

Sensing and revision in a modal logic of belief and action

Andreas Herzig¹ and Dominique Longin²

Abstract. We propose a modal logic of belief and actions, where action might be nondeterministic, and there might be misperception.

The agent must be able to revise his beliefs, because (contrarily to knowledge) observations might be inconsistent with his beliefs. We propose a new solution in terms of successor state axioms, which does not resort to orderings of plausibility. Our solution allows for regression in the case of deterministic actions.

1 Introduction

Since the beginning of the 90ies, solutions to the frame problem have been extended to cover perception [17, 13, 20, 23]. In these approaches perception has been analyzed in terms of actions. To such perception actions one can oppose *uninformative actions*, which are actions whose outcome is not perceived by the agent. (When the agent learns that such an action has occurred, he is nevertheless able to predict its results according to the action laws.) It is noted in several places (e.g. [17],[21, footnote 10],[10]) that actions can be analysed as a sequence of uninformative actions and perception actions. For example, the action of tossing a coin can be decomposed into the uninformative action of tossing without observing the result – eyes shut –, followed by the perception action of checking the result. The most important class of uninformative actions are ontic actions (physical actions), which are actions that can be described without referring to belief.

Perception actions are reduced to actions of observing that some proposition is true: that the light in some room is on, that tossing a coin resulted in heads, etc. We call such actions *observation actions*. We suppose that they do not change the environment, but only the agent’s mental state. (For the sake of simplicity we suppose that there is only one agent.)

When reasoning about observations one has to distinguish what is true from what is believed by an agent: it might be the case that some proposition A is true, but the agent is not aware of it. Therefore we suppose that in every situation (alias possible world) w the agent entertains a set of beliefs $B(w)$.

Now suppose some action a occurs, resulting in a new situation w' . What is $B(w')$ like? If a is uninformative, then $B(w')$ should only depend on $B(w)$, a and the action laws: the agent predicts the result of a using the action laws for a . Indeed, apart from the mere execution of a the agent should learn nothing about a ’s particular effects that hold in w' .

According to this account, observation actions are uninformative: to learn that the observation of A has occurred means to learn that A . All the relevant information is thus encoded in the notification

of the action occurrence. Then to take into account a amounts to incorporate A into $B(w)$.³

Take the action of tossing a coin. When the agent is notified the occurrence of the tossing action, as he cannot observe the effects of toss, he predicts them in an a priori way, according to his mental state and the action laws. The agent can thus be said to “mentally execute” toss. Hence afterwards he believes that Heads \vee Tails holds, but neither believes Heads nor Tails. When the agent subsequently learns that the coin fell heads (being notified that Heads has been observed) then he moves to believing that Heads.

In consequence we can restrict our attention to uninformative actions. We focus on the following type of scenarios:

- in a given situation w the agent entertains a set of beliefs $B(w)$;
- some action a occurs, resulting in a new situation w' ;
- the agent is notified that some a' has occurred (where $a' \neq a$ if there is misperception);
- the agent does not learn which of the effects of a hold in w ;
- the agent takes into account the occurrence of a' by appropriately changing $B(w)$, and forms the new set $B(w')$ that he holds in w' .

We have thus generalized our account to allow for misperception.

Most of the approaches in the reasoning about actions domain are formulated in terms of a modal operator of knowledge [14, 17, 20, 23, 21, 10]. Knowledge being viewed as true belief, in such approaches *surprises* are impossible: if an agent knows that A then A must be true; as observations don’t change the environment, A still holds after any observation; hence $\neg A$ can only be observed if there is misperception, but in this case the agent realizes that, and immediately rejects the input.

It follows that two operations are enough to implement knowledge change: updates *à la* Katsuno-Mendelzon (KM-updates) [11] to take into account uninformative actions, and expansions *à la* Alchourrón-Gärdenfors-Makinson (AGM-expansions) [1] to take into account perception actions.

