Belief Revision with Sensing and Fallible Actions

James P. Delgrande
School of Computing Science
Simon Fraser University
Burnaby, B.C.
V5A 1S6 Canada
jim@cs.sfu.ca

Hector J. Levesque
Dept. of Computer Science
University of Toronto
Toronto, ON
M5G 3G4 Canada
hector@ai.toronto.edu

Abstract

An agent will generally have incomplete and possibly inaccurate knowledge about its environment. In addition, such an agent may receive erroneous information, perhaps in being misinformed about the truth of some formula. In this paper we present a general approach to reasoning about action and belief change in such a setting. An agent may carry out actions, but in some cases may inadvertently execute the wrong one (for example, pushing an unintended button). As well, an agent may sense whether a condition holds, and may revise its beliefs after being told that a formula is true. Our approach is based on an epistemic extension to basic action theories expressed in the situation calculus, augmented by a plausibility relation over situations. This plausibility relation can be thought of as characterising the agent's overall belief state; as such it keeps track of not just the formulas that the agent believes to hold, but also the plausibility of formulas that it does not believe to hold. The agent's belief state is updated by suitably modifying the plausibility relation following the execution of an action. We show that our account generalises previous approaches, and fully handles belief revision, sensing, and erroneous actions.

Introduction

An agent may interact with its environment in various ways. It may carry out physical actions, and thereby effect change in the environment; it may carry out sensing actions, and learn properties of the current state of the environment; or it may be informed of some aspect of the environment. In general, an agent will not have complete knowledge of its environment, and its knowledge may be inaccurate. In this latter case, an agent may begin with incorrect beliefs, or an action may produce unintended results; or an agent may be incorrectly informed of some fact. Consequently, it is crucial that the agent maintain as accurate a corpus of beliefs as possible, and be able to recover from erroneous beliefs.

In this paper, we develop a general model of an agent that is able to reason and maintain its stock of beliefs in such scenarios. This approach is developed within the framework of the situation calculus (Levesque *et al.* 1998; Reiter 2001), specifically the epistemic extension presented in (Scherl and Levesque 2003). As well, it incorporates notions from belief revision; in particular, we make extensive use of the notion of plausibility, taken from ranking functions (or ordinal conditional functions) (Spohn 1988).

This work generalises previous work in that it integrates possibly-fallible actions, belief revision (via informing actions), and sensing. The overall approach is one that has received extensive treatment in the belief revision community: we associate with an agent a *belief state* that consists not just of a set of contingent beliefs, but also a plausibility ordering over other potential beliefs, expressed in terms of an ordering over situations. Consequently, if an agent discovers that its beliefs are incorrect, then the plausibility ordering provides a principled means for modifying its beliefs.

Our approach is based on the situation calculus, which provides a full account of reasoning about action. Actions are described in terms of their preconditions and their effects, exploiting Reiter's solution to the frame problem (Reiter 2001). We augment this by including the case where an agent may intend to execute one action but inadvertently executes another. (For example the agent may accidentally press a wrong button.) Consequently we allow that the agent's beliefs may evolve according to one sequence of actions (the actions it believes that it executed) while the world evolves in a different direction (according to the actions that the agent actually executes). This also has an epistemic component, in that the agent may be aware of such alternatives, and so in executing an action will keep track of such (according to the agent, counterfactual) possibilities.

If this was all there were to the story, then the agent's beliefs would simply diverge more and more from the real situation. However, the agent may carry out sensing actions; such actions are, by definition, with respect to the actual situation, and so via sensing the agent may correct incorrect beliefs. As well, we also allow that an agent may be informed of some fact. The idea here is that if the agent is informed that ϕ, it will amend its beliefs so that it accepts ϕ. This operation is exactly that of belief revision (Gärdenfors 1988; Peppas 2008). A key point is that an agent may be informed of some formula, ϕ, and later of some other formula ψ that conflicts with ϕ; in this case the agent would nonetheless maintain a consistent set of beliefs (except in the limiting case where ψ is inconsistent).

This approach extends previous work in several respects. It provides a complete integration of an account of reasoning about action with belief revision. In so doing, it allows arbitrary action sequences involving (possibly fallible) physical actions, sensing actions, and informing actions. We retain the results of basic action theories (Reiter 2001), and so inherit the formal results attending such theories. The approach itself is very general – for example it will be seen to straightforwardly incorporate the general notion of conditionalisation of (Spohn 1988). In this paper we focus on a specific, arguably useful, account, and then at the conclusion indicate ways in which the approach may be generalised.

The next section reviews related work. We next give an informal description of our approach, followed by the formal development. We then explore properties of the approach, after which we compare it with related work. In the final sections we discuss future research and conclude.

Background

The next subsection gives a brief introduction to the situation calculus, including extensions to handle an agent's beliefs and sensing. The following subsection gives an introduction to belief revision, including ordinal conditional functions and a specific approach that we incorporate here.

The situation calculus

The language $\mathcal{L}$ of the situation calculus (McCarthy and Hayes 1969) is first-order with equality and many-sorted, with sorts for actions, situations, and objects (everything else). A *situation* represents a world history as a sequence of actions. There is a set of initial situations corresponding to the ways the domain might be initially. The actual initial state of the domain is represented by the distinguished situation constant, S_0. The term $do(a, s)$ denotes the unique situation that results from an agent doing action a in situation s. The term $do(\sigma, s)$ where σ is the sequence $\langle a_1, \dots, a_n \rangle$ abbreviates $do(a_n, do(\dots, do(a_1, s) \dots))$. Initial situations are defined as those without a predecessor:

$$Init(s) \doteq \neg \exists a, s'. \, s = do(a, s').$$

In general, the situations can be structured into a set of trees, where the root of each tree is an initial situation and the edges are actions. Predicates and functions whose values may change from situation to situation (and whose last argument is a situation) are called *fluents*.

To axiomatise a dynamic domain in the situation calculus, we use *basic action theories* (Reiter 2001) consisting of (1) initial state axioms, which describe the initial states of the domain including the initial beliefs of the agents; (2) precondition axioms, which specify the conditions under which each action can be executed;[1] (3) successor state axioms, which describe how each fluent changes as the result of actions; (4) sensing axioms for each action, described below; (5) unique names axioms for the actions; and (6) domain-independent foundational axioms (we adopt the ones given in (Levesque *et al.* 1998) which accommodate multiple initial situations). By axiomatising a domain in these terms, we obtain a number of advantages in reasoning about action and change, including a simple solution to the frame problem (Reiter 2001).

[1]For simplicity, we ignore these here and simply assume that all actions are always possible.

