Belief Change in the Context of Fallible Actions and Observations

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

Abstract

We consider the iterated belief change that occurs following an alternating sequence of actions and observations. At each instant, an agent has some beliefs about the action that occurs as well as beliefs about the resulting state of the world. We represent such problems by a sequence of ranking functions, so an agent assigns a quantitative plausibility value to every action and every state at each point in time. The resulting formalism is able to represent fallible knowledge, erroneous perception, exogenous actions, and failed actions. We illustrate that our framework is a generalization of several existing approaches to belief change, and it appropriately captures the non-elementary interaction between belief update and belief revision.

Introduction

Several formalisms have been introduced for reasoning about belief change in the context of actions and observations, including (Shapiro *et al.* 2000; Herzig, Lang, & Marquis 2004; Jin & Thielscher 2004). Roughly speaking, agents perform belief update following actions and agents perform belief revision following observations. Existing formalisms for the most part have treated actions and observations independently, with little explicit discussion about the interaction between the two. In this paper, we consider the belief change that occurs due to an alternating sequence of actions and observations. We are interested in action domains where an agent may have erroneous beliefs, both about the state of the world as well as the action history.

Let K denote the initial beliefs of an agent, represented as a set of possible worlds. For $1 \leq i \leq n$, let A_i denote an action and let O_i denote an observation. Informally, we are interested in sequences of the form

$$K \diamond A_1 * O_1 \diamond \cdots \diamond A_n * O_n \quad (1)$$

where $\diamond$ is an update operator and $*$ is a revision operator. Such sequences may contain conflicting information. For example, the observation O_n may not be possible following the actions $A_1, \ldots, A_n$. In this case, there are two options.

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

In order to determine which option is preferable for a specific problem, an agent needs to be able to compare the plausibility of O_n with the plausibility of each A_i.

Expressions of the form (1) have previously been considered in (Hunter & Delgrande 2005), under the assumption that ontic action histories are infallible and recent observations take precedence over older observations. Clearly, there are action domains in which these assumptions are not reasonable. In this paper, we propose a more flexible approach in which actions and observations are both represented by Spohn-style ranking functions. When presented with conflicting information, an agent appeals to the relative plausibility of each action and observation.

This paper makes several contributions to existing work on epistemic action effects. The main contribution is a formal mechanism for representing fallible beliefs about action histories. Existing formalisms are unable to compare the plausibility of an action occurrence with the plausibility of a state of the world. By using the same formal tool to represent beliefs about actions and states, we explicitly address the manner in which prior action occurrences are postulated or retracted in response to new observations. A second contribution is a flexible treatment of unreliable observations, in which new information need not always be incorporated. We formulate all of our results in a simple transition system framework that makes our treatment of action effects explicit and easy to compare with more elaborate action formalisms.

Preliminaries

We are interested in action domains that can be described by a finite set of fluent symbols $\mathbf{F}$, a finite set of action symbols $\mathbf{A}$, and a transition system describing the effects of actions(Gelfond & Lifschitz 1998). A *state* is an interpretation over $\mathbf{F}$; we identify the state s with the set of fluent symbols that are true in s.

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

We restrict attention to *deterministic* transition systems, i.e. we assume that $\langle s, A, s' \rangle \in R$ and $\langle s, A, s'' \rangle \in R$ implies $s' = s''$. We also assume that $\mathbf{A}$ always contains a distinguished null action symbol denoted by λ.

A *belief state* is a set of states, informally the set of states that an agent considers possible. An *observation* is also a set of states. The observation α provides evidence that the actual state is in α. We assume the reader is familiar with the AGM approach to belief revision (Alchourrón, G ardenfors, & Makinson 1985) and the Katsuno-Mendelzon approach to belief update (Katsuno & Mendelzon 1992). Our approach differs in that we consider belief revision and belief update with respect to sets of states, rather than formulas. Moreover, we define update with respect to an action with effects given by an underlying transition system.

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 $s' \in \alpha\}$.*

We introduce a common-sense example, to which we will return after introducing our formal machinery.

Example Consider an 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 single 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 retract actions over time in response to new observations. We remark that, in order to represent this kind of reasoning, we need to be able to compare the plausibility of action occurrences at different points in time.

Ranking Functions over Actions and States

Plausibility Functions

At each point in time, an agent needs a plausibility ordering over all actions and a plausibility ordering over 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. In this paper, we restrict attention to plausibility functions over finite sets.

