

Belief Update Revisited

Jérôme Lang

IRIT – CNRS / Université Paul Sabatier
31062 Toulouse Cedex, France
lang@irit.fr

Abstract

Although many papers about belief update have been written, its precise scope still remains unclear. In this paper we aim at identifying this scope, and we show that belief update is a specific case of feedback-free action progression. This strong connection with the field of reasoning about action leads us to reconsider belief update and investigate new issues, especially reverse update, which is to regression what update is to progression.

1 Introduction

In papers about belief change, one often reads sentences such as “belief revision consists in incorporating some new information about a static world, while belief update consists in incorporating into a belief base about an old state of the world a notification of some change in the world”. However, the distinction is not so simple. To assess the scope of belief change operators, we need to be able to talk about the properties of the system (the world and the available actions) and the properties of the agent’s state of knowledge, as in the taxonomy for reasoning about action and change from [Sandewall, 1995]. However, unlike reasoning about action, belief change processes have never (as far as we know) been analyzed from the point of view of such a taxonomy. A first step is taken towards this direction (for belief revision only) in [Friedman and Halpern, 1996].

Elaborating a full taxonomy for belief change and assessing the scope of each approach is an ambitious project that definitely needs more than one short article. Here we focus on a specific class of belief change operators, namely belief update. This paper aims at identifying its precise scope, i.e. the conditions (expressed by properties of the world and of the agent’s beliefs) under which update is a suitable process for belief change. The main claim of the paper (discussed in Section 3) is that updating a knowledge base by α corresponds to progressing it by a specific “purely physical”, feedback-free action “make α true” whose precise meaning depends on the chosen update operator. This in turn raises the following question, addressed in Section 4: if update is progression, are there belief change operators corresponding to regression? Section 5 briefly considers the question of whether updates

can be compiled into propositional action theories. Related issues are discussed in Section 6.

2 Background on belief update

Let L^V be the propositional language generated from a *finite* set of propositional variables V , the usual connectives and the Boolean constants \top, \perp . $S = 2^V$ is the set of *states* (i.e., propositional interpretations). For any $\varphi \in L^V$, $Mod(\varphi)$ is the set of states satisfying φ . For any $X \subseteq S$, $for(X)$ is the formula of L^V (unique up to logical equivalence) such that $Mod(for(X)) = X$. If $X = \{s\}$, we write $for(s)$ instead of $for(\{s\})$. We use $\varphi \oplus \psi$ as a shorthand for $\varphi \leftrightarrow \neg\psi$.

A belief update operator \diamond is a mapping from $L^V \times L^V$ to L^V , i.e., it maps two propositional formulas φ (the initial belief state) and α (the “input”) to a propositional formula $\varphi \diamond \alpha$ (the resulting belief state). We recall here the Katsuno-Mendelzon (KM for short) postulates for belief update [Katsuno and Mendelzon, 1991].

- U1 $\varphi \diamond \alpha \models \alpha$.
- U2 If $\varphi \models \alpha$ then $\varphi \diamond \alpha \equiv \varphi$.
- U3 If φ and α are both satisfiable then $\varphi \diamond \alpha$ is satisfiable.
- U4 If $\varphi \equiv \psi$ and $\alpha \equiv \beta$ then $\varphi \diamond \alpha \equiv \psi \diamond \beta$.
- U5 $(\varphi \diamond \alpha) \wedge \beta \models \varphi \diamond (\alpha \wedge \beta)$.
- U6 If $\varphi \diamond \alpha \models \beta$ and $\varphi \diamond \beta \models \alpha$ then $\varphi \diamond \alpha \equiv \varphi \diamond \beta$.
- U7 If φ is complete then $(\varphi \diamond \alpha) \wedge (\varphi \diamond \beta) \models \varphi \diamond (\alpha \vee \beta)$.
- U8 $(\varphi \vee \psi) \diamond \alpha \equiv (\varphi \diamond \alpha) \vee (\psi \diamond \alpha)$.

Although we have recalled all postulates for the sake of completeness, we should not accept them unconditionally. They have been discussed in several papers, including [Herzig and Rifi, 1999] in which it was argued that not all these postulates should be required, and that the “uncontroversial” ones (those deeply entrenched in the very notion of update and satisfied by most operators studied in the literature) are (U1), (U3), (U8), and (U4) to a lesser extent. We therefore call a *basic update operator* any operator \diamond from $L^V \times L^V$ to L^V satisfying at least (U1), (U3), (U4) and (U8). In addition, \diamond is said to be *inertial* if it also satisfies (U2), and \diamond is a *KM update operator* if it satisfies (U1)-(U8)¹. In the paper we refer to some specific update operators such as the

¹(U5), (U6) and (U7) are much more controversial than the other

PMA [Winslett, 1990] that we have no space to recall (see for instance [Herzig and Rifi, 1999]).

