Linköping Electronic Articles in
Computer and Information Science
Vol. 1(1996): nr 12

Reasoning about actions and change with ramification

Erik Sandewall

Department of Computer and Information Science
Linkping University
Linkping, Sweden

Linköping University Electronic Press
Linköping, Sweden

http://www.ep.liu.se/ea/cis/1996/012/

Published on October 1, 1996 by
Linköping University Electronic Press
581 83 Linköping, Sweden

Linköping Electronic Articles in
Computer and Information Science
ISSN 1401-9841
Series editor: Erik Sandewall

Typeset by the author using LaTeX
Formatted using étendu *style*

Recommended citation:
<Author>. <Title>. Linköping Electronic Articles in Computer and Information Science, Vol. 1(1996): nr 12. http://www.ep.liu.se/ea/cis/1996/012/. *October 1, 1996.*

This URL will also contain a link to the author's home page.

Abstract

The paper analyzes and compares methods for ramification, or for update, which are based on minimization of change while respecting static domain constraints. The main results are:

- An underlying semantics for propagation oriented ramification, called *causal propagation semantics*
- *Comparative assessments* of a number of previously proposed entailment methods for ramification, in particular, methods based on minimization of change after partitioning of the state variables (fluents). This includes the methods previously proposed by del Val and Shoham, by Kartha and Lifschitz, and by our group.
- Assessment of two *ranges of sound applicability* for one of those entailment methods, MSCC, relative to causal propagation semantics.
- Identification of an essential feature for minimization-based approaches to ramification, namely *changeset-partitioning.*

This paper has been accepted to the KR'96 conference, and is concurrently published in its conference proceedings:
L. Aiello, J. Doyle, and S. Shapiro (eds). Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference. Published by Morgan-Kaufmann Publishers, San Francisco, California, USA.

1 The Systematic Approach to Actions and Change

The class of problems being addressed here will first be defined in abstract terms, and will then be discussed from a more applied point of view.

1.1 Logics of Actions and Change

A *chronicle* is a description, using logical formulae, of a number of actions, their general effects, and observations of some properties of the world while the actions are being performed. In simple cases, a chronicle consists of the following kinds of formulae:

action statements, for example $D(20, 25, \mathit{make\text{-}coffe})$, saying that the current agent performs the act of making coffee from time 20 to time 25;

observations, for example $H(25, \mathit{coffee\text{-}hot})$, saying that there is hot coffee at time 25;

action laws, for example $D(s, t, \mathit{make\text{-}coffee}) \Rrightarrow H(t, \mathit{coffee\text{-}hot})$, saying that if coffee is made from time s to time t, then hot coffee exists at time t;

timing statements, for example $\boldsymbol{t}_2 < \boldsymbol{t}_4 - 1$, relating the values of two expressions using temporal constants. Such statements are useful when the timepoint arguments of D and H are constant symbols and not explicit timepoints (for example, $\boldsymbol{t}_2$ rather than 25).

The letter D stands for *Do* and H for *Holds.* To accommodate multi-valued properties, we consider the general form of H to be $H(t, f, v)$, saying that the "feature" f has the value v at time t. Then $H(t, p)$ is an abbreviation for $H(t, p, \mathtt{T})$ for the case where p can only have the value $\mathtt{T}$ or $\mathtt{F}$[1]. The terms *feature* and *property* will be used synonymously here.

In more complex cases, chronicles may contain observations that refer to multiple timepoints, action statements that refer to conditional or composite actions, and other ontologies for time such as branching time. Also, action symbols and features/properties may have objects as arguments, for example

$$H(25, \mathit{stateof}(\mathit{tralite4}), \mathtt{Green}),$$

expressing that the state of traffic light number 4 is $\mathtt{Green}$ at time 25. Logics whereby one can express and reason about chronicles are called *logics of actions and change.*

[1] In earlier publications I wrote $D(s, t, E)$ as $[s, t]\,E$, and $H(t, f, v)$ as $[t]\,f \mathrel{\hat{=}} v$.

1.2 Research Problem

The notation for chronicles that is suggested by these examples can of course be understood as the language of a first-order logic. The central research problem for logics of actions and change is that *the set of intended models for a chronicle does not agree with the set of Tarskian models as provided by first-order logic*, henceforth referred to as *classical models.* The main research task is therefore to find and to analyze suitable complements to the machinery of classical logic whereby intended and actual model sets agree, so that one obtains all the intended conclusions and only them.

The research problem and the research task can of course be equivalently formulated in terms of the "set of conclusions" instead of the "set of models". Possible conclusions from a given chronicle include derived facts (syntactically similar to observations) about the state of the world at other times or in other properties than those that were directly observed.

Different variants of this research problem are obtained by varying the assumptions that lead to the mismatch between intended and classical models. The basic assumption is *inertia* or "no change without a known reason". In the simplest case it means that actions only change those properties of the world which are indicated in the action law, whereas other properties remain constant. This simple case is referred to as *strict inertia.* Less trivial cases are obtained e.g. if certain static constraints are known for the state of the world, and if it is then necessary to assume that actions cause additional changes, besides those indicated by the action law, in order that the situation resulting from a particular action shall satisfy the static constraints. This case is referred to as *ramification*, and it will be addressed towards the end of the present article.

Non-deterministic actions pose particular problems in the context of the inertia assumption. One way of dealing with them is to introduce an additional *occlusion predicate* $X(f,t)$, expressing that the value of the feature f may change during the timestep $[t-1,t]$.

1.3 Applications

The research on logics for actions and change has its origin in artificial intelligence, where it is historically known under the name of the "frame problem" [4]. Two kinds of uses have been foreseen in A.I. One is for *common-sense reasoning* based, in particular, on information given in natural language. Statements about actions and about the state of the world at different times are commonplace in natural-language text, and understanding them is one essential aspect of understanding the text as a whole.

The other usage is for *cognitive robotics*, that is, for allowing an intelligent robot to reason about its own actions and its effects, as well as about the effects of natural events and of actions taken by other robots in the same habitat. In this case, the actions are implemented as programs in the robot controller system, and reasoning about the actions is a particular kind of reasoning about programs. It is particular in the sense that one reasons not only about the program itself, but also about spontaneous events in the world where the robot (and indirectly, the program) is set to operate.

