Reasoning about Persistence: A Theory of Actions

Yan Zhang and Norman Y. Foo
Knowledge Systems Group
Basser Department of Computer Science
University of Sydney
NSW 2006, Australia

Abstract

Winslett proposed a method for reasoning about action called the *possible models approach* (PMA). The PMA successfully removed the major difficulty manifested by Ginsberg and Smith's *possible worlds approach* (PWA). In this paper, we show that Winslett's PMA fails to solve the frame and ramification problems for some actions, as does the PWA. From this observation, we classify actions as *definite* and *indefinite*, and find that, in general, the PMA is not appropriate for both definite and indefinite actions. We propose a new approach to formalize actions based on *persistence*. We compare our approach with the PMA in detail, and show that our new formalization can avoid the problems in the PMA and PWA in most cases, and give more intuitive results for reasoning about action, regardless of whether the action is definite or indefinite.

1 Introduction

Winslett proposed a method for reasoning about action, which she called a *possible models approach* (PMA) [Winslett, 1988]. The PMA was devised to remove some difficulties in Ginsberg and Smith's *possible worlds approach* (PWA) for reasoning about action [Ginsberg and Smith, 1987a; Ginsberg and Smith, 1987b]. As argued by Winslett, although the PWA is an elegant, simple and powerful technique for reasoning about action, it fails to solve the *frame problem* (What facts about the world remain true when an action is performed?), *ramification problem* (What facts about the world must change when an action is performed?) and *qualification problem* (When is it reasonable to assume that an action will succeed?), if the PWA is forced to operate with incomplete information [Winslett, 1988].

The basic idea of the PWA is to take a logical theory as the *description of the world*. The effect of an action is modeled by incorporating a set of formulas F (that specify the effects of the action) into the world description T. The description of the possible worlds after the action are the *maximal subsets* S of $T \cup F$ that are consistent and $F \subseteq S$ [Ginsberg and Smith, 1987a].

The PMA, on the other hand, distinguishes the *state of the world* from a theory that describes the world. The state of the world is identified with a Herbrand model. This distinction leads to different effects in representing the *frame principle* of actions between the PWA and PMA. In particular, the frame principle of actions says that there are *minimal changes* in the world when an action is performed. The PWA translates this into "minimal changes in the description of the world when an action is performed." However, as has been shown by Winslett [Winslett, 1988], since a minimal change in the state of the world does not necessarily correspond to a minimal change in the description of the world, the PWA may not give a correct handling of incomplete information. On the other hand, the PMA represents the minimal change on the state of the world rather than on the description of the world when an action is performed, and thus can avoid the problems with the PWA.

Does the PMA always give satisfactory results for reasoning about action? In this paper, we show that the PMA is sometimes incorrect in that it may give unintuitive results relative to the frame problem, and for its dual, the ramification problem. In developing our alternative approach we found it helpful to classify actions into *definite* and *indefinite* actions to systematize a comparison between the PMA and our new proposal. The PMA is shown to be inadequate for both definite and indefinite actions in the general case. Our proposal for reasoning about actions is based on *persistence*. We show that our approach can avoid the problems with the PMA and PWA in most cases, and yield more intuitive results, regardless of whether the action is definite or indefinite.

The paper is organized as follows. In the following section we briefly review the PMA, and illustrate the problem with the PMA by an intuitive example. We then start to develop our new approach. In section 3 we introduce the concepts of definite and indefinite actions, which are helpful for our discussion. We define the persistence set in section 4, and based on this definition we give a precise representation of the effect of action. In section 5 we present more examples to show how our approach works for reasoning about action. We then compare our approach with the PMA in detail from a logical viewpoint in section 6. Finally, we discuss the related work with our approach presented in this paper.

2 The Problem with the PMA

We first briefly review Winslett's PMA. Let T be a first order theory, which is treated as a description of the world, and C be a subset of formulas in T representing the domain constraints about the world called *protected formulas*. To reason about the effects of performing an action with post-condition F (a first order formula), the PMA considers the effect of the action on each possible state of the world, that is, on each model of T. The PMA changes the truth valuations of the ground atoms in each model as little as possible in order to make both F and the protected formulas of T true in that model. The possible states of the world after the performance of the action are all those models thus produced. Winslett made a Herbrand universe assumption so that models are simply subset of the Herbrand base.