The picture is different in the case of belief change, because beliefs can be contradicted by observations: I believe that I have a coin in my pocket, but on checking I find out I don’t; I believe my watch is waterproof but when trying it out it isn’t, etc. It is non-trivial to extend the above solutions to handle such examples. Expansion operations do not suffice: we need belief *revision* operations *à la* AGM.

More generally, the problem arises as soon as the agent believes some action is inexecutable and nevertheless learns that it has occurred. In this case the agent must first revise his current beliefs by the preconditions of the action, and then apply the action laws associated to the action.

Many authors raise the issue and are aware of the difficulties (e.g.

¹ Institut de Recherche en Informatique de Toulouse, France, email: Andreas.herzig@irit.fr

² Laboratoire Travail et Cognition, Université Toulouse-le-Mirail, France, email: longin@univ-tlse2.fr

³ This does not hold for all perception actions, such as testing if a proposition is true. Such tests can nevertheless be reduced to uninformative actions, see Sect. 2.5.

[17, 6]), but the only proposal up to now is that of Shapiro *et al.* [21], which is based on orderings of plausibility.

In the sequel we shall do without such a device. Our approach is characterized by the following hypotheses:

- (H1) All atomic actions are either ontic or observation actions.
- (H2) Actions might be non-deterministic.
- (H3) There might be misperception and non-perception of action occurrences.
- (H4) Uninformative actions do not affect the agent’s cognition. Hence we exclude actions such as modifying the agent’s memory.
- (H5) The action laws are known by the agent.

In Sect. 2 we introduce a logic for belief and action, which is similar to Segerberg’s DDL [19, 18]. Then we focus on uninformative and observation actions, and show how updating and revision can be done (Sect. 3, Sect. 4, Sect. 5). Finally we discuss related work (Sect. 6).

2 Dynamic Doxastic Logic

2.1 Belief

We suppose that our language contains a modal operator of belief Bel . The formula $BelA$ is read “the agent believes that A ”. $BellfA$ is defined as $BellfA \stackrel{\text{def.}}{=} BelA \vee Bel\neg A$, and can be read “the agent believes A or believes $\neg A$ ”, or more shortly “the agent knows⁴ whether A is true or not”.

We adopt the modal logic KD45 as the logic of belief, i.e. we suppose agents do not entertain inconsistent beliefs, and are aware of their beliefs and disbeliefs.

2.2 Actions

We use a simple version of PDL [8] to speak about actions. Actions are noted a, a', b, \dots . The empty action is noted λ . To each action a there is associated a modal operator $After_a$. The formula $After_aA$ reads “ A is true after a ”. $After_a\perp$ expresses that a is inexecutable. An example of a formula involving belief and action is $Bel\neg After_a\perp \wedge After_a\perp$, expressing that the agent believes that a can be executed, while this is not the case. The operator $Feasible_a$ is introduced as an abbreviation: $Feasible_aA \stackrel{\text{def.}}{=} \neg After_a\neg A$. $Feasible_a\top$ expresses that a is executable.

We adopt the standard axiomatics of PDL, which for our fragment is nothing but the multimodal logic K. ($After_a$ corresponds to the Dynamic Logic operator $[a]$, and $Feasible_a$ to $\langle a \rangle$.)

2.3 Possible worlds semantics

We adopt the standard possible worlds semantics, with models having a set of situations W , and accessibility relations R_{Bel} and R_a respectively associated to the modal operators Bel and $After_a$. We view the belief state of an agent in a given situation w as a set of possible worlds $R_{Bel}(w) = \{v : wR_{Bel}v\}$, and $v \in R_{Bel}(w)$ means that the situation v is compatible with the agent’s beliefs. $R_a(w) = \{w' : wR_a w'\}$ is the set of possible results w' of action a when applied in w .

R_{Bel} is reflexive, transitive and euclidean. In a situation $w \in W$, the set of situations $R_{Bel}(w) = \{u \in W : wR_{Bel}u\}$ is called the *belief state* of the agent in w .

⁴ We use the term ‘knows’ here because ‘the agent believes whether A ’ sounds odd.

