Transition cascade semantics and first assessments results for ramification
Preliminary report

Erik Sandewall
Email: erisa@ida.liu.se

Abstract

This article reports on assessment results for several approaches to the ramification problem. Two types of assessments are reported: (1) *Assessment of soundness* for one minimization based and one causation-based method; (2) *Relative range assessments* for a number of minimization based methods.

Assessment of soundness is based on an underlying semantics for ramification. We propose, define, and use a *transition cascade* semantics for this purpose. Transition cascade means that an action is viewed as consisting of an initial state transition which represents the invocation of the action, followed by a succession of other state transitions which may be understood as representing a causal chain. The assessment of an entailment method specifies restrictions on the invocation and causation relations which guarantee that the entailment method is sound.

Relative range assessments compare entailment methods pairwise and specify whether the set of selected models obtained by one is a subset of the set of selected models obtained by the other.

IDA Technical Report 1996
LiTH-IDA-R-96-19
ISSN-0281-4250

Department of Computer
and Information Science,
Linköping University,
S-581 83 Linköping, Sweden

1 Validation of theories of action and change

In earlier publications addressing logics of actions and change, we have developed an underlying semantics for the special case of strict inertia, and used it for assessment of the range of applicability of a number of entailment methods. The case of strict inertia does not allow for ramification, qualifications, or surprises. The present article addresses the problem of ramification using the same systematic methodology as was previously used for strict inertia. We start with a discussion of methodology.

1.1 Before the advent of validation: example based methodology

There has been much research in recent years on methods for reasoning about actions and change, and on finding solutions of the so-called "frame problem". New variants of nonmonotonic logics for common-sense reasoning have been proposed, only to be quickly refuted by counterexamples. Unfortunately, the results that have been obtained in this fashion are notoriously unreliable. According to a common research methodology in the area, the evidence in favor of a proposed logic or theory should consist of intuitive plausibility arguments and a small number of scenario examples for which the theory is proven (or claimed) to give the intended conclusions and no others. This prevailing methodology was summarized by Davis [6] as follows:

> *The basic approach used here, as in much of the research on automating commonsense reasoning, is to take a number of examples of commonsense inference in a commonsense domain, generally deductive inference; identify the general domain knowledge and the particular problem specification used in the inference; develop a formal language in which this knowledge can be expressed; and define the primitives of the language as carefully and precisely as possible...*

This methodology has been effective for identifying a range of important problems, such as the qualification and ramification problems. However, since it relies on examples, it is not sufficient for obtaining reliable results of a general nature. One also needs to *validate* a proposed approach and prove that it is correct for the intended type of reasoning about a whole class of scenarios.

1.2 Validation of nonmonotonic theories

The quotation from Davis indicates also how each proposed approach to common-sense reasoning must specify the syntax of its language of logical formulae, the conventions for how to use that language for describing actual scenarios, and the methods for drawing conclusions from a scenario description. I shall use the term *method* or (more explicitly) *entailment method* for the formally defined mapping from scenario descriptions to sets of conclusions (or equivalently, models), and reserve the use of the word *logic* for the syntax and the formal semantics which may be used by a particular entailment method, or by several methods jointly. The logic used by an entailment method will be called its *base logic*. The combination of an entailment method, its base logic, and the conventions for their use will be referred to as a *theory*, for example, a theory of actions and change.

We shall only be concerned with validation approaches which apply specifically to the entailment method, while taking the underlying base logic and conventions as given.

The validation of an entailment method must rely, besides the base logic, on a formal definition of the function Σ from a *scenario description* Υ, expressed using logical formulae, to the set $\Sigma(\Upsilon)$ of its *intended models*. If $[\![\Upsilon]\!]$ represents the set of classical models, then it is commonly assumed that $\Sigma(\Upsilon) \subseteq [\![\Upsilon]\!]$, reflecting the possibility of conclusions by default. Once the intended models function has been defined, one can analyze a proposed entailment method in terms of its formulation as a *selection function* S from scenario descriptions to sets of models. The validation consists of proving that the set $S(\Upsilon)$ of *selected models* is equal to $\Sigma(\Upsilon)$ for an interesting and sufficiently large class of scenario descriptions Υ.

The validation criterion can of course be formulated as well in terms of sets of conclusions: if $Cn(\Upsilon)$ is the set of classical conclusions from the given scenario description, $C(\Upsilon)$ is the set of actual conclusions using a proposed entailment method and $Ci(\Upsilon)$ is the set of intended conclusions according to the chosen formal definition, then usually $Cn(\Upsilon) \subset Ci(\Upsilon)$ especially when a nonmonotonic logic is being used, and one requires from C that $C(\Upsilon) = Ci(\Upsilon)$. In this article I use the model-theoretic formulation since it is the more convenient one for analyzing logics based on *preferential semantics*, where the set of selected models is chosen (in the simplest case) as $Min(\ll, [\![\Upsilon]\!])$, i.e. as the set of those classical models which are minimal with respect to some preference relation $\ll$. Many of the proposed entailment methods for action and change have this character.

This framework for validation depends on the fact that the base logic is held fixed, so that the notion of intended models can be defined in the terms of that logic. Validation and comparison of whole theories, containing both a logic and an entailment method, is a more complex affair and will not be discussed here.

1.3 Approaches to the validation of theories of actions and change

In the case of reasoning about actions and change, the scenario description is a *chronicle* containing statements about the state of the world at various points in time (the *observations*), about specific actions and their temporal order (the *schedule* of actions), as well as *action laws* defining the effects of actions. Models represent possible histories of the world; those models which satisfy all the statements but violate the "frame" assumption of persistence or inertia are classical but unintended models. Additional types of axioms may also be included for special purposes.

Lin and Shoham [22] were probably the first to report a validation of a previously proposed theory of action. They showed that chronological minimization of change is correct for prediction in deterministic situation-calculus theories with strict inertia, i.e. no surprises or qualification, but allowing for ramification. The assumption of determinism applies for the ramification side-effects as well, so the set of intended models is required to be a singleton in their result.

A number of later papers have obtained validations of other theories, and for significantly broader classes of chronicles. Three main approaches have been followed.

1.3.1 Systematic methodology - "Features and fluents"

The present author has used a *systematic methodology* whereby one can obtain precise results on the *range of applicability* of a number of previously proposed entailment methods, as well as some new ones [32, 33, 34, 35]. These results apply, right from the start, to a fairly large class of logics for actions and change, allowing for chronicles with nondeterministic actions and discrete metric time. This includes integer time as a special case, and allows one to represent actions with extended duration, which may be characterized not only in terms of their final effects, but also with respect to changes within their interval of duration. It also allows chronicles where the order of the actions is only partially specified, like in nonlinear planning [28]. Several of the validations generalize trivially to the case of deterministic ramification, but ramification in nondeterministic chronicles is more difficult and was not addressed in those articles.

We have also shown in [34] that these results generalize easily to the case of branching metric time as well, which means that situations calculus is subsumed as a simple special case. Yi [40] generalized some of these results to apply for scenarios with concurrency in the context of explicit metric time.

These validation results are *assessments*, which identify the range of correct applicability of each of a number of methods. We initially define a relatively large class of chronicles, and subdivide the class using a *taxonomy* in terms of a number of properties that chronicles may have. The complete assessment of an entailment method consists of both a lower bound and an upper bound. The lower bound is a family of chronicles within the taxonomy for which the method is provably correct; the upper bound is a family (larger or equal to the lower bound) outside which it is provably impossible that the method is correct. In this way it has sometimes been possible to identify the exact range of applicability of previously proposed theories of action.

The definition of the taxonomy is based on an *underlying semantics* which captures basic notions of intelligent agents, and defines intended models in its terms. This underlying semantics can be viewed as a kind of abstract simulator of the interaction between an intelligent robot and its environment. The same underlying semantics is used both for defining the intended-models function Σ and the taxonomy of chronicles.

1.3.2 Validation relative to monotonic formalizations

Other validation results for chronicles in situation calculus with strict inertia have been reported by Lifschitz [18], Reiter [27], and Kartha [12]. Validations for theories with ramification have been reported by Lin and Reiter [21] and by Peppas and Wobcke [26]. Validations for the generalization of situation calculus to theories with concurrent single-timestep actions have been reported by Lin and Shoham [23] and by Baral and Gelfond [2].

Different researchers have also used different validation criteria. Lin and Shoham [22], Lifschitz [18], and Reiter [27] relate a nonmonotonic theory over a certain set of "common-sense" axioms, to a monotonic theory over a larger set of axioms which is assumed to capture the full intentions. In other words, they define $\Sigma(\Upsilon)$ as $[\![\mathbf{G}(\Upsilon)]\!]$, where $\mathbf{G}$ is an operation which extends or transforms the given set of logical formulae.

1.3.3 Validation relative to an action description language

Lifschitz and Gelfond have introduced a separate specification language $\mathcal{A}$ for defining the effects of actions, and used it to define the intended conclusions for reasoning about actions [10]. Kartha [12, 13] applied the same criterion to the theories previously proposed by Baker [1], Pednault [25], and Reiter [27]. Baral and Gelfond [2] generalized $\mathcal{A}$ to the case of concurrent single-timestep actions and used it for validations as well.

