Using Ranking Functions to Determine Plausible Action Histories

Aaron Hunter and James P. Delgrande
School of Computing Science
Faculty of Applied Sciences
Simon Fraser University
{amhunter, jim}@cs.sfu.ca

Abstract

We use ranking functions to reason about belief change following an alternating sequence of actions and observations. At each instant, an agent assigns a plausibility value to every action and every state; the most plausible world histories are obtained by minimizing the sum of these values. Since plausibility is given a quantitative rank, an agent is able to compare the plausibility of actions and observations. This allows action occurrences to be postulated or refuted in response to new observations. We demonstrate that our formalism is a generalization of our previous work on the interaction of revision and update.

1 Introduction

When reasoning about epistemic action effects, it is useful to draw a distinction between *ontic actions* and *epistemic actions*. Ontic actions are actions that change the state of the world, whereas epistemic actions are actions that change the beliefs of an agent without changing the world. Several formalisms have been proposed to represent action domains in which an agent may perform both ontic actions and epistemic actions [SPLL00; HLM04]. These formalisms have focused primarily on the epistemic effects of a single action. In this paper, we consider belief change in the context of alternating sequences of ontic and epistemic actions. The formalism that we introduce is a generalization of our work in [HD05].

Informally, we are interested in alternating sequences of updates and revisions of the form

$$K \diamond A_1 * O_1 \diamond \cdots \diamond A_n * O_n$$

where each A_i is an ontic action and each O_i is an observation represented by a set of possible worlds. We are particularly interested in the case where action histories and observation histories may both be fallible. In this context, it is necessary for an agent to have some means for resolving conflicts between observations and perceived action histories. For example, if no world in O_n is possible following the action sequence $A_1, \ldots, A_n$, then there are two options.

1. Reject O_n.
2. Accept O_n, and modify $A_1, \ldots, A_n$.

The first option intuitively corresponds to the case where O_n is less plausible than the action history, and the second option corresponds to the case where it is more plausible. In order to determine which option is preferable for a specific problem, an agent effectively needs to be able to compare the plausibility of O_n with the plausibility of each A_i. This kind of comparison is only possible if there is a single plausibility ranking over actions and observations.

In this paper, we propose that Spohn-style ranking functions can be used to define a flexible formalism for reasoning about belief change over alternating sequences of actions and observations. The idea is simply to give a subjective ranking of actions and observations at each point in time. By looking at this sequence of rankings, an agent is able to determine the most plausible world histories. We demonstrate the utility of our new formalism by example and by comparison with related formalisms. We also demonstrate that this is indeed a generalization of our previous work, and we suggest that this more general approach makes the role of our so-called interaction postulates more explicit.

2 Preliminaries

We introduce some terminology and formal machinery that is commonly used for reasoning about action effects [GL98]. We are interested in action domains that can be described by a set of fluent symbols $\mathbf{F}$ and a set of action symbols $\mathbf{A}$. Informally, fluent symbols represent properties of the world that may change in response to the execution of the actions in $\mathbf{A}$. Formally, the effects of actions are given by transition systems.

Definition 1 *A transition system is a pair* $\langle S, R\rangle$ *where* $S \subseteq 2^{\mathbf{F}}$ *and* $R \subseteq S \times \mathbf{A} \times S$.

A transition system is simply a directed graph where the nodes represent states and the edges are labeled with action symbols. In this paper, we assume that every action is always executable, so we restrict attention to transition systems where every state has an outgoing edge for each action symbol. We also restrict attention to actions with deterministic effects.

We define a *belief state* to be a set of interpretations over $\mathbf{F}$, informally the set of interpretations that an agent considers possible. An *observation* is also a set of interpretations.

The observation α is interpreted to provide evidence that the actual world is in α.

The process in which an agent changes their beliefs in response to a predicted change in the state of the world is called *belief update*. One of the standard approaches to belief update is given in [KM92], where a set of formulas is updated by another formula. By contrast, we define belief update operators that map a belief state and an action to a new belief state. In particular, a transition system defines a belief update operator as follows.

