Underlying Semantics for the Assessment of Reiter's Solution to the Frame Problem

Tania Bedrax-Weiss

CIRL and University of Oregon,
1269 University of Oregon,
Eugene OR 97403-1269.
tania@cirl.uoregon.edu

Leopoldo E. Bertossi

Departamento de Ciencia de la Computacion,
Escuela de Ingenieria,
Pontificia Unviersidad Catolica de Chile,
Casilla 306, Correo 22, Santiago, Chile.
bertossi@ing.puc.cl

Abstract

This paper presents an application of Sandewall's methodology [20] for assessing nonmonotonic entailment criteria for reasoning about action and change. We establish the correctness of Reiter's general solution to the frame problem [16] for a broad and well characterized class of problems. This is done by: (1) identifying a nonmonotonic entailment criterion corresponding to Reiter's solution; (2) discovering ontological and epistemological assumptions that we believe underly the domains of possible applicability of Reiter's approach; (3) defining an underlying semantics corresponding to these assumptions; and (4) providing a semantic correspondence between the models selected by the entailment criterion and the models determined by the underlying semantics.

We also compare our methodology to Kartha's evaluation [7] of Reiter's general solution with respect to the action logic semantics [4].

1 Introduction

We establish the correctness of Reiter's general solution to the frame problem [16] based on the situation calculus [9] for a broad and well characterized class of problems according to Sandewall's methodology [20]. Reiter's solution consists of a compilation of a preliminary axiomatization into a successor state axiom. The solution is monotonic only after the compilation, but it is nonmonotonic if deduction is performed directly on the preliminary axiomatization. We define the corresponding nonmonotonic entailment criterion and we assess this criterion with respect to Sandewall's underlying semantics. The process is as follows. Given a preliminary axiomatization describing the effects of actions (the starting point of Reiter's solution), we translate this axiomatization into a chronicle in Sandewall's temporal logic. We also identify a class of problems in Sandewall's taxonomy of epistemological and ontological assumptions and we define the corresponding underlying semantics. This semantics is determined by all possible simulations arising from the given chronicle. Finally, we prove there is a correspondence between the models obtained from Reiter's solution and the intended models obtained from the underlying semantics (see Figure 1).

Section 2 presents a summary and discussion of Reiter's solution to the frame problem in the situation calculus. Sandewall's methodology is briefly presented in Section 3. In Section 4 we assess Reiter's solution according to Sandewall's methodology. In Section 5, we compare our methodology to Kartha's evaluation [7] of Reiter's general solution with respect to action logic semantics [4]. Finally, we present our conclusions and sketch directions for future work.

2 Reiter's General Solution

Reiter [16] proposes a formalism to reason about action in the situation calculus [9]. Building on work by Haas [6], Pednault [10], and Schubert [21], he syntactically transforms a preliminary axiomatization that describes the preconditions and the direct effects for each action into a specification that completely describes each fluent at states that result from action executions. The preliminary specification is given in the situation calculus with predicates, fluents, actions, and a binary function symbol, *do*. The function symbol *do*

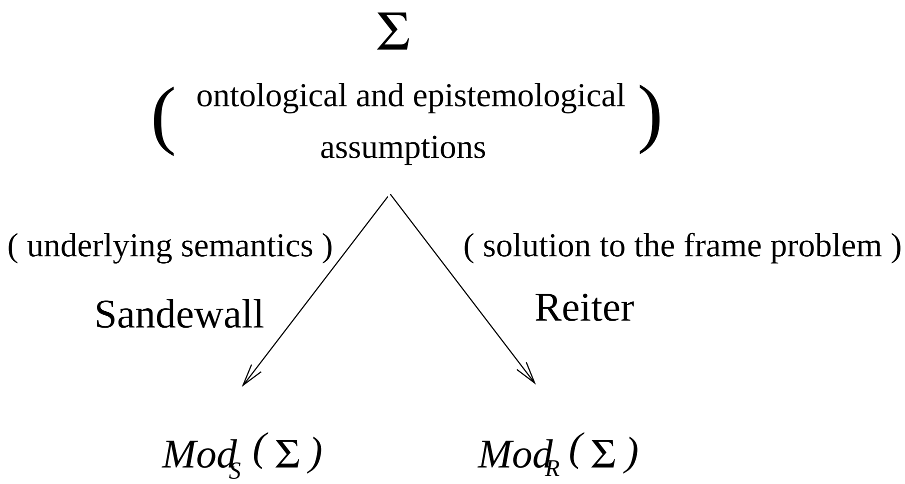

Figure 1: Correctness of Reiter's solution

takes as arguments an action and a situation and denotes the situation that results from executing the action in that situation. A distinguished predicate, *Poss*, is introduced in the language to specify the conditions under which actions are possible to execute. If *Poss* is satisfied in the state the action is to be executed, we will say that the action is a legal action in that state. Otherwise, we will say it is an illegal action in that state.

2.1 Axiomatization

The preliminary axiomatization consists of knowledge about S_0, the initial state of the world, which need not be complete, state independent knowledge and the following axioms:

Action Precondition Axioms: For each action name A we will have a precondition axiom of the form:

$$\pi_A(\bar{x}, s) \supset Poss(A(\bar{x}), s). \tag{1}$$

Positive Effects Axioms: For all pairs formed by a fluent R and an action name A, where the action causes fluent R to become true, an axiom of the form:

$$Poss(A(\bar{x}), s) \wedge \varphi(\bar{x}, \bar{y}, s) \supset R(\bar{y}, do(A(\bar{x}), s)). \tag{2}$$

That is, if action A is possible, and the preconditions for fluent R represented by the metaformula $\varphi(\bar{x}, \bar{y}, s)$ are true at state s, then fluent R becomes true in the successor state $do(A(\bar{x}), s)$ obtained after executing action A in state s. Here, $\bar{x}$ and $\bar{y}$ are parameters of the action and the fluent, respectively.

Negative Effects Axioms: For all pairs formed by a fluent R and an action name A, where the action causes fluent R to become false, an axiom of the form:

$$Poss(A(\bar{x}), s) \wedge \psi(\bar{x}, \bar{y}, s) \supset \neg R(\bar{y}, do(A(\bar{x}), s)). \tag{3}$$

The syntactic completion process performed on this preliminary specification is the following. For each fluent we combine the positive effect axioms that contain that fluent into a general single positive effect axiom. To see how, assume that we have only two positive effects axioms for fluent R:

$$Poss(A_1(\bar{x}), s) \wedge \varphi_1(\bar{x}, \bar{y}, s) \supset R(\bar{y}, do(A_1(\bar{x}), s))$$

and

$$Poss(A_2(\bar{z}), s) \wedge \varphi_2(\bar{z}, \bar{y}, s) \supset R(\bar{y}, do(A_2(\bar{x}), s)).$$

We may combine them into one general positive effect axiom for R:

$$Poss(a, s) \wedge [(\exists\, \bar{x})(a = A_1(\bar{x}) \wedge \varphi_1(\bar{x}, \bar{y}, s)) \vee (\exists\, \bar{z})(a = A_2(\bar{z}) \wedge \varphi_2(\bar{z}, \bar{y}, s))] \\ \supset R(\bar{y}, do(a, s)).$$

Thus, for each fluent, R, we obtain a general positive effect axiom of the form:

$$Poss(a, s) \wedge \chi(a, s) \supset R(do(a, s)).$$

In this axiom, a is a variable for actions that is considered (tacitly) universally quantified. Similarly, we obtain a general negative effect axiom for each fluent of the form:

$$Poss(a, s) \wedge \phi(a, s) \supset \neg R(do(a, s)).$$

The general effect axioms do not solve the frame problem. Reiter relies on a "Generalized Completeness Assumption" to do so. The GCA says that the general effect axioms for a fluent provide all the conditions under which a fluent may change its truth value. This assumption is made at a metalevel, but it inspires a syntactical transformation of the specification so that it remains first-order. After this assumption is made, we obtain the following successor state axiom for each fluent, R:

Successor State Axiom:

$$Poss(a, s) \supset [R(do(a, s)) \equiv [\chi(a, s) \vee (R(s) \wedge \neg\phi(a, s))]]. \tag{4}$$

This axiom is tacitly universally quantified over actions. Intuitively, it says that if action a is possible, then fluent R is true at the successor state that results from the execution of action a if and only if a is one of the actions causing R to be true (for which the fluent preconditions are true), or R was already true when a was executed and a is not one of the actions falsifying R. The successor state axioms provide, according to Reiter, a simple solution to the frame problem. There is one successor state axiom for each fluent and the successor state axioms will be short provided few actions change the fluent's truth value (which, he argues, is usually the case). In order for things to work properly, unique names axioms [13] are required for actions and states.

