

Preliminary Insights on Temporal Approximation

Aarati Parmar

Department of Computer Science,
 Gates Building, 2A wing
 Stanford University, Stanford, CA 94305-9020, USA
 aarati@cs.stanford.edu

Abstract

In this paper we present some preliminary formulations of temporal approximations and refinements. These are respectively the processes of abstracting away details about temporal objects, and the processes of adding more detail to them. Intelligent robots operating in the world will not only need to reason temporally, but they will need to be able to approximate or refine their reasoning as necessary in order to operate efficiently. The four kinds of refinements/approximations we study are treating ramifications as internal events, already proposed in (McCarthy, 2002), elaborations of narratives, expansions of events, and increasing the predictive capacity of theories.

1 Introduction

Various aspects of temporal reasoning for AI have been formalized, including inertia, hypotheticals, concurrency, etc. In this paper we enumerate and give preliminary formalizations of different kinds of *temporal approximations and refinements*. We define *temporal approximation* as a transformation of a theory of action and change to another which is simpler and/or more salient to the task at hand. A *temporal refinement* is the opposite transformation and will add more detail about a sequence of events, or provide explanations for them. We believe that these mechanisms will be necessary for intelligent robots operating in the real world; there is so much information available that a robot will need to be able to abbreviate it in order to operate efficiently. Robots will also need to be able to refine their theories of the world to accommodate new phenomena.

The four different kinds of temporal approximations/refinements we explore in this paper are:

1. **Ramifications as internal events.** This has already been introduced in (McCarthy, 2002), and is the idea that we can interpret static constraints as resolutions of *internal events*. We interpret the internal events as imbalances in the system which move it towards a steady state; a static constraint simply describes this steady state where nothing changes.¹ McCarthy uses his buzzer example

¹We contrast internal events with external events, which are defined by (McCarthy, 2002) as events outside of the given system which are not explained by the system and usually interpreted as actions.

to demonstrate that not only is this concept a nice way to reify and elaborate state changes, but a necessary one when phenomena do not reach a steady state. More will be said in Section 3.

2. **Elaboration of a narrative.** One can further refine a narrative by adding more events to it, either between known events or at the same time as them. We will call this an example of a *dense refinement* because we are packing in extra event occurrences (and whatever facts follow from these events happening) between known events. For example, in the Missionaries and Cannibals problem (McCarthy, 1998), one could further elaborate that *Cannibal1* sneezed while rowing from *bank1* to *bank2*, and after one of the crossings, all of the characters stopped and took a nap. A diagram of the general idea is shown in Figure 1. It would also be interesting to go in the opposite direction; that is, how to eliminate irrelevant facts about a sequence of events to get a simpler, but still consistent story. We address this kind of refinement in Section 4.




Figure 1: A diagram of a dense refinement. Circles correspond to situations, arrows are transitions between them. The e_i s refer to events occurring at the situation.

3. **Expansion of events.** An event, say the act of buying a box of tissue, can be expanded into an entire sequence of events: walking into a store, finding the tissue, putting it on the counter, and paying for it (McCarthy, 1999). We call this an *expansive refinement*. It differs from the dense refinement explained earlier in that one situation is mapped to a sequence of situations. Figure 2 should demonstrate how it is different. In general a situation s can be blown up into a more detailed sequence of situations, which together elaborate the simpler truths of s . Section 5 has more on this.




Figure 2: A diagram of an expansive refinement. s_1 is expanded to a sequence of situations s'_1, \dots, s'_6 .