1.3.4 Postulate oriented approaches

Yet another approach, which does not really count as validation, is to constrain the possible choices of the set of intended conclusions through a number of *postulates* for the relation between Γ and $\Sigma(\Gamma)$, or between $[\![\Gamma]\!]$ and $\Sigma(\Gamma)$. Del Val and Shoham [7] use the postulates of Katsumo and Mendelzon [15] as the basis, and Peppas and Wobcke [26] use an extended set of postulates for epistemic entrenchment as the basis for the correctness proof for their system. Since postulates are in principle written to constrain the choice of $\Sigma(\Gamma)$ but not to specify it completely, it seems that their use will lead to intrinsically weaker results.

1.3.5 Equivalence results

Given this variety of validation criteria one may also consider to what extent they are equivalent and exchangable. A first step towards such an analysis has been reported by Thielscher [37], who identified the relationship between the $\mathcal{A}$ language of Gelfond and Lifschitz and a subset of the underlying semantics that is used in the present paper.

2 The systematic methodology

The introduction and use of an underlying semantics raises questions of how a proposed logical theory is related to the *meaning* of the formulae that it uses, and of what that very meaning is. Let us now address how these questions are viewed in the systematic methodology.

This methodology proposes to use a two-step process for going from common sense to logic. The first step attempts to capture some characteristics of common sense in terms of a formal definition of $\Sigma(\Upsilon)$; the second step tries to find a theory whose selection function S satisfies $S(\Upsilon) = \Sigma(\Upsilon)$.

What is then the evidence that a proposed definition of Σ itself captures common sense reasoning in an adequate manner?

2.1 The use of a progression of underlying semantics

In our view, the formalization of common sense can only be achieved as the limiting case in a progression of successively more refined functions Σ_1, Σ_2, and so on. To illustrate the need for several such Σ_i, consider the "stolen car scenario", a very simple example of common-sense reasoning originally proposed by Kautz in [16]. In this chronicle it is assumed that if a certain car is left in the garage overnight then it will still be there the next morning. It is also observed that after the car has been left in the garage for two nights, it is no longer there. What are the common-sense conclusions, or equivalently what should be the set $\Sigma(\Upsilon)$ of intended models for this chronicle?

One possible ontology assumes simple persistence without any exceptions or modifications. I will refer to this case as *strict inertia.* Then the given chronicle is inconsistent, since the car is supposed to remain in the garage after two nights, but it is also stated not to be there. The set of intended models for the stolen car scenario under strict inertia is therefore empty. Another possible ontology allows persistence to be modified by "surprises", and at the same time one makes the *epistemological assumption* that $\Sigma(\Upsilon)$ must be non-empty except in the case where the set $[\![\Upsilon]\!]$ of classical models is itself empty. The set of surprises is minimized subject to the restriction $\Sigma(\Upsilon) \neq \varnothing$. In this case, the given chronicle has two intended models: one where the car vanishes surprisingly during the first night, and one where the car vanishes surprisingly during the second night.

The well-known "qualification" and "ramification" problems concern additional ontological characteristics of common-sense reasoning (CSR) besides surprises. It appears that full CSR involves a considerable number of such difficulties: not only surprises, qualification, and ramification, but also concurrent events, causation chains of events, creation and termination of objects, unreliable observations, and so forth. All of them represent challenges for a logic of common-sense reasoning. However, none of the proposed theories in the current literature gets even close to addressing all of these phenomena, although several of them propose to address more than one phenomenon by one single technical method. In particular, some of the proposed preference relations on models aspire to establish inertia, to minimize side-effects (ramification), and to minimize surprises all at the same time.

It is difficult to compare logics that address different problems and sometimes don't even state exactly what assumptions they make. What I propose, therefore, is to introduce one definition $\Sigma_{\mathbf{IA}}$ for strict inertia, a modified definition $\Sigma_{\mathbf{IAS}}$ for inertia with surprises but without any of the other characteristics, and so on. Unlike the proposed theories, the definitions of the various Σ_i should represent all the characteristics explicitly, so that one can clearly discern each of them.

This proposal has sometimes evoked impatience: why bother with a very simple case such as $\Sigma_{\mathbf{IA}}$ i.e. strict inertia, when other researchers have already thought about more advanced problems such as qualification and ramification? The answer is that the systematic assessment of a theory of action is considerably more complex than finding a few examples where it works, and the assessment is a non-trivial undertaking even for the case of strict inertia. One ought not have to apologize for the strategy of obtaining a solid analysis of the simple cases before one proceeds to the more difficult ones, even though one impatiently feels that the directly useful results will be found in the later stages of the process.

In this perspective, the reason for using the formal definition $\Sigma_{\mathbf{IA}}$ corresponding to strict inertia is not that it is claimed to correspond exactly to common sense – it clearly does not – but because it is the *first step* on a *well structured path* towards the goal of characterizing common sense reasoning; a goal which is very ambitious and which therefore we can not expect to achieve in the near future. For each choice Σ_i of the intended-models function, one obtains a well-defined assessment problem for one or more proposed theories of action, and this is a problem that can be analyzed in an entirely formal fashion.

Although each intended-models function Σ_i is to be *defined* in a formal fashion, the *choice* of that definition for the purpose of CSR is properly outside the scope of formal methods. The traditional methodology in A.I. would leave it to the discretion of the researcher: everyone presumably knows what common sense is? One may speculate whether methods of cognitive psychology could be used for testing the psychological validity of one or the other intended-models functions, but that is not a topic of the present article.

2.2 The methodology for assessment of range of applicability

Consider a given choice of ontological assumptions, in particular the case of strict inertia represented by its intended-models function $\Sigma_{\mathbf{IA}}$, and consider also a proposed entailment method viewed as a selection function S. A *direct validation* of S with respect to $\Sigma_{\mathbf{IA}}$ will simply verify whether $S(\Upsilon) = \Sigma_{\mathbf{IA}}(\Upsilon)$ for every chronicle Υ that is correctly formed with respect to the chosen ontological assumptions, and the outcome will state whether S is or is not always correct. However, even if S is not always correct, one may be interested in its *range of applicability*, i.e. in a family $\mathcal{Z}$ of chronicles such that $\Upsilon \in \mathcal{Z} \Rightarrow S(\Upsilon) = \Sigma_{\mathbf{IA}}(\Upsilon)$. Such results, which I call assessments, have a potential practical interest: it is not evident that a broader-range theory will always be preferred, since an entailment method with a more narrow range of applicability might allow a more efficient implementation. However in order to make use of it, one must have precise knowledge of whether it is correctly applicable for the application one has at hand.

Range of applicability results have a theoretical interest as well: as we proceed to successive new definitions Σ_{i} of the intended models, a selection function that agrees completely with one version Σ_{i} can by definition not agree completely with the next one. An understanding of selection functions that are almost correct for one stage may therefore be useful for the next stage of analysis.

The range of applicability may be characterized in terms of various properties of the chronicle, for example whether it contains a nondeterministic action, whether it contains observations for time greater than zero, and so on. I will refer to such properties as *sub-characteristics* since they provide a finer-grain categorization in addition to the one that is obtained from the main characteristics for inertia, surprises, ramification, etc. The collection of characteristics and sub-characteristics together constitute a *taxonomy* of chronicles which may be used as the coordinate system for specifying and comparing the range of applicability of various theories of action.

Notice what is the logical order in which the various definitions must be made:

1. Define the reportoire of main characteristics.
2. Select some class of chronicles $\mathcal{Z}_i$ in terms of the chosen main characteristics.
3. Choose a definition of syntactic structure for chronicles in $\mathcal{Z}_i$.

4. Define the underlying semantics for $\mathcal{Z}_i$, and in particular the selection function Σ_i from well-formed chronicles to their intended model-sets.

5. The underlying semantics is also used for the formal definition of the sub-characteristics which generate a taxonomical structure within $\mathcal{Z}_i$.

6. Proposed theories of action for $\mathcal{Z}_i$ are expressed as selection functions and are assessed with respect to the selection function Σ_i. The results are expressed in terms of the taxonomy defined by the sub-characteristics.

7. Algorithms and other implementation aspects are developed for the theories that have been assessed.

Each definition of Σ_i is merely an approximation to real common sense. When the entailment methods and implementations for one choice of $\mathcal{Z}_i$ are sufficiently well understood, one proceeds to a more refined class and repeats the process. One hopes of course that the analysis of a later class can make use of the results from earlier classes in the progression.

The earlier papers by us and our colleagues showed how this methodology could be applied to the case of strict inertia, allowing for nondeterminism, metric linear time, metric branching time, and concurrent actions. In the present article we begin to extend the same approach to the case of ramification, using a causation-oriented underlying semantics.

3 Underlying semantics for strict inertia

Some of the difficulties of the "frame problem" – reasoning about actions and change – relate to how the effects of one particular action are described in terms of the logic at hand. Ramification is an example of such an action-level problem: it has to do with how action laws can be restricted to only specifying the "main effects" or "most significant effects" of actions, whereas other effects are obtained indirectly from more global knowledge.

Other difficulties relate to the structure of the chronicle as a whole, including both action laws, schedule, and observations, as well as other kinds of information which may also be available. The problem that straightforward chronological minimization does not work correctly for postdiction is an example of such a chronicle-level difficulty. This problem was solved by the introduction of filtering [30]. The correctness and full generality of that solution was formally proved in [34]. A reformulation of filtering in terms of circumscription, called nested circumscription, was presented by Lifschitz in [19]. Dealing with surprises offers another example of chronicle-level difficulties: presumably, one wishes to make a trade-off between possible surprises early or late during a chronicle, and not merely minimize them chronologically.

