Context-Sensitive Event Occurrence Minimisation

Murray Shanahan

Department of Computing,
Imperial College of Science, Technology & Medicine,
180, Queen's Gate, London SW7 2BZ, ENGLAND
email: mps@doc.ic.ac.uk

Abstract

When reasoning about an incompletely known narrative of events, it's necessary to employ event occurrence minimisation in order to get the expected results from standard temporal projection techniques. However, a blanket default assumption that the known event occurrences are the only event occurrences is rarely legitimate. A finer grain of minimisation is required, which takes into account the possibility of unknown events when the blanket default assumption is inapplicable. This paper outlines a mechanism for achieving this, based on John McCarthy's recent proposal for formalising context.

Introduction

A *narrative* is a course of events about which we may have incomplete information. Some formalisms for reasoning about action are narrative based, such as the Event Calculus [Kowalski & Sergot, 1986], [Shanahan, 1995]. Other formalisms, in particular the Situation Calculus [McCarthy & Hayes, 1969], which is based on a branching tree of hypothetical situations rather than a single narrative line, need to be extended to cope with narrative information [Pinto & Reiter, 1993], [Miller & Shanahan, 1994].

In either case, the need arises to minimise event occurrences. When we describe a narrative of events, we generally want to assume that no events occur other than those which we know about. This enables us to employ standard temporal projection techniques, including the commonsense law of inertia, to determine what fluents hold at what times. Without event occurrence minimisation, we would have to admit the possibility of unknown events intervening in the narrative and perturbing the course of history. This would prevent us from drawing the conclusions we intuitively require.

On the other hand, we need to avoid over-zealous minimisation. The assumption that no events occur other than those we know about is frequently wrong. For example, on the basis of a noise coming from the adjacent office, I may assume that someone is in there. But because my knowledge of what is happening on the other side of a brick wall is incomplete, it wouldn't be safe for me to apply the commonsense law of inertia to assume that they are still there ten minutes later.

This paper exploits some of McCarthy's ideas about context [McCarthy, 1993] to make event occurrence minimisation more sensitive, permitting honest ignorance about the occurrence of certain events at the same time as allowing default assumptions about the non-occurrence of others. This is achieved by allowing multiple narratives of events, and associating separate narratives with separate contexts. Then minimisation takes place within certain contexts but not within others.

1. Event Occurrence Minimisation

Let's begin with a brief outline of a framework for describing narratives and for drawing conclusions about those narratives via standard Situation Calculus domain theories. This outline will follow the approach in [Miller & Shanahan, 1994]. A special sort for time points is introduced, in addition to the usual sorts for situations, fluents, and actions. The predicate Happens is used to represent event occurrences. The formula Happens(e,t) represents that an event of type e happens at time point t.[1] For example, to represent that the sun rose at 5 O'clock, we could write,

Happens(5,Sunrise).

Situation Calculus domain theories are presented in the usual way familiar from the literature, using a Result function. The term Result(e,s) denotes the situation that obtains after an event (or action) of type e occurs in situation s. Fluents will be reified through a Holds predicate. The formula Holds(f,s) denotes that fluent f holds in situation s. For example, to represent the effect of sunrise and sunset on the prevailing light, we could write,

Holds(Light,Result(Sunrise,s))

¬ Holds(Light,Result(Sunset,s)).

We need a solution to the frame problem. I will adopt Baker's circumscriptive solution [1991], although any other approach would serve equally well. We have the following frame axiom.

$$[\mathrm{Holds}(f,\mathrm{Result}(e,s)) \leftrightarrow \mathrm{Holds}(f,s)] \leftarrow \neg\mathrm{Ab}(e,f,s) \quad \text{(F1)}$$

According to Baker's circumscription policy, the predicate Ab is minimised and the Result function is allowed to vary. Also, Baker's approach requires an "existence of situations axiom", which guarantees that a situation exists for every possible combination of fluents. For further details of Baker's technique, the reader is referred to [Baker, 1991].

[1] Throughout the paper, predicate, function and constant symbols will begin with upper-case letters, and variables will begin with lower-case letters. All variables are universally quantified unless otherwise indicated.