Moore (Moore 1985) defined a possible-worlds semantics for a logic of knowledge in the situation calculus by treating situations as possible worlds. This was adapted to basic action theories in (Scherl and Levesque 1993; Scherl and Levesque 2003) using two special predicates, *SF* and *B*. In this account, each action is assumed to return a (binary) sensing result, and $SF(a, s)$ holds when action a returns sensing value 1 in situation s. The sensing axioms in a basic action theory are used to specify the conditions under which $SF(a, s)$ holds. (Actions that return no useful sensing information are simply axiomatised as always returning 1.)

The B predicate is the usual belief accessibility relation and $B(s', s)$ holds when the agent in situation s thinks that situation s' might be the actual situation. This is considered to be a fluent.[2] Scherl and Levesque present a successor state axiom for B that characterises how actions, including actions with sensing information, affect the beliefs of the agent. Here is a variant of their successor state axiom:[3]

$$\begin{aligned} &B(s', do(a, s)) \equiv \\ &\quad \exists s^* [B(s^*, s) \wedge s' = do(a, s^*) \wedge (SF(a, s^*) \equiv SF(a, s))]. \end{aligned}$$

So the situations s' that are B-related to $do(a, s)$ are the ones that result from doing action a in a previously related situation s^*, such that the sensor associated with action a has the same value in s^* as it does in s.

Belief itself is then defined as an abbreviation for truth in all accessible situations:[4]

$$Bel(\phi[now], s) \ \doteq \ \forall s'. B(s', s) \supset \phi[s'].$$

Scherl and Levesque show that various modal logics of belief (with and without introspection) result from imposing properties on the B fluent in the initial state.

As is typical of logics of belief, sentences of the logic are interpreted as being simply true or false, not necessarily believed by any agent; cases of belief are expressed explicitly using a belief operator. As well, every formula of the logic with a free situation variable can be an argument of this belief operator and therefore an object of belief.

Belief change

The area of *belief change* (Peppas 2008) studies how an agent may modify its belief state in the presence of new information. We focus on *belief revision*, in which the agent is given new information, represented by a formula ϕ, about a (static) domain. This formula is to be incorporated into the agent's set of beliefs. We also later briefly consider *belief update*. In belief update, the agent is again given information to be incorporated into its beliefs, but in this case the

[2]For this reason the order of the situation arguments is reversed from the usual convention in modal logic.

[3]Free variables are assumed to be universally quantified from outside. Also, if ϕ is a formula with a single free situation variable, $\phi[t]$ denotes ϕ with that variable replaced by situation term t.

[4]*now* is used conventionally as a placeholder for a situation argument. Instead of writing $\phi[now]$, we occasionally omit the situation argument completely.

new information concerns a change in the domain due to some actions being carried out.

Belief change functions are usually regarded as being guided, or characterised, by various *rationality postulates*. The *AGM approach* (Gärdenfors 1988) provides the best-known set of postulates for belief revision. An agent's beliefs are modelled by a deductively-closed set of sentences, K, called a *belief set*. Thus a belief set satisfies $K = Cn(K)$, where $Cn(K)$ is the deductive closure of K. $K_\perp$ is the inconsistent belief set (i.e. $K_\perp = \mathcal{L}$). Revision is modeled as a function from belief sets and formulas to belief sets. The operation $K + \phi$ is the *expansion* of K by ϕ, and is defined as $Cn(K \cup \{\phi\})$. The AGM revision postulates are as follows.

(K*1) $K * \phi = Cn(K * \phi)$

(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 $\nvdash \neg\phi$

(K*6) If $\phi \equiv \psi$ then $K * \phi \equiv K * \psi$

(K*7) $K * (\phi \wedge \psi) \subseteq K * \phi + \psi$

(K*8) If $\neg\psi \notin K * \phi$ then $K * \phi + \psi \subseteq K * (\phi \wedge \psi)$

For purposes of later comparison, we also list the update postulates of (Katsuno and Mendelzon 1992), rephrased in terms of belief sets. A formula is said to be *complete* just if it implies the truth or falsity of every other formula. In (K⋄8), $[K]$ is the set of maximum consistent theories containing K.

(K⋄0) $K \diamond \phi = Cn(K \diamond \phi)$

(K⋄1) $\phi \in K \diamond \phi$

(K⋄2) If $\phi \in K$ then $K \diamond \phi = K$

(K⋄3) If ϕ and K are satisfiable then so is $K \diamond \phi$

(K⋄4) If $\phi \equiv \psi$ then $K \diamond \phi = K \diamond \psi$

(K⋄5) $K \diamond (\phi \wedge \psi) \subseteq K \diamond \phi + \psi$

(K⋄6) If $\psi \in K \diamond \phi$ and $\phi \in K \diamond \psi$ then $K \diamond \phi = K \diamond \psi$

(K⋄7) If K is complete then

$$K \diamond (\phi \vee \psi) \subseteq Cn(K \diamond \phi \cup K \diamond \psi)$$

(K⋄8) $K \diamond \phi = \bigcap_{w \in [K]} w \diamond \phi$

(Spohn 1988) explores *ordinal conditional functions* (also called *plausibility orderings*) which provide a basis for specifying a wide range of revision and other belief change functions. In a slight simplification of Spohn's approach, we take an ordinal conditional function (OCF) to be a mapping κ from the set of interpretations to the nonnegative integers such that some interpretation is assigned the number 0. For an interpretation w, $\kappa(w)$ is the *rank* of w. Intuitively, the rank of an interpretation represents its degree of plausibility. The lower an interpretation's rank, the more plausible that interpretation is. A formula ϕ is in the belief set $Bel(\kappa)$ just if every interpretation of rank 0 is a model of ϕ; that is:

$$Mod(Bel(\kappa)) = \{w \mid \kappa(w) = 0\}.$$

Given an OCF κ, we extend this function to a ranking on sentences (or sets of sentences) as follows:

$$\kappa(\phi) = \begin{cases} \infty & \text{if } \vdash \neg\phi \\ \min\{\kappa(w) \mid w \models \phi\} & \text{otherwise} \end{cases}$$

With ordinal conditional functions, the focus of revision shifts from the notion of *belief set*, the set of formulas that the agent believes, to that of a *belief state*, given by a plausibility ordering. Spohn gives a very general notion of belief change called the (ϕ, d)-conditionalisation of κ, in which formula ϕ is accepted to a degree d, in that $\kappa(\neg\phi)$ is assigned a ranking of d.

