

Considerations on Belief Revision in an Action Theory

James Delgrande

School of Computing Science
Simon Fraser University
Burnaby BC, V5A 1S6, Canada
jim@cs.sfu.ca

Among the many and varied areas that Vladimir Lifschitz has worked on is reasoning about action and change, in particular with respect to action languages, where an action language in turn is based on the underlying semantic notion of a transition system. Transition systems have been shown to be an elegant, deceptively simple, yet rich framework from which to address problems of action consequence, causality, planning and the like. In this paper I consider a problem in the interaction between reasoning about action, observations, and the agent's knowledge, specifically when an observation conflicts with the agent's knowledge; and so the agent must revise its knowledge. In particular, it is shown how an agent's initial belief set may be propagated through an action sequence so that, in contrast to previous work, for a revision one does not need to refer back to the initial state of the agent.

1 Introduction

An agent acting in some domain will generally have incomplete and possibly incorrect knowledge regarding that domain. Semantically, an agent's knowledge K may be characterised by a set of interpretations, or possible worlds, consisting of those worlds that, insofar as the agent is concerned, could be the actual world. Assuming that there are no other agents (including a "nature" that might provide stochastic events), there are two ways in which the agent's knowledge may evolve. First, it may execute actions. In this case, the agent's new beliefs can be semantically characterised by the image of each world in K under the executed action. Hence the agent will believe that, if the action preconditions hold, then after execution of the action, the action's effects will hold. One means of specifying the effects of actions is a *transition system*, where a transition system is a directed graph in which vertices are labelled by possible states of the world, and directed edges are labelled by the actions that take one state of the world to another. So after executing action a , the agent's beliefs will be characterized by the a -accessible worlds from worlds in K .

Second, the agent may *sense* or *observe* the environment. Such knowledge-producing actions don't alter the (external) domain, but they do give the agent information concerning the domain. Assuming that the agent's beliefs are in fact *knowledge* (and so the real world is among those in K), the sensing action

that reports that ϕ is true can be characterised semantically by intersecting the K worlds with the ϕ -worlds. This works well enough if the agent’s beliefs are correct, but in the case where the agent believes that ϕ is false yet senses that ϕ is true, it will fall into inconsistency.

The obvious solution is to *revise* the agent’s beliefs by the result of sensing, rather than simply *expanding* the agent’s beliefs by the new information. Such a solution of course presupposes the existence of a revision function whereby, from any knowledge base and formula for revision, a new knowledge base can be determined. However, this still does not resolve all difficulties.

In particular, it seems that if an agent executes an interleaved series of actions and observations, then its beliefs should be determined by iteratively executing the resulting series of actions and belief revisions. That is, one would expect that a *Markovian* process, in which the agent’s next belief state depends solely on its previous state together with the relevant action, would be sufficient. However, as we describe in the next section, this naïve approach may lead to incorrect results. These problems have been resolved in an approach called *belief evolution* but, in doing so, the desirable Markovian aspect is lost, in that one must refer back to the agent’s initial state of beliefs in order to accommodate the result of an observation.

The goal of this paper is to address this issue. That is, naïvely executing a sequence of actions and revisions leads to problems; and the solution, belief evolution, gives the correct result, but necessitates that each observation must be projected back to the initial state. Here we give a procedure that is equivalent to belief evolution but in which one doesn’t have to project each observation back to the initial state. The solution is quite intuitive: in addition to keeping track of the effects of actions via a transition system, one also keeps track of (i.e. *progresses*) information obtained from observations.

In the next section we present the formal framework and discuss related work. Following this we informally but more concretely sketch the problem and solution. The next section gives the formal details. We conclude with a discussion of related work and suggestions for future work.

2 Background

2.1 Transition Systems

The basic definitions of a transition system are taken from [10]. An *action signature* is a pair $\langle \mathbf{F}, \mathbf{A} \rangle$ where \mathbf{F}, \mathbf{A} are disjoint non-empty sets of symbols. \mathbf{F} is the set of *fluent symbols* and \mathbf{A} is the set of *action symbols*. For simplicity we assume that \mathbf{F} and \mathbf{A} are finite.

The fluent symbols in \mathbf{F} are propositional variables. The action symbols in \mathbf{A} denote the actions that an agent may perform. The effects of actions are specified by a transition system.

Definition 1. *A transition system T for an action signature $\sigma = \langle \mathbf{F}, \mathbf{A} \rangle$ is a pair $\langle S, R \rangle$ where*

1. S is a set of propositional interpretations of \mathbf{F} ,
2. $R \subseteq S \times \mathbf{A} \times S$.

S is called the set of *states* and can be regarded as specifying the set of *possible worlds*. R is the *transition relation*. If $(s, a, s') \in R$, then s' is a possible state that could occur as a result of the action a being executed in state s . For formula ϕ , $\|\phi\|$ is the set of states at which ϕ is true, that is $\|\phi\| = \{s \in S \mid s \models \phi\}$.