It is important to point out that Reiter's solution to the frame problem [16], considers only deterministic actions, that is, actions whose effects are completely known and where the preconditions of the actions are valid in the state they are executed in. Nonetheless, there is nothing in the logic to prevent the execution of illegal actions (actions whose preconditions are not satisfied in the state they are executed).[1] Also note that Reiter's solution as presented originally, [16] does not consider explicit ramifications but they are not disallowed if they appear embedded in the effect axioms.

Example

Let us consider a library database. For simplicity, we will only consider the actions of classifying a book (assigning a new ID to a newly acquired book) and deleting a book (if a classified book is missing we want to make it unavailable). Assume there is a table with the books that have not yet been classified, *unclassified*, and a table with the classified books *classified*. There will also be a table that records the books that are *missing*. When the books are classified they are assigned a unique ID number which encodes the ISBN number. Books to be classified have to be chosen from the *unclassified* books and books that are *deleted* need to have been *classified* and *missing*.

Preliminary Axiomatization

- Knowledge about the initial situation and the objects in the domain,
- Precondition Axioms for Actions,

$$unclassified(isbn, s) \supset Poss(ClassifyBook(isbn, id), s)$$

$$classified(isbn, id, s) \wedge missing(id) \supset Poss(DeleteBook(id), s)$$

- Effect Axioms for Actions,

$$\begin{array}{l} Poss(ClassifyBook(isbn, id), s) \supset \\ \quad classified(isbn, id, do(ClassifyBook(isbn, id), s)) \wedge \end{array}$$

[1] In real applications one actually extends the language to restrict the solution to accessible states, that is, states that can be reached after a finite sequence of legal actions.

$$\neg \mathit{unclassified}(isbn, do(ClassifyBook(isbn, id), s))$$

$$Poss(DeleteBook(id), s) \supset \\ \quad \neg \mathit{classified}(isbn, id, do(DeleteBook(id), s))$$

By syntactic transformation of the above preliminary specification of effects of actions, we obtain the final specification.

Syntactic Transformation: Combining the axioms for each fluent to characterize exactly under which conditions the fluent is true in a given successor situation,

$$(\forall\, a, s).Poss(a, s) \supset \\ \quad \mathit{classified}(isbn, id, do(a, s)) \equiv a = ClassifyBook(isbn, id) \vee \\ \quad\quad (\mathit{classified}(isbn, id, s) \wedge a \neq DeleteBook(id))$$

$$(\forall\, a, s).Poss(a, s) \supset \\ \quad \mathit{unclassified}(isbn, do(a, s)) \equiv (\forall\, id).a \neq ClassifyBook(isbn, id)$$

$$(\forall\, a, s).Poss(a, s) \supset missing(id, do(a, s)) \equiv TRUE$$

■

2.2 Reiter's Nonmonotonic Entailment Criterion

Reiter's solution to the frame problem is monotonic only in the sense that we obtain a first-order specification after the transformation has been performed on the preliminary axiomatization. This method is reminiscent of Clark's completion of logic programs [2]. ¿From this point of view, Reiter's solution represents a syntactic method of selection of the intended models and it is non-monotonic: the class of models selected via the syntactic transformation of a more restricted preliminary axiomatization need not be a subclass of the class of selected models of the less restricted axiomatization.

Example

We revisit our example to illustrate how Reiter's solution may be nonmonotonic. A logic is nonmonotonic if adding additional axioms to the hypothesis set invalidates a proof step. We will add additional action axioms to the axiomatization and show how it invalidates a consequence of the original axiomatization.

Let Σ be the preliminary axiomatization of our example. Assume that $\mathit{classified}(\mathit{isbn}1, \mathit{id}1, s1)$ holds. Given the following effect axioms for actions,

$$\begin{aligned}&\mathit{Poss}(\mathit{ClassifyBook}(\mathit{isbn}, \mathit{id}), s) \supset \\ &\quad \mathit{classified}(\mathit{isbn}, \mathit{id}, \mathit{do}(\mathit{ClassifyBook}(\mathit{isbn}, \mathit{id}), s)) \wedge \\ &\quad \neg \mathit{unclassified}(\mathit{isbn}, \mathit{do}(\mathit{ClassifyBook}(\mathit{isbn}, \mathit{id}), s))\end{aligned}$$

$$\begin{aligned}&\mathit{Poss}(\mathit{DeleteBook}(\mathit{id}), s) \supset \\ &\qquad \neg \mathit{classified}(\mathit{isbn}, \mathit{id}, \mathit{do}(\mathit{DeleteBook}(\mathit{id}), s))\end{aligned}$$

we may ask if $\mathit{classified}(\mathit{isbn}1, \mathit{id}1, \mathit{do}(\mathit{RemoveBook}(\mathit{id}1), s1))$ follows. Given that *RemoveBook* is not one of the actions that modify *classified* we may infer that *classified* will continue to hold in $\mathit{do}(\mathit{RemoveBook}(\mathit{id}1), s1)$.

On the other hand, if we add the following effect axiom for the new action *RemoveBook* and assume the precondition axiom is the same as the precondition axiom for *DeleteBook*,

$$\begin{aligned}&\mathit{Poss}(\mathit{RemoveBook}(\mathit{id}), s) \supset \\ &\quad \neg \mathit{classified}(\mathit{isbn}, \mathit{id}, \mathit{do}(\mathit{RemoveBook}(\mathit{id}), s))\end{aligned}$$

and we ask if $\mathit{classified}(\mathit{isbn}1, \mathit{id}1, \mathit{do}(\mathit{RemoveBook}(\mathit{id}1), s1))$ follows, we infer that it does not because *RemoveBook* is one of the actions that falsify *classified*.

■

Thus, adding axioms to the theory not necessarily grows the set of consequences of the theory. We proceed to define the notion of the nonmonotonic entailment criterion corresponding to Reiter's solution.

Assume that we have a set of formulas Σ corresponding to a preliminary axiomatization of a domain as described at the beginning of Section 2.1. Let $\mathcal{R}(\Sigma)$ denote Reiter's syntactic completion of Σ, as described above. A particular entailment criterion, $\approx\!\!|_R$, is associated with this completion process with a natural definition: for a formula φ in the language of Σ,

$$\Sigma \mathrel{|\!\!\approx_R} \varphi \quad \Leftrightarrow \quad \mathcal{R}(\Sigma) \models \varphi.$$

On the right-hand side of the equivalence we have the classical entailment criterion, i.e. defined in terms of a "surface semantics" for a first-order predicate calculus. This surface semantics is essentially given in terms of models for first-order languages that represent worlds that evolve in time. Thus, they consider all the possible histories of the world. The criterion $\models\!\!\approx_R$ on the left-hand side is nonmonotonic and can be interpreted as a model preference criterion that selects some of the models of Σ in the class determined by the surface semantics in order to check the validity of φ.

In order to assess Reiter's solution to the frame problem, we need to characterize the class of non-monotonic models selected by Reiter's entailment criterion. If we characterize these "preferred models" in a constructive way we will be able to compare the construction with the construction given by Sandewall's underlying semantics. We show that it is possible to define these models in a constructive way that relies in the successor state axioms.

The preferred models of Σ can be inductively characterized as follows. In the base case (no actions are executed), select all first-order models which satisfy the initial domain description. In the inductive step we make use of the successor state axioms. There can be two cases depending on whether the action is legal or illegal in that state. If a legal action is executed, select all first-order models such that for each pair (fluent, action), if the successor state axiom for the fluent mentions the action, the value of the fluent in the selected model is true if and only if in equation (4) a appears in χ or if it was true and a does not appear in ϕ, and it is negative otherwise. It can be checked that the successor state axiom for a fluent assigns only one value to the fluent at any given state [16], so there is at least one model for any consistent domain description.

If the action is illegal (its preconditions are not satisfied in the state the action is executed), nothing can be derived. We interpret this as a selection of all possible models of the "successor situation". As we have said before, there is no way to exclude this case from consideration since there is nothing in the logic that restricts the execution to legal actions only.[2] If we introduce a predicate of executability or accessibility of states [18] that restricts actions to this case, the definition of the selected models restricted to this predicate remains unchanged, with the exception of illegal actions. Observe that the models of Σ we have characterized are exactly those of $\mathcal{R}(\Sigma)$.

[2]Notice that *do* is a total function [18].

The class of models selected by Reiter's entailment criterion will be compared to those defined according to Sandewall's methodology in the following sections.

