

However, such PDL assertions only assert what can change, not what *doesn't change*. We impose a much stronger interpretation on $\mathcal{A}^{\mathrm{ND}}$ assertions.

The set of path descriptions make up the second part of a specification of a particular problem and provide information about the *actual* execution of the system (for instance, in a prediction task we might express initial conditions and list the actions that occur.) For example, we might have the following statements:

$$\textbf{initially}\ \mathit{HaveCoin}$$

$$\mathit{Heads}\ \textbf{after}\ \mathit{flip}$$

which state that initially the agent had the coin and a flip resulted in heads. We discuss below how such "observations" are used in the reasoning process.

We will describe the formal semantics of $\mathcal{A}^{\mathrm{ND}}$ by describing the class of MPL-models that are *faithful* to a given action theory (i.e., capture the *intended* system dynamics of the theory). Finally, we describe how to solve the frame problem using our methodology. We take the frame problem to be that of concisely expressing the intended dynamics of an action theory, as described above. Given a compact and natural user specification of an action theory, we describe the construction of an MPL theory that is satisfied only by the intended (or faithful) models of the theory, thus providing the desired syntactic characterization. Our construction procedure draws much from Reiter's [16] *explanation closure*. However, because *correlations* of effects are possible in nondeterministic settings, Reiter's method is not directly applicable. The expressive power of MPL can be used to deal with this situation.

We present formal criteria reflecting the assumptions above, and show that our procedure results in theories satisfying these properties. We also show that our solution can be thought of as a generalization of Reiter's solution — both coincide on theories with deterministic actions.

2 Nondeterministic Transition Systems

Our semantics of action will be described in terms *nondeterministic transition systems*. A transition system consists of two main components: a set W of possible states, and a *transition function* τ that describes the possible successors of a state after executing an action in that state. We provide informal descriptions of the connectives of MPL and their semantics. Due to space limitations, precise details are deferred to the full version of this paper.

An *MPL-model* is a transition system with additional function π that maps each world to a truth assignment over primitive propositions. We evaluate formulae with respect to *paths* in a model, or trajectories through the state space arising through the execution of some sequence of actions. These paths have the form:

$$x = w_0 \stackrel{a_0}{\rightarrow} w_1 \stackrel{a_1}{\rightarrow} w_2 \ldots$$

Such a path describes the evolution that starts at w_0 and passes through states $w_1, w_2, \ldots$ as actions $a_0, a_1, \ldots$ are executed. We note that such a path corresponds to the *actual* occurrence of the actions in question having the *actual* outcomes $w_1, w_2, \ldots$ listed. It does not rule out the fact that other actions *might* have been executed, or that the actual actions *might* have had different outcomes. We define *First*(x) as w_0, the initial state in the path, and use $x \cdot y$ to denote the *concatenation* of two paths x and y (assuming the last state of x and the first of y coincide).

The language $\mathbf{L}_{MPL}$ is constructed using the usual classical modal connectives together with the connectives $\langle\langle\alpha\rangle\rangle$ and E, where α is any program term.[3] Formulae are evaluated with respect to a path in a recursive manner. Loosely, atomic propositions are evaluated according to the truth assignment at *First*(x). Thus, formulae that do not involve modal operators describe the first state of the path. The $\langle\langle\alpha\rangle\rangle$

[3]A program term is a program in the sense of PDL [2]; for our purposes here, it suffices to treat α as a sequence of actions $a_1 \ldots a_n$.

modality describes the remainder of the path. Roughly, $\langle\langle\alpha\rangle\rangle\varphi$ is true at a path x if $x = y \cdot z$ where y is a path corresponding to some execution of α, and φ holds at z. Intuitively, $\langle\langle\alpha\rangle\rangle\varphi$ is true just when α is executed and *actually* results in an outcome at which φ is true. The E modality describes *alternative paths* that start from the same initial state. The formula $\mathsf{E}\varphi$ holds at x just when there exists a path y, such that *First*$(y) =$ *First*(x) and φ holds at y. Finally, we say that a model M satisfies a theory T if every path in M satisfies all the formulae in T.