The notation for chronicles that was exemplified above allows one to characterize the duration of an action in terms of metric time. The consequent side of an action law is not restricted to specifying the final result of an action; it may also characterize the duration of the action (as a constraint on the combination of s and t), and the state of the world at intermediate points during the execution of the action. Logics of actions and change are therefore potential candidates for reasoning about *hard real-time* programs. However, that aspect has not been investigated very much yet, and will not be addressed in the present article.

Typically, one wishes action laws to be as compact as possible, so ideally they should only have to specify some of the changes that are due to an action, leaving other resulting changes to be inferred by logical means. This topic is closely related to the *update problem* in databases.

1.4 Use of Preferential Nonmonotonic Logics

The standard approach in this research is to use *nonmonotonic logics*. The following notation will be used. Let Υ be a chronicle, represented as a set of logical formulae of the kinds exemplified above. Then $[\![\Upsilon]\!]$ will refer to the set of classical models for Υ, and $Mod(\Upsilon)$ will refer to the set of intended models. A *unary selection function* is a function from a set of models to a subset of that set. A selection function S is *correct* for a chronicle Υ iff $Mod(\Upsilon) = S[\![\Upsilon]\!]$. Selection functions are therefore the primary instrument for reconciling intended and classical model sets.

A *preference relation* is a strict partial order on models. Many simple cases of selection functions can be defined directly in terms of a preference relation $\ll$, so that the set of selected models is

$$Min(\ll, [\![\Upsilon]\!]),$$

meaning the set of $\ll$-minimal classical models. More sophisticated selection functions can be obtained in the following fashions:

Binary selection functions: the premises are divided into two *partitions*, so that the chronicle Υ is in fact a pair $\langle\Upsilon_1, \Upsilon_2\rangle$. The set of classical models is defined by

$$[\![\langle\Upsilon_1, \Upsilon_2\rangle]\!] = [\![\Upsilon_1]\!] \cap [\![\Upsilon_2]\!],$$

and the set of selected models is defined as

$$S([\![\Upsilon_1]\!], [\![\Upsilon_2]\!]),$$

where S is a function which maps two sets of models to a subset of their intersection.

Reformulation: the set of selected models is defined as

$$[\![G(\Upsilon)]\!],$$

where G is a *reformulation function*, that is, a syntactic transformation on the set Υ of formulae which satisfies $G(\Upsilon) \supseteq \Upsilon$. The goal of reducing the set of models is achieved by minimization in the case of filtering, and by adding more formulae to the originally given set, in the case of reformulation.

Combinations and generalizations of the methods described so far, that is, minimization, k-ary selection functions, and reformulation. In particular, it is possible and often useful to let a chronicle consist of more than two partitions, corresponding to the use of a k-ary selection function.

One particular use of binary selection functions is for *filtering*, where the set of selected models for $\langle \Upsilon_1, \Upsilon_2 \rangle$ is

$$Min(\ll, [\![\Upsilon_1]\!]) \cap [\![\Upsilon_2]\!].$$

1.5 Entailment Methods and Range of Applicability

In earlier publications [6, 5, **?**], I have argued that the appropriate question with respect to a method based on a non-monotonic logic is not "is it correct", but "what is its range of applicability". In other words, one identifies the set of conditions that must be satisfied by a chronicle Υ in order that the set of selected models according to the given entailment method shall agree with the set of intended models for the same chronicle. I refer to this as the *systematic approach* to the study of applied non-monotonic logics.

The reasons for analyzing the range of applicability should be fairly evident. One is not always interested in finding and using the method with the largest range of applicability: it is to be expected that there are some methods with a fairly narrow range of applicability, but which allow more efficient implementation. If such a method exists, and its range of applicability does include those chronicles that may arise in a given application, then one will be happy to use that method for the application at hand.

The range of applicability may be significantly improved through the use of the following concepts. A *pretransformation* $\mathbf{G}$ is a mapping from chronicles to chronicles such that $Mod(\Upsilon) = Mod(\mathbf{G}(\Upsilon))$. An *entailment method* is a pair $\langle \mathbf{G}, S \rangle$ where $\mathbf{G}$ is a pretransformation and S is a selection function whose arity corresponds to the structure of the chronicle. If S is unary, then an entailment method $\langle \mathbf{G}, S \rangle$ assigns the set $S[\![\mathbf{G}(\Upsilon)]\!]$ of models to each chronicle Υ. It is *correct* for the chronicle Υ iff

$$S[\![\mathbf{G}(\Upsilon)]\!] = Mod(\Upsilon).$$

Similarly, if S is a binary selection function, and $\mathbf{G}(\langle \Upsilon_1, \Upsilon_2 \rangle) = \langle \Upsilon_1', \Upsilon_2' \rangle$, then $\langle \mathbf{G}, S \rangle$ is correct for $\langle \Upsilon_1, \Upsilon_2 \rangle$ iff

$$S(\langle [\![\Upsilon_1']\!], [\![\Upsilon_2']\!] \rangle) = Mod(\langle \Upsilon_1, \Upsilon_2 \rangle).$$

Pretransformations may be used, for example, for adding timing statements which exclude concurrent execution of actions, if the intended models only allow sequentially executed actions.

In order to identify the range of applicability for an entailment method, one needs the following formal tools:

an underlying semantics, providing a formal definition of the set of intended models corresponding to each chronicle;

a taxonomy of chronicles, providing a "coordinate system" whereby one can characterize the range of applicability.

In the earlier publications I have defined the underlying semantics and the taxonomy for the case of strict inertia, together with precise formulations of the syntax and of the classical semantics, which are of course also needed.

The *purpose of the present paper* is to give an introduction to the systematic approach. I shall first *summarize* some of the earlier results, and describe the *conceptual framework* within which they have been developed. After that, I present some *new results* concerning generalizations to the case of inertia with ramification.

2 The Precise Structure of Chronicles

Ordinary logic analyzes what conclusions can be drawn from an unstructured set of premises (axioms). Since logics of actions and change tend to need partitioned premises, we introduce a general framework for accomodating such partitioning. Formally, a *chronicle in $\mathcal{K}$* is a tuple

$$\langle \mathcal{K}, \mathtt{Diw}, \textsc{scd}, \textsc{obs} \rangle$$

where $\mathcal{K}$ is a code indicating the epistemological assumption of *complete knowledge about the actions* that accompanies the chronicle, `Diw` is a *world description* containing e.g. action laws, SCD is a set of formulae for the *schedule* characterizing the actions and their temporal relationships using action statements and timing statements, and OBS is a set of *observations* as introduced above.