4. **Increasing the predictive capacity of theories.** Any common-sense theory T describing aspects of the world will be inherently incomplete with respect to some detail or other. That is, T will explain some aspects of the world, but there will always be other phenomena T cannot explain. In order to reason about these phenomena, T will have to treat them as a “black box,” that is, it will only be able to accept observations about the inexplicable phenomena, as opposed to predicting or explaining them. For example, (McCarthy, 2002) relates the story of the stuffy room, where the room is stuffy iff both vents *vent1* and *vent2* in the room are blocked. The article includes an elaboration where Pat hates to be in a stuffy room and will unblock *vent2* to air out the room. Mike on the other hand gets cold when a vent is unblocked and will block it. Without knowledge of Pat and Mike’s preferences, our theory T will have no ability to infer or explain Pat and Mike’s actions, and will have to take their actions as given. However, if the axioms about their behavior are added to T , T will predict what will happen. By adding these axioms we have pushed back our curtain of ignorance about the rest of the world and increased our predictive power, depicted in Figure 3. We address this more in Section 6.




Figure 3: The intuitions behind increasing predictive capacity of theories. The unshaded box represents the part of the world T can predict. The shaded areas are the parts T cannot explain. T can be expanded to predict phenomena in the lightly shaded region by adding more occurrence axioms. External events (lightning bolt) emanate from the shaded areas and cannot be explained by T .

We can view the temporal approximations/refinements above as a transformation on a theory that will induce some other transformation on the (intended) models of our theory. We hope to introduce formulations that require only very simple operations on the theory to get the intended models (like adding axioms).

As for language, we use the situation calculus augmented with event occurrences. We note that these approximations should be formalizable in other languages that represent time and change; we only use this variant of situation calculus because of our familiarity with it and the utility of its features. Our exposition shares many of the motivations as in (McCarthy, 2002), although the technical details differ considerably.

2 Preliminaries

We use the situation calculus which consists of situations, events, and fluents. Situations are snapshots of the world which have some temporal extent. *Fluents* are facts which hold at a particular situation, and are formalized as relation symbols with one free situation variable. $Result(e, s)$ is used to denote the situation resulting from event e occurring in s . For maximum flexibility of interpretation we do not include any of the arboreal axioms used in various versions (like Toronto’s) of situation calculus, nor do we include the induction axiom, although we may find them useful for specific problems. Instead, we add two new symbols to the language of situation calculus: $Occurs(e, s)$ and $Next(s)$. $Occurs(e, s)$ is a predicate that asserts that event e occurs in situation s , while $Next(s)$ points to the resulting situation after whatever events occur at s . For example if $Occurs(fall(domino), s)$ then $Upright(domino, s)$ but $\neg Upright(domino, Next(s))$.

Right off we can see that $Occurs$ allow us to express concurrent events. Also, $Occurs(e, s)$ differs from $Result(e, s)$ in that $Occurs(e, s)$ asserts e actually does happen in s , as opposed to the result if e happened in s . It may seem incongruous to allow both notions of actual and hypothetical events, but we believe this is needed to represent (necessarily) incomplete knowledge of the world; we can talk about a distinguished actual set of events while still representing different hypothetical sequences within the language, as pointed out by (Pinto, 1998). For example, we can assert, without inconsistency, both

$$Occurs(book\text{-}ticket(Edmonton), \\ Result(accept(reviewers, paper)), s) \\ \wedge \\ Occurs(book\text{-}ticket(Las\ Vegas), \\ Result(reject(reviewers, paper)), s),$$

reflecting the author’s plans to go to Vegas if she does not have to go to AAAI to discuss this paper. But we cannot express what events actually do happen in reality. $Occurs(e, s)$ asserts e occurs in s , but not that s itself is actual – it could be hypothetical. In order to delineate an actual sequence of situations we would have to add a predicate such as $Actual(s)$, introduced in (Pinto and Reiter, 1993).

Our language is that of standard situation calculus, which includes the three mutually exclusive sorts: *Sit*, *Eve*, and *Flu*. Throughout this paper we take s to be a situation variable, and e an event variable.

\mathcal{L} consists of finitely many relational (*fluent*) symbols $F(s)$. We also include a relation $Occurs(e, s)$. In addition to the usual function $Result : Sit \times Eve \rightarrow Sit$, we also include a function $Next : Sit \rightarrow Sit$.

