

Iterated Belief Change: A Transition System Approach

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 use a transition system approach to reason about the evolution of an agent's beliefs as actions are executed. Some actions cause an agent to perform belief revision and some actions cause an agent to perform belief update, but the interaction between revision and update can be non-elementary. We present a set of basic postulates describing the interaction of revision and update, and we introduce a new belief evolution operator that gives a plausible interpretation to alternating sequences of revisions and updates.

1 Introduction

Formalisms for reasoning about action effects typically focus on the representation of actions that change the state of the world. However, several formalisms have been introduced for reasoning about actions that change the beliefs of an agent without altering the state of the world [Shapiro *et al.*, 2000; van Ditmarsch, 2002; Herzig *et al.*, 2004]. In order to reason about multiple actions in this context, it is necessary to consider sequences of alternating belief revisions and belief updates. However, to date there has been little explicit discussion about the formal properties of such sequences.

This paper makes two contributions to the existing work on epistemic action effects. The first contribution is the development of a transition system framework suitable for reasoning about belief change. The new framework provides a simple tool for reasoning about revision and update in a single formalism, and it facilitates the treatment of conditional updates. The second contribution is the presentation of a principled approach to the interaction between revision and update. There are plausible examples in which agents appear to revise a prior belief state in response to a new observation; such examples are difficult to represent in existing formalisms.

2 Background and Motivation

2.1 Belief Change

We distinguish two kinds of belief change. Belief revision occurs when an agent receives new information about a static world. The original approach to belief revision is the AGM approach [Alchourron *et al.*, 1985]. Due to limitations on

space, we do not review the approach here; instead, we simply assume that the reader is familiar with the AGM postulates. Belief update, on the other hand, is the process in which the beliefs of an agent are modified in response to a world that has changed. One standard approach to belief update is the Katsuno and Mendelzon approach, which follows the AGM tradition by introducing a set of rationality postulates for belief update [Katsuno and Mendelzon, 1992].

2.2 Reasoning about Action

We introduce some standard terminology for describing transition systems [Gelfond and Lifschitz, 1998]. An action signature is a pair $\langle \mathbf{F}, \mathbf{A} \rangle$ of non-empty sets, respectively called the set of *fluent symbols* and the set of *action symbols*. Informally, fluent symbols are propositional variables representing properties of the world that may change over time and action symbols are atomic symbols representing deterministic actions that may be performed. A *state* is an interpretation over \mathbf{F} , and a *transition system* is a directed graph where each node is labeled by a state and each edge is labeled by a set of action symbols. The edges in a transition system indicate how the fluents change values in response to the execution of individual actions. We use the capital letter A , possibly with subscripts, to range over actions. The notation \bar{A} will be used to denote a finite sequence of actions of indeterminate length.

A *belief state* is a set of states. We can think of a belief state as expressing a proposition. Informally, a belief state is the set of states that an agent considers possible. In this paper, we use lower case Greek letters to denote belief states. We are interested in belief change resulting from two distinct kinds of actions: *ontic actions* and *epistemic actions*. Ontic actions are actions that change the state of the world, whereas epistemic actions change the beliefs of an agent without altering the state of the world. Informally, after executing an ontic action, an agent should update the current belief state on a point-wise basis. On the other hand, if an agent executes an epistemic action, then the current belief state should be revised.

2.3 The Basic Problem

As noted above, belief update occurs when an ontic action is performed. Hence, we define belief update operators that take two arguments: a belief state and an ontic action. Epistemic actions are identified with *observations*, which are simply sets of interpretations. As a result, revision operators also

take two arguments, each of which is a set of interpretations. Informally, the observation α provides evidence that the actual world is in α . Let \diamond be an update operator and let $*$ be a revision operator. We are interested in giving a reasonable interpretation to sequences of the form

$$\kappa \diamond A_1 * \alpha_1 \diamond \dots \diamond A_n * \alpha_n.$$

There are intuitively plausible examples in which applying the operators iteratively results in an unsatisfactory result. In the next section, we introduce one such example.

2.4 Illustrative Example

We extend the litmus paper problem originally presented in [Moore, 1985]. In the original problem, there is a beaker containing either an acid or a base, and there is an agent holding a piece of litmus paper that can be dipped into the beaker to determine the contents. The litmus paper turns red if it is placed in an acid and it turns blue if it is placed in a base. We extend the problem by admitting the possibility that the paper is not litmus paper, instead it is just plain white paper.

Boutilier points out that the standard approach to belief change provides an unintuitive representation of this problem [Boutilier, 1995]. One issue is that the standard approach does not allow conditional action effects; an agent simply updates the initial belief state by the new color of the litmus paper. Intuitively, this seems incorrect because there are actually two independent belief changes that occur. First, the agent dips the paper in the beaker and projects each possible world to the outcome of the dipping action. Second, the agent looks at the paper and observes the new color. Hence, the problem involves an update followed by a revision.

Even if the belief change is broken into two steps, the standard approach is limited in that an agent can only revise the current belief state. Suppose that an agent initially believes that the paper is litmus paper, but then it remains white after dipping it in the beaker. In this case, the agent should conclude that the paper was never litmus paper to begin with. This indicates that it is sometimes necessary for agents to revise prior belief states in the face of new knowledge.

We will return to this example periodically as we introduce formal machinery that provides a more natural representation.

3 A Transition System Approach

3.1 Metric Transition Systems

Standard transition systems do not provide a useful basis for performing belief revision, because belief revision generally requires some notion of plausibility among states or formulas. In order to define a revision operator, we introduce a distance function between states. A *metric* over $2^{\mathbf{F}}$ is a function d that maps each pair of states to a non-negative real number, and satisfies the following conditions:

1. $d(w_1, w_2) = 0$ iff $w_1 = w_2$
2. $d(w_1, w_2) = d(w_2, w_1)$
3. $d(w_1, w_2) + d(w_2, w_3) \geq d(w_1, w_3)$.

We will generally be concerned with metrics that return only integral distances, so from here on we will use the term *metric* to refer to integer-valued metrics.

Definition 1 A *metric transition system* is a triple $\langle S, R, d \rangle$ where

1. $S \subseteq 2^{\mathbf{F}}$
2. $R \subseteq S \times \mathbf{A} \times S$
3. d is a metric on S

Informally, if w_1 is close to w_2 , then an agent considers w_1 to be a plausible alternative to w_2 .

3.2 Belief Update

In this section, we define belief update with respect to a transition system. Recall that we update a belief state by an action with effects given by a transition system. This contrasts with the standard approach, in which a belief state is updated by a formula. The advantage of our approach is that it facilitates the representation of actions with conditional effects.

Intuitively, after executing an action A , an agent updates the belief state by projecting every state s to the state s' that would result if they executed A in s .

Definition 2 Let $T = \langle S, R, d \rangle$ be a metric transition system. The *update function* $\diamond : 2^S \times \mathbf{A} \rightarrow 2^S$ is defined as follows

$$\alpha \diamond A = \{f \mid (e, A, f) \in R \text{ for some } e \in \alpha\}.$$

Note that the distance function d does not play any role in belief update.

We briefly illustrate that transition systems can be used to define a standard update operator in which a belief state is updated by a formula rather than an action. Given \mathbf{F} , let \mathbf{A} be the set of action symbols of the form A_ϕ , where ϕ is a conjunction of literals over \mathbf{F} . Define $T_{\mathbf{F}}$ to be the transition system with $S = 2^{\mathbf{F}}$ and $RsA_\phi s'$ just in case $s' \models \phi$ and s' agrees with s on every atom that does not occur in ϕ . Informally, the action A_ϕ corresponds to an update by ϕ .

Proposition 1 The update operator obtained from $T_{\mathbf{F}}$ satisfies the Katsuno and Mendelzon postulates.

3.3 Belief Revision

With each metric transition system T , we associate a revision function. The revision function associated with T is the distance-based revision function defined in [Delgrande, 2004]. We choose this approach because it requires the introduction of relatively little formal machinery and, provided that d satisfies some natural conditions, this revision operator satisfies the AGM postulates.

Definition 3 Let $T = \langle S, R, d \rangle$ be a metric transition system. The *revision function* $*$: $2^S \times 2^S \rightarrow 2^S$ is defined as follows

$$\kappa * \alpha = \{w \in \alpha \mid \exists v_1 \in \kappa, \forall v_2 \in \alpha, \forall v_3 \in \kappa, d(w, v_1) \leq d(v_2, v_3)\}.$$

Hence, if an agent is in belief state κ , then $\kappa * \alpha$ is the set of all worlds that are minimally distant from some world in κ .

3.4 Litmus Paper Revisited

We revisit the litmus paper problem in the context of metric transition systems. The problem can be represented with the following action signature:

$$\langle \{Red, Blue, Acid, Litmus\}, \{dip\} \rangle.$$




Figure 1: Extended Litmus Paper Problem

Intuitively, *Red* is true if the paper is red, *Blue* is true if the paper is blue, *Acid* is true if the beaker contains an acid, and *Litmus* is true if the paper is litmus paper. The effects of dipping are given by the transition system in Figure 1. We can extend this to a metric transition system by defining d to be the Hamming distance between states. Note also that we assume all actions are executable in all states, but self-loops are omitted from Figure 1 for ease of readability.

Recall that we are interested in an agent that initially believes they are holding a piece of white litmus paper. Hence, the initial belief state E is

$$E = \{\{Litmus\}, \{Litmus, Acid\}\}.$$

After dipping the paper, we update the belief state:

$$E \diamond dip = \{\{Litmus, Blue\}, \{Litmus, Red, Acid\}\}.$$

At this point, the agent observes that the paper is neither blue nor red. This observation is represented by the following set:

$$O = \{\emptyset, \{Litmus\}, \{Acid\}, \{Litmus, Acid\}\}.$$

The naive suggestion is to simply revise $E \diamond dip$ by O , giving

$$E' = \{\{Litmus\}, \{Litmus, Acid\}\}.$$

However, this is not a plausible final belief state. Given the transition system in Figure 1, it is clear that white litmus paper can not remain white after a dipping action. Hence, successively applying the update and revision operators obtained from the transition system does not yield a plausible result.

4 Rationality Postulates

In this section, we give a set of rationality postulates specifying some natural properties that we expect to hold when an update is followed by a revision. The postulates are not overly restrictive and they do not provide a basis for a categorical semantics; they simply provide a point for discussion and comparison. Our underlying intuition is that the most recent observation should always be incorporated, provided that it is consistent with the history of actions that have been executed. Hence, the postulates are only appropriate for action domains in which there are no failed actions.

Assume a fixed propositional signature F . Let κ and α be sets of worlds, and let \bar{A} be a sequence of actions. We adopt the shorthand notation $\kappa \diamond \bar{A}$ as an abbreviation for the sequential update of κ by each element of \bar{A} . The following postulates describe some basic properties of the interaction between an update operator \diamond and a revision operator $*$.

1. If $(2^F \diamond \bar{A}) \cap \alpha \neq \emptyset$, then $(\kappa \diamond \bar{A}) * \alpha \subseteq \alpha$
2. If $(2^F \diamond \bar{A}) \cap \alpha = \emptyset$, then $(\kappa \diamond \bar{A}) * \alpha = \kappa \diamond \bar{A}$
3. $(\kappa \diamond \bar{A}) \cap \alpha \subseteq (\kappa \diamond \bar{A}) * \alpha$
4. If $(\kappa \diamond \bar{A}) \cap \alpha \neq \emptyset$, then $(\kappa \diamond \bar{A}) * \alpha \subseteq (\kappa \diamond \bar{A}) \cap \alpha$
5. $(\kappa \diamond \bar{A}) * \alpha \subseteq 2^F \diamond \bar{A}$

We give some intuitive motivation for each postulate.

Postulate 1 is a straightforward AGM-type assertion that α must hold after revising by α , provided α is possible after executing \bar{A} . Postulate 2 handles the situation where it is impossible to be in an α -world after executing \bar{A} . In this case, we simply discard the observation α . These postulates together formalize the underlying assumption that there are no failed actions.

Taken together, Postulates 3 and 4 assert that revising by α is equivalent to taking the intersection with α , provided the intersection is non-empty. These postulates are similar to the AGM postulates asserting that revisions correspond to expansions, provided the observation is consistent with the knowledge base.

Postulate 5 provides the justification for revising prior belief states in the face of new knowledge. The postulate asserts that, after revising by α , we must still have a belief state that is a possible consequence of executing \bar{A} . In some cases, the only way to assure that α holds after executing \bar{A} is to modify the initial belief state. We remark that the postulates do not indicate how the initial belief state should be modified.

5 Belief Evolution

5.1 Representing Histories

Transition systems are only suitable for representing Markovian action effects; that is, effects that are determined entirely by the current state and the action executed. However, examples like the litmus paper problem indicate that sometimes agents need to look at prior belief states as well. Hence, even if ontic action effects are Markovian, it does not follow that changes in belief are Markovian. As such, we need to introduce some formal machinery for representing histories.

Definition 4 A belief trajectory of length n is an n -tuple $\langle \kappa_0, \dots, \kappa_{n-1} \rangle$ of belief states.

Intuitively, a belief trajectory is an agent's subjective view of how the world has changed. We remark that a belief trajectory represents the agent's current beliefs about the world

history, not a historical account of what an agent believed at each point in time.

We will also be interested in observation trajectories and action trajectories, each of which is simply another n -tuple.

Definition 5 An observation trajectory of length n is an n -tuple $OBS = \langle OBS_1, \dots, OBS_n \rangle$ where each $OBS_i \in 2^S$.

Definition 6 An action trajectory of length n is an n -tuple $ACT = \langle ACT_1, \dots, ACT_n \rangle$ where each $ACT_i \in \mathbf{A}$.

Each set OBS_i is interpreted to be evidence that the actual world is in OBS_i at time i . An action trajectory is a history of the actions an agent has executed. Note that, as a matter of convention, we start the indices at 0 for belief trajectories and we start the indices at 1 for observation and action trajectories. The rationale for this convention will be clear later. We also adopt the convention hinted at in the definitions, whereby the n^{th} component of an action trajectory OBS will be denoted by OBS_n , and the n^{th} component of an action trajectory ACT will be denoted by ACT_n .

We define a notion of consistency between action trajectories and observation trajectories. The intuition is that an action trajectory ACT is consistent with an observation trajectory OBS if and only if each observation OBS_i is possible, given that the actions $(ACT_j)_{j \leq i}$ have been executed.

Definition 7 Let $ACT = \langle ACT_1, \dots, ACT_n \rangle$ be an action trajectory and let $OBS = \langle OBS_1, \dots, OBS_n \rangle$ be an observation trajectory. We say that ACT is consistent with OBS if and only if there is a belief trajectory $\langle \kappa_0, \dots, \kappa_n \rangle$ such that, for all $i \geq 1$,

1. $\kappa_i \subseteq OBS_i$
2. $\kappa_i = \kappa_{i-1} \diamond ACT_i$.

If ACT is consistent with OBS , we write $ACT||OBS$.

A pair consisting of an action trajectory and an observation trajectory gives a complete history of all actions that have occurred. As such, it is useful to introduce some terminology.

Definition 8 A world view of length n is a pair $W = \langle ACT, OBS \rangle$, where OBS is an observation trajectory and ACT is an action trajectory, each of length n .

If $ACT||OBS$, we say that $\langle ACT, OBS \rangle$ is consistent.

5.2 A New Belief Change Operator

We introduce a new operator \circ that takes two arguments: a belief state and a world view. Roughly speaking, we would like $\kappa \circ \langle ACT, OBS \rangle$ to be the belief trajectory that results from the initial belief state κ and the alternating action-observation sequence $ACT_1, OBS_1, \dots, ACT_n, OBS_n$. We call \circ a belief evolution operator because it takes a sequence of actions, and returns the most plausible evolution of the world.

The formal definition of \circ is presented in the following sections. The definition relies on a fixed revision operator $*$ and a fixed update operator \diamond . As such, it might be more accurate to adopt notation of the form $\circ_{*,\diamond}$, but we opt for the less cumbersome \circ and assume that the underlying operators are clear from the context. We remark, in particular, that every finite metric transition system generates a unique belief evolution operator. However, it is worth noting that the definition of \circ does not rely on any specific approach to revision.

The action domains of interest for belief evolution will be those in which it is reasonable to assume that action trajectories are correct and actions are successful. This is intuitively plausible in a single agent environment, because it simply amounts to assuming that an agent has complete knowledge about the actions that they have executed. Hence, in the definition of \circ , the belief trajectory returned will always be consistent with the actions that have been executed.

5.3 Infallible Observations

In this section, we assume that observations are always correct. Formally, this amounts to a restriction on the class of admissible world views. In particular, we need not consider inconsistent world views.

We need to introduce a bit of notation.

Definition 9 Let T be a transition system; let $\bar{A} = A_0, \dots, A_n$; and let α be a set of states. Then $\alpha^{-1}(\bar{A})$ denotes the set of all w such that there is a path from w to an element of α following the edges A_0, \dots, A_n .

Hence, $\alpha^{-1}(\bar{A})$ is the set of states that can precede a world in α , given that the sequence \bar{A} has been executed.

For illustrative purposes, it is useful to consider world views of length 1. Suppose we have an initial belief state κ , an ontic action A and an epistemic action α . Without formally defining the belief evolution operator \circ , we can give an intuitive interpretation of an expression of the form

$$\kappa \circ \langle \langle A \rangle, \langle \alpha \rangle \rangle = \langle \kappa_0, \kappa_1 \rangle.$$

The agent knows that the actual world is in α at the final point in time, so we must have $\kappa_1 \subseteq \alpha$. Moreover, the agent should believe that κ_1 is a possible result of executing A from κ_0 . In other words, we must have $\kappa_0 \subseteq \alpha^{-1}(A)$. All other things being equal, the agent would like to keep as much of κ as possible; therefore, the natural solution is the following:

1. $\kappa_0 = \kappa * \alpha^{-1}(A)$,
2. $\kappa_1 = \kappa_0 \diamond A$.

This procedure can be applied to world views of length greater than 1. The idea is to trace every observation back to a precondition on the initial belief state. After revising the initial belief state by all preconditions, each subsequent belief state can be determined by a standard update operation.

We have the following formal definition for \circ . In the definition, let \overline{ACT}_i denote the subsequence of actions ACT_1, \dots, ACT_i .

Definition 10 Let κ be a belief state, let ACT be an action trajectory of length n and let OBS be an observation trajectory of length n such that $ACT||OBS$. Define

$$\kappa \circ \langle ACT, OBS \rangle = \langle \kappa_0, \dots, \kappa_n \rangle$$

where

1. $\kappa_0 = \kappa * \bigcap_i OBS_i^{-1}(\overline{ACT}_i)$
2. for $i \geq 1$, $\kappa_i = \kappa_{i-1} \diamond ACT_1 \diamond \dots \diamond ACT_i$.

We remark that the intersection of observation preconditions in the definition of κ_0 is non-empty, because $ACT||OBS$.

The following propositions are immediate, and they demonstrate that \circ subsumes both revision and update. In each proposition, we assume that $ACT||OBS$.

Proposition 2 Let κ be a belief state, let $ACT = \langle A \rangle$ and let $OBS = \langle 2^F \rangle$. Then

$$\kappa \circ \langle ACT, OBS \rangle = \langle \kappa, \kappa \diamond A \rangle.$$

In the following, we assume that λ is a null action that never changes the state of the world.

Proposition 3 Let κ be a belief state, let $ACT = \langle \lambda \rangle$ and let $OBS = \langle \alpha \rangle$. Then

$$\kappa \circ \langle ACT, OBS \rangle = \langle \kappa * \alpha, \kappa * \alpha \rangle.$$

Hence, both revision and update can be represented through the \circ operator. As such, it is reasonable to adopt the following notation for trajectories of length 1:

$$\kappa \diamond ACT * OBS = \kappa \circ \langle \langle ACT \rangle, \langle OBS \rangle \rangle.$$

Proposition 4 If $ACT \parallel OBS$, the operators $*$ and \diamond (defined above) satisfy the interaction postulates (1)-(5).

The three preceding propositions demonstrate the suitability of \circ as a natural operator for reasoning about the interaction between revision and update.

5.4 Fallible Observations

We address fallible observations by allowing inconsistent world views. If a world view is inconsistent, then there is no initial belief state that supports all of the observations, so some observations need to be ignored. To deal with inconsistency, we adopt the convention previously explored in [Nayak, 1994] and [Papini, 2001], in which more credence is given to the most recent observations. We demonstrate that \circ can be extended naturally to represent belief change under this convention. In order to state the extended definition, we need to be able to extract a maximally consistent sub-view from an inconsistent world view.

Definition 11 Let $W = \langle ACT, OBS \rangle$ be a world view of length n . Define $\tau(W) = \langle ACT, OBS' \rangle$, where $OBS' = \langle OBS'_1, \dots, OBS'_n \rangle$ is defined by the following recursion.

- If $OBS_n^{-1}(ACT) \neq \emptyset$ then $OBS'_n = OBS_n$, otherwise $OBS'_n = 2^F$.
- For $i < n$, if

$$OBS_i^{-1}(\overline{ACT_i}) \cap \bigcap_{j>i} (OBS'_j)^{-1}(\overline{ACT_j}) \neq \emptyset$$

$$\text{then } OBS'_i = OBS_i, \\ \text{otherwise } OBS'_i = 2^F.$$

The observations in $\tau(W)$ are determined by starting with the most recent observation, then working backwards through the observations from most recent until the initial observation. At each point, we keep an observation if it is consistent with the observations that followed it; otherwise, we discard the observation as incorrect.

The following properties are immediate.

- If W is a world view, then $\tau(W)$ is consistent.
- If W is a consistent world view, then $\tau(W) = W$.

Recall that the original definition of \circ applied only to consistent world views. By passing through τ , we can extend the definition to apply to arbitrary world views.

Definition 12 Let κ be a belief state, and let W be a world view of length n . If W is inconsistent, then $\kappa \circ W = \kappa \circ \tau(W)$.

We could equivalently have stated a single definition for \circ by passing all world views through τ . We have presented the definition in two cases in order to highlight the distinct treatment of fallible observations. The extended definition is still satisfactory from the perspective of the rationality postulates.

Proposition 5 The operators $*$ and \diamond (obtained from \circ) satisfy the interaction postulates (1)-(5).

Thus far, applying the \circ operator requires tracing action preconditions back to the initial state for revision, then applying action effects to get a complete history. If we are only concerned with the final belief state, then there are many cases in which we do not need to go to so much effort.

Proposition 6 Let κ be a belief state, let ACT be an action trajectory of length n and let α be a belief state such that $\alpha \subseteq \kappa \diamond ACT$. If OBS is the observation trajectory with $n - 1$ null observations followed by α , then the final belief state in $\kappa \circ \langle ACT, OBS \rangle$ is $(\kappa \diamond ACT) * \alpha$.

The proposition indicates that, given a single observation that is consistent with the actions that have been executed, we can simply revise the outcome of the actions and we get the correct final belief state.

5.5 Litmus Paper Concluded

We conclude the litmus paper example by giving a plausible treatment based on a belief evolution operator. The world view $WV = \langle \langle dip \rangle, \langle O \rangle \rangle$ represents a dipping action followed by the observation that the paper is still white. If \circ is defined from the metric transition system in Figure 1, the final belief state in $E \circ WV$ is given by

$$\begin{aligned} E * O^{-1}(dip) \diamond dip &= E * \{\emptyset, \{Acid\}\} \diamond dip \\ &= \{\emptyset, \{Acid\}\}. \end{aligned}$$

This calculation is consistent with our original intuitions, in that the agent revises the initial belief state before updating by the *dip* action. This ensures that we will have a final belief state that is a possible outcome of dipping. Moreover, the initial belief state is revised by the pre-image of the final observation, which means it is modified as little as possible while still guaranteeing that the final observation will be feasible. Note also that the final belief state given by this calculation is intuitively plausible. It simply indicates that the contents of the beaker are still unknown, but the agent now believes the paper is not litmus paper. Hence, a belief evolution operator employs a plausible procedure and returns a desirable result.

6 Relationship with Iterated Revision

If the null action is the only action permitted, then belief evolution is closely related to iterated revision. We briefly consider the suitability of belief evolution operators for reasoning about iterated revision. In the following proposition, let $\bar{\lambda}$ denote a sequence of null actions of indeterminate length.

Proposition 7 For any κ and OBS , there is a unique belief state κ' such that $\kappa \circ \langle \bar{\lambda}, OBS \rangle = \langle \kappa', \dots, \kappa' \rangle$.

This result is consistent with the view that belief evolution operators return a trajectory representing an agent’s current beliefs about the evolution of the world. We remark that, in general, κ' is not obtained by successively revising by the elements of OBS . Moreover, we claim that this is appropriate for action domains in which recency determines the plausibility of an observation.

Darwiche and Pearl present a set of four postulates for iterated revision [Darwiche and Pearl, 1997]. It is easy to verify that many belief revision operators do not satisfy the Darwiche-Pearl postulates if applied in succession. However, if we define the iterated revision $\kappa * OBS_1 * \dots * OBS_n$ to be the unique belief state in $\kappa \circ \langle \bar{\lambda}, OBS \rangle$, then we get the following result.

Proposition 8 *Iterated revision, as defined by any belief evolution operator \circ , satisfies all four Darwiche-Pearl postulates.*

Hence, given any belief revision operator, we can formulate an adequate approach to iterated revision by passing to the corresponding evolution operator.

7 Discussion

We make some brief remarks about related formalisms for reasoning about epistemic action effects. Our formalism is similar to the multi-agent belief structures of [Herzig *et al.*, 2004] in that both approaches combine revision with update based on actions with conditional effects. However, multi-agent belief structures do not consider any non-elementary interaction between revision and update. On the other hand, the Situation Calculus approach of [Shapiro *et al.*, 2000] implicitly handles the interaction of revision and update correctly. The “revision actions” employed in that formalism are provably equivalent to belief evolution operators. Correctly identifying these actions as evolution operators improves the intuitive plausibility of iterated belief change in the Situation Calculus approach.

The framework presented in this paper requires an agent to believe the most recent observation. A more general approach would attach a plausibility value to every observation, and then an agent would keep the most plausible observations. We are currently working on such a generalization. Similarly, we are looking at attaching plausibilities to action histories in order to represent domains in which an agent is uncertain about the actions that have been executed.

8 Conclusion

We have presented a transition system framework for reasoning about the epistemic effects of actions. We identified ontic action effects with belief update and epistemic action effects with belief revision, and we focused on the interaction between iterated update and revision operators. We illustrated by example that the interaction between update and revision can be non-elementary, so we proposed a set of rationality postulates restricting the interaction and demonstrated that the postulates all hold in our transition system framework. By contrast, existing formalisms for reasoning about epistemic action effects either ignore the interaction between

revision and update or they deal with it implicitly. Hence, our formalism contributes to the existing work on epistemic action effects in two ways. First, it is able to provide simple object level representations of action domains in which prior belief states need to be revised. Second, it provides an explicit treatment of the interaction between revision and update, which has not always been salient in related formalisms.

References

- [Alchourron *et al.*, 1985] 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.
- [Boutilier, 1995] C. Boutilier. Generalized update: Belief change in dynamic settings. In *Proceedings of IJCAI 1995*, pages 1550–1556, 1995.
- [Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997.
- [Delgrande, 2004] J.P. Delgrande. Preliminary considerations on the modelling of belief change operators by metric spaces. In *Proceedings of NMR 2004*, pages 118–125, June 2004.
- [Gelfond and Lifschitz, 1998] M. Gelfond and V. Lifschitz. Action languages. *Linköping Electronic Articles in Computer and Information Science*, 3(16):1–16, 1998.
- [Herzig *et al.*, 2004] A. Herzig, J. Lang, and P. Marquis. Revision and update in multi-agent belief structures. In *Proceedings of LOFT 6*, 2004.
- [Katsuno and Mendelzon, 1992] 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.
- [Moore, 1985] R.C. Moore. A formal theory of knowledge and action. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Commonsense World*, pages 319–358. Ablex Publishing, 1985.
- [Nayak, 1994] A.C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.
- [Papini, 2001] 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.
- [Shapiro *et al.*, 2000] 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.
- [van Ditmarsch, 2002] H.P. van Ditmarsch. Descriptions of game actions. *Journal of Logic, Language and Information (JoLLI)*, 11:349–365, 2002.