We can define the usual dynamic logic modalities in MPL. We define $\langle\alpha\rangle\,\varphi$ as $\mathsf{E}\langle\langle\alpha\rangle\rangle\varphi$, and $[\alpha]\,\varphi$ as $\neg\langle\alpha\rangle\,\neg\varphi$. It is easy to verify that $\langle\alpha\rangle\,\varphi$ holds at path x just when some execution of α at *First*(x) could lead to a path where φ holds, and that $[\alpha]\,\varphi$ holds at x when all executions of α from *First*(x) lead to paths where φ hold. We note that the truth of certain formulae, in particular $\mathsf{E}\phi$, $\langle\alpha\rangle\,\phi$ and $[\alpha]\,\phi$ depend only on the current state, *First*(x), and not the current path x.

3 The Frame Problem and its Solution

In this section, we describe how $\mathcal{A}^{\mathrm{ND}}$ is used to specify nondeterministic actions in a natural and compact fashion and specify the intended semantics of our language using MPL-models. We then introduce a procedure that, given an action theory, constructs an MPL-theory that captures this semantics. In particular, it deals successfully with the frame problem.

3.1 Nondeterministic Action Specification

To describe the dynamics of a given domain, we assume that each action a is specified by a set of statements in the following natural form:

$$\begin{array}{l} a \textbf{ causes } \rho^a_{1,1} \,||\, \ldots \,||\, \rho^a_{1,k^a_1} \textbf{ when } D^a_1 \\ \ldots \\ a \textbf{ causes } \rho^a_{n,1} \,||\, \ldots \,||\, \rho^a_{n,k^a_n} \textbf{ when } D^a_n \end{array} \tag{3}$$

Each D^a_i is an arbitrary (consistent) proposition and each $\rho^a_{i,j}$ is a (consistent) conjunction of literals. We require, for each action a, that $n, k^a_i > 0$ and that $D^a_i \vdash \neg D^a_j$ if $i \neq j$. For any set of actions $\mathcal{A}$, a theory consisting of a set of such axioms for each $a \in \mathcal{A}$ is dubbed an *action theory for* $\mathcal{A}$. An action theory is *complete for* a if $\vdash \vee\{D^a_i\}$; and it is *complete for* $\mathcal{A}$ if it is complete for each $a \in \mathcal{A}$.

Intuitively, the conditions D^a_i are *discriminants* that describe the various circumstances under which action a can have different possible effects. For example, the action of dropping an object has quite different effects if the object in question is fragile or not, so one should describe one set of effects relative to the proposition *fragile* and another relative to $\neg$*fragile*. These discriminants must be disjoint and, in our initial presentation, we assume they are exhaustive (i.e., the action theory is complete), deferring the general case to Section 3.3. Each proposition $\rho^a_{i,j}$ describes a *possible effect* of action a under condition D^a_i. Intuitively, a particular occurrence of a could potentially cause any of these possible effect propositions to hold. Because we only consider definite action specifications in this paper, possible effects are conjunctions of literals. As we describe below, their effect on the state of the world is unambiguous and (under our assumptions) completely known. For instance, Axiom (1) asserts that should an agent be holding a coin, flipping that coin has two possible outcomes, heads or tails. (See [10] for a probabilistic action representation similar to this.) In the full version of the paper, we elaborate on an additional type of statement that describes *necessary* effects of actions, such as Axiom (2). Roughly, our treatment adds the necessary effects to each of the possible effects. In our example, we add the literal $\neg$*HaveCoin* to each of the two possible effects in Axiom (1).

Since possible effects are conjunctions of literals, we will sometimes treat a term $\rho^a_{i,j}$ as if it were a set of literals rather than a conjunction. Each formula of the form (3) is called an *action clause* and describes the possible effects of action a under condition D^a_i. We use $atm(a,i,j)$ to denote the set of atoms occurring among the literals in $\rho^a_{i,j}$, and $atm(a,i)$ to denote $\cup_j atm(a,i,j)$.