We copy much of the notation used in (McIlraith, 2000). Here we only explain the notation, as we use different axioms in the sections below. $\gamma_F^+(e, s)$, $\gamma_F^-(e, s)$, $\nu_F^+(s)$ and $\nu_F^-(s)$ are all abbreviations for *simple formulae*, which is a formula of situation calculus in which the only situation variable in the entire expression is the free variable s . Essentially, a simple formula $\phi(s)$ expresses some property ϕ of a situation s , instead of a relation between situations. $\gamma_F^+(e, s)$ abbreviates the conditions for F to hold in $Result(e, s)$, while $\gamma_F^-(e, s)$ the conditions for F not to hold. $\nu_F^+(s)$ are the conditions for F to hold in s , while $\nu_F^-(s)$ are those for $\neg F(s)$.

We ignore $Poss(e, s)$ and assume events are axiomatized so that if the preconditions for e are not met, then nothing changes: $F(Result(e, s)) \iff F(s)$.

Of course, to prevent inconsistencies we cannot have the conditions cause both $F(s)$ and $\neg F(s)$. This is the *consistency assumption* referred to in (McIlraith, 2000) and is formulated in terms of our language in (1), for each fluent F :

$$\neg[(\gamma_F^+(e, s) \vee \nu_F^+(Result(e, s))) \wedge (\gamma_F^-(e, s) \vee \nu_F^-(Result(e, s)))] \quad (1)$$

Finally, we use some notation and terminology to make our statements more concise. A *sequence of situations* is a sequence s_1, \dots, s_n , abbreviated \bar{s} , such that $s_{i+1} = Next(s_i) \vee (\exists e)s_{i+1} = Result(e, s_i)$. We define an order $<$ on situations s and t , where $s < t$ iff there is a sequence of situations s_1, \dots, s_n , $n > 1$ where $s_1 = s$ and $s_n = t$. Note that $<$ is irreflexive. A *static constraint* is an axiom relating only facts about one situation s ; it is of the form $(\forall s)\phi(s)$, where ϕ is a simple formula.

3 Ramifications as Internal Events

We explore this kind of refinement first, to acquaint the reader further with our intended interpretations of $Occurs$ and $Next$. The idea already explored in (McCarthy, 2002) is to treat static constraints as internal events. Our methodology treats imbalances in a static constraint as internal events which move in the direction of rectifying the imbalance.² So

²Our intuition comes from calculus. The forces (internal events) acting on a system can be interpreted as a gradient (imbalance in the system). Hence the progression of the system through time can be represented as a set of differential equations. Usually the solution to the differential equations has some steady-state equilibrium. So in these terms, our task is, given the steady-state of a system, to discover the gradient field, or equivalently forces, which moved the system to the equilibrium. There are many possible gradient functions, but we believe that the syntactic form of

for example, the Stuffly Room problem has a static constraint of the form:

$$Blocked(vent1, s) \wedge Blocked(vent2, s) \implies Stuffly(s) \quad (2)$$

This would be transformed to the *event occurrence axiom*

$$[Blocked(vent1, s) \wedge Blocked(vent2, s) \wedge \neg Stuffly(s)] \implies Occurs(Becomes-Stuffly, s),$$

along with the *event effect axiom*

$$Occurs(Becomes-Stuffly, s) \implies Stuffly(Next(s)).$$

We see that this surgery on ramifications allows us to represent directionality of ramifications *without* resorting to special implications or higher-order reasoning. We also note that the implication in (2) is interpreted by humans as more than material implication; it is causal, as noted in (McIlraith, 2000). To illustrate this we perform a similar surgery on a logically equivalent form of (2):

$$[Blocked(vent1, s) \wedge Blocked(vent2, s) \wedge \neg Stuffly(s)] \implies Occurs(Becomes-Unblocked(vent2), Next(s)).$$

This however corresponds to the spurious change-minimizing model of the Stuffly Room problem.