Transition systems can be represented as directed graphs, where each node is labeled with a state and each edge is labeled with an element of \mathbf{A} . We also define for $X \subseteq S$

$$\begin{aligned} a(X) &= \{s' \in S \mid \exists s \in X \text{ and } (s, a, s') \in R\} \\ a^{-1}(X) &= \{s \in S \mid \exists s' \in X \text{ and } (s, a, s') \in R\}. \end{aligned}$$

For $s \in S$, $a(s)$ is used to abbreviate $a(\{s\})$ and similarly for $a^{-1}(s)$. This extends in the obvious way to formulas; that is, we use $a(\phi)$ to mean $a(\|\phi\|)$. Finally, for a sequence of actions $\vec{a} = \langle a_1, \dots, a_n \rangle$ we extend the above so that we have

$$\begin{aligned} \vec{a}(S) &= a_n(\dots a_1(S) \dots) \\ \vec{a}^{-1}(S) &= a_1^{-1}(\dots a_n^{-1}(S) \dots) \end{aligned}$$

If there is exactly one possible resulting state s' when a is executed in s for every $s \in S$ and $a \in \mathbf{A}$, then T is *deterministic*. A minor but key point concerning deterministic transition systems is that, while $a(s)$ is a singleton, $a^{-1}(s)$ may not be. We assume throughout this paper that a transition system is deterministic, and that every action is executable in every state. If the transition system does not specify the results of a particular action in a given state, we assume that the state does not change when that action is executed. This is equivalent to adding self loops for every action at every state where no transition is given.

2.2 Belief Revision

The central and best-known approach to belief change is the AGM approach [1, 9, 19], named after the developers of the approach. In the AGM approach, beliefs of an agent are modelled by a deductively closed set of formulas, or *belief set*. Thus a belief set is a set of formulas K such that $K = \mathcal{Cn}(K)$ where $\mathcal{Cn}(K)$ is the closure of K under classical logical consequence. It is assumed that the underlying logic subsumes classical propositional logic. Formally, a revision operator $*$ maps a belief set K and formula ϕ to a revised belief set $K * \phi$. The AGM postulates for revision are as follows; the operator $+$ stands for *expansion*, where $K + \phi$ is defined to be $\mathcal{Cn}(K \cup \{\phi\})$.

- (K*1) $K * \phi = \mathcal{Cn}(K * \phi)$
- (K*2) $\phi \in K * \phi$
- (K*3) $K * \phi \subseteq K + \phi$

- (K*4) If $\neg\phi \notin K$ then $K + \phi \subseteq K * \phi$
- (K*5) $K * \phi$ is inconsistent only if ϕ is inconsistent
- (K*6) If $\phi \equiv \psi$ then $K * \phi = K * \psi$
- (K*7) $K * (\phi \wedge \psi) \subseteq K * \phi + \psi$
- (K*8) If $\neg\psi \notin K * \phi$ then $K * \phi + \psi \subseteq K * (\phi \wedge \psi)$

See [9, 19] for a discussion and motivation of these postulates.

Katsuno and Mendelzon [14] have shown that a necessary and sufficient condition for constructing an AGM revision operator is that there is a function that associates a total preorder on the set of possible worlds with any belief set K , as follows:¹

Definition 2. A faithful assignment is a function that maps each belief set K to a total preorder \preceq_K on possible worlds such that for any possible worlds w_1, w_2 :

1. If $w_1, w_2 \in \text{Mod}(K)$ then $w_1 \approx_K w_2$
2. If $w_1 \in \text{Mod}(K)$ and $w_2 \notin \text{Mod}(K)$, then $w_1 \prec_K w_2$.

The resulting preorder is called the *faithful ranking* associated with K . Intuitively, $w_1 \preceq_K w_2$ if w_1 is at least as plausible as w_2 according to the agent. Katsuno and Mendelzon then provide the following representation result, where $t(W)$ is the set of formulas of classical logic true in W :

Theorem 1. A revision operator $*$ satisfies postulates (K*1)–(K*8) iff there is a faithful assignment that maps each belief set K to a total preorder \preceq_K such that

$$K * \phi = t(\min(\text{Mod}(\phi), \preceq_K)).$$

Thus the revision of K by ϕ is characterised by those models of ϕ that are most plausible according to the agent.