3 Sandewall's Methodology

3.1 Overview

We use Sandewall's methodology to validate Reiter's solution to the frame problem and to assess the range of applicability for a particular class of problems. One of the first things we need to do according to the methodology is to identify assumptions that have to be made about the domains that can be successfully subjected to Reiter's solution. In his book [20], Sandewall has identified several assumptions about possible domains for reasoning about action providing a taxonomy for them such that a problem belongs to a class in the taxonomy if the class captures all of the assumptions about the problem. In his taxonomy, he distinguishes two kinds of assumptions: *epistemological* and *ontological* assumptions. Epistemological assumptions include all of the assumptions that can be made about the description of the domain (i.e., how complete is the knowledge we have about the particular domains, how actions are described, etc.). Ontological assumptions include all of the assumptions made about the particular domain (i.e., whether actions are sequential or concurrent, deterministic or non-deterministic, whether changes occur within the execution of the action or there are delayed effects, etc.).

Once the assumptions about the domains have been identified as a class in his ontological taxonomy, one proceeds to choose a language for the description of problems of this class and completes the identification in the epistemological taxonomy according to the assumptions made about the description of the domain. The class corresponding to the problem description is the class composed of the ontological and epistemological classes. Reiter's solution to the frame problem requires the use of the language of the situation calculus which (given a particular domain) uniquely determines the corresponding class.

Once the class is identified, we need to do is to define the semantics for the class (if it has not previously been defined by Sandewall). The semantics is defined as a simulation in an inhabited dynamical system populated by an acting element, the *ego* and a reacting element, the *world*. They interact given a description of the domain in Sandewall's language, called a *chronicle*. The definition of the semantics for the class corresponding to Reiter's solution starts by showing how we translate Reiter's preliminary axiomatization into a chronicle in Sandewall's language. We proceed to define the simulation whereby the ego chooses actions described in the chronicle and the world reacts to the actions according to the corresponding "action law" of the chronicle. The simulation is recorded in *developments* from which the *intended models* are recovered. The assessment of Reiter's solution consists of comparing the intended models obtained from the simulation with the set of models obtained according to Reiter's entailment criterion defined in the previous section. The process of assessing Reiter's solution is summarized in Figure 2.

3.2 Inhabited Dynamical Systems

Sandewall's methodology is intended to address knowledge representation and reasoning issues in inhabited and inert dynamical systems where agents inhabiting the systems perform actions that may produce changes in the system. The system is just a collection of objects which have properties that if subject to change, we will call *fluents*. In our case, the system is inhabited because there will be some agent performing actions in the simulation. Sandewall distinguishes two levels of description of such a system. The lower level is the *material level* and describes the system as it "really is". The higher level is the *image level* and it is an abstraction (for instance, with discrete values or qualitative values, or partial interpretations) of the material level. The image level can be seen as the result of a perception function operating on the material level.

Agents in the system consist of two parts, an ego and a vehicle. The vehicle contains the perception function and performs the actions. The ego chooses actions to execute. The world consists of the ego's vehicle and the material system. Thus, actions that the ego chooses, are realized by the world and the fluent values change according to the "laws" governing the material system.

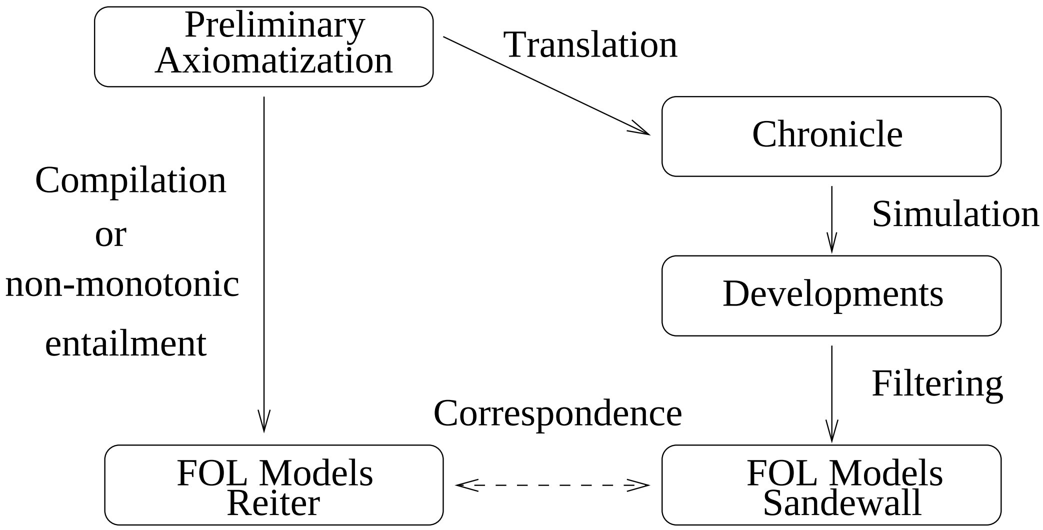

Figure 2: Assessment

3.3 Underlying Semantics

The precise evolution of the scenario is determined by the underlying semantics and depends on the (ontological and epistemological) assumptions about the domain. There are several semantical layers ranked by level of expressiveness. The simpler one is the *encapsulated semantics* in which each action is characterized by the values of fluents before and after the action is executed but they are unknown during the execution of the action. This is the semantics we are interested in because it corresponds to the semantics of the situation calculus, where the duration of executions of actions are unspecified and the values of fluents are known at all times when actions are not being executed (i.e., in all situations).[3] The *trajectory semantics* extends the encapsulated semantics to include information about exactly when (within the execution of the action) the changes are produced.

In the encapsulated semantics, the reacting element of the interaction or the world, W, is characterized by the pair of indexed sets $\langle \texttt{Infl}, \texttt{Rstat} \rangle$, where $\texttt{Infl}(E, r)$ is the set of fluents that are changed by action E in state r and where $\texttt{Rstat}(E, r)$ is the set of (partial) states assigning values only to fluents in $\texttt{Infl}(E, r)$. The states in Sandewall's semantics are mappings from fluents to values of fluents. They are time-independent since they are not associated to timepoints until they arise in the simulation, i.e., one state may arise at different timepoints in different simulations. Contrast this with situations in Reiter's solution where the situation is different if it appears after a different sequence of actions has been executed.

If we execute action E, the set of states resulting from the execution is: $\{r \oplus r' \mid r' \in \texttt{Rstat}(\mathrm{E,r})\}$, where r is the current state and r' is the partial state restricted to the fluents in $\texttt{Infl}(E, r)$. If the action is completely specified and it is deterministic, there will be only one partial state r' in $\texttt{Rstat}(E, r)$and the result of the execution is a single state $r \oplus r'$. The symbol $\oplus$ is borrowed from the Z notation in software specification for overriding the values of whatever is at the left of the symbol for whatever is at the right of the symbol. That is, values of r are overridden by the corresponding values in r'. Inertia is represented by $\oplus$ since the fluents that are not in r' remain unchanged. It is clear that the underlying (epistemological and ontological)

[3] Extensions to the situation calculus that consider durations, concurrency, and time may be found in [5] and [12]. Since Reiter does not consider these extensions, we do not consider them in our analysis.

assumptions come into play here, where for instance, some pairs may be missing if we don't have complete information on action executions.

We construct the pair $\langle \texttt{Infl}, \texttt{Rstat} \rangle$ from the action descriptions in the chronicle and according to the ontological and epistemological assumptions. For each action E in the chronicle and each state r, we include in $\texttt{Infl}(E, r)$ all of the fluents that are affected by the action in the state and we include in $\texttt{Rstat}(E, r)$ the values that are assigned to the fluents in $\texttt{Infl}(E, r)$ by the action law in the chronicle.

Once we have defined $\langle \texttt{Infl}, \texttt{Rstat} \rangle$, we obtain the set of models of the chronicle through the simulation which is (roughly) given by the following interaction between the ego and the world. The scenario evolves in time. If the ego chooses to execute some action, E, at time $\mathbf{s}$ (whether or not is in accordance with the schedule given, will be filtered after all possible simulations are generated), where the state at time $\mathbf{s}$ is r, the world responds by choosing an arbitrary timepoint $\mathbf{t}$, greater than $\mathbf{s}$, such that E is executed between $\mathbf{s}$ and $\mathbf{t}$. Recall that the duration of the action is unspecified in the encapsulated semantics, so the world is free to choose any "future timepoint" $\mathbf{t}$. For each fluent of r not in $\texttt{Infl}(E, r)$ its value at all times between $\mathbf{s}$ and $\mathbf{t}$ are the same as its value at $\mathbf{s}$ and for each fluent in $\texttt{Infl}(E, r)$, its value at $\mathbf{t}$ will be the value determined by its value in r', but its value at any time between $\mathbf{s}$ and $\mathbf{t}$ can be chosen arbitrarily. The world advances time and at the new time, $\mathbf{t}$, the ego may or may not choose to execute another action. If the ego chooses to execute another action, the world reacts in a similar way but this time getting the information from the pairs $\langle \texttt{Infl}, \texttt{Rstat} \rangle$, corresponding to the new action and the new state. If the ego chooses not to execute any actions, the world only advances time, propagating all fluent values.

