On the Existence and Formalization of Natural Events

Javier Pinto
Depto. de Ciencia de la Computación,
Pontificia Universidad Católica de Chile,
Casilla 306, Santiago 22, Chile.
email: jpinto@ing.puc.cl

Abstract

In logical theories of change it is common to consider two types of objects that provoke change: 1) *actions* that are carried out by agents, and 2) *events* which occur externally, without agent intervention. In this paper we argue for the existence of a special type of event, which we call *natural events*. An event is called natural if its occurrence is the result of a natural process. For example, if a ball is dropped it is natural that it hits the ground. Thus, the event of the ball hitting the ground is considered natural. A natural event can be prevented by actions or events that intervene and change the nature of the process leading to the event. If no such intervening events occur, then the natural event *has to* occur. Hence, this third type of "change provoking" entity, the *natural event*, may be the result of some agent's prior action(s). Thus, in general, natural events need not be external.

In this paper, we discuss the need to introduce explicitly the notion of *natural event* in a logic of change and present a formalization based on the situation calculus.

1 Introduction and Motivation.

Consider a cook preparing Spaghetti al Pesto. Berta, the cook, has the raw spaghetti and all the ingredients to prepare the sauce. She needs a set of utensils and boiling water. In order to get the boiling water, she fills a pot with water and puts it on a hot plate. The hot plate is heating at its max setting and Berta starts preparing and mixing the ingredients for the sauce. After the water boils, Berta puts the raw spaghetti into the pot and continues with the sauce. She manages to have the sauce prepared just as the spaghetti is cooked. She drains the spaghetti and serves the meal.

In order to describe the complex events involved in examples of this sort, we need a language that allows for the representation of properties of the world, actions and events that may occur and action laws describing the effects of actions on the world's properties. Traditionally, theories of action and change have focused on discrete change. An important characteristic of these theories is that change is considered to be an immediate effect of action or event occurrences. An obvious limitation of this approach is that these theories cannot deal with worlds in which actions may have *delayed effects*, which occur when "[c]hanges due to an action may occur at a later time" [Sandewall, 1994a].

Consider the example of Berta, almost all actions/events that occur when she prepares a meal have a direct and immediate effect on the world. For example, cutting a mushroom in half or pouring sauce over the spaghetti produce definite effects on the world. A different kind of action is one where Berta puts the pot with water on a hot plate. This action, along with the fact that the heater is on at its max setting, ensures Berta that the water will eventually boil. A similar example is Sandewall's "European bathtub example" where the action of turning on the bathtub tap will cause the fluent "floor wet" to become true a while later.

The focus of this paper is on the representational requirements for a logical language to represent actions with delayed effects. We argue for the need to incorporate the notion of *natural events* in theories of action and change.

In this paper, we use the extended situation calculus presented in [Pinto and Reiter, 1995]. This situation calculus is a first order language with equality and sorts $\mathcal{A}, \mathcal{S}, \mathcal{F}, \mathcal{T}, \mathcal{D}$ for actions (including events and natural events), situations, fluents (world's dynamic properties), time (the non-negative reals) and other domain objects respectively. Distinguished symbols: S_0, a situation constant denoting the starting situation. *do*, a function symbol that takes an action and a situation and yields a situation (also called *result* in other formalizations of the situation calculus). *holds*, a predicate symbol that takes a fluent and a situation and is true if the property denoted by the fluent is true in the situation. Operator $<$ for situations such that $s_1 < s_2$ is true if there is a path in the situation tree starting in s_1 and leading to s_2.

In the situation calculus, the initial situation S_0 and the function symbol *do* generate a tree of situations rooted at S_0. The branching factor of the tree equals the cardinality of $\mathcal{A}$. This tree has the property that there is a unique path that leads from S_0 to any situation, and no path leads to S_0. Also, we make use of Pinto and Reiter's notion of *actual line* [Pinto and Reiter, 1995], by which a time line is added to the situation calculus.

The idea is to select a single rooted branch in the situation calculus tree as the set of *real* situations, such that all actions mentioned in it are regarded as having *occurred.* (For further details see [Pinto and Reiter, 1995; Pinto, 1994a]).

Also, we assume that the frame problem is solved using Reiter's approach. Thus, for each fluent symbol in the domain we derive a successor state axiom of the form:

$$Poss(a, s) \supset [holds(R, do(a, s)) \equiv \gamma_R^+(a, s) \vee holds(R, s) \wedge \neg\gamma_R^-(a, s)].$$

This derivation is possible only after making the so called *explanation closure* assumptions. See [Reiter, 1991] for details.