Actions might be nondeterministic (because the R_a are not necessarily functions), and they might be inexecutable (when there is no w' such that $wR_a w'$).

2.4 Misperception

In most of the related approaches [17, 21, 23] it is supposed that actions are public: when a occurs then its occurrence is correctly notified to the agent. This means that (1) if an agent believes that some action a occurred then a indeed occurred (correctness), and (2) if a occurred then the agent believes a occurred (completeness).

We suppose here that the agent might instead perceive some other action b . The atomic proposition $perc(a, b)$ expresses that the occurrence of a has been notified as the occurrence of b by the agent. Note that setting b to the empty action allows to simulate the case where the agent is unaware that an action has occurred. The other way round, setting a to the empty action allows to simulate illusions.

Note that in terms of Sandewall’s systematic approach [16], most of the approaches in reasoning about actions suppose that knowledge is explicit, accurate and correct (Sandewall’s class \mathcal{K}). Hence there is no misperception or illusion.

Several approaches to misperception exist in the literature, e.g. [2, 4]. In [4] a classification is given. Among the three cases there, we can account here for the case where observations do not agree with the effects that actions are supposed to have, and the case where new observations indicate unpredicted change (conflicting with the principle of inertia). Among the different revision strategies that are discussed, the one we adopt here is that of preferring the last observation when constructing the new belief state.

2.5 Observation actions

We note $observe(A)$ the action of observing that A . Observation actions can be characterised by the following *logical axioms*⁵ (see [9, 10]).

$$\begin{aligned}
A &\rightarrow Feasible_{observe(A)}\top && \text{(TestAct}_1\text{)} \\
\neg A &\rightarrow After_{observe(A)}\perp && \text{(TestAct}_2\text{)} \\
After_{observe(A)}A &&& \text{(TestAct}_3\text{)} \\
C &\rightarrow After_{observe(A)}C \text{ if } C \text{ is objective} && \text{(TestAct}_4\text{)} \\
Feasible_{observe(A)}C &\rightarrow After_{observe(A)}C && \text{(TestAct}_5\text{)}
\end{aligned}$$

An *objective formula* is a formula without occurrences of the doxastic modal operator Bel . The first two axioms together say that $observe(A)$ is executable iff A is true. Therefore learning that $observe(A)$ has been executed amounts to learning that A . (TestAct₃) says that A holds after observing that A . Together with the more general principles of sections Sect. 3 and Sect. 4 it will guarantee that observing that A leads to believing that A . (TestAct₄) expresses that observation actions are perception actions, and the last says they are deterministic.⁶

Observation actions behave as expansions in the AGM-theory: $observe(A)$ makes shrink the belief state by ‘throwing out’ those possible situations where A is false.

⁵ Note that logical axioms are known by the agent. (This is obtained by the necessitation rule of KD45.)

⁶ $observe(A)$ is similar to the PDL test “ $A?$ ”. The difference is that for the latter $After_{A?}C$ is defined as $A \rightarrow C$. Hence such tests validate $B \rightarrow After_{A?}B$ for every formula B . However, consider $B = \neg BellfA$: intuitively, the formula $\neg BellfA \rightarrow After_{A?}\neg BellfA$ should not be valid. Therefore such a principle must be restricted to objective B ’s, which is what we did here.

Nondeterministic composition of $observe(A)$ and $observe(\neg A)$ can ‘simulate’ the action $testIf(A)$ of testing-if A : testing if the coin is heads amounts to nondeterministically choose between $observe(A)$ and $observe(\neg A)$ and execute the chosen action. Therefore testing-if can be viewed as an abbreviation of testing-that:

$$After_{testIf(A)} B \stackrel{\text{def.}}{=} After_{observe(A)} B \wedge After_{observe(\neg A)} B$$

It can then be proved that $Feasible_{testIf(A)} \top$ holds, as well as $After_{testIf(A)} BelIf A$ and $A \rightarrow After_{testIf(A)} Bel A$. Note that while $observe(A)$ is an uninformative action, $testIf(A)$ is not.