This interaction is recorded in a structure called the *development*. The interaction takes place from one breakpoint to the next during which no information can be obtained, i.e., observations may take place only at breakpoints. Formally, a *development* is a tuple $\langle \mathcal{B}, M, R, \mathcal{A}, \mathcal{C} \rangle$, where:

$\mathcal{B}$: is the set of interpretations of the form $\{0, \ldots, k\}$ for breakpoints (which in the case of single-step actions coincide with the timepoints), where k corresponds to the current state,

M: is the valuation for timepoints (assigns members of $\mathcal{B}$ to timepoints) and objects,

R: is the history of the fluent changes from timepoint 0 to timepoint k, actually, we consider this to be an indexed set of assignments of values to fluents,

$\mathcal{A}$: is the set of actions that have been executed,

$\mathcal{C}$: is the set of actions that currently are being executed. If only one action can be executed at a time, it will always contain a single action.

There are a few things we would like to point out. First of all, R is a total function that given a fluent and a timepoint, determines the value of the fluent at that timepoint. Thus, R is a complete history. Also, in $\mathcal{A}$, actions are time stamped indicating the starting and termination times, while in $\mathcal{C}$, actions are only time stamped with the starting times.

More precisely, the moves of the ego and the world are recorded in the developments as follows: If the ego chooses an action, ϵ, from the chronicle at time $\mathbf{s}$, the world updates $\mathcal{C}$ with the tuple $\langle \mathbf{s}, \epsilon \rangle$ and reacts as follows: (1) it updates R according to $\texttt{Rstat}(\epsilon, r)$ where r is the state corresponding to time $\mathbf{s}$ by overriding any previous fluent values with new ones found in $\texttt{Rstat}(\epsilon, r)$; (2) it updates $\mathcal{B}$ by adding a new timepoint $\mathbf{t}$; (3) it removes $\langle \mathbf{s}, \epsilon \rangle$ from $\mathcal{C}$; and (4) adds $\langle \mathbf{s}, \epsilon, \mathbf{t} \rangle$ to $\mathcal{A}$.

This interaction, whereby the ego makes a move and the world reacts to the move, goes on repeatedly for all actions defined in the chronicle and for all histories that arise from the moves of the world. This simulation is defined for every initial state of the particular domain[4] and every valuation of the timepoints that satisfies the time restrictions imposed by the chronicle (e.g., single-step actions). This generates all the possible histories of the scenario, given all the possible executions of actions in all possible initial states. If the chronicle contains a schedule, the developments obtained through the simulation are filtered to select only those that are consistent with the schedule. There is a similar filtering process with the observations. From these developments we recover the intended models.

[4]There are many ways to complete the initial state if the state is not completely specified.

4 Analysis of Reiter's Solution

In order to compare the intended models to those selected by Reiter's formalism, we must first identify the class of problems in Sandewall's taxonomy for which Reiter's solution is applicable. Then we obtain a translation of Reiter's preliminary axiomatization into a scenario description in Sandewall's language. This translation will define a chronicle which contains several kinds of axioms: observations (corresponding to the description of the initial domain description), action laws (corresponding to a translation of the effect axioms), and, possibly, a schedule (corresponding to a sequence of actions that will define a future state). We give a brief description of this translation later on.

The intended models are obtained from this chronicle according to the underlying semantics and a filtering process that selects the set of models consistent with the observations and the schedule of the chronicle. Finally, we state the semantic correspondence between the models selected by Reiter's formalism and Sandewall's underlying semantics.

4.1 Ontological and Epistemological Assumptions

In a philosophical approach to the situation calculus, to Reiter's solution as presented in [16], and to problems with inertia (worlds resisting change), we discover the ontological and epistemological assumptions which characterize the class of problems for which Reiter's solution is intended to be applicable. Epistemological assumptions are those that refer to the description of the domain and ontological assumptions are those that refer to the properties of the domain, as in [20].

We start by identifying the ontological assumptions that underly the domains describable in the situation calculus: the assumption of *single step actions* since actions are executed from one situation to another, and of *sequential actions* since actions are executed one at a time. We must mention though, that we will consider time and not situations as our temporal framework (actions are executed from time t to time $t + 1$).[5] We will interpret, then, each situation as corresponding to a timepoint, and such that the ini-

[5] We do this translation to follow Sandewall's initially proposed language, but we could have done everything with situations.

tial situation corresponds to timepoint 0, and situation s_i corresponds to timepoint i. The actual translation is presented in the following section.

With respect to the ontological assumptions about the domains that may be formalized in Reiter's framework, we must mention that there is a restricted case of *inertia*: nothing may ever change unless an action is executed. If the action that is being executed is legal the only changes that are recorded are those explicit in the action description, but if the action is illegal, the successor state is undetermined and we must allow for things to change unexpectedly.

Reiter's solution considers *deterministic actions*, that is, actions whose effects are completely determined by the state in which the action is performed. Nonetheless, the results of the actions are state dependent (they depend on the situation the action is executed in). This ontological assumption is what Sandewall calls *alternative results*.

Reiter's formalism allows different kinds of changes, but they are all observed immediately after the execution of the action (i.e., there are no delayed effects). Sandewall identifies two types of (dependency) changes in fluent values *structural dependency* and *local dependency*. Structural dependency is present when there are properties of some object defined in terms of properties of other objects. Local dependencies are restrictions on the values of properties of the same object. These kinds of dependencies are not excluded in Reiter's formalism and may appear in the effect axioms. If there appear separately, they may be compiled into the precondition or effect axioms [8, 12].

Summarizing, we have identified the following ontological assumptions: *inertia* (restricted to legal actions), *single-step actions* (integer time), *alternative results of actions, deterministic actions* (restricted to legal actions), *local or structural dependencies between fluents, and sequential actions.* [6]

With respect to the epistemological assumptions, (i.e., assumptions that refer to the kind of knowledge that is available and how it is represented), we found that there may be *incomplete knowledge about the initial state*, *each change is completely characterized* in the sense that all actions that produce a change are known (GCA), *the axiomatization is consistent*, and *actions are*

[6] Note that there are no surprises (i.e., unspecified fluent changes except in the illegal case) or miracles (the only actions that are allowed to occur are those specified in the preliminary axiomatization of the domain).

defined in terms of preconditions, so the effects of an action are completely known only if it is legal. If the action is illegal, the action's effects are not determined.

4.2 Translation into Sandewall's Language

Sandewall defines an Elementary Feature Logic which is basically a first-order temporal language (he calls it logic because he later defines different entailment criteria on this language) [20] to describe a particular scenario of a domain. The entailment criteria to reason in this scenario are defined separately by giving its intended set of models. In order to assess Reiter's solution we translate the preliminary axiomatization into a chronicle in Sandewall's language and define the sets of intended models according to Reiter's entailment criterion.

The translation is straightforward so we will only briefly sketch how this is done. We will interpret each situation as corresponding to a timepoint, and such that the initial situation corresponds to timepoint 0, and situation s_i corresponds to timepoint i. For each statement of the form $R(x, S_0)$ we introduce an observation statement in Sandewall's language of the form $[0]R(x)$, where this stands for $R(x)$ is true at time 0. For each statement φ, that does not have a state variable as argument, i.e. state independent knowledge, we introduce an observation statement of the form $[]\varphi$, meaning for all t, φ. To translate the effect and precondition axioms into action laws we need to compile for each action, the effect axioms and precondition axioms that mention that action to obtain a first-order equivalent "action law":

Take action $A(\bar{x})$ and all the axioms of the form (2) and (3) in Section 2.1, that mention the action and write an axiom of the following form, where we have already substituted $Poss(A(\bar{x}), s)$ by its definition (or the left-hand side of equation (1) in Section 2.1)

$$
\begin{aligned}
&[t, t+1]A(\bar{x}) \Rrightarrow \\
&\qquad [t]\pi_A(\bar{x}) \rightarrow [t]\chi' \rightarrow [t+1]R{:=}\text{true} \wedge \\
&\qquad [t]\pi_A(\bar{x}) \rightarrow [t]\phi' \rightarrow [t+1]R{:=}\text{false} \wedge \\
&\qquad [t]\pi_A(\bar{x}) \rightarrow [t]\varepsilon' \rightarrow [t+1]Q{:=}\text{true} \wedge \ldots .
\end{aligned}
$$

The symbol := has an implicit meaning of possible change, that is, it means that the fluent will have the value at the right-hand side of the symbol (in this case, true) regardless of what its value was at time t. The formulas χ', ϕ', and ε' are the natural translations of their counterparts in the situation calculus.

If given in the preliminary axiomatization, the sequence of ground actions T of the form $do(A_{n-1}, \ldots, do(a_1, do(A_0, S_0)) \ldots)$ will be translated to:

$$[0]A_0, [1]A_1, \ldots, [n-1]A_{n-1}.$$

Example

Let us return to the example in Section 2.1 and let us rewrite the axioms according to the translation we have just described.