From the description of a narrative of events in terms of Happens formulae, we want to be able to use a standard Situation Calculus domain theory to draw conclusions about which fluents hold at any given time point. To do this, we introduce a function State, such that the term State(t) denotes the situation which obtains at time t. We then have the following axiom.

$$\begin{aligned} &State(t1) = Result(e1,State(t2)) \leftarrow \\ &\quad [Happens(e1,t2) \wedge t1 \geq t2 \wedge \\ &\quad\quad \neg \exists e2,t3\ [Happens(e2,t3) \wedge \\ &\quad\quad\quad [e2 \neq e1 \vee t2 \neq t3] \wedge t2 < t3 < t1]] \end{aligned} \qquad (N1)$$

In order to be able to draw any useful conclusions from this axiom, Happens will have to be minimised, and State allowed to vary. To see why, consider the Yale Shooting scenario [Hanks & McDermott, 1987]. We will represent the actual narrative of events — Load, Sneeze, Shoot.[2] The usual approach is to examine the properties of a Result term, namely,

Result(Shoot,Result(Sneeze,Result(Load,S0))).

But with the narrative approach, we have sentences such as the following.

Happens(Load,5)

Happens(Sneeze,10)

Happens(Shoot,15)

The effects of loading and shooting are captured by the following sentencs. I also introduce an Unload action, for reasons that will become apparant.

$$Holds(Loaded,Result(Load,s)) \qquad (Y1)$$

$$\neg Holds(Alive,Result(Shoot,s)) \leftarrow Holds(Loaded,s) \qquad (Y2)$$

$$\neg Holds(Loaded,Result(Unload,s)) \qquad (Y3)$$

$$Holds(Alive,S0) \qquad (Y4)$$

$$\neg Holds(Loaded,S0) \qquad (Y5)$$

Without the minimisation of event occurrences, we cannot draw the intuitive conclusion from these sentences that Alive does not hold at time 20. This is because we cannot show, for example, that ¬ Happens(Unload,12). We want to assume by default that the description of a narrative of events includes all the events relevant to the conclusions we expect to draw. We want this to be a <u>default</u> assumption so that if we later learn that indeed the gun <u>was</u> unloaded at time 12, our representation is sufficiently elaboration tolerant to be able to absorb this new fact without leading to contradiction, and without needing to be reconstructed from scratch.

Our requirements can be met by incorporating the minimisation of Happens into our circumscription policy. Furthermore, as shown in [Miller & Shanahan, 1994], this will not interfere with the minimisation required to overcome the frame problem. However, there's a problem. If we minimise event occurrences indiscriminately then we don't allow for events we simply don't know about. To take the example from [Miller & Shanahan, 1994], suppose Mary is in London and Joe is in New York, and we want to formally express Mary's knowledge and formalise what conclusions she can reasonably draw from it.

Mary knows what's going on in London, but has at best a sketchy picture of what's happening in New York. Suppose, modifying an example due to John McCarthy, that Mary is stacking blocks in London and Joe is stacking blocks in New York. Mary knows all the events that take place that are relevant to her block stacking. Suppose we have the following.

Happens(Move(A,B),10)

Happens(Move(C,A),20)

And suppose want to conclude from this that Holds(On(A,B),State(30)), using a standard Blocks World theory about moving blocks, and some solution to the frame problem, such as Baker's, to ensure that moving C doesn't affect A. Also, we need to minimise Happens, representing the assumption that Mary knows all the events that are relevant to her, to eliminate the possibility that A gets moved between 20 and 30 by another event.

How do we ensure that the latter minimisation doesn't also allow Mary to conclude that nothing happens in New York? Suppose blocks A, B, and C are in London and that blocks D, E, and F are in New York. Mary might know that Holds(On(D,E),State(10)). But she doesn't know what stacking events are taking place in New York, so she doesn't want to conclude that ¬ Happens(Move(D,F),11). But minimising Happens will allow this conclusion to be drawn, forcing her to conclude that D stays on E forever, even though Joe is busy moving blocks around in New York.