2.6 Ontic actions

We suppose that to each ontic action a there is associated a set of *effect laws* and a set of *executability laws*. The former are of the form $A \rightarrow After_a C$ and the latter are of the form $A \rightarrow Feasible_\alpha \top$ where A and C must be *factual*, i.e. without any modal operator.

For example, the (ontic) toss action is executable if one has a coin: $HasCoin \rightarrow Feasible_{toss} \top$, and has the effects $Heads \vee Tails$ and $\neg HasCoin$: $After_a((Heads \vee Tails) \wedge \neg HasCoin)$.

3 Updating

Semantically a (non-deterministic) action is a relation R_a between possible situations – alias possible worlds –, where $w' \in R_a(w)$ means that w' is a possible result of a when applied in w . We view the belief state of an agent in a given situation w as a set of possible worlds $R_{Bel}(w)$, and $v \in R_{Bel}(w)$ means that the situation v is compatible with the agent’s beliefs. The occurrence of an action makes the current situation w evolve to a new situation $w' \in R_a(w)$. What can we say about $R_{Bel}(w')$, i.e. the agent’s belief state at w' ?

First of all, it might be the case that the agent is not notified of a , but some other action b , as expressed by the atomic proposition $perc(a, b)$. How should the agent take into account such an action occurrence b ? Following Moore [14] and Scherl and Levesque [17], the agent’s belief state $R_{Bel}(w')$ in w' results from applying action b to all possible worlds in $R_{Bel}(w)$ (“mentally executing b ”), and collecting the resulting situations:

$$R_{Bel}(w') = \bigcup_{v \in R_{Bel}(w)} R_b(v)$$

This looks fine, but there is a problematic case here when $R_b(v) = \emptyset$ for every $v \in R_{Bel}(w)$: then $R_{Bel}(w') = \emptyset$, which would mean that the agent ends up with an inconsistent set of beliefs. This contradicts our hypothesis that beliefs are consistent (axiom D). Under such a proviso our axiom for updates is the generalization of the successor state axiom for knowledge of [17] to belief and non-deterministic actions:

$$(perc(a, b) \wedge \neg After_a \perp \wedge \neg Bel After_b \perp) \rightarrow (Feasible_a Bel A \leftrightarrow Bel After_b A) \quad (SSA_1)$$

where a is an uninformative action. It says that the agent cannot observe anything after a is performed: indeed, for any formula A , if he cannot predict before a is performed that A will hold after a is performed, then he will not know A after a is performed.

Consider e.g. $a = \text{toss}$, and suppose $perc(\text{toss}, \text{toss})$ and $HasCoin \wedge Bel HasCoin$ hold. It follows from the executability laws for toss that $\neg After_{\text{toss}} \perp$ and $\neg Bel After_{\text{toss}} \perp$ (the latter by necessitation and axiom D). We therefore have for $A = \text{Heads}$: $Feasible_{\text{toss}} Bel Heads \leftrightarrow Bel After_{\text{toss}} Heads$. From the left to the

right, (SSA_1) expresses that for uninformative actions there are no a posteriori beliefs the agent didn’t already hold a priori: if after some execution of toss the agent believes that the coin is heads (i.e. $Feasible_{\text{toss}} Bel Heads$) then – as he had no means to check whether Heads is true – he must have believed before tossing that the coin is biased, i.e. $Bel After_{\text{toss}} Heads$. Reading the equivalence from the right to the left, consider $A = \neg Bel If Heads$. (SSA_1) expresses that the agent’s uncertainty about the nondeterministic result of toss is preserved through its execution: before executing toss the agent ignores whether heads or tails will result: $Bel After_{\text{toss}} \neg Bel If Heads$ and this disbelief $\neg Bel If Heads$ is preserved through the execution of toss : $Feasible_{\text{toss}} Bel \neg Bel If Heads$, which is equivalent to $Feasible_{\text{toss}} \neg Bel If Heads$.