OCFs (and their qualitative counterpart, total preorders over interpretations) have been used to study *iterated belief revision*, a topic not addressed in the AGM approach. The best-known approach is that of (Darwiche and Pearl 1997). They propose the following postulates:

(DP1) If $\psi \vdash \phi$, then $(K * \phi) * \psi = K * \psi$

(DP2) If $\psi \vdash \neg\phi$, then $(K * \phi) * \psi = K * \psi$

(DP3) If $\phi \in K * \psi$, then $\phi \in (K * \phi) * \psi$

(DP4) If $\neg\phi \notin K * \psi$, then $\neg\phi \notin (K * \phi) * \psi$

They provide a representation result in terms of total preorders over interpretations and give a concrete operator in terms of a ranking function.

Other approaches to iterated revision include (Boutilier 1996; Nayak *et al.* 2003; Jin and Thielscher 2007). We focus on (Nayak *et al.* 2003) which, in the next section, we adapt for our approach. They propose a *conjunction postulate* which can be expressed as[5]

(CNJ) If $\phi \wedge \psi \nvdash \perp$ then $K * \phi * \psi = K * (\phi \wedge \psi)$

They show that this postulate along with the AGM and DP postulates are strong enough to uniquely determine a revision function. If an agent's epistemic state is given by a total preorder $\preceq$ over interpretations, then the preorder following revision by ϕ, $\preceq_\phi$ is given by:

If $w_1 \models \phi$ iff $w_2 \models \phi$ then $w_1 \preceq w_2$ iff $w_1 \preceq_\phi w_2$

If $w_1 \models \phi$ and $w_2 \not\models \phi$ then $w_1 \prec_\phi w_2$

Thus interpretations that agree on the truth of ϕ retain their relative rankings, but interpretations in which ϕ is true are ranked as more plausible than those in which it is false.

The Approach

Our approach elaborates on the Scherl-Levesque scheme in two different ways, by incorporating plausibilities and allowing mistaken actions.

First, an agent may associate a *plausibility* ranking to a situation. Plausibility values are nonnegative integers, where a lower value means that a situation is considered to be more plausible. We can use $B(s', n, s)$ to indicate that in situation s the agent considers s' to have plausibility n. Tuples of the form $B(s', 0, s)$ correspond to the Scherl-Levesque $B(s', s)$ and are used to determine what is believed. Tuples of the form $B(s', n, s)$ where $n > 0$ correspond to counterfactual possibilities that may need to be reconsidered as new information is acquired.

[5](Nayak *et al.* 2003) treat revision as a unary function. The difference is immaterial for our purposes.

Mistaken actions

Our second elaboration is that we allow that an agent may inadvertently execute a physical action other than the one that it intends.

Consider the following example: There are three switches, centre, left, and right. All switches are in the *off* position. If the agent flips the left switch, it will naturally believe that the left switch is *on*. But suppose instead that the agent attempts to flip the centre switch, but inadvertently flips the left one (because the lighting is bad, say). In the resulting situation, the agent has flipped the left switch as before, but it believes that the left switch is *off*.

The conclusion: when there can be mistaken actions like these, the actual physical actions that occur are not enough to determine the situations the agent considers possible; we need to also consider what the agent *believed* it was doing. This will lead us to consider a four-place fluent $B(s', n, \sigma, s)$ where the extra σ argument represents the sequence of actions that the agent believed it was performing at the time.

Sensing and informing

Next, assume that the agent senses whether the left switch is *on*. Sensing is with respect to the actual situation, and so the agent would learn that the left switch was *on*, contrary to its beliefs. The agent would then re-rank its plausibilities. As a result, not only does the agent believe that the left switch is *on* following the sensing action, but it also believes that before that, it flipped the left switch, not the centre one.

We also allow that an agent may be informed of some formula ϕ. If the agent is informed that ϕ, then the agent will believe that ϕ is the case. This corresponds exactly to belief revision, but in a reasoning-about-action framework. Being informed of ϕ has the same effect as sensing a formula: The situations in which ϕ is true (effectively) have their plausibilities uniformly decreased until some ϕ-situations have plausibility 0; at the same time $\neg\phi$ situations have their plausibility increased, so that no $\neg\phi$-situation has plausibility 0.

To be sure, sensing and being informed are very closely related, a fact that we formally establish later. Sensing is more familiar in the reasoning about action community, while being informed is standard in the belief change community; hence we include both. Being informed differs from sensing, in that sensing results are determined by the actual situation. Hence successive sensing actions for the same fluent will yield same result. This is not the case with being informed, where the agent may be informed that ϕ and immediately after that $\neg\phi$.

Adjusting plausibilities

There are various ways in which an agent may adjust its plausibilities. We give a specific, "preferred" approach, and work with it here. We later indicate how the approach may be generalised.

Consider where the agent is informed or senses that ϕ is true. We assume that if two situations agree on the truth value of ϕ that their *relative* plausibility remains unchanged. That is, since there is nothing that favours one situation over the other, the plausibility of the two situations will change by the same amount. As well, if the agent learns that ϕ is true, then the minimally-ranked ϕ situations will be assigned plausibility 0; consequently other ϕ situations will have their plausibility adjusted down by the same amount. We must also guarantee that no $\neg\phi$ situation has plausibility 0. To this end, we stipulate that the minimum resulting plausibility of a $\neg\phi$ situation is greater than the maximum resulting plausibility of a ϕ situation. This is a strong requirement, and it indicates that the agent places a great deal of faith in ϕ being true. However, it pays off, in that it is intuitive, conceptually simple, and easy to work with. It also corresponds to the operator describe in (Nayak *et al.* 2003), discussed earlier, and so it also provides the appealing characteristic axiom (CNJ).

Implicit in this approach is another assumption, of *recency*. This says that an agent will put greatest faith in the most recently-gleaned item of information. Thus if it is informed first that ϕ is true and then next that ϕ is false, it will believe that ϕ is false. This assumption is very commonly made in the belief change community, although it is not uncontentious (Delgrande *et al.* 2006). Again, it makes the subsequent development simpler.

We describe our approach formally in the next section. However, we emphasise that the framework is very general. It is straightforward, for example, to encode a different approach to modifying plausibilities, or to assert that results of sensing actions are more reliable than those of other action types, or that recency does not hold.

The Formal Account

In this section, we describe the new theory of belief and belief change in the presence of mistaken actions and later outline how it may be generalised.

Mistaken actions

We first introduce the predicate $Alt(a_1, a_2, s)$ to mean that an agent doing physical action a_2 in situation s might think it is doing a_1. (Equivalently, an agent believing it is doing a_1 in situation s might be inadvertently doing a_2 instead.) This predicate is a fluent and different application domains will have different ideas about when an action can be mistaken in this way for another.