3 Belief update as action progression

The difference between the scope of revision and that of update is often expressed as a static vs. dynamic opposition. However, *nothing in the AGM theory of belief revision implies that we should restrict its application to static worlds.* As remarked in [Friedman and Halpern, 1999], what is essential in belief revision is not that the world is static, but that *the language used to describe the world is static.* Thus, if an evolving world is represented using time-stamped propositional variables of the form v_t (v true at time t), we can perfectly revise a belief set by some new information about the past or the present, and infer some new beliefs about the past, the present, or even the future.

Example 1 *On Monday, Alice is the head of the computer science lab while Bob is the head of the math lab. On Tuesday I learned that one of them resigned (but I don't know which one). On Wednesday I learn that Charles is now the head of the math lab, which implies that Bob isn't. (It is implicit that heads of labs tend to keep their position for quite a long time.) What do I believe now?*

Example 1 contains a sequence of two “changes”. Both are detected by observations, and the whole example can be expressed as a revision process (with time-stamped variables). Let us identify Monday, Tuesday and Wednesday by the time stamps 1, 2 and 3. On Monday I believe A_1, B_1 , as well as the persistency laws $A_1 \leftrightarrow A_2, A_2 \leftrightarrow A_3, B_1 \leftrightarrow B_2$ etc., therefore I also believe A_2, B_2 etc.: I expect that Alice and Bob will remain the heads of their respective labs on Tuesday and Wednesday. The revision by $\neg A_2 \vee \neg B_2$ (provided that the revision operator minimizes change) leads me to believe $A_1, B_1, A_2 \oplus B_2, A_3 \oplus B_3$ etc.: on Tuesday, I still believe that Alice and Bob were heads of their labs on Monday, and that now *exactly one* of them is. Then the revision by $\neg B_3$ (at time 3) makes me believe $A_1, B_1, A_2, \neg B_2, A_3, \neg B_3$: on Wednesday, I understand that Bob was the one to resign on Tuesday, and therefore that Alice was still head of the CS lab on Tuesday, and is still now².

Now, the fact that belief revision can deal with (some) evolving worlds suggests that claiming that belief update is the right belief change operation for dealing with evolving worlds is insufficient and ambiguous. The literature on belief update abounds in ambiguous explanations such as “update consists in bringing the knowledge base up to date when the world is described by its changes”³. There is a first ambiguity in the expressions “describing the world by its changes”

ones (see [Herzig and Rifi, 1999]); they characterize the specific class of updates based on a similarity-based minimization process (which is known to lead to several counterintuitive results).

²Note that this scenario is also a case for belief extrapolation [Dupin de Saint-Cyr and Lang, 2002], which is a particular form of time-stamped revision.

³This formulation appears in [Katsuno and Mendelzon, 1991], which may be one of the explanations for such a long-lasting ambiguity.

or “notification of a change”, as “change” actually has to be understood as “possibility of change” (we’ll come back to this point). But the main problem is the status of the input formula α . To make things clear, here is an example.

Example 2 *My initial belief is that either Alice or Bob is in the office (but not both). Both tend to stay in the office when they are in. Now I see Bob going out of the office. What do I believe now?*

Trying to use belief update to model this example is hopeless. For all common update operators seen in the literature, updating $A \oplus B$ by $\neg B$ leads to $\neg B$, and not to $\neg A \wedge \neg B$. Indeed, because of (U8), we have $(A \oplus B) \diamond \neg B \equiv [(A \wedge \neg B) \diamond \neg B] \vee [(\neg A \wedge B) \diamond \neg B] \equiv (A \wedge \neg B) \vee (\neg A \wedge \neg B) \equiv \neg B$. The only way to have $\neg A \wedge \neg B$ as the result would be to have $(A \wedge \neg B) \diamond \neg B \equiv \neg A \wedge \neg B$, which can hold only if there is a causal relationship between A and B , such as B becoming false entails A becoming false – which is not the case here.

Example 2 definitely deals with an evolving world and contains a “notification of change”, and still it cannot be formulated as a belief update process. On the other hand, like Example 1, it can be perfectly expressed as is a time-stamped belief revision process⁴.