Definition 2 *Let $T = \langle S, R \rangle$ be a transition system. The update function $\diamond : 2^S \times \mathbf{A} \rightarrow 2^S$ is given by $\alpha \diamond A = \{s \mid \langle s', A, s \rangle \in R$ for some $s' \in \alpha\}$.*

The process in which an agent changes their beliefs in response to new information about a static world is called *belief revision*, and one of the standard approaches is the AGM approach [AGM85]. Again, we diverge slightly from the standard approach in that we do not deal with formulas; instead we think of revision as an operation in which a belief state and an observation are mapped to a new belief state.

3 Motivating Example

We briefly introduce a simple, commonsense example in which an agent needs to compare the plausibility of certain actions with the plausibility of observations. We will return to this example periodically as we introduce the formal machinery.

We consider a simple action domain involving four agents: Bob, Alice, Eve, and Trent. Bob places a chocolate chip cookie on his desk and then leaves the room; he believes that no one is likely to eat his cookie while he is gone. At time 1, Bob knows that Alice is at his desk. At time 2, Bob knows that Eve is at his desk. At time 3, Trent comes and tells Bob that a bite has been taken from the cookie on his desk.

Given the preceding information, Bob can draw three reasonable conclusions: Alice bit the cookie, Eve bit the cookie, or Trent gave him poor information. If Bob has no additional information about the world, then each conclusion is equally plausible. However, we suppose that Bob does have some additional information. In particular, suppose that Alice is a close friend of Bob and they have shared cookies in the past. Moreover, suppose that Bob believes that Trent is always honest. Bob's additional information about Alice and Trent provides a sufficient basis for determining which of the three possible conclusions is the most plausible.

Informally, at time 2, Bob believes that his cookie was unbitten at all earlier points in time. After Trent tells him the cookie is bitten, he must determine the most plausible world history consistent with this information. In this case, the most plausible solution is to conclude that Alice bit the cookie. Note that this conclusion requires Bob to alter his subjective view of the action history. There is a non-monotonic character to belief change in this context, because Bob may be forced to postulate and suppress actions over time in response to new observations. The ramifications of changing the action history are determined by the underlying transition system.

4 Plausibility Functions

At each point in time, an agent needs a plausibility ordering over all actions and all states. Moreover, in order to resolve inconsistency at different points in time, each of the plausibility orderings must be comparable. One natural way to create mutually comparable orderings is by assigning quantitative plausibility values to every action and state at every point in time. Towards this end, we define plausibility functions.

Definition 3 *Let X be a non-empty set. A plausibility function over X is a function $r : X \rightarrow \mathbf{N}$.*

If r is a plausibility function and $r(x) \leq r(y)$, then we say that x is at least as plausible as y. We remark that we will typically be interested in plausibility functions over finite sets, where there is always a non-empty set of maximally plausible elements.

Plausibility functions are inspired by Spohn's ordinal conditional functions [Spo88], but there are some important differences. First, we allow plausibility functions over an arbitrary set X, rather than restricting attention to propositional interpretations. This allows us to treat partially observable actions in the same manner that we treat observations. Another important difference is that ordinal conditional functions must always assign rank 0 to a non-empty subset of elements of the domain. Plausibility functions are not restricted in this manner; the minimal rank for a given plausibility function may be greater than 0. This distinction is based on our underlying intuition that some observations provide more reliable information than others.

In order to illustrate the application of plausibility functions, we continue our simple example.

Example (cont'd) We describe how the cookie problem can be represented with plausibility functions.

Let $\mathbf{F} = \{BiteTaken\}$ and let $\mathbf{A} = \{BiteAlice, BiteEve\}$. Both actions have the same effect, namely they both make the fluent $BiteTaken$ become true. We represent the problem with 3 plausibility functions: a_1, a_2, and o_3.