The rest of the paper is organized as follows: In section 2 we describe and informally justify the ontological assumptions on which we base our proposal. In section 3 we introduce and formalize the notion of *natural event.* Also, we discuss how these natural events interact with Reiter's solution to the frame problem. Finally, in section 4 we discuss our future work and conclude.

2 Ontological Assumptions.

As basic framework, we use the *discrete situation calculus* [Lin and Reiter, 1994] with reified fluents and extended with the notion of an *actual time line* [Pinto and Reiter, 1993; Pinto, 1994a]. It is essential to realize that our discussion below is quite independent of the use of the situation calculus. In the same manner as the *frame problem* is a problem which is independent of the formalism used to model actions and change, the notion of *natural event* is not an accessory of the situation calculus.

2.1 Instantaneous Actions vs. Actions with Durations.

Recently, there has been a growing interest in the development of formal frameworks to assess the degree of applicability of logics of action and change. The two most influential are Sandewall's [Sandewall, 1994b; Sandewall, 1994a] and Gelfond and Lifschitz's [Gelfond and Lifschitz, 1993; Lifschitz, 1993] proposals. Both approaches differ in many respects. In Sandewall's approach, one is given a *base logic* in which one is able to specify the properties of the world, the values of these properties at different points in time and a set of action laws. Sandewall provides a taxonomy of worlds based on ontological and epistemological assumptions which are made explicit at a meta-level. The semantics of the logic depends upon the set of assumptions made.

Gelfond and Lifschitz's approach is based on a taxonomy of *action description languages* whose basic element is the language **A** [Gelfond and Lifschitz, 1993]. Each language in the taxonomy corresponds to one class of worlds in Sandewall's taxonomy, and each language has its own semantics, also based on some ontological and epistemological assumptions. An advantage of Sandewall's approach is that the ontological and epistemological assumptions are left explicit in his framework. In contrast, in Gelfond and Lifschitz's approach these assumptions are left implicit. In [Thielscher, 1994], Thielscher establishes the relationship of the base language **A** with Sandewall's taxonomy. A more general study of the relationship between both assessment methodologies needs to be developed.

Interestingly, both approaches model actions as entities that have an inherent duration. In fact, Sandewall specifies action occurrences as statement of the form $[\mathtt{s},\mathtt{t}]\mathtt{A}(\ldots)$, where $\mathtt{s}$ and $\mathtt{t}$ are variable terms for timepoints. All effects of actions arise immediately after the end of the action occurrence. Thus, the effects of $\mathtt{A}(\ldots)$'s occurrence will hold immediately after t. On the other hand, Lifschitz introduced the **AI** language to the taxonomy of action description languages started with Gelfond and Lifschitz's **A** language. In **AI**, as in Sandewall's base logic, actions occur in intervals and their effects show after the action is finished.

Thus, it appears that there is a strong trend towards axiomatizations that force actions to have durations. However, we argue that this is the wrong choice. From an strictly formal point of view, which choice is made seems to be irrelevant. In fact, most logical approaches to model time take it to be a linear structure isomorphic to the natural numbers or to the non-negative real numbers. Thus, any interval may be modeled by its endpoints. Similarly, any point can be modeled as a duration-less interval[1].

At a pragmatic level, however, there are several advantages to model actions as instantaneous. The most important is that it leads to succinct representations. In particular, if one considers a pair of actions a_1 and a_2, then these actions can either be concurrent or one occurs after the other. Thus, there are only three possible combinations. On the other hand, if a_1 and a_2 have durations, then they may occur concurrently, they may overlap, they may occur in sequence with time in between, etc. (see [Allen, 1983]). Each of these situations has to be considered when analyzing the effects of actions in the world. Another advantage of duration-less actions is that modeling concurrency is relatively straightforward, since there is only one way in which concurrency may arise.