The key point is (U8) which, by requiring that all models of the initial belief set be updated separately, *forbids us from inferring new beliefs about the past from later observations*: indeed, in Example 2, belief update provides no way of eliminating the world $(A, \neg B)$ from the set of *previously* possible worlds, which in turn, does not allow for eliminating $(A, \neg B)$ from the list of possible worlds *after* the update: if $(A, \neg B)$ is a possible world at time t , then its update by $\neg B$ *must* be in the set of possible worlds at time $t + 1$. In other terms, update fails to infer that Alice wasn't in the office and still isn't.

Belief update fails as well on Example 1: updating $A \wedge B \wedge \neg C$ by $\neg A \vee \neg B$ gives the intended result, but only by chance (because the agent's initial belief state is complete). The second step fails: with most common update operators, updating $(A \oplus B) \wedge \neg C$ by $\neg B \wedge C$ leads to $\neg B \wedge C$, while we'd expect to believe A as well.

The diagnosis should now be clear: the input formula α is not a mere observation. An observation made at time $t + 1$ leads to filter out some possible states at time $t + 1$, which in turn leads to filter out some possible states at time t , because the state of the world at time t and the state of the world at time $t + 1$ are correlated (by persistence rules or other dynamic rules⁵.) And finally, the successor worlds (at time $t + 1$) of these worlds at time t that update failed to eliminate can

⁴Note that without time stamps (and in particular within the framework of belief update), we cannot distinguish between “ B has become false” (as in “I see Bob go out of the office”) and “the world has evolved in such a way that B is now false” (as in “I now see Bob out of his office”). Anyway, for Example 2, the expected outcome is the same in both cases (provided that A and B are expected to persist with respect to the granularity of time considered).

⁵The only case where belief update could be compatible with interpreting α as an observation would therefore be the case where not the faintest correlation exists between the state of the world at different time points; in this case, we would have $\varphi \diamond \alpha \equiv \alpha$ whenever α is consistent – a totally degenerate and uninteresting case.

not be eliminated either. Such a backward-forward reasoning needs a proper generalization of update (and of revision), unsurprisingly called *generalized update* [Boutilier, 1998].

Now, α has to be understood as an *action effect*, and update is a particular form of *action progression* for *feedback-free actions*. Action progression (as considered in the literature of reasoning about action and logic-based planning) consists in determining the belief state obtained from an initial belief state after a given action is performed.

This connection between belief update and action progression was first mentioned in [del Val and Shoham, 1994], who argue that updating an initial belief state φ by a formula α corresponds to one particular action; they formalize such actions in a formal theory of actions based on circumscription, and their framework for reasoning action is then used to derive a semantics for belief update. The relationship between update and action progression appears (more or less explicitly) in several other papers, including [Liberatore, 2000b], who expresses several belief update operators in a specific action language. Still, the relationship between update and action progression still needs to be investigated in more detail.

We first need to give some background on reasoning about action. Generally speaking, an action A has two types of effects: an *ontic (or physical) effect* and an *epistemic effect*. For instance, if the action consists in tossing a coin, its ontic effect is that the next value of the fluent *heads* may change, whereas its epistemic effect is that the new value of the fluent is observed (this distinction between ontic and epistemic effects is classical in most settings). Complex actions (with both kinds of effects) can be decomposed into two actions, one being ontic and feedback-free, the other one being a purely epistemic (sensing) action.

The simplest model for a purely ontic (i.e., feedback-free) action A consists of a *transition graph* R_A on S .⁶ $R_A(s, s')$ means that s' is accessible from s after A . $R_A(s) = \{s' \mid R_A(s, s')\}$ is the set of states that can obtain after performing A in s . If $R_A(s)$ is a singleton for all s then A is *deterministic*. If $R_A(s) = \emptyset$ then A is *inexecutable* in s . A is *fully executable* iff $R_A(s) \neq \emptyset$ for every s .

Let A be a purely ontic action modelled by a transition graph R_A on S . For any formula $\varphi \in L^V$, the *progression of φ by A* is the propositional formula (unique up to logical equivalence) whose models are the states that can obtain after performing A in a state of $Mod(\varphi)$: $prog(\varphi, A)$ is defined by

$$prog(\varphi, A) = for \left(\bigcup_{s \models \varphi} R_A(s) \right) \quad (1)$$

Lastly, for any action A , $Inv(A)$ is the set of *invariant states* for A , i.e. the set of all states s such that $R_A(s) = \{s\}$.

Clearly enough, (1) is identical to (U8). Therefore, for any update operator (and more generally any operator satisfying (U8)) and any input formula α , *updating by α is an action progression operator*. This raises several questions: (a) Which action is this exactly? (b) What is the class of actions that correspond to updates? (c) If update is progression, are there belief change operators corresponding to regression?

⁶More sophisticated models may involve graded uncertainty such as probabilities, delayed effects etc.