Correspondingly, an underlying semantics must characterize the intended models of a chronicle using both the action level and the chronicle level. In this section I shall show how both of these levels are modelled in the underlying semantics for strict inertia. When we proceed to ramification, which is the topic of the present article, it is only the action level that needs to be changed, whereas the chronicle level can be retained unchanged – ramification is intrinsically an action-level phenomenon. Moreover, even the action level for ramification is constructed as a moderate generalization of what was introduced for strict inertia.

We shall start with a basic state-transition semantics which is very simple, and which in particular does not at all reflect the concept of inertia (persistence) which is characteristic of the "frame problem". We then proceed to slightly more detailed semantics which capture the essentials of the frame problem for strict inertia. The further generalization to ramification will follow in the next section.

3.1 State-transition semantics

The following basic concepts are used. Let $\mathcal{R}$ be the set of possible states of the world, formed as the Cartesian product of the finite range sets of a finite number of *state variables.* Also, let $\mathcal{E}$ be the set of possible actions, and let the *main next-state function* $\mathbf{N}(E, r)$ be a function from $\mathcal{E} \times \mathcal{R}$ to non-empty subsets of $\mathcal{R}$. The function $\mathbf{N}$ is intended to indicate the set of possible result states if the action E is performed when the world is in state r. The assumption $\mathbf{N}(E, r) \neq \emptyset$ expresses that every action E can always be executed in every starting state r. (The symbol E is derived from "event"; we consider actions and events to be closely related).

Now, consider a chronicle consisting of action laws, schedule, and observations, where the action laws specify a next-state function $\mathbf{N}$ using some appropriate notation, and where time is assumed to be taken as the non-negative integers. An interpretation for such a chronicle shall be chosen as a history over time, that is, a mapping from timepoints to states. (Assignments of values to constant symbols etc. are also needed, but are of marginal importance for the present purpose). The set of intended models for a chronicle is defined using an *ego-world game*, as follows. First, the observations are removed from the chronicle, so that only the action laws and the schedule remain. The following process is performed using the observation-free chronicle: Start with time = zero and the world in an arbitrary state. At each point in time, perform the move of the "ego" or agent, and the move of the world. The move of the ego is dictated by the schedule, and consists of selecting the action that is indicated by the schedule for the current point in time, or no action if nothing is specified. The move of the world is as follows: if the state of the world at the current time t is r, and the ego has just chosen the action E, then select some r' which is a member of $\mathbf{N}(E, r)$, make it the state of the world at time $t + 1$, and increase the current time by 1. If no move is required, assign the state r to the next timepoint $t + 1$ as well.

The desribed procedure will non-deterministically generate a set of possible histories of the world. The given definition applies trivially if all actions are specified explicitly by the schedule. Less trivial cases arise if actions are specified in other ways, for example as a disjunction between actions ("either action A or B is performed"), or with an imprecisely specified timing. However, those are technical details, and the general idea remains the same.

Among the world histories that are generated by the ego-world game, one then selects a subset consisting of those histories where all the observations are satisfied. That subset is the set of intended models for the given chronicle.

In this case, which is the simplest one of all, the set of intended models will of course be exactly the set of models for the chronicle, using a conventional concept of models. We now proceed to some less trivial cases, where the characteristic aspects of the frame problem enter the picture.

3.2 Partial-state transition semantics

The purpose of the partial-state transition semantics is to capture the basic assumption of inertia ("no change except as the explicit result of an action") while retaining the general flavor of the state-transition semantics. We use $\mathcal{R}$, $\mathcal{E}$, and $\mathbf{N}$, like before, but we also introduce an *occlusion function* $\mathbf{X}(E, r)$ which is a function from $\mathcal{E} \times \mathcal{R}$ to subsets of the set of state variables. The occlusion function is intended to specify, for a given action E and starting state r, the set of those state variables whose value is affected when the action E is performed in the state r. In other words, it must be the case that any state variable which is not a member of $\mathbf{X}(E, r)$ has the same value in r and in all the members of $\mathbf{N}(E, r)$. (Thus, it would be correct to select $\mathbf{X}(E, r)$ as a the set of all state variables, but it is intended that smaller sets shall be chosen whenever possible).

3.2.1 Action laws and auxiliary next-state functions

The functions $\mathbf{N}$ and $\mathbf{X}$ together determine the action laws down to the level of semantic equivalence: if these two functions are given, then the action law for an action E shall be constructed so that for each $r \in \mathcal{R}$, it specifies what $\mathbf{N}(E, r)$ says about the possible new values of the state variables in $\mathbf{X}(E, r)$, and it shall not say anything about the new values of other state variables. (Notice, however, that the new values of state variables in $\mathbf{X}(E, r)$ may equally well *depend on* the values of state variables outside $\mathbf{X}(E, r)$ in the starting state r of the action).

Since we are only interested in the set of new states that are permitted by the action law, and not in the actual syntactic form of that action law, it is not necessary to introduce any assumptions about the choice of logical language or about the exact structure of action laws. Using a much simpler approach, we shall introduce two additional, similar but distinct functions $\mathbf{N}'$ and $\mathbf{N}^\circ$ which are derived from the combination of $\mathbf{N}$ and $\mathbf{X}$, and which will be the basis for defining and analyzing the nonmonotonic entailment methods. Of these functions, $\mathbf{N}^\circ$ will be used for strict inertia, and $\mathbf{N}'$ for the analysis of ramification.

The function $\mathbf{N}'$, the *stated next-state function*, is obtained in the following manner. For every result state $r' \in \mathbf{N}(E, r)$, identify the set of all states which agree with r' for all state variables in $\mathbf{X}(E, r)$, but which have arbitrary values elsewhere. Construct $\mathbf{N}'(E, r)$ as the union of all those sets. It follows immediately that $\mathbf{N}'(E, r) \supseteq \mathbf{N}(E, r)$.

This construction should be understood as follows. The action law expresses, in the form of logic formulae, all what $\mathbf{N}$ tells us about the values of state variables in $\mathbf{X}(E, r)$, but nothing about the values of other state variables. We construct the set of models for these logical statements as a subset of $\mathcal{R}$, but using the same vocabulary (the same set of names for state variables) as is used in $\mathcal{R}$. $\mathbf{N}'(E, r)$ has that set of models as its value.

The function $\mathbf{N}'$ that is obtained in this way is therefore a direct counterpart of the usual, logical formulation of an action law which is written as concisely as possible.

Similar to $\mathbf{N}'(E, r)$, the *concise next-state function* $\mathbf{N}^\circ(E, r)$ is obtained as follows. For every E and r, $\mathbf{N}^\circ(E, r)$ is the set of partial states obtained from $\mathbf{N}(E, r)$ by restricting each of its members to be defined for exactly the state variables in $\mathbf{X}(E, r)$.

Notice, therefore, that $\mathbf{N}'$ and $\mathbf{N}^\circ$ are uniquely determined by $\mathbf{N}$ and $\mathbf{X}$, where $\mathbf{X}$ represents the choice as to which state variables to document in the action law. For a given $\mathbf{N}$ and for each choice of $\mathbf{X}$, one obtains a unique $\mathbf{N}'$ and a unique $\mathbf{N}^\circ$. (We do not write $\mathbf{X}$ as an index on $\mathbf{N}$ because we are going to need another index below, and double indices are cumbersome).

Notice, finally, that $\mathbf{N}(E, r) = \emptyset \Leftrightarrow \mathbf{N}'(E, r) = \emptyset$, and similarly for $\mathbf{N}^\circ$. In other words, our assumption of the universal applicability of actions is preserved in the transformations from $\mathbf{N}$ to $\mathbf{N}'$ and $\mathbf{N}^\circ$.

3.2.2 The revised ego-world game

Once the definition of $\mathbf{N}^\circ$ has been obtained, the partial-state transition semantics is similar to the one defined in the previous section. The following changes are made relative to the case of strict inertia. The move of the ego works like before. The move of the world is now as follows: if the state of the world at the current time t is r, and the ego has just chosen the action E, then select some r' which is a member of $\mathbf{N}^\circ(E, r)$, and construct the state of the world at time $t + 1$ as $r \oplus r'$, that is, as the new state obtained by taking for each state variable the value assigned by r' whenever defined, and otherwise the value assigned by r. At the same time, the current value of time is increased by 1.

3.3 Trajectory semantics

Many applications of actions and change require the use of metric time, and of actions where the world goes through several successive states within the duration of the action. In particular, this is necessary for dealing with concurrent actions in a systematic manner. A corresponding underlying semantics is offered by the *trajectory semantics*, where a trajectory is defined as a finite sequence of partial states over a fixed subset of the state variables, and the function $\mathbf{N}^\circ$ of the previous section is replaced by a function which maps every pair (E, r) to a set of trajectories. The game is modified accordingly: if the world is in the state r and the ego has selected the action E, then the move of the world is to select one member of the set of trajectories for the current (E, r), and to append it to the present history. Thus, the chosen trajectory specifies the outcome of the action both with respect to its duration, and with respect to the sequence of states that the world is in within that duration. The generalization to continuous time is straight-forward.

3.4 Previous range-of-applicability results

In earlier publications, we have identified the range of applicability of a number of previously proposed entailment methods with respect to the trajectory semantics. In all cases it was possible to obtain both a lower bound and an upper bound on the range of applicability, and in many of the cases these bounds are identical, which means that the range of applicability has been determined completely.

For the details about the trajectory semantics and the range-of-applicability results, please refer to [34] and to the summary in [35].

4 Ramification

We proceed now to ramification, which is the main topic of the present article. We shall first make a general review of the problem and of current approaches. Then we develop an appropriate underlying semantics, and analyze current approaches in its terms.