Let us describe how to refine a theory with static constraints to include internal events. Let F be a fluent symbol. T_{sc} is a first-order theory which satisfies the following properties:

1. T_{sc} contains effect axioms of the form:

$$\begin{aligned} \gamma_F^+(e, s) &\implies F(Result(e, s)) \\ \gamma_F^-(e, s) &\implies \neg F(Result(e, s)) \end{aligned} \quad (3)$$

2. T_{sc} also has static constraints of the form:

$$\begin{aligned} \nu_F^+(s) &\implies F(s) \\ \nu_F^-(s) &\implies \neg F(s) \end{aligned} \quad (4)$$

where each implication is to be read causally. Note that we can have non-causal constraints by letting $F \equiv \perp$.

3. We also include the proper UNAs for events.
4. Finally, we assume the *consistency assumption* given above in (1).

From T_{sc} we can compute a more refined theory $\mathcal{D}(T_{sc})$ in the same language of T_{sc} , extended with the relation $Occurs(e, s)$ and function $Next(s)$. $\mathcal{D}(T_{sc})$ will reinterpret T_{sc} , replacing static constraints with internal events. For each fluent F we also add the events $on(F)$ and $off(F)$ to our language. $\mathcal{D}(T_{sc})$ then can be constructed from T_{sc} in the following way:

the static constraints gives a clue as to the direction of the gradient (discussed next).

1. $\mathcal{D}(T_{sc})$ contains the same effect axioms as in T_{sc} .
2. For each static constraint $\nu_F^+(s) \implies F(s)$ in T_{sc} we have in $\mathcal{D}(T_{sc})$

$$Occurs(on(F), s) \iff_{def} \nu_F^+(s) \wedge \neg F(s)$$

and for each $\nu_F^-(s) \implies \neg F(s)$, $\mathcal{D}(T_{sc})$ contains

$$Occurs(off(F), s) \iff_{def} \nu_F^-(s) \wedge F(s).$$

Hence we have explicitly represented the imbalance in a static constraint by the occurrence of an event.

3. The same unique names axioms as in T_{sc} , extended to cover the new events $on(F)$ and $off(F)$.
4. The consistency assumption (1).
5. We must also relate $Occurs$ to $Next$. For each fluent F , we formalize the definition of $on(F)$ and $off(F)$ by the “event successor state axiom”:

$$F(Next(s)) \iff Occurs(on(F), s) \vee \neg Occurs(off(F), s) \wedge F(s)$$

Given that $\mathcal{D}(T_{sc})$ is meant to represent the progression of the system in T_{sc} from any s to its steady state along the trajectory $Next(\dots(Next(s)))$, it makes sense to have a function which points to the situation where steady state is reached, if it exists. We define a function $Next^*(s)$, which points to the first situation after or including s in which the system has reached quiescence:

$$Next^*(s) \iff_{def} (\mu t > s) \neg (\exists e) Occurs(e, t) \quad (5)$$

A logically equivalent definition of $Next^*$ is $(\mu t > s) (\forall F) (F(Next(t)) \iff F(t))$. The example of the buzzer shows that it is possible that $Next^*$ can be undefined. But when it is defined, it definitely points to some steady state equilibrium.

The question is whether it is the same as that prescribed by T_{sc} , especially under various solutions to the frame problem. We can attempt to answer this question in the context of (McIlraith, 2000)’s closed form solution to theories with ramifications. Our theory T_{sc} will fit (McIlraith, 2000)’s syntactic requirements, as long as it is *solitary stratified*. That is, we can assign a level to each fluent F so that for every implication of the form $\mathcal{D}(s) \implies F(s)$ and $\mathcal{E}(s) \implies \neg F(s)$ in T_{sc} , if the fluent F is of level i , then the fluents in the formulas $\mathcal{D}(s)$ and $\mathcal{E}(s)$ all have levels strictly less than i . If F has level 1, $\mathcal{D}(s)$ and $\mathcal{E}(s)$ cannot reference any fluents and cannot say anything about the state of s . From now on we write a fluent F of level i as F_i .