3.2 Closure of Action Theories

Given a complete action theory, one would like to ask certain queries about the effects of actions or properties of action sequences. Although action theories seem to express the desired information, we must make precise the effect of an action (or lack thereof) on every proposition in order to make complete predictions.

Our approach is to provide an *interpretation* of action theories in MPL. The interpretation $I(T)$ is an MPL theory, which has precise semantics. We want to ensure that this interpretation fully captures our intuitions, namely, that the possible effects listed are the *only* possible effects; and each effect describes the *only* propositions that change. Such an interpretation allows the user to specify action theories without circumscribing all the possibilities explicitly.

We make this intuition precise by defining the class of models that are *faithful* to an action theory. Intuitively, each world w satisfies exactly one discriminant D^a_i for any action a. As such, each effect $\rho^a_{i,j}$ is possible at w. A faithful model has a transition from w (under a) corresponding to each such effect, and no transitions that do not reflect one of these effects. Formally, a transition $w \xrightarrow{a} v$ *corresponds* to an effect $\rho^a_{i,j}$ just when v satisfies $\rho^a_{i,j}$ and agrees with w on all atoms $p \notin atm(a,i,j)$. We say that a model M is *faithful* to an action theory T, if for each w, a, i such that $M, w \models D^a_i$, we have that:

- each transition $w \xrightarrow{a} v$ in M corresponds to some $\rho^a_{i,j}$, and
- for each $\rho^a_{i,j}$, there is a transition $w \xrightarrow{a} v$ that corresponds to it.

As an example, consider the following axioms describing the possible triggering of an unreliable burglar alarm:

$$\textit{force}\ \textbf{causes}\ \top \,||\, \textit{Open} \,||\, \textit{Open} \wedge \textit{Alarm}\ \textbf{when}\ \neg\textit{Open} \tag{4}$$

$$\textit{force}\ \textbf{causes}\ \top\ \textbf{when}\ \textit{Open} \tag{5}$$

Suppose a model of this theory contains a world w satisfying $\neg$*Open*, $\neg$*Alarm* and some "irrelevant" proposition, *Raining* (R). If the model is faithful to the action theory containing this clause, then the model must have exactly three (classes of) transitions associated with w under action *force*: one to a world satisfying exactly the same propositions as w (*Open* and *Alarm* remain false, while R remains true); one to a world satisfying the same propositions except making *Open* true (in particular, *Alarm* remains false); and one making both *Open* and *Alarm* true (again R remains true). The action *force* can map w to no other situations. In contrast, a world v satisfying $\neg$*Open* and *Alarm* has only two possible transitions: if the door fails to open, *Alarm* still persists (forcing does not deactivate the alarm).

While faithful models give an intuitive semantics for an action theory T, we must also provide a logical, syntactic characterization of this semantics. We do so by constructing a compact MPL-theory $I(T)$ that deals with the frame problem, and whose only models are faithful to the original theory T. As such, we may reason about the intended interpretation of T directly within the language of MPL.

The interpretation $I(T)$ of T is formed in two steps. The first step is based on Reiter's proposal: if atom p is not mentioned in $\rho^a_{i,1}, \ldots, \rho^a_{i,k^a_i}$, then it should persist when D^a_i is true and action a carried out. The first part of our procedure deals with these "easy" cases of persistence.