The world-description component of the scenario is again a tuple whose first element characterizes the type of world. For the case of strict inertia, the world description will be assumed to be a twotuple $\langle \mathbf{IA}, \textsc{law} \rangle$, were again **IA** is a code indicating the ontological assumption of strict inertia, and LAW is the set of action laws. A more general class of world descriptions is formed as threetuples $\langle \mathbf{IAD}, \textsc{law}, \textsc{dcs} \rangle$, where **IAD** is a code for inertia with ramification, LAW is as before, and DCS is a set of *static domain constraints* of the general form

$$\forall t[\alpha],$$

where α is an expression formed only using the H predicate, and with t as the first argument. The following is an example of a static domain constraint:

$$\forall t[H(t, f_4, \mathtt{F}) \Rightarrow H(t, f_6, \mathtt{G})],$$

where of course `F` must be a possible value for the feature f_4, and `G` must be a possible value for the feature f_6. This domain constraint says that in any state where f_4 has the value `F`, f_6 must have the value `G`.

The epistemological code $\mathcal{K}$ and the ontological codes **IA** and **IAD** belong to a broader reportoire of codes for characterizing epistemological and ontological assumptions. Each choice of assumptions and codes is associated with its own definition of intended models and, therefore, its own analysis of range of applicability for entailment methods.

The standard use of domain constraints is for ramification: the action law is supposed to specify the "essential" changes that are due to the action, and other changes are obtained implicitly as the (or a) minimal set(s) of additional changes that must be assumed in order to avoid getting into a state that violates some of the domain constraints.

The special case where DCS is the empty set reobtains strict inertia: since there are no domain constraints, there is never any reason to force additional value changes, so the only changes of feature values are those that are expressly stated in the action laws for expressly known actions. Therefore, a world description $\langle \mathbf{IA}, \text{LAW} \rangle$ is equivalent to one of the form $\langle \mathbf{IAD}, \text{LAW}, ? \rangle$.

For simplicity, we address only the "propositional" case where feature symbols f and action symbols E appear without arguments. In many applications it is necessary to reason about multiple objects with similar properties, for example several cars that move in the same environment, as well as actions that take one or more objects as arguments. However, this case is very often only "quasi-propositional" in the sense that each chronicle binds the object domain to a specific, finite size. The action laws need to be specified for arbitrary object domains, since they should be reusable in several chronicles, but reasoning about a particular chronicle is still essentially propositional. In [5] we have worked out the relatively straightforward details of this generalization.

From a formal point of view, action laws will be understood as definitions of abbreviations, and not as first-class logical formulae. The action law in the introductory example,

$$D(s, t, \textit{make-coffee}) \Rrightarrow H(t, \textit{coffee-hot}),$$

simply states that e.g. the action statement

$$D(20, 25, \textit{make-coffee})$$

is an abbreviation for

$$H(25, \textit{coffee-hot}).$$

The purpose of dealing with action statements and action laws in this way is to avoid the need of saying, in formulae, that there are no other actions besides those stated in SCD. If LAW is a set of action laws in this sense, and SCD is a set of action and timing statements, then LAW[SCD] will represent the result of expanding all formulae in SCD using the definitions in LAW. From the point of view of a selection function, a $\mathcal{K}$-**IA** chronicle therefore contains two premise sets, LAW[SCD] and OBS. Its classical model set is $[\![\text{LAW[SCD]}]\!] \cap [\![\text{OBS}]\!]$. An entailment method $\langle \mathbf{G}, S \rangle$ for $\mathcal{K}$-**IA** chronicles assigns

$$S([\![\text{LAW}'[\text{SCD}']]\!], [\![\text{OBS}']\!])$$

as the set of selected models for a given chronicle

$$\langle \mathcal{K}, \langle \mathbf{IA}, \text{LAW} \rangle, \text{SCD}, \text{OBS} \rangle,$$

provided that

$$\langle \mathcal{K}, \langle \mathbf{IA}, \text{LAW}' \rangle, \text{SCD}', \text{OBS}' \rangle = \mathbf{G}(\langle \mathcal{K}, \langle \mathbf{IA}, \text{LAW} \rangle, \text{SCD}, \text{OBS} \rangle).$$

Similarly, an entailment method $\langle \mathbf{G}, S \rangle$ for $\mathcal{K}$-**IAD** chronicles requires S to be ternary, and assigns

$$S([\![\text{LAW}'[\text{SCD}']]\!], [\![\text{OBS}']\!], [\![\text{DCS}']\!])$$

as the set of selected models for a given chronicle

$$\langle \mathcal{K}, \langle \mathbf{IAD}, \text{LAW}, \text{DCS} \rangle, \text{SCD}, \text{OBS} \rangle,$$

provided that

$$\langle \mathcal{K}, \langle \mathbf{IAD}, \text{LAW}', \text{DCS}' \rangle, \text{SCD}', \text{OBS}' \rangle = \mathbf{G}(\langle \mathcal{K}, \langle \mathbf{IA}, \text{LAW}, \text{DCS} \rangle, \text{SCD}, \text{OBS} \rangle).$$

3 Trajectory Semantics

The *underlying semantics* for a class of chronicles defines the set of *intended models* for chronicles in that class. We shall first define it for $\mathcal{K}$-**IA** chronicles, and later generalize it to $\mathcal{K}$-**IAD** chronicles.

3.1 Simple IDS Worlds

A similarity type σ is a mapping from feature symbols to the set of possible values for the feature. The set of feature symbols f that are defined by a given σ will be written $\mathcal{F}_\sigma$. A *state* for a given similarity type σ is a mapping which to each feature symbol f assigns a corresponding value i.e. a member of $\sigma[f]$. The set $\mathcal{R}_\sigma$ for given σ is the set of possible states for σ.

Let also $\mathcal{E}$ be a set of *action symbols*, like *make-coffee* in the trivial example. A *complete trajectory* v is a finite non-empty sequence of states in $\mathcal{R}_\sigma$. A *simple IDS world* is a mapping W from $\mathcal{E} \times \mathcal{R}_\sigma$ to non-empty sets of complete trajectories.

A *finite history* from time 0 to n (now) is a mapping from integers in the interval $[0, n]$ to complete states.

