The University of Sydney

Possible Worlds Semantics for Events

Technical Report Number 472

September, 1993

Norman Y Foo and Yan Zhang

ISBN 0 86758 690 7

Basser Department of Computer Science
University of Sydney NSW 2006

Possible Worlds Semantics For Events

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

e-mail: norman@cs.su.oz.au yan@cs.su.oz.au

Abstract

Various theories of events or actions have been proposed to account for commonsence conclusions. The typical way in which this proceeds is for a scenario to be invented and a logic tried against it. Then someone else may perturb the scenario a bit and the old logic fails, while a new one may succeed. The basic difficulty with this methodology is its ad hoc character. To overcome this, recent work [LS91] has tried to delineate classes of monotonic theories that the non-monotonic ones are to be measured against. We believe a better methodology is to go the whole hog and admit model classes as the arbiter. To this end we outline some postulates about a possible worlds semantics that is suitable for evaluating non-monotonic theories of events.

Computer Science Technical Report 478 November 1993
University of Sydney

1. Introduction

The theory of actions or events has occupied the attention of many researchers because it is central to the understanding of how intelligent agents can reason about changes in the world that originate from agent activity. Standard texts in artificial intelligence devote substantial portions to formalisms, both procedural and declarative, for specifying and reasoning about events. Among the established formalisms are STRIPS [FN71] and its successors, the situation calculus [MH69] and various non-monotonic logics [GI87], and more recently update [KM91] and belief revision theories [GA88]. Typical of the non-monotonic approaches is the concern for extending incomplete knowledge by either modifying the inference relation in classical logic to a non-monotonic one as in Reiter’s default logic [RE80], or by postulating some essentially second-order minimization as in McCarthy’s circumscription [MC86]. The measure of success of such proposals rest on "benchmarks" such as the infamous Yale Shooting Problem [HM87], and indeed there is a collection of such benchmarks that is claimed to be a good test of proposed formalisms. A major difficulty with this claim is that it appears to a characteristic of work in this area that it progresses by the invention of (sometimes rather contrived) scenarios that formalism A that is in vogue cannot fully solve, but a new formalism B can -- that is, until formalism C comes along with its own scenarios. Unlike the classical sciences of physics or even biology, event theory does not appear to have a community consensus about what are the appropriate experimental frames to decide rival claims. We believe that this is due in part to a confusion between the roles of ontology and epistemology that the classical sciences have been careful to separate.

As enunciated by several philosophers of science like Giere [GI88] a common way by which scientific theory formation proceeds is for scientists to first delimit the ontology -- to decide what things and relations that are to be admitted or recognized (or measured) -- before any theory is attempted. The success or failure of the theory is relative to this ontology. That does not appear to be the way event theorists have hitherto approached their formalisms. In the so-called qualification problem manifested by the notorious "Potato in the Tailpipe" vexation, the failure of successive non-monotonic theories to account for why the car will not start is, we argue, an indication of changing ontology. Physicists, despite their desire for a "theory of everything", operate in practice with different theories and indeed experimental techniques for different domains, and even in the same domain they use different theories to deal with various levels of refinement. Thus, the band theory of solids is the effective one for semi-conductor junctions, but the hybrid parameter theory is the one useful for transistor action. The operating dictum is: a class of theories for a class of ontologies. Lately, a number of papers have appeared that suggest that at least some event theorists are beginning to recognise this dictum. In [LS91,DS92] Lin, Shoham and DelVal have attempted to characterize the class of *monotonic* theories relative to which a non-monotonic theory of events is correct -- i.e., they make the same predictions or postdictions. But a monotonic theory, especially one that is (ground) complete is so close to being essentially a (ground) model that one wonders why the comparison is not made with a model instead. The view that we should first delimit or define the class of *systems* that is of interest is in fact propounded by Sandewall [SA93], Zeigler [ZE84], and Peppas [PE93]. The description of this class can be done in any way that can communicate its scope and general properties. If it is done using set theory, then we may have a traditional class of Tarski models. Often the separation between the ontology and the theory is implicit, as in many engineering "models" which are really formal theories employing algebra and differential equations, and where distinctions are made between observable and no-observable quantities.