2. More Sensitive Minimisation

The way to avoid the over-zealous minimisation of event occurrences is,

- To recognise that Mary's and Joe's activities constitute separate narratives.
- To minimise the occurrence of events only within those narratives for which we want to assume complete knowledge, permitting honest ignorance about the rest.

But how do we go about this? A first attempt at a solution assumes that geographical separation corresponds to narrative separation [Miller & Shanahan, 1994]. Let's look at this attempt briefly. In this solution, the term Location(e) denotes the geopgraphical location of an event of type e. A new predicate Happens* is introduced for events which are part of a narrative for which complete knowledge is assumed. In this particular example, knolwedge of the London narrative is assumed to be complete, whilst knowledge of the New York narrative is not. We have,

$$[Happens(e,t) \wedge Location(e) = London] \rightarrow Happens^*(e,t). \qquad (G1)$$

Now instead of minimising Happens, we will minimise Happens*. This has the effect of minimising the set of

[2] I prefer a Sneeze action to the usual Wait action, because in narrative terms waiting is best represented simply as a pause.

events that happen in London. In other words, it minimises the set of e's and t's such that Happens(e,t) ∧ Location(e) = London. But we still need another axiom which says that actions only affect fluents in the same geographical locations as themselves. Assuming that Location(f) denotes the location of fluent f, we have,

$$\text{Location}(e) = \text{Location}(f) \leftarrow \neg\,[\text{Holds}(f,\text{Result}(e,s)) \leftrightarrow \text{Holds}(f,s)]. \quad \text{(G2)}$$

To see how this works, let's consider the two block stacking narratives again. Further extending the application of the Location function, we have,

$$\text{Location}(\text{Move}(x,y)) = c \leftarrow [\text{Location}(x) = c \wedge \text{Location}(y) = c]$$

$$\text{Location}(\text{On}(x,y)) = c \leftarrow [\text{Location}(x) = c \wedge \text{Location}(y) = c]$$

$$\text{Location}(A) = \text{Location}(B) = \text{Location}(C) = \text{London}$$

$$\text{Location}(D) = \text{Location}(E) = \text{Location}(F) = \text{NewYork}.$$

We also need uniqueness-of-names axioms for locations, blocks, actions and fluents. Now from the minimisation of Happens* and Axiom (G1), we can show,

$$[\text{Happens}(e,t) \wedge \text{Location}(e) = \text{London}] \rightarrow [[e=\text{Move}(A,B) \wedge t=10] \vee [e=\text{Move}(C,A) \wedge t=20]].$$

From which we can conclude, using Axiom (G2), that whatever happens in New York does not affect what holds in London. In particular, given the frame axiom (F1) and the minimisation of Ab, we can show that,

$$\neg\,[\text{Happens}(e,t) \wedge 10 < t \leq 30 \wedge \text{Ab}(e,\text{On}(A,B),\text{State}(t))].$$

Which in turn gives, from Axioms (F1) and (N1),

$$\text{Holds}(\text{On}(A,B),\text{State}(30)).$$

We can draw this conclusion, although we don't know the sequence of events between time 10 and time 30, because we do know all the events between those times that can affect the relevant fluent. This is one conclusion we wanted for Mary. We also wanted her to be unable to conclude anything about what was happening in New York. Recall that Mary knew that a Move(D,E) action took place at time 10. But, as we require, we <u>cannot</u> show from the minimisation of Happens* that,

$$[\text{Happens}(e,t) \wedge \text{Location}(e) = \text{NewYork}] \rightarrow [e=\text{Move}(D,E) \wedge t=10].$$

Which means that, as required, there are there are models in which we have, for example,

$$\text{Happens}(e,t) \wedge \text{Location}(e) = \text{NewYork} \wedge e=\text{Move}(D,F) \wedge t=11.$$

3. Problems with the Naïve Approach

So far so good. The solution works for this example. But the solution is nowhere near as general as we would like. To begin with, the identification of geographical separation with narrative separation is obviously naïve. Suppose, for example, that the narratives we are interested in are telephone calls, rather than block stackings. Mary and Jane could be in adjacent rooms in London talking respectively to Joe and Fred in adjacent rooms in New York. But the events in Mary and Joe's conversation comprise one narrative and those in Jane and Fred's a different narrative.[3]

Furthermore, there are plenty of events going on at the atomic level which are in the same geographical location as Mary which she knows nothing about, and which she doesn't need to know anything about in order to reason about her block stacking. Similarly, there may be things happening just outside the window not significantly further away than the blocks she is stacking, which she doesn't need to know about to reason correctly about the blocks (leaves blowing, people walking along the street, and so on).