A *simple IDS game* is played between an *ego* and a simple IDS world, in the following fashion. At each step in the game, there is an integer *current time* n and a *current history* h which is a mapping from the closed interval $[0, n]$ to $\mathcal{R}_\sigma$. Initially, the current time is 0, the current history assigns some initial state to 0, and the ego has the move. Each time the ego has the move, it may choose to invoke an action, i.e., a member of $\mathcal{E}$, or to do nothing. If the ego invokes an action E at time n with the current history h, the world's move is to select some trajectory $v \in W(E, h(n))$, that is, one of the trajectories that the world admits for the present state of the world, $h(n)$, and the currently chosen action E. The world then increases n by the length of v, and concatenates v at the end of h. If the ego chooses not to make any move, then the world defines $h(n+1)$ as $h(n)$ and increases n by 1. In either case the move is then returned to the ego.

This definition of a world and of the ego-world game is a specialization, with a trivial reformulation, of how those concepts were defined in [5]. Note that the definition of simple IDS worlds is a straight-forward variant and generalization of finite-state automata, with the following specific properties:

- Structured states, since a state is assumed to be a vector of "state variables".
- Non-determinism: several "moves" are possible for the world.
- Metric time: the result of an action is a sequence of states rather than a single state.

The game is open to choice both for the ego (which has a choice of actions at each step) and for the world (which may have more than one trajectory for a given action and starting state). Different choices result in different games. Each game results in an (infinite) *history* H, that is, a mapping from the non-negative integers to states in $\mathcal{R}_\sigma$. An *interpretation* is a pair $\langle M, H \rangle$, where M assigns values to temporal constant symbols, and H is a history. The truth of an observation or a timing statement in an interpretation is defined in the obvious way. An action statement $D(s, t, E)$ is said to *be satisfied* in an interpretation $\langle M, H \rangle$ iff there is some game that results in the history H, and the action E was invoked at time $M[s]$ in that game, and the world then selected a trajectory of length $M[t] - M[s]$. A set of action statements is said to be satisfied in an interpretation $\langle M, H \rangle$ iff there is some game, resulting in H, such that all the given action statements and no others are satisfied for $\langle M, H \rangle$ in that game.

The set of *intended models* $Mod(\Upsilon)$ for a given **IA** chronicle

$$\Upsilon = \langle \mathcal{K}, \langle \mathbf{IA}, \textsc{law} \rangle, \textsc{scd}, \textsc{obs} \rangle$$

is defined as the set of those interpretations $\langle M, H \rangle$

- which are obtained from a game between a world W which is correctly characterized by LAW, and an arbitrary ego, from an arbitrary starting state; and
- which satisfy all the observations in OBS and all the timing statements in SCD; and
- where the set of action statements in SCD is satisfied.

A chronicle contains a world description based on action laws LAW, but the definition of intended models for such a chronicle uses the ego-world game, where the world is defined as a mapping W assigning sets of trajectories to action symbols. It remains to relate the action-law characterization LAW and the trajectory characterization W of a world. The important difference, of course, is that the action law only specifies those features that *change* in the course of the action, or even a subset of them, whereas the trajectory is constructed out of complete states and therefore specifies the new values for all the features in $\mathcal{F}$.

3.2 World Descriptions for Strict Inertia

In order to define how action laws LAW correctly characterize a world W, we shall complement *logical world descriptions* which consist of logic formulae, with *set-theoretic world descriptions.* In either case, a world description is a more compact version of the original trajectory characterization. There are straight-forward ways of relating worlds (W) and set-theoretic world descriptions, and the latter are in turn related (by a simple syntactic transformation) to action laws in the logical world description.

We first consider world descriptions of type **IA**, that is, world descriptions which are based on an assumption of strict inertia. First some auxiliary definitions. If $F \subseteq \mathcal{F}_\sigma$, then a *partial trajectory* for F is a finite nonempty sequence of partial states each of which is defined exactly for the features in F. A *world description* for **IA**, then, is a tuple

$$\langle \mathbf{IA}, \mathit{Infl}, \mathit{Rtrajs} \rangle,$$

where **IA** is a code indicating the type of world description, *Infl* ("influenced features") is a mapping which for each combination of action symbol and starting state returns a subset of $\mathcal{F}$, and *Rtrajs* ("reduced trajectories") is a mapping from action symbols times states such that $Rtrajs(E, r)$ is a non-empty set of partial trajectories for $Infl(E, r)$.

Let $r \ddagger F$ be the partial state obtained from the total state r and the set F of features by restricting r so that it is only defined for members of F. Also, let $v \ddagger F$ be the partial trajectory obtained from the total trajectory v by restricting each element to F. For a given, simple IDS world W and a given function *Infl*, the *corresponding* **IA** world description $\langle \mathbf{IA}, Infl, Rtrajs \rangle$ is obtained by choosing

$$Rtrajs(E, r) = \{v \ddagger Infl(E, r) \mid v \in W(E, r)\},$$

that is, by restricting each trajectory in $W(E, r)$ to those features indicated in $Infl(E, r)$.

Let $\oplus$ be the override operator, so that if r is a total state and r' is a partial state, then $r \oplus r'$ is the total state assigning the same value as r' whenever defined, and otherwise the value assigned by r. If r is a total state and v' is a partial trajectory, then $r \oplus v'$ is the total trajectory defined similarly.

Now, if W is a simple IDS world, and the world description $\langle \mathbf{IA}, Infl, Rtrajs \rangle$ is formed from W, then this world description *conserves* W iff

$$W(E, r) = \{r \oplus v' \mid v' \in Rtrajs(E, r)\},$$

that is, if one can take the partial trajectories in $Rtrajs(E, r)$, extend every partial state in each of those trajectories by assigning, for each feature that does not have a value, the value that it has in r, and if one then gets back the original world W.

Evidently, for a given W, if *Infl* is chosen so that it at least contains all members f of $\mathcal{F}_\sigma$ which change their value in some trajectory in $W(E, r)$, then the corresponding world description $\langle \mathbf{IA}, Infl, Rtrajs \rangle$ conserves W.

Let a world W be given, and choose a corresponding world description $\langle \mathbf{IA}, Infl, Rtrajs \rangle$ which conserves W. In [5] we defined, in syntactic terms, the transformation from such a world description to an equivalent world description of the form $\langle \mathbf{IA}, \text{LAW} \rangle$, where LAW is a set of action laws as exemplified in section 1 here. We say that $\langle \mathbf{IA}, \text{LAW} \rangle$ *correctly characterizes* a world W iff it is formed, in the way defined in [5], from a world description that conserves W.

3.3 Summary of transformations

The successive shifts of representation that were described above can be summarized as follows.