To take another example, suppose $Bel(\neg Heads \wedge \neg Tails)$ holds, and suppose that a toss -action takes place but the agent isn’t notified. Hence $perc(a, \lambda)$ holds, and $After_{\text{toss}} Bel(\neg Heads \wedge \neg Tails)$.

As said in section 2, observations are uninformative actions, and (SSA_1) applies to observation actions, too: if after observing that Heads I believe that $Heads \wedge A'$ for some A' , and $observe(Heads)$ is executable, then I believe before $observe(Heads)$ that $Heads \wedge A'$ will be true afterwards.

Remark 1 *The only case where the \rightarrow direction of the equivalence in (SSA_1) cannot be accepted is when a erases all or part of the memory of the agent (e.g. taking off the batteries of a robot). The \leftarrow direction is counter-intuitive only if the agent knows that a adds unjustified information to his memory. This is the case e.g. when he is hypnotized or takes drugs. We have excluded such extreme cases by hypothesis (H4).*

Finally, as announced in the end of section 2, from (SSA_1) , $(TestAct_3)$ and standard modal principles we can prove the following theorem:

Theorem 1 $After_{observe(A)} Bel A$

4 Revising

4.1 Enabling actions

Suppose the agent believes a is inexecutable, and learns that a has nevertheless occurred. Axiom (SSA_1) says nothing about that case. Such surprising occurrences of actions are indeed problematic, because the agent is unable to just mentally execute a , and must first change his beliefs about a ’s preconditions.

We propose to formalize the operation of changing beliefs about preconditions by the mental execution of a particular ontic action whose effect is to make the executability preconditions of a true. This makes sense when applied to possible worlds: if an agent believes a to be inexecutable but nevertheless learns that a has happened, he adjusts each of his possible worlds w so as to enable executability of a in w .

Formally, we associate to every atomic action a an action $enable_a$, and we say that $enable_a$ makes a executable.

We postulate that $enable_a$ can occur in every situation. Hence we have the executability law

$$Feasible_{enable_a} \top \quad (Exec_{enable_a})$$

This means that for every action there is at least one situation where it is executable. Note that this excludes from our actions the action $observe(\perp)$ which is never executable.

Let a be any ontic action. According to our definition, the set of executability laws for a has the form $\{A_1 \rightarrow Feasible_a \top, \dots, A_n \rightarrow Feasible_a \top\}$

As $enable_a$ makes a executable, the set of executability laws for a determines the following effect law for $enable_a$:

$$After_{enable_a}(A_1 \vee \dots \vee A_n) \quad (Eff_{enable_a})$$

For example, for the toss action, if the agent believes there is no coin, and nevertheless learns that toss has been executed, then he enables toss in his possible worlds: $After_{enable_{toss}} HasCoin$.

What about the observation actions? The executability precondition of $observe(A)$ being A , to make A executable amounts to making A true in the actual situation. Hence we have in this case the axiom

$$After_{enable_{observe(A)}} A \quad (TestAct_6)$$

4.2 The axiom for revision

We are now able to postulate the following axiom for uninformative actions, which applies when revision is needed:

$$(perc(a, b) \wedge BelAfter_b \perp \wedge \neg After_a \perp) \rightarrow (Feasible_a BelA \leftrightarrow BelAfter_{enable_b} After_b A) \quad (SSA_2)$$

Semantically, this means that when the agent is notified that b has occurred, and believes that b is inexecutable, then the possible situations after a are obtained by:

1. enabling b in every possible situation;
2. applying b to these situations;
3. collecting the resulting situations.

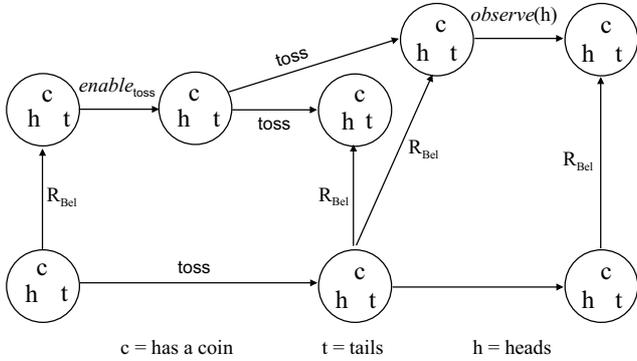


Figure 1. the toss example

Let us illustrate (SSA_2) by our running example. One of the possible Kripke models is given in figure 1. Suppose initially the agent ignores that there is a coin: $HasCoin \wedge Bel \neg HasCoin$. The tossing action is therefore executable, but the agent believes it isn't.