1. a_1 is a plausibility function over actions at time 1
2. a_2 is a plausibility function over actions at time 2
3. o_3 is a plausibility function over observations at time 3

Informally, each function should obtain a minimum value at the event that Bob considers the most plausible at the given point in time. Since Bob initially believes that no one will eat his cookie, both a_1 and a_2 should obtain a minimum value at the null action λ. The observation that the cookie has been bitten at time 3 is represented by defining o_3 with a minimum at the set of worlds where the cookie has a bite out of it. The additional soft constraints are used to determine the magnitude of the values for each event. Define a_1 and a_2 by the values in the following table.

	λ	$BiteAlice$	$BiteEve$
a_1	0	1	10
a_2	0	10	2

The fact that Alice is more likely to bite the cookie is represented by assigning a low plausibility value to $BiteAlice$ at time 1. Define o_3 as follows.

	$\emptyset$	$\{BiteTaken\}$
o_3	10	0

Hence, the observation $\{BiteTaken\}$ is assigned the minimum plausibility value, and the only alternative observation is assigned a very high plausibility value. This reflects the fact that Trent's report is understood to supersede the assumption that Alice and Eve do not bite the cookie.

5 Graded World Views

Graded world views are a formal tool for determining a maximally plausible world history, given an alternating sequence of ontic actions and epistemic actions. Intuitively, a graded world view simply consists of a sequence of plausibility functions. In the general case, we need two plausibility functions at each point in time. One function assigns a plausibility value to every action symbol, and the other function assigns a plausibility value to every state. The following definitions extend the observation trajectories and action trajectories of [HD05].

Definition 4 *A graded observation trajectory of length n is an n-tuple of plausibility functions over $2^{\mathbf{F}}$.*

Definition 5 *A graded action trajectory of length n is an n-tuple of plausibility functions over $\mathbf{A}$.*

Using graded trajectories, we get the following notion of a graded world view.

Definition 6 *A graded world view is a pair $\langle ACT, OBS\rangle$ where ACT is a graded action trajectory and OBS is a graded observation trajectory of the same length.*

Informally, at each point in time, an action is performed and it is followed by an observation. At time i, the most plausible action is given by the i^{th} plausibility function in ACT and the most plausible observation is given by the i^{th} plausibility function in OBS.

We remark briefly on the intuition behind graded action trajectories. The plausibility value assigned to A represents the plausibility that A is executed at a given instant. Hence, the lowest plausibility values will be assigned to actions that an agent actually performs. The highest values will be assigned to exogenous actions that an agent believes are unlikely to occur. In this paper, we do not explicitly consider failed actions. Instead, we simply note that failed actions can be added to our formalism by allowing non-deterministic actions and attaching a plausibility value to possible effects, as in [Bou95]. We leave such an extension for future work.

Implicitly, the initial belief state in every graded world view is $2^{\mathbf{F}}$. However, if the initial plausibility function in ACT assigns a small minimum value to the null action λ, then one can think of the initial element of OBS as the initial belief state. In this manner, the plausibility of the initial belief state is treated in exactly the same manner as the plausibility of any subsequent observation.

We formally define the notion of a history.

Definition 7 *Let $T = \langle S, R\rangle$ be a transition system. A history over T is a tuple $\langle w_0, A_1, \dots, A_n, w_n\rangle$ where for each i:*

1. *$w_i \in S$,*
2. *$A_i \in \mathbf{A}$,*
3. *$\langle w_i, A_i, w_{i+1}\rangle \in R$.*

A history is simply an alternating sequence of interpretations and actions that represents a possible evolution of the world. Let $HIST_n$ denote the set of histories involving n actions.

Given a graded world view, the main computational task for an agent is to determine the most plausible histories. This is similar to the process of belief extrapolation with mixed scenarios [DdSCL02], with two main differences. First, in belief extrapolation, there is a single plausibility ordering over histories rather than $2n$ orderings over actions and states. Second, belief extrapolation operators are intended for action domains in which individual fluents may change values in an arbitrary manner. In our framework, every change must be caused by some action defined by the underlying transition system.

The plausibility of a history with respect to a graded world view is calculated by summing the plausibility at each instant.