We start with an auxiliary definition. We define the conditions $\textit{Pos}(a, l)$ under which a literal l might become true when action a is performed:

$$\textit{Pos}(a, l) \quad \equiv_{\mathrm{df}} \quad \bigvee\{D^a_i : l \in \rho^a_{i,j} \ \text{ for some } \ j \leq k^a_i\} \tag{6}$$

As usual, we take $\vee\emptyset \equiv \bot$. Thus, if a literal l fails to appear in a possible effect of a in any action clause, l cannot possibly be caused by a (although it may persist). We require that a literal remain true after performance of action a if it is not among the effects of a. That is, for every literal l and action a of interest, we require an axiom

$$l \wedge \neg\textit{Pos}(a, \neg l) \supset [a]l \tag{7}$$

In our example, $\textit{Pos}(\textit{force}, \textit{Alarm}) \equiv \neg\textit{Open}$ and $\textit{Pos}(\textit{force}, \neg\textit{Alarm}) \equiv \bot$. Thus the two axioms expressing the persistence of *Alarm* are:

$$\begin{aligned} \textit{Alarm} &\supset [\textit{force}]\textit{Alarm}, \\ \neg\textit{Alarm} \wedge \textit{Open} &\supset [\textit{force}]\neg\textit{Alarm} \end{aligned}$$

These axioms state that *Alarm* always persist after *force*, and that $\neg$*Alarm* persists when the door is open. Note that these axioms do not characterize whether $\neg$*Alarm* persist when the door is closed — this handled by the second part of our procedure.

The second part of the procedure is more subtle and is needed because of multiple possible outcomes and correlations among action effects. Recall that the intent of an action clause is based on two intuitions: first, the possible outcomes of an action are exactly those that are explicitly mentioned in the clause; and second, only propositions that are explicitly mentioned in this outcome change value. Persistence cannot be restricted to unmentioned propositions when multiple outcomes are possible. In the alarm example, should the door fail to open, alarm is unaffected and must persist; but if the door opens the alarm *may* be triggered.

Again, we start with a preliminary definition. We define a condition that describes when a particular outcome of an action a, say $\rho^a_{i,j}$ actually occurs at a given state (on the "actual" trajectory, path or course of events):

$$\mathit{Occ}(a,i,j) \equiv_{\text{df}} \langle\langle a\rangle\rangle \rho^a_{i,j} \wedge \bigwedge_{P\in atm(a,i)-atm(a,i,j)} \{P \equiv \langle\langle a\rangle\rangle P\} \tag{8}$$

The formula $\mathit{Occ}(a,i,j)$ specifies not only that action a occurs and that it has the effect $\rho^a_{i,j}$, but also that all other atoms mentioned in the ith action clause for a persisted. Thus, the implicit persistence of atoms mentioned in the action clause under consideration is made explicit by this formula — only the outcomes that influence these atoms can cause a change in that atom. For each action clause

$$a \textbf{ causes } \rho^a_{i,1} \,||\, \ldots \,||\, \rho^a_{k^a_i} \textbf{ when } D^a_i$$

we assert two axioms:

$$D^a_i \wedge \langle\langle a\rangle\rangle\top \supset \mathit{Occ}(a,i,1) \vee \cdots \vee \mathit{Occ}(a,i,k^a_i) \tag{9}$$
$$D^a_i \supset \mathsf{E}\mathit{Occ}(a,i,1) \wedge \cdots \wedge \mathsf{E}\mathit{Occ}(a,i,k^a_i) \tag{10}$$

Axiom (9) asserts that if the actual world satisfies the discriminant D^a_i and the current path is such that action a occurs at this state, then *one* of the outcomes $\rho^a_{i,j}$ is realized and all other atoms occurring in the action clause persist (in other words, *only* the outcomes $\rho^a_{i,j}$ are possible). Axiom (10) ensures that for any world satisfying D^a_i and possible effect $\rho^a_{i,j}$, there is a path rooted at that world, with initial action a, that realizes that effect (in other words, *all* of the outcomes $\rho^a_{i,j}$ are possible). Consider again (4); its syntactic interpretation (using the obvious abbreviations) is:

$$\begin{aligned}
\neg O \wedge \langle\langle f\rangle\rangle\top \;\; &\supset \;\; (O \equiv \langle\langle f\rangle\rangle O \;\wedge\; A \equiv \langle\langle f\rangle\rangle A) \vee \\
&\quad\;\; (\langle\langle f\rangle\rangle O \;\wedge\; A \equiv \langle\langle f\rangle\rangle A) \;\vee\; \langle\langle f\rangle\rangle(O \wedge A) \\
\neg O \;\; &\supset \;\; \mathsf{E}(O \equiv \langle\langle f\rangle\rangle O \;\wedge\; A \equiv \langle\langle f\rangle\rangle A) \;\wedge \\
&\quad\;\; \mathsf{E}(\langle\langle f\rangle\rangle O \;\wedge\; A \equiv \langle\langle f\rangle\rangle A) \;\wedge\; \mathsf{E}\langle\langle f\rangle\rangle(O \wedge A)
\end{aligned}$$

Thus, if a *force* action occurs when *Open* is false, either the door and alarm will persist as they were (door fails to open), or the door will open and the alarm will persist (alarm fails), or the door opens and the alarm sounds. The first axiom ensures *only* these three possibilities arise, and the second ensures that *each* of these outcomes is in fact possible.

It is Axioms (9) and (10) that deal with the "correlation problem." This is due to our treatment of each possible *outcome* in the axiom, instead of dealing with each proposition individually as in (7). This shows the key divergence from Reiter's methodology. These axioms also show the need for the expressive power of MPL — equivalent axioms cannot be stated in PDL without explicit enumeration of all truth assignments to propositions in $atm(a, i)$.

Given an action theory T, the *interpretation of* T, denoted $I(T)$, consists of one axiom of form (7) for each literal l and action a, and one axiom of form (9) and one of form (10) for each action clause in T.

3.3 Properties of the Interpretation

In this section, we briefly summarize some of the formal properties of our interpretation procedure. These are described in more detail in the full paper.

In the description of the interpretation procedure above, we have assumed that action theories are complete: for each action a, its possible effects under all circumstances are listed (i.e., the set of discriminants D_i^a is logically exhaustive). However, this need not be the case. In fact, if the effect of action a is not listed for a certain condition C, we expect that the action has no effect when C holds. In other words, all propositions persist under a. In the full paper we show that this intuition is captured by adding the axiom $\langle a \rangle \top$ to $I(T)$ for each action $a \in \mathcal{A}$. The resulting theory is equivalent to the one obtained by adding "a **causes** $\top$ **when** C" to T. As such, axioms such as Axiom (5) above can be left unstated by the user.

The motivation for our interpretation procedure was the desire to capture faithful models of the action theory. In other words, we would like $I(T)$ to be faithful to T. Recall that a faithful model is one in which the admitted transitions are exactly those explicitly described by the action theory. It is not hard to verify that, in fact, $I(T)$ is faithful in this sense.

Theorem 1 *$M \models I(T)$ if and only if M is faithful to T.*

That our syntactic interpretation captures our prior intuitions is thus verified in a formal and precise way.

Another intuition underlying our solution to the frame problem (as well as the solutions proposed by many others) is that the interpretation should be unambiguous. Intuitively, a theory is unambiguous if for any completely specified state of affairs, it determines the precise effect of any action. Lin and Shoham [11] formalize this idea for the situation calculus in deterministic settings. Unfortunately, their formalization renders any theory with nondeterministic actions as "incomplete". In the full paper we describe a more robust notion, that of *unambiguous theories*, suitable for general transition systems, by appeal to a "canonical" model of a theory. The canonical model can be viewed as a complete representation of the system dynamics associated with a given theory. As a consequence of Theorem 1 we have:

Theorem 2 *$I(T)$ is unambiguous.*

If our action theory contains just one possible effect in each action clause, then, according to our interpretation of action theories, it is *deterministic*. More precisely, the intended model of such a theory is deterministic. In this case our solution to the frame problem is equivalent in some sense to that proposed by Reiter. We start by reconstructing Reiter's solution in MPL. Reiter essentially assumes that a primitive proposition P changes value only when the action clause specifically mentions the change. He asserts the following clause for each proposition P and action a:[4]

$$\langle a \rangle\, P \;\equiv\; \textit{Pos}(a, P) \vee (P \wedge \neg \textit{Pos}(a, \neg P)) \tag{11}$$

Since T is deterministic, it has just one outcome in each action clause; therefore if $\textit{Pos}(a, l)$ is true, P must be true after doing a. Axiom (11) states that P is true after action a is executed if and only if a causes P to be true, or P was true beforehand and a does not cause P to be false.