For simplicity, we assume that only physical actions can be mistaken for others and that sensing actions and informing actions are always unambiguous. In the very simplest case, no physical action is fallible either: this can formalised by the following axiom:[6]

- $Alt(a', a, s) \equiv a' = a.$

With no fallible actions at all, our characterisation of belief to follow will reduce to previous accounts. However, here are more complex examples:

- $Alt(open(x), a, s) \equiv a = open(x) \vee (x = door \wedge a = null).$

[6] To stick to the syntactic requirements of a basic action theory, two axioms are needed: one concerning initial states and one concerning successor states. To simplify here, we combine them into a single axiom with a universally quantified situation argument.

In this case, an agent believing it has just opened x will either be correct or, when x is the door, the *null* action (an action with no effects) may actually happen instead.

- $Alt(push(x), a, s) \equiv a = push(y) \wedge |x - y| \leq 1$.
  An agent believing it has pushed button x may in fact have pushed some button y that is close to x.

We can also formalise cases where there is context dependency in what actions are fallible. For example,

- $Alt(push(x), a, s) \equiv a = push(y) \wedge$
  $LightOn(s) \supset (x = y) \wedge$
  $\neg LightOn(s) \supset |x - y| \leq 1$.
  When the light is on, the agent will always push the correct button; but when the light is off, the agent can be off by at most one button.

In the last case, actions that change the *LightOn* fluent (such as turning on the light) also end up changing when an action can be mistaken for another.

Our account of belief change

We need to formalise an accessibility relation $B(s', n, \sigma, s)$, where n is the level of plausibility and σ is a sequence of actions that the agent believes it was performing at the time.

We first introduce some abbreviations that will be useful in working with plausibilities. The first, $Ht(d, s)$ asserts that d is the range of plausibility values in B in situation s:

$$Ht(d, s) \doteq \exists s' B(s', d-1, \sigma, s) \wedge \forall s'', d', \sigma' (B(s'', d', \sigma', s) \supset d' < d). \quad (1)$$

That is, if m is the maximum plausibility in B at s then $Ht(m+1, s)$ is true. The second abbreviation is used to assert that the minimum plausibility of ϕ in B at situation s is d.

$$Min(\phi, d, s) \doteq \exists s' B(s', d, \sigma, s) \wedge \phi[s'] \wedge \forall s'', d', \sigma' ((d' < d \wedge B(s'', d', \sigma', s)) \supset \neg\phi[s'']) \quad (2)$$

We also use a third abbreviation, which states that the sensing action a has the same result on situations s and s':

$$Ag(a, s', s) \doteq SF(a, s') \equiv SF(a, s). \quad (3)$$

These abbreviations allow a more compact expression of our successor state axiom for B, to follow.

There are distinct action types: physical actions that change the state of the world, sensing actions, wherein the agent learns the truth value of some formula, and informing actions, wherein the agent is told that some formula is true. These are exhaustive and mutually exclusive. Thus, if $Action(\cdot)$ is true of exactly the action terms, we have:

Axiom 1

$$Action(a) \equiv PhysAct(a) \dot\vee SensingAct(a) \dot\vee InfAct(a)$$

We also assume that if an agent is in an initial situation, then it believes that it is in an initial situation, in that it is not possible that there is a plausible situation that is not initial:

Axiom 2

$$Init(s) \wedge B(s', n, \sigma, s) \supset Init(s') \wedge \sigma = \langle\,\rangle$$

Initially, plausibilities depend on the first argument to B:

Axiom 3

$$Init(s) \wedge B(s', n_1, \sigma, s) \wedge B(s', n_2, \sigma, s) \supset n_1 = n_2$$

We follow Scherl and Levesque in stipulating that the B fluent is Euclidean and transitive for initial situations, and for fixed plausibility and action sequences:

Axiom 4

$$Init(s) \wedge B(s', n, \langle\rangle, s) \supset .\forall s'', m. B(s'', m, \langle\rangle, s') \equiv B(s'', m, \langle\rangle, s)$$

This axiom reflects the fact that we use a "global" notion of plausibility; see the Related Work section for a discussion.

Our successor state axiom for B is somewhat long, to account for how plausibilities evolve under different action types.

Axiom 5

$$\begin{aligned}
&B(s', n, \sigma, do(a, s)) \equiv \\
&\quad \exists s^*, n^*, a^*, \sigma^*, a_i, d. \\
&\quad\quad B(s^*, n^*, \sigma^*, s) \wedge s' = do(a^*, s^*) \wedge \sigma = \sigma^* \cdot a_i \wedge \\
&\quad\quad [PhysAct(a) \wedge Alt(a_i, a, s) \wedge && (4)\\
&\quad\quad\quad (a^* = a_i \wedge n = n^*) \vee && (5)\\
&\quad\quad\quad (a^* \neq a_i \wedge Alt(a_i, a^*, s^*) \wedge && (6)\\
&\quad\quad\quad\quad n = n^* + d \wedge Ht(d, s))\ ] \\
&\quad\quad \vee \\
&\quad\quad [SensingAct(a) \wedge a^* = a \wedge a = a_i \wedge && (7)\\
&\quad\quad\quad ((\ Ag(a, s^*, s) \wedge n = n^* - d \wedge \\
&\quad\quad\quad\quad Min(Ag(a, s^*, s), d, s)) \vee \\
&\quad\quad\quad (\neg Ag(a, s^*, s) \wedge n = n^* + d \wedge Ht(d, s)))\ ] \\
&\quad\quad \vee \\
&\quad\quad [InfAct(a) \wedge a^* = a \wedge a = a_i \wedge a = inf_\phi \wedge && (8)\\
&\quad\quad\quad ((\ \phi(s^*) \wedge n = n^* - d \wedge \\
&\quad\quad\quad\quad Min(\phi, d, s)) \vee \\
&\quad\quad\quad (\neg\phi(s^*) \wedge n = n^* + d \wedge Ht(d, s)))\ ]
\end{aligned}$$

In the first part of the axiom, $B(s', n, \sigma, do(a, s))$ holds just if s' is the result of some other action a^* being carried out in a situation s^* where s^* is plausible according to s. To say more about these terms, and how the new plausibility n is calculated from the previous one n^*, one has to look at the type of the action a.

Consider where a is a physical action (line 4). Here is where we use the σ term representing the sequence of actions believed at the time to have taken place. We have that $\sigma = \sigma^* \cdot a_i$, so a_i is the most recent action believed to have taken place, and it must be an alternative to the actual action a. In line 5, we consider the plausibility of a situation resulting from doing this a_i, and it remains what it was (independent of a). In line 6, we consider the plausibility of an unlikely situation resulting from doing an action other than a_i; and here it will be given a plausibility value greater than any existing plausibility value. (This allows for a later scenario in which the agent determines that it couldn't have executed a_i after all, in which case a^* may be a candidate alternative action.)