Formally, we say that two models M_1 and M_2 *differ* on a ground atom p if p appears in exactly one of M_1 and M_2. Let $Diff(M_1, M_2)$ denote the set of all different ground atoms between M_1 and M_2, M a model of T, and F a first order formula. $Incorporate(F, M)$ is the set of all models M' such that

1. F and the protected formulas of T are true in M';
2. for any model M'' satisfying 1, $Diff(M, M'') \subseteq Diff(M, M')$ implies $M'' = M'$.

Then Winslett defined the possible states of the world resulting from applying an action with post-condition F as

$\bigcup_{M \in Models(T)} Incorporate(F, M)$.

The PMA uses the principle of minimal change on models to represent the frame principle of actions when an action is performed. The question is: is it always reasonable? Let us consider the following example.

Example 1. Imagine an instructional computer system in which all programs have three statuses called *running*, *terminated*, and *abnormally terminated* (*ab-terminated*). At any time point, each program has only one of the three statuses. We also assume that for any program, the only occurrence of abnormal termination of the program is caused by system errors. The program executes the next instruction on receiving the input "*continue*". Formally, we have two domain constraints (i.e., protected formulas in Winslett's terminology) to represent this world:

(1) $\forall xy\ (running(x, y) \oplus terminated(x, y) \oplus$ $ab\text{-}terminated(x, y))$,

(2) $\forall xy\ (system\text{-}ok(y) \supset \neg ab\text{-}terminated(x, y))$,

where $\oplus$ is exclusive or, $running(x, y)$ means that program x is running on system y (similarly for the meanings of $terminated(x, y)$ and $ab\text{-}terminated(x, y)$), and $system\text{-}ok(y)$ means that system y is at a normal status. Suppose the current unprotected formulas are

(3) $system\text{-}ok(a)$,

(4) $running(p, a)$.

From the definition of the PMA, the current description of the world T consists of (1) $-$ (4), and the only model S_0 of T consists of ground atoms (3) and (4), i.e., $S_0 = \{system\text{-}ok(a), running(p, a)\}$. By the Herbrand model convention the negative ground atoms such as $\neg terminated(p, a)$ and $\neg ab\text{-}terminated(p, a)$ are omitted. Now consider the action *continue*. After this action, what is the status of program p? Intuitively, after a while, the status of p has two possibilities $-$ one is still running, the other is terminated (since there is no any explicit information to infer that system errors may have occurred after *continue*, p's status will not be abnormally terminated). So the post-condition of *continue* is simply $running(p, a) \vee terminated(p, a)$. Now what can we say for the result of *continue*? From the previous definition, it is easy to show that

$\bigcup_{M \in Models(T)} Incorporate(Post(continue), M) = S_0$.

That is, the only possible state of the world resulting from action *continue* is that nothing has changed, i.e., the program will run forever if no system error occurs. However, this is unreasonable from our intuition. It is not difficult to show that the PWA also gives the same result: the only possible world to add $running(p, a) \vee terminated(p, a)$ in T is T itself. $\Box$

The reason for this unintuitive result is that the PMA represents the frame principle of actions by the minimal change on models when an action is performed, and this is not always correct. This can be easily seen if we consider the *persistence* of facts in the state of the world when an action is performed. Informally, we say that a fact *persists* if given that no relevant action has occurred over a stretch of time, the fact does not change its truth value over that time [Shoham, 1988]. In our example, the effect of action *continue* is $running(p, a) \vee terminated(p, a)$, and the initial state of the world is $\{system\text{-}ok(a), running(p, a)\}$. The only fact that persists is $system\text{-}ok(a)$, and since $running(p, a) \vee terminated(p, a)$ may or may not cause the truth value of $running(p, a)$ to change, we have no any reason to say that $running(p, a)$ *must* persist, where the PMA does just because the state $\{system\text{-}ok(a), running(p, a)\}$ satisfies the formula $running(p, a) \vee terminated(p, a)$ before *continue* is performed. The assumptions underlying the PMA can therefore be sometimes too strong.