What we would like to capture somehow is the idea that a narrative has to encompass all the events which are relevant to a particular set of fluents, that is to say, all those events which affect those fluents. A natural starting point is to think of each narrative as a different context, in the sense of McCarthy [1993]. In the present example, the New York narrative is one context — certain things are true in that context, and certain actions and fluents are part of that context — and the London narrative is a separate context, with its own fluents, actions and true formulae. Then minimisation of event occurrences becomes context-sensitive. We will want to assume complete knowledge about certain contexts, but not about others.

The same event could be a part of many different narratives, and correspondingly same facts can be true in many different contexts. The 1989 earthquake in San Francisco, for example, certainly affected many people's otherwise independent lives. Furthermore, separate narratives can intersect briefly, as would happen if Joe telephoned Mary from New York and interrupted her block stacking. This event becomes part of both narratives.

The temporary intersection of otherwise separate narratives is potentially the source of a serious difficulty in correctly formalising the context-sensitive minimisation of event occurrences, because we cannot say in advance which events need to be included in any given context. This is reminiscent of the difficulty of finding a solution to the frame problem which can deal with ramifications. In effect, we would like an axiom that enshrines the following principle. Any event which, if it were included in a given context, would affect one or more fluents in that context, should be included in that context. And every other event should be excluded.

4. Context-Sensitive Minimisation

Here's an attempt to get what we want. The basic ideas about context are adopted from [McCarthy, 1993], although McCarthy's original intention was to apply them to a somewhat different problem, namely that of the lack of "generality" of formal theories in AI. Because nested contexts are not used, no special semantic machinery, such as that described in [Buvac & Mason, 1993], is required.

[3] Note that the narratives in question here are not the stories that Mary, Jane, Joe and Fred are telling each other, but are the sequences of events that comprise the speech acts in their conversations.

Standard first-order predicate calculus model theory is enough. The term Value(c,x) denotes the value of object x in context c. The formula In(c,x) represents that object c is part of context c, and the formula Ist(c,p) denotes that p is true in conext c. First, we have a contextualised version of Axiom (N1).

$$\begin{aligned}&\text{Value}(c,\text{State}(t1)) = \\ &\quad \text{Result}(e1,\text{Value}(c,\text{State}(t2))) \leftarrow \\ &\quad\quad [\text{Ist}(c,\text{Happens}(e1,t2)) \wedge t1 \geq t2 \wedge \\ &\quad\quad\quad \neg\exists e2,t3\ [\text{Ist}(c,\text{Happens}(e2,t3)) \wedge \\ &\quad\quad\quad\quad [e2 \neq e1 \vee t2 \neq t3] \wedge t2 < t3 < t1]]\end{aligned} \tag{Nc1}$$

To cope with the initial situation, we will have to bring in another axiom from [Miller & Shanahan, 1994], which is the companion axiom to (N1), but which accounts for times before any events have occurred. The formula Initially(f) is used to represent that fluent f holds in the initial situation. In its original form, the axiom is as follows.

$$\begin{aligned}&\exists s\ [[\text{Initially}(f) \leftrightarrow \text{Holds}(f,s)] \wedge \\ &\quad \forall t1[\neg\ \exists e,t2\ [\text{Happens}(e,t2) \wedge t2 < t1] \rightarrow \\ &\quad\quad \text{State}(t1) = s]]\end{aligned} \tag{N2}$$