In this paper we propose a possible worlds description of the kind of ontology that we believe is suitable for classifying different classes of systems relative to which different classes of event theories can be evaluated. Possible worlds have a long and respectable history and the version that we use here is Kripke semantics as treated in many books [HC84]. We presume some knowledge of this kind of semantics as well as first-order logic, but nothing deep will be required. The basic idea is to let each state of the system be a possible world, and to let each event label an arc of a state transition regarded as a component of the accessibility relation for that event. States may be viewed as ground-complete sets of literals, i.e., for each state s, each ground literal l either l or notl holds in s. This is the formalization of the idea that ideally states capture completely the configurations that are admissible in the chosen ontology, where the latter is reflected by the choice of first-order predicates (function symbols, constants, etc) to describe the system.

In the sections to follow we will briefly review some standard material essential to the subsequent definitions that illustrate some ways to define system classes that are candidates for recent event theories.

In later papers we will address questions about the relationships between theories and this kind of semantics. Here we will merely hint at these relationships because a proper development involves details that will exceed the scope of this exposition. However, using this semantics we will in the near future address questions like the correctness diagnosis of theories, what are frame and qualification problems, and the nature of system invariants and event conditions.

2. Possible Worlds

The definitions that follow are completely routine [HC84] but reiterated for completeness. The formal language used in association with possible worlds is first order (f.o.) logic with no function symbols, which is presumed to be familiar to the reader, augmented with a modal operator. It is also presumed that f.o. formulas are understood [HC84].

Definition

A model M is a quadruple (S,R,V,D) where

i. S is a set of worlds or states;

ii. R is the union of binary relations $R sub e$ on S X S, each e $member$ E being a ground event. Intuitively, $R sub e$ are the arcs that label the state transitions due to event e;

iii. D is a domain of individuals;

iv. and V is a function that: associates with each n-ary predicate symbol P an (n+1)-ary relation consisting of tuples ($d sub 1 , .. ,d sub n$, s) where $d sub 1 , .. ,d sub n$ are elements of D, and s is a state in S. For each given s it is convenient to denote this relation by $P sub s$ ($d sub 1 , .. ,d sub n$) which is a "curried" form of V. Intuitively, V says of each interpretation of P in state s which tuples in D satisfy it. Additionally, we stipulate that each element in D is named by a constant that V associates with it. In exactly the standard way, V treats variables by associating each variable x some element in D. Denote this element by V(x).

Evaluation of formulas in M

These recursive definitions are again routine and included for completeness. In the first two, read tuples in place of x and d, and we conflate d on the left (a constant symbol) with its occurrence on the right (an element in D).

i. M,s $dturn$ P(d) if $P sub s (d)$ holds.

ii. M,s $dturn$ P(x) if $P sub s (V(x))$ holds.

iii. M,s $dturn$ x = y if $V(x) = V(y)$ holds

iv. M,s $dturn$ $all (x) alpha (x)$ if for all $V prime1$ differing from V in at most its association of x with V(x), M,s $dturn$ $alpha$ (x).

v. M,s $dturn$ $alpha˜ra˜ beta$ if M,s $dturn$ $alpha$ implies M,s $dturn$ $beta$.

vi. M,s $dturn$ $not˜alpha$ if M,s $ndturn$ $alpha$.

vii. M,s $dturn$ [e] $alpha$ if for all t such that s$R sub e$t, M,t $dturn$ $alpha$.

The usual abbreviations and and or are assumed. In line with most work in event theory we assume that each event e has associated with it a pre-condition pre(e) and a post-condition post(e), where each is a f.o. formula.

Definition

If M,s $dturn$ pre(e), the event e is applicable to s.

Transition Axiom

For all states s and events e, pre(e) ra [e]post(e)

From the semantic evaluations above this amounts to M,s $dturn$ pre(e) and s$R sub e$t implies M,t $dturn$ post(e). (For readers familiar with work in dynamic logic or Floyd-Hoare logic [GO82] this is the same definition that is used for relating pre- and post-conditions for programs. Indeed, it is a bit of a mystery why the similar terminology has not prompted more cross-fertilization between programming language semantics and event theory.) In the AI literature this axiom is also called the Causality Axiom, an appropriate name if one considers the post-condition to be caused by the event.

Definition

When s$R sub e$t, call s the source state and t a target state of e.

Application Axiom

For all states s and events e, s$R sub e$t implies M,s $dturn$ pre(e) and M,t $dturn$ post(e).