Plausibility functions are inspired by Spohn's ordinal conditional functions (Spohn 1988), with two main 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. The second 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.

We introduce some notation related to plausibility functions. For any plausibility function r, let $\min_r$ denote the minimum value obtained by r and let $Bel(r) = \{w \mid r(w) = \min_r\}$. The *normalization* of a plausibility function r is the function r' with minimum zero, defined by $r'(w) = r(w) - \min_r$. The *degree of strength* of a plausibility function r is the span between the plausibility of the minimally ranked elements and the non-minimally ranked elements. Formally, the degree of strength of r is the least n such that $\min_r + n = r(v)$ for some $v \notin Bel(r)$. There are two natural interpretations of the degree of strength of a plausibility function r over a set of states. If we think of r as an initial belief state, then r represents the belief that the actual state is in $Bel(r)$ and the degree of strength of r is an indication of how strongly this is believed. If we think of r as an observation, then the degree of strength is a measure of reliability. We remark that Spohn defines the degree of strength of a subset of X, rather than the degree of strength of a ranking function. Our definition coincides with Spohn's definition if we identify the degree of strength of r with Spohn's degree of strength of the set $Bel(r)$. Hence, we use the same conception of degree of strength, but we are only interested in the strength of belief in the minimally ranked elements.

Graded World Views

Informally, a *graded world view* represents an agent's subjective view of the evolution of the world. Before defining graded world views, we need to define the notion of a *history* over a transition system.

Definition 4 *Let $T = \langle S, R \rangle$ be a transition system. A history of length n is a tuple $\langle w_0, A_1, \ldots, A_n, w_n \rangle$ where for each i: $w_i \in S$, $A_i \in \mathbf{A}$, and $\langle w_i, A_i, w_{i+1} \rangle \in R$.*

Let $HIST_n$ denote the set of histories of length n.

We are interested in action domains where an agent is uncertain not only about the state at each point in time, but also uncertain about the action that has been executed. As such, at each time i, we use a plausibility function over $2^{\mathbf{F}}$ to represent an agent's beliefs about the state of the world and we use a plausibility function over $\mathbf{A}$ to represent an agent's beliefs about the action that occurs.

Definition 5 *A graded world view of length* n *is a* $(2n+1)$-*tuple*

$$\langle OBS_0, ACT_1, OBS_1, \ldots, ACT_n, OBS_n \rangle$$

where each OBS_i *is a plausibility function over* $2^{\mathbf{F}}$ *and each* ACT_i *is a plausibility function over* $\mathbf{A}$.

At time i, the most plausible actions are the minimally ranked actions of ACT_i and the most plausible states are the minimally ranked states of OBS_i. As such, we take OBS_0 to represent the initial belief state, and each subsequent OBS_i to represent a new observation. If $ACT = \langle ACT_1, \ldots, ACT_n \rangle$ and $OBS = \langle OBS_0, \ldots, OBS_n \rangle$, then we write $\langle ACT, OBS \rangle$ as a shorthand for the graded world view $\langle OBS_0, ACT_1, OBS_1, \ldots, ACT_n, OBS_n \rangle$.

Given a graded world view $\langle ACT, OBS \rangle$, we would like to be able to determine the plausibility of a history h. However, $\langle ACT, OBS \rangle$ does not provide sufficient information to pick out a unique plausibility function over histories. For example, a graded world view does not indicate the relative weight of recent information versus initial information.

Although a graded world view does not define a unique plausibility function over histories, we can define a general notion of consistency between graded world views and plausibility functions on histories. Let $\langle r_0, \ldots, r_n \rangle$ be a sequence of plausibility functions over $X_0, \ldots, X_n$, respectively. Let r be a plausibility function over $X_0 \times \cdots \times X_n$. We say that r is *consistent with* $\langle r_0, \ldots, r_n \rangle$ if, for every i and every $x_i, x_i' \in X_i$

$$r_i(x_i) < r_i(x_i')$$
$$\Longleftrightarrow$$
$$r(\langle x_0, \ldots, x_i, \ldots, x_n \rangle) < r(\langle x_0, \ldots, x_i', \ldots, x_n \rangle)$$

So r is consistent with $\langle ACT, OBS \rangle$ just in case r increases monotonically with respect to each component of $\langle ACT, OBS \rangle$. Any plausibility function r that is consistent with $\langle ACT, OBS \rangle$ provides a potential candidate ranking over histories.