In the contextualised version, the predicate Initially has been reified in order to make it context dependent. The contextualised version is,

$$\begin{aligned}&\exists s\ [[\text{Ist}(c,\text{Initially}(f)) \leftrightarrow \text{Holds}(f,s)] \wedge \\ &\quad \forall t1\ [\neg\ \exists e,t2\ [\text{Ist}(c,\text{Happens}(e,t2)) \wedge t2 < t1] \rightarrow \\ &\quad\quad \text{Value}(c,\text{State}(t1)) = s]]\end{aligned} \tag{Nc2}$$

Now we no longer expect conclusions about what holds at a given time, in other words conclusions of the form Holds(f,State(t)). Instead we expect contextualised conclusions of the form Holds(f,Value(c,State(t))). A fluent that holds in one context may be undefined in another.

Narratives will be described as in the decontextualised version, in terms of Initially and Happens. But in addition, the fluents in the domain can be assigned to any number of contexts using the predicate In. Let's consider Mary's and Joe's narratives again. As well as Happens and Initially facts, the narrative description will incorporate In facts, such as the following. The context of Mary's narrative is called "Mary".

$$\text{In}(c,\text{On}(x,y)) \leftarrow \text{In}(c,x) \wedge \text{In}(c,y)$$

$$\text{In}(\text{Mary},A)$$

$$\text{In}(\text{Mary},B)$$

$$\text{In}(\text{Mary},C)$$

In will be minimised at a lower priority than Ab.[4] Then, axioms are required which link decontextualised Initially and Happens facts to contexts, so that Axioms (Nc1) and (Nc2) will apply to them.

$$\begin{aligned}&\text{Ist}(c,\text{Happens}(e,t)) \leftarrow \\ &\quad \exists f,s\ [\text{Ab}(e,f,s) \wedge \text{Happens}(e,t) \wedge \text{In}(c,f)]\end{aligned} \tag{Nc3}$$

$$\text{Ist}(c,\text{Initially}(f)) \leftarrow [\text{Initially}(f) \wedge \text{In}(c,f)] \tag{Nc4}$$

[4] This minimisation is for convenience and is not intended to reflect any law of common sense.

Instead of minimising Happens, we want to minimise the events that occur in each context which can affect fluents in that context. This is done by minimising the predicate Happens*, which is constrained by the following axiom.

$$\begin{aligned}&\text{Happens}^*(c,e,t) \leftarrow \\ &\quad \exists f,s\ [\text{Ist}(c,\text{Happens}(e,t)) \wedge \text{Ab}(e,f,s) \wedge \text{In}(c,f)]\end{aligned} \tag{Nc5}$$

Happens* is minimised at a lower priority than both Ab and In, and Ist is allowed to vary. In effect, Axioms (Nc3) and (Nc5) serve the same role as Axioms (G1) and (G2) in the naïve attempt to get more selective minimisation presented in Section 2.

5. A Worked Example

To see how the above logical machinery works, I will now work through an example of two separate block stacking narratives. Let's assume that Joe and Mary are in adjacent rooms, rather than distant cities. In Room 1, Mary is stacking blocks A, B, and C. At the same time in Room 2, Joe is stacking blocks D, E, and F. Mary knows nothing of what Joe is doing.

The fluents known to be involved in the two contexts are as follows.

$$\text{In}(c,\text{On}(x,y)) \leftarrow \text{In}(c,x) \wedge \text{In}(c,y)$$

$$\text{In}(c,\text{Clear}(x)) \leftarrow \text{In}(c,x)$$

In(Mary,A)	In(Mary,B)
In(Mary,C)	In(Joe,D)
In(Joe,E)	In(Joe,F)

For the initial situation, we have,

Initially(On(A,C))	Initially(On(C,Table))
Initially(On(B,Table))	Initially(Clear(A))
Initially(Clear(B))	Initially(Clear(Table))
Initially(On(D,E)).	