Various researchers have argued that, in order to address iterated belief revision, it is more appropriate to consider *belief states* (also called *epistemic states*) as objects of revision. A general framework was proposed in [25], while the first systematic exploration of this notion in Artificial Intelligence is reported in [6]. A belief state \mathcal{K} effectively encodes information regarding how the revision function itself changes under a revision. The belief set corresponding to belief state \mathcal{K} is denoted $\text{Bel}(\mathcal{K})$. A revision operator $*$ now maps a belief state \mathcal{K} and formula ϕ to a revised belief state $\mathcal{K} * \phi$. The AGM revision postulates can be reformulated by replacing instances of a revision, say $K * \phi$, by the belief set $\text{Bel}(\mathcal{K} * \phi)$, and replacing (K*6) by:

- (K*6) If $\mathcal{K}_1 = \mathcal{K}_2$ and $\phi \equiv \psi$ then $\text{Bel}(\mathcal{K} * \phi) = \text{Bel}(\mathcal{K} * \psi)$.

Darwiche and Pearl [6] extend the Katsuno and Mendelzon results as follows:

¹ In [14] (and [6], below), an agent's beliefs are represented by a formula rather than a belief set; hence they adopt a different, but equivalent, set of postulates. Since we deal with finite languages, the difference is immaterial.

Definition 3. A faithful assignment is a function that maps each belief state \mathcal{K} to a total preorder $\preceq_{\mathcal{K}}$ on possible worlds such that for any possible worlds w_1, w_2 :

1. If $w_1, w_2 \models \text{Bel}(\mathcal{K})$ then $w_1 =_{\mathcal{K}} w_2$
2. If $w_1 \models \text{Bel}(\mathcal{K})$ and $w_2 \not\models \text{Bel}(\mathcal{K})$, then $w_1 \prec_{\mathcal{K}} w_2$

Theorem 2. A revision operator $*$ satisfies postulates (K*1)–(K*8) iff there exists a faithful assignment that maps each belief state \mathcal{K} to a total preorder $\preceq_{\mathcal{K}}$ such that

$$\text{Mod}(\mathcal{K} * \phi) = \min(\text{Mod}(\phi), \preceq_{\mathcal{K}}).$$

In the above theorem, $\text{Mod}(\mathcal{K})$ is understood as $\text{Mod}(\text{Bel}(\mathcal{K}))$.

2.3 Reasoning about Action and Knowledge

Gelfond and Lifschitz [10] provide an overview of action languages and transition systems. Detailed, specific approaches can be found in, for example, [12, 11]. The frameworks developed in these (and other related) papers specify syntactic constructs comprising an *action language*, where the semantics of these constructs can be given in terms of a transition system. In such languages one can then specify the preconditions and effects of actions, along with other related notions, such as persistence of fluent values.

In [2] the base action language \mathcal{A} is extended to allow observations of the truth values of fluents and the actual occurrence of actions, while the entailment relation allows for querying regarding types of hypothetical reasoning. This approach is extended in turn in [3], where it is employed in diagnostic reasoning in an action framework. In [17, 24], the action language \mathcal{A} is extended to handle an agent’s knowledge and sensing actions. The representation of this in terms of a transition system is clear and intuitive: The agent’s beliefs are modelled by a set of states, or possible worlds, K ; and in sensing that a formula ϕ is true, all $\neg\phi$ worlds are removed from this set. Since it is assumed that an agent’s beliefs are accurate, in that the real world is among those in K , sensing leads to a monotonic refinement of the agent’s knowledge. (If nondeterministic actions are allowed, as in [17], the agent’s ignorance may increase but it will never be incorrect.)

There has been similar research in other formalisms for reasoning about action, most notably in the *situation calculus* [16]. In brief, the situation calculus is a first-order² theory in which fluents take a situation argument, and a function $do(a, s)$ yields the situation resulting from executing action a in situation s . Hence $Red(A, s)$ indicates that object A is *Red* in situation s , while $\neg Red(A, do(\text{paintBlue}, A, s))$ indicates that A is not *Red* in the situation resulting from s in which the *paintBlue* action is executed.

Scherl and Levesque [21] extend the basic theory to account for an agent’s knowledge, basically by axiomatising the modal logic S5 within the situation

² Well, not quite, since there is an induction axiom on situations.

calculus, as well as sensing. [23] extends this approach to allow for a restricted version of belief revision, in that a formula for revision must be true in the domain. Thus, in [23] if one revises by ϕ and then $\neg\phi$, inconsistency results.

3 Revision in Transition Systems

3.1 The Problem

It would seem to be a straightforward matter to address belief revision in a transition system framework, by simply combining an account of reasoning about knowledge with an account of belief revision. This in turn could be accomplished by simply specifying a faithful assignment, as in Definition 2, for every set of states. And indeed it is straightforward to incorporate belief revision in such a manner. However, as [13] points out, iteratively determining the results of actions and observations may yield incorrect results.