The ramification problem for reasoning about actions and change has been described as follows by Ginsberg and Smith, [11]:

> The difficulty is that it is unreasonable to explicitly record all of the consequences of an action, even the immediate ones. For example [...] For any given action there are essentially an infinite number of possible consequences that depend upon the details of the situation in which the action occurs.

Of course, those "infinite number of consequences" have to be specified to the logical system in some way; the logic can not do magic. More precisely, ramification should provide a solution for at least the following two problems:

- If some regularities in the world at hand are reflected in the consequent side of several of the action laws, then it makes sense to factor out the common parts and to represent them once and for all as "domain constraints". Instead of just applying the action law to determine the effects of an action, one applies the action law plus the domain constraints.

- Consider applications where every scenario describes a number of separate but interconnected objects, different scenarios involve different configurations, and each action has its immediate effects on one or a few of the objects, but indirect effects on other objects which are connected to the first ones, in some sense of the word "connected". Then, it would be completely unreasonable to let the action laws contain different cases which enumerate the possible configurations. Instead, action laws should only specify the "immediate" or "primary" effects of the action, and logical inference should be used for tracing how some changes "cause" other changes across the structure of the configuration at hand.

In both cases, action laws are supposed to specify some of the action's effects, which in some sense are the "primary" effects, and the deductive system shall rely on additional information for inferring other changes that are due to the action. One common approach is to use *domain constraints* that specify what are possible states of the world, and to assume that the total set of changes in state variables that result from the action, consists of those changes that are explicitly stated in the action law, plus a minimal set of other changes that are necessary in view of the domain constraints. Notice that such an approach only makes sense in the second of the two items that were just mentioned: one can imagine that the causational propagation across the scenario configuration obeys some minimization criterion, but if domain constraints have merely been factored out of action laws then there is no particular reason for minimization of the changes.

The *update problem* is in principle the same problem: given a world state before an action, some information about the world state after the action, and certain general restrictions on permissible world states, which are the possible or preferred world states resulting from the action? The main difference is one of presentation: the ramification problem is presented in the context of reasoning about actions and change, where the actions themselves as well as their timing and sequencing is explicit in the logical language being used. For the update problem, by contrast, one generally considers a single update, and leaves aside the question of timing and sequencing of the multiple actions. In othe words, one concentrates on the action-level problem, while disregarding the chronicle-level problems as defined at the beginning of section 3, above.

Even among the methods that use minimization of change, there are a number of different variants which have been proposed by different authors, and which differ in their details. So far, these methods have only been developed using the example-based methodology, where each proposal is motivated on the basis of (1) intuitive and informal reasons why a particular method is the "natural" one, and (2) selected test examples where competing methods do not give intended results. We shall briefly summarize some of these approaches, without going too deeply into the pros and cons of each, and describing only those approaches that will be the subject of assessment later on in this paper. There is no need to be exhaustive since our purpose here is to *characterize properties* of some of these approaches, and not to advocate one of them over the others. The list is not exhaustive; we summarize primarily those approaches which are the subject of assessment below.

4.1 Plain minimization of change

The original method, plain minimization of change, is originally due to Winslett [39], at least in the context of the A.I. literature. The idea is simple: identify a set of *static domain constraints*, that is, restrictions on the possible states of the world. For a given action, the effects specified by the action law are obligatory; a minimal set of additional changes is selected so that the domain constraints are satisfied.

It was observed already by Winslett that plain minimization of change sometimes obtains unintended models, because change in a state variable that was not intended to change may offset change (within the restrictions of the domain constraints) of some other state variables where change is indeed intended. Winslett suggested to solve this problem by imposing a priority ordering on the state variables, but this line of research has apparently not been pursued. The following approaches have been developed instead.

4.2 Categorization of state variables

Categorization of state variables is an extension of plain minimization of change: instead of establishing certain changes in the world, and then minimizing changes across all (other) state variables, one divides the state variables into two or more classes which are treated differently with respect to minimization.

One instance of this concept was the introduction of *occlusion* by Sandewall [30, 33], also represented by the *persists* predicate in [29], the *release* operation of Kartha and Lifschitz [14] and the *persists* predicate of del Val and Shoham in [8]. In the simplest case, one separates the state variables into two classes A and B, where the action laws specifying the effects of an action only specify changes in class A, whereas domain constraints are allowed to refer to state variables in both classes. Then, change is minimized for state variables in class B, whereas those in class A are allowed to vary freely to the extent permitted by the action laws.

Another categorization method is to single out some state variables as *dependent* on others, and to stipulate that minimization of change does not apply to dependent state variables. The definitions of the dependencies are part of the domain constraints. This approach was first proposed by Lifschitz as a distinction between frame and non-frame fluents [17], although in implementation terms it goes back to the early STRIPS system [9].

4.3 Explicit specification of dependencies or causation

Plain minimization of change is symmetric: a constraint of the form $a \Rightarrow b$, if imposed on a state where a is true and b is false, can be satisfied either by making a false, or by making b true. This does not always correspond to the intentions in the application at hand, since sometimes the influence applies only one way.

The distinction between dependent and independent state variables is often sufficient for dealing with this problem. However, Thielscher [38] has shown an example, involving relays, where the distinction between dependent and independent variables is claimed not to be sufficient. In his paper, and in concurrent papers by other authors [24, 20], a "causal" approach is proposed. The basic idea is to express *causal dependencies*, which say that change in one state variable necessitates or enables a change in another state variable. (These links may of course depend on the old and the new value of the first variable, and on current values of other components of the state).

These causal approaches have become popular recently, but one should realize that similar methods have been proposed earlier by Cordier and Siegel [4, 5] and by Brewka and Hertzberg [3].

However, the proposed causal approaches also have some problems. Their fundamental limitation is that they assume the explicitly stated effects of the action (the effects mentioned in the action law) to be the first ones in a causal chain, and that indirect effects are downstream in that chain. For a counterexample, consider the case of a lamp la which is turned on independently by two different switches: $la \leftrightarrow sw_1 \vee sw_2$. la, sw_1, and sw_2 are the only state variables, and there is no other domain constraint. The action of "turning on the lamp" is defined simply by stating that la becomes true. Presumably, ramification shall allow us to conclude that either sw_1 or sw_2 has become true as well. However, a causal approach will observe that the causal link is from the switches to the lamp, and not vice versa, so it will not deal correctly with this example.

We submit that this is not an odd example: quite generally, it is desirable to characterize actions in terms of their ultimate and main effects, and not in terms of the specifics of how to perform them. Reasoning from main effect to instrumental operations, and from there to other effects of the instrumental operations is therefore an important part of the ramification problem. The underlying semantics that is proposed in the next section is intended to capture that intuition.

It may be argued that the causational approaches can be used in such a way that the formal causation goes in the direction that is opposite to actual causation. However, in this case the approach loses a lot of its intuitive strength, and the need for a formal validation against an underlying semantics becomes even more acute than before.

5 Transition cascade semantics

In the case of partial-state transition semantics, there is a function $\mathbf{X}$ which identifies all those state variables that are affected for a given combination of action and starting state. The intuition is that action laws shall specify the new values of all those state variables, and that any other state variable is understood to keep its old value.

In the case of ramification, one wishes $\mathbf{X}$ to be even more restrictive, so that it only identifies *some of* the state variables that change, whereby the action law can be made even more compact. Changes in other state variables are to be derived from more general information.

In order to capture these intuitions, it is natural to use an underlying semantics that has some notion of direct versus indirect change. We propose to use the following *transition cascade semantics*, which is obtained through an additional modification of the state transition semantics described above. The domain $\mathcal{R}$ of world states and the domain $\mathcal{E}$ of actions are defined like before. A binary non-reflexive *causal transition relation* C between states is introduced; if $C(r, r')$ then r' is said to be the *successor* of r. A state r is called *stable* iff it does not have any successor. The set $\mathcal{R}_c$ of *admitted states* is chosen as a subset of $\mathcal{R}$ all of whose members are stable with respect to

C. (It may or may not be chosen to contain all the stable states). The identification of a set of admitted states is equivalent to the introduction of domain constraints.

Furthermore, we assume an *action invocation relation* $D(E, r, r')$ where $E \in \mathcal{E}$ is an action, r is the state where the action E is invoked, and r' is the new state where the instrumental part of the action has been executed. For every E and r there must be at least one r' such that $D(E, r, r')$, that is, every action is always invokable.

An *action system* is a tuple $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$, and is intended as follows. For any state $r \in \mathcal{R}_c$, consider a state r_1 such that $D(E, r, r_1)$, and a sequence of states

$$r_1, r_2, ... r_k$$

where $C(r_i, r_{i+1})$ for every i between 1 and $k - 1$, and r_k is a stable state. Such a sequence is called a *transition chain*, and r_k is considered as a *result state* of the action E in the sitution r.

We intend, of course, that it shall be possible to characterize the result states r_k in terms of E and r, but without referring explicitly to the details of the intermediate states. The following assumptions (or something similar) are needed in order to make this work as intended.

Definition 1 If three states r, r_i, and r_{i+1} are given, we say that the pair r_i, r_{i+1} **respects** r iff

$$r_i(f) \neq r_{i+1}(f) \Rightarrow r_i(f) = r(f),$$

for any state variable f that is defined in $\mathcal{R}$. Then, an action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ is said to be **respectful** iff, for every $r \in \mathcal{R}_c$, every $E \in \mathcal{E}$, r is respected by every pair r_i, r_{i+1} in every transition chain, and the last element of the chain is a member of $\mathcal{R}_c$. $\Box$

This condition amounts to a "write-once" or "single-assignment" property: if the action E is performed from state r, the world may go through a sequence of states, but in each step from one state to the next, there cannot be changes in state variables which have already changed previously in the sequence, nor can there be any additional change in a state variable that changed already in the invocation transition from r to r_1.