Suppose the coin is tossed resulting in Heads, and suppose the agent is correctly notified that toss has been executed: $perc(toss, toss)$. In a first step the agent enables tossing by making its executability condition $HasCoin$ true, and then mentally executes toss. Putting these two actions together produces the resulting belief state, which is composed of a world where heads holds, and another one where tails holds. Syntactically, from (SSA_2) and the laws for

toss we obtain $After_{toss} Bel(Heads \vee Tails)$, i.e. the agent believes the coin fell either heads or tails.

When the agent subsequently perceives that heads holds (via learning the occurrence of $observe(Heads)$) he then eliminates the world where tails holds from his belief state.

From SSA_1 and (SSA_2) one can derive a principle of *doxastic determinism*:

Theorem 2 $Feasible_a BelA \rightarrow After_a BelA$

5 Preserving facts

Given our successor state axioms we can reuse non-epistemic solutions to the frame problem.

Just as Scherl and Levesque have applied Reiter's solution [17] we use the solution of [3] in order to stay within propositional logic.

Which truths can be preserved after the performance of an uninformative action? Our key concept is that of the *influence of an action*. If there exists a relation of influence between the action and an atom p , then p cannot be preserved. The relation $a \rightsquigarrow p$ is read "the action a influences the truth value of p ". In our example, $\rightsquigarrow = \{toss \rightsquigarrow Heads, toss \rightsquigarrow Tails, toss \rightsquigarrow HasCoin\}$. Note that \rightsquigarrow is in the metalanguage. We extend \rightsquigarrow to formulas by stipulating that $a \rightsquigarrow A$ if there is an atom p occurring in A such that $a \rightsquigarrow p$.

The concept of influence (or dependence) is close to notions that have recently been studied in the field of reasoning about actions in order to solve the frame problem, e.g. Sandewall's [16] occlusion, Thielscher's [22] influence relation, or the 'possibly changes' operators of Giunchiglia *et al.* [7].

The preservation of formulas that are not influenced by an action is formalized by the influence-based *logical axiom*

$$A \rightarrow After_a A \text{ if } a \not\rightsquigarrow A \text{ and } A \text{ is factual} \quad (Preserv)$$

This expresses that if a does not influence A then A is preserved. The restriction that A be factual avoids e.g. $Feasible_a \top \rightarrow After_a Feasible_a \top$, which is not necessarily the case because a might modify the executability preconditions of a' .

6 Discussion and related work

We have defined a modal logic of belief and nondeterministic actions where the agent's beliefs about the action laws might be inaccurate. Our central axioms (SSA_1) and (SSA_2) have the form of successor state axioms. When actions are deterministic, (SSA_1) is exactly the syntactic counterpart of the successor state axiom of [17].

In our framework belief-contravening information can be restricted to learning that some action a has been executed. Inconsistency with the agent's beliefs means that the agent believes a to be inexecutable, and learns that a has occurred. We have shown that such a revision operation can be implemented by an updating operation enabling the execution of a . Our second axiom (SSA_2) is a new solution that does not resort to orderings of plausibility.

6.1 Regression.

When restricted to deterministic actions our axioms allow for regression. In the case of nondeterministic actions it is not clear how this could be done. An alternative is to use the famous modular completeness result due to Sahlqvist [15], which applies here almost immediately (because our axioms are of the required form). We thus

get for free soundness and completeness results, as well as a tableau algorithm. If the tableau algorithm terminates then we get a decision procedure for our logic. We are currently working on that, aiming at applying recent results on modal axioms of confluence and permutation (of which our SSA_1 and SSA_2 are instances).