Example 1. Assume that there is a beaker containing either an acid (A) or a base ($\neg A$).³ As well, there is a piece of paper that the agent believes is litmus paper (L). If litmus paper is dipped in acid it turns red (R), and if it is dipped in a base it turns blue (B). The paper is initially white ($\neg R \wedge \neg B$). The only action is to dip the litmus paper into the liquid. The agent initially believes that it is holding litmus paper and that the paper is white.

Consider where the agent dips the paper into the beaker, after which it observes that the paper is white. There are AGM revision operators where in the resulting belief state the agent believes that it is holding a piece of white litmus paper. For example consider the Dalal [5] revision operator: The agent’s beliefs are characterized by the possible worlds $\overline{ARBL}, \overline{ARBL}$; the image of these worlds following a *dip* action is $ARBL, \overline{ARBL}$. The agent observes that $\neg R \wedge \neg B$ is true. The Dalal operator selects the closest worlds of the formula for revision to the worlds characterising the agent’s beliefs, where “closeness” is defined by the Hamming distance. In this case the closest worlds are $\overline{ARBL}, \overline{ARBL}$, so the agent believes that it is holding white litmus paper. The same result obtains with the Satoh [20] operator.⁴ Clearly this is an undesirable state, and instead the agent should believe that the paper is not litmus paper.

To deal with such problems, [13] proposes a new operation, *belief evolution*, defined in terms of existing action-update and revision operators. Example 1 illustrates that, if an agent has belief set K , then characterising the result of an action a followed by a revision by ϕ as $K * a(\phi)$ may lead to problems. In belief evolution, the action a followed by a revision by ϕ is expressed as $a(K * a^{-1}(\phi))$. That is, one doesn’t revise by ϕ following the action a ; rather, given the pair (a, ϕ) , one considers those (initial) worlds which, after a is executed, yield ϕ worlds. This set of worlds (that is, $a^{-1}(\phi)$) is used for revising the initial belief

³ To simplify matters, we disallow the possibility of a neutral pH. We obtain the same results if the example is extended to allow that the liquid may be neutral.

⁴ This isn’t quite an AGM operator as the last postulate may not hold.

set K , and the result is progressed via a . It is shown that this formulation gives the appropriate results. So in the preceding example, the agent will indeed believe that the paper is not litmus paper. This operator naturally extends to a sequence of actions $\vec{a} = \langle a_1, \dots, a_n \rangle$, so that a revision by ϕ following this sequence of actions would be computed via the extension $\vec{a}(K * \vec{a}^{-1}(\phi))$.

This is fine, but it implies that the *Markovian* aspect of the naïve procedure (viz. $\vec{a}(K) * \phi$) is lost. Rather, a formula for revision ϕ is projected back through the n actions, a revision is carried out, and the result progressed through the n actions. Informally, it seems at best inconvenient to have to refer to the initial belief state in order to carry out a revision at the end of a sequence of actions; as well, it seems to be unnecessarily inefficient. (It can also be noted that switching to a situation calculus setting, for example, doesn't help, in that *basic action theories* of the situation calculus are set up so that truth of fluents at a situation can be determined by the initial state along with the actions executed.)

3.2 Revision in the Context of Sensing

However, let's (re)examine what's going on in Example 1. The definition of a faithful ranking, Definition 2, stipulates that a total preorder is associated with every set of possible worlds. Arguably, Example 1 illustrates that this is inappropriate in an action context, in that the result of an action will give information that will affect the results of revision. Thus, in Example 1 the agent should not believe that the paper is litmus paper after the dipping action, because the result of the dipping rules this out: litmus paper cannot be white when dipped into a solution known to be non-neutral. Consequently revision, in one form or another, needs to take this into account.

From the consideration that Definition 2 is inappropriate in an action setting, it becomes clear why belief evolution works: one refers to only a single revision operator, that defined by \preceq_K . For an observation that follows a sequence of actions, one maps everything back to the original belief set, carries out revision with respect to the initial belief set (and so with respect to the ranking \preceq_K), and then progresses the results through the given action sequence.

The question arises as to whether one might be able to do better than with the approach to belief evolution and fortunately, the answer is “yes”. The obvious way to proceed is adopt the more general notion of belief state. Then, given an action a in the initial state, one progresses not just the agent's contingent knowledge K , but also the belief state \mathcal{K} . It turns out that things are a bit more complicated than what this “obvious” solution suggests. Nonetheless, in the next section we show how one can carry out actions and revisions in a Markovian manner by keeping track of information gained from sensing.

4 The Approach

In what follows, we assume that we are given a deterministic transition system $T = \langle S, R \rangle$ over a finite action signature $\sigma = \langle \mathbf{F}, \mathbf{A} \rangle$. As well, we assume that

we are given a belief state \mathcal{K} where $Bel(\mathcal{K}) = K \subseteq S$ is a fixed set of states representing the agent's initial set of beliefs and $\preceq_{\mathcal{K}}$ is a faithful ranking associated with \mathcal{K} giving the agent's initial revision function. We will henceforth understand an agent's (contingent) beliefs to be characterised by a set of worlds; thus the agent's beliefs following a revision will similarly be a set of worlds.