While Reiter's original formulation in SC can have only deterministic actions, this determinism must be made explicit in MPL. Thus, we assert, for each proposition P and action a, the following axiom:

$$\langle a \rangle\, P \;\equiv\; [a]\, P \tag{12}$$

This axiom states that if P is possible in some outcome of a if and only if it is true in all outcomes of a.[5] Let $EC(T)$ denote the collection of axioms of form (11) and (12) for each proposition P and action a. Clearly, $EC(T)$ embodies the essence of Reiter's solution, expressed in the language of MPL.

Theorem 3 *Let T be a deterministic action theory. $M \models I(T)$ iff $M \models EC(T)$.*

This result demonstrates that our proposal for nondeterministic actions is akin to a "conservative extension" of Reiter's solution for deterministic settings: our solution coincide with Reiter's in situations where both apply. Furthermore, it shows that one may use PDL to express the closure of deterministic action theories — the full expressive power of MPL is not required. We give a more detailed comparison with Reiter's proposal in the full paper, and describe the use of *schematic* instances of axioms such as (7) to compactly express our MPL theory.

4 Reasoning about Actions

Given the semantics of actions above, it remains to be seen exactly what role it should play in reasoning about action. The methodology embodied by our solution to the frame problem requires that the action theory be treated somewhat differently than *observations*, or specific constraints over the actual course of events. One reason for this is the distinct roles played by action clauses and observation statements.

Action clauses impose constraints on the dynamics of the system. In particular, they specify properties of the outcomes of *any* (actual or hypothetical) execution of

[4]Reiter actually asserts one clause for each proposition, utilizing quantification over actions to express the "schema" shown here. Should we think of these axioms as schemata, or introduce quantification over actions, the number of axioms we introduce is comparable.

[5]It also ensures that a can be executed in any state of the system, as we have assumed.

an action under given conditions. A model of an action theory is such that every world and every path satisfies these constraints. The frame problem, on our view, is simply a question of elaborating these constraints so that there is no ambiguity in the models of an action theory. In other words, the closed theory completely (and uniquely) determines possible evolutions of the system. For example, the clause

$$\textit{force}\ \textbf{causes}\ \top \,||\, \textit{Open} \,||\, \textit{Open} \wedge \textit{Alarm}\ \textbf{when}\ \neg \textit{Open}$$

ensures that any state satisfying $\neg$*Open* has the possible outcomes listed, regardless of whether that state is ever reached.

Observations, in contrast, specify only properties of some actual path (e.g., that certain actions *actually* occurred and had certain *actual outcomes*). For instance, one might assert that forcing the door occurred while the alarm was silent, and resulted in the alarm sounding. One may then be interested in the other implications of these facts on the actual course of events. To determine what might cause or be caused in such a trajectory, we require some specification of the system's dynamics; we are guaranteed that the only predictions one should make are of those facts that are entailed by the observations *given* the constraints on possible trajectories. The role of observations is to rule out certain paths as "unactualized."

To specify such observations we add *path statements* to our language $\mathcal{A}^{\mathrm{ND}}$. In particular, we are interested in statements that describe the actual path. These take the general form (following [3]):

$$\phi\ \textbf{after}\ \alpha \tag{13}$$

Intuitively, this states that a propositional formula ϕ holds after the program term α (for our purposes, a sequence of actions) is executed. In this example, we might might take the set of observation O to consist of the following statements:

$$\textbf{initially}\ \neg \textit{Open} \wedge \neg \textit{Alarm} \tag{14}$$

$$\textit{Alarm}\ \textbf{after}\ \textit{force} \tag{15}$$

where the "**initially** ϕ" is an abbreviation of (13) when α is the empty sequence of actions. These observations state that in the initial state the door is closed and the alarm is off, but after forcing the door the alarm is triggered. The interpretation of such observations is quite straightforward. For each "ϕ **after** α" in O, we add the MPL-formula $\langle\!\langle \alpha \rangle\!\rangle \phi$ to its interpretation $I^o(O)$.