6.2 Public actions

Almost all the approaches suppose that actions are public. It has been relaxed in [5], where drawbacks of the earlier solution in [12] are pointed out. The solution of [5] corresponds to our case where $perc(a, \lambda)$ holds.

6.3 Revision: the approach of Shapiro *et al.*

In [21], Shapiro *et col.* add to the Scherl and Levesque framework a revision-like operation based on plausibility orderings. They define $BelA$ as truth of A in the most plausible among the possible situations. If a sensing action eliminates the most plausible of the possible situations, then previously less plausible situations become the most plausible ones. The plausibility ordering should be kept fixed.

While being intuitively appealing, their solution has several drawbacks. (1) As the authors note, it is restricted to deterministic actions. (2) “The specification of [the plausibility ordering] over the initial situation is the responsibility of the axiomatizer of the domain.” [21] This is particularly demanding because (3) in order to guarantee that after a the set of possible situations is nonempty, the authors require the set of possible situations to contain enough situations initially, restricting thus the agent’s ‘doxastic freedom’. (4) As pointed out in [5], such a solution to the problem of revision might endanger the solution to the frame problem. It seems to be fair to say that specifying a satisfactory plausibility ordering is a delicate task, involving a lot of imponderabilities in what concerns the relative plausibility of independent propositions. (5) The approach is unsatisfactory when applied to communication. Consider the following example: agent k is competent at p , and j is not. Agent i is completely ignorant initially: hence all possible situations are equally plausible for i . Then (under adequate hypotheses of cooperation) we can expect that when j asserts p , then i adopts p , i.e. $After_{asserts(j,p)} Bel p$. Moreover, as all situations were equally plausible, p holds in every situation possible for i . Therefore when subsequently k asserts $\neg p$, i will unavoidably move to an empty set of possible situations. (6) Action occurrences are supposed to be perceived correctly and completely (and the agent is aware of that). Therefore wrong beliefs can only come from the initial situation, and the doxastic concept of [21] turns out to be quite close to knowledge.

6.4 Segerberg

The approach of Segerberg [19, 18] is similar in spirit to ours. He has a successor state axiom for expansion [19, axiom #12], but no such account for revision.

6.5 The AGM postulates

The normative framework for belief revision being the AGM theory [1], which of their postulates do we satisfy? With a similar encoding as that of Shapiro *et col.* it can be shown that we satisfy the basic postulates (K^*1) – (K^*4), and (K^*6). (The names of the postulates are as in [21]). If we define update actions as in [21] we satisfy the update postulates ($K\triangleright 1$), ($K\triangleright 2$), ($K\triangleright 4$), and ($K\triangleright 5$) just as there. If we

define updating by A as $enable_{observe(A)}$ then we moreover satisfy ($K\triangleright 3$).