Question (a) first. As argued above, (U8) and (1) mean that the action is feedback-free. This comes down to saying that belief update assumes *unobservability*: the set of possible states after A is performed is totally determined by the set of possible states before it is performed and the transition system corresponding to A . In other words, *what you foresee is what you get*: once we have decided to perform A , waiting until it has actually been performed will not bring us any new information. Note that using update in Example 2 would correspond to performing an action whose effect is to make Bob go out of his office (when he is initially not in the office, this action has no effect). Likewise, in Example 1, updating $A \oplus B \wedge \neg C$ by $\neg B \wedge C$ corresponds to performing the action “demote Bob from his position and appoint Charles instead”.

Therefore, updating by α is a purely ontic (feedback-free) action. Can we now describe this action in more detail? (U1) means that the action of updating by α has to be understood as “make α true”. More precisely, due to the absence of feedback reflected by (U8), updating φ by α could be understood as a dialogue between an agent and a robot: “All I know about the state of the world is that it satisfies φ . Please, go to the real world, see its real state, and whatever this state, act so as to change it into a world satisfying α , following some rules” (given that the robot does not communicate with the agent once it is the real world.) The rules to be followed by the robot are dictated by the choice of the update operator \diamond . If \diamond satisfies (U2), then the rules state that if the α is already true then the robot must leave the world as it is. If \diamond is the PMA [Winslett, 1990], then the rules are “make α true, without changing more variables than necessary”. More generally, when \diamond is a Katsuno-Mendelzon operator, associated with a collection of similarity preorders (one for each world), the robot should make α true by changing s into one of the states that are most similar to it⁷. When \diamond is WSS [Winslett, 1990; Herzig, 1996] or the MPMA [Doherty *et al.*, 1998], then the rules are “make α true, without changing the truth values of a given set of variables (those that do not appear in α , or those that play no role in α).” And so on.

Writing things more formally: given an update operator \diamond and a formula α , let $update(\diamond, \alpha)$ be the ontic action whose transition graph is defined by: for all $s, s' \in S$,

$$s' \in R_{update(\diamond, \alpha)}(s) \text{ iff } s' \models for(s) \diamond \alpha$$

The following characterizations are straightforward, but worth mentioning (and they will be useful later).

Proposition 1 *Let \diamond satisfy (U8).*

1. $\varphi \diamond \alpha \equiv prog(\varphi, update(\diamond, \alpha))$;
2. \diamond satisfies (U1) if and only if for any formula $\alpha \in L^V$ and any $s \in S$, $R_{update(\diamond, \alpha)}(s) \subseteq Mod(\alpha)$;
3. \diamond satisfies (U2) if and only if for any formula $\alpha \in L^V$, $Inv(update(\diamond, \alpha)) \supseteq Mod(\alpha)$;

⁷Note that, as argued in [Peppas *et al.*, 1996], this similarity has here be understood as an *ontological* notion (s being closer to s_1 than to s_2 may, in practice, reflect that from s it is easier to go to s_1 than to s_2) and not as an epistemic notion of similarity, as it would be the case for belief revision. This agrees with our view of $update(\diamond, \alpha)$ as an ontic action.

4. \diamond satisfies (U3) if and only if for any satisfiable formula $\alpha \in L^V$, $update(\diamond, \alpha)$ is fully executable.

Note that (U1) and (U2) imply that for all $s \in S$, $R_{update(\diamond, \alpha)}(s) \subseteq Inv(update(\diamond, \alpha))$. The other postulates do not have any direct effect on the properties of $update(\diamond, \alpha)$ considered as an isolated action, but they relate different actions of the form $update(\diamond, \alpha)$. Noticeably, requiring (U4) corresponds to the equality between $update(\diamond, \alpha)$ and $update(\diamond, \beta)$ when α and β are logically equivalent. The characterizations of (U5), (U6) and (U7) in terms of reasoning about action do not present any particular interest.

Let us now consider question (b). Obviously, given a fixed update operator \diamond satisfying (U1), (U3), (U4) and (U8), some fully executable actions are not of the form $update(\diamond, \alpha)$ (this is obvious because there are 2^{2^n} actions of the form $update(\diamond, \alpha)$ and 2^{n+2^n-1} fully executable actions, where $n = |V|$.) Now, what happens if we allow \diamond to vary? The question now is, what are the actions that can be expressed as $update(\diamond, \alpha)$, for some update operator \diamond and some α ?