1. Start with a world $W(E, r)$.

2. Select a suitable *Infl* function for W.

3. Construct *Rtrajs* from W and *Infl*, and obtain the set-theoretic world description $\langle \mathbf{IA}, \mathit{Infl}, \mathit{Rtrajs} \rangle$. Require that it conserves W.

4. Convert it to the corresponding logical world description $\langle \mathbf{IA}, \text{LAW} \rangle$, where the action laws in LAW are considered as abbreviations.

5. Use the logical world description (repeatedly) in chronicles of the form
$$\Upsilon = \langle \mathcal{K}, \langle \mathbf{IA}, \text{LAW} \rangle, \text{SCD}, \text{OBS} \rangle.$$

6. The set of intended models $\mathit{Mod}(\Upsilon)$ for such a chronicle is defined using W, SCD, and OBS.

7. The set of selected models for the chronicle according to an entailment method $\langle \mathbf{G}, S \rangle$ is defined as $S([\![\text{LAW}'[\text{SCD}']]\!], [\![\text{OBS}']\!])$, where $\mathbf{G}$ transforms SCD to SCD$'$, and similarly for the others. The particular $\mathbf{G}_{\mathsf{seq}}$ used below adds some more formulae to SCD, and leaves the other components unchanged.

This definition is usable even if Υ is given, so that one starts in step 5, since the previous steps are reversible. A given LAW identifies a corresponding W uniquely.

4 Simple Approaches to Strict Inertia

4.1 Preference-Order Compatibility

For the characterization of range of applicability, we need the following concepts. If r is a state and v is a trajectory of length j, then $r;v$ is the trajectory of length $j+1$ obtained by adding r as the first element in v. (In other words, the semicolon is used as a "cons" operator).

Let $\ll_m$ be a preference relation on interpretations $\langle M, H \rangle$ which is well-defined even if H is a finite history, and which is always false when comparing histories of unequal length. We wish to use $\ll_m$ for comparing members of $\mathit{Rtrajs}(E, r)$, but there is a technical problem since $\ll_m$ compares interpretations and not trajectories.

Technically, therefore, M_0 is defined as a placeholder valuation which assigns the value 0 to every temporal constant symbol. The *corresponding model function* for a world W is defined by

$$\mathit{Cmf}(E, r) = \{\langle M_0, (r; v) \rangle \mid v \in W(E, r)\}.$$

More general versions of *Cmf* will appear later on, so this is just the simplest case for the purpose of introduction. It works as follows: for given E and r we construct one interpretation for each trajectory in

$W(E,r)$, but we prefix it with r, and add a harmless M component. Then, a simple IDS world W is said to be *m-compatible* iff

$$I \not\ll_m I'$$

whenever I and I' are two members of $Cmf(E,r)$ for some E and r. It follows at once that strictly deterministic worlds, defined as those where $W(E,r)$ is a singleton for each combination of E and r, are m-compatible for arbitrary $\ll_m$.

4.2 Prototypical Chronological Minimization, PCM

The method of prototypical chronological minimization (PCM) is a slightly corrected variant of an entailment method that was first proposed by Kautz [2] and Shoham [8]. The preference relation $\ll_{\text{pcm}}$ is defined so that $\langle M,H\rangle \ll_{\text{pcm}} \langle M',H'\rangle$ iff $M = M'$ and there is some $\mathtt{t}$ such that

1. $H(t) = H'(t)$ for all $t < \mathtt{t}$;
2. $\mathtt{Diff}(H(\mathtt{t}-1), H(\mathtt{t})) \subset \mathtt{Diff}(H'(\mathtt{t}-1), H'(\mathtt{t}))$;

where

$$\mathtt{Diff}(r,r') = \{f \mid r[f] \neq r'[f]\},$$

that is, the set of features where r and r' differ. The entailment method PCM for $\mathcal{K}$-**IA** is defined as $\langle \mathbf{G}_{\text{seq}}, S_{\text{pcm}}\rangle$, where $\mathbf{G}_{\text{seq}}$ adds sequentiality premises to the schedule part of the given chronicle, and

$$S_{\text{pcm}}(A,B) = Min(\ll_{\text{pcm}}, A \cap B).$$

Assessment: it was proved in [6] that PCM is correct for all chronicles Υ satisfying the following three assumptions:

- The world represented in the LAW part of Υ must be PCM-compatible.
- The world represented in the LAW part must also be *equidurational*, meaning that the set of possible durations of an action must be independent of the starting state.
- The OBS partition may not contain any observations for times other than 0.

We notice that if $\langle \mathbf{IA}, \mathit{Infl}, \mathit{Rtrajs}\rangle$ is formed from and conserves W, then W is PCM-compatible iff *Rtrajs* is. The first and second requirements are *ontological* ones, in the sense that they are requirements on the world that is involved in the scenario. It was proved in [5] that these two requirements are not only sufficient but also necessary, in the sense that for any world W that does not satisfy both of these requirements, there exists some schedule Υ using W for which PCM does not give the intended set of models.

The third requirement is *epistemological*, in the sense that it specifies restrictions on our knowledge about the world rather than on the world itself. It is sufficient but not necessary.

All three requirements are undesirable from the point of view of practical usage. The third requirement can be removed using the technique of *filtering*. This is the topic of the next subsection.

4.3 Prototypical Chronological Minimization with Filtering, PCMF

The entailment method PCMF for $\mathcal{K}$-**IA** is defined as $\langle \mathbf{G}_{\text{seq}}, S_{\text{pcmf}} \rangle$, where $\mathbf{G}_{\text{seq}}$ is as before, and

$$S_{\text{pcmf}}(A, B) = Min(\ll_{\text{pcm}}, A) \cap B.$$

It was proved in [6] that PCMF is correct for all chronicles Υ satisfying the first two assumptions for the correctness of PCM, that is,

- The world represented in the LAW part of Υ must be PCM-compatible.
- The world represented in the LAW part must also be *equidurational*, meaning that the set of possible durations of an action must be independent of the starting state.

It was also shown that these requirements are necessary as well as sufficient. The removal of the third requirement of PCM means that PCMF is also applicable for *postdiction* and not only for prediction. Postdiction arises e.g. in diagnostic reasoning.

5 Strict Inertia with Occlusion

5.1 Occlusion

The restriction of PCM-compatibility is harmless for deterministic actions, but it excludes many cases of nondeterministic actions. Several methods have been proposed for lifting this restriction. One such technique is to use *occlusion*, which means effectively that the *Infl* component of the set-theoretic world description is represented explicitly in the models, so that it is also accessible for preference relations and selection functions. The following is an outline of how the framework of the previous section has to be revised for this purpose.