Definition 8 *Let $\langle ACT, OBS\rangle$ be a graded world view. The plausibility of a history $h = \langle w_0, A_1, \dots, A_{n-1}, w_n\rangle$ with respect to $\langle ACT, OBS\rangle$ is the sum*

$$plaus(h) = \sum_{i=1}^{n} ACT_i(A_i) + OBS_i(w_i).$$

It is useful to introduce an operator that maps graded world views to the set of histories that are assigned the minimum sum of plausibility values.

Definition 9 *Let WV denote the set of graded world views of length n for a fixed action signature. Define $\Phi : WV \rightarrow 2^{HIST_n}$ as follows:*

$$\Phi(\langle ACT, OBS\rangle) = \{h \mid plaus(h) \leq plaus(g) \text{ for all } g \in HIST_n\}.$$

We revisit the earlier example with this new notation.

Example (cont'd) In order to give a complete representation of the problem, we need to define a graded world view. We define a graded action trajectory and a graded observation trajectory by extending the tables given previously.

	λ	$BiteAlice$	$BiteEve$
a_0	0	10	10
a_1	0	1	10
a_2	0	10	2
a_3	0	10	10

Note that we now have plausibility functions at time 0 and at time 3. At each of these times, the null action is given the minimum plausibility to reflect that no action occurs. We need to add time 0 in order to restrict the initial belief state to $\emptyset$. Strictly speaking, we do not need to add time 3, we

	$\emptyset$	$\{BiteTaken\}$
o_0	0	10
o_1	0	0
o_2	0	0
o_3	10	0

simply add it to remain consistent with the timeline in the initial problem description.

We are interested in finding $\Phi(\langle a_0, \ldots, a_3\rangle, \langle o_0, \ldots, o_3\rangle)$. By inspection, we find that the minimum plausibility is obtained by the following history:

$$h = \langle \emptyset, \lambda, \emptyset, BiteAlice, BiteTaken, \lambda, BiteTaken, \lambda, BiteTaken\rangle.$$

This history represents the sequence of events in which Alice bites the cookie at time 1. Intuitively, this is the correct solution: given the choice between Alice and Eve, Bob believes that Alice is the one who is more likely to help herself to the cookie.

6 Basic Properties

6.1 Pointwise Dominance

Suppose that the underlying set $\mathbf{F}$ of fluent symbols and the underlying set $\mathbf{A}$ of action symbols are both finite. Let $W = \langle ACT, OBS\rangle$ be a graded world view with

$$ACT = \langle ACT_1, \ldots, ACT_n\rangle$$

and

$$OBS = \langle OBS_1, \ldots, OBS_n\rangle.$$

The simplest way to find a plausible world history is to simply take the most plausible action and most plausible worlds at each point in time. The following definition makes this notion more precise.

Definition 10 *Let $h = \langle w_0, A_1, \ldots, A_n, w_n\rangle$. We say h is a pointwise minimum for $\langle ACT, OBS\rangle$ if, for all i,*

1. *for all $A \in \mathbf{A}$, $ACT_i(A_i) \leq ACT_i(A)$, and*
2. *for all $w \in 2^{\mathbf{F}}$, $OBS_i(w_i) \leq OBS_i(w)$.*

Note that histories are restricted in that each world must be the outcome of the preceding action. As such, it is possible that a graded world view will have no pointwise minimum. However, if there are any pointwise minima, then clearly they will be the most plausible histories. We state this simple fact more formally.

Proposition 1 *Let $W = \langle ACT, OBS\rangle$ be a graded world view and let M be the set of pointwise minima for W. If $M \neq \emptyset$, then $\Phi(W) = M$.*

This observation suggests that checking for pointwise minima may be a good first step in the search for plausible world histories. Finding pointwise minima is not easy in the general case.

Proposition 2 *For a fixed graded world view W, determining if W has a pointwise minimum is NP-complete.*

6.2 Equivalence

Clearly it is possible for two distinct graded world views to have the same set of minimally ranked world histories. In fact, it is possible for two distinct graded world views to induce the same preference ordering over histories. In this section, we define a natural equivalence relation over graded world views with an eye towards categorical representations. We start by defining a relation on plausibility functions.