Proposition 2 *Let A be a fully executable ontic action such that $R_A(s) \subseteq Inv(A)$ for all $s \in S$. Then there exists a KM-update operator, and a formula α , such that $A = update(\diamond, \alpha)$.*

The proof is constructive: α is taken such that $Mod(\alpha) = Inv(A)$, and the collection of faithful orderings in the sense of [Katsuno and Mendelzon, 1991] is defined by $s_1 <_s s_2$ if and only if $s = s_1 \neq s_2$ or $(s \neq s_1, s \neq s_2, s_1 \in R_A(s), s_2 \notin R_A(s))$; and $s_1 \leq_s s_2$ iff not $(s_2 <_s s_1)$.

From Propositions 1 and 2 we get

Corollary 1 *Let A be an ontic action. There exists a KM-update operator \diamond , and a formula α such that $A = update(\diamond, \alpha)$, if and only if A is fully executable and $R_A(s) \subseteq Inv(A)$ for all $s \in S$.*

A variant of Proposition 2 (and Corollary 1) can be obtained by not requiring $R_A(s) \subseteq Inv(A)$: in that case there exists an update operator \diamond satisfying all the KM postulates except (U3), and a formula α such that $A = update(\diamond, \alpha)$. α can be taken as \top and $s \leq_s s_2$ iff $s_1 \in R_A(s)$ or $s_2 \notin R_A(s)$.

4 Reverse update

Now, question (c). Is there a natural notion which is to action regression what update is to progression? The point is that we do not have one, but two notions of action regression. The *weak progression regression* (or *weak preimage*) of ψ by A is the formula whose models are the states from which the execution of A possibly leads to a model of ψ , while the *strong progression regression* (or *strong preimage*) of ψ by A is the formula whose models are the states from which the execution of A certainly leads to a model of ψ :

$$\begin{aligned} reg(\psi, A) &= form(\{s, R_A(s) \cap Mod(\psi) \neq \emptyset\}) \\ Reg(\psi, A) &= form(\{s, R_A(s) \subseteq Mod(\psi)\}) \end{aligned}$$

See for instance [Lang *et al.*, 2003] for the interpretation of these two notions of regression in reasoning about action. This naturally leads to two notions of reverse update.

Definition 1 *Let \diamond be an update operator.*

- the weak reverse update \odot associated with \diamond is defined by: for all $\psi, \alpha \in L^V$, for all $s \in S$,
 $s \models \psi \odot \alpha$ iff for $(s) \diamond \alpha \not\models \neg \psi$
- the strong reverse update \otimes associated with \diamond is defined by: for all $\psi, \alpha \in L^V$, for all $s \in S$,
 $s \models \psi \otimes \alpha$ iff for $(s) \diamond \alpha \models \psi$

Intuitively, weak reverse update corresponds to (deductive) postdiction: given that the action “make α true” has been performed and that we now know that ψ holds, what we can say about the state of the world before the update was performed is that it satisfied $\psi \odot \alpha$. As to strong reverse update, it is an abductive form of postdiction, better interpreted as goal regression: given that a rational agent has a goal ψ , the states of the world in which performing the action “make α true” is guaranteed to lead to a goal states are those satisfying $\psi \otimes \alpha$.

The following result shows that \odot and \otimes can be characterized in terms of \diamond :

Proposition 3

1. $\psi \odot \alpha \models \varphi$ iff $\neg \varphi \diamond \alpha \models \neg \psi$;
2. $\varphi \models \psi \otimes \alpha$ iff $\varphi \diamond \alpha \models \psi$;

As a consequence of Proposition 3, $\psi \odot \alpha$ is the *weakest* formula φ such that $\neg \varphi \diamond \alpha \models \neg \psi$, and $\psi \otimes \alpha$ is the *strongest* formula φ such that $\varphi \diamond \alpha \models \psi$.

Example 3 *Let $\diamond = \diamond PMA$ [Winslett, 1990]. Let b and m stand for “the book is on the floor” and “the magazine is on the floor”. The action $update(\diamond, b \vee m)$ can be described in linguistic terms by “make sure that the book or the magazine is on the floor”. Then $b \odot (b \vee m) \equiv b \vee (\neg b \wedge \neg m) \equiv b \vee \neg m$, which can be interpreted as follows: if we know that the book is on the floor after $update(\diamond, b \vee m)$ has been performed, then what we can say about the previous state of the world is that either the book was already on the floor (in which case nothing changed) or that neither the book nor the magazine was on the floor (and then the update has resulted in the book being on the floor). On the other hand, $b \otimes (b \vee m) \equiv b$: if our goal is to have the book on the floor, the necessary and sufficient condition for the action $update(\diamond, b \vee m)$ to be guaranteed to succeed is that the book is already on the floor (if neither of them is, the update might well leave the book where it is and move the magazine onto the floor).*

An interesting question is whether weak and strong reverse update can be characterized by some properties (which then would play the role that the basic postulates play for “forward” update). Here is the answer (recall that a basic update operator satisfies U1, U3, U4 and U8).