We write (McIlraith, 2000)’s closed form solution of T_{sc} as a theory $\mathcal{S}(T_{sc})$, which will have successor state axioms of the form:

$$F_i(Result(e, s)) \iff \gamma_{F_i}^+(e, s) \vee \nu_{F_i}^+(Result(e, s)) \vee F_i(s) \wedge \neg \gamma_{F_i}^-(e, s) \wedge \neg \nu_{F_i}^-(Result(e, s)) \quad (6)$$

with the same UNAs and the consistency assumption. The solitary stratification guarantees that the terms on the right hand side containing $Result$ ($\nu_{F_i}^+$ and $\nu_{F_i}^-$), can be rewritten, using the successor state axioms for the fluents in them, to not mention $Result$. Then these axioms will correspond to true successor state axioms which can be properly progressed and regressed to one’s heart’s content.

We can apply a similar transformation to $\mathcal{D}(T_{sc})$, to end up with a theory $\mathcal{S}(\mathcal{D}(T_{sc}))$ that has the form:

$$\begin{aligned} F_i(Result(e, s)) &\iff \gamma_{F_i}^+(e, s) \vee F_i(s) \wedge \neg \gamma_{F_i}^-(e, s) \\ Occurs(on(F_i), s) &\iff \nu_{F_i}^+(s) \wedge \neg F_i(s) \\ Occurs(off(F_i), s) &\iff \nu_{F_i}^-(s) \wedge F_i(s) \\ F_i(Next(s)) &\iff Occurs(on(F_i), s) \vee \neg Occurs(off(F_i), s) \wedge F_i(s) \end{aligned} \quad (7)$$

$\mathcal{S}(T_{sc})$ computes all changes due to e , both from effects and ramifications, immediately into $Result(e, s)$. $\mathcal{S}(\mathcal{D}(T_{sc}))$ on the other hand first only computes effects into $Result(e, s)$, and then lets ramifications propagate as internal events, until it reaches quiescence in the situation $Next^*(Result(e, s))$.

When T_{sc} is solitary stratified, we can show $Next^*$ will be defined in $\mathcal{S}(\mathcal{D}(T_{sc}))$, and in fact will be $Next^n$, where n is the highest level of stratification. The question is however whether the steady state that is reached in $\mathcal{S}(\mathcal{D}(T_{sc}))$ after an external event e is applied to the system is the same as $Result(e, s)$ in $\mathcal{S}(T_{sc})$. That is, is it the case that:

$$\begin{aligned} \mathcal{S}(T_{sc}) \models F(Result(e, s)) &\iff \\ \mathcal{S}(\mathcal{D}(T_{sc})) \models F(Next^*(Result(e, s))) &? \end{aligned}$$

For level 1 fluents this can be shown to hold, but for fluents at greater levels it is not clear because the internal events may progress and interact in strange ways. Timing issues may also arise, in that the values of fluents along the progression of $Next$ are never all at the required values at the same time for another fluent to change its value. This richer treatment of static constraints needs to be studied further.

4 Dense Refinements

In this section we define a *dense refinement* as when more intervening and even co-occurring events are added to a given sequence of events. Adding intervening events can add more detail to a story, but should not refute any parts of the story. We formalize it as follows:

Let \bar{s} and \bar{t} be sequences of situations, and T and T' two theories, such that $\mathcal{L}(T) \subseteq \mathcal{L}(T')$. We say \bar{t} in T' *densely refines* \bar{s} in T if we can find a mapping $\sigma : \bar{s} \rightarrow \bar{t}$ such that:

1. $(\forall s, s' \in \bar{s}) (s < s' \implies \sigma(s) < \sigma(s'))$: σ is order-preserving.
2. $(\forall s \in \bar{s}) (T \models Occurs(e, s) \implies T' \models Occurs(e, \sigma(s)))$: The events in \bar{s} must also happen in \bar{t} .

3. $(\forall s_1, s_2 \in \bar{s})(\sigma(s_2) = \text{Result}(e, \sigma(s_1)) \implies s_2 = \text{Result}(e, s_1))$: Our refinement T' of T may transform situation transitions of the form $s' = \text{Result}(e, s)$ to those of the form $\sigma(s') = \text{Next}(\sigma(s))$, because T' is a more powerful predictor. That is, rather than having to be told that s' is the result of e in s , T' can predict this. (We study this idea further in Section 6.)