3 Definite and Indefinite Actions

Consider a first order language L with equality. Let $F(p)$ be a first order formula, where p is a tuple of non-equality predicate symbols occurring in the formula. Let I be a Herbrand interpretation of the language, and $|I|$ the universe of I. For simplifying our discussion, in this paper, we *fix* the universe of interpretations in language L. That is, we only consider the class of interpretations in which for any two interpretations I_1, I_2 of L, $|I_1| = |I_2|$, and each function symbol (constant is a 0-ary function symbol) of L has the same interpretation in I_1 and I_2 respectively.

Definition 1 A formula $F(p)$ is *definite* if for any two interpretations I_1, I_2 such that $I_1 \models F(p)$ and $I_2 \models F(p)$, and for each $p_i \in p$ and each tuple of constants (functions) c which has the same arity as p_i, $I_1 \models p_i(c)$ iff $I_2 \models p_i(c)$. $\Box$

Intuitively, a formula F is definite if for any two interpretations (remember that we fix the universe of interpretations of L) in which F is satisfied, each predicate symbol occurring in F has the same extension in these two interpretations. We say that a formula is *indefinite* if it is not definite and not a tautology. It is not difficult to decide if a formula is indefinite.

Proposition 1 The following results hold:

1. $\exists x[\neg]p(x)$ is indefinite, where p is a non-equality predicate symbol,
2. for any formulas F_1 and F_2, where $F_1 \not\equiv F_2$, $F_1 \vee F_2$ is indefinite if $F_1 \vee F_2$ is not a tautology.

The notation $[\neg]$ means that the negation sign $\neg$ may or may not appear. $\Box$

From the concepts of definite and indefinite formulas, we classify actions into two kinds, – definite and indefinite actions. Herbrand models represent states of the world and it is assumed that there is a set of domain constraints C about the world, which are satisfied by each state of the world. An action a is defined by a pre-condition $Pre(a)$ and a post-condition $Post(a)$, where $Pre(a)$ and $Post(a)$ are first order formulas. We say that a is *applicable* to a state S of the world if $S \models Pre(a)$. Informally, action a changes the state of the world from one to another, and in the resulting state of the world, $Post(a)$ is satisfied. The aim of this paper is to give a reasonable representation for the possible resulting states of the world when an action is performed on a state of the world. We also assume that the precondition $Pre(a)$ is *complete*, in the sense that the action can definitely be performed if the pre-condition is satisfied. The post-condition of the action, however, will not be assumed complete, since there may be many ramifications of an action in any particular situation.

Definition 2 An action a is *definite* if

1. $Post(a)$ is a definite formula, and
2. for any formula F satisfying $\{Post(a)\} \cup C \models F$ but $C \not\models F$, where F is not followed from the expansion rule or the $\exists$-introduction rule (i.e., not followed via $A \models A \vee B$ or $A(c) \models \exists x A(x)$), F is definite. $\Box$

Following Definition 2, we say an action a is *indefinite* if $Post(a)$ is indefinite; or for some formula F satisfying Condition 2 of Definition 2, F is indefinite.

4 A Persistence-based Theory of Actions

In this section, we define a precise representation for the possible states of the world resulting from actions based on *persistence*. Let S_0 be the current state of the world, C be the set of domain constraints and a be an action.

Definition 3 Suppose action a is applicable to S_0, i.e., $S_0 \models Pre(a)$. Let

1. $\Delta_0 = S_0 - \{f \mid \{Post(a)\} \cup C \models \neg f\}$,
2. $\Delta_1 = \Delta_0 - \{f \mid$ there are ground atoms $f_1, \cdots, f_n$, where $f_i \in S_0$ or $\neg f_i \in S_0$[1] for each $i \in N = \{1, \cdots, n\}$, such that $\{Post(a)\} \cup C \models [\neg]f \vee \bigvee_{i \in N}[\neg]f_i$ but $C \not\models [\neg]f \vee \bigvee_{i \in N}[\neg]f_i$, and for any proper subset M of N, $\{Post(a)\} \cup C \not\models [\neg]f \vee \bigvee_{j \in M}[\neg]f_j\}$,
3. $\Delta_i = \Delta_{i-1} - \{f_1, \cdots, f_k \mid$ there exists some subset $F \subseteq S_0 - \Delta_{i-1}$ such that $(F \cup \Delta_{i-1} - \{f_1, \cdots, f_k\}) \cup C \models f_1 \wedge \cdots \wedge f_k$ but $(\Delta_{i-1} - \{f_1, \cdots, f_k\}) \cup C \not\models f_1 \wedge \cdots \wedge f_k\}$.