An agent's beliefs may change in one of two ways. It may execute an action, in which case each state in K is mapped to another state. Given the fact that we have a deterministic transition system, there is no uncertainty as to the outcome of an action execution (although, of course, in some states an action may be nonexecutable). Second, the agent may sense that a formula ϕ is true. We assume that the agent believes that sensing is correct, and so the agent will consequently accept that ϕ is true. In this case, the agent would revise its beliefs according to an appropriate faithful ranking and, given this faithful ranking, there is no uncertainty as to the outcome. The key point then is that following any action sequence, we wish to be able to revise the agent's beliefs; and to do this we require, directly or implicitly, a ranking function over possible worlds. We obtain such a ranking function by suitably progressing not just the agent's contingent beliefs, but also the agent's ranking function $\preceq_{\mathcal{K}}$.

The following definition specifies how a plausibility ranking will change following the execution of an action. A complication is that a possible world $w \in a(S)$ may be the image of more than one world in S , in that $a^{-1}(w)$ may not be unique. Consequently, the rank of $w \in a(S)$ is taken as the rank of the minimally-ranked world in $a^{-1}(w)$.

Definition 4.

Let $T = \langle S, R \rangle$ be a transition system over $\sigma = \langle \mathbf{F}, \mathbf{A} \rangle$, and let $\preceq_{\mathcal{K}}$ be a total preorder over S . For $a \in \mathbf{A}$, define $w_1 \preceq_{a(\mathcal{K})} w_2$ iff

$$\exists w'_1, w'_2 \in S \text{ where } w'_1 \preceq_{\mathcal{K}} w'_2 \text{ and } w'_1 \in \min(a^{-1}(w_1), \preceq_{\mathcal{K}}), w'_2 \in \min(a^{-1}(w_2), \preceq_{\mathcal{K}})$$

Then we can define $Mod(a(\mathcal{K}) * \phi)$ as $\min(Mod(\phi), \preceq_{a(\mathcal{K})})$.

It follows straightforwardly that $\preceq_{a(\mathcal{K})}$ is a total preorder over S with minimum worlds given by $Mod(a(\mathcal{K}))$; ⁵ we omit the proof.

What would be ideal is if Definition 4 preserved \preceq -relations under actions, and so one obtained that

$$Mod(a(\mathcal{K} * \phi)) = Mod(a(\mathcal{K}) * a(\phi)). \tag{1}$$

However, it is easily shown that this is not the case, as illustrated by Figure 1. On the left hand side we have the total preorder induced by \mathcal{K} , and on the right hand side we have the image of this preorder following action a . In both cases the lower a world in a ranking, the more plausible that world is; hence the agent's (contingent) beliefs before and after execution of action a are given by the consequences of $p \wedge q$.

⁵ Recall that for belief state \mathcal{K} we define the set of models of the agent's beliefs as $Mod(\mathcal{K}) = Mod(Bel(\mathcal{K}))$.

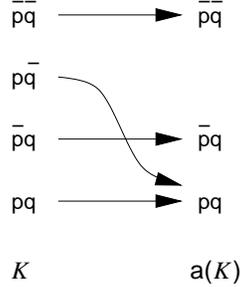


Figure 1

Let $\phi = (p \equiv \neg q)$. Then we have that

$$\begin{aligned} \text{Mod}(a(\mathcal{K} * \phi)) &= a(\{\neg pq\}) = \{\neg pq\} \\ \text{Mod}(a(\mathcal{K}) * a(\phi)) &= a(\mathcal{K}) * \{pq, \neg pq\} = \{pq\} \end{aligned}$$

This example shows that neither containment in (1) holds.

However, consider again Example 1, but in the context of Definition 4. Since the liquid is either an acid or a base, then *if* the paper is litmus paper, then it is *impossible* for it to be white after the dipping action. Hence, white-litmus-paper worlds should not take part in any subsequent revision, since they have been ruled out by the action. This helps explain why belief evolution provides the correct results: for an observation (here *white*, expressed as $\neg R \wedge \neg B$) that is inconsistent with the agent's beliefs following an action, it must be the case that the agent's original beliefs are inconsistent with the inverse image of the observation under the action. That is, we have the elementary result:

Proposition 1. *Let $T = \langle S, R \rangle$ be a transition system over $\sigma = \langle \mathbf{F}, \mathbf{A} \rangle$, $a \in \mathbf{A}$, and let \mathcal{K} be a belief state.*

If $\text{Mod}(a(\mathcal{K})) \cap \|\phi\| = \emptyset$ then $\text{Mod}(\mathcal{K}) \cap a^{-1}(\|\phi\|) = \emptyset$.