But, we will *not* allow transformations in the other direction: if T can predict that s' (and associated events) is the next situation after s , then T' must as well, and not resort to suddenly “not knowing” and using *Result*. This is what the above formula states.

4. $(\forall \phi \in \mathcal{L}(T))(\forall s \in \bar{s})(T \models \phi(s) \iff T' \models \phi(\sigma(s)))$. Note that only the situations in the image of σ must agree with their mapped situations in \bar{s} , and only over the formulas in $\mathcal{L}(T)$. It might be interesting to relax this restriction by replacing the equivalence with a right implication.

Intuitively, the subsequence $\sigma(\bar{s})$ is a *skeleton* of \bar{t} which corresponds to the original sequence \bar{s} modulo the language $\mathcal{L}(T)$ and possibly some extra events. Requirement 4 will indirectly allow only *consistent* insertion of events between those in \bar{s} . For example, let \bar{s} be the sequence of boat-crossings to get three missionaries and three cannibals across a river subject to the usual constraints. We cannot consistently elaborate s by inserting one boat-crossing between two others, because then the boat will be on the wrong side of the river, and either the inserted event or the one after it will be inconsistent, so that the fluents across the sequences will not match up.

What is important now is to check that this formulation of dense refinements is correct, and what mileage we can get out of having such relations in our theory. Instead of having two separate theories T and T' espouse these different sequences, we could talk about both sequences *within the theory* by putting T and T' in different *contexts* within one larger theory. It would be interesting then to see how *lifting axioms* could move between the two representations. It would also be nice to characterize the forms of T and T' that allow easy specification of dense refinements.

5 Expansive Refinements

We can conceive of a refinement of events where one event is expanded to a sequence of events, as in the tissue buying example above. So for example, say we have the trivial sequence of situations $\bar{s} = s_1$, and $\text{Occurs}(\text{Buys}(\text{tissue}), s_1)$. This act can be refined as the sequence of events

$$\begin{aligned} & \text{Occurs}(\text{Enter}(\text{store}), s'_1) \wedge \\ & \text{Occurs}(\text{Get}(\text{tissue}), s'_2) \wedge \\ & \text{Occurs}(\text{Put}(\text{tissue}, \text{counter}), s'_3) \wedge \\ & \text{Occurs}(\text{Pay}(\text{clerk}, \text{tissue}), s'_4). \end{aligned}$$

Instead of being an event, the concept “buys tissue” becomes a fact (fluent) (process?) about the situations s'_1, \dots, s'_4 .

The immediate formalization of this idea seems trivial; simply map the situation s_1 by fiat to the sequence of situations s'_1, \dots, s'_4 . In order to group these situations s'_1, \dots, s'_4

together we should label them all as situations during which the tissue-buying is happening. However it seems there is at least a little fine structure which is being ignored. We would like to know *how we know* when an object should be treated as an event, versus a fluent. Syntax gives us no hints: if we rewrite fluents $F(s)$ to the form $\text{Holds}(f, s)$, they are no different (syntactically speaking) from $\text{Occurs}(e, s)$, except that we imagine fluents as having duration and inertia, while events are point-like and cause change.

However there is a well-known transformation between events and fluents: (Pinto, 1994) shows that a point-like event can always be transformed to a one with duration by mapping it to a fluent whose truth is turned on and off by point-like events. So we have in our arsenal at least one known mechanism for expanding events to those with duration, which might as well be treated as a fluent.

Studying the tense of the verb encoding the event gives us important clues as well. When “buy” is in the simple tense, as in “Bob buys a box of tissue” the event is point-like, and definitely of the form $\text{Occurs}(e, s)$. On the other hand, the statement “Jane is buying a box of tissue,” where “buy” is in the present progressive tense, corresponds more to the sequence of situations s'_1, \dots, s'_4 . Other properties of verbs may be useful in at least deciding whether a verb corresponds to an event or sequence of events. *Stative verbs* (van Eijck and Kamp, 1997) are those which describe state rather than an event. Examples include know, like, wish, etc. *Dynamic verbs* on the other hand describe events. The progressive tense only applies to dynamic verbs, and seems (to the author at least) to change them to stative ones.