The *persistence set* $\Delta(a, S_0)$ of S_0 with respect to the action a is defined as

$\Delta(a, S_0) = \bigcap_{i=0}^{\infty} \Delta_i$. $\Box$

[1] By the Herbrand model convention the negative ground atoms are omitted. However, for convenience in our discussion, we assume $\neg f_i \in S_0$ if $f_i \notin S_0$, where f_i is in the Herbrand base.

We explain this definition as follows. Let a be an action applicable to S_0. Obviously, after a's performance, a fact f in S_0 *must* change if f is inconsistent with $Post(a)$ with respect to C, i.e., $\{Post(a)\} \cup C \models \neg f$. We call such a fact *non-persistent*. On the other hand, if f or $\neg f$ appears in a disjunction which is entailed by $Post(a)$ (or together with other formulas, i.e., C), we say that f is *indefinitely affected* by a. The intuitive meaning of an indefinite effect is that after performing a the satisfaction of $Post(a)$ in the resulting state may or may not cause a change of the truth value of f, but we do not know which is the case. Finally, if there exists a subset F of S_0 in which some fact is non-persistent or indefinitely affected by a, and f is entailed by F (or together with C), then f is *implicitly affected* by a because after a's performance, f may lose its justification F. In this case, if there is no other justification for f, it is incautious to assume that f is persistent. We call those facts that are either implicitly or indefinitely affected by a *mutable*. Thus, the persistent facts in S_0 are those that are not non-persistent or not mutable. This is shown by the following diagram.

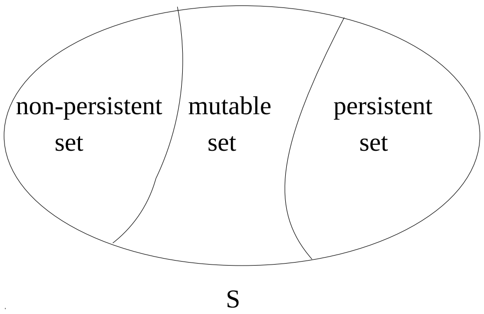

Figure 1: A state of the world to which a is applicable.

Intuitively, the persistence set $\Delta(a, S_0)$ is the maximal set that includes all facts in S_0 which *must* persist when a is performed, and further, it excludes any fact in S_0 whose truth value *must* or *might* change *only* after a's performance. In fact, this persistence definition is based on a conservative principle. As discussed in Section 2, it is generally assumed that if an action is irrelevant to some fact, that fact will persist during the performance of that action. Unfortunately, there is no purely logical method which is strong enough to formalize the *irrelevance* of actions in the general case [Shoham, 1987]. But we can resort to the dual. It is possible to identify the widest possible scope of action relevance, and this is through non-persistent and mutable sets. More details about persistence can be found in [Zhang and Foo, 1992].

Proposition 2 For any state of the world S_0 and any action a which is applicable to S_0, the persistence set $\Delta(a, S_0)$ is unique. □

The representation of the persistence-based formalization of actions can then be expressed as follows.

Definition 4 We say that a Herbrand interpretation S is a *possible resulting state* of the world after action a is performed on the state of the world S_0, if

1. $S \models C$,
2. $S \models Post(a)$, and
3. $\Delta(a, S_0) \subseteq S$.

We define the set of all such resulting states as the *possible states* of the world resulting from applying action a on S_0, denoted as $Result(a, S_0)$. Call this way of determining the possible resulting states the *Persistence Set Approach*, PSA in brief. □

5 More Examples

Let us consider the example presented in Section 2 once again. From Definition 2 and Definition 3, we know that *continue* is indefinite, and the persistence set $\Delta(continue, S_0) = \{system\text{-}ok(a)\}$ (recall that $S_0 = \{system\text{-}ok(a), running(p, a)\}$), thus there are two possible states of the world

$S = \{system\text{-}ok(a), running(p, a)\}$, and
$S' = \{system\text{-}ok(a), terminated(p, a)\}$,

which are the desired results for this example. The PSA can differ from the PMA even for definite actions, as is shown in the next example.