With belief evolution, an observation ϕ is projected back (via a^{-1}) to the original (pre-action) state of affairs, where revision takes place, thereby adjusting the agent's original, erroneous, set of beliefs. The result is then progressed under the action a , (trivially) yielding a result compatible with a . As described, implementing a reasonable notion of revision then involves regressing an action sequence to the initial state, carrying out revision, and then progressing the same sequence of actions. The fact that one can avoid this effort hinges on the following result.⁶

Proposition 2. $\|\phi\| \cap a(S) = a(a^{-1}(\phi))$

⁶ The following results (clearly) assume the same antecedent conditions as Proposition 1. We omit their statement, for reasons of perspicuity.

On the left hand side we have those ϕ -worlds that are compatible with the execution of action a . The right hand side says that this is just the a -regression of ϕ -worlds followed by their a -progression. The next result says that the regression/progression process in belief evolution can be restricted to the formula for revision:

Proposition 3. $Mod(a(\mathcal{K} * a^{-1}(\phi))) = Mod(a(\mathcal{K}) * a(a^{-1}(\phi)))$

Corollary 1. $Mod(\vec{a}(\mathcal{K} * \vec{a}^{-1}(\phi))) = Mod(\vec{a}(\mathcal{K}) * \vec{a}(\vec{a}^{-1}(\phi)))$

Proposition 3 is just a weaker version of (1). On the left hand side we have the expression for belief evolution. On the right hand side we have the progression of the belief state, $a(\mathcal{K})$, revised by the regression/progression of the formula for revision. The corollary notes that this, unsurprisingly, extends to a sequence of actions.

Combining the previous two propositions yields the following:

Proposition 4. $Mod(a(\mathcal{K} * a^{-1}(\phi))) = Mod(a(\mathcal{K}) * (\|\phi\| \cap a(S)))$

Corollary 2. $Mod(\vec{a}(\mathcal{K} * \vec{a}^{-1}(\phi))) = Mod(\vec{a}(\mathcal{K}) * (\|\phi\| \cap \vec{a}(S)))$

This shows that belief evolution (on the left hand side of the equality) can be expressed as a simple revision (tracked via Definition 4) following a sequence of actions. The (modest) complication is that one needs to also keep track of the context implicit in the sensing action ($a(S)$), and limit the ϕ worlds involved in the revision to these worlds. That is, the proposition shows belief evolution can be computed by progressing the initial revision function, along with the set of possible states (i.e. $a(S)$).

This then restores the Markovian aspect of iterated actions and revisions, in that the right hand side of Proposition 4 makes reference to only the state of affairs following execution of the action a . Conceptually this seems to be more compelling than the original statement of belief evolution (given on the left hand side of Proposition 4). As well, it may offer pragmatic advantages. Consider the following informal argument: Assume that computing action effects, including inverse actions, and revisions have equal, unit cost, and other operations (such as intersection) have negligible cost. Given an initial belief state \mathcal{K} , the cost of computing the action/observation sequence $\langle a_1, \phi_1, \dots, a_n, \phi_n \rangle$ will have cost $n(n+2)$ via belief evolution, but $3n$ via the right hand side of Proposition 4. Hence we go from a quadratic to a linear number of operations. (This argument is admittedly a bit misleading, since belief revision is generally of greater complexity than propositional reasoning [8]; and both approaches to computing the agent's beliefs following an action/observation sequence require n revisions. However, it indicates that in an implementation there may nonetheless be advantages to the latter approach, particularly in an approach to revision with good expected computational properties, such as in Horn formula belief sets [7].)

The above results apply to an agent's belief set. The question arises as to whether one may be able to say something about the agent's belief state. The next result extends Proposition 4 from belief sets to belief states, with respect to the specific approach to revision of [18].

Proposition 5. *Let $T = \langle S, R \rangle$ be a transition system over $\sigma = \langle \mathbf{F}, \mathbf{A} \rangle$. Assume that faithful rankings satisfy the additional conditions:*

- (CR1) *If $w_1, w_2 \models \phi$, then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K} * \phi} w_2$*
- (CR2) *If $w_1, w_2 \not\models \phi$, then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K} * \phi} w_2$*
- (Lex) *If $w_1 \models \phi$ and $w_2 \not\models \phi$ then $w_1 \prec_{\mathcal{K}} w_2$*

$$\text{Then } a(\mathcal{K} * a^{-1}(\phi)) = a(\mathcal{K}) * (\|\phi\| \cap a(S)).$$

The semantic conditions (CR1), (CR2), and (Lex) characterise the specific approach discussed in [18]. (CR1), (CR2) are the first two (of four) conditions for iterated revision from [6], although they were originally employed in the quantitative setting for *ordinal conditional functions*, in [25]. (Lex) in turn implies the conditions (CR3), (CR4) in [6].