Definition 11 *Let P_1 and P_2 be plausibility functions over a set X. We say that $P_1 \cong P_2$ if, for every $x, y \in X$,*

$$P_1(x) - P_1(y) = P_2(x) - P_2(y).$$

It is clear that $\cong$ is an equivalence relation. It is also clear that this relation can be extended to graded world views.

Definition 12 *Let WV_1 and WV_2 be graded world views over histories for a fixed action signature. We say that $WV_1 \cong WV_2$ if, for every pair of histories g and h,*

$$plaus_1(g) - plaus_1(h) = plaus_2(g) - plaus_2(h).$$

Let P_1 be a plausibility function. We say that P_2 is obtained from P_1 by a *translation* if there is some n such that, for all x, $P_1(x) = P_2(x) + n$. It is easy to see that, whenever $P_1 \cong P_2$, it must be the case that P_2 is a obtained by a translation on P_1. If a graded world view WV_2 is obtained from WV_1 by uniformly translating every component, then clearly $WV_1 \cong WV_2$. However, it is straightforward to construct equivalent graded world views that are not obtained by translations.

7 Comparison with Related Formalisms

7.1 Representing Belief States

We introduce some notation that allows belief states to be represented by plausibility functions. If K is a set of interpretations and c is an integer, let $K \uparrow c$ denote function defined as follows:

$$K \uparrow c\,(w) = \begin{cases} 0 & \text{if } w \in K \\ c & \text{otherwise} \end{cases}$$

If c is a positive integer, then $K \uparrow c$ denotes a plausibility function in which the elements of K are the most plausible, and everything else is equally implausible. Plausibility functions of the form $K \uparrow c$ will be called *simple*.

If $c < 0$, then $K \uparrow -c$ does not actually define a plausibility function. However, allowing negative values leads to a simple symmetry in our notation. In the following proposition, $\overline{K}$ denotes the complement of K.

Proposition 3 *For any belief state K and positive integer c*

$$K \uparrow c \cong \overline{K} \uparrow -c.$$

Note that translating $\overline{K} \uparrow -c$ by an integer greater than c gives another equivalent plausibility function. However, this equivalence does not mean that $K \uparrow c$ is interchangeable with translations of $\overline{K} \uparrow -c$ in a given world view. The magnitude of the largest plausibility value is different, which can be significant when determining minimal sums.

Suppose that

$$ACT = \langle ACT_1, \ldots, ACT_n\rangle$$

and

$$OBS = \langle OBS_1, \ldots, OBS_n \rangle$$

where each ACT_i and OBS_i is simple, with maximum plausibility c. Hence, we essentially have belief states with no plausibility ordering. In this case, it is easy to show that

$$\langle w_0, A_1, \ldots, A_n, w_n \rangle \in \Phi(\langle ACT, OBS \rangle)$$

if and only if

$$|\{A_i | A_i \in ACT_i\}| + |\{w_i | w_i \in OBS_i\}|$$

is maximal among all histories. In other words, the most plausible histories are those that agree with $\langle ACT, OBS \rangle$ at a maximal number of components.

The case in which there is no plausibility ordering is not very interesting from the perspective of belief change. However, it is easy to see that AGM belief revision operators can also be represented. In particular, let r be a plausibility function over X with minimum value $\min_r$. For any n, let $r[n]$ denote the set of complete, consistent theories that are satisfied by some I with $r(I) \leq n$.

Proposition 4 *The collection $\mathcal{R} = \{r[n] \mid n \geq \min_r\}$ is a system of spheres centered on $r[\min_r]$.*

Now, by applying well-known results of Grove [Gro88], it is easy to construct a graded world view of length 2 corresponding to any AGM revision operator. In particular, if only null actions are permitted and the second observation is simple, then we essentially have AGM revision. This relationship is not surprising, since plausibility functions clearly induce an ordering over the the set of possible worlds.

We remark that graded world views bear a resemblance to the generalized belief change framework proposed by Liberatore and Schaerf [LS00]. However, there are some important distinctions. The Liberatore-Schaerf approach associates a "penalty" with state change, which is minimized when determining plausible models. As such, it is difficult to represent problems where non-null actions are strictly more plausible than null actions. By contrast, graded world views have no implicit preference for null actions. Moreover, since we define actions with respect to a transition system, graded world views are more suitable for the representation of actions with conditional effects.