Proposition 4 *\odot is the weak reverse update associated with a basic update operator \diamond if and only if \odot satisfies the following properties:*

W1 $\neg \alpha \odot \alpha \equiv \perp$;

W3 if α is satisfiable then $\top \odot \alpha \equiv \top$;

W4 if $\psi \equiv \psi'$ and $\alpha \equiv \alpha'$ then $\psi \odot \alpha \equiv \psi' \odot \alpha'$;

W8 $(\psi \vee \psi') \odot \alpha \equiv (\psi \odot \alpha) \vee (\psi' \odot \alpha)$.

In addition to this, \diamond satisfies (U2) if and only if \odot satisfies

W2 $(\psi \odot \alpha) \wedge \alpha \equiv \psi \wedge \alpha$.

Note that (W4) and (W8) are *exactly* the same properties as (U4) and (U8), replacing \diamond by \odot . Properties (U5), (U6) and (U7) do not seem to have meaningful counterparts for \odot (and anyway, as already argued, these three postulates are too controversial).

We have a similar result for strong reverse update:

Proposition 5 \otimes is the strong reverse update associated with a basic update operator \diamond if and only if \otimes satisfies the following properties:

S1 $\alpha \otimes \alpha \equiv \top$;

S3 if α is satisfiable then $\perp \otimes \alpha \equiv \perp$;

S4 if $\psi \equiv \psi'$ and $\alpha \equiv \alpha'$ then $\psi \otimes \alpha \equiv \psi' \otimes \alpha'$;

S8 $(\psi \wedge \psi') \otimes \alpha \equiv (\psi \otimes \alpha) \wedge (\psi' \otimes \alpha)$.

In addition to this, \diamond satisfies (U2) if and only if \otimes satisfies

S2 if $\psi \models \alpha$ then $\psi \models \psi \otimes \alpha$.

Note that, unlike weak reverse update, strong reverse update does generally not satisfy modelwise decomposability (U8/W8), but a symmetric, conjunctive decomposability property (S8).

Moreover, if \diamond is a basic update operator then

SIW if α is satisfiable then $\psi \otimes \alpha \models \psi \odot \alpha$

Note that (SIW) fails without (U3). Example 3 shows that the converse implication of (SIW) does not hold in general. Finally, \otimes and \odot coincide if and only if $update(\diamond, \alpha)$ is deterministic.

One may wonder whether reverse update has something to do with erasure [Katsuno and Mendelzon, 1991]. An erasure operator \blacklozenge is defined from an update operator \diamond by $\psi \blacklozenge \alpha \equiv \psi \vee (\psi \diamond \neg \alpha)$. Erasing by α intuitively consists in making the world evolve (following some rules) such that after this evolution, the agent no longer believes α . A quick look suffices to understand that erasure has nothing to do with weak and strong reverse update. Erasure corresponds to action progression for an action $erase(\alpha)$ whose effect is epistemically negative (*make α disbelieved*). This implies in particular that $\top \blacklozenge \top$ is always unsatisfiable (\top cannot be made disbelieved) whereas $\top \odot \top \equiv \top \otimes \top \equiv \top$.

Pursuing the investigation on reverse update does not only have a theoretical interest: weak (deductive) reverse update allows for postdiction, and strong (abuctive) reverse update allows for goal regression (when the actions performed are updates) and is therefore crucial if we want to use an update-based formalism for planning (see [Herzig *et al.*, 2001]).

5 Update and propositional action theories

Now that we know that $update(\diamond, \alpha)$ is an action progression operator, can it be described by a propositional action theory such as those described in [Giunchiglia *et al.*, 2004]? Recall that a propositional action theory Σ_A for action A is a propositional formula Σ_A built on the language generated by the time-stamped propositional variables $V_t \cup V_{t+1}$, where $V_t = \{v_t | v \in V\}$ and $V_{t+1} = \{v_{t+1} | v \in V\}$, such that, for each $s_t \in 2^{V_t}$ and $s'_{t+1} \in 2^{V_{t+1}}$, $s_t \cup s'_{t+1}$ holds if and only if $s' \in R_A(s)$.