This approach [18] gives very high priority to the formula for revision, in that in revising by ϕ , all ϕ worlds are ranked as being more plausible than any $\neg\phi$ world. In return, one obtains an operator that is conceptually simple and, for the present case, compatible with an action framework.

5 Discussion

In this paper we have reformulated an approach to belief revision in an action formalism, *belief evolution*, so that one doesn't need to refer to the initial state of the agent and the domain. To this end, a number of limiting assumptions were made: it was assumed that action preconditions and effects are known, and that actions are deterministic. As well, observations are correct. A fully general account of belief change will need to relax these assumptions of course; however the idea in this paper is to attempt to gain insight into the interaction of actions and observations by considering a more manageable scenario.

5.1 Related Work

The assumptions made here are the same as those made in the account of belief change in the situation calculus in [23]. Interestingly, the problem observed in Example 1 doesn't arise in [23], and we obtain that after observing that the paper is white, the agent believes that the paper is not litmus paper. In the situation calculus, situations are indexed by actions; and in those situations where the agent is holding litmus paper, the paper will turn either red or blue, and those situations where it is not litmus paper, the colour will remain unchanged. That is, since situations record the actions taken, actions are implicitly part of the subject of the revision process. As a result, in observing that $\neg R \wedge \neg B$ is true, it must be that the paper is not litmus paper.

In contrast, in the transition system approach, the link between an action and the resulting state is not retained, in that for a world $w \in a(\mathcal{K})$, $a^{-1}(w)$ is not guaranteed to be a singleton, as it is in the situation calculus. The result is that in the naïve approach, the agent may end up believing that it is holding

white litmus paper, since the link between litmus paper and colour after dipping is lost. In essence, the situation calculus approach is finer grained and more nuanced than extant approaches involving transition systems.⁷ So this raises the question of whether one should opt for a theory such as the situation calculus, over a simpler approach such as with a transition system.

While there are advantages to the situation calculus, particularly in terms of expressibility, there are also advantages to the simpler approach, namely one might expect it to be easier to implement and to be more computationally efficient, simply because less information is kept track of. Moreover, with the situation calculus, at least with basic action theories, one essentially still needs to regress to the initial state. While there will be cases where one may need to reason among earlier circumstances (e.g.[4]), if one need not, then there will presumably be advantages computationally to adopting a simpler, Markovian, approach.

Regarding other work, the setting described in [3] is broadly similar to that here. In [3], given an action language framework, diagnostic problem solving, based on a set of observations, is developed. That is, given a definition of a *system*, and a set of observations inconsistent with this system, the notion of a *diagnosis* is defined to account for the abnormal behaviour. The setting of [3] is much more specific than that of the present paper, but the overall direction and goals are similar. Consequently, it would be an interesting question to see how one approach may benefit the other. Thus the present approach may help explicate the notion of observation in [3]; as well, by incorporating belief states, one might obtain a more fine-grained diagnostic approach that could accommodate observations and beliefs of varying reliability. On the other hand, by bringing in specific notions (such as that of components, or distinguished sets of literals, or specific action types) from a diagnostic setting, one might be able to provide an account for revision in a specific, practical setting.

5.2 Extending the approach

There are several ways in which the approach may be generalised. To begin, we've essentially been working with a single observation following an action sequence. The natural extension to a set of observations interleaved with actions implicitly assumes that the observations are, if not correct, at least consistent, taken as a group. Thus, still assuming deterministic actions, the most straightforward means of dealing with a set of observations is to first project all observations to the final state. Under the assumption that this set is consistent, the intersection of the resulting sets of worlds corresponds to the conjunction of these (projected) observations, which can then be used for a revision.

Clearly however, observations may not be correct or compatible, and in belief revision in general for example, the agent may be informed of something that is

⁷ Which is to say, there is no obstacle to formulating a situation-calculus like approach involving transition systems, where one keeps track of trajectories of actions and states of the world.

later contradicted. Hence in general one would need to be able to deal with an alternating sequence of actions and observations, where observations may not be correct. Given such a set of (possible incorrect or incompatible) observations, one can still project (or progress) each observation individually through the full sequence of actions, yielding a set of possibly incompatible observations, together with the result of progressing the agent's beliefs. One can then in some fashion *merge* the resulting observations, to give an overall summation of the observations that can then be used to revise the agent's beliefs. This leaves open the question of how to best merge the agent's observations, but in this case there is an extensive literature (for representative work, see for example [15]) and one can treat the problem independently from reasoning about action.⁸

A second extension is to allow more general types of actions. It may be that nondeterminism can be directly incorporated into the approach, provided that a (nondeterministic) action results in equally-probably outcomes. Thus, in the case where the agent tosses a fair coin for example, there are two possible outcomes; while the final outcome is unknown, both outcomes will have the same plausibility. Hence for revision, it may be that in this situation there is no real complication beyond keeping track of the multiple outcomes. If a nondeterministic action allows outcomes of differing likelihood, for example in allowing some actions to occasionally fail or to yield unexpected outcomes, then it is not clear how such outcomes should be reconciled with the agent's plausibility ordering. A second way in which actions can be generalised is in allowing *exogenous actions*. In this case, an approach such as [22] may work here.