The analysis is different for sensing actions (line 7). First, if an agent believes that it has executed a sensing action, then indeed it has executed that action, and so we have $a = a^*$. Sensing is like belief revision, in that if the result of sensing is consistent with the agent's beliefs, then after sensing, the agent will believe what it did before, along with the sensing result. However, if sensing conflicts with the agent's beliefs, then we want the agent's beliefs to be replaced by those beliefs characterised by the most plausible situations in which the sensing result holds. As well, those situations that conflict with the sensing result will become less plausible than any situation that agrees with the sensing result. These cases are handled by the disjunction following line 7.

For informing actions (line 8), the agent is told that a formula ϕ is true, with the result that the agent's beliefs are revised by ϕ. To this end, we assume that for each formula ϕ there is a corresponding informing-of-ϕ action, written inf_ϕ. Moreover, if the agent believes that it was informed of ϕ then this was indeed the case, so $a = a^*$. Then the agent's plausibilities are modified as with sensing, but depending on the truth of ϕ at each situation s^*.

Belief defined

The preceding section specifies how the belief fluent B evolves under various actions. The agent's beliefs are then characterised by the most plausible accessible situations:

$$Bel(\phi, \sigma, s) \doteq \forall s'. B(s', 0, \sigma, s) \supset \phi[s']. \qquad (9)$$

This is a variant of the definition given in (Shapiro *et al.* 2011), but incorporating the σ term. Note that this definition requires appropriate σ and s terms. If the σ and s are different enough (for example, involving sequences of actions of different length), there will be no accessible s', and so all ϕ will be believed. On the other hand, in many cases, we end up investigating what is believed where σ and s have the same sequence of actions (that is, where what is believed to have occurred is correct). We can handle this common occurrence by overloading the B operator:

$$B(s', s) \doteq B(s', 0, seq(s), s) \qquad (10)$$

where $seq(s)$ means the sequence of actions leading to s:

$$(Init(s) \supset seq(s) = \langle\rangle) \wedge (seq(do(a, s)) = seq(s) \cdot a).$$

It is not too hard to show that this common case has all the properties of the Scherl-Levesque B operator.[7] Thus, overloading the Bel operator

$$Bel(\phi, s) \doteq \forall s'. B(s', s) \supset \phi[s']. \qquad (11)$$

leads to a version of belief that precisely mirrors the one formalised by Scherl and Levesque.

Properties

In this section, we use Σ to denote a basic action theory containing the five axioms above and using all the abbreviations introduced (including both versions of B and Bel).

The first property is that, with the exception of alternative actions, a plausibility value is functionally determined by the agent's intended action. Consider where Σ entails $B(s^*, n, \sigma, s)$ and the agent intends to execute action a_i. We obtain:

[7] This holds under the assumption that the Alt relation is reflexive. This is not a very restrictive assumption, and we would expect well-behaved basic action theories to have this property.

Theorem 1

$$\begin{aligned}&\Sigma \cup \{\, B(do(a_1^*, s^*), n_1, \sigma \cdot a_i, do(a_1, s)),\\ &\qquad B(do(a_2^*, s^*), n_2, \sigma \cdot a_i, do(a_2, s))\,\}\\ &\models\\ &\quad (SensingAct(a_i) \vee InfAct(a_i) \vee (a_i = a_1^* \equiv a_i = a_2^*))\\ &\qquad \supset n_1 = n_2\end{aligned}$$

If a_i is a sensing or informing action, all the named actions in the added premisses are the same, and trivially $n_1 = n_2$. Otherwise, if $a_i = a_1^*$ and $a_i = a_2^*$ then $n_1 = n_2$, and similarly, if $a_i \neq a_1^*$ and $a_i \neq a_2^*$ then $n_1 = n_2$. Note that these results are independent of what the agent actually executes, that is, a_1 or a_2. The exceptional case, for example $a_i = a_1^*$ and $a_i \neq a_2^*$ is where the agent intends to execute a_i, and the plausibility of the situation resulting from a_i ($= a_1^*$) is assigned one value, whereas a situation resulting from an alternative action to a_i (viz. a_2^*) is assigned a different value. In this latter instance, it is easy to show that $n_1 < n_2$.

As in Scherl and Levesque, the Euclidean and transitivity properties of B extend to arbitrary situations:

Theorem 2

$$\Sigma \models B(s', n, \sigma, s) \supset . \forall s''. B(s'', n, \sigma, s') \equiv B(s'', n, \sigma, s)$$

Positive and negative introspection with respect to the agent's beliefs then follow:

Theorem 3

$$\Sigma \models Bel(\phi, \sigma, s) \supset Bel(Bel(\phi, now), \sigma, s)$$
$$\Sigma \models \neg Bel(\phi, \sigma, s) \supset Bel(\neg Bel(\phi, now), \sigma, s)$$

It follows straightforwardly from the definition of B that an agent believes the results of sensing or being informed:

Theorem 4

1. $\Sigma \models \phi[s] \supset Bel(\phi, \sigma, do(sense_\phi, s))$
2. $\Sigma \models \neg\phi[s] \supset Bel(\neg\phi, \sigma, do(sense_\phi, s))$
3. $\Sigma \models Bel(\phi, \sigma, do(inf_\phi, s))$

If an agent believes ϕ to hold, then it believes it will believe ϕ after sensing ϕ:

Theorem 5

$$\Sigma \cup \{Bel(\phi, \sigma, s)\} \models Bel(Bel(\phi, do(sense_\phi, now)), \sigma, s)$$

That is, the agent has faith in its beliefs, in that it believes that its beliefs will be borne out by observation. Of course, in the actual situation it may be the case that ϕ is not true, and sensing then would uncover this fact.

If an agent believes that it will believe ϕ after executing action a, then it will in fact believe ϕ after executing action a, and vice versa:

Theorem 6

$$\begin{aligned}&\Sigma \cup \{PhysAct(a) \vee InfAct(a)\} \models\\ &\quad Bel(\phi, \sigma \cdot a, do(a, s)) \equiv Bel(Bel(\phi, do(a, now)), \sigma, s)\end{aligned}$$

An example

Consider the following example, which illustrates various aspects of our approach. Imagine we have some object that may be red (R) or not, an action that paints the object red (pR) and a corresponding sensing action (sR). It is possible that in attempting to paint the object red, the action may fail and the agent execute the $null$ action.