Example 2. Consider the *living room domain* which was first proposed by Ginsberg and Smith [Ginsberg and Smith, 1987a] and reconsidered by Winslett [Winslett, 1988; Winslett, 1991]. Ignoring the detailed semantic description of this domain (see [Ginsberg and Smith, 1987a; Winslett, 1988]), here we just summarize its set of domain constraints C as follows:

(1) $duct(x) \equiv (x = duct1 \vee x = duct2)$
(2) $location(x) \equiv (duct(x) \vee x = floor)$
(3) $(on(x, y) \wedge on(x, z)) \supset (y = z)$
(4) $(on(x, y) \wedge on(z, y)) \supset (z = x \vee y = floor)$
(5) $(duct(y) \wedge \exists x on(x, y)) \equiv blocked(y)$
(6) $(blocked(duct1) \wedge blocked(duct2)) \equiv stuffy(room)$
(7) $on(x, y) \supset (location(y) \wedge \neg location(x))$
(8) $\exists y on(x, y) \vee location(x)$

Suppose the initial description of the world T is $C \cup \{on(TV, duct1), on(plant, duct2), on(magazine, floor), \exists y on(newspaper, y)\}$. (Remember that the PWA and the PMA do not exclude incomplete information in the world description). From (4) and (8) we know that $T \vdash on(newspaper, floor)$. Furthermore, it is easy to see that the only Herbrand model S_0 of T [2] is

$\{on(TV, duct1), on(plant, duct2),$
$on(newspaper, floor), on(magazine, floor),$
$blocked(duct1), blocked(duct2),$
$stuffy(room)\}$.

So in the initial state of the world, the newspaper is on the floor. Now consider an action to move TV from $duct1$ to $floor$ (i.e., $moving(TV, floor)$. Its pre-condition is $on(TV, duct1)$ and post-condition is $on(TV, floor)$). According to the PWA, the only possible world with this action is

$$T' = T \cup \{on(TV, floor)\} - \{on(TV, duct1)\}.$$

Now the question is: is the newspaper still on the floor? What we can get from T' is
$T' \models on(newspaper, floor) \vee on(newspaper, duct1)$.

Thus from the PWA, there are two possible positions of the newspaper – one says that after action $moving(TV, floor)$ the newspaper is still on the floor, which is the intuitive result from the frame principle; the other says that after the action, the newspaper flies to $duct1$. How can we explain the second possible position of the newspaper? One explanation is that this change is caused by powerfully sucking air through the windows [Winslett, 1988]. But why can't the magazine change its position for the same reason?

The PSA, on the other hand, avoids this unintuitive result. From Definition 3, the persistence set of S_0 with respect to $moving(TV, floor)$ is

$\{on(plant, duct2), on(magazine, floor),$
$on(newspaper, floor), blocked(duct2)\}$.

So the only resulting state S after this action is

$\{on(TV, floor), on(plant, duct2),$
$blocked(duct2), on(magazine, floor),$
$on(newspaper, floor)\}$,

and this is our desired result.

It is not difficult to show that $moving(TV, floor)$ is a definite action. By applying the PMA, however, what we get are not only the desirable state S, but also other five undesirable states [Winslett, 1988]. Winslett argued that there was a way to introduce some prioritization on predicates, such that the undesired states can be eliminated by minimizing certain predicates according to the prioritization [Winslett, 1991]. Winslett also showed that this idea is closely related to Lifschitz's prioritized circumscription [Lifschitz, 1985]. However, this approach is not suitable for indefinite actions. Considering Example 1 (ignoring the detail) it is not difficult to show that applying the PMA, regardless of how prioritization on predicates are defined, the desirable results are not obtainable [Zhang and Foo, 1992]. □

6 Comparison with the PMA

In previous sections, we have shown that the PSA can give more intuitive results than the PMA for both definite and indefinite actions. Sometimes the PMA gives

[2] Following Winslett, we omit the "location" and "duct" atoms in the model for brevity.

more possibilities than expected (i.e., move TV from $duct1$ to the $floor$ in Example 2), and sometimes it gives less (i.e., the *continue* action in Example 1). Our objective in this section is to provide a comparison of logical properties of the PMA with the PSA.