The definition of an interpretation is revised so that it is a triple $\langle M, H, X \rangle$, where the first two components are like before, and the third component is a mapping from features times timepoints to truth-values. The formulation of the action laws for a given world description $\langle \mathbf{IA}, \mathit{Infl}, \mathit{Rtrajs} \rangle$ is revised so that if an action statement $D(s, t, E)$ applies in an interpretation $\langle M, H, X \rangle$, and the feature f is a member of $\mathit{Infl}(E, H(s))$, then $X(f, u)$ is forced to be true for every u in the interval $[s+1,t]$. The intention is that if $X(f,u)$ is true then one makes an exception from the principle of inertia so as to allow a change in the value of f between time $u-1$ and time u.

This extension of the structure of classical models makes it possible for preference relations and selection functions to use their X component. At the same time, there is a structural mismatch since intended models still have the simpler structure $\langle M, H\rangle$. We therefore introduce the coercion function Ci such that

$$\mathsf{Ci}(W) = \{\langle M, H\rangle \mid \exists X[\langle M, H, X\rangle \in W]\}.$$

An entailment method $\langle \mathbf{G}, S\rangle$ is then defined to be correct for a chronicle Υ of the usual $\mathcal{K}$-**IA** structure iff

$$\mathsf{Ci}(S([\![\text{LAW}[\text{SCD}]]\!], [\![\text{OBS}]\!])) = \mathit{Mod}(\Upsilon).$$

The remaining technical details of this machinery can be found in [5]. Notice that action laws can only force $X(f,u)$ to be true at certain points; they can not force it to be false at all other points. Therefore, one needs some other device, such as minimization, in order to avoid X being true at unintended points.

5.2 Chronological Minimization of Occlusion and Change, CMOC

The entailment method CMOC for $\mathcal{K}$-**IA** (chronological minimization of occlusion and change) makes use of occlusion in order to eliminate the restrictions limiting nondeterminism. The preference relation $\ll_{\text{cmoc}}$ is defined so that $\langle M, H, X\rangle \ll_{\text{cmoc}} \langle M', H', X'\rangle$ iff $M = M'$ and there is some $\mathtt{t}$ such that

1. $H(t) = H'(t)$ and $X(t) = X'(t)$ for all $t < \mathtt{t}$;
2. Either of the following holds:
   (a) $X(\mathtt{t}) \subset X'(\mathtt{t})$, or
   (b) $X(\mathtt{t}) = X'(\mathtt{t})$, and $\mathtt{Diff}(H(\mathtt{t}-1), H(\mathtt{t})) - X(\mathtt{t}) \subset \mathtt{Diff}(H'(\mathtt{t}-1), H'(\mathtt{t})) - X(\mathtt{t})$.

Here, $X(t)$ represents the set of those features f for which $X(f,t)$ is true. The effect of $\ll_{\text{cmoc}}$ is to chronologically minimize changes in *unoccluded* features. The entailment method CMOC is defined using filtering as $\langle \mathbf{G}_{\text{seq}}, S_{\text{cmoc}}\rangle$, where $\mathbf{G}_{\text{seq}}$ is like before, and

$$S_{\text{cmoc}}(A, B) = \mathit{Min}(\ll_{\text{cmoc}}, A) \cap B.$$

The assessment of this entailment method is as follows. It was proved in [6] that CMOC is correct for all equidurational chronicles, and that this requirement is not only sufficient but also necessary, in the same sense of necessary as for PCM.

5.3 Other Entailment Methods

A number of other entailment methods have been assessed in [5]. Some of them involve pretransformations that are more complex than $\mathbf{G}_{\mathrm{seq}}$; others involve more complex preference relations. Their assessments often involve preference-order compatibility or other similar constraints.

6 Inertia with Ramification

The framework that has been described so far will now be extended to the case of inertia with ramification. These are new results.

6.1 World Descriptions for Inertia with Ramification

A logical world description for a world with ramification is $\langle \mathbf{IAD}, \text{LAW}, \text{DCS} \rangle$, that is, a threetuple as already specified above. Here LAW is a set of action laws of the same kind as for the **IA** case, and DCS is a set of static domain constraints. The corresponding set-theoretic formulation of an **IAD** world description is now introduced as a tuple $\langle \mathbf{IAD}, \mathit{Infl}, \mathit{Rtrajs}, \mathcal{R}_C, \mathit{Dep} \rangle$, where the first three components are like before. $\mathcal{R}_C$ is a subset of $\mathcal{R}_\sigma$, and is intended as the set of states which are acceptable in view of the static domain constraints. Thus, $\mathcal{R}_C$ is effectively the set of models for DCS, but represented simply as states and without the temporal dimension.

Furthermore, Dep is a function from action symbols times states to sets of features, such that $\mathit{Dep}(E,r)$ and $\mathit{Infl}(E,r)$ are disjoint sets. It is intended that $\mathit{Dep}(E,r)$ shall contain those features whose values are determined by the static domain constraints, so that the changes in these features need not be specified by the action laws and should not be subject to minimization.

The functions Infl and Dep together divide $\mathcal{F}$ into three partitions:

occluded features: members of $\mathit{Infl}(E,r)$;

dependent features: members of $\mathit{Dep}(E,r)$;

remanent features: not members of either $\mathit{Infl}(E,r)$ or $\mathit{Dep}(E,r)$.

In rough terms, dependent features correspond to the "non-frame" features of Kartha and Lifschitz, and occluded features correspond to their "frame, released" features [1]. For a more precise comparison, and for a motivation for the use of these partitions, please refer to [7].

In the case of **IA**, we were able to replace a world W by a world description $\langle \mathbf{IA}, \text{LAW} \rangle$ that conserved W, and then use the world description as the basis for the logical treatment. Therefore, the step from W to the world description was independent of the preference relation. This is not the case for **IAD** world descriptions: the basic idea here is to allow, for a given W, a function $\mathit{Infl}(E,r)$ which is so restrictive that it omits some features which actually change their values in some trajectories in $W(E,r)$, but to rely on the static domain constraints and the minimization of change for recovering the lost information.