- Fluent: R
- Actions: pR, sR, $null$
- Initial States: S_0, S_1
- Basic Action Theory: Σ contains the five axioms and

  $R(S_1) \wedge \neg R(S_0)$
  $R(do(a, s)) \equiv a = pR \vee R(s)$
  $B(s, 0, \langle\rangle, S_0) \equiv s = S_0 \vee s = S_1$
  $Alt(a', a, s) \equiv a' = a \vee (a' = pR \wedge a = null)$
  $SF(a, s) \equiv a \neq sR \vee R(s)$

  Then we have the following scenarios:

1. Initially the agent doesn't know the colour of the object.

   $\Sigma \models \neg Bel(R, S_0) \wedge \neg Bel(\neg R, S_0)$.

   After sensing the agent knows the block is not red.

   $\Sigma \models Bel(\neg R, do(sR, S_0))$.

2. The agent believes it paints red and does paint red. In terms of the B fluent, we obtain:

   $\Sigma \models B(do(pR, S_0),\ 0, \langle pR\rangle, do(pR, S_0))$
   $\Sigma \models B(do(pR, S_1),\ 0, \langle pR\rangle, do(pR, S_0))$
   $\Sigma \models B(do(null, S_0), 1, \langle pR\rangle, do(pR, S_0))$
   $\Sigma \models B(do(null, S_1), 1, \langle pR\rangle, do(pR, S_0))$

   Hence the agent believes that the object is red.

   $\Sigma \models Bel(R, do(pR, S_0))$.

   The agent also retains as an implausible belief (via B instances with plausibility 1) that it executed the *null* action.

3. The agent believes it paints red but the action fails.

   $\Sigma \models B(do(pR, S_0),\ 0, \langle pR\rangle, do(null, S_0))$
   $\Sigma \models B(do(pR, S_1),\ 0, \langle pR\rangle, do(null, S_0))$
   $\Sigma \models B(do(null, S_0), 1, \langle pR\rangle, do(null, S_0))$
   $\Sigma \models B(do(null, S_1), 1, \langle pR\rangle, do(null, S_0))$

   The agent's beliefs are the same here as in the previous case. In particular we obtain

   $\Sigma \models Bel(R, \langle pR\rangle, do(null, S_0))$,

   and so the agent believes it was successful in the paint action even though it failed.

   If the agent next senses the R fluent, each of the previous B instances has the following (respective) image:

   $\Sigma \models B(do(\langle sR, pR\rangle, S_0),\ 2, \langle sR, pR\rangle, do(\langle sR, null\rangle, S_0))$
   $\Sigma \models B(do(\langle sR, pR\rangle, S_1),\ 2, \langle sR, pR\rangle, do(\langle sR, null\rangle, S_0))$
   $\Sigma \models B(do(\langle sR, null\rangle, S_0), 0, \langle sR, pR\rangle, do(\langle sR, null\rangle, S_0))$
   $\Sigma \models B(do(\langle sR, null\rangle, S_1), 3, \langle sR, pR\rangle, do(\langle sR, null\rangle, S_0))$

   Sensing reveals that R is false; the least B-accessible situation in which R is false is assigned plausibility 0. The B-accessible situations in which R is true have their plausibility increased by the (current) height of the plausibility ordering, which is 2 here. The outcome is as desired: The agent believed it had painted the object red; sensing shows that the object is non-red. In the most plausible situation, the paint action fails, and consequently after sensing the agent believes that the painting action failed:

   $\Sigma \models Bel(\neg R, \langle sR, pR\rangle, do(\langle sR, null\rangle, S_0))$

4. Concerning revision and sensing, we have the following examples. First, an agent will believe the block to be red if so informed:

   $\Sigma \models Bel(R, do(inf_R, S_0))$.

   This can be corrected after sensing:

   $\Sigma \models Bel(\neg R, do(\langle sR, inf_R\rangle, S_0))$.

   However, informing can override sensing:

   $\Sigma \models Bel(R, do(\langle inf_R, sR\rangle, S_0))$.

This last result is a consequence of our assumption of *recency*. It is quite possible that one may not want such a result, preferring for example that sensing always overrule being informed. Such behaviour is easily obtained by adjusting the B fluent successor state axiom. We discuss such modifications and enhancements in the penultimate section.

Belief change operators

Let Σ be some appropriately-defined action theory. The agent's belief set is defined by:

$$BS(\Sigma) = \{\phi \mid \Sigma \models Bel(\phi, S_0)\}$$

Belief revision can be defined as follows:

$$BS(\Sigma * \phi) = \{\psi \mid \Sigma \models Bel(\psi, do(inf_\phi, S_0))\} \qquad (12)$$

We obtain the following result directly from (Nayak *et al.* 2003).

Theorem 7 *For any action theory Σ and formula ϕ, the AGM Postulates (K*1)-(K*8) are satisfied when $*$ is defined as in (12). Moreover, the Darwiche-Pearl Postulates (DP1)-(DP4) are satisfied, as is (CNJ).*

We can similarly examine sensing actions. In the case of sensing a formula ϕ, the agent will believe either ϕ or $\neg\phi$ depending on whether or not ϕ is true in the situation at hand. Assume that for any ϕ, there is a sensing action $sense_\phi$ such that $(SF(sense_\phi, s) \equiv \phi(s))$ is entailed. Define:

$$BS(\Sigma *_s \phi) \ = \ \{\psi \mid \Sigma \models Bel(\psi, do(sense_\phi, S_0))\} \qquad (13)$$

The following result is straightforward:

Theorem 8 *For any action theory Σ and formula ϕ*

1. *$\phi \in BS(\Sigma *_s \phi)$ iff $\Sigma \models \phi(S_0)$.*
2. *If $\Sigma \models \phi(S_0)$ then $BS(\Sigma *_s \phi) = BS(\Sigma * \phi)$*
   If $\Sigma \models \neg\phi(S_0)$ then $BS(\Sigma *_s \phi) = BS(\Sigma * \neg\phi)$

Katsuno-Mendelzon style update doesn't make much sense from the point of view of the agent. Recall that in update, a formula ϕ is recorded as being true following the execution of some action, and the task is to determine what else is true. In our framework, an agent is fully aware of the effects of the actions it believes that it has executed; and so its beliefs are simply the image of its previous beliefs under this intended action. To say that ϕ is true in an update is either redundant, since the agent knows the effects of its intended actions, or, in the case where the agent is informed about ϕ, corresponds to a revision.