Despite the fact that the basic motivations of these two approaches are quite different we can examine the relations between them by considering the *mutable sets* obtained from these two different approaches. Let S_0 be an initial state of the world, C the set of domain constraints of the world, a an action applicable to S_0, and $\Delta(a, S_0)$ the persistence set of S_0 with respect to a. Let $\Gamma(a, S_0)$ be the set of *non-persistent* facts of S_0 (the facts that must change), i.e., $\Gamma(a, S_0) = \{f \mid f \in S_0 \text{ and } \{Post(a)\} \cup C \models \neg f\}$. Obviously, for any $S \in Result(a, S_0)$ and $S' \in Incorporate(Post(a), S_0)$, $\Gamma(a, S_0) \cap S = \emptyset$ and $\Gamma(a, S_0) \cap S' = \emptyset$. The mutable set obtained from the PSA is

$$\delta_{\mathrm{PSA}}(a, S_0) = S_0 - \Delta(a, S_0) - \Gamma(a, S_0),$$

and its analog obtained from the PMA, $\delta_{\mathrm{PMA}}(a, S_0)$, is

$$S_0 - \bigcap Incorporate(Post(a), S_0) - \Gamma(a, S_0).$$

We know that in the PMA, $\bigcap Incorporate(Post(a), S_0)$ is the set of all facts that are true in every possible model. So $S_0 - \bigcap Incorporate(Post(a), S_0) - \Gamma(a, S_0)$ is the set of mutable facts corresponding to the PMA.

Definition 5 Let f be a positive or negative ground atom f, Σ a set of positive or negative ground atoms. We say that Σ is a *support set* of f with respect to C, where C is a set of formulas, if

1. Σ and C are consistent,
2. $\Sigma \cup C \models f$, where $\Sigma \not\models f$ and $C \not\models f$,
3. for any $f' \in \Sigma$, $\Sigma - \{f'\} \cup C \not\models f$. □

Definition 6 Let S_0 be a state of the world, C the set domain constraints of the world, $F \subseteq S_0$ and $f \in S_0$. We say that F *influences* f, if there exists a support set Σ of f with respect to C in S_0, i.e., $\Sigma \subseteq S_0$, such that $F \subseteq \Sigma$. □

Directly from this definition, we know that if F influences f, then for any non-empty subset F' of F, F' influences f. We say that F *uniquely influences* f if there exists the *unique* minimal support set Σ of f with respect to C in S_0 such that $F \subseteq \Sigma$. So far, we have the following result.

Proposition 3 Let a be a definite or indefinite action applicable to S_0. Then

1. for each fact $f \in \delta_{\mathrm{PSA}}(a, S_0)$, f is influenced by some subset of $\Gamma(a, S_0)$ or indefinitely affected by a;
2. if $\delta_{\mathrm{PMA}}(a, S_0) - \delta_{\mathrm{PSA}}(a, S_0) \neq \emptyset$, then for each fact $f \in \delta_{\mathrm{PMA}}(a, S_0) - \delta_{\mathrm{PSA}}(a, S_0)$, f is *neither* uniquely influenced by any subset of $\Gamma(a, S_0)$ *nor* indefinitely affected by a. □

This proposition simply says that each mutable fact obtained by the PSA is based on some *logical relevance*, i.e., influence or indefiniteness. However, in the PMA, some fact may be mutable for no reason. This is why the PMA may have some unintended effects in reasoning about an action.

7 Related Work

In this paper, we presented a formalization of actions based on persistence, which we call the *Persistence Set Approach* (PSA), and compared it with the PMA in detail. We showed that the PSA provides a unified framework for representing effects of definite and indefinite actions, and argued that it is conceptually simple and plausible for reasoning about action. We notice that the computation of persistence sets seems difficult generally. In the full version of this paper ([Zhang and Foo, 1992]) we present an algorithmic description and a computational analysis of the PSA necessary for its implementation. We argue that it is possible to get a computationally tractable method for some special cases, i.e., only definite actions occur or all constraints are Horn clauses. More details about the computation of persistence sets will be considered in our further work.