Mary's actions are as follows.

$$\text{Happens}(\text{Move}(A,B),10)$$

$$\text{Happens}(\text{Move}(C,A),20)$$

The following is the Blocks World theory we will use.

$$\begin{aligned}&\text{Holds}(\text{On}(x,y),\text{Result}(\text{Move}(x,y),s)) \leftarrow \\ &\quad [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s)]\end{aligned} \tag{B1}$$

$$\begin{aligned}&\neg\ \text{Holds}(\text{On}(x,z),\text{Result}(\text{Move}(x,y),s)) \leftarrow \\ &\quad [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s) \wedge y \neq z]\end{aligned} \tag{B2}$$

$$\begin{aligned}&\text{Holds}(\text{Clear}(z),\text{Result}(\text{Move}(x,y),s)) \leftarrow \\ &\quad [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s) \wedge y \neq z]\end{aligned} \tag{B3}$$

$$\begin{aligned}&\neg\ \text{Holds}(\text{Clear}(y),\text{Result}(\text{Move}(x,y),s)) \leftarrow \\ &\quad [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s) \wedge y \neq \text{Table}]\end{aligned} \tag{B4}$$

Now we can show from the Blocks World theory and (F1) that,

$$\begin{aligned}&\text{Ab}(\text{Move}(x,y),\text{On}(x,z),s)) \leftarrow \\ &\quad [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s) \wedge \\ &\quad \text{Holds}(\text{On}(x,z),s) \wedge z \neq y]\end{aligned}$$

Which gives, from (Nc5) and what we know about the contexts in this domain,

$$\text{Ist}(c,\text{Happens}(\text{Move}(x,y),t)) \leftarrow \exists s\ [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s) \wedge \text{Holds}(\text{On}(x,z),s) \wedge z \neq y \wedge \text{Happens}(\text{Move}(x,y),t) \wedge \text{In}(x,c) \wedge \text{In}(y,c)].$$

Baker's existence of situations axiom ensures that for any x, y, and z,

$$\exists s\ [\text{Holds}(\text{Clear}(x),s) \wedge \text{Holds}(\text{Clear}(y),s) \wedge \text{Holds}(\text{On}(x,z),s) \wedge z \neq y].$$

And therefore,

$$\text{Ist}(c,\text{Happens}(\text{Move}(x,y),t)) \leftarrow [\text{Happens}(\text{Move}(x,y),t) \wedge \text{In}(x,c) \wedge \text{In}(y,c)].$$

This gives, as we might expect,

$$\text{Ist}(\text{Mary},\text{Happens}(\text{Move}(A,B),10))$$

$$\text{Ist}(\text{Mary},\text{Happens}(\text{Move}(C,A),20)).$$

And consequently,

$$\text{Happens*}(\text{Mary},\text{Move}(A,B),10)$$

$$\text{Happens*}(\text{Mary},\text{Move}(C,A),20).$$

Then, through the minimisation of Ab and of Happens*, we have,

$$[\text{Ist}(\text{Mary},\text{Happens}(e,t)) \wedge \exists f,s\ [\text{Ab}(e,f,s) \wedge \text{In}(\text{Mary},f)]] \leftrightarrow [[e=\text{Move}(A,B) \wedge t=10] \vee [e=\text{Move}(C,A) \wedge t=20]].$$

Now we can show, using (Nc1) and (Nc2), the Blocks World axioms, and (F1), that Holds(On(A,B),Value(Mary,State(30))), which is the desired result. However, although we have Initially(On(D,E)), and whatever actions of Joe's on blocks D, E, and F we described, we would not be able to conclude anything of the form Holds(On(x,y),Value(Mary,State(t))) where x and y were D, E, or F. Similarly, we cannot conclude Holds(On(A,B),Value(Joe,State(30))).

If we want to include information about one narrative explicitly in another, we can using the Ist predicate. For example, if Mary learns that D was on E in the initial situations, we can add,

$$\text{Ist}(\text{Mary}, \text{Initially}(\text{On}(D,E))).$$

However, this will not allow any further conclusions to be drawn from within the context of Mary's narrative about D, E or F. We can't show, for example, that Holds(On(D,E),Value(Mary,State(5))), although we can show that Holds(On(D,E),Value(Joe,State(5))). This is because the minimisation of Happens* only rules out models with extra events in the Mary's narrative which affect fluents included in the context of her narrative. It cannot rule out the possibility of extra events which affect fluents outside the context of her narrative. If we wanted to, of course, we could always include D, E, and F in Mary's context.