Chronicle

- Observations about the initial state and objects in the domain,
- Action Laws,

$$\begin{aligned}
&[t, t+1]\mathit{ClassifyBook}(\mathit{isbn}, \mathit{id}) \Rrightarrow \\
&\qquad [t]\mathit{unclassified}(\mathit{isbn}) \rightarrow [t+1]\mathit{classified}(\mathit{isbn}, \mathit{id}) \wedge \\
&\qquad\qquad [t+1]\neg \mathit{unclassified}(\mathit{isbn})
\end{aligned}$$

$$\begin{aligned}
&[t, t+1]\mathit{DeleteBook}(\mathit{id}) \Rrightarrow \\
&\quad [t]\mathit{classified}(\mathit{isbn}, \mathit{id}) \wedge [t]\mathit{missing}(\mathit{id}) \rightarrow [t+1]\neg \mathit{classified}(\mathit{isbn}, \mathit{id})
\end{aligned}$$

The models of this chronicle are recovered from the underlying semantics determined by the developments generated by the simulation of executions of the actions *ClassifyBook* and *DeleteBook*. ■

In the following section we define the underlying semantics for chronicles obtained through this translation process from preliminary axiomatizations. We also define the set of models of the chronicle and we compare these sets of models to those selected by Reiter's nonmonotonic entailment criterion as defined in Section 2.2.

4.3 Assessment

According to Sandewall's methodology described in Section 3, we identified the class of problems we intend to solve with Reiter's solution. This class was defined according to the ontological and epistemological assumptions. We need to characterize the intended models that correspond to this class by means of the underlying semantics.

4.3.1 Underlying Semantics for Reiter's Solution

The simulation (interaction between the ego and the world) determines all the possible histories of the particular instance of the domain for a sequence of actions. More precisely, given: (a) an arbitrary chronicle, (b) a sequence of actions, (c) an ego, and (d) a world (all of these satisfying the ontological and epistemological assumptions), the interaction will determine the possible developments for that chronicle. The intended models will be obtained from these developments.

In the case of Reiter's solution, we are given an arbitrary preliminary axiomatization. We translate this axiomatization into a chronicle as described in Section 4.2. It follows, that whatever fluents are true in the initial situation are true at time 0. The initial development will register the fluent values in its history component, R_0. The updating of this development proceeds according to the interaction between the ego and the world.

We construct the pair $\langle \mathtt{Infl}, \mathtt{Rstat} \rangle$ from the chronicle with the action descriptions and according to the encapsulated semantics as in Section 3.3. These pairs will be defined only if the preconditions for E are valid in the state r. We need to extend the encapsulated semantics to allow for illegal actions. How will the world interpret a move when the ego invokes an illegal action? There are two alternatives. One is to ignore any effects of whatever illegal actions the ego has invoked (in the spirit of the STRIPS [3] semantics) and the other is to produce some unknown change or set of changes. In the first case, we would set $\mathtt{Rstat}(\mathrm{E,r})= \emptyset$ and $\mathtt{Infl}(\mathrm{E,r})= \emptyset$ for each illegal action (illegal with respect to state r), whereas in the second case, for each illegal action, we would assign to $\mathtt{Rstat}(E, r)$ a set of partial states where each partial state contains some combination of changes, and to $\mathtt{Infl}(E, r)$ the set of all fluents (since all of them may change). Reiter, in his extensions [17, 19, 18, 11] to the initial proposal [16] uses induction to restrict the treestructured models so that only legal situations (those that arise from legal executions of actions or the initial situation) appear in the tree. Since we are analyzing his initial proposal, we choose the second alternative in the spirit of having a more realistic representation for unknown effects —something may or may not change and we do not know what it will be.

With this in mind, the moves of the ego and the world are recorded in the developments as follows. If the ego chooses an action, ϵ specified in the chronicle, at time t, the world updates $\mathcal{C}$ with the tuple $\langle t, \epsilon \rangle$ and reacts to the action by branching out on each element of $\texttt{Rstat}(\epsilon, r)$ and updating each branch as described in Section 3.3. The only difference is that he is forced, by the assumptions, to pick $t + 1$ as the ending time of the interval in which the action is executed if t is the starting time.

The interaction goes on repeatedly for all actions defined in the chronicle and for all histories that arise up to a certain timepoint. Interactions of this sort take place for every initial history, R_0, consistent with the partial description of the initial state of the particular domain. This collection of interactions generate all the possible evolutions of the chronicle.

The developments are tree structures. Since the initial state of the domain description is not completely specified, we will have several tree structures corresponding to the evolution of the world, namely, the trees whose history component of the initial development is consistent with the given description of the initial state. The models defined by the underlying semantics are obtained from the developments. They are first-order logic models of the history slices of the developments at all timepoints. The models of the chronicle are the models of the underlying semantics that are consistent with the observations and the schedule of the chronicle (see Figure 3).

Notice that the underlying semantics defined in this way is independent of Reiter's entailment criterion, since it depends only on the assumptions of the domain. It is also easy to check that,

Remark 1 *The underlying semantics completely captures the assumptions we have identified in Section 4.1 and it rules out the others listed in Sandewall's book (Chapters 1 and 2).*

The construction of the underlying semantics via the translation, simulation, and filtering can be seen as selecting some preferred models from the models given by the surface semantics. Thus, a new nonmonotonic entailment criterion arises, which we shall refer to as $\approx_S$.

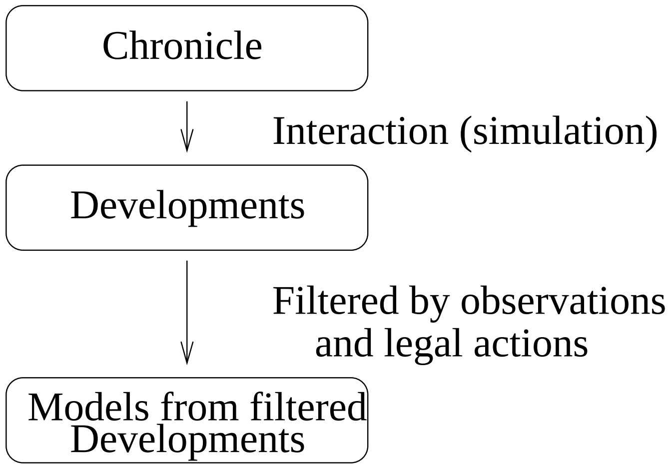

Figure 3: Models of the chronicle

Example

We return to our canonical example to illustrate the sets $\mathtt{Infl}(E, r)$ and $\mathtt{Rstat}(E, r)$. Assume we have one book with corresponding ISBN number $isbn1$ and ID number $id1$. Figures 1 and 2 describe the contents of $\mathtt{Infl}(E, r)$ and $\mathtt{Rstat}(E, r)$ when $E = \mathit{ClassifyBook}$. To improve the reading of the tables, we introduce the following abbreviations: $\mathit{unclassified}(isbn1)$ will be abbreviated by $u(isbn1)$, $\mathit{classified}(isbn1, id1)$ by $c(isbn1, id1)$, and $\mathit{missing}(id1)$ by $m(id1)$.

One development may look like this:

init $\langle\{0\}, \varnothing, \{\langle t \mapsto 0, u(isbn1), \neg c(isbn1, id1), \neg m(id1)\rangle\}, \varnothing, \varnothing\rangle$

ego $\langle\{0\}, \varnothing, \{\langle t \mapsto 0, u(isbn1), \neg c(isbn1, id1), \neg m(id1)\}, \varnothing,$
$\{\langle 0, \mathit{ClassifyBook}(isbn1, id1)\rangle\}\rangle$

world $\langle\{0, 1\}, \varnothing, \{\langle t \mapsto 0, u(isbn1), \neg c(isbn1, id1), \neg m(id1)\rangle,$
$\langle t \mapsto 1, \neg u(isbn1), c(isbn1, id1), \neg m(id1)\rangle\},$
$\{\langle 0, \mathit{ClassifyBook}(isbn1, id1), 1\rangle\}, \varnothing\rangle$

ego $\langle\{0, 1\}, \varnothing, \{\langle t \mapsto 0, u(isbn1), \neg c(isbn1, id1), \neg m(id1)\rangle,$
$\langle t \mapsto 1, \neg u(isbn1), c(isbn1, id1), \neg m(id1)\rangle\},$
$\{\langle 0, \mathit{ClassifyBook}(isbn1, id1), 1\rangle\}, \{\langle 1, \mathit{DeleteBook}(id1)\rangle\}\rangle$

etc.