We write $\mathcal{R}_C^*$ for the set of all (finite) trajectories and (infinite) histories consisting exclusively of states in $\mathcal{R}_C$. A world W is said to *preserve* a set $\mathcal{R}_C \subseteq \mathcal{R}_\sigma$ iff $W(E,r) \subseteq \mathcal{R}_C^*$ for every $r \in \mathcal{R}_C$ and every action symbol E.

Let $\mathcal{R}_C$ and a world W preserving $\mathcal{R}_C$ be given. For given functions *Infl* and *Dep* of the kinds defined above, the *corresponding* **IAD** world description is the fivetuple

$$\langle \mathbf{IAD}, \mathit{Infl}, \mathit{Rtrajs}, \mathcal{R}_C, \mathit{Dep} \rangle,$$

where *Rtrajs* is obtained from W just like before, by restricting all trajectories to their occluded features. Formally,

$$\mathit{Rtrajs}(E,r) = \{v \ddagger \mathit{Infl}(E,r) \mid v \in W(E,r)\}.$$

In the corresponding logic-formula variant of the world description,

$$\langle \mathbf{IAD}, \textsc{law}, \textsc{dcs} \rangle,$$

$\textsc{dcs}$ is obtained from $\mathcal{R}_C$, and $\textsc{law}$ is obtained from the other components in such a way that interpretations have the form $\langle M, H, X, Y \rangle$, where M, H, and X are like before, and Y is analogous to X but for the *Dep* component. Thus, $Y(f,t)$ shall be true iff the feature f is a member of $\mathit{Dep}(E,r)$ where E is being executed at time t. The definition of the coercion function Ci is adjusted accordingly.

Intended models are defined like before, using W. Selection functions for $\mathcal{K}$-**IAD** shall take three arguments, and assign

$$S([\![\textsc{law}[\textsc{scd}]]\!], [\![\textsc{obs}]\!], [\![\textsc{dcs}]\!])$$

as the set of selected models for a given chronicle $\langle \mathcal{K}, \langle \mathbf{IA}, \textsc{law}, \textsc{dcs} \rangle, \textsc{scd}, \textsc{obs} \rangle$. Clearly $[\![\textsc{dcs}]\!]$ consists of the infinite members of $\mathcal{R}_C^*$.

6.2 Preferential Conservation

The preferential conservation concept combines the functions of "conservation" and "compatibility" which were used for **IA** ontology. The basic idea is that a world description is m-conservative for a given W iff one can reconstruct the originally given world using minimization with the preference relation $\ll_m$. Notice that it is the combination of a world and a world *description* (that is, essentially, a set of action laws for the world) that is m-conservative, and not the world itself.

The specifics are as follows. Let

$$\mathtt{Diw} = \langle \mathbf{IAD}, \mathit{Infl}, \mathit{Rtrajs}, \mathcal{R}_C, \mathit{Dep} \rangle$$

be a world description corresponding to a world W that preserves $\mathcal{R}_C$. For each combination of E and r, we are to construct two sets of finite models, each finite model being a tuple $\langle M, H, X, Y\rangle$ where H is a history for a finite interval $[0,k]$ of time. The given world description is m-conservative iff those two sets are equal.

The first set is defined using a generalization of the corresponding model function *Cmf*. If F is a set of feature symbols, $F^{\bullet}$ is a function such that $F^{\bullet}(f,t)$ is true iff $f \in F$, for arbitrary integer $t > 0$. Then the function *Cmf* is adapted to read

$$Cmf(E,r) = \{\langle M_0, (r;v), Infl(E,r)^{\bullet}, Dep(E,r)^{\bullet}\rangle \mid v \in W(E,r)\}.$$

Comparing this to the old definition in Subsection 4.1, we notice that the new one depends on `Diw` and not only on W.

Secondly, for given `Diw` with the structure indicated above, we define the *expanded trajectories function* as

$$Xtf(E,r) = \{\langle M_0, (r;v), Infl(E,r)^{\bullet}, Dep(E,r)^{\bullet}\rangle \mid$$
$$v \ddagger Infl(E,r) \in Rtrajs(E,r) \wedge v \in \mathcal{R}_C^{*}\}.$$

In other words, as we went from W to *Rtrajs* by removing the values for all features outside $Infl(E,r)$, we now go back by allowing those features to vary freely within only the constraints of $\mathcal{R}_C$. It follows at once that

$$Cmf(E,r) \subseteq Xtf(E,r)$$

for every E and r. We say that the given world description is *m-conservative* for a given world iff

$$Cmf(E,r) = Min(\ll_m, Xtf(E,r))$$

for all E and r. This means that the preference relation $\ll_m$ is able to recover the information that was lost when W was "projected" to obtain *Rtrajs*. Notice that the property of being m-conservative is a property of the world *description*, and not of the world itself, since it depends on the choice of *Infl*, $\mathcal{R}_C$, and *Dep*.

6.3 Chronological Preferential Entailment

The entailment methods MRC and MRCC are natural generalizations of CMOC allowing ramification. MRC stands for *Minimal Remanent Change*, and MRCC stands for *Minimal Remanent Change with Changeset-partitioning*; they are defined as follows. The preference relation $\ll_{\texttt{mrc}}$ is defined so that

$$\langle M, H, X, Y\rangle \ll_{\texttt{mrc}} \langle M', H', X', Y'\rangle$$

iff $M = M'$ and there is some $\mathtt{t}$ such that all of the following hold:

1. $H(t) = H'(t)$, $X(t) = X'(t)$, and $Y(t) = Y'(t)$ for all $t < \mathtt{t}$.

2. Either of the following holds:

(a) $X(\mathtt{t}) \subset X'(\mathtt{t})$;

(b) $X(\mathtt{t}) = X'(\mathtt{t})$ and $Y(\mathtt{t}) \subset Y'(\mathtt{t})$;

(c) $X(\mathtt{t}) = X'(\mathtt{t})$, $Y(\mathtt{t}) = Y'(\mathtt{t})$, and
$\mathtt{Diff}(H(\mathtt{t}-1), H(\mathtt{t})) - R(\mathtt{t}) \subset \mathtt{Diff}(H'(\mathtt{t}-1), H'(\mathtt{t})) - R(\mathtt{t})$,
where $R(t) = X(t) \cup Y(t)$.