An *aggregate plausibility function* is a function $plaus$ that maps every graded world view to a plausibility function on histories. An aggregate plausibility function $plaus$ is *admissible* if, for every $\langle ACT, OBS \rangle$, the function $plaus(\langle ACT, OBS \rangle)$ is consistent $\langle ACT, OBS \rangle$.

We provide some examples. Note that aggregate plausibility functions return a function as a value; we can specify the behaviour of an aggregate by specifying a plausibility value for each pair consisting of a graded world view and a history. Let $h = \langle w_0, A_1, \ldots, A_n, w_n \rangle$. One admissible aggregate is obtained by taking the sum of plausibility values.

$$sum(ACT, OBS)(h) = \sum_{i=1}^{n} ACT_i(A_i) + OBS_i(w_i)$$

A weighted sum can be used to reflect the relative importance of different time points. For each i, let b_i be a positive integer.

$$sum_w(ACT, OBS)(h) = \sum_{i=1}^{n} ACT_i(A_i) + b_i \cdot OBS_i(w_i).$$

By setting $b_i = 2^i$, the aggregate function sum_w can be used to represent a strict preference for recent information. The functions sum and sum_w are just two simple examples; many more examples can be defined by specifying aggregate functions that increase monotonically with each component.

Example (cont'd) 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 need to define plausibility functions a_1, a_2 over actions and plausibility functions o_1, o_2, o_3 over states. The function a_i gives the plausibility ranks for each action at Time i and the function o_i gives the plausibility ranks for each state at Time i. Since Bob believes that no one will eat his cookie, every a_i obtains a minimum value at the null action λ. Define a_1, 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 lower plausibility value to $BiteAlice$ at Time 1.

The plausibility function o_0 represents the initial state, so it should assign a minimum value to the state where the cookie is unbitten. The plausibility function o_2 represents Trent's report that the cookie has been bitten. We remark that we will generally treat reported information as an observation, and we will use the degree of strength of the reported information as an indication of the reliability of the source. In this case, the degree of strength of o_2 is an indication of trust in Trent. Define o_0, o_1, o_2 as follows.

	$\emptyset$	$\{BiteTaken\}$
o_0	0	9
o_1	0	0
o_2	9	0

Note that the degree of strength of o_2 is higher than the degree of strength of a_1 or a_2. This reflects the fact that Trent's report is understood to supersede the assumption that Alice and Eve do not bite the cookie. Graded world views have been defined precisely for this kind of comparison between action plausibilities and state plausibilities.

If we use the aggregate function sum, then we are interested in finding the minimal sum of plausibilities over $\langle o_0, a_1, o_1, a_2, o_2 \rangle$. By inspection, we find that the minimum plausibility is obtained by the following history:

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

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 more plausible culprit.

We remark that graded world views bear a resemblance to the generalized belief change framework proposed by Liberatore and Schaerf (2000). However, 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, our approach differs in that we allow actions with conditional effects given by a transition system.

Subjective Probabilities

One issue that arises from our definition of a graded world view is the fact that it is not clear how plausibility values should be assigned in practical problems. We address this problem by illustrating a correspondence between plausibility functions and *probability functions*. We simplify the discussion by restricting attention to rational-valued probability functions as follows.

Definition 6 *Let X be a non-empty set. A probability function over X is a function $Pr : X \rightarrow \mathbf{Q}$ such that*

- *for all $x \in X$, $0 \leq Pr(x) \leq 1$*
- $\sum_{x \in X} Pr(x) = 1$.

At a common-sense level, it is clear what it means to say that "action A occurred at time t with probability p." By contrast, the problem with plausibility values is that there is no obvious sense of scale; it is difficult to assign numerical plausibility values, because the numbers have no clear meaning. We illustrate how probability functions can be translated uniformly into plausibility functions, thereby giving a sense of scale and meaning to plausibility values.

Let Pr be a probability function over a finite set X. Let Q denote the least common denominator of all rational numbers $\frac{p}{q}$ such that $Pr(x) = \frac{p}{q}$ for some $x \in X$. Define the plausibility function r as follows.

1. If $Pr(x)$ is minimal, set $r(x) = Q$.
2. Otherwise, if $Pr(x) = \frac{p}{Q}$, then set $r(x) = Q - p$.

Hence, every probability function can be translated into a plausibility function.