■

Starting state r	$\mathtt{Infl}(\mathit{ClassifyBook}, r)$
$\{u(isbn1), \neg c(isbn1, id1), \neg m(id1)\}$	$\{c(isbn1, id1), u(isbn1)\}$
$\{u(isbn1), c(isbn1, id1), \neg m(id1)\}$	$\{c(isbn1, id1), u(isbn1), m(id1)\}$
$\{u(isbn1), c(isbn1, id1), m(id1)\}$	$\{c(isbn1, id1), u(isbn1), m(id1)\}$
$\vdots$	$\vdots$

Figure 4: The fluents influenced by action *ClassifyBook*.

Starting state r	$\texttt{Rstat}(\mathit{ClassifyBook},r)$
$\{u(\mathit{isbn}1), \neg c(\mathit{isbn}1, \mathit{id}1), \neg m(\mathit{id}1)\}$	$\{\neg u(\mathit{isbn}1), c(\mathit{isbn}1, \mathit{id}1)\}$
$\{u(\mathit{isbn}1), c(\mathit{isbn}1, \mathit{id}1), \neg m(\mathit{id}1)\}$	$\langle\{\neg u(\mathit{isbn}1)\}, \{u(\mathit{isbn}1)\}\ldots\rangle$
$\{u(\mathit{isbn}1), c(\mathit{isbn}1, \mathit{id}1), m(\mathit{id}1)\}$	$\langle\{m(\mathit{isbn}1)\}, \{\neg m(\mathit{isbn}1)\}\ldots\rangle$
$\vdots$	$\vdots$

Figure 5: The partial states resulting from the execution of *ClassifyBook*.

4.3.2 Semantic Correspondence

In this section we state the correspondence between the first-order models (the interpretations of the fluents, in particular) of both frameworks — Reiter's and ours based on Sandewall's methodology. This result applies regardless of whether the executed actions are legal or illegal.

Given a preliminary axiomatization Σ in Reiter's formalism, we want to prove that the intended models of the corresponding chronicle correspond to the models obtained from Reiter's non-monotonic entailment criterion acting on this axiomatization. In order for things to work properly we demand from the preliminary axiomatization the following *consistency property.*

Definition 2 *A preliminary axiomatization Σ, is said to have the* consistency property *if it is not the case that $Poss(a, s)$ is true and that the fluent preconditions for both positive and negative change (φ and ψ in equation (3) and equation (2) in Section 2.1) for all fluents affected by a are valid in s.*

All we are asking for is that the fluent value be determined either positively or negatively in each situation.

Let us consider a chronicle Υ, arising from the translation of a preliminary axiomatization Σ of a particular domain description in Reiter's formalism (se Section 2.1. If Σ has the consistency property, the chronicle Υ will be well defined in the sense that at least one model will be selected. If a sequence of actions is specified along with Σ, Υ will contain a schedule corresponding to that sequence of transactions. If Υ contains a schedule T, the underlying semantics provides the nonempty set of models for Υ according to the execution of T. If Υ doesn't contain a schedule T, the underlying semantics provides models for all possible schedules arising from the actions defined in Υ.

Note that the models of Σ according to Reiter's nonmonotonic entailment criterion[7], $Mod(\mathrm{R}(\Sigma))$, is a class of models of Σ of the form $M = \langle$O,S,A, $do^M, S_0^M, F_1^M, \ldots\rangle$, where

O is the object domain

S is the situation domain (e.g., contains S_0)

A is the action domain

do^M is a function that maps actions and situations to new situations

$S_0{}^M$ is the initial situation

$F_1{}^M$ is a fluent such that all the arguments except the last one range over O, it stands for a relation on these arguments, and the last argument ranges over S.

These are just standard first-order models of Σ. The models according to the underlying semantics are defined as follows.

Definition 3 *$Mod(\Upsilon)$ is the tuple $\langle M,R \rangle$, where*

M is the interpretation for objects and timepoints

R is the history component of the development determining what is true at each timepoint according to the ego/world interaction.

If $Mod(\mathrm{R}(\Sigma))$ are the models according to Reiter's nonmonotonic entailment criterion, then, on one hand, we consider $Mod(\mathrm{R}(\Sigma))$ and if there is a particular sequence of actions, T, we restrict the models to the situation that results from the execution of T. On the other hand, we define the developments of Υ as constructed above for the underlying semantics. We filter the developments to select only those whose history components are consistent with the observations. If there is a particular sequence of actions, or schedule, T, we filter the developments again in order to select only the developments consistent with the schedule T, i.e., containing each action in T in their $\mathcal{A}$ component at the same timepoints (see Section 4.3.1). The models are obtained after all necessary filterings are done.

[7]See Section 2.2.

Theorem 4 *Let Σ be a preliminary axiomatization that has the consistency property. Let Υ be the corresponding chronicle. $Mod(\Upsilon)$ at each time slice corresponds to $Mod(R(\Sigma))$ at each situation.*

The theorem is proved inductively comparing the interpretations of the $n\text{–}th$ sections of the developments with the interpretations of situation s_n. Since we use time as our general framework, at each timepoint, the interpretations in $Mod(\mathrm{R}(\Sigma))$ assign values to each situation corresponding to that timepoint, whereas the interpretations in $Mod(\Upsilon)$ assign values to whole histories. This forces us to compare the interpretations of the fluents at each time slice of the history.[8]

¿From the theorem we can obtain a soundness and completeness result:

Corollary 5 *(a) For any sentence φ,*

$$\Sigma \models\!\!\approx_R \varphi \quad \textit{iff} \quad \Sigma \models\!\!\approx_S \varphi$$

(b) If T is specified in addition to Σ,

$$\Sigma \models\!\!\approx_R \varphi(T) \quad \textit{iff} \quad \Sigma \models\!\!\approx_S \varphi(T).$$

Here, $\models\!\!\approx_R$ is the entailment criterion corresponding to Reiter's nonmonotonic solution defined in Section 2.2 and $\models\!\!\approx_S$ is the nonmonotonic entailment criterion corresponding to the underlying semantics as defined in Section 4.3.1. The theorem and the corollary state that Reiter's solution to the frame problem is correct and complete for the wide class of domains (and their preliminary axiomatizations) that satisfy the identified ontological and epistemological assumptions (Section 4.1).

5 Comparison to Kartha's Evaluation

Kartha [7] has already presented an analysis of Reiter's solution with respect to Gelfond and Lifschitz's action logic $\mathcal{A}$ [4]. He proves that Reiter's solution captures the semantics of $\mathcal{A}$ starting from a description (chronicle) in $\mathcal{A}$ and then translating it into Reiter's formalism. This translation is sound and complete.

[8]We are grateful to Matt Ginsberg for pointing this out to us.

Although our work is similar to Kartha's in the sense that we both characterize and compare the models for a given scenario description, there are several differences. In $\mathcal{A}$ the semantics of a description is fixed, i.e., the entailment criterion is the same for all descriptions, which leaves no option but to translate from $\mathcal{A}$ into Reiter's formalism. Sandewall's language, on the other hand, allows the entailment criterion to vary since it is independent of the language used to describe the scenario. Because of this, we can define corresponding non-monotonic entailment criteria in Sandewall's framework and perform the translation in either direction. Sandewall's methodology provides a certain degree of flexibility that is not present in Kartha's approach.

There is another difference with respect to both methodologies. Sandewall gives a general language and a general list of assumptions which he restricts when he identifies exactly the class of problems to be addressed. Any assessment based on $\mathcal{A}$, on the other hand, must extend the language and the semantics (if necessary). It is not always clear how to extend the language to capture the desired effect, while performing restrictions is usually pretty straightforward. So the difference is one of narrowing or extending the language and if one is extending the language, one must define the corresponding "fixed" semantics for the new language as Kartha does.

Another difference is that Sandewall analyses a criterion with respect to the ontological and epistemological assumptions that underlie a particular domain while Kartha's evaluation considers the models given by two different entailment criteria. In this sense, Sandewall's methodology may be said to check the models selected by the entailment criterion against "reality", whereas Kartha is forced to compare two different abstractions of reality (that may or may not be adequate).

Our work is a little bit less restrictive than Kartha's work because our work evaluates Reiter's framework for all domains that can be formalized in his framework, while Kartha restricts his analysis to those expressible in the language $\mathcal{A}$. The difference appears to be that the language $\mathcal{A}$ considers named actions and fluents with no quantification over any of these, so it can be seen as a propositional case. Reiter's framework is full first-order logic.

6 Conclusions

We have made several contributions: (1) we identified a nonmonotonic entailment criterion corresponding to Reiter's syntactic compilation and we inductively defined the preferred models; (2) we identified the ontological and epistemological assumptions that underlie the situation calculus and Reiter's solution in general; (3) we defined an underlying semantics that corresponds to the class of assumptions that considers legal and illegal actions, thereby extending Sandewall's work; (4) we proved a semantic correspondence between Reiter's models and the models of the underlying semantics, which holds for legal and illegal actions; (5) as a consequence, we obtained that Reiter's solution is correct and complete for a broad class of problems; (6) we showed that Sandewall's methodology is appropriate for this kind of systematic assessments; and (7) we show that this methodology has advantages over Kartha's assessment of Reiter's solution based on Lifschitz's action logic semantics.