7.2 Representing Belief Evolution Operators

Belief evolution operators have been introduced to represent sequences of alternating updates and revisions. We briefly sketch the approach, and refer the reader to [HD05] for the details. Let $A = A_1, \ldots, A_n$ be a sequence of action symbols and let $O = O_1, \ldots, O_n$ be a sequence of observations. Given an initial belief state K, the evolution operator $\circ$ roughly corresponds to the following iterated belief change:

$$K \circ \langle A, O \rangle = K \diamond A_1 * O_1 \diamond \cdots \diamond A_n * O_n.$$

Simply performing the updates and revisions in succession gives unintuitive results. As a result, we have specified a number of so-called interaction postulates, and the definition of $\circ$ is constructed in a manner that assures the postulates must hold.

There are two underlying assumptions in belief evolution that are not required in a representation by graded world views.

1. The plausibility of an observation is determined by recency.
2. The action history is assumed to be correct.

Both of these assumptions can be represented in a graded world view by setting up the plausibility functions appropriately. In particular, for each i, we define

$$OBS_i = O_i \uparrow 2^i.$$

By incrementing the plausibility of false observations exponentially, we can assure that recent observations will be given greater credence. The fact that action histories must be correct is represented by setting

$$ACT_i = A_i \uparrow 2^{n+1}$$

for every i. Recall that $\circ$ is defined with respect to an update operator $\diamond$ and a revision operator $*$. As a result, in order to represent $\circ$ in a graded world view, we also need to encode the plausibility ordering implicit in $*$. Omitting the details of the construction, we get the following result.

Proposition 5 *Let $\circ$ be a belief evolution operator obtained from $\diamond$ and $*$. There is a graded world view W_{ev} such that, if $K \circ \langle O, A \rangle = \langle K_0, \ldots, K_n \rangle$, then*

$$\langle w_0, A_1, \ldots, A_n, w_n \rangle \in \Phi(W_{ev})$$
$$\Longleftrightarrow$$
$$\textit{for each } i,\ w_i \in K_i.$$

Proposition 5 demonstrates that graded world views can represent any belief evolution operator. So, the interaction postulates for belief evolution will be satisfied by a graded world view whenever the plausibility functions are defined as above. Hence, from the perspective of graded world views, the role of the interaction postulates is essentially to restrict the admissible plausibility functions.

7.3 Comparison With Belief Extrapolation

As noted earlier, the motivation underlying our formalism is similar to the motivation underlying belief extrapolation operators. In this section, we demonstrate some expressive differences between the two formalisms. In the interest of space, we refer the reader to [DdSCL02] for the required background on belief extrapolation. We remark that we will abuse notation by equating the trajectories of belief evolution with histories.

We need to give some simple terminology used in belief extrapolation. A *scenario* is a tuple of formulas. For t less than the length of Σ, let $\Sigma(t)$ denote the t^{th} formula in the scenario Σ. We say that a history $\langle w_0, \ldots, w_n \rangle$ satisfies a scenario Σ if $w_i \models \Sigma_i$ for each $i \leq n$. The set of histories satisfying Σ is denoted by $\mathit{Traj}(\Sigma)$.

We are interested in determining if all belief extrapolation operators can be represented by graded world views. First, we need to formalize the problem more precisely.

Definition 13 *Let $\updownarrow$ be a belief extrapolation operator. We say that $\updownarrow$ is representable if, for every scenario Σ of length n, there is a graded world view $\langle OBS, ACT \rangle$ of length n such that*

$$\mathit{Traj}(\Sigma \updownarrow) = \Phi(\langle OBS, ACT \rangle).$$

If $\updownarrow$ is representable, then the behavior of $\updownarrow$ can be simulated with graded world views.

The following proposition indicates that belief extrapolation operators have an expressive advantage.