As a consequence of these definitions, we obtain that if $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ is a respectful action system, if $r \in \mathcal{R}_c$, $E \in \mathcal{E}$, and $D(E, r, r')$ holds for some r', then all the transition chains that emerge from (E, r) are finite and cycle-free. The set of result states will be denoted $\mathbf{N}(E, r)$, using the same notation as in the state-transition semantics. Thus, the main next-state function $\mathbf{N}(E, r)$ which was previously taken as given, is here taken as derived from the elementary relations D and C in the action system. It also follows from the definition that $\mathbf{N}(E, r) \neq \emptyset$.

Respectful action systems are intended to capture the basic intuitions of actions with indirect effects which are due to causation, as follows. Suppose the world is in a stable state r, and an action E is invoked. The immediate effect of this is to set the world in a new state, which is not necessarily stable. If it is not, then one allows the world to go through the necessary sequence of state transitions, until it reaches a stable state. *That whole sequence* of state transitions is together viewed as the action, and the resulting admitted state is viewed as the result state of the action.

The assumption of sequentiality of actions within a chronicle is then taken to mean that the agent or "ego" abstains from invoking any new action before the previous action has run to completion. The exact interpretation of the state transitions may vary. The most obvious is to think of them as spontaneous transitions in the physical world, so that the world actually traverses all the states in the sequence without external intervention. In this case, it is natural to derive a trajectory semantics from the multiple transition semantics. One may also think of the transitions as more or less logical conclusions about a world where all changes happen (or are modelled as happening) concurrently. Finally, one may include the case where additional interventions by the agent (viewed as a controller) are necessary along the path from the starting state to the result state of the action.

For example, in the lamp scenario, if the state variables represent the first switch, the second switch, and the lamp, respectively, the invocation relation D will contain argument sequences such as (E_1, FFF, TFF) for turning on the first switch; the causal transition relation will contain the argument sequence (TFF, TFT) expressing that if the first switch is on, then the lamp goes on as well. The action E_2 which is instrumentally defined as turning on either of the two switches is true for both $D(E_2, FFF, TFF)$ and $D(E_2, FFF, FTF)$.

The idea with ramification, now, is that one does not wish to specify *all* the effects of an action in the action law; one wishes to restrict the action law to the most important effects, and other effects should be deducible if needed. In other words, the action law shall not specify all components of the result state, and not even all the components that change, but only some of them. Other changes are deduced using information about the relations C and D.

Notice that it is possible, but not necessary, to let the action law characterize the invocation relation D. Alternatively, the action law may characterize some of the later changes along the chain, namely, if the indirect changes are considered to be the most important ones in terms of describing the action. For example, the action law for E_2 in the second-previous paragraph may specify the change in the state variable indicating that "the lamp is lit", leaving the "upstream" change of one or the other of the switches to be inferred. In general, the non-explicit effects of the action may be a combination of changes that are causationally "downstream" as well as "upstream" of the explicit ones.

6 Minimization approaches to ramification

With this underlying semantics, we shall begin the analysis of some of the approaches to ramification which have been proposed in recent years. We first address methods for ramification using minimization and categorization.

6.1 Measure systems

The following notation is introduced in order to obtain a common framework within which various minimization-oriented methods can be expressed. A *measure system* is a threetuple

$$\langle \mathcal{Q}, d, \prec \rangle, \tag{1}$$

where $\mathcal{Q}$ is a domain of *quantities*, d is a *difference function* $\mathcal{R} \times \mathcal{R} \to \mathcal{Q}$, $\prec$ is a strict partial order on $\mathcal{Q}$, and where $q \not\prec d(r,r)$ for all r and q. Notice that even $d(r,r') = d(r',r)$ is not required, which is why we call this a "difference" and not a "distance". We shall say that r' is *less remote* than r'' from r if $d(r,r') \prec d(r,r'')$.

Now, let $\mathbf{N}$, $\mathbf{X}$, and a measure system $m = \langle \mathcal{Q}, d, \prec \rangle$ be given, which also means that $\mathbf{N}'$ is defined. Then $\mathbf{N}'_m$ shall be the function such that

$$\mathbf{N}'_m(E,r) \subseteq \mathbf{N}'(E,r) \cap \mathcal{R}_c \tag{2}$$

and where $\mathbf{N}'_m(E,r)$ consists exactly of those members of $\mathbf{N}'(E,r) \cap \mathcal{R}_c$ which are *minimally remote from r according to the measure system m*. Formally, for any member r' of $\mathbf{N}'_m(E,r)$, there must not be any other member r'' of

$\mathbf{N}'(E,r) \cap \mathcal{R}_c$ which is less remote than r' from r, and any $r' \in \mathbf{N}'(E,r) \cap \mathcal{R}_c$ having that property must be a member of $\mathbf{N}'_m(E,r)$.

This concept of measure system provides a framework within which it is straightforward to define a number of minimization-of-change methods. For example, the simple case where one partitions state variables into "occluded" and "unoccluded" ones, and minimizes the set of changes in the unoccluded state variables, is obtained by defining $\mathcal{Q}$ as the power set of the set of state variables, $d(r,r')$ as the set of unoccluded state variables where r and r' assign different values, and $\prec$ as the strict subset relation.

A particularly simple case is *strict ramification*, which we shall denote **SR**, where no preference is imposed, and which is written

$$\mathbf{N}'_{SR}(E,r) = \mathbf{N}'(E,r) \cap \mathcal{R}_c. \tag{3}$$

It is easily verified that the approximation $\mathbf{N}'_m$ for arbitrary m satisfies all the Katsuno-Mendelzon postulates [15] **U1** through **U8**, except **U2**. The appropriateness of **U2** has already been cast in doubt by del Val and Shoham in [8], where they observed that **U2** is not well chosen in the case of nondeterministic actions.

6.2 Comparisons of measure systems

With the given definitions, a measure system m is *correct* for a main next-state function $\mathbf{N}$ and occlusion function $\mathbf{X}$ iff $\mathbf{N}(E,r) = \mathbf{N}'_m(E,r)$ for every action E and every $r \in \mathcal{R}_c$. The *range of applicability* of a measure system m is the set of pairs $\langle \mathbf{N}, \mathbf{X} \rangle$ where m is correct.

Notice that in the present section we do not relate the minimization method to the underlying semantics, since we take the function $\mathbf{N}(E,r)$ as given. In the next section, where we start using the underlying semantics, $\mathbf{N}(E,r)$ it will instead be considered as defined along the lines that were introduced in section 5.

If one can not find a correct measure system for a given application, then presumably one prefers a measure system that is *sound* in the sense that it does not allow any unwarranted conclusions. In our framework, the measure system m is sound for a main next-state function $\mathbf{N}$ and occlusion function $\mathbf{X}$ iff $\mathbf{N}(E,r) \subseteq \mathbf{N}'_m(E,r)$ for every action E and every $r \in \mathcal{R}_c$.

It is evident that no measure system is correct for all $\mathbf{N}$, and that straight ramification is the only measure system that is sound for all $\mathbf{N}$. (The latter follows by choosing $\mathbf{N}(E,r)$ so that it equals $\mathbf{N}'(E,r) \cap \mathcal{R}_c$, so that no minimization is desired). We have compared different minimization methods with respect to how restrictive they are: a minimization method based on measure system m is *more restrictive* than one based on measure system m' iff $\mathbf{N}'_m(E,r) \subseteq \mathbf{N}'_{m'}(E,r)$ for all E and for all $r \in \mathcal{R}_c$. Thus, a more restrictive method is one which obtains a smaller set of preferred models, and thereby a larger set of conclusions, and a smaller set of pairs $\langle \mathbf{N}, \mathbf{X} \rangle$ for which it is sound. Clearly, straight ramification is the unique least restrictive method. We write $m < m'$ when m is more restrictive than m'.

The following methods are of particular interest. In all cases, the proposed categorization is allowed to depend on the choice of E and r; it does not have to be fixed relative to them.

MSC: Minimal secondary change. State variables are divided into two categories: *occluded* ones specified by $\mathbf{X}$, and *secondary* ones which are the rest. Action laws specify the possible new values for occluded state variables. For given E and r, one selects those members of $\mathbf{N}'(E,r) \cap \mathcal{R}_c$ which minimize the set of secondary state variables that change compared to r. This is the method proposed by del Val and Shoham in [8].

MSCC: Minimal secondary change with changeset-partitioning. Similar to MSC for deterministic actions, but differs for nondeterministic actions. In MSCC, one divides $\mathbf{N}'(E,r) \cap \mathcal{R}_c$ into partitions each consisting of those proposed new states which are equal with respect to occluded state variables. In each of the partitions, one selects the subset of those members which are minimally remote from r, using the same criterion as for MSC. Then one forms the union of those closest subsets. It follows at once that MSCC is less restrictive than MSC, although more restrictive than straight ramification.