7 Future Work

There are still several issues that need to be addressed. We have only considered Reiter's original solution [16], not its extension to some kinds of ramifications [8] and the restriction to accessible states [18]. In [19, 17], Reiter formalizes the evolution of a database according to his solution to the frame problem [16]. In [1], this formalization was positively assessed with Sandewall's methodology. In that paper, there was a restriction to ground databases, i.e., without null values. It seems interesting to extend their results to null values by making a logical reconstruction of the database as described by Reiter [14, 15]. Other extensions to databases may also be considered, such as, concurrent actions and active databases.

8 Acknowledgments

We are grateful to Ray Reiter, Erik Sandewall, and Javier Pinto for stimulating conversations and advice. This research has been supported partly by DIUC, SAREC/CONICYT, Fundacion DICTUC, and FONDECYT (#193-

0554) grants, partly by the Air Force Office of Scientific Research under contract 92-0693 and by ARPA/Rome Labs under contracts F30602-91-C-0036 and F30602-93-C-00031. Tania Bedrax-Weiss is specially thankful David Etherington and to everyone at CIRL for their support during the development of this work. Leopoldo Bertossi visited Linköping University in Sweden and is grateful to Erik Sandewall for his hospitality.

References

[1] T. Bedrax-Weiss and L. Bertossi. An assessment of Reiter's general monotonic solution to the frame problem and its applicability to database updates. In C. Schwind, editor, *Workshop of Logic and Change, ECAI-94*, pages 5–18, Amsterdam, The Netherlands, August 1994.

[2] K.L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 292–322, New York, 1978. Plenum Press.

[3] R. Fikes and N. Nilsson. STRIPS: A new approach to application of theorem proving procedures to problem solving. In *Readings in Planning*, 1971.

[4] M. Gelfond and V. Lifschitz. Describing action and change by logic programs. In *Joint International Conference and Symposium on Logic Programming*, 1992.

[5] M. Gelfond, V. Lifschitz, and A. Rabinov. What are the limitations of the situation calculus? In *AAAI Symposium of Logical Formalization of Commonsense Reasoning*, pages 59–69, Stanford, CA, 1991.

[6] A. Haas. The case for domain-specific frame axioms. In F. Brown, editor, *The frame problem in artificial intelligence. Proceedings of the 1987 Workshop*, pages 343–348, Los Altos, CA, 1987. Morgan Kaufmann Publishers, Inc.

[7] G. N. Kartha. Soundness and completeness theorems for three formalizations of action. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, volume 1, pages 724–729, Chambéry, France, 1993.

[8] F. Lin and R. Reiter. State constraints revisited. *Journal of Logic and Computation*, 4(5):655–678, 1994. Special Issue on Action and Processes.

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

[10] E. Pednault. ADL: Exploring the middle ground between STRIPS and the situation calculus. In H. Levesque R. Brachman and R. Reiter, editors, *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, pages 324–332, San Mateo, CA, 1989. Morgan Kaufmann Publishers, Inc.

[11] J. Pinto and R. Reiter. Adding a time line to the situation calculus. In *Working Notes of Common Sense'93, The Second Symposium on Logical Formalizations of Commonsense Reasoning*, pages 172–177, Austin,TX, January 1993.

[12] Javier Pinto. *Temporal Reasoning in the Situational Calculus*. PhD thesis, Department of Computer Science, University of Toronto, 1994.

[13] R. Reiter. Equality and domain closure in first-order databases. *Journal of the ACM*, 27(2):235–249, April 1980.

[14] R. Reiter. Towards a logical reconstruction of relational databases theory. In J. Mylopoulos and J. Schmidt, editors, *On Conceptual Modelling: Perspectives from Artificial Intelligence, Databases and Programming Languages*, pages 191–233. Springer-Verlag, 1984.

[15] R. Reiter. A sound and sometimes complete query evaluation algorithm for relational databases with null values. *Journal of the ACM*, 2(33):349–370, 1986.

[16] R. Reiter. The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In V. Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, 1991.

[17] R. Reiter. Formalizing database evolution in the situation calculus. In *Proceedings of the Fifth Generation Computer Systems*, Tokyo, Japan, June 1992.

[18] R. Reiter. Proving properties of states in the situation calculus. *Artificial Intelligence*, 64(2):337–351, 1993.

[19] R. Reiter. On specifying database updates. *Journal of Logic Programming*, 25(1):53–91, 1995.

[20] E. Sandewall. *Features and Fluents. The Representation of Knowledge about Dynamical Systems. Volume I.* Oxford Logic Guides No. 30. Oxford University Press, 1994.

[21] L. Schubert. Monotonic solution of the frame problem in the situation calculus: an efficient method for worlds with fully specified actions. In H. Kyburg, R. Loui, and G. Carlson, editors, *Knowledge Representation and Defeasible Reasoning*, pages 23–67, Boston, MA, 1990. Kluwer Academic Press.

A Proofs

Since the developments are complete histories of the evolution of the scenario we need to distinguish between the components that record the complete history and the slices of the components at specific timepoints. We define $R_{0:k}$ to be the history of the development from time 0 to time k and R_k to be the slice at time k. All other components of the development are defined analogously.

Lemma 6 *The ego/world interaction generates a collection of infinite trees. Moreover, if the number of fluents in the initial state is infinite, the collection is infinite, if the number of fluents* $\mid \mathcal{F} \mid$, *is finite, the collection is finite and of cardinality* $2^{|\mathcal{F}|}$.

Proof: Assume that we have already proven the structures generated are trees. It is easy to see that if we have a finite number of fluents, namely, $\mid \mathcal{F} \mid$ fluents, we will generate $2^{|\mathcal{F}|}$ different trees, one for each possible combination of truth values for the fluents. If the number of fluents is infinite, we will have an infinite collection of trees.

We now proceed to prove that the structures generated by the simulation are indeed, trees. The proof will be given by induction on the ego/world interaction. The nodes of the tree will be the developments that correspond to the initial state or those updated by the world, where the last component, $\mathcal{C}$, is empty. The arcs of the tree will be labeled with the development where the ego has recorded it's choice of action, that is, where the $\mathcal{C}$ component is not empty.

- *Base case* At time 0, the initial development is of the form $\langle \mathrm{B}_0, \varnothing, \mathrm{R}_0, \varnothing, \varnothing \rangle$, which can be identified with the root of the tree, since there is only one such node for each initial state because the state is completely specified.

- *Inductive case* At time t, we have a tree of depth p $t+1$, whose leaves are of the form $\langle \mathcal{B}_{0:t}, \varnothing, R_{0:t}, \mathcal{A}_{0:t}, \varnothing \rangle$. Any arc originating from such nodes will be labeled with $\langle \mathrm{B}_{0:t}, \varnothing, \mathrm{R}_{0:t}, \mathrm{A}_{0:t}, \langle t, \tau \rangle \rangle$. This arc will arrive in the development $\langle \mathrm{B}_{0:t} \cup \{t+1\}, \varnothing, \mathrm{R}_{0:t+1}, \mathrm{A}_{0:t} \cup \{\langle t, \tau, t+1 \rangle\}, \varnothing \rangle$, where R_{t+1} results from overriding the values of the fluents in R_t according to the world, W. There will be at least one arc for each τ defined in W. By Lemma 7 (given below), the arcs will arrive in different nodes. This argument is valid regardless of whether the action is legal or not, because if the action is legal, there exists only one arc departing from the node labeled with the action; if the action is illegal, there may be more than one, but the histories of the nodes they each arrive in are different. The only other possibility where the arcs may collapse is if the action is state independent ($Poss(\tau, s) \equiv TRUE$) and there is more than one arc with the same label. In this case, the histories will still be different because if there is more than one arc with the same label, the previous action must have been an illegal action that produced a different history in each node.