Consider however the scenario where an agent is seen to execute some action, and it is subsequently learned that the agent believes ϕ. One can then ask: what else might the agent believe? This can be determined via the following notion of *external update*. Given an action theory Σ, call ϕ *compatible* with action a in Σ if for some a' we have $Alt(a', a, S_0)$ and $Bel(\phi, \langle a' \rangle, do(a, S_0))$ holds. Then for ϕ compatible with a in Σ, define:

$$BS(\Sigma \diamond_a \phi) = \{\psi \mid \forall a' \text{ s.t. if } \Sigma \models Alt(a', a, S_0) \land Bel(\phi, \langle a' \rangle, do(a, S_0)) \text{ then } \Sigma \models Bel(\psi, \langle a' \rangle, do(a, S_0))\} \quad (14)$$

So the agent has been seen to execute action a; and one also knows (perhaps by being told by the agent) that the agent believes ϕ. The operator $\diamond_a$ specifies what else the agent may believe, by considering what the agent would believe if it had inadvertently executed the wrong action. We obtain:

Theorem 9 *For any action theory Σ, action a, and formula ϕ compatible with a in Σ, the KM Postulates (K$\diamond$0), (K$\diamond$1), (K$\diamond$3)-(K$\diamond$8) are satisfied when $\diamond_a$ is defined as in Definition 14.*

Related work

As noted, the closest work to the present is (Shapiro *et al.* 2011). This approach is expressed in terms of an epistemic extension to the situation calculus, where an agent acts in an environment and can sense whether a condition holds. An agent has a plausibility ordering over (accessible) situations. Because actions always execute predictably and there are no informing actions, this allows a significant simplification in dealing with sensing results that conflict with the agent's beliefs: If the agent senses that ϕ holds, then all accessible $\neg\phi$ situations are eliminated from the plausibility ordering. Thus as the agent carries out sensing actions, the number of accessible situations monotonically decreases. The evolution of the agent's beliefs can be regarded as a process of correcting the agent's initially-incorrect beliefs, interspersed with physical actions that map each accessible situation to its image under the action. If an agent were to receive two conflicting sensor reports, it would fall into inconsistency.

Shapiro *et al.* consider the modelling of belief change by adjusting plausibilities (as we do here) but they reject this option. One of their goals (and ours) is to handle positive and negative introspection properly. They argue that any reasonable scheme for updating plausibilities of accessible situations conflicts with an account of introspection. An account of fully-introspective belief requires that the belief accessibility relation be transitive and Euclidean, which they express as follows:

$$\exists n B(s', n, s) \supset (\forall s'', m.\ B(s'', m, s') \equiv B(s'', m, s)) \quad (15)$$

Their interpretation of $B(s', n, s)$ is that "in s, the agent thinks s' (is) possible with ... plausibility of n" (p. 8). However, we suggest that (15) is not entirely compatible with this interpretation, and that in fact it is an overly strong condition to require for an introspective agent.

Consider an example with fluents p and q, where p means my laptop is working and q means the weather is fine. Let p, q be true at S_0; $\neg p$, q be true at S_1; and $\neg p$, $\neg q$ be true at S_2. Suppose my laptop is very reliable, but the weather is highly variable. So assume that we have $[B(s', 0, S_0) \equiv (s' = S_0)]$, $B(S_1, 5, S_0)$ and $B(S_2, 6, S_0)$. Thus I believe in S_0 that my laptop is working and it is fine; it is highly implausible that my laptop isn't working and it is fine, it is only slightly more implausible that my laptop isn't working and it is not fine. Formula (15) stipulates that $B(S_2, 6, S_1)$, that is, in a situation where I believe my laptop isn't working and it is fine, it is highly implausible that my laptop isn't working and it is not fine. But this seems incorrect; the relative difference in plausibility between the two situations is not all that great, and one would expect something more like $B(S_2, 1, S_1)$, contradicting (15).[8]

Other related work has been carried out with respect to *action languages* (Gelfond and Lifschitz 1998), where the underlying semantic structure is a *transition system*. A transition system is a directed graph, where vertices represent states of the world, and (labelled) edges give action transitions. An action language is, not surprisingly, a language for describing the effect of actions, where meaning can be attached to sentences of the language in terms of an underlying transition system. An epistemic extension is defined by characterising an agent's beliefs by a subset of the vertices; and the evolution of the agent's beliefs following an action a is given in the expected way by following a-labelled edges. (Lobo *et al.* 2001) and (Son and Baral 2001) describe action languages for expressing the state of an agent's knowledge, and how it evolves following physical and sensing actions.

(Hunter and Delgrande 2011) consider the problem of sensing and acting where the agent's initial beliefs may be incorrect. They show that an iterative process of successively determining the effects of actions and observations leads to difficulties. Instead they propose a notion of *belief evolution* where, if a sensing action is inconsistent with the agent's beliefs, both the result of the sensing action and the agent's beliefs are regressed to the initial state, the agent's (initial) beliefs are revised, and the result projected back to the current time point. A point of contrast with the present approach is that in an action language, one keeps track of an evolving set of states of the world. In the situation calculus, one keeps track of a set of situations, and so implicitly a set of histories.

There has also been a substantial amount of work on dealing with uncertain observations and actions. Much of this work is based on ordinal conditional functions (OCFs). (Boutilier 1998) presents a model that encompasses classical accounts of revision and update. Part of the motivation of this work is that an agent making an observation may not know what event caused that observation, and so exceptional action effects may be modelled. Then ranked explanations possibly involving past (exogenous) events are used to account for observations. In the approach, an event maps each world to a κ-ranking over worlds. Each world has associated with it an event ordering that describes the plausibility of

[8] Note that in our approach, we are interested in introspection of beliefs held by the agent (so where $n = m = 0$ in (15)). Formula (15) doesn't hold; rather we have the weaker result given in Theorem 2.

event occurrences. (Goldszmidt and Pearl 1992) also studies belief update modelled via OCFs. (Boutilier *et al.* 1998) presents a semantic approach, again based on ordinal conditional functions, that accommodates unreliable observations. Somewhat similarly, (Bacchus *et al.* 1999) consider noisy observations in the situation calculus. (Laverny and Lang 2005) brings together many of the aspects that we have considered, including unreliable observations and normal and exceptional action effects, via *knowledge based programs*, again founded on OCFs.

There has also been work in dynamic epistemic logic that is reminiscent of our setting in which an agent may inadvertently execute the wrong action. (Baltag *et al.* 1998) considers the case where an agent or group of agents believes that some event occurred, but another group of agents believes that some other event was what occurred. In this case, the actions are public or semi-public announcements, and an agent's knowledge strictly increases; hence such change is more like expansion than revision.