This is needed to prevent state pairs from being spuriously attributed to event e. The implication should not in general be reversed because state t may by happenstance satisfy the post-condition of e but may not in fact be reachable from s.

All models M are required to satisfy the above two axioms. In the jargon of modal logic, we regard these axioms as defining a standard frame. Additional assumptions about the relations R_e or states s are called postulates. Unlike the axioms, postulates are introduced in an attempt to to provide models that can be used to characterize various non-monotonic and other logics.

Definition

A constraint is a f.o. formula ϕ such that ϕ is not a tautology and M,s $\models$ ϕ for all s in S. Denote the set of all constraints by C. If this is finite (as we shall assume here), C will also be a formula that represents all constraints.

A few remarks about pre- and post-conditions are necessary to ensure that the first part of the development below is placed in proper context. In the usual treatment of pre-conditions (similar remarks apply to post-conditions) an event e is specified *generically* using free variables that get instantiated to particular names when a specific event of that type is applied to a particular state. For instance the "stack" event in a blocks world may have pre-condition schema $\forall X \neg on(X,Y)$ when the generic event stack(Z,Y) is specified. However, for the purpose of this paper, events are specific until we consider otherwise. This means that "stack" events for stack(a,b) and stack(c,d) are considered to be different even though they arise from the same schema. The generalization to generic events specified by schemas, while straightforward, is a distraction at this point.

3. Inertia and Minimal Change

As a first application of this semantics we will consider some approaches to reasoning about state and knowledge change. These are typified by the idea of *updates* [KM91] on the one hand and by *belief revision* [GA88]on the other. The cited references show that the notion of *minimal change* is central to both state and belief change. In state change the idea seems to be one that can be described as a kind of inertia -- the components of states that do not have to change as a result of an event do not in fact change. In belief revision the idea is that in revising beliefs to accept new beliefs, we should discard as few old beliefs as we can. These ideas are easy to capture as additional postulates in our semantics. Indeed the aim of this paper is to suggest that the various approaches to event theory be validated against similar postulates about possible worlds by exhibiting some interesting examples.

First a few preliminaries. For convenience we will denote a clause $l_1 \vee \ldots \vee l_n$ by $\langle l_1 , \ldots , l_n \rangle$.

Definition

If $GAMMA$ $dturn$ <$l sub 1 , ... , l sub n$> but $GAMMA$ $ndturn$ <$l sub i sub 1 ,.., l sub i sub k >$ for no proper subclause of the ground clause $<l sub 1 , ... , l sub n >$, say that $<l sub 1 , ... , l sub n >$ is $GAMMA$-minimal. When n > 1, also say that $<l sub 1 , ... , l sub n >$ is unsubsumed.

Definition

Ram(e) = {$<l sub 1 , ... , l sub n >$ | $<l sub 1 , ... , l sub n >$ is post(e) cup C -minimal, and C $ndturn$ $<l sub 1 , ... , l sub n >$}.

Ram(e) is what is usually called the ramifications of the event e, with post(e) being the immediate effects.

Some observations:

i. If $<l sub 1 , ... , l sub n >$ is in Ram(e) then C $dturn$ post(e) ra $<l sub 1 , ... , l sub n >$.

ii. If $<l sub 1 , ... , l sub n >$ is in Ram(e) then in every state s such that M,s $dturn$ post(e), M,s $dturn$ l for at least one literal l of this clause.

iii. Ram(e) is satisfiable when post(e) cup C is consistent. By virtue of (ii), every satisfying assignment of truth values to the literals of Ram(e) is a partial specification of a state. The determination of how many such assignments there are is a #p-complete problem.

Definition

If s is a state, Lit(s) = {l | l is a ground literal and M,s $dturn$ l}. A set $DELTA$ $subset$ Lit(s) is a support for l $member$ Lit(s) if (i) l $nomem$ $DELTA$, (ii) C $ndturn$ l and (iii) C cup $DELTA$ $dturn$ l. It is a minimal support if whenever $DELTA prime1$ $subset$ $DELTA$ is also a support for l, then $DELTA prime1$ = $DELTA$.

Definition

A subset $GAMMA$ $subset$ Lit(s) is support stable in s if every l $member$ $GAMMA$ that has a support has a minimal support $DELTA$ $subset$ $GAMMA$.