Here concludes the proof that the simulation generates a collection of trees, whose cardinality depends on the cardinality of the fluents of the scenario. ■

Lemma 7 *If there are two arcs with different labels departing from the same node, the nodes which they arrive at are different.*

Proof: Assume we have two arcs departing from the node $\langle \mathrm{B}_{0:t}, \varnothing, \mathrm{R}_{0:t}, \mathrm{A}_{0:t}, \varnothing \rangle$. Assume the first one is $\langle \mathrm{B}_{0:t}, \varnothing, \mathrm{R}_{0:t}, \mathrm{A}_{0:t}, \langle t, \tau_1 \rangle \rangle$ and the second one is $\langle \mathrm{B}_{0:t}, \varnothing, \mathrm{R}_{0:t}, \mathrm{A}_{0:t}, \langle t, \tau_2 \rangle \rangle$. The first case corresponds to the case where the ego chose to execute action τ_1, so the world records the interaction in the new node $\langle \mathrm{B}_{0:t+1}, \varnothing, \mathrm{R}'_{0:t+1}, \mathrm{A}_{0:t} \cup \{\langle t, \tau_1, t+1 \rangle\}, \varnothing \rangle$. The second case corresponds to the case where the ego chose to execute τ_2 so the new node will be $\langle \mathrm{B}_{0:t+1}, \varnothing, \mathrm{R}''_{0:t+1}, \mathrm{A}_{0:t} \cup \{\langle t, \tau_2, t+1 \rangle\}, \varnothing \rangle$. If these two were equal, we should have $\tau_1 = \tau_2$, but this would mean that the two arcs have the same label, which contradicts the hypothesis. Thus, two arcs with different labels arrive in different nodes. ■

Theorem 8 *Let Σ be a preliminary axiomatization that has the consistency property. Let Υ be the corresponding chronicle. $Mod(\Upsilon)$ at each time slice corresponds to $Mod(R(\Sigma))$ at each situation.*

Proof: The following proof is independent of whether Υ contains a schedule or not. If Υ contains a schedule, $Mod(\Upsilon)$ will be the set of models that results from filtering the set of models obtained without considering the schedule to obtain only those models consistent with the schedule. That is, if $Mod(\Upsilon)$ is the set of models obtained without the schedule, and the models of the schedule are $[\![\mathtt{sch}]\!]$, then, the set of models of the chronicle with the schedule is $Mod(\Upsilon) = Mod(\Upsilon \backslash \mathtt{sch}) \cap [\![\mathtt{sch}]\!]$.

We need to prove that for each situation there is a corresponding development, such that the $\langle M, R \rangle$ components of the development correspond to the interpretation of the state of the situation and the $\mathcal{A}$ component corresponds to the sequence of actions that give rise to that situation, and vice-versa. This will be done inductively. We denote a sequence of actions $A_0, A_1, \ldots, A_k$ by simply A_k.

Base Case: no action has been executed (the initial situation).

One direction: The set of models of the initial developments corresponds to the set of models of the initial situation.

The proof of this case rests on the following claim: there is a nonempty set of developments that correspond to the initial situation with cardinality equal to the power set of the number of fluents that are not specified in the initial situation. Recall that the initial state is not completely specified since Reiter's solution does not require a completely specified initial situation. For each fluent specified in the initial situation in Σ, there is a consistent translation into an observation in the chronicle Υ. According to Lemma 6 the ego/world interaction defines a collection of developments for each chronicle for each possible initial state. Since there is a development for every possible initial state, there are developments consistent with the initial observations. Filtering the intended models with the models of the observations at time 0 yields a collection of models where the first-order (FOL) interpretations of fluents are the same as the FOL interpretations of fluents in the initial situation. Moreover, the cardinality of the set of intended models $\langle \mathrm{M}^i, \mathrm{R}_0^i \rangle$ is two to the power set of the number of unspecified fluents in the initial situation since all unspecified values are allowed to vary.

Other direction: The set of models of the initial situation corresponds to the set of models of the initial developments.

For each interpretation of each fluent in the pair $\langle \mathrm{M}^i, \mathrm{R}_0^i \rangle$, there are two possibilities, either the fluent is specified in the initial observations or it is not. If it is, then it is specified in the initial situation and since the observations are consistent with the initial situation, the FOL interpretations of the fluent are the same. If the fluent is not specified in the observations, it is not specified in the initial situation and therefore its value is allowed to vary. Thus, there is a FOL interpretation for the initial situation which assigns the same value as the value assigned by this pair.

Inductive Case: execution of action, a after A_{k-1}.

There are two different cases depending on whether a is legal or not after the execution of A_{k-1}.

Case I: a is legal.

One direction: The set of models of the development up to time $k+1$ corresponds to the set of models of situation s_{k+1}.

Let $s_{k+1} = do(a, s_k)$ and let $\langle \mathcal{B}_{0:k}, \varnothing, R_{0:k}, \mathcal{A}_{0:k}, \varnothing \rangle$ be the development corresponding to s_k according to the inductive hypothesis. It suffices to show that the values of fluents in s_{k+1} are the same as the values of fluents of the history components of the developments up to time $k+1$. Assume that fluent f is true in s_{k+1}. According to Reiter's non-monotonic entailment criterion, there are two possibilities, either f was not true in s_k and a changed its value or f was true in s_k and a does not change the value of f.[9]

[9]Recall that the models selected by Reiter's entailment criterion are exactly those determined by the successor state axiom which has the form of equation (4) in Section 2.1.

a) a changes the value of f to true. In this case, according to the way the successor state axiom was obtained, there exists a positive effect axiom in the preliminary axiomatization for the pair (a, f). Thus, by construction of W of the underlying semantics, $f \in \texttt{Infl}(a, R_k)$. Moreover, f is assigned the value true in $\texttt{Rstat}(a, R_k)$. Thus, in every development that is generated by the choice of the ego to execute a, the history component of the development assigns the value true to f.

b) a does not change the value of f and f was true in the previous situation. In this case, there does not exist a negative effect axiom in the preliminary axiomatization for the pair (a, f). Thus, by construction of W according to the underlying semantics, $f \notin \texttt{Infl}(a, R_k)$. Moreover, the value of f is propagated from the value at timepoint k. Since f was true at timepoint k because of the inductive hypothesis, f continues to be true at timepoint $k+1$ in every development consistent with situation s_{k+1}.

Since fluent f only true when either a) or b) holds, the correspondence also holds when f is false.

Other direction: The set of models of situation s_k corresponds to the set of models of the development up to time k.

Let $\langle \mathrm{B}_{0:k+1}, \varnothing, \mathrm{R}_{0:k+1}, \mathrm{A}_{0:k+1}, \varnothing \rangle$ be a development up to time $k+1$ after the execution of a. For each fluent f in slice R_{k+1} we have two possibilities, either f is true or false at timepoint $k+1$. If f is true, then we again have two possibilities, either $f \in \texttt{Infl}(a, R_k)$ of $f \notin \texttt{Infl}(a, R_k)$ and f was true at timepoint k.

a) a changes the value of f to true. Then, there exists a positive effect axiom in Σ for the pair (a, f), and the successor state axiom for f determines that f is true in situation s_{k+1}.

b) a does not change the value of f and f was true at the previous timepoint. Then, there does not exist a negative effect axiom in Σ for the pair (a, f), and the successor state axiom for f determines that since f was true in s_k (by inductive hypothesis) then f continues to be true in s_{k+1}.

The case where f is false is proven analogous.

This concludes the proof for the legal case.

Case II: a is illegal.

One direction: The set of models of the development up to time k corresponds to the set of models of situation s_k.

Let $s_{k+1} = do(a, s_k)$ and let $\langle \mathrm{B}_{0:k}, \varnothing, \mathrm{R}_{0:k}, \mathrm{A}_{0:k}, \varnothing \rangle$ be the development corresponding to s_k according to the inductive hypothesis. We must prove that the values of fluents in s_{k+1} are the same as the values of fluents of the history components of the developments up to time $k+1$. Recall (Section 2.2) that in the illegal case, Reiter's non-monotonic entailment criterion selected all possible models for the "successor situation". Thus, it suffices to prove the existence of some development corresponding to each of these interpretations. But, this result follows directly from the construction of the developments in this case.

Other direction: The set of models of situation s_k corresponds to the set of models of the development up to time k.

Recall that in the illegal case, the developments branched out once for each possible combination of changes (Section 4.3.1) . Thus, we need to prove the existence of some interpretation, according to Reiter's non-monotonic entailment criterion, for the "successor situation". Since we are comparing FOL interpretations for the fluents, and fluents can only take values true or false and each "successor situation" and development will contain a different combination of fluent values by construction of these structures, it suffices to observe that the cardinality of the set of models for slice k of the developments equals the cardinality of the interpretations for the "successor situation" because the set of fluents is the same since the developments are generated from the chronicle which is constructed from Σ.

This concludes the proof of the theorem.

$\blacksquare$

Corollary 9 *(a) For any sentence φ,*

$$\Sigma \mathrel{|\!\approx_R} \varphi \quad \textit{iff} \quad \Sigma \mathrel{|\!\approx_S} \varphi$$

(b) If T is specified in addition to Σ,

$$\Sigma \mathrel{|\!\approx_R} \varphi(T) \quad \textit{iff} \quad \Sigma \mathrel{|\!\approx_S} \varphi(T).$$

Proof: It follows directly from the theorem. $\blacksquare$