6. Causal Relevance

The solution to the over-zealous minimisation of event occurrences presented in Sections 4 and 5 above presupposes that the set of fluents included in a context is closed with respect to causal relevance. For example, suppose the context of Joe's narrative includes the fluent LightOn, denoting that the light in his office is on. Joe can switch the light on and off via a Toggle action. Imagine a situation in which the light is broken, but Joe doesn't know that it's broken. The Toggle action has no effect under these circumstances. So there is a fluent, which we can call Broken, which affects the outcome of Joe's Toggle actions, but whose value is unknown to Joe.

The formal apparatus of Sections 4 and 5 conflates two uses of the idea of context in a way which prevents us from dealing correctly with examples such as this. We need to separate an agent's epistemic context — what he or she knows — from a fluent's causal context — the set of fluents that are causally relevant to it. The fluents that are causally relevant in a context are those on which the outcome of events in that context depend. To effect this separation, let's reserve the predicate In for epistemic context, and try to formalise the conditions under which the epistemic context is closed with respect to causal relevance.

Informally, we are interested in cases where a fluent f1 is abnormal with respect to e in one situation s1 but not in another s2. Then there must be some fluent whose value is different in s1 and s2, which is causally relevant to f1. To pick out such fluents, we will want to focus on pairs of situations s1 and s2 which differ as little as possible, and yet which yield different results for e. Fluents on which s1 and s2 disagree must be included in the same context as f1 if that context is to be complete.

First, we have a definition which will help us pick out such pairs of situations. The formula DisagreeOn(s1,s2,f) represents that situations s1 and s2 have different values for fluent f.

$$\text{DisagreeOn}(s1,s2,f) \equiv_{def} \neg\ [\text{Holds}(f,s1) \leftrightarrow \text{Holds}(f,s2)]$$

The formula AsCloseTo(s1,s2,s3) represents that the set of fluents on which s1 and s2 disagree is at least as small as the set of fluents on which s1 and s3 disagree.

$$\text{AsCloseTo}(s1,s2,s3) \equiv_{def} \forall f\ [\text{DisagreeOn}(s1,s3,f) \rightarrow \text{DisagreeOn}(s1,s2,f)]$$

Now we can pick out the fluents which can influence the outcome of an event of type e which in turn affects a fluent in a given context c. They are the members of the smallest possible set of fluents that will distinguish two situations in which e has a different effect on the same fluent. Note that the effect of e on f in situation s1 is different from its effect on f in situation s2 if we can show $\text{Ab}(e,f,s1) \wedge \neg\ \text{Ab}(e,f,s2)$. We have,

$$\text{Relevant}(c,f1) \equiv_{def} \exists s1,s2,f2,e,t\ [\text{In}(c,f2) \wedge \text{Ist}(c,\text{Happens}(e,t)) \wedge \text{Ab}(e,f2,s1) \wedge \neg\ \text{Ab}(e,f2,s2) \wedge \text{DisagreeOn}(s1,s2,f1) \wedge \forall s3\ [[\text{Ab}(e,f2,s1) \wedge \neg\ \text{Ab}(e,f2,s3)] \rightarrow \text{AsCloseTo}(s1,s2,s3)]]$$

We are now in a position to confine the minimisation of event occurrences to contexts which are closed with respect to causal relevance, in other words to contexts which

include every fluent which is causally relevant to any other fluent in the context. First, we define a closure in this sense.

$$\text{Closed}(c) \equiv_{def} \neg \exists f\, [\text{Relevant}(c,f) \wedge \neg \text{In}(c,f)]$$

Now, we modify Axiom (Nc5) to take account of complete contexts only.

$$\text{Happens}^*(c,e,t) \leftarrow \quad \text{(Nc6)}$$
$$\exists f,s\, [\text{Closed}(c) \wedge \text{Ist}(c,\text{Happens}(e,t)) \wedge \text{Ab}(e,f,s) \wedge \text{In}(c,f)]$$

As before, Happens* is minimised with a lower priority than Ab or In, and Ist is allowed to vary. Effectively, there is no event occurrence minimisation in non-complete contexts.