The intuition behind this alternative is as follows. Consider an action which is nondeterministic already from the point of view of the action law, that is, the action law specifies several alternative outcomes of the action. Then one would like to minimize changes for each of the outcomes separately. For example, the action of playing a game on a gambling machine may result in different states in the machine; let us say just *win* and *lose* for simplicity. Suppose that the action law only specifies this outcome, and there is a static domain constraint specifying that *win* is accompanied by *money-falling-out*. Then MSC will always prefer *lose* over *win* because it minimizes the set of side-effects. MSCC, by comparison, considers *win* and *lose* to be equally preferred. Changeset-partitioning was first proposed by Sandewall in [31].

MRC: Minimal remanent change. State variables are divided into three disjoint categories: occluded, remanent, and dependent ones. The occluded state variables are those whose new values are specified in the action laws, like before. Changes in remanent state variables are minimized, whereas changes in dependent state variables are disregarded for the purpose of minimization. This is essentially the method proposed by Kartha and Lifschitz in [14], if one identifies "occluded" with "frame released", "remanent" with "frame unreleased", and "dependent" with "non-frame" state variables.

MRCC: Minimal remanent change with changeset-partitioning. Uses the same threeway partitioning as in MRC, but minimizes changes separately for each outcome in occluded state variables, like in MSCC.

The intuition behind the three-way partitioning in MRC and MRCC is as follows: even if one chooses to divide the state variables into two categories, one where change is minimized and one where it is not, there may still be two distinct reasons why a state variable is in the latter category. On one hand, there are those state variables where the outcome or set of possible outcomes is already specified in the action law, so that further minimization is not needed and often not desired (because it may remove some of the intended outcomes of nondeterministic actions). On the other hand, there are those state variables which are functionally dependent on some of the others, and where the dependency is specified by the static domain constraints. In our formulation of the problem, the first subgroup is indicated by $\mathbf{X}$, since $\mathbf{X}$ is used in two ways: it defines $\mathbf{N}'$ from $\mathbf{N}$, and it is part of the definition of minimization. The other subgroup can not be included in $\mathbf{X}$ – the formal account in [36] introduces the required notation.

One might think that since the dependent state variables are functionally determined by the others, it would not hurt to bring them along in the minimization. The problem is, as was shown already by Winslett in the 'Aunt Agatha' scenario, that there may be undesired tradeoffs so that an unintended model containing a change in a remanent state variable may be be accepted because it contains a smaller set of changes in dependent variables. In such scenarios it is necessary to exclude dependent state variables from minimization.

Finally, the list should also include the classical method:

MC: Minimization of change. No partitioning of the state variables. One selects the subset of $\mathbf{N}'(E, r) \cap \mathcal{R}_c$ consisting of those members minimizing the set of state variables where the value differs from r. This is the method originally proposed by Winslett [39].

The formal definitions of these methods in terms of measure systems are straightforward, and can be obtained from [36].

6.3 Relative range results

Recalling that $m \leq m'$ means that m is more restrictive than m', and writing **SR** for straight ramification, we have proved the following relative range results:

Theorem 2

$$\mathbf{MC} \leq \mathbf{MSCC} \leq \mathbf{SR} \tag{4}$$

$$\mathbf{MSC} \leq \mathbf{MSCC} \leq \mathbf{SR} \tag{5}$$

$$\mathbf{MRC} \leq \mathbf{MRCC} \leq \mathbf{SR} \tag{6}$$

The proofs are found in our technical report, which is also available on-line [36]. There does not seem to be any similar result relating e.g. **MSCC** and **MRCC**.

These results suggest at once that, as long as the knowledge of the application domain does not give better evidence, one ought to prefer methods with changeset-partitioning over methods without, that is, MSCC over MSC and MRCC over MRC. This is because one will then obtain a method with a larger range of sound applicability. In other words, although possibly one will fail to obtain some intended conclusions, at least there is a smaller possibility that the system actually infers non-intended conclusions.

Preferably, however, one should be able to use more specific information about the application domain at hand, and obtain an entailment method which is more restrictive but in a way that retains soundness, and ideally achieves completeness.

7 Assessments of entailment methods based on minimization of change

After having related the various minimization methods to each other, we proceed now to the analysis of how they relate to the underlying semantics. The following results are preliminary ones.

The general problem is viewed as follows. Let the following be given:

- A respectful action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$. Its corresponding main next-state function will be written $\mathbf{N}$, as before.
- An occlusion function, $\mathbf{X}$.

- A measure system $m = \langle \mathcal{Q}, d, \prec \rangle$.

Informally, it is still the case that $\mathbf{X}$ indicates what are the state variables whose new values are indicated in the action law. The functions $\mathbf{N}$ and $\mathbf{X}$ together define the stated next-state function $\mathbf{N}'$, as described above. The measure system is said to be *correct* for the action system and the occlusion function iff

$$\mathbf{N}'_m(E, r) = \mathbf{N}(E, r)$$

for every E and r. Likewise, the measure system is said to be *sound* for the action system and the occlusion function iff

$$\mathbf{N}'_m(E, r) \supseteq \mathbf{N}(E, r)$$

for every E and r. Notice how the entities in this equation have been derived: $\mathbf{N}$ from the action system alone, $\mathbf{N}'$ from $\mathbf{N}$ and $\mathbf{X}$, and $\mathbf{N}'_m$ from $\mathbf{N}'$ and m.

In particular, given one of the specific measure systems and entailment methods that were described above, for example MRCC, a particular partitioning of the state variables is said to be *correct* for the action system, the occlusion function, and the method at hand iff the corresponding measure system is correct for them. Soundness is defined similarly.

The corresponding *assessment problem* for an entailment method is defined as the following question: what are the requirements on the combination of a respectful action system, an occlusion function, and a partitioning whereby the given entailment method is correct (or sound)?

A broader question, which also is of interest, is for what classes of action systems and occlusion functions is it possible to find a correct measure system at all.

7.1 Occlusion of invocation-changed state variables

It turns out that fairly strong restrictions are required in order to get any assessment results at all, and that it is quite easy to construct cases where apparently there is no correct partitioning. The following defines a class of action systems and occlusion functions where the soundness of MSCC has been proved. We consider separately the two cases where action laws describe invocations and where they describe downstream effects.

Definition 3 A respectful action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ is said to have **fixed control** iff there is some subset $\mathcal{F}_c$ of the set $\mathcal{F}$ of state variables, such that

$$D(E, r, r') \Rightarrow dom(r - r') \subseteq \mathcal{F}_c$$

and

$$C(r, r') \Rightarrow dom(r - r') \cap \mathcal{F}_c = \varnothing$$

The set $\mathcal{F}_c$ is called the *control set.* □

Here and henceforth the following conventions are used: $-$ denotes set subtraction, mappings are viewed as sets of argument/value pairs, and $dom(r)$ maps a mapping (in this case, a partial state) to its domain.

Members of the control set will be called *controlled* state variables; the others will be called *obtained* state variables.

This definition means that invocation transitions can only influence state variables in the control set, and causal transitions can never influence state variables in the control set. Thielscher's relay example in [37] is an example of a system which does not have fixed control.

Definition 4 Given a state domain $\mathcal{R}$, a state r'' is said to be a **hybrid** between two states r and r' *with respect to* a set $F \subseteq \mathcal{F}$ of state variables iff

(1) for every state variable f, $r''(f) = r(f) \vee r''(f) = r'(f)$, and

(2) if $f \in F$ then $r'(f) = r''(f)$, and

(3) r'' is different from both r and r'. □

Definition 5 A respectful action system is said to be **hybrid free** with respect to a set F of state variables iff no two members r and r' of $\mathcal{R}_c$ have a hybrid with respect to F which is also a member of $\mathcal{R}_c$. □

It follows that if an action system is hybrid free with respect to a set F, then it is hybrid free with respect to any superset of F. (Any hybrid with respect to the superset would be a hybrid with respect to F). We immediately obtain:

Theorem 6 **MSCC** is sound for the combination of a respectful action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ and an occlusion function $\mathbf{X}$ if the following conditions are satisfied:

1. The action system has fixed control with a control set $\mathcal{F}_c$
2. The action system is hybrid free with respect to $\mathcal{F}_c$
3. If $D(E, r, r')$ then $dom(r - r') \subseteq \mathbf{X}(E, r)$

Proof We are to prove that $\mathbf{N}'_{MSCC}(E,r) \supseteq \mathbf{N}(E,r)$. Suppose the opposite is the case, that is, for some E and r there is some $r^* \in \mathbf{N}(E,r)$ which is not a member of $\mathbf{N}'_{MSCC}(E,r)$. Since it is a member of $\mathbf{N}(E,r)$, r^* must be a member of $\mathbf{N}'(E,r)$ and of $\mathcal{R}_c$, so its non-membership of $\mathbf{N}'_{MSCC}(E,r)$ must be due to the existence of some $r' \in \mathbf{N}'(E,r) \cap \mathcal{R}_c$ which is less remote than r^* from r.

By the definition of MSCC, r' must agree with r^* in all occluded state variables. By assumptions 1 and 3, r^* agrees with r in all non-occluded controlled state variables, so r' must do so as well. Therefore, r^* and r' agree in all state variables in $\mathcal{F}_c$. But r' differs from r in a subset of the state variables where r^* differs from r. Therefore, either r' is a hybrid between r and r^* with respect to $\mathcal{F}_c$, or $r' = r$. The former case is inconsistent with assumption 2 given that $r' \in \mathcal{R}_c$. The latter, bizarre case requires that r and r^* are equal in all controlled state variables, which means that $N(E,r) = \{r\}$ and $r^* = r$, making it impossible for any r' to be less remote from r than r^* is. □

7.2 Occlusion of non-invoked state variables

We proceed now to the case where action laws characterize other changes of state variables than those involved in the invocation.