One way of introducing actions with durations in our framework is to define them in terms of instantaneous actions as follows. We consider two separate sorts: $\mathcal{A}$ of instantaneous actions, and $\mathcal{A}_d$ of actions with durations. Furthermore, we introduce three function symbols: $start_d, end_d : \mathcal{A}_d \rightarrow \mathcal{A}$ and $ing : \mathcal{A}_d \rightarrow \mathcal{F}$. For example, given the action with duration *paint*, the instantaneous actions $start_d(paint)$ and $end_d(paint)$ denote the actions associated with the start and the end of the action with duration. In addition, $ing(paint)$ is a fluent that is true while *painting* holds. We also require:

$$holds(ing(a_d), do(start_d(a_d), s)),$$
$$\neg holds(ing(a_d), do(end_d(a_d), s)),$$
$$Poss(end_d(a_d), s) \supset holds(ing(a_d), s),$$

[1] The interval temporal logic of Allen's [Allen, 1984] considers a structure of time in which there are indivisible time intervals which cannot be defined as pairs of points, so called *moments*.

$$Poss(start_d(a_d), s) \supset \neg holds(ing(a_d), s).$$

Thus, actions with durations can be represented with two instantaneous actions (that start and end the action with duration) along with a fluent that holds in the situations in which the action is in progress. In this approach, the function symbol *do* only applies to instantaneous actions. Thus, we consider the sort $\mathcal{A}_d$ to be independent of sort $\mathcal{A}$. If one were to modify the language to include actions with duration as valid arguments for *do* terms, one would have to account for the possibility of these actions no to finish. A proposal along this lines was advanced by Gelfond, Lifschitz and Rabinov in a non-discrete situation calculus framework [Gelfond *et al.*, 1991].

2.2 Discrete Events Assumption.

Another important assumption that we make is to consider the set of actions to be discrete. Thus, in any finite interval of time the number of events that may occur is also finite. In the discrete situation calculus, this assumption is formalized with the second order induction axiom. Thus, if two situations are in the same branch of a tree structure, then one leads to the other through a finite set of actions/events.

This assumption may seem too strong. In fact, contradictions arise if one tries to encode Zeno's paradox in the extended situation calculus. However, the paradox seems to arise due to the mathematical model being used as opposed to any realization of the phenomena in nature. For example, if a ball bounces at time 1, then 1.5, then 1.75, etc., we we'll conclude that at time 2 it will have made infinitely many bounces. This is what a simple minded model of the bouncing ball would conclude. In reality such a thing will never occur.

The most important advantage of the discrete-ness of events assumption is that it allows us to use induction to prove properties of states (see [Reiter, 1993]).

3 Processes and Natural Events.

In this paper we do not formally define the notion of a process. Informally, we consider a process to be a *development*. That is, an ordered set of actions or events, ruled by a set of behavioral constraints.

For instance, the process of preparing *Spaghetti al Pesto* consists of a sequence of actions and events. Each event or action is constrained by specific rules that establish the appropriateness of each step. A process of this nature may be regarded as a *complex action* [Lespérance *et al.*, 1994]. Complex actions are built from primitive actions by combining them using a set of action operators that include *sequencing*, *non-deterministic choice*, etc. Thus, it would seem that *processes* are not primitive objects, but an artifact that allows a higher level of abstraction. However, there is an aspect that is not accounted for. Assume that I throw a ball to my friend on the other side of the yard. The throwing action initiates a process in which the ball describes a motion determined by the physical characteristics of the environment and the specifics of the throwing action. The process that is initiated has to be described in such a way that the process is assumed to continue until it is stopped by some agent's action, some external event or by the process itself. But, we must ensure that the process goes on if no perturbing action occurs. How do we model this process-initiating actions?

Consider our cook, Berta, trying to synthesize a plan to prepare Spaghetti al Pesto. Berta uses an electric kettle to boil the water required for the cooking. The kettle is already full of water. Now, assume that Berta can perform the action *PlugKettle*. What are the effects of plugging the kettle? A naive answer would be that after plugging the kettle, we have boiling water; that is:

$$holds(waterInKettle, s) \supset \\ holds(boilingWater, do(PlugKettle, s)).$$

Unfortunately, this is inadequate. This sentence can be interpreted as saying that the water boils immediately after the kettle is plugged. However, this is not correct. The water will boil a few minutes after the kettle is plugged (assuming the kettle remains plugged and power is always available). Since the sentence above does not mention explicit time, it cannot possibly model the phenomenon at hand. Alternatively, we could consider that $PlugKettle$ is an action that lasts for a period of time. This possibility has several drawbacks. For instance, the action would have to occur for a period of time equal to the time taken for the water to boil. If the kettle were unplugged before the boiling occurs, then the action cannot be said to have occurred. Another alternative is to write an axiom like:

$$occurs(PlugKettle, s) \supset \\ occurs(waterBoils, do(PlugKettle, s)),$$

where occurs is defined as in [Pinto and Reiter, 1995]. Here, we allow for a lapse of time between the $PlugKettle$ action and the $waterBoils$ event. However, this alternative is also problematic, since it forces a sequence of events where the water boils after the plugging event without allowing for intervening events or actions. Thus, Berta would be paralyzed until the water boils. A more appropriate possibility would be:

$$occurs(PlugKettle, s) \supset \\ (\exists\, s').s' > s \wedge occurs(waterBoils, s')),$$

which states that in some situation after the occurrence of $PlugKettle$ the water boils. Unfortunately, this too is inappropriate. The reason is that it is not necessarily the case that the water will boil after the kettle is plugged in. In fact, some agent may unplug the kettle before the water boils. Then, what are the effects of plugging the kettle? The answer is determined by the level of abstraction that one is prepared to accept for the axiomatization of a particular domain. For instance, if we are thinking about a household robot, we may simply state that the plugging event sets up the conditions necessary to start a process. In general, we may say that:

The direct effects of actions/events may be of two types: 1) *fluent changing* effects, which correspond to changes on the fluent values provoked by the action. For example, after releasing the ball I had in my hand, my hand becomes empty. 2) *process triggering* effects, which correspond to conditions that are set for a process to get started. For example, after releasing the ball I had in my hand, the ball starts moving.

In state based representations the world is considered to be *static*. That is, the world does not change while it is in a given state. Furthermore, state changes arise due to actions or external events. It is not possible, then, to represent change due to ongoing natural processes. For example, if a ball is falling, its position changes; but no actions or events are taking place. Our solution to this representational problem is to use a hybrid representation. We divide properties of the world according to whether or not they can be considered to be discrete. For instance, in the water boiling example, the property *there is water in the kettle* may be regarded as discrete, whereas *the temperature of the water is 100 degrees Celsius* may be regarded as non-discrete. For example, the fluent $Heating$ may be used to indicate whether the water in the kettle is heating. Thus, $holds(Heating, s)$ would be true of those situations in which the kettle is plugged. Furthermore, we could write successor state axioms like:

$$\begin{aligned}Poss(a,s) \supset [&holds(Heating, do(a,s)) \equiv\\ &holds(Heating, s) \wedge a \neq Unplug \vee\\ &a = PlugKettle].\end{aligned}$$

Now, assume that S denotes some state such that $holds(Heating, S)$ is true. Furthermore, assume that no agents perform actions. Then, what happens? The answer should be that the water boils. Thus, the event $WaterBoils$ should *occur* when no agents prevent such an occurrence. In order to infer that $WaterBoils$ occurs we would like to write a formula of the following form:

$$\varphi(s) \supset occurs(WaterBoils, s).$$

Where φ is a formula in which there is no situation term other than s. Thus, we would like to specify conditions that a situation must meet in order for a natural event (such as $WaterBoils$) to arise.

In this axiomatization we distinguish natural events from the other events (or actions) introducing the predicate $natural \subseteq \mathcal{A}$. The most important characteristic of natural events is that when the conditions for their occurrence arise, they must occur. In this framework, a complete specification of the predicate *natural* has to be provided by the axiom writer. For instance, $WaterBoils$ should be *natural*. In order to deal with natural events that arise as the result of natural processes, we need a notion of metric time. As in [Pinto and Reiter, 1995] we incorporate metric time into the situation calculus. We consider that situations have temporal extent. Thus, a situation may start at a certain time and end at another. Thus, rather than having actions with durations, we have situations with durations. We consider time to be like the positive reals. A situation starts at a time point and its ending depends on the action being performed in that situation. We introduce the function symbol *end*, which takes a situation and an action and denotes a time point. Thus, $end(do(PlugKettle, s), WaterBoils)$ would denote the time at which the situation resulting from $PlugKettle$ is finished by the occurrence of the action $WaterBoils$. As a general *occurrence axiom* for natural events, we will write:

$$\begin{aligned}&[natural(a) \wedge Poss(a,s) \wedge\\ &\quad \neg(\exists\, a')\,[occurs(a', s) \wedge end(s,a') < end(s,a)]] \supset\\ &\qquad occurs(a,s).\end{aligned}$$

Here, a refers to a natural event and $end(s,a)$ denotes the time at which a would occur given the conditions at the start of s. Hence, the axiom states that if a natural event is possible in a given situation and no other action or event occurs before $end(s,a)$, then the natural event will occur. This axiom ensures us that when the proper conditions for a natural event to occur are present, then the event actually occurs.