REFERENCES

- [1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson, ‘On the logic of theory change: Partial meet contraction and revision functions’, *The Journal of Symbolic Logic*, **50**(2), (June 1985).
- [2] F. Bacchus, J.Y. Halpern, and H. Levesque, ‘Reasoning about noisy sensors in the situation calculus’, in *Proc. 14th Int. Joint Conf. on Artificial Intelligence (IJCAI’95)*, pp. 1933–1940, (1995).
- [3] Marcos A. Castilho, Olivier Gasquet, and Andreas Herzig, ‘Formalizing action and change in modal logic I: the frame problem’, *Journal of Logic and Computation*, **9**(5), 701–735, (1999).
- [4] Silvia Coradeschi, ‘Reasoning with misperception in the Features and Fluents Framework’, in *Proc. 12th European Conf. on Artificial Intelligence (ECAI-96)*, pp. 657–661, (1996).
- [5] Robert Demolombe and Maria del Pilar Pozos Parra, ‘Formalisation de l’évolution de croyances dans le Calcul des Situations’, in *Proc. of Modèles Formels de l’Interaction (MFI’01)*, eds., B. Chaib-draa and P. Enjalbert, volume 2, pp. 205–217, (2001).
- [6] Gerbrandy, *Bisimulations on Planet Kripke*, Ph.D. dissertation, University of Amsterdam, 1999.
- [7] E. Giunchiglia, G. N. Kartha, and V. Lifschitz, ‘Representing action: indeterminacy and ramifications’, *AI Journal*, **95**, (1997).
- [8] David Harel, ‘Dynamic logic’, in *Handbook of Philosophical Logic*, eds., D. Gabbay and F. Guenther, volume II, D. Reidel Publishing Company, (1984).
- [9] Andreas Herzig, Jérôme Lang, and Thomas Polacsek, ‘A modal logic for epistemic tests’, in *Proc. Eur. Conf. on Artificial Intelligence (ECAI’2000)*, Berlin, (August 2000).
- [10] Andreas Herzig, Jérôme Lang, Dominique Longin, and Thomas Polacsek, ‘A logic for planning under partial observability’, in *Proc. Seventeenth National Conf. on Artificial Intelligence (AAAI-2000)*, (2000).
- [11] Hirofumi Katsuno and Alberto O. Mendelzon, ‘Propositional knowledge base revision and minimal change’, Technical Report KRR-TR-90-3, Dep. of Computer Science, University of Toronto, (March 1990).
- [12] Yves Lespérance, Hector J. Levesque, and Raymond Reiter, ‘A Situation Calculus approach to modeling and programming agents’, in *Foundations and theories of Rational Agents*, eds., A. Rao and M. Wooldridge. Kluwer Academic Publishers, (1999).
- [13] Hector J. Levesque, ‘What is planning in the presence of sensing?’, in *Proc. of the 13th Nat. Conf. on Artificial Intelligence (AAAI-96)*, Portland, Oregon, (August 1996).
- [14] Robert C. Moore, ‘A formal theory of knowledge and action’, in *Formal Theories of the Commonsense World*, eds., J.R. Hobbs and R.C. Moore, 319–358, Ablex, Norwood, NJ, (1985).
- [15] H. Sahlqvist, ‘Completeness and correspondence in the first and second order semantics for modal logics’, in *Proc. 3rd Scandinavian Logic Symposium*, ed., S. Kanger, volume 82 of *Studies in Logic*, (1975).
- [16] E. Sandewall, ‘The range of applicability of some nonmonotonic logics for strict inertia’, *J. of Logic and Computation*, **4**(5), 581–615, (1994).
- [17] Richard Scherl and Hector J. Levesque, ‘The frame problem and knowledge producing actions’, in *Proc. Nat. Conf. on AI (AAAI’93)*, pp. 689–695. AAAI Press, (1993).
- [18] Krister Segerberg, ‘Two traditions in the logic of belief: bringing them together’. to appear.
- [19] Krister Segerberg, ‘Belief revision from the point of view of doxastic logic’, *Bulletin of the IGPL*, **3**, 534–553, (1995).
- [20] S. Shapiro, Yves Lespérance, and Hector J. Levesque, ‘Specifying communicative multi-agent systems’, in *Agents and Multi-Agent Systems - Formalisms, Methodologies, and Applications*, eds., W. Wobcke, M. Pagnucco, and C. Zhang, pp. 1–14. Springer-Verlag, LNAI 1441, (1998).
- [21] S. Shapiro, M. Pagnucco, Y. Lespérance, and H. J. Levesque, ‘Iterated belief change in the situation calculus’, in *Proc. Seventh Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2000)*, pp. 527–538, (2000).
- [22] Michael Thielscher, ‘Computing ramifications by postprocessing’, in *Proc. 14th Int. Joint Conf. on Artificial Intelligence (IJCAI’95)*, pp. 1994–2000, Montreal, Canada, (1995).
- [23] Michael Thielscher, ‘Representing the knowledge of a robot’, in *Proc. KR’00*, eds., A. Cohn, F. Giunchiglia, and B. Selman, pp. 109–120. Morgan Kaufmann, (2000).