If we don't care about the size of the action theory, the answer is obviously positive: just take

$$\Sigma_{update(\diamond, \alpha)} = \bigwedge_{s \in S} for(s)_t \rightarrow (for(s) \diamond \alpha)_{t+1}.$$

Then for all s_t, s'_{t+1} , we have $s_t \cup s_{t+1} \models \Sigma_\alpha$ if and only if $for(s') \models for(s) \diamond \alpha$, and, equivalently, $\varphi \diamond \alpha \models \psi$ if and only if $\varphi_t \wedge \Sigma_\alpha \models \psi_{t+1}$.

The question is now whether $update(\diamond, \alpha)$ can be described by a *succinct* action theories. The answer depends on the update operator. Existing complexity results on belief update enable us to give a quick negative answer for many update operators. Note first that $prog(\varphi, \alpha) \models \psi$ if and only if $\varphi_t \wedge \Sigma_\alpha \models \psi_{t+1}$; therefore, given a propositional action theory Σ_A and two formulas φ, ψ , deciding whether $prog(\varphi, \alpha) \models \psi$ is in **coNP** (actually **coNP**-complete). Now, for many update operators, including the PMA, deciding whether $\varphi \diamond \alpha \models \psi$ is Π_2^P -complete [Eiter and Gottlob, 1992; Liberatore, 2000a]. Let \diamond be one of these update operators. If there were a way of expressing $update(\diamond, \alpha)$ (for all α) by a polynomially large action theory, then we would have $\Pi_2^P = \text{coNP}$, and the polynomial hierarchy would collapse at the first level, which is very unlikely.

On the other hand, a positive answer is obtained for the following update operators, for which deciding whether $prog(\varphi, \alpha) \models \psi$ is only **coNP**-complete: the MPMA [Doherty *et al.*, 1998], WSS [Winslett, 1990], WSS^\downarrow and WSS_{dep}^\downarrow [Herzig and Rifi, 1999]. All those update operators consists first in forgetting a set of variables (for instance, those that appear in the input formula α , or those in which α depends etc.), and then expanding the result by α . For instance, for $\diamond = \diamond_{WSS}$: $\Sigma(\alpha) = \alpha_{t+1} \wedge \bigwedge_{v \notin Var(\alpha)} v_t \leftrightarrow v_{t+1}$.

6 Summary and conclusion

Let us try to summarize what we have said so far. Both revision and update deal with dynamic worlds, but they strongly differ in the nature of the information they process. Belief revision (together with the introduction of time stamps in the propositional language) aims at correcting some initial beliefs about the past, the present, and even the future state of the world by some newly *observed* information about the past or the present state of the world. Belief update is suitable only for (some specific) action progression without feedback: updating φ by α corresponds to progressing (or projecting forward) φ by the action $update(\diamond, \alpha)$, to be interpreted as *make α true*. The “input formula” α is the effect of the action $update(\diamond, \alpha)$, and definitely not an observation. Expressed in the terminology of [Sandewall, 1995], the range of applicability of update is the class $K_p\text{-IA}$: correct knowledge⁸, no observations after the initial time point, inertia if (U2) is assumed, and alternative results of actions.

Assessing the scope of belief update is all the more important as several variations on belief update have recently appeared in the literature. [Liu *et al.*, 2006] lift belief update to description logics. [Eiter *et al.*, 2005] gives a framework for changing action descriptions, which is actually closer to a revision process than to update; however, *updating* an action by

⁸However, this point is somewhat debatable: update would work as well if we don't assume that the agent's initial beliefs is correct – of course, in this case the final beliefs may be wrong as well.

changing physically its transition graph would be meaningful as well, and deserves further study.

In complex environments, especially planning under incomplete knowledge, actions are complex and have both ontic and epistemic effects; the belief change process then is very much like the feedback loop in partially observable planning and control theory: perform an action, project its effects on the current belief state, then get the feedback, and revise the projected belief state by the feedback. Clearly, update allows for projection only. Or, equivalently, if one chooses to separate the ontic and the epistemic effects of actions, by having two disjoint sets of actions (ontic and epistemic), then ontic actions lead to projection only, while epistemic actions lead to revision only. Therefore, if one wants to extend belief update so as to handle feedback, there is no choice but integrating some kind of revision process, as in several recent works [Boutilier, 1998; Shapiro and Pagnucco, 2004; Jin and Thielscher, 2004; Hunter and Delgrande, 2005]. Another possibility is to generalize update so that it works in a language that distinguishes facts and knowledge, such as epistemic logic S5: this *knowledge update* process is investigated in [Baral and Zhang, 2005]. Here, effects of sensing actions are handled by *updating* (and not revising) formulas describing the agent’s knowledge. Such a framework takes the point of view of a modelling agent O who reasons on the state of knowledge of another agent ag . Thus, for instance, updating a S5 model by $K_{ag}\varphi$ means that the O updates her beliefs about ag ’s knowledge; considering ag ’s mental state as part of the *outside world* for agent O , this suits our view of update as a feedback-free action for O (updating by $K_{ag}\varphi$ corresponds as “make $K_{ag}\varphi$ true”, which can for instance be implemented by telling ag that φ is true).