Given the pair (T, O) where T is an action theory and O is a set of observations, we might want to draw conclusions about the actual state of affairs. To do so we will often pose a query, or sentence Q that describes some feature of the actual path. For example,

$$\textit{Open}\ \textbf{after}\ \textit{force}$$

asks if the door is open after being forced. We say that (T, O) satisfies a query Q if

$$I(T) \models \bigwedge I^o(O) \supset I^o(Q) \tag{16}$$

Thus, the query is accepted if every path permitted by the intended dynamics (i.e., in a model of $I(T)$) that satisfies O also satisfies Q. In our example, the query is accepted, since the semantics of MPL ensures

$$I(T) \models \neg \mathit{Open} \wedge \neg \mathit{Alarm} \wedge \langle\langle \mathit{force} \rangle\rangle \mathit{Alarm} \supset \langle\langle \mathit{force} \rangle\rangle Open$$

(where T is theory described by (4) and (5)). Note that if we drop $\neg$*Alarm* from (14), the prediction "*Open* **after** *force*" is no longer valid (for the alarm may have been on before the force action and, although the door might have failed to open, the alarm persists, explaining (15)).

We should note that many other types of queries are possible, including hypothetical queries. These queries ask about hypothetical, unactualized paths. For example, "what would have happened if instead of forcing the door I would have disabled the alarm first?" In the full paper, we deal with such hypothetical queries in detail, in addition to the more elaborate types of queries that are expressible in MPL, including reasoning about compound actions (i.e., *programs* that involve loops, if-then-else statements, etc.).

In the full version paper we also make a careful comparison of our approach to that of Gelfond and Lifschitz [3]. It is quite is easy check that deterministic theories in our language are expressible in their language $\mathcal{A}$, and as we show, their semantics of deterministic action theories is very close to ours. We note, however, that our semantics for observations is rather different, primarily because of the presence of nondeterministic actions. In deterministic setting, observations simply limit the possible candidate initial states. Once determined, all queries can be answered, simply because an initial state uniquely determines the outcome of any sequence of actions. In our nondeterministic setting, the situation is more complex. As our example above shows, even though the observation (14) specifies a unique initial state (in terms of the propositions in our example), observation (15) provides *additional* information. In particular, observations are used to resolve uncertainty about the actual outcomes of actions.

5 Discussion

We have presented a logic and methodology for the representation of actions with nondeterministic effects, and described a solution to the frame problem in such settings. The formal solution matches our semantic intuitions about the intended constraints on the system dynamics, and we showed that our solution extends Reiter's (deterministic) solution in the sense that the two solutions coincide on deterministic theories. A key feature of of our approach is its exploitation of the expressive power of MPL to deal with correlations among action effects.

The fairly straightforward treatment of nondeterminism is in part due to our insistence that possible outcomes be made explicit in action descriptions. An action a that has two possible effects A and B is written

$$a \ \textbf{causes} \ A \,||\, B \ \textbf{when} \ \top$$

Our interpretation procedure then "closes" each of these effects (more or less) separately, with no ambiguity in the possible transitions that result. It is important to contrast such a *definite* specification with an *indefinite* action description like

$$a \textbf{ causes } A \vee B \textbf{ when } \top$$

From a logical perspective, the requirement that the disjunction $A \vee B$ be true after performing action a can be fulfilled in several ways, even if all other literals are fixed. Possible ways of achieving this effect are A (letting B persist), B (letting A persist), and $A \wedge B$. If we consider nondeterministic models, there are seven classes of models (corresponding to nonempty subsets of these three choices) that match this specification, three of which take a to be deterministic.