We make the assumption that no natural events occur concurrently. In general, this assumption is too strong. However, it is not difficult to incorporate concurrency into the language, this is matter for future work (see [Pinto, 1994b]).

To describe some important properties of natural actions we introduce a new predicate $\Pi_t \subseteq \mathcal{A} \times \mathcal{S} \times \mathcal{T}$ which is domain dependent. $\Pi_t(a,s,t)$ is true whenever the conditions at the time $start(s)$ are such that at time t action a would occur.

For example, assume that $holds(Heating, S_0)$ is true. Assuming that no other natural processes are active, our axiomatization should yield:

$$(\exists\, t) t > start(S_0) \wedge \Pi_t(WaterBoils, S_0, t)$$

as true. Thus, some time after the start of S_0 it is possible that the water will boil. The exact value of t should be obtained from an axiomatization of Π_t. The axiomatization need not be complete and several models of the axiomatization may arise as a consequence. The axiomatization could be based on qualitative information or on a quantitative physical model of the heat transfer process. For example, the value of t for which Π_t is made true may correspond to a solution to a system of equations describing a given phenomenon.

In the same manner as with *natural*, we assume that we have a complete axiomatization for Π_t. Based on the predicates just introduced, we now describe some essential properties of natural events.

First, if a situation s starts at time t_s, then the situation s would be ended by some natural action a at the first time greater than t_s for which $\Pi_t(a,s,t)$ was true. This is written as follows:

$$\begin{aligned}&natural(a) \wedge \Pi_t(a,s,t) \wedge\\ &\quad \neg(\exists\, t')[start(s) < t' < t \wedge \Pi_t(a,s,t')] \supset\\ &\qquad end(s,a) = t.\end{aligned}$$

Also, we use the predicate Π_t to completely characterize the predicate *Poss*. Thus, unlike the case with discrete non-*natural* actions, we may write a general form for a definition of *Poss* for natural events:

$$natural(a) \supset Poss(a, s) \equiv [\Pi_t(a, s, end(s, a)) \wedge \neg(\exists\, a', t').natural(a') \wedge \Pi_t(a', s, t') \wedge start(s) < t' \leq end(s, a)].$$

That is, natural action a is possible in state s if, given the conditions at the start of s, the action is possible (as described by Π_t) at the time $end(s, a)$, and no other natural action could have occurred between the start of s and $end(s, a)$.

Another essential aspect of any *natural* action a is described by the following axiom:

$$[natural(a) \wedge Poss(a, s) \wedge Poss(a', s) \wedge \neg natural(a')] \supset end(s, a') < end(s, a).$$

Thus, if a is a *natural* action that is possible in situation s, then no other action is possible after the time at which a would occur (i.e., $end(s, a)$). An immediate consequence of the above axioms is:

Observation 3.1

$$[natural(a) \wedge natural(a') \wedge Poss(a, s) \wedge Poss(a', s)] \supset a = a'.$$

Thus, at most one natural action is possible in a given situation.

Clearly, if we extend the action language with a *concurrency operator* the axiomatization will have to be revised and this observation withdrawn.

In a given situation several processes may be active and many *natural events* may be scheduled to occur in the future due to them. However, in our formalization, the natural event that is scheduled to occur first is possible in the current situation and all other are not possible. All other natural events would be possible in future situations as long as their processes remain active. Whether or not a process is active in a given situation is encoded by a set of fluents (e.g. *Heating*). Thus, to determine if a process remains active after an action or event takes place we make use of Reiter's successor state axioms.

4 Summary, Final Remarks and Conclusions

In this paper we argue for the need to incorporate the notion of *natural events* or *natural actions* into the situation calculus. This need arises because many events are the natural result of ongoing processes. We present a formalization of natural actions based on a situation calculus extended with occurrences. A complete axiomatization along with some examples can be found in [Pinto, 1995].