Proposition 6 *There is a belief extrapolation operator* $\updownarrow$ *that is not representable.*

We remark that the proof of Proposition 6 is constructive and it demonstrates that there is a simple, concrete, inertial extrapolation operator that is not representable. Intuitively, the distinction is that a belief extrapolation operator is based on an ordering over trajectories rather than several orderings over actions and states.

There is also a sense in which graded world views are more expressive than belief extrapolation operators. In particular, they provide a mechanism for handling unreliable observations. One of the main assumptions underlying belief extrapolation is that every observation should be incorporated in the new scenario. By contrast, we are interested in applications where some observations may be incorrect. For example, the cookie problem can easily be modified by assigning a very high value to the observation reported by Trent. This would reflect the fact that Trent is not a reliable source of information, and it would lead to plausible histories in which that observation is simply ignored. In particular, if Trent is not reliable, then the most plausible history is the one in which nobody bites the cookie.

8 Discussion

We have introduced a formalism for reasoning about sequences of actions and observations. The formalism uses ranking functions at each instant to determine the most plausible action or observation, and determines the most plausible histories by summing over all instants. The formalism provably subsumes belief revision and belief evolution. The relationship with belief extrapolation is more subtle, with each formalism having expressive advantages and disadvantages.

The generality of graded world views can be seen in comparison with iterated revision. Papini illustrates two different approaches to iterated revision, one which gives greater credence to recent information and one which gives greater credence to old information [Pap01]. In the same manner, given a sequence of actions and observations $A_1, O_1, \ldots, A_n, O_n$, any ordering of the actions and observations may be used to resolve conflicts. Clearly, any such ordering can be represented by a graded world view by assigning maximal plausibility values that increase as powers of 2. Moreover, repetitions of observations and actions can be used to make democratic decisions based on, for example, the number of times that a given observation occurs. For this kind of reasoning, the full generality of arbitrary plausibility functions is useful.

We consider graded world views to be the most general possible extension of belief evolution. The interaction postulates of belief evolution essentially formalize the fact that action histories are infallible. Hence, from the perspective of graded world views, the postulates serve to restrict the admissible plausibility functions. Alternative postulates could be proposed to give different restrictions for a different class of problems. In the future, we would like to look for a most general set of postulates for which we could prove a representation result for graded world views.

References

[AGM85] C.E. Alchourron, P. Gardenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Bou95] C. Boutilier. Generalized update: Belief change in dynamic settings. In *Proceedings of IJCAI 1995*, pages 1550–1556, 1995.

[DdSCL02] F. Dupin de Saint-Cyr and J. Lang. Belief extrapolation (or how to reason about observations and unpredicted change). In *Proceedings of KR 2002*, pages 497–508, 2002.

[GL98] M. Gelfond and V. Lifschitz. Action languages. *Linköping Electronic Articles in Computer and Information Science*, 3(16):1–16, 1998.

[Gro88] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[HD05] A. Hunter and J.P. Delgrande. Iterated belief change: A transition system approach. In *Proceedings of IJCAI05*, 2005.

[HLM04] A. Herzig, J. Lang, and P. Marquis. Revision and update in multi-agent belief structures. In *Proceedings of LOFT 6*, 2004.

[KM92] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gardenfors, editor, *Belief Revision*, pages 183–203. Cambridge University Press, 1992.

[LS00] P. Liberatore and M. Schaerf. Brels: A system for the integration of knowledge bases. In *Proceedings of KR2000*, pages 145–152. Morgan Kaufmann Publishers, 2000.

[Pap01] O. Papini. Iterated revision operations stemming from the history of an agent's observations. In H. Rott and M. Williams, editors, *Frontiers in Belief Revision*, pages 279–301. Kluwer Academic Publishers, 2001.

[SPLL00] S. Shapiro, M. Pagnucco, Y. Lesperance, and H.J. Levesque. Iterated belief change in the situation calculus. In *Proceedings of KR 2000*, pages 527–538. Morgan Kaufmann Publishers, 2000.

[Spo88] W. Spohn. Ordinal conditional functions. a dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics, vol. II*, pages 105–134. Kluwer Academic Publishers, 1988.