The distinction between nondeterminism and indefinite action specification is important: nondeterminism corresponds to inherent *uncertainty* about the outcomes of an action (at least, given the level of detail one is willing to model), while indefinite specifications denote a certain *ignorance* about these outcomes. Once we make this distinction and introduce the representational tools (e.g., the $\langle a \rangle$ operator in MPL or $||$ in $\mathcal{A}^{\text{ND}}$) to capture it, the treatment of nondeterministic actions becomes clear. On our view, the frame problem is one of completing the specification of the *known* system dynamics. This allows one to adopt the convention that only the known effects of an action need be specified; *unmentioned* aspects as treated as unchanging. When action descriptions are definite (even if nondeterministic), possible action outcomes are known; indefinite descriptions do not fix the possible outcomes, and much weaker predictions are the result. For example, the effect $A \vee B$ does not ensure that A is even a possible outcome. The resolution of such ignorance about action outcomes is a problem separate from the frame problem.

Attempts to represent nondeterministic actions using indefinite specifications (e.g., using disjunctions) are forced to make some choice about which transitions are possible. (see, for example, [8] where some minimal change that satisfies the disjunction is used). While this convention is tenable, unfortunately it restricts the expressiveness of the action language. By representing nondeterminism and uncertainty using syntactic constructs meant for indefinite specifications, one loses the ability to express true ignorance of action effects.

There are a number of avenues that remain to be explored. In the full paper, we describe special treatments of actions with uncorrelated effects and with independent "aspects." Future research includes the application of our methodology in more general settings, including dealing with actions with defeasible effects and actions that affect the agent's information state. In particular, we take the qualification problem to be intimately related to the knowledge of conditions under which an action has certain effects, and, given an appropriate semantics for belief, can separate the qualification problem from the frame problem without difficulty. We believe that our methods can be extended quite easily to these cases. We are currently exploring the ramification problem as well.

References

[1] A. B. Baker. Nonmonotonic reasoning in the framework of the situation calculus. *Artificial Intelligence*, 49:5–23, 1991.

[2] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *J. Comp. and Sys. Sci.*, 18:194–211, 1979.

[3] M. Gelfond and V. Lifschitz. Representing actions in extended logic programming. In *Proc. 10'th Conf. on Logic Programming*, pp. 559–573, 1992.

[4] S. Hanks and D. McDermott. Default reasoning, nonmonotonic logics and the frame problem. In *AAAI-86*, pp. 328–333, 1986.

[5] D. Harel, D. Kozen, and R. Parikh. Process logic: Expressiveness, decidability, completeness. *J. Comp. and Sys. Sci.*, 25:144–170, 1982.

[6] G. N. Kartha. Soundness and completeness theorems for three formalizations of action. In *IJCAI-93*, pp. 724–729, 1993.

[7] G. N. Kartha. Two counterexamples related to Baker's approach to the frame problem. *Artificial Intelligence*, 69:379–392, 1994.

[8] G. N. Kartha and V. Lifschitz. Actions with indirect effects (preliminary report). In *KR-94*, 1994.

[9] H. A. Kautz. The logic of persistence. In *AAAI-86*, pp. 401–405, 1986.

[10] N. Kushmerick, S. Hanks, and D. Weld. An algorithm for probabilistic least-commitment planning. In *AAAI-94*, pp. 1073–1078, 1994.

[11] F. Lin and Y. Shoham. Provably correct theories of action (preliminary report). In *AAAI-91*, pp. 349–354. 1991.

[12] D. G. Luenberger. *Introduction to Dynamic Systems: Theory, Models and Applications*. Wiley, 1979.

[13] J. McCarthy. Applications of circumscription to formalizing commonsense reasoning. *Artificial Intelligence*, 28:89–116, 1986.

[14] J. McCarthy and P.J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463–502, 1969.

[15] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.

[16] R. 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)*, pp. 359–380. Academic Press, 1991.

[17] Y. Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press, 1988.