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50(2), 510–530 (1985)
2. Baral, C., Gelfond, M., Proveti, A.: Representing actions: Laws, observations and hypotheses. *Journal of Logic Programming* 31(1-3), 201–243 (1997)
3. Baral, C., McIlraith, S.A., Son, T.C.: Formulating diagnostic problem solving using an action language with narratives and sensing. In: *Proceedings of the International Conference on the Principles of Knowledge Representation and Reasoning*. pp. 311–322 (2000)
4. Boutilier, C.: Generalized update: Belief change in dynamic settings. In: *Proceedings of the International Joint Conference on Artificial Intelligence*. pp. 1550–1556 (1995)
5. Dalal, M.: Investigations into theory of knowledge base revision. In: *Proceedings of the AAAI National Conference on Artificial Intelligence*. pp. 449–479. St. Paul, Minnesota (1988)
6. Darwiche, A., Pearl, J.: On the logic of iterated belief revision. *Artificial Intelligence* 89, 1–29 (1997)

⁸ In belief revision, a common assumption, even in static domains, is that the most recent observation should take precedence. The suggestion here is that, unless there are specific reliabilities attached to observations, there is no good reason to prefer one observation over another.

7. Delgrande, J., Peppas, P.: Revising Horn theories. In: Proceedings of the International Joint Conference on Artificial Intelligence. pp. 839–844. Barcelona, Spain (2011)
8. Eiter, T., Gottlob, G.: On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57(2-3), 227–270 (1992)
9. Gärdenfors, P.: *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA (1988)
10. Gelfond, M., Lifschitz, V.: Action languages. *Electronic Transactions on AI* 3 (1998)
11. Giunchiglia, E., Lee, J., Lifschitz, V., McCain, N., Turner, H.: Nonmonotonic causal theories. *Artificial Intelligence* 153(1-2), 49–104 (2004)
12. Giunchiglia, E., Lifschitz, V.: An action language based on causal explanation: Preliminary report. In: Proceedings of the AAAI National Conference on Artificial Intelligence. pp. 623–630 (1998)
13. Hunter, A., Delgrande, J.P.: Iterated belief change due to actions and observations. *Journal of Artificial Intelligence Research* 40, 269–304 (2011)
14. Katsuno, H., Mendelzon, A.: Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3), 263–294 (1991)
15. Konieczny, S., Pino Pérez, R.: Merging information under constraints: A logical framework. *Journal of Logic and Computation* 12(5), 773–808 (2002)
16. Levesque, H., Pirri, F., Reiter, R.: Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science* 3(18) (1998)
17. Lobo, J., Mendez, G., Taylor, S.: Knowledge and the action description language $\mathcal{A}_{\mathcal{K}}$. *Theory and Practice of Logic Programming* 1(2), 129–184 (2001)
18. Nayak, A.C., Pagnucco, M., Peppas, P.: Dynamic belief revision operators. *Artificial Intelligence* 146(2), 193–228 (2003)
19. Peppas, P.: Belief revision. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) *Handbook of Knowledge Representation*, pp. 317–359. Elsevier Science, San Diego, USA (2008)
20. Satoh, K.: Nonmonotonic reasoning by minimal belief revision. In: Proceedings of the International Conference on Fifth Generation Computer Systems. pp. 455–462. Tokyo (1988)
21. Scherl, R., Levesque, H.: Knowledge, action, and the frame problem. *Artificial Intelligence* 144(1-2), 1–39 (2003)
22. Shapiro, S., Pagnucco, M.: Iterated belief change and exogeneous actions in the situation calculus. In: Proc. ECAI04 (2004)
23. Shapiro, S., Pagnucco, M., Lespérance, Y., Levesque, H.J.: Iterated belief change in the situation calculus. *Artificial Intelligence* 175(1), 165–192 (2011)
24. Son, T., Baral, C.: Formalizing sensing actions: A transition function based approach. *Artificial Intelligence* 125(1-2), 19–91 (2001)
25. Spohn, W.: Ordinal conditional functions: A dynamic theory of epistemic states. In: Harper, W., Skyrms, B. (eds.) *Causation in Decision, Belief Change, and Statistics*, vol. II, pp. 105–134. Kluwer Academic Publishers (1988)