We think that the PSA can be applied as a general methodology for modeling state change. There are a number of issues related to the PSA that we are considering. Here we list some of them:

- the problem of multiple extensions
- persistence updating in knowledge bases
- representing temporal persistence

The PSA, at least in general, induces multiple extensions from a single description of the world (i.e., Example 1). In order to avoid conflicting extensions (or unintuitive resulting states) it is necessary to provide more domain information in the PSA for some applications. It seems feasible to introduce some perference policy into the PSA to handle the problem of multiple extensions (i.e., the idea of *prioritized circumscription* [Lifschitz, 1985] or *eptistemic entrenchment* [Gardenfors and Makinson, 1988]).

The other related work is knowledge base updates. While there are many approaches in this area, most of them have some limitations [Katsuno and Mendelzon, 1991]. One of the difficulties is to update knowledge base with indefinite information (i.e., something like $p(a) \vee q(b)$ or $\exists x p(x)$). As the PSA provides a unified viewpoint for processing definite and indefinite information, we argue that the idea of persistence may provide new insights into knowledge base updates. Furthermore, it seems that the PSA also provides a unified representation for upadtes and actions. In [Zhang and Foo, 1993b], we show that under the persistent semantics, given a specific update operator some related action (or a sequence of actions) can be generated.

The PSA presents a principle of *persistence* for reasoning about change, which is quite different from the principle of *minimality* that is widely used in current approaches. It would not be difficult to combine the persistence idea into other formalisms. In [Zhang and Foo, 1993a] we presented a persistence-based formalization within the situation calculus framework [McCarthy and Hayes, 1969] to represent temporal persistence, and showed that the Yale Shoolting Problem [Hanks and McDermott, 1987] is solved in a natural way, that is quite different from those minimality-based methods.

Acknowledgements

This research is supported in part by a grant from the Australian Research Council. The first author is supported by a SOPF scholarship from the Australian Government. Discussions with members of the Knowledge Systems Group of the University of Sydney have improved the presentation.

References

[Gardenfors and Makinson, 1988] P. Gardenfors and D. Makinson. Revision of knowledge systems using epistemic entrenchment. In Moshe Y. Vardi, editor, *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–96, Los Altos, California, 1988. Morgan Kaufmann.

[Ginsberg and Smith, 1987a] M. L. Ginsberg and D. E. Smith. Reasoning about action I: A possible worlds approach. In F.M. Brown, editor, *Proceedings of the 1987 Workshop on the Frame Problem in Artifical Intelligence*, pages 233–258, Los Altos, California, 1987. Morgan Kaufmann.

[Ginsberg and Smith, 1987b] M. L. Ginsberg and D. E. Smith. Reasoning about action II: The qualification problem. In F.M. Brown, editor, *Proceedings of the 1987 Workshop on the Frame Problem in Artifical Intelligence*, pages 259–287, Los Altos, California, 1987. Morgan Kaufmann.

[Hanks and McDermott, 1987] S. Hanks and D. McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379–412, 1987.

[Katsuno and Mendelzon, 1991] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[Lifschitz, 1985] V. Lifschitz. Computing circumscription. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 121–127, 1985.

[McCarthy and Hayes, 1969] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltxer and D. Mitchie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[Shoham, 1987] Y. Shoham. What is the frame problem? In F.M. Brown, editor, *Proceedings of the 1987 Workshop on the Frame Problem in Artificial Intelligence*, pages 5–21, Los Altos, California, 1987. Morgan Kaufmann.

[Shoham, 1988] Y. Shoham. *Reasoning about Change.* MIT Press, Cambridge, Massachusetts, 1988.

[Winslett, 1988] M. Winslett. Reasoning about action unsing a possible models approach. In *Proceedings of AAAI-88*, pages 89–93, 1988.

[Winslett, 1991] M. Winslett. Circumscriptive semantics for updating knowledge bases. *Annals of Mathematics and Artificial Intelligence*, 3:429–450, 1991.

[Zhang and Foo, 1992] Y. Zhang and N. Foo. On formalizing things that persist. Technical Report, University of Sydney, October, 1992.

[Zhang and Foo, 1993a] Y. Zhang and N. Foo. Persistent extension and temporal projection. unpublished manuscript, March, 1993.

[Zhang and Foo, 1993b] Y. Zhang and N. Foo. Updates and actions. Technical Report, University of Sydney, (to appear), 1993.