Some observations:

Support stable sets may have literals that have no support. These literals correspond to self-justifying or directly observable facts. Support sets need not in general be logically closed. It is important to note that our definition of support sets does not consider non-literals as possible members. While it is possible to extend consideration to formulas like praq, the added complexity is a distraction at this stage in our development.

Examples

1. Let Lit(s) and $GAMMA$ both be {p,q} and C be $not p or q$. Then q has minimal support {p} and p has no support. Trivially $GAMMA$ is support stable.

2. Let Lit(s) be {p,q,r}, $GAMMA$ be {q,r} and C be {$not p or q,˜not q or r$}. Then q has minimal support {p} and r has minimal support {q}. Because p $nomem$ $GAMMA$, $GAMMA$ is not support stable.

3. Let Lit(s) be {p,q,r}, $GAMMA$ be {p,r} and C be {praq, qrar, rrap}. Then $GAMMA$ is support stable because {p} and {r} are mutual minimal supports for each other even though {q} is not in $GAMMA$. Clearly $GAMMA$ is not logically closed. This example illustrates the fact that the *circular* supports accepted by many foundationalists need not remain intact in support stable sets, but of course they are recovered by logically closing such sets. It is also seen in this example that "hidden" chains ({q} here), need not be made explicit because any cycle that is a constraint is reducible to equivalences between pairs in the cycle.

Lemma

Unions of support stable sets are support stable. Hence, there is a unique largest support stable set in any given set.

Definition

Let $GAMMA sub 0 , GAMMA sub 1 , ... $ be a sequence of sets such that $GAMMA sub 0$ = l, where l is a ground literal, and $GAMMA sub i+1$ = $GAMMA sub i$ cup {l | l $member$ $DELTA$, $DELTA$ a minimal support of m $member$ $GAMMA sub i$}. A foundation for l based on this sequence is F(l) = $union from {i member omega} GAMMA sub i$.

Some observations:

Generally there is more than one foundation for each l . There is always some foundation for l even if it has no support, viz., {l} itself. Such l may be considered to be self-supported.

Lemma

Every foundation F(l) is support stable.

We are now ready to exhibit the additional postulates promised. In event theory there are two extremes of change. In one, called *coherent* change, objects or beliefs denoted by literals are not linked to others by the notion of support. In other words, even if a constraint like p ra q is present and both p and q hold in state s, there is no causal or other material connection between p and q, so that changing p to notp in a next state because of an event will not force a change in q to notq. This is in contrast to *foundational* change where some such link is present. The intuitive notion of some literals being dependent on others is captured by the definition of F(l) above. Each view, coherent and foundational, has its own version of inertia and minimality.

Observation (ii) following the definition of Ram(e) above, while showing that for an unsubsumed clause in Ram(e) at least one target state t exists such that at least one literal of the clause holds in t, does not force the existence of a target state for each literal in the clause. If we desire a model in which this is forced, the condition can be expressed as follows.

(Weak) Permissive State Postulate

If <$l sub 1 , .. , l sub n$> is in Ram(e), then for each literal $l sub i$ in the clause <$l sub 1 , .. , l sub n$> there is a target state t such that M,t $dturn$ $l sub i$.

Note that the states corresponding to the different literals in the clause need not all be different. If we desire them to be distinct, we have the strong permissive form.

(Strong) Permissive State Postulate

As above, but each literal holds in a separate target state.

Some remarks on permissivness are helpful. For weak permissiveness, we observe that is not the case that every satisfying truth assignment for Ram(e) fails to satisfy a particular literal l in a given unsubsumed clause $l~or~D$ where D is the rest of the clause. For if so Ram(e) $dturn$ $not l$, and then by resolving $l~or~D$ with $not l$ we have Ram(e)

$dturn$ D, which contradicts the minimality of $l~or~D$. Hence every such literal l is satisfiable by some assignment and weak permissiveness is feasible. One might worry whether strong permissiveness is also feasible. It would be infeasible if for some pair p and q of literals in a given unsubsumed clause $p~or~q~or~D$ every satisfying truth assignment for Ram(e) either satisfies or falsifies p and q jointly. If this is so then Ram(e) cup {$p~and~not q$} is not satisfiable. But this means that Ram(e) $dturn$ $not p~or~q$. Then by resolution, Ram(e) $dturn$ $q~or~D$, contradicting the minimality of $p~or~q~or~D$. Hence strong permissiveness is feasible.