Definition 7 A respectful action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ is **unary controlled** iff $D(E', r, r')$ implies that $dom(r - r')$ is a singleton, for any E, r, and r'. □

Note that this property is realistic for low-level actions in engineered systems, where controlled state variables represent the actuators.

Definition 8 Consider the combination of a respectful action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ having fixed control, and an occlusion function $\mathbf{X}$. A state $r' \in \mathcal{R}_c$ is called a **shortcut** for (E, r) in this combination, where $r \in \mathcal{R}_c$, if and only if there is some $r'' \in N(E,r)$ such that

(1) for every state variable f, $r'(f) = r(f) \vee r'(f) = r''(f)$, and

(2) r' equals r in all controlled state variables, and

(3) r' equals r'' for all state variables in $\mathbf{X}(E,r)$. □

Example 9 The following is an example of a shortcut. Consider an electric circuit consisting of a number of switches, lamps, and relays, where the switches and only them are controlled. In the current state r of the circuit, changing a particular switch s will turn on a lamp l which is presently off; there may also be other effects. Let r'' be the resulting state. An action E is defined so that $\mathbf{X}(E, r) = \{l\}$ and r'' is a member of $\mathbf{N}(E, r)$, that is, the action is defined in terms of the effect of turning on the lamp, and not in terms of changing the switch. However, there is also a sequence of other actions where different switches are changed in succession, and which ends up in a state r' where all switches are in the same position as in the present state r, the lamp l is on anyway, and only a subset of the changes from r to r'' occur in r'. $\Box$

It would seem that the property of having a shortcut is not a particularly natural one for an action system to have. However, there is an important special case, namely for the combination of an action system and an occlusion function where an action possibly has a null stated effect, that is, there is some $r'' \in N(E, r)$ for which r and r'' are equal for all occluded state variables. In this case, r itself is a shortcut for (E, r).

We immediately obtain:

Theorem 10 **MSCC** is sound for the combination of a respectful action system $\langle \mathcal{R}, \mathcal{E}, C, \mathcal{R}_c, D \rangle$ and an occlusion function $\mathbf{X}$ if the following conditions are satisfied:

1. The action system has fixed control with a control set $\mathcal{F}_c$.
2. The action system is hybrid free with respect to $\mathcal{F}_c$.
3. $\mathbf{X}(E, r)$ is disjoint from $\mathcal{F}_c$ for every E and $r \in \mathcal{R}_c$.
4. The action system is unary controlled.
5. The combination of the action system and the occlusion function does not have any shortcut.

Proof We are to prove that $\mathbf{N}'_{MSCC}(E, r) \supseteq \mathbf{N}(E, r)$. Suppose the opposite is the case, that is, for some E and r there is some $r^* \in \mathbf{N}(E, r)$ which is not a member of $\mathbf{N}'_{MSCC}(E, r)$. Since it is a member of $\mathbf{N}(E, r)$, r^* must be a member of $\mathbf{N}'(E, r)$ and of $\mathcal{R}_c$, so its non-membership of $\mathbf{N}'_{MSCC}(E, r)$ must be due to the existence of some $r' \in \mathbf{N}'(E, r) \cap \mathcal{R}_c$ which is less remote than r^* from r.

Consider the causation chain leading up to r^*, and its first element r_1. Because of the unary control property, r and r_1 will differ in exactly one state variable; this state variable is a controlled one because of assumption 1, and it is not occluded because of assumption 3. There are two possibilities with respect to r':

(1) r' equals r_1 in all controlled state variables. Then r' is a hybrid between r and r^*, violating assumption 2.

(2) r' equals r in all controlled state variables. Since r' is a member of $\mathbf{N}'(E, r)$, it must agree with some member r'' of $\mathbf{N}(E, r)$ in all occluded state variables. Therefore, r' is a shortcut for (E, r), which violates assumption 5. $\Box$

The results of these theorems do not apply in general for MSC because of interference between primary outcomes of the action. Generalization of these results to MRCC encounters technical complications in the generalization of the "hybrid" concept.

7.3 Comparison with correctness criteria for strict inertia

The assessments for entailment methods for strict inertia in [34] made use of the concept of minimal-change-compatible, which in the present context may be equivalently defined as follows:

Definition 11 An action system is **minimal-change-compatible** iff it is not the case, for any r, E, r', and r'', that

$$r' \in \mathbf{N}(E, r)$$

$$r'' \in \mathbf{N}(E, r)$$

and r'' is a hybrid between r and r' with respect to $\varnothing$. $\Box$

We showed in [34] that PCMF (prototypical minimization of change with filtering) is correct for all minimal-change-compatible chronicles with single-timestep actions. We observe that if an action system is hybrid free with respect to $\varnothing$, then it is both minimal-change-compatible and hybrid free with respect to an arbitrary set F.

7.4 Examples

The soundness property means that the entailment method MSCC does not miss any intended models, but it still allows for the possibility of including some unintended models. The following two examples illustrate how unintended models may be included in $\mathbf{N}'_{MSCC}(E, r)$.

Example 12 For the first example, suppose $\mathcal{R}$ is selected as $\{T, F\} \times \{1, 2\} \times \{1, 2\}$, where the first component is controlled and always occluded, and the other two are secondary. Suppose further that the causal transition function C is true for the following pairs:

$$C(T11, T12)$$
$$C(T22, T21)$$
$$C(F12, F11)$$
$$C(F21, F22)$$

and that $T12, T21, F11, F22$ are in $\mathcal{R}_c$. The invocation relation D is supposed to contain

$$D(E_1, F11, T11)$$
$$D(E_1, F22, T22)$$

as well as some suitable transitions when the controlled state variable is initially T. It is clear that $\mathbf{N}(E_1, F11) = \{T12\}$ and $\mathbf{N}'_{MSCC}(E_1, F11) = \{T12, T21\}$, so soundness but not correctness is obtained. $\Box$

The problem here, of course, is that the entailment method MRCC is based on a notion of minimal change which does not use any of the causal transition information from C.

Example 13 The following example illustrates another possibility for the introduction of unintended models. Suppose now that $\mathcal{R}$ is selected as $\{T, F\} \times \{T, F\} \times \{1, 2\}$, where the first two components are controlled and the third one is obtained. Suppose further that the transition function C satisfies

$$C(TF1, TF2),$$

and that $FF1, FT1, TF2, TT1$ are in $\mathcal{R}_c$. The invocation relation D is supposed to contain

$$D(E_1, FF1, TF1)$$
$$D(E_1, FT1, TT1)$$

as well as some suitable invocation transitions for the other members of $\mathcal{R}_c$. Then $\mathbf{N}(E_1, FF1) = \{TF2\}$. In this case $\mathbf{X}(E_1, FF1)$ needs only contain the first state variable, so both the second and the third state variable can be selected as secondary. Using MSCC with that assumption one obtains $\mathbf{N}'_m(E_1, FF1) = \{TF2, TT1\}$, so again soundness but not correctness is obtained. $\Box$

In this case, the unintended model entered because $\mathbf{X}(E, r)$ is only required to be a subset of the controlled state variables. This opens the door for admitting some model which satisfies the action laws with respect to occluded variables, but which also includes some change in a controlled, non-occluded variable (in this case, the second state variable), provided that it thereby obtains a smaller set of changes in remanent state variables.

The source of incompleteness that was illustrated here can be removed by requiring that all controlled state variables must be occluded, that is, their new value must be stated in the action law. This means effectively that the action law must not merely specify the new values for some state variables that change their value, but also specify explicitly and un-overridably that some state variables do not change their value in the current action. An additional possibility is to introduce a fourth category for those state variables which *must not* change as the result of an action.

8 Causal approaches to ramification

We proceed now to a first account and analysis of so-called causal approaches to ramification. In view of the causation orientation of the underlying semantics that is being used here, one would suppose that causal approaches should enjoy fewer restrictions and a broader range of applicability than the minimization oriented approaches.

Several proposals for causation-oriented methods have been published, first from European research [4, 3, 5] and then from work done in America [38, 24, 20]. Unfortunately, they all use fairly specialized logics, which adds to the difficulties of analyzing them. As a first step, we shall therefore define our own variant of these causal approaches, called **influence-restricted secondary change**, abbreviated **IRSC**, which will then be used for the assessment. IRSC is similar in particular to the approach of Thielscher [38], but not equivalent to it.

The entailment methods MRC and MRCC assumed a partitioning of the state variables into three categories: occluded, remanent, and dependent ones. For IRSC, we only retain the distinction between occluded and secondary state variables, like in MSC and MSCC. We also dispense with all minimization. The occlusion function $\mathbf{X}$ is used for characterizing the choice of occluded state variables.

Instead of minimization, IRSC uses an *influence relation* $\mathbf{I}(f, f', r)$ which says that f potentially influences f' in the causal transition chains that emerge from the state r. This relation is a strict partial order with respect to its first two arguments, and we write its non-strict transitive closure as $\mathbf{I}^*$. If F is a set of state variables, then $\mathbf{I}(F,r)$ denotes the superset of F which also contains all state variables f' which satisfy $\mathbf{I}^*(f,f',r)$ for some f in F.