In this paper we did not address the problem of modeling continuous change. Instead, we were able to isolate our axiomatization from the continuous change using the predicate $\Pi_t(a, s, t)$, which is true of a natural event a if it would occur at time t in situation s (provided that no intervening events occur). The encoding of the information to determine the times for which $\Pi_t(a, s, t)$ is true may be done in different ways. For instance, it may be specified as the solution to differential equations characterizing the process. In fact, this is the way in which we have chosen to encode such information in [Pinto, 1995]. As alternative, t can also come from a solution to a qualitative process description.

Currently, we are working on the development of algorithms to reason with natural events. In particular, we are interested in *temporal projection* and *explanation* tasks.

For our future work, we need to provide a means to assess the applicability of our proposal. Before we can attempt to perform this assessment, we need to extend existing formal framework for this task. Thus, we should extend Sandewall's [Sandewall, 1994b; Sandewall, 1994a] or Gelfond and Lifschitz's [Gelfond and Lifschitz, 1993; Lifschitz, 1993] proposals to include *natural events.*

References

[Allen, 1983] James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

[Allen, 1984] James F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154, 1984.

[Gelfond and Lifschitz, 1993] Michael Gelfond and Vladimir Lifschitz. Representing action and change by logic programs. *The Journal of Logic Programming*, 17:301–322, 1993.

[Gelfond *et al.*, 1991] M. Gelfond, V. Lifschitz, and A. Rabinov. What are the Limitations of the Situation Calculus? In *Working Notes, AAAI Spring Symposium Series. Symposium:Logical Formalization of Commonsense Reasoning.*, pages 59–69, March 1991.

[Lespérance *et al.*, 1994] Yves Lespérance, Hector Levesque, Fangzhen Lin, Daniel Marcu, Raymond Reiter, and Richard Scherl. A Logical Approach to High-Level Robot Programming – A Progress Report. In Benjamin Kuipers, editor, *Control of the Physical World by Intelligent Systems, Papers from the 1994 AAAI Fall Symposium*, pages 79–85. AAAI, New Orleans, November 1994.

[Lifschitz, 1993] Vladimir Lifschitz. A language for describing actions. In *The Second Symposium on Logical Formalizations of Commonsense Reasoning*, pages 103–113, January 1993.

[Lin and Reiter, 1994] Fangzhen Lin and Raymond Reiter. State constraints revisited. *Journal of Logic and Computation*, 4(5), 1994. Special Issue on Action and Processes.

[Pinto and Reiter, 1993] Javier Pinto and Raymond Reiter. Adding a time line to the situation calculus. In *The Second Symposium on Logical Formalizations of Commonsense Reasoning*, pages 172–177, January 1993.

[Pinto and Reiter, 1995] Javier Pinto and Raymond Reiter. Reasoning about time in the situation calculus. To appear in Annals of Mathematics and Artificial Intelligence, special festschrift issue in honour of Jack Minker., 1995.

[Pinto, 1994a] Javier Pinto. *Temporal Reasoning in the Situation Calculus.* PhD thesis, Department of Computer Science, University of Toronto, Toronto, Ontario, Canada, February 1994. Also published as Technical Report, Dept. of Computer Science, University of Toronto (KRR-TR-94-1), Feb. 1994.

[Pinto, 1994b] Javier A. Pinto. Concurrent events: Sinergy and cancellation of effects. In *Proceeding of the 1994 ECAI-Workshop on Logic and Change.*, pages 105–110, 1994.

[Pinto, 1995] Javier Pinto. Integrating discrete and continuous change in a logical framework. Submitted for publication., 1995.

[Reiter, 1991] Raymond Reiter. *The Frame Problem in the Situation Calculus: A Simple Solution (sometimes) and a completeness result for goal regression*, pages 359–380. Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy. Academic Press, San Diego, CA, 1991.

[Reiter, 1993] Raymond Reiter. Proving Properties of States in the Situation Calculus. *Artificial Intelligence*, 64(2):337–351, December 1993.

[Sandewall, 1994a] Erik Sandewall. Features and fluents, a systematic approach to the representation of knowledge about dynamical systems. Technical report, Department of Computer and Information Science, University of Linköping, Sweden, May 1994. To be published as a book (Oxford University Press).

[Sandewall, 1994b] Erik Sandewall. The range of applicability of some nonmonotonic logics for strict inertia. *Journal of Logic and Computation*, 1994. To appear.

[Thielscher, 1994] Michael Thielscher. An analysis of systematic approaches to reasoning about actions and change. Unpublished Manuscript, 1994.