Another dimension to explore is the relation between the situation s_1 and s'_1, \dots, s'_4 . Intuitively, the sequence of situations *realizes* the event $\text{Occurs}(\text{Buys}(\text{Tissue}), s_1)$. That is, it provides some justification or explanation for the event. This idea of *realizability* is the same as that used in constructive mathematics: briefly, an object p *realizes* a first order formula ϕ if it is some effective procedure for showing ϕ holds. p can be a proof of ϕ , or a program that computes ϕ , etc. A key goal of exploring such *expansive refinements* would be to define some kind of common sense notion of a sequence of situations realizing another.

Finally, we can go in the opposite direction and regard the approximation as a compression of situations which are too close and similar (or irrelevant): if the granularity (Hobbs, 1985) between a set of situations is below a threshold, we might as well compress them into one. s_1 then is a compression of s'_1, \dots, s'_4 with respect to certain metrics, yet to be discovered.

6 Increasing the Predictive Capacity of Theories

In this section we attribute some special intuitive interpretations to $\text{Occurs}(e, s)$, $\text{Next}(s)$ and $\text{Result}(e, s)$. In particular we contrast the progression functions Next and Result . As mentioned before, Result is inherently by syntax a hypothetical function, in that it describes *the result of what would happen were e to be applied to s* , but it makes no statement by itself as to whether it does happen in s .

$Occurs(e, s)$ on the other hand, does provide this functionality, and in conjunction with $Next(s)$ will give us the resulting situation if e occurred in s .³ $Next$ “knows” what events occur in s (thanks to $Occurs$), and can progress s without any help or outside information.

We further stress this point by considering the discussion of internal versus external events given in (McCarthy, 2002). External events do not have occurrence axioms in the theory, while internal events do. Hence external events can only be introduced into the theory by $Result$ and can never be predicted (which would require the use of $Occurs$ and $Next$). External events are therefore identified as being actions performed by actors. In fact, the initiators of an external event are almost always termed as actors, because their volitions are not something we can predict.

We stress this difference because it is crucial part of the semantics behind the refinement we describe in this section, that of increasing the predictive capacity of theories. Generally, a sequence of events will have to be “spoon-fed” to a theory T as a series of statements $s_2 = Result(e_1, s_1)$, $s_3 = Result(e_2, s_2)$, \dots , $s_n = Result(e_{n-1}, s_{n-1})$. Then T 's usual task is to entail what facts are true in the sequence of situations s_1, \dots, s_n , either with the use of successor state axioms or some non-monotonic reasoning.

But never does T try to *predict* this action sequence. Of course, in many domains, prediction may not make any sense. But in many narratives many events that are given can be predicted by T , if *only it knew more about the actors and their motivations in the narrative*. Let us present the example of the stuffy room with Pat and Mike formally.

A non-predictive theory T might contain axioms:

$$\begin{aligned}
Blocked(v, Result(e, s)) &\iff \\
&(\exists p)e = Block(p, v) \vee \\
&Blocked(s) \wedge \neg(\exists p)e = Unblock(p, v) \\
Stuffy(s) &\iff \\
&Blocked(vent1, s) \wedge Blocked(vent2, s) \\
Blocked(vent1, s_0) \wedge Blocked(vent2, s_0) &
\end{aligned} \tag{8}$$