Then MRC is defined as $\langle \mathbf{G}_{\mathtt{seq}}, S_{\mathtt{mrc}} \rangle$ using

$$S_{\mathtt{mrc}}(A, B, C) = Min(\ll_{\mathtt{mrc}}, A \cap C) \cap B.$$

The preference relation $\ll_{\mathtt{mrcc}}$ is defined in a similar fashion, except that the last subitem in the definition is

- $X(\mathtt{t}) = X'(\mathtt{t})$, $Y(\mathtt{t}) = Y'(\mathtt{t})$, $H(\mathtt{t}) \ddagger X(\mathtt{t}) = H'(\mathtt{t}) \ddagger X(\mathtt{t})$, and $\mathtt{Diff}(H(\mathtt{t}-1), H(\mathtt{t})) - R(\mathtt{t}) \subset \mathtt{Diff}(H'(\mathtt{t}-1), H'(\mathtt{t})) - R(\mathtt{t})$.

The selection function $S_{\mathtt{mrcc}}$ and the entailment method MRCC itself are defined as analogous with MRC.

The basic idea in these two entailment methods is to do chronological minimization of changes in remanent features, while not minimizing changes in occluded or dependent features. Changes in occluded features are supposed to be specified by the action laws, and are not to be minimized. This is instrumental for dealing systematically with nondeterministic actions.

Changes in dependent features are forced by the domain constraints, and are also not subject to minimization in these methods. One might think that minimization of changes in dependent features is harmless, since such features are typically completely determined by the values of other features. However, well-known scenario problems such as the Aunt Agatha Livingroom [9] and the Double Switch scenarios [3] have shown that one may obtain unintended tradeoffs between changes in remanent and in dependent features, thereby introducing unintended selected models.

These two entailment methods give the same result for strictly deterministic actions. They differ for non-deterministic actions, where MRCC will minimize changes in remanent features separately for each outcome in occluded features, whereas MRC will minimize changes in remanent features jointly for all outcomes in occluded features.

The MRC entailment method is a generalization of the defined semantics for the language $\mathcal{AR}_0$ by Kartha and Lifschitz [1].

The assessment proofs for PCM and CMOC can easily be adapted to show that MRC is correct for equidurational chronicles containing MRC-conservative world descriptions, and similarly for MRCC. Neither one of MRC-conservativeness or MRCC-conservativeness subsumes the other, since it is easy to construct a world description that is MRC-conservative but not MRCC-conservative, and vice versa [7].

The case of strictly inert world descriptions can be obtained easily, by choosing $Dep(E, r) = ?$ for all E and r, and $\mathcal{R}_C = \mathcal{R}_\sigma$. It is trivial that every strictly inert world description that conserves a world W is MRC- and MRCC-conservative for W. Therefore, every world has *some* MRC- and MRCC-conservative description, which can be obtained by including sufficiently many features in $\mathit{Infl}(E, r)$. In particular, if one includes all features whose values may change in the course of the action, strict inertia is obtained, and both MRC- and MRCC-conservation is achieved trivially. From a practical point of view, however, one presumably wishes to keep $\mathit{Infl}(E, r)$ as small as possible, in order to make the corresponding action law as compact as possible.

In this context, it was demonstrated in [7] that for single-timestep actions, MRCC has an interesting monotonicity property: if

$$\langle \mathbf{IAD}, \mathit{Infl}, \mathit{Rtrajs}, \mathcal{R}_C, \mathit{Dep} \rangle$$

is an MRCC-conservative world description corresponding to a world W, Infl' satisfies $\mathit{Infl}(E, r) \subseteq \mathit{Infl}'(E, r)$ for all E and r, and Rtrajs' is formed from Infl', then

$$\langle \mathbf{IAD}, \mathit{Infl}', \mathit{Rtrajs}', \mathcal{R}_C, \mathit{Dep} \rangle$$

is also MRCC-conservative with respect to W. (The generalization to actions with arbitrary finite duration seems to be straight-forward). This result means that if one starts with the full world W and gradually reduces the extent of *Infl*, then the resulting world descriptions are at first MRCC-conservative, but if one reaches a point where the world description is not MRCC-conservative, and one then continues to reduce the extent of *Infl*, one will never again get back to an MRCC-conservative world description. MRC does not have the same monotonicity property.

In other words, the concept of a "smallest MRCC-conservative *Infl* function" is well defined for each choice of W, $\mathcal{R}_C$, and *Dep*. This result is of interest since in practice, one likes to involve as few features as possible in the action laws. It would be convenient to have an algorithm which takes a specification of W and obtains, in the form of action laws, the smallest MRCC-conservative action laws by minimizing the extent of *Infl*. The monotonicity result means that it is meaningful to look for such an algorithm.

7 The Systematic Approach: Current Developments

In this paper, I have tried to present, in very condensed form, the *conceptual framework* that is used in the *systematic approach* to logics of actions and change. The complete presentation of that framework can be obtained in the monograph [5]. The basic approach is presently being extended in a number of directions, allowing for concurrency, ramification, faulty observations, etc. There is also considerable work on implementation methods for some of the assessed entailment methods. The present paper has only showed current work in one direction, namely for ramification.

An up-to-date overview of the state of the art in this research can be obtained from the following WWW page:

`http://www.ida.liu.se/labs/rkllab/groups/FF/`

which contains an overview of current results and links to current articles.

References

[1] G. Neelakantan Kartha and Vladimir Lifschitz. Actions with indirect effects (preliminary report). In *International Conference on Knowledge Representation and Reasoning*, pages 341–350, 1994.

[2] Henry Kautz. The logic of persistence. In *National Conference on Artificial Intelligence*, pages 401–405, 1986.

[3] Vladimir Lifschitz. Restricted monotonicity. In *National Conference on Artificial Intelligence*, pages 432–437, 1993.

[4] John McCarthy and Patrick Hayes. Some philosophical problems from the viewpoint of artificial intelligence. In *Machine Intelligence, Vol. 4*, pages 463–502. Edinburgh University Press, 1969.

[5] Erik Sandewall. *Features and Fluents. The Representation of Knowledge about Dynamical Systems. Volume I.* Oxford University Press, 1994.

[6] Erik Sandewall. The range of applicability of some nonmonotonic logics for strict inertia. *Journal of Logic and Computation*, 4(5), October 1994.

[7] Erik Sandewall. Systematic comparison of approaches to ramification using restricted minimization of change. Technical Report LiTH-IDA-R-95-15, IDA, 1995.

[8] Yoav Shoham. *Reasoning about Change.* MIT Press, 1988.

[9] Marianne Winslett. Reasoning about actions using a possible models approach. In *National Conference on Artificial Intelligence*, pages 89–93, 1988.