In fact, these two postulates are merely extreme ends of a more general situation if consistency of phi with Ram(e) is sufficient to guarantee the existence of a target state in which phi holds. Call this the:

Complete Target State Postulate

An immediate consequence of this postulate is the following. If <$l sub 1 , .. , l sub n$> is in Ram(e) then for every subclause <$l sub i sub 1 , .. , l sub i sub k$> such that Ram(e) $ndturn$ <$not l sub i sub 1 , .. , not l sub i sub k$> there is a target state t with M,t $dturn$ $l sub i sub j$ for 1 le j le k.

Here are some intuitively persuasive postulates that attempt to capture system "laziness".

Coherent Inertia Postulate

For all ground literals l, events e, states s and t such that M,s $dturn$ pre(e) M,s $dturn$ l and Ram(e) $ndturn$ $not l$ imp $exist$t (s$R sub e$t and M,t $dturn$ l).

Coherent Minimality Postulate

For all ground literals l, events e, states s and t such that M,s $dturn$ pre(e) allt (s$R sub e$t imp M,t $dturn$ l) imp (M,s $dturn$ l or Ram(e) $dturn$ l).

Corollary

Coherent inertia is equivalent to coherent minimality.

Inertia is not as strong as it could be. For instance, if the literals p, q and r held in s and Ram(e) is exactly $not r$, then it is admissible under the coherent inertia postulate for there to be two target states in one of which $not p$ holds and in the other $not q$ holds.

Intuitively, there should be only one target state, since there is no reason for p or q to change. So, why is the $exist$ introduced in the postulate (it is responsible for not "nailing down" a unique target state)? The reason is that if Ram(e) has unsubsumed clauses, e.g., $q˜or˜r$, then although Ram(e) $ndturn$ $not q$ and Ram(e) $ndturn$ $not p$, we cannot have both p and q hold in the target state. Generally, we can do better than the looseness in the coherent minimality postulate. Given an unsubsumed clause of n literals in Ram(e), any n-1 subset of its literals in negated form will not be inconsistent with it, and if these negated forms hold in the source state (and are not contradicted by Ram(e)) then a strong form of inertia will preserve them in a single target state. In the special case of definite events (see below), the strong form makes eminent sense. What is desired is a kind of *narrowing* or the relation $R sub e$ so that its target set is as small as possible. While it is possible to define this (indeed enough hints were given to do this) for this paper we will do so only for the definite events that will be discussed later.

Definition

The certainty set Cer(s,e) of a state s under event e is the subset of Lit(s) such that l $member$ Cer(s,e) iff M,t $dturn$ l for every t such that s$R sub e$ t.

Definition

The uncertainty set Unc(s,e) of a state s under event e is the subset of Lit(s) such that l occurs as a disjunct of an unsubsumed clause in Ram(e).

Definition

Given state s, denote the largest support stable subset of Lit(s) in s that is consistent with Ram(e) by $sigma$(s,e).

Properties of $sigma$(s,e):

If l $nomem$ $sigma$(s,e) then either $sigma$(s,e) cup {l} is not consistent with Ram(e) or $sigma$(s,e) cup {l} is not support stable in s. In the former case, Ram(e) $dturn$ not l. In the latter case consider the foundations F(l) of l. If there is one F(l) that is consistent with Ram(e) then F(l) cup $sigma$(s,e) is support stable (it being a union of two support stable sets) and also consistent with Ram(e). This violates the definition of $sigma$(s,e) being the largest set with these two properties. Thus all foundations F(l) of l $nomem$ $sigma$(s,e) are inconsistent with Ram(e). We therefore have (the if direction follows from the definition):

Theorem

l $nomem$ $sigma$(s,e) iff either l or every F(l) is inconsistent with Ram(e).

Analogous to the coherent minimality postulate, when objects in states bear a foundational relationship to one another whenever they are in a logical support relation, we have the following:

Foundational Inertia Postulate

For all ground literals l, events e, states s and t such that M,s $dturn$ pre(e) M,s $dturn$ l and l $member$ $sigma$(s,e) imp $exist$t (s$R sub e$t and M,t $dturn$ l).

Foundational Minimality Postulate

For all ground literals l, events e, states s and t such that M,s $dturn$ pre(e) allt (s$R sub e$t imp M,t $dturn$ l) imp (M,s $dturn$ l or $not l˜nomem˜sigma (s,e)$.