Then, given a sequence of situations

$$\begin{aligned}
s_1 &= Result(UnBlock(Pat, vent2), s_0), \\
s_2 &= Result(Block(Mike, vent2), s_1), \\
s_3 &= Result(UnBlock(Pat, vent2), s_2), \dots
\end{aligned}$$

the only extra information T can provide is that $Stuffy \wedge Blocked(vent2) \wedge Blocked(vent1)$ in s_2, s_4, \dots and $\neg Stuffy \wedge \neg Blocked(vent2) \wedge Blocked(vent1)$ in s_1, s_3, \dots . But if T had more axioms about Pat and Mike's preference, it could predict this narrative. Consider adding the event occurrence and successor state axioms in (9) to T :

$$\begin{aligned}
Stuffy(s) &\iff Occurs(UnBlock(Pat, vent2), s) \\
\neg Stuffy(s) \wedge \neg Blocked(vent2, s) &\iff \\
&Occurs(Block(Mike, vent2), s) \\
\neg Stuffy(s) \wedge \neg Blocked(vent1, s) &\iff \\
&Occurs(Block(Mike, vent1), s) \\
Blocked(v, Next(s)) &\iff \\
&(Blocked(s) \wedge \neg(\exists p)Occurs(Block(p, v), s)) \vee \\
&(\exists p)Occurs(UnBlock(p, v), s)
\end{aligned} \tag{9}$$

Now, the sequence $s_0, Next(s_0), Next^2(s_0), Next^3(s_0), \dots$ will mirror and predict the narrative $s_0, s_1, s_2, s_3, \dots$. The motivations of Pat and Mike are nothing but black boxes in T , but are revealed and then predicted along the $Next$ sequence in $T \cup (9)$, all by adding a few axioms using $Occurs$ and $Next$.

Another *Drosophila* is flipping a fair two-sided coin. We can write and solve a complicated set of physics equations that will describe and predict the result of the coin flip, but for most of us it is easier and more practical to treat the physics as a black box, and either take the coin's end state as given, or as the result of a probabilistic process. One can use our approach to provide foundations for how non-deterministic (probabilistic) methods are often just an approximation to a more complicated deterministic inference.

7 Conclusions and Discussion

In this paper we have guided the reader through preliminary formalizations of some temporal approximations and refinements that will be necessary for intelligent robots operating in the world. Hopefully we have illustrated how these mechanisms implement some important and interesting modes of reasoning. Of course we do not mean to tease the reader with our very incomplete results, and mean to say much more about all of these concepts in the near future.

8 Acknowledgments

The ideas in this paper are the fruits of extensive discussions with John McCarthy and Sheila McIlraith.

References

- Hobbs, J. R. (1985). Granularity. In *International Joint Conference on Artificial Intelligence (IJCAI'85)*, pages 432–435.
- McCarthy, J. (1998). Elaboration Tolerance⁴. In *In Proceedings of the Fourth Symposium on Logical Formalizations of Common Sense Reasoning*.
- McCarthy, J. (1999). Personal communication.
- McCarthy, J. (2002). Actions and Other Events in Situation Calculus⁵. In *Proceedings of KR 2002*. To be published.
- McIlraith, S. (2000). An axiomatic solution to the ramification problem. *Artificial Intelligence*, 116:87–121.

³Note that $Occurs$ still is not meant to assert that e actually happens in the world; s itself could be hypothetical.

⁴<http://www-formal.stanford.edu/jmc/elaboration.html>

⁵<http://www-formal.stanford.edu/jmc/sitcalc/sitcalc.html>

Pinto, J. (1994). *Temporal Reasoning in the Situation Calculus*⁶. PhD thesis, Dept. of Computer Science, Univ. of Toronto.

Pinto, J. and Reiter, R. (1993). Temporal reasoning in logic programming: A case for the situation calculus. In *Proceedings of the Tenth International Conference on Logic Programming*, pages 203–221.

Pinto, J. A. (1998). Occurrences and Narratives as Constraints in the Branching Structure of the Situation Calculus⁷. *Journal of Logic and Computation*, 8:777–808.

van Eijck, J. and Kamp, H. (1997). Representing discourse in context. In van Benthem, J. and ter Meulen, A., editors, *Handbook of Logic and Language*, pages 180–237. North-Holland.

⁶<http://www.cs.toronto.edu/cogrobo/jpThesis.ps.Z>

⁷<http://citeseer.nj.nec.com/241020.html>