Example (cont'd) In the cookie example, the given plausibility functions are obtained by starting with the following probability functions and then normalizing.

	λ	$BiteAlice$	$BiteEve$
a_1	.5	.45	.05
a_2	.5	.15	.35

	$\emptyset$	$\{BiteTaken\}$
o_0	.9	.1
o_1	.5	.5
o_2	.1	.9

This perspective on plausibility functions also provides some justification for the use of the aggregate function sum. In particular, if we assume that the subjective probability functions are *independent*, then the probability of a given sequence of events is determined by taking a product. In the cookie example, we can compare the probability of Alice biting the cookie versus Eve biting the cookie:

1. $Pr(\langle \emptyset, BiteAlice, BiteTaken, \lambda, BiteTaken \rangle)$
   $= .9 \times .45 \times .5 \times .5 \times .9 = .091125$
2. $Pr(\langle \emptyset, \lambda, \emptyset, BiteEve, BiteTaken \rangle)$
   $= .9 \times .5 \times .5 \times .35 \times .9 = .070875$

It is easy to check that the history where Alice bites the cookie is actually the most probable history. So, in this example, the minimally ranked history according to the aggregate function sum is also the most probable history according the sequence of probability functions. This is a general property of our translation: maximizing probability over independent probability functions corresponds to minimizing the sum over plausibility values. In the interest of space, we omit the proof of this property.

Note that we have explicitly assumed plausibility functions in a graded world view are independent. The likelihood of a given action at a fixed point in time does not depend on the state of the world. We leave the treatment of conditional events for future work.

Graded World Views as Epistemic States

Unless otherwise indicated, we assume that plausibility values are assigned to histories by the aggregate function sum. Although this is not the only approach to combining plausibility functions, it provides a categorical example that allows us to ground the discussion.

Graded world views can be defined that simply pick out a distinguished set of elements of the domain. If α is a subset of X and c is an integer, let $\alpha \uparrow c$ denote the function defined as follows:

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

Plausibility functions of the form $\alpha \uparrow c$ will be called *simple*. If X is a set of states, then simple plausibility functions correspond to belief states; if X is a set of actions, then simple plausibility functions pick out the actions that are believed to have occurred.

More generally, every graded world view defines an epistemic state in the sense of Darwiche and Pearl(1997). In particular, we can define a pre-order $\preceq$ on states as follows. For any states s_1, s_2, define $s_1 \preceq s_2$ if and only if

$$sum(\langle ACT, OBS \rangle)(h_1) \leq sum(ACT, OBS)(h_2)$$

for some pair of histories h_1, h_2 with terminal states s_1, s_2, respectively. For any graded world view $\langle ACT, OBS \rangle$, we let $Bel(ACT, OBS)$ denote the minimal states according to this ordering.

If we think of graded world views as epistemic states, then we can define standard belief change operations in a more familiar manner. If r_A is a plausibility function over actions and r_S is a plausibility function over states, we define $\bullet$ as follows:

$$\langle ACT, OBS \rangle \bullet \langle r_A, r_S \rangle = \langle ACT \cdot r_A, OBS \cdot r_S \rangle$$

where $\cdot$ denotes concatenation. We can identify belief update with the concatenation of a single action and we can identify belief revision with the concatenation of a single observation. This new approach to update and revision is demonstratably more expressive than the standard approaches since it incorporates varying degrees of reliability.

Consider an expression of the form $\langle INIT \rangle \bullet \langle r_A, r_S \rangle$. In this context, $INIT$ represents the initial beliefs of an agent, r_A represents an agent's beliefs about the action that has been executed, and r_S represents the observed state of the world. Assume the minimal elements have degrees of strength $deg(INIT)$, $deg(r_A)$, and $deg(r_S)$ respectively. Varying the magnitudes of these values allows us to capture several different underlying assumptions.

1. Fallible initial beliefs: $deg(INIT) < deg(r_A)$ and $deg(INIT) < deg(r_S)$.
2. Erroneous perception: $deg(r_S) < deg(INIT)$ and $deg(r_S) < deg(r_A)$.
3. Fallible action history: $deg(r_A) < deg(INIT)$ and $deg(r_A) < deg(r_S)$.

By manipulating degrees of strength in this manner, we can also represent exogenous actions, additive evidence, and noisy observations.

Comparison with Existing Formalisms

Single Shot Belief Change

In this section, we consider graded world views from the perspective of single-shot belief change. The initial epistemic state is given by a graded world view, and we are interested in the belief change that occurs when a single action or observation is added. We first consider the case of a single ontic action.

Proposition 1 *Let $\langle ACT, OBS \rangle$ be a graded world view and let 0 denote a constant plausibility function on states. For any plausibility function r over $\mathbf{A}$ with $Bel(r) = \{A\}$,*

$$Bel(\langle ACT, OBS \rangle \bullet \langle r, 0 \rangle) = Bel(\langle ACT, OBS \rangle) \diamond A.$$

Hence, for a single action, we need not consider the entire action history. If we are only interested in the most plausible outcome states, we need only consider the belief state defined by the initial graded world view.

In the case of a single observation, we are interested in comparing graded world views with AGM revision operators. There are two senses in which graded world views are clearly more expressive than AGM operators. First, when we add a new observation to a graded world view, we get a new pre-order over states; so graded world views are able to represent more sophisticated approaches to iterated revision. Second, graded world views allow unreliable observations that need not be incorporated into an agent's beliefs. Both of these distinctions are unimportant if we restrict attention to a single observation with a sufficiently high degree of reliability. Under these restrictions, graded views can essentially be captured by AGM revision operators. We make this claim precise in the next propositions.

First, we prove that every plausibility function defines a system of spheres. For any n, let $r[n]$ denote the set of complete, consistent theories over $\mathbf{F}$ that are satisfied by some I with $r(I) \leq n$.

Proposition 2 *Let r be a plausible function over a finite action signature. The collection $\mathcal{R} = \{r[n] \mid n \geq \min_r\}$ is a system of spheres centered on $r[\min_r]$.*

By combining this with Grove's representation result (Grove 1988), we can prove that a graded world view can be represented by an AGM revision operator when we restrict attention to observations with high degree of strength.

Proposition 3 *Let $\langle ACT, OBS \rangle$ be a graded world view. There is an AGM revision function $*$ and a natural number n such that, for any plausibility function r over states with degree of strength at least n,*

$$\begin{aligned} & Bel(\langle ACT, OBS \rangle \bullet \langle \lambda \uparrow n, r \rangle) \\ = \; & Bel(\langle ACT, OBS \rangle) * Bel(r). \end{aligned}$$

Proposition 3 states that, for a single observation, the most plausible states can be determined by AGM revision. The converse is also true.

Proposition 4 *Let $*$ be an AGM revision operator and let K be a belief state. There is a graded world view $\langle ACT, OBS \rangle$ with $Bel(\langle ACT, OBS \rangle) = K$ and a natural number n such that, for every non-empty observation α,*

$$K * \alpha = Bel(\langle ACT, OBS \rangle \bullet \langle \lambda \uparrow n, r \rangle)$$

where r is any plausibility function over states where the minimal ranked elements α have degree n.

Propositions 3 and 4 illustrate that, if we are only interested in the final belief state, then graded world views are equivalent to AGM revision when we restrict attention to a single observation with a sufficiently high degree of reliability.

Spohn uses ranking functions to define a more general form of belief change called *conditionalization* (Spohn 1988). The idea is that new evidence is presented as a pair (α, m), where m indicates the strength of the observation α. Given a plausibility function r over states, the conditionalization $r_{(\alpha,m)}$ is a new plausibility function where the minimal α-worlds receive plausibility 0 and the rank of non α-worlds is shifted upwards by m. We can define conditionalization in terms of graded world views by defining the following plausibility function

$$r_C(\alpha, m)\,(w) = \begin{cases} 0 & \text{if } w \in \alpha \\ m + \min(\alpha) & \text{if } w \notin \alpha \end{cases}$$

where $\min(\alpha)$ denotes the minimum value that r assigns to an element of α.

Proposition 5 *For any r, α, m, the ranking function $r_{(\alpha,m)}$ is the normalization of the function $r + r_C(\alpha, m)$.*

It is easy to generalize Proposition 5 to the case of graded world views. Briefly, since $\langle ACT, OBS \rangle$ defines a plausibility function over final states, we can simply consider the corresponding function r_C.

We remark that the epistemic extension of the Fluent Calculus provides an axiomatic treatment of plausibility minimization (Jin & Thielscher 2004). However, since the Fluent

Calculus is not concerned with partially observable actions, it is not possible to represent graded world views in the current framework. This would be an interesting topic for future research.

Iterated Epistemic Action Effects

Iterated epistemic action effects are treated in (Hunter & Delgrande 2005). It is pointed out that alternating sequences of updates and revisions should not be computed by naively applying successive operations; instead a set of rationality postulates is presented, and so-called *belief evolution operators* are introduced to give a reasonable treatment of the interaction between update and revision. Let K be a belief state, let $\mathsf{A} = A_1, \ldots, A_n$ be a sequence of action symbols and let $\mathsf{O} = O_1, \ldots, O_n$ be a sequence of observations. A belief evolution operator $\circ$ is defined such that we have the following informal correspondence

$$K \circ \langle \mathsf{A}, \mathsf{O} \rangle \approx K \diamond A_1 * O_1 \diamond \cdots \diamond A_n * O_n.$$

There are two underlying assumptions in the definition of belief evolution. First, there is a reliability ordering over all observations. Second, the action history is assumed to be correct.

Let $K, \mathsf{O}, \mathsf{A}$ be as above, define $\theta(K, \langle \mathsf{A}, \mathsf{O} \rangle) = \langle ACT, OBS \rangle$ where:

1. $OBS_0 = K \uparrow 1$
2. For $1 \leq i \leq n$, $OBS_i = O_i \uparrow 2^i$
3. For $1 \leq i \leq n$, $ACT_i = A_i \uparrow 2^{n+1}$

Proposition 6 *For any $K, \mathsf{O}, \mathsf{A}$ as above, $K \circ \langle \mathsf{A}, \mathsf{O} \rangle = Bel(\theta(K, \langle \mathsf{A}, \mathsf{O} \rangle))$.*

We remark that this translation is not surjective; there are graded world views that do not correspond to any belief evolution operator. Also, note that if we restrict attention to null actions, this translation gives Nayak's lexicographic iterated revision operator (Nayak 1994).

Another approach to reasoning about the evolution of an agent's beliefs in a non-static world is provided by *belief extrapolation* (Dupin de Saint-Cyr & Lang 2002). Belief extrapolation operators are intended to capture the manner in which an agent's beliefs should evolve in response to unpredicted change. A belief extrapolation operator $\updownarrow$ maps a sequence of observations to a set of preferred histories. We say that a belief extrapolation operator is *representable* if there is a graded world view with the same set of preferred histories.

Proposition 7 *There is a belief extrapolation operator $\updownarrow$ that is not representable.*

The proof of Proposition 7 is constructive and it demonstrates that there is a simple, concrete, extrapolation operator that is not representable. The key point is that a belief extrapolation operator is defined with respect to a single ordering over histories, whereas a graded world view is defined with respect to several independent orderings over states. We remark that belief extrapolation operators are not strictly more expressive than graded world views, however. Belief extrapolation does not assume an underlying transition system, so it does not allow for actions with conditional effects.

Conclusion

We have introduced a formalism for reasoning about sequences of actions and observations. The formalism uses Spohn-style ranking functions at each instant to determine the most plausible action or observation, and determines the most plausible histories by an aggregate function over all instants. We have proved that the formalism subsumes belief revision, belief evolution, and conditionalization. Moreover, it is suitable for the representation of fallible beliefs, erroneous perception, exogenous actions, and failed actions. We have used transition systems for the representation of actions in order to facilitate comparison with a wide range of action formalisms. In future work, we will be interested in axiomatizing the belief change that is permitted by the class of admissible aggregate functions.

References

Alchourrón, C.; G ardenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50(2):510–530.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89(1-2):1–29.

Dupin de Saint-Cyr, F., and Lang, J. 2002. Belief extrapolation (or how to reason about observations and unpredicted change). In *Proceedings of KR2002*, 497–508.

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

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

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

Hunter, A., and Delgrande, J. 2005. Iterated belief change: A transition system approach. In *Proceedings of IJCAI05*, 460–465.

Jin, Y., and Thielscher, M. 2004. Representing beliefs in the fluent calculus. In *Proceedings of ECAI04*.

Katsuno, H., and Mendelzon, A. 1992. On the difference between updating a knowledge base and revising it. In G ardenfors, P., ed., *Belief Revision*. Cambridge University Press. 183–203.

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

Nayak, A. 1994. Iterated belief change based on epistemic entrenchment. *Erkenntnis* 41:353–390.

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

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