Corollary

Foundational inertia is equivalent to foundational minimality.

Foundational minimality says that the literals that hold in all the target states t of event e applied to s are precisely those that already held in s, or else their negations have no foundation in s. In the latter case, these negations notl do not survive in any t, hence by state completeness of ground literals, the corresponding l must hold in every t.

In a recent approach to reasoning about actions based on persistence, Zhang and Foo [ZF93] use essentially a version of foundational inertia as the guiding principle. The difference is that they permit partial state descriptions, but it is not hard to see that these define *sets* of states in our present context. The exact connection is developed in a forthcoming paper where we explore the relationship between *theories* and our semantic model. It suffices here to say that the persistent set is almost the set Cer(s,e) defined above, and hence very conservative.

4. Sharpening Some Properties

Much of the work on event theory assumes that Ram(e) does not suffer from unsubsumed clauses. Because this is so widespread, it deserves a name.

Definition

An event e is definite if Ram(e) does not have any unsubsumed clauses.

Another property that is related to definiteness is determinism.

Definition

An event e applied to state s is deterministic if there is a unique t such that $s R_e t$. Otherwise it is nondeterministic.

Here is an easy consequence of the definitions so far.

Lemma

If e is deterministic then Ram(e)is definite when the strong permissiveness postulate holds.

For the converse, we need a new postulate.

Transition Narrowing Postulate

R_e is restricted to those relations which have minimal target sets for each given source state.

The notion of minimality here is cardinality. This postulate is overlaid on top of the others in the sense that it has to be consistent with them. We have purposely left room for more then one minimum, and remain uncommitted about further possible restrictions that may sharpen "minimum" to "least". Now, suppose Ram(e) is definite, and for simplicity finite (this is not essential). Then say it is $\{l_1, \ldots, l_n\}$. For any source state s, if M,s $\models$ l and $\neg l$ is not one of the literals l_i in Ram(e), by coherent minimality there is at least one target state t in which l holds. The narrowing postulate forces all such l's to be in a *single* target t (in which Ram(e) must hold by the transition axiom), and hence e is deterministic. Hence we have:

Lemma

If Ram(e) is definite, then coherent minimality and transition narrowing implies e is deterministic.

Now we can indicate how another approach to reasoning about actions can be related to these ideas. Winslett [WI88] used a Hamming-like metric on possible models of (partial) theories to select the models that survive an event. Her events are all definite. Similar to our earlier remarks on the Zhang and Foo persistent set approach, the difference between that development and ours here is precisely that between a theory and a set of states that it describes. If each state s in Winslett's theory were subject to coherent minimality and transition narrowing, the aggregate target states are almost her possible models. We say almost because our states are complete. We will show in a later paper how to reconcile them.

Finally, to illustrate the potentiality for this semantics, we address frame axioms. Recall that a situation [MH69] a partial state description in which the truth values of certain fluents (propositional constants, ground literals) are known. The problem in a situation calculus is to provide a syntactic f.o. accounting of situation evolution caused by event sequences. The standard blocks world example of how to specify a stack (X,Y) action is holds(clear(Y) and held(X), s) ra holds(on(X,Y), do(stack(X,Y), s)). Generally pre(e) and post(e) for event e are specified as holds(pre(e), s) ra holds(post(e), do(e,s)). The frame problem is the necessity of also specifying the unchanged fluents that are not affected by the event, e.g., the configuration of other blocks, or the color of any block. Using coherent minimality and transition narrowing as semantic principles gives us a precise standard against which we can assess any purported solution to the frame problem. If we agree to an ontology that, say, accepts coherent minimality but not transition narrowing, then the frame problem is different from the usual one. The utility of semantics is to make explicit the nature of the problems and the existential assumptions that are involved in non-monotonic and related logics.

5. Conclusion and Further Pointers

This work is currently being exploited in several directions. One has already been indicated, i.e. the precise characterization of the semantic assumptions behind contemporary approaches to reasoning about actions, updates and belief revision. Another is the relationship between generic event schemas and their ground instances. This impacts some of previous work on the extraction of system laws from event conditions and will lead to a re-formulation of those ideas. Yet another is the diagnosis of theory failure. An example of what can be done here is to consider pre(e) as absolute and semantically correct. Then a theory of the system may have Pre(e) as a guessed at approximation to pre(e). Under what postulates and what experiments can the incorrectness of Pre(e) be diagnosed?