The set of selected new states according to the IRSC method, using an influence relation $\mathbf{I}$, will be written $\mathbf{N}'_{\mathbf{I}}(E,r)$; it is defined as follows:

$$\mathbf{N}'_{\mathbf{I}}(E,r) = \{r' \in \mathbf{N}'(E,r) \cap \mathcal{R}_c \mid dom(r'-r) \subseteq \mathbf{I}(\mathbf{X}(E,r),r)\},$$

with the same usage for $-$ and dom as before. This definition says that $\mathbf{N}'_{\mathbf{I}}(E,r)$ is the set of all $r' \in \mathbf{N}'(E,r) \cap \mathcal{R}_c$ which only differ from r in occluded state variables (whose new values are explicitly stated in the action law) and in state variables that are reachable by $\mathbf{I}^*$ from the occluded ones.

As usual, the assessment problem is as follows: given a respectful action system, an occlusion function $\mathbf{X}$, and an influence relation $\mathbf{I}$, to identify restrictions on these structures which guarantee that $\mathbf{N}(E,r) = \mathbf{N}'_{\mathbf{I}}(E,r)$ for all E and r. In particular, we are interested in cases where $\mathbf{X}$ and $\mathbf{I}$ have been derived in a systematic way from the action system.

The following definition captures an obvious way of constructing $\mathbf{I}(f,f',r)$ from the action system and $\mathbf{X}$:

Definition 14 A combination of a respectful action system, an occlusion function $\mathbf{X}$, and an influence relation $\mathbf{I}$ is **invocation based** iff it satisfies the following two conditions for an arbitrary action symbol E and arbitrary state $r \in \mathcal{R}_c$:

1. If $D(E,r,r')$, then $dom(r'-r) \subseteq \mathbf{X}(E,r)$.
2. Consider any transition chain

$$r_1, r_2, \ldots r_k$$

where $D(E,r,r_1)$, $C(r_i,r_{i+1})$, and $r_k \in \mathcal{R}_c$, as usual. Write $r_0 = r$. Whenever, for some $i \geq 0$, a state variable f has different values in r_i and r_{i+1}, and a state variable f' has different values in r_{i+1} and r_{i+2}, it must be true that $\mathbf{I}(f,f',r)$. $\Box$

We immediately obtain:

Theorem 15 **IRSC** is sound if used for an invocation based combination of an action system, an occlusion function, and an influence relation.

Proof Let r_k be an arbitrary member of $\mathbf{N}(E, r)$, which is determined as such by the transition chain

$$r_1, r_2, ...r_k.$$

Since the combination is invocation based, it follows that $dom(r_1 - r) \subseteq \mathbf{X}(E, r)$, and then that

$$dom(r_k - r) \subseteq \mathbf{I}(\mathbf{X}(E, r), r).$$

Therefore $r_k \in \mathbf{N}'_I(E, r)$, which means that

$$\mathbf{N}(E, r) \subseteq \mathbf{N}'_I(E, r).$$

This concludes the proof. $\Box$

Thus, unlike the result for the minimization-based approaches, IRSC does not require a strict separation between controlled and obtained state variables. It is possible for a state variable to change in the invocation for some pair (E, r), and to change in the causation chain for some other (E, r). This is typical of the relay example that was put forward by Thielscher in [38].

However, there are still two problems with IRSC. First, this result only proves soundness and not completeness, so in gross terms it is no better than MRCC that was analyzed above. The other problem is that it IRSC in a invocation based combination can only work if the action law expresses the invoked changes, that is, those obtained by the D relation. It does not work at all if action laws describe actions in terms of their indirect changes, as in the two-switch lamp example.

There are several possiblities for how one would modify IRSC or its usage in order to eliminate this restriction. As usual, the difficulty is not in proposing the method, but in analyzing its properties. This is a topic for further research.

9 Conclusion

As background, we have reviewed current approaches to underlying semantics and range-of-applicability results for logics of actions and change with strict inertia. After that, we have defined and proposed a causation-oriented *transition cascade semantics* for ramification, and we have reported relative range results for minimization-based approaches. Finally, based on the same underlying semantics, we have reported on soundness results both for some of the minimization-based approaches and for a prototypical causation oriented approach to ramification.

The general observation from this study is that one needs to impose fairly strong restrictions and still one obtains relatively weak results (only soundness, no semantic completeness). This suggests that contemporary approaches to ramification are not sufficiently well motivated, and that additional, application-specific information should be used to modify the present methods or to influence how they are used.

References

[1] Andrew B. Baker. Nonmonotonic reasoning in the framework of the situation calculus. *Artificial Intelligence*, 49:5–23, 1991.

[2] Chitta Baral and Michael Gelfond. Representing concurrent actions in extended logic programming. In *International Joint Conference on Artificial Intelligence*, pages 866–871, 1993.

[3] Gerhard Brewka and Joachim Hertzberg. How to do things with worlds: On formalizing actions and plans. *Journal of Logic and Computation*, 1993.

[4] Marie-Odile Cordier and Pierre Siégel. A temporal revision model for reasoning about world change. In *International Conference on Knowledge Representation and Reasoning*, pages 732–739, 1992.

[5] Marie-Odile Cordier and Pierre Siégel. Prioritized transitions for updates. In C. Froidevaux and J. Kohlas, editors, *Proc. European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 142–151. Springer Verlag, 1995.

[6] Ernest Davis. *Representation of Commonsense Knowledge.* Morgan-Kaufmann, 1990.

[7] Alvaro del Val and Yoav Shoham. Deriving properties of belief update from theories of actions. In *National Conference on Artificial Intelligence*, pages 584–589, 1992.

[8] Alvaro del Val and Yoav Shoham. Deriving properties of belief update from theories of actions (II). In *International Joint Conference on Artificial Intelligence*, pages 732–737, 1993.

[9] Richard E. Fikes and Nils J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.

[10] Michael Gelfond and Vladimir Lifschitz. Representing actions in extended logic programs. In *International Conference on Logic Programming*, pages 559–573, 1992.

[11] Matthew L. Ginsberg and David E. Smith. Reasoning about actions I: A possible worlds approach. In Frank Brown, editor, *Proceedings of the 1987 Workshop on the Frame Problem in Artificial Intelligence*, pages 233–258. Morgan Kaufmann, Lawrence, Kansas, 1987.

[12] G. Neelakantan Kartha. Soundness and completeness theorems for three formalizations of action. In *International Joint Conference on Artificial Intelligence*, pages 724–729, 1993.

[13] G. Neelakantan Kartha. On the range of applicability of baker's approach to the frame problem. In *National Conference on Artificial Intelligence*, 1996.

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

[15] Hirofumi Katsuno and Alberto O. Mendelzon. A unified view of propositional knowledge base update. In *International Joint Conference on Artificial Intelligence*, pages 1413–1419, 1989.

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

[17] Vladimir Lifschitz. Frames in the space of situations. *Artificial Intelligence*, 46:365–376, 1990.

[18] Vladimir Lifschitz. Toward a metatheory of action. In *International Conference on Knowledge Representation and Reasoning*, pages 376–386, 1991.

[19] Vladimir Lifschitz. Nested abnormality theories. *Artificial Intelligence*, 74:351–365, 1995.

[20] Fangzhen Lin. Embracing causality in specifying the indirect effects of actions. In *International Joint Conference on Artificial Intelligence*, 1995.

[21] Fangzhen Lin and Ray Reiter. State constraints revisited. *Journal of Logic and Computation*, 4(5), October 1994.

[22] Fangzhen Lin and Yoav Shoham. Provably correct theories of action (preliminary report). In *National Conference on Artificial Intelligence*, pages 349–354, 1991.

[23] Fangzhen Lin and Yoav Shoham. Concurrent actions in the situation calculus. In *National Conference on Artificial Intelligence*, pages 590–595, 1992.

[24] Norman McCain and Hudson Turner. A causal theory of ramifications and qualifications. In *International Joint Conference on Artificial Intelligence*, 1995.

[25] Edwin P.D. Pednault. ADL: Exploring the middle ground between STRIPS and the situation calculus. In *International Conference on Knowledge Representation and Reasoning*, pages 324–332, 1989.

[26] Pavlos Peppas and Wayne Wobcke. On the use of epistemic entrenchment in reasoning about action. In *European Conference on Artificial Intelligence*, pages 403–407, 1992.

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

[28] Earl Sacerdoti. The non-linear nature of plans. In *International Joint Conference on Artificial Intelligence*, pages 206–214, 1975.

[29] Erik Sandewall. Non-monotonic entailment for reasoning about time and action, part I: Sequential actions. Technical report, IDA, 1988.

[30] Erik Sandewall. Filter preferential entailment for the logic of action in almost continuous worlds. In *International Joint Conference on Artificial Intelligence*, pages 894–899, 1989.

[31] Erik Sandewall. Features and fluents. Review version of 1991. Technical Report LiTH-IDA-R-91-29, IDA, 1991.

[32] Erik Sandewall. Nonmonotonic temporal logics and autonomous agents: each contributes to the rigorous basis of the other. In O. Herzog et al, editor, *Grundlagen und Anwendungen der Künstlichen Intelligenz. Proc. German A.I. conference.* Springer Verlag, 1993.

[33] Erik Sandewall. The range of applicability of nonmonotonic logics for the inertia problem. In *International Joint Conference on Artificial Intelligence*, 1993.

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

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

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

[37] Michael Thielscher. An analysis of systematic approaches to reasoning about actions and change. In P. Jorrand, editor, *International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA)*, Sofia, Bulgaria, 1994. World Scientific Publishing Co.

[38] Michael Thielscher. Computing ramifications by postprocessing. In *International Joint Conference on Artificial Intelligence*, 1995.

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

[40] Choong-Ho Yi. Towards the assessment of logics for concurrent actions. In *AAAI 1995 Spring Symposium*, 1995.