Future Work

We have presented a specific approach that addresses and integrates possibly-failing actions, sensing, and revision. In this section we briefly suggest some ways in which the framework may be generalised or in which one may obtain a more nuanced approach.

First, we have assumed that (physical) action execution, sensing, and being informed are all equally reliable, in that each action type brings about the same change in the plausibility ranking. Clearly, these classes of actions need not be equally reliable, and one might, for example, want to stipulate that the result of any sensing action is strictly more reliable than any informing event. This is most easily accommodated by generalising the notion of a plausibility value to that of a vector. For example, plausibility could be represented by a pair $(n, m,)$, where number comparison would be done lexicographically: $(n, m) \leq (n', m')$ iff $n \leq n'$ and if $n = n'$ then $m \leq m'$. Then, for sensing to strictly be more reliable than informing, a sensing action would modify the first value in an ordered pair, while informing would modify the second. The result would be that, for example, if the agent sensed that a light was on, but was then informed that it was off, it would continue to believe that the light was on.

Second, we have adapted the approach of (Nayak *et al.* 2003) to deal with failing actions, sensing, and revision. The result is a very strong commitment to the truth of any particular outcome. Thus, if informed that a light is on, the agent's plausibilities will be such that every resulting situation in which the light is on will be more plausible than any resulting situation in which it is off. A more nuanced approach can be obtained by simply modifying the predicate Ht. In fact, it can be observed that the full range of conditionalisation functions given in (Spohn 1988) may be captured by suitably modifying Ht. For example, by replacing (1) by $Ht(d, s) \doteq d = 1$ one obtains the concrete operator described in (Darwiche and Pearl 1997).

This leads to the consideration that one could encode a general notion of *reliability*, again by modifying Ht. For example, there is no obstacle to expressing that a sensor reading comes with a certain amount of noise, or that one sensor is more reliable than another, or that an informing action is not entirely reliable. Thus for example, in sensing the distance to an object, a sensor may give a result that is accurate to some given tolerance. Or one could express that an agent doesn't believe that a light is on, and after being informed by an unreliable agent that the light is on, it still believes that the light is off, but with less confidence than previously.

These considerations extend to possibly-failing actions, which can be generalised to provide a more fine-grained account. Thus the notion of action alternatives can be generalised from an absolute notion to a relative one. For example, if the agent intends to push a particular button, then it may be more likely to inadvertently push a button closer to the intended button than one further away. This can be handled by having different *Alt* actions be assigned different plausibilities in (6).

Last, as an account of reasoning about action, there are numerous ways in which the approach could be extended For example, it would be useful to handle action effects from other agents, including exogenous actions. As well, nondeterministic actions could be taken into account.

Conclusion

We have developed a general model of an agent that is able to reason and maintain its belief state in a setting in which it may inadvertently execute the wrong action, and in which it may sense or be informed about its environment. This approach is developed within the epistemic extension of the situation calculus. It also incorporates notions from belief revision, specifically, we use a plausibility ordering to represent the agent's belief state.

We describe actions in terms of their preconditions and their effects, making use of Reiter's solution to the frame problem. We augment this by admitting the case where an agent may intend to execute one action but inadvertently executes another. Consequently the agent's beliefs may evolve according to one sequence of actions (the actions it believes that it executed) while the world evolves in a different direction (given by the actions that the agent actually executes). The agent may carry out sensing actions, and via such actions the agent may correct its beliefs. As well, an agent may revise its beliefs by being informed of a fact.

This approach extends previous work in several respects. It provides an integration of reasoning about action with sensing and belief revision. In so doing, it allows arbitrary action sequences involving (possibly fallible) physical actions, sensing actions, and informing actions. We retain the results of basic action theories, and so inherit the formal results attending such theories. While we present a specific approach, the overall framework is quite general – for example it is straightforward to extend the notion of change to that of Spohn's conditionalisation. Consequently the approach may serve as a platform from which to examine various issues in belief change and reasoning about action – for example the relation between sensing and being informed, or the interplay of knowledge gained from various action types.

Acknowledgements

We thank Yongmei Liu (who suggested that the *Alt* predicate be situation dependent) and Bryan Renne for useful discussions concerning dynamic epistemic logic. We also thank the reviews for their pertinent and helpful comments. Financial support was gratefully received from the Natural Sciences and Engineering Research Council of Canada.

References

F. Bacchus, J.Y. Halpern, and H.J. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2):171–208, 1999.

A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, TARK '98, pages 43–56. Morgan Kaufmann, 1998.

Craig Boutilier, Nir Friedman, and Joseph Y. Halpern. Belief revision with unreliable observations. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 127–134, 1998.

C. Boutilier. Iterated revision and minimal revision of conditional beliefs. *Journal of Logic and Computation*, 25:262–305, 1996.

Craig Boutilier. A unified model of qualitative belief change: A dynamical systems perspective. *Artificial Intelligence*, 98(1-2):281–316, 1998.

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

J. Delgrande, D. Dubois, and J. Lang. Iterated revision as prioritized merging. In P. Doherty, J. Mylopoulos, and C. Welty, editors, *Proceedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning*, pages 210–220, Lake District, UK, 2006. The AAAI Press/The MIT Press.

P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.

Michael Gelfond and Vladimir Lifschitz. Action languages. *Electronic Transactions on AI*, 3, 1998.

Moisés Goldszmidt and Judea Pearl. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning*, pages 661–672, 1992.

Aaron Hunter and James P. Delgrande. Iterated belief change due to actions and observations. *Journal of Artificial Intelligence Research*, 40:269–304, 2011.

Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.

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

Noël Laverny and Jérôme Lang. From knowledge-based programs to graded belief-based programs part I: On-line reasoning. *Synthese*, 147(2):277–321, 2005.

H.J. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18), 1998.

J. Lobo, G. Mendez, and S.R. Taylor. Knowledge and the action description language $\mathcal{A}_\mathcal{K}$. *Theory and Practice of Logic Programming*, 1(2):129–184, 2001.

J. McCarthy and P.J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In D. Michie and B. Meltzer, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

R.C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.

Abhaya C. Nayak, Maurice Pagnucco, and Pavlos Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003.

P. Peppas. Belief revision. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, San Diego, USA, 2008.

R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, Cambridge, MA, 2001.

Richard B. Scherl and Hector J. Levesque. The frame problem and knowledge-producing actions. In Richard Fikes and Wendy Lehnert, editors, *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 698–695, Menlo Park, California, 1993. American Association for Artificial Intelligence, AAAI Press.

R.B. Scherl and H.J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.

Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.

T. Son and C. Baral. Formalizing sensing actions: A transition function based approach. *Artificial Intelligence*, 125(1-2):19–91, 2001.

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