We hope that the reader has been persuaded that this semantics is worthy of deeper investigation. One may still be sceptical if it was noticed that nothing in this paper used any profound properties of modal semantics. Why not do all this in a first-order model theory? The answer to this has two parts. First, it is much more cumbersome to do the same thing in f.o. model theory despite the isomorphism. But perhaps more important is the second. The power of the method has not yet been fully revealed. As an example, a construction called submodel generation, analogous to the subspace of states reachable by a class of events, is available to model (mixed) event sequences. As another example, the method of filtrations is available to make precise the nature of localness of event influences.

Possible worlds has been taken seriously by programming language semantics researchers for a long time. It is now appropriate for AI to make the linkage.

6. Acknowledgements

This research is supported in part by grants from the Australian Research Council. The second author is partially supported by a SOPF scholarship from the University of Sydney.

Pavlos Peppas has contributed to many discussions out of which these ideas have crystallized. Other members of the Knowledge Systems Group have, by their challenges, forced us to make distinctions that we might not have otherwise made.

7. References

[DS92] del Val, A. and Shoham, Y., "Deriving Properties of Belief Update from Theories of Action", Proc. 10th National AI Conference, AAAI'92, Morgan Kaufmann Publishers, San Mateo, CA., 1991.

[FN71] Fikes, R. and Nilsson, N., "STRIPS: a New Approach to the Application of Theorem Proving to Problem Solving", Artificial Intelligence, 2 (3/4), pp. 189-208, 1971.

[GA88] Gardenfors, P., "Knowledge in Flux", Bradford Books, MIT Press 1988, Cambridge, Massachusetts.

[GI88] Giere, R.N., "Explaining Science", University of Chicago Press, Chicago 1988.

[GI87] Ginsberg, M.L., (ed) "Readings in Non-monotonic Reasoning", Morgan Kaufmann Publishers, San Mateo, CA., 1987.

[GO82] Goldblatt, R., "Axiomatising the Logic of Computer Programming", LNCS 130, Springer-Verlag, Berlin, 1982.

[HM87] Hanks, S. and McDermott, D., "Non-monotonic Logic and Temporal Projection", Artificial Intelligence, 33, pp. 379-412, 1987.

[HC84] Hughes, G.E. and Cresswell, M.J., "A Companion to Modal Logic", Methuen 1984, London.

[KM91] Katsuno, H. and Mendelzon, A.O., "Onthe Difference between Updating a Knowledge Base and Revising It", Proc. 2nd International Conference on Principles of Knowledge Representation and Reasoning, KR'91, Morgan Kaufmann Publishers, San Mateo, CA., 1991.

[LS91] Lin, F. and Shoham, Y., "Provably Correct Theories of Action", Proc. 9th National AI Conference, AAAI'91, Morgan Kaufmann Publishers, San Mateo, CA., 1991.

[MC86] McCarthy, J., "Applications of Circumscription to Formalizing Common-Sense Knowledge", Artificial Intelligence, 28, pp. 86-116, 1986. Also in Ginsberg [GI87].

[MH69] McCarthy, J. and Hayes, P., "Some Philosophical Problems from the Standpoint of Artificial Intelligence", in Machine Intelligence 4, ed. B. Meltzer and D. Michie, pp. 463-502, Edinburgh University Press, Edingurgh, 1969. Also in Ginsberg [GI87].

[PE93] Peppas, P., C.S. PhD Thesis, University of Sydney, 1993.

[RE80] Reiter, R., "A Logic for Default Reasoning", Artificial Intelligence, 13, pp. 81-132, 1980. Also in Ginsberg [GI87].

[SE93] Sandewall, E., "Features and Fluents", forthcoming book, 1993.

[WI88] Winslett, M., "Reasoning about Actions using a Possible Models Approach", Proc 7th National AI Conference, AAAI'88, Morgan Kaufmann Publishers, San Mateo, CA., 1988.

[ZF93] Zhang, Y. and Foo, N., "Reasoning about Persistence: A Theory of Actions", (to appear), Proceedings of the 13th International Joint Conference on Artificial Intelligence, IJCAI'93, France.

[ZE84] Zeigler, B.P., "Multifacetted Modelling and Discrete Event Simulation", Academic Press, Orlando, Florida 1984.