Working through examples of the use of these axioms reveals that they yield the results we expect for some cases, such as Joe's office light, but not for others. Using a natural representation of the office light domain, we can show Closed(Joe) when the light cannot be broken, but ¬ Closed(Joe) when the light can be broken but Broken is not in Joe's epistemic context.

However, examples like the parallel block stacking narratives don't work out so well. The trouble is that we can't prove ¬ Relevant(Joe,Clear(A)), since the axioms don't rule out the possibility that Happens(Move(A,D)). Then, since we don't have In(Joe,A), we cannot prove Closed(Joe). A similar argument can be applied for Closed(Mary). The upshot of all this is that minimising Happens* is ineffectual for Mary's and Joe's narratives.

This difficulty is reminiscent of the difficulty Hayes had in an early attempt to solve the frame problem using a notion related to causal relevance [Hayes, 1971]. The trouble is that the set of fluents that can potentially affect any given fluent is enormous. We want to keep this set as small as possible. In the case of the block stacking narratives, we come unstuck because of the potential effect of an action which is in fact impossible. It may be necessary to include explicit facts about causal separation to overcome this.

8. Concluding Remarks

The techniques presented in this paper go some way towards solving the problem of over-zealous event occurrence minimisation. The use of context is a promising direction for research in this area. However, a completely general solution to the problem is still pending.

The issue of over-zealous minimisation is not confined to event occurrences. It is a general problem in default reasoning. For example, the minimisation of the Ab predicate to overcome the frame problem assumes that the known effects of an action are all the relevant effects of an action. Often this is the right default assumption to make. But sometimes all the relevant effects are not known. For example, if a bomb explodes in the street below my office, all sorts of fluents relevant to me will be affected, and it is very hard for me to predict exactly how. A general framework is required for performing default reasoning at a level of granularity which enables an agent to be honest about his or her ignorance. Context may be the key to achieving this within a standard circumscriptive framework.

Acknowledgements

Thanks to Sasa Buvac, Tom Costello and John McCarthy. The work reported here was mostly carried out during a visit to Stanford University in April 1994. The author is supported by an EPSRC fellowship.

References

[Baker, 1991] A.B.Baker, Nonmonotonic Reasoning in the Framework of the Situation Calculus, *Artificial Intelligence*, vol 49 (1991), pages 5-23.

[Buvac & Mason, 1993] S.Buvac and I.A.Mason, Propositional Logic of Context, *Proceedings AAAI 93*, pages 412-419.

[Hanks & McDermott, 1987] S.Hanks and D.McDermott, Nonmonotonic Logic and Temporal Projection, *Artificial Intelligence*, vol 33 (1987), pages 379-412.

[Hayes, 1971] P.Hayes, A Logic of Actions, in *Machine Intelligence 6*, ed. D.Michie, Edinburgh University Press (1971), pages 495-520.

[Kowalski & Sergot, 1986] R.A.Kowalski and M.J.Sergot, A Logic-Based Calculus of Events, *New Generation Computing*, vol 4 (1986), pages 67-95.

[McCarthy, 1993] J.McCarthy, Notes on Formalizing Context, *Proceedings IJCAI 93*, pages 555-560.

[McCarthy & Hayes, 1969] J.McCarthy and P.J.Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in *Machine Intelligence 4*, ed. D.Michie and B.Meltzer, Edinburgh University Press (1969), pages 463-502.

[Miller & Shanahan, 1994] R.S.Miller and M.P.Shanahan, Narratives in the Situation Calculus, *The Journal of Logic and Computation*, vol 4, no 5 (1994), pages 513-530.

[Pinto & Reiter, 1993] J.Pinto and R.Reiter, Temporal Reasoning in Logic Programming: A Case for the Situation Calculus, *Proceedings of the Tenth International Conference on Logic Programming* (1993), pages 203-221.

[Shanahan, 1995] M.P.Shanahan, A Circumscriptive Calculus of Events, to appear in *Artificial Intelligence*, vol 75, no 2 (1995).