

Belief Revision and Progression of Knowledge Bases in the Epistemic Situation Calculus

Christoph Schwering
RWTH Aachen University
Aachen, Germany
schwering@kbsg.rwth-aachen.de

Gerhard Lakemeyer
RWTH Aachen University
Aachen, Germany
gerhard@kbsg.rwth-aachen.de

Maurice Pagnucco
University of New South Wales
Sydney, Australia
morri@cse.unsw.edu.au

Abstract

Fundamental to reasoning about actions and beliefs is the *projection problem*: to decide what is believed after a sequence of actions is performed. *Progression* is one widely applied technique to solve this problem. In this paper we propose a novel framework for computing progression in the epistemic situation calculus. In particular, we model an agent’s preferential belief structure using conditional statements and provide a technique for updating these conditional statements as actions are performed and sensing information is received. Moreover, we show, by using the concepts of *natural revision* and *only-believing*, that the progression of a conditional knowledge base can be represented by only-believing the revised set of conditional statements. These results lay the foundations for feasible belief progression due to the *unique-model property* of only-believing.

1 Introduction

Fundamental to reasoning about actions and beliefs is the *projection problem*: to decide what is believed after a sequence of actions is performed. There are two popular ways to solve this problem: *regression* rewrites a query about the future to a query about the initial situation only; *progression* changes the knowledge base to reflect the effects of the actions. Regression usually becomes infeasible when dealing with very long action sequences. A long-lived system— for instance, a domestic service robot— hence must progress its mental state once in a while. In particular, such a robot may continually acquire new information about its environment, which may or may not be consistent with what the robot believed or had learned before. When progressing its knowledge base, the robot needs to carefully *revise its beliefs* to handle these potentially conflicting pieces of information.

The following running example will illustrate our approach. Suppose our robot is carrying an object which it believes to be quite robust but not made of metal. Hence, when the robot drops the item, it believes the object is still intact. When a clinking noise occurs afterwards, perhaps indicating that the object is broken or is made of metal, this may change; as the robot considers fragility more plausible than the object

being metallic, it now believes that the object is broken. If then the robot inspects the object and it turns out to be fine after all, the previous belief is given up again and the robot assumes the object neither broken nor metallic (implicitly assuming the clink was due to something else).

Progression is a very intuitive way to implement projection. It has attracted a lot of attention in the reasoning about action community, perhaps most notably in the seminal work by Lin and Reiter [Lin and Reiter, 1997] in the situation calculus [McCarthy, 1963; Reiter, 2001]. While Lin–Reiter progression has been transferred to one of the epistemic extensions of the situation calculus [Lakemeyer and Levesque, 2009], it has not yet been studied in any of the situation calculus dialects that deal with belief change. In fact, most variants of the situation calculus take the Scherl–Levesque view of sensing [Scherl and Levesque, 2003], where sensing is assumed to be correct and thus cannot be revised at all.

In this paper, we propose a logical framework where the robot’s preferential belief structure is modelled using counterfactual conditionals. Sensing tells the agent new information, which may turn out wrong later and then becomes subject to belief revision. We propose a solution to the projection problem in this setting by belief progression. More precisely, we show that, by using the concepts of *natural revision* [Boutilier, 1993] and *only-believing* [Schwering and Lakemeyer, 2014], the progression of a conditional knowledge base can be represented by only-believing the revised set of conditional statements. The connection to only-believing is particularly attractive due to its *unique-model property*, which lays the foundations for feasible belief progression.

The next section discusses approaches related to our proposal here. In Section 3 we introduce our new logic, before we define Reiter’s concept of basic action theories in this framework and examine our running example in Section 4. Section 5 presents the main result of this paper: *how to revise and progress a conditional knowledge base*. In Section 6 we compare our work with standard belief revision frameworks before concluding.

2 Related Work

The situation calculus is perhaps the most thoroughly studied action formalism, although there are other significant approaches such as the event calculus [Kowalski and Sergot,

1989], the fluent calculus [Thielscher, 1999], and the family of action languages \mathcal{A} [Gelfond and Lifschitz, 1993]. A number of belief revision extensions of the situation calculus have been proposed [Shapiro *et al.*, 2011; Demolombe and Pozos Parra, 2006; Delgrande and Levesque, 2012; Fang and Liu, 2013; Schwering and Lakemeyer, 2014]. Most of these do not address the issue of faulty sensors with the exception of [Delgrande and Levesque, 2012; Fang and Liu, 2013], who both achieve this through plausibility updating schemes. However, these formalisms are quite heavyweight and leave open the projection problem which we claim is the key to implementation. The only available solution to the belief projection problem is by regression [Schwering and Lakemeyer, 2015], but it assumes correct sensors. Another framework to deal with faulty sensors is the Bayesian approach by [Bacchus *et al.*, 1999].

Belief revision has also been addressed in dynamic epistemic logic [van Benthem, 2007], where revised beliefs are reduced to initial beliefs in a regression-like fashion.

We use Boutilier’s natural revision [Boutilier, 1993]. While several plausibility updating schemes have been proposed, for example [Spohn, 1988; Nayak *et al.*, 2003], and despite legitimate criticism [Booth and Meyer, 2006], we choose natural revision because, as we shall see later, it agrees well with only-believing. Only-believing [Schwering and Lakemeyer, 2014] determines a unique epistemic model for a conditional knowledge base. It is related to only-knowing [Levesque and Lakemeyer, 2001] and System Z [Pearl, 1990].

Progression was first studied by Lin and Reiter in the situation calculus [Lin and Reiter, 1997]. Roughly speaking, their idea is to progress a knowledge base by replacing all relevant predicates with existentially quantified second-order variables. Their purpose is to “memorize” what was true before the progression. The original predicates are then re-introduced and equated with a formula about the second-order variables only. We will re-visit this idea in Section 5.

3 The Logic

In this section we introduce a novel logic that combines reasoning about action and belief revision. This first-order modal language is a variant of \mathcal{ES} [Lakemeyer and Levesque, 2011]. Actions may lead to a revision of the agent’s beliefs, which follows the rules of natural revision [Boutilier, 1993]. It also integrates an operator, called only-believing [Schwering and Lakemeyer, 2014], which uniquely determines the agent’s beliefs for a given conditional knowledge base.

3.1 The Language

The language consists of formulas over *fluent* or *rigid predicates* and *rigid terms*. The truth value of fluents may vary as the result of actions, but rigid do not.

The set of *terms* is the least set which contains infinitely many first-order variables and is closed under the application of infinitely many function symbols of any arity.

The set of well-formed *formulas* is the least set that contains $H(t_1, \dots, t_k)$, $(t_1 = t_2)$, $\neg\alpha$, $(\alpha \wedge \beta)$, $\forall x.\alpha$, $[t]\alpha$, $\Box\alpha$, $\mathbf{B}\alpha$, and $\mathbf{O}(\alpha, \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\})$, where H is a k -ary predicate symbol, t_i and t are terms, x is a variable, and

$\alpha, \beta, \phi_i, \psi_i$ are formulas. TRUE, FALSE, $(\alpha \vee \beta)$, $(\alpha \supset \beta)$, $(\alpha \equiv \beta)$, and $\exists x.\alpha$ are the usual abbreviations.

We read $[t]\alpha$ as “ α holds after action t ” and $\Box\alpha$ as “ α holds after any sequence of actions.” $\mathbf{B}\alpha$ is read as “ α is believed.” Conditionals $\phi \Rightarrow \psi$ are understood counterfactually [Lewis, 1973]: “if ϕ was true, then ψ would be true.” The only-believing operator $\mathbf{O}(\alpha, \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\})$ means that α is known, the conditionals $\phi_i \Rightarrow \psi_i$ are believed, and this is *all* that is known or believed. The purpose of only-believing is to uniquely determine a belief structure.

There are two special fluent predicates: $Poss(t)$ represents the precondition of action t ; $IF(t)$ represents the new information learned by the agent through action t .

By α_t^x we mean the formula α with t substituted for all free occurrences of x . We sometimes write \bar{t} for t_1, \dots, t_k .

A formula with no $[t]$ or \Box is called *static*. A formula with no \mathbf{B} or \mathbf{O} is called *objective*. A formula with no free variable is called a *sentence*.

3.2 The Semantics

We now give a possible-worlds semantics for this language. An interpretation of a sentence α consists of an epistemic state f , a world w , and a sequence of executed actions z . We write $f, w, z \models \alpha$ to say that the interpretation satisfies the sentence. We take as fixed in the domain of discourse the set of all ground terms denoted by R . Since R is countable, we can handle quantification by substitution. By R^* we denote the set of all sequences of ground terms, including the empty sequence $\langle \rangle$. The action sequence z initially starts with $\langle \rangle$ and then grows deterministically with each executed action r to $z \cdot r$. A *world* w maps all atomic sentences $H(r_1, \dots, r_k)$ and action sequences $z \in R^*$ to truth values $\{0, 1\}$, and satisfies the rigidity constraint: if H is rigid, then $w[H(r_1, \dots, r_k), z] = w[H(r_1, \dots, r_k), z']$ for all $z, z' \in R^*$. An *epistemic state* f maps each plausibility (taken from \mathbb{N}) to a set of worlds considered possible at this plausibility level. A value of 0 indicates the highest possible plausibility. A world may occur at multiple plausibility levels; usually we will indeed have $f(0) \subseteq f(1) \subseteq \dots$, that is, $f(p)$ contains all worlds at least as plausible as p .

We begin with the objective semantics:

1. $f, w, z \models H(r_1, \dots, r_k)$ iff $w[H(r_1, \dots, r_k), z] = 1$;
2. $f, w, z \models (r_1 = r_2)$ iff r_1 and r_2 are identical;
3. $f, w, z \models \neg\alpha$ iff $f, w, z \not\models \alpha$;
4. $f, w, z \models (\alpha \wedge \beta)$ iff $f, w, z \models \alpha$ and $f, w, z \models \beta$;
5. $f, w, z \models \forall x.\alpha$ iff $f, w, z \models \alpha_r^x$ for all $r \in R$;
6. $f, w, z \models [r]\alpha$ iff $f, w, z \cdot r \models \alpha$;
7. $f, w, z \models \Box\alpha$ iff $f, w, z \cdot z' \models \alpha$ for all $z' \in R^*$.

Before we proceed with the semantics of beliefs, we need to formalize the revision of an epistemic state. When the agent is informed that α holds, we promote the most plausible α -worlds to the highest plausibility level and shift all other worlds down by one level. This revision scheme is known as *natural revision* [Boutilier, 1993]. Intuitively, it is the minimal change required of the plausibility ordering to ensure belief in α . We denote the result of this revision by $f * \alpha$:

Definition 1 Given an epistemic state f and a sentence α , let $p^* = \min\{p \mid f, w, \langle \rangle \models \alpha \text{ for some } w \in f(p)\} \cup \{\infty\}$ be the first plausibility level¹ consistent with α and let $W = \{w \mid w \in f(p^*) \text{ and } f, w, \langle \rangle \models \alpha\}$ be the α -worlds from that level. For convenience, we let $f(-1)$ stand for $\{\}$. Then the *revision of f by α* is denoted by $f * \alpha$ and is defined as:

- if $p^* = \infty$ then $(f * \alpha)(p) = \{\}$ for all $p \in \mathbb{N}$;
- if $W \neq f(p^*) \setminus f(p^* - 1)$ then
 - $(f * \alpha)(p) = f(p - 1) \cup W$ for all $0 \leq p \leq p^*$;
 - $(f * \alpha)(p) = f(p - 1)$ for all $p > p^*$;
- otherwise
 - $(f * \alpha)(p) = f(p - 1) \cup W$ for all $0 \leq p \leq p^*$;
 - $(f * \alpha)(p) = f(p)$ for all $p > p^*$.

The plausibility of a world w is the minimal p such that $w \in f(p)$. In $f * \alpha$, the most plausible α -worlds from f are shifted to the first plausibility level. All other worlds are made less plausible by one level. The second and third case only differ in that the latter skips $f(p^*)$ to avoid $(f * \alpha)(p^*) = (f * \alpha)(p^* + 1)$ when $f(p^* - 1) \cup W = f(p^*)$.

Any action r provides the agent with the (perhaps vacuously true) information that $IF(r)$ holds. We therefore account for r in the epistemic state by revising by $IF(r)$ and then applying the effects of r to all worlds in f :

Definition 2 The *progression of a world w by z* is a world w_z such that $w_z[\rho, z'] = w[\rho, z \cdot z']$ for all atomic sentences ρ and action sequences z' . The *progression of an epistemic state f* is denoted by f_z and is defined inductively by

- $f_{\langle \rangle} = f$;
- $f_{z \cdot r}(p) = \{w_r \mid w \in (f_z * IF(r))(p)\}$ for all $p \in \mathbb{N}$.

With these definitions in hand, we are ready to proceed with the epistemic semantics:

8. $f, w, z \models \mathbf{B}\alpha$ iff $f_z, w', \langle \rangle \models \alpha$ for all $w' \in f_z(0)$;
9. $f, w, z \models \mathbf{O}(\alpha, \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\})$ iff for some $p_1, \dots, p_m \in \mathbb{N} \cup \{\infty\}$,¹ for all $p \in \mathbb{N}$,
 - (a) $f_z, w', \langle \rangle \models (\alpha \wedge \bigwedge_{i:p_i \geq p} (\phi_i \supset \psi_i))$ iff $w' \in f_z(p)$;
 - (b) for all $p_i > p$, for all $w' \in f_z(p)$, $f_z, w', \langle \rangle \not\models \phi_i$;
 - (c) for all $p_i = p$, for some $w' \in f_z(p)$, $f_z, w', \langle \rangle \models \phi_i$.

In the following, we sometimes omit f or w in $f, w, z \models \alpha$ when it is irrelevant to the truth of α . We also may omit z when $z = \langle \rangle$. A set of sentences Σ *entails* a sentence α iff for all f and w , if $f, w \models \beta$ for all $\beta \in \Sigma$, then $f, w \models \alpha$. We write $\Sigma \models \alpha$, and abbreviate $\models \alpha$ when $\Sigma = \{\}$.

3.3 Some Properties

Since $f_{\langle \rangle} = f$, our definition of only-believing is identical with the one in [Schwering and Lakemeyer, 2014] when $z = \langle \rangle$. Therefore the following theorem carries over to our logic:

Theorem 3 ([Schwering and Lakemeyer, 2014]) *Let $\Gamma = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ and let α, ϕ_i, ψ_i be objective. Then there is a unique f such that $f \models \mathbf{O}(\alpha, \Gamma)$.*

¹In Definition 1, $p^* = \infty$ means that all plausibility levels are inconsistent with α . In Rule 9, $p_i = \infty$ analogously indicates that all plausibility levels are inconsistent with the antecedent ϕ_i .

In fact, the following straightforward procedure generates the epistemic state that satisfies $\mathbf{O}(\alpha, \Gamma)$ [Schwering and Lakemeyer, 2014]. Initially let $p_1 := 0, \dots, p_m := 0$. Then let p run from 0 to m and repeat the following two steps:

- Let $f(p) := \{w \mid w \models (\alpha \wedge \bigwedge_{i:p_i \geq p} (\phi_i \supset \psi_i))\}$.
- For all i , if there is no $w \in f(p)$ such that $w \models \phi_i$, let $p_i := p + 1$.

Then let $f(p) := f(m)$ for all $p > m$. Finally let $p_i := \infty$ for all $p_i > m$. Then $f \models \mathbf{O}(\alpha, \Gamma)$ for the plausibilities p_1, \dots, p_m . Observe that then $f(0) \subseteq f(1) \subseteq \dots$ holds.

Since our objective semantics is the same as the one for \mathcal{ES} [Lakemeyer and Levesque, 2004], its theorems carry over to our logic:

Theorem 4 *Let $\models_{\mathcal{ES}}$ denote the entailment relation of \mathcal{ES} . For any objective sentence α , $\models \alpha$ iff $\models_{\mathcal{ES}} \alpha$.*

This correspondence does not hold for knowledge or belief because our notion of informing differs from the sensing concept predominant in the situation calculus: \mathcal{ES} and its descendants, including the belief revision variant [Schwering and Lakemeyer, 2014], follow Scherl and Levesque [Scherl and Levesque, 2003] and define sensing to be *always correct*—a strong assumption we do not make here. We hence resort to the weaker concept of *informing* where new information may contradict older information.

4 Basic Action Theories

To axiomatize a dynamic domain we use the modal variant of Reiter's *basic action theories* [Reiter, 2001; Lakemeyer and Levesque, 2011]. A basic action theory over a finite set of fluent predicates \mathcal{F} consists of a static and a dynamic part. In the context of a basic action theory a formula is called *fluent* when it is objective, static, and all predicates are either from \mathcal{F} or rigid.

The dynamic axioms express when an action is executable (Σ_{pre}), how actions change the truth values of fluents (Σ_{post}), and which belief actions produce (Σ_{info}):²

- Σ_{pre} contains a single sentence $\Box Poss(a) \equiv \pi$ where π is a fluent formula;
- Σ_{post} contains a sentence $\Box[a]F(\vec{x}) \equiv \gamma_F$ for all $F \in \mathcal{F}$ where γ_F is a fluent formula;
- Σ_{info} contains a single sentence $\Box IF(a) \equiv \varphi$ where φ is a fluent formula.

The sentences in Σ_{post} are called *successor state axioms* because they relate the state after an action a to the one before a . They incorporate Reiter's solution to the *frame problem* [Reiter, 2001]. The *informed fluent* axiom Σ_{info} is to axiomatize the information an action tells the agent. We refer to the dynamic axioms as Σ_{dyn} .

The static part of a basic action theory expresses what the agent believes to be true: Σ_{bel} contains finitely many belief conditionals $\phi \Rightarrow \psi$ where ϕ and ψ are fluent sentences.

²We assume \Box has lower and $[t]$ has higher precedence than logical connectives and that all first-order variables are quantified from outside. So $\Box[a]F(\vec{x}) \equiv \gamma_F$ stands for $\forall a. \forall \vec{x}. \Box((a)F(\vec{x}) \equiv \gamma_F)$.

The *projection problem* in this setting is to decide if $\mathbf{O}(\Sigma_{dyn}, \Sigma_{bel}) \models \alpha$ holds,³ where α may involve actions and/or beliefs. In the next section we present a solution to the projection problem by *progression*, which modifies the initial beliefs to take into account the actions' effects.

Example

The example from Section 1 can be modelled as a basic action theory as follows. There are two rigid predicates, F and M , for the object being fragile or metallic, respectively. There is one fluent predicate, B , which indicates whether or not the object is broken. The action drop ⁴ causes the object to break if it is fragile. The clinking noise is represented by the action clink , which informs that the object is broken or metallic ($B \vee M$). Lastly, the inspect action tells us that the object is not broken ($\neg B$). We do not model any preconditions for simplicity. This translates to the following dynamic axioms:

$$\begin{aligned}\Sigma_{pre} &= \{\Box Poss(a) \equiv \text{TRUE}\}; \\ \Sigma_{post} &= \{\Box B \equiv a = \text{drop} \wedge F \vee B\}; \\ \Sigma_{info} &= \{\Box IF(a) \equiv (a = \text{clink} \supset B \vee M) \wedge \\ &\quad (a = \text{inspect} \supset \neg B)\}.\end{aligned}$$

The robot believes that the object is neither fragile nor metal, and it generally considers it more likely that the object is fragile than being metal. Furthermore we are absolutely certain that the object is not broken in the beginning. Thus we have:

$$\begin{aligned}\Sigma_{bel} &= \{\text{TRUE} \Rightarrow \neg F \wedge \neg M, \\ &\quad F \vee M \Rightarrow F \wedge \neg M, \\ &\quad B \Rightarrow \text{FALSE}\}.\end{aligned}$$

By the construction of Theorem 3, $f \models \mathbf{O}(\Sigma_{dyn}, \Sigma_{bel})$ iff

$$\begin{aligned}f(0) &= \{w \mid w \models \Sigma_{dyn} \wedge \neg B \wedge \neg F \wedge \neg M\}; \\ f(1) &= \{w \mid w \models \Sigma_{dyn} \wedge \neg B \wedge ((\neg F \wedge \neg M) \vee (F \wedge \neg M))\} \\ &= \{w \mid w \models \Sigma_{dyn} \wedge \neg B \wedge \neg M\}; \\ f(p) &= \{w \mid w \models \Sigma_{dyn} \wedge \neg B\} \quad \text{for all } p \geq 2.\end{aligned}$$

Notice that the effect of $B \Rightarrow \text{FALSE}$ is to assert $\neg B$ at all plausibility levels. We will now examine how belief changes after dropping the object, hearing a clink, and inspecting the object.

After dropping the object, we believe it to be still intact, that is, $\mathbf{O}(\Sigma_{dyn}, \Sigma_{bel}) \models [\text{drop}]\mathbf{B}(\neg B \wedge \neg F \wedge \neg M)$. This is because drop triggers no revision and hence:

$$\begin{aligned}f_{\text{drop}}(0) &= \{w \mid w \models \Sigma_{dyn} \wedge \neg B \wedge \neg F \wedge \neg M\}; \\ f_{\text{drop}}(1) &= \{w \mid w \models \Sigma_{dyn} \wedge (B \equiv F) \wedge \neg M\}; \\ f_{\text{drop}}(p) &= \{w \mid w \models \Sigma_{dyn} \wedge (B \equiv F)\} \quad \text{for all } p \geq 2.\end{aligned}$$

When we hear a clink after dropping the object, the revision by $IF(\text{clink})$ promotes the most plausible ($B \vee M$)-worlds to the first plausibility level. The first ($B \vee M$)-worlds come from $f_{\text{drop}}(1)$, namely those $w \in f_{\text{drop}}(1)$ with $w \models F$. Hence we have:

$$\begin{aligned}f_{\text{drop-clink}}(0) &= \{w \mid w \models \Sigma_{dyn} \wedge B \wedge F \wedge \neg M\}; \\ f_{\text{drop-clink}}(1) &= \{w \mid w \models \Sigma_{dyn} \wedge (B \equiv F) \wedge \neg M\}; \\ f_{\text{drop-clink}}(p) &= \{w \mid w \models \Sigma_{dyn} \wedge (B \equiv F)\} \quad \text{for all } p \geq 2.\end{aligned}$$

³We identify a finite set of sentences with their conjunction.

⁴We use sans-serif font for ground terms.

Since for all $w \in f_{\text{drop-clink}}(0)$, $w \models B \wedge F \wedge \neg M$, we have $\mathbf{O}(\Sigma_{dyn}, \Sigma_{bel}) \models [\text{drop}][\text{clink}]\mathbf{B}(B \wedge F \wedge \neg M)$.

When we now inspect the object, we revise by $IF(\text{inspect})$, which promotes the first $\neg B$ -worlds from $f_{\text{drop-clink}}$ to the first plausibility level. The first $\neg B$ -worlds come from $f_{\text{drop-clink}}(1)$, namely those $w \in f_{\text{drop-clink}}(1)$ with $w \models \neg F$. For $z = \text{drop} \cdot \text{clink} \cdot \text{inspect}$ we hence have:

$$\begin{aligned}f_z(0) &= \{w \mid w \models \Sigma_{dyn} \wedge \neg B \wedge \neg F \wedge \neg M\}; \\ f_z(1) &= \{w \mid w \models \Sigma_{dyn} \wedge (B \equiv F) \wedge \neg M\}; \\ f_z(p) &= \{w \mid w \models \Sigma_{dyn} \wedge (B \equiv F)\} \quad \text{for all } p \geq 2.\end{aligned}$$

Observe that the revision by $IF(\text{inspect})$ undoes the previous revision by $IF(\text{clink})$, that is, $f_z = f_{\text{drop}}$, and thus: $\mathbf{O}(\Sigma_{dyn}, \Sigma_{bel}) \models [\text{drop}][\text{clink}][\text{inspect}]\mathbf{B}(\neg B \wedge \neg F \wedge \neg M)$. In particular, we believe the object is not metallic, due to natural revision. In approaches where sensing is assumed to be correct, like [Shapiro *et al.*, 2011; Schwering and Lake-meyer, 2014], such a conclusion would not have been possible because they would have ruled out all worlds contradicting ($B \vee M$) and $\neg B$, so only M -worlds would be left.

5 Progression

In this section we present a form of progression which captures the beliefs after an action using only-believing. We need to take into account both the epistemic revision effect of the action and its physical effects. To this end, we first define revision of only-believing and show its correctness with respect to the semantics (Definition 5 and Theorem 6). Then we integrate the physical effects of actions to obtain a form of progression of basic action theories (Definition 7 and Theorem 8).

Definition 5 Let $\Gamma = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ and let $\alpha, \beta, \phi_i, \psi_i$ be objective. Let f be the epistemic state such that $f \models \mathbf{O}(\alpha, \Gamma)$ for plausibilities p_1, \dots, p_m . Let

$$\begin{aligned}p^* &= \min\{p \mid w \models \beta \text{ for some } w \in f(p)\} \cup \{\infty\}; \\ \gamma_p &= \begin{cases} \text{FALSE} & \text{if } p = -1; \\ \bigwedge_{i:p_i \geq p} (\phi_i \supset \psi_i) & \text{otherwise;} \end{cases} \\ \delta &= \begin{cases} \gamma_{p^*-1} \vee (\gamma_{p^*} \wedge \beta) & \text{if } \alpha \wedge \gamma_{p^*} \wedge \neg \gamma_{p^*-1} \not\models \beta; \\ \gamma_{p^*-2} \vee (\gamma_{p^*} \wedge \beta) & \text{if } \alpha \wedge \gamma_{p^*} \wedge \neg \gamma_{p^*-1} \models \beta \text{ and } p^* > 0; \\ \text{FALSE} & \text{if } \alpha \wedge \gamma_{p^*} \wedge \neg \gamma_{p^*-1} \models \beta \text{ and } p^* = 0. \end{cases}\end{aligned}$$

Then the *revision of $\mathbf{O}(\alpha, \Gamma)$ by β* is denoted by $\mathbf{O}(\alpha, \Gamma) * \beta$ and defined as $\mathbf{O}(\alpha, \Gamma')$ where

$$\begin{aligned}\Gamma' &= \{\text{TRUE} \Rightarrow \beta\} \cup \\ &\quad \{\phi_i \wedge \neg \beta \Rightarrow \psi_i \mid p_i < p^*\} \cup \\ &\quad \{\phi_i \Rightarrow \psi_i \mid p_i \geq p^*\} \cup \\ &\quad \{\neg \delta \Rightarrow \gamma_{p^*}\}.\end{aligned}$$

$\mathbf{O}(\alpha, \Gamma) * \beta$ thus means to believe β , but if β turns out to be wrong, we return to the old beliefs $\phi_i \wedge \neg \beta \Rightarrow \psi_i$. The conditional $\neg \delta \Rightarrow \gamma_{p^*}$ ensures that when also all $\phi_i \wedge \neg \beta \Rightarrow \psi_i$ turn out to be wrong, we return to the same beliefs as before the revision. Intuitively, a world satisfies δ if before the revision its plausibility was at least $p^* - 1$ or if its plausibility was

p^* and it satisfies β . (The different definitions for δ handle cases where plausibility levels concur.) Therefore, $\neg\delta \Rightarrow \gamma_{p^*}$ means that if β and everything we considered at least as plausible as $p^* - 1$ before the revision turned out to be wrong, then *after* the revision we believe *the same as before* the revision if β and all beliefs at least as plausible as $p^* - 1$ had turned out to be wrong. Notice that $\mathbf{O}(\alpha, \Gamma) * \beta$ can be generated using first-order reasoning.

We now prove that this revision matches natural revision that is used in the semantics:

Theorem 6 *Let $f \models \mathbf{O}(\alpha, \Gamma)$. Then $f * \beta \models \mathbf{O}(\alpha, \Gamma) * \beta$.*

Proof Sketch. The proof is lengthy, so we sketch the main idea. The idea is to show that $f * \beta$ satisfies the right-hand side of Rule 9 for the following plausibilities. For $\text{TRUE} \Rightarrow \beta$, the plausibility is 0. For each $\phi_i \wedge \neg\beta \Rightarrow \psi_i$, the plausibility is $p_i + 1$. For each $\phi_i \Rightarrow \psi_i$, the plausibility is 0 if $p_i = p^*$ and $w \models \phi_i \wedge \beta$ for some $w \in f(p^*)$; it is $p_i + 1$ if $w \not\models \beta$ for some $w \in f(p^*) \setminus f(p^* - 1)$; and p_i otherwise. For $\neg\delta \Rightarrow \gamma_{p^*}$, the plausibility is $p^* + 1$ if $\gamma_{p^*} \wedge \neg\gamma_{p^* - 1} \not\models \beta$; p^* otherwise. It is then tedious but straightforward to show that Rule 9 is satisfied. Crucial for this is that $\models (\neg\delta \supset \gamma_{p^*}) \equiv \gamma_{p^*}$, which allows to rearrange the conditionals from level p^* because $\neg\delta \Rightarrow \gamma_{p^*}$ takes their place. \square

We are now ready to define the progression of a basic action theory with conditional beliefs $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})$. Given an action r , we first revise the theory by $IF(r)$ and then handle the effects of r on the fluents. The revision is captured by $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}}) * \varphi_r^a$ where φ is the informed fluent axiom ($\square IF(a) \equiv \varphi \in \Sigma_{\text{dyn}}$). (The reason for taking φ_r^a instead of $IF(r)$ is to keep the belief conditionals fluent.) In the following we show how to handle the physical effect of r . Let the set of fluents be $\mathcal{F} = \{F_1, \dots, F_n\}$ and let $\mathcal{P} = \{P_1, \dots, P_n\}$ be rigid predicates of corresponding arity which do not otherwise occur in $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})$. We denote by $\alpha_{\vec{P}}^{\vec{F}}$ the formula obtained from replacing each F_i with P_i . The progression of a basic action theory is then defined as follows:

Definition 7 Let Σ'_{bel} be the revised belief conditionals, that is, $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}}) * \varphi_r^a = \mathbf{O}(\Sigma_{\text{dyn}}, \Sigma'_{\text{bel}})$. The *progression of $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})$ by r* is then denoted by $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})_r$ and is defined as $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma''_{\text{bel}})$ where

$$\Sigma''_{\text{bel}} = \Sigma'_{\text{bel}}^{\vec{F}} \cup \{ \neg(\forall \vec{x}. F(\vec{x}) \equiv \gamma_{F_r^a}^{\vec{F}}) \Rightarrow \text{FALSE} \mid F \in \mathcal{F} \}.$$

The intuition behind the definition is as follows. Each new predicate P_i captures the pre- r truth value of F_i . Now, when an action r is executed, we first revise by the information φ_r^a produced by r which gives us the new beliefs Σ'_{bel} . The beliefs $\Sigma'_{\text{bel}}^{\vec{F}}$ represent the same belief structure as Σ'_{bel} , except that each F_i is renamed to P_i . Adding a conditional $\neg(\forall \vec{x}. F(\vec{x}) \equiv \gamma_{F_r^a}^{\vec{F}}) \Rightarrow \text{FALSE}$ finally has the effect of requiring $\forall \vec{x}. F(\vec{x}) \equiv \gamma_{F_r^a}^{\vec{F}}$ at every plausibility level, which leads to F taking the correct post- r value. Note that the revision of a basic action theory again is a basic action theory.

We say α is \mathcal{P} -free if no predicate in α is from \mathcal{P} . The following theorem establishes the correctness of progression:

Theorem 8 *Let α be \mathcal{P} -free and without \mathbf{O} .*

Then $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}}) \models [r]\mathbf{B}\alpha$ iff $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})_r \models \mathbf{B}\alpha$.

Proof Sketch. The proof proceeds in two steps. Firstly, it is shown that $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma'_{\text{bel}})$ and $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma''_{\text{bel}})$ (as defined in Definition 7) lead to the same belief structure. Intuitively, this is because in $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma''_{\text{bel}})$ each F_i is renamed by P_i , so the additional conditionals in Σ''_{bel} do not affect the plausibilities. Secondly, if $f \models \mathbf{O}(\Sigma_{\text{dyn}}, \Sigma'_{\text{bel}})$ and $g \models \mathbf{O}(\Sigma_{\text{dyn}}, \Sigma''_{\text{bel}})$, then f and g are bisimilar in the sense that for all p , for each $w \in f(p)$, there is some $w' \in g(p)$ such that w_r and w' agree on all truth values except for the predicates from \mathcal{P} , and vice-versa. Then one can show by induction that any two states bisimilar in this sense satisfy the same \mathcal{P} -free sentences. \square

Our definition of progression is closely related to Lin and Reiter's progression [Lin and Reiter, 1997]. While they use existentially quantified second-order variables to memorize the pre- r truth value of each F_i , we use new rigid Skolem-predicates P_i . This is weaker than Lin-Reiter progression in the sense that $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}}) \not\models [r]\mathbf{O}(\Sigma_{\text{bel}}, \Sigma_{\text{dyn}})_r$. However, second-order logic would have led to a considerably more complex definition of the \mathbf{O} operator.

Example

We now examine the progression of $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})$ by drop and then by clink. The epistemic states we obtain match the ones from Section 4 except for the newly added predicate P_B to memorize the old value of B .

We first consider the progression by drop. As the revision by φ_{drop}^a just adds another conditional to $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}}) * \varphi_{\text{drop}}^a$ whose antecedent and consequent both are equivalent to TRUE , we proceed with the progression of the physical effects. By Definition 7, the resulting conditionals are:

$$\begin{aligned} \Sigma'_{\text{bel}} = \{ & \text{TRUE} \Rightarrow \neg F \wedge \neg M, \\ & F \vee M \Rightarrow F \wedge \neg M, \\ & P_B \Rightarrow \text{FALSE}, \\ & \neg(B \equiv \text{drop} = \text{drop} \wedge F \vee P_B) \Rightarrow \text{FALSE} \}. \end{aligned}$$

Then $g \models \mathbf{O}(\Sigma_{\text{dyn}}, \Sigma'_{\text{bel}})$ iff $g(p) = \{w \mid w \models \kappa_p\}$ where

$$\begin{aligned} \kappa_0 &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge (B \equiv F \vee P_B) \wedge \neg F \wedge \neg M \\ &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge \neg B \wedge \neg F \wedge \neg M; \\ \kappa_1 &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge (B \equiv F) \wedge \neg M; \\ \kappa_p &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge (B \equiv F) \quad \text{for all } p \geq 2. \end{aligned}$$

Observe that the only difference between f_{drop} from Section 4 and g is the additional restrictions on P_B , and hence both epistemic states