Acknowledgements

Thanks to Andreas Herzig for years of discussions about the very meaning of belief update (without which I would never have thought of writing this paper). Thanks as well to Meghyn Bienvenu, Florence Dupin de Saint-Cyr, Sébastien Konieczny, Fangzhen Lin, and Pierre Marquis.

References

- [Baral and Zhang, 2005] C. Baral and Y. Zhang. Knowledge updates: Semantics and complexity issues. *Artificial Intelligence*, 164(1-2):209–243, 2005.
- [Boutilier, 1998] C. Boutilier. A unified model of qualitative belief change: a dynamical systems perspective. *Artificial Intelligence*, 1-2:281–316, 1998.
- [del Val and Shoham, 1994] A. del Val and Y. Shoham. Deriving properties of belief update from theories of action. *J. of Logic, Language, and Information*, 3:81–119, 1994.
- [Doherty et al., 1998] P. Doherty, W. Łukasiewicz, and E. Madalińska-Bugaj. The PMA and relativizing change for action update. In *Proc. KR’98*, pages 258–269.
- [Dupin de Saint-Cyr and Lang, 2002] F. Dupin de Saint-Cyr and J. Lang. Belief extrapolation (or how to reason about observations and unpredicted change). In *Proc. KR2002*, pages 497–508.
- [Eiter and Gottlob, 1992] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates and counterfactuals. *Artificial Intelligence*, 57:227–270, 1992.
- [Eiter et al., 2005] T. Eiter, E. Erdem, M. Fink, and J. Senko. Updating action domain descriptions. In *Proc. IJCAI 2005*, pages 418–423.
- [Friedman and Halpern, 1996] N. Friedman and J.Y. Halpern. Belief revision: A critique. In *Proc. KR’96*, pages 421–431.
- [Friedman and Halpern, 1999] N. Friedman and J.Y. Halpern. Modelling beliefs in dynamic systems. part ii: revision and update. *JAIR*, 10:117–167, 1999.
- [Giunchiglia et al., 2004] F. Giunchiglia, J. Lee, V. Lifschitz, N. Mc Cain, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 2004.
- [Herzig, 1996] A. Herzig. The PMA revisited. In *Proceedings of KR’96*, pages 40–50.
- [Herzig and Rifi, 1999] A. Herzig and O. Rifi. Propositional belief update and minimal change. *Artificial Intelligence*, 115:107–138, 1999.
- [Herzig et al., 2001] A. Herzig, J. Lang, P. Marquis, and T. Polacsek. Actions, updates, and planning. In *Proceedings of IJCAI-01*, pages 119–124.
- [Hunter and Delgrande, 2005] A. Hunter and J. Delgrande. Iterated belief change: a transition system approach. In *Proc. IJCAI05*.
- [Jin and Thielscher, 2004] Y. Jin and M. Thielscher. Representing beliefs in the fluent calculus. In *Proc. ECAI-04*, pages 823–827.
- [Katsuno and Mendelzon, 1991] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. KR’91*, pages 387–394.
- [Lang et al., 2003] J. Lang, F. Lin, and P. Marquis. Causal theories of action: a computational core. In *Proc. IJCAI-03*, pages 1073–1078.
- [Liberatore, 2000a] P. Liberatore. The complexity of belief update. *Artificial Intelligence*, 119:141–190, 2000.
- [Liberatore, 2000b] P. Liberatore. A framework for belief update. In *Proc. JELIA-2000*.
- [Liu et al., 2006] H. Liu, C. Lutz, M. Milicic, and F. Wolter. Updating description logic aboxes. *Proc. KR2006*, 46–56.
- [Peppas et al., 1996] P. Peppas, A. Nayak, M. Pagnucco, N. Foo, R. Kwok, and M. Prokopenko. Revision vs. update: taking a closer look. In *Proc. ECAI96*.
- [Sandewall, 1995] E. Sandewall. *Features and Fluents*. Oxford University Press, 1995.
- [Shapiro and Pagnucco, 2004] S. Shapiro and M. Pagnucco. Iterated belief change and exogenous actions in the situation calculus. In *Proc. ECAI04*.
- [Winslett, 1990] M. Winslett. *Updating Logical Databases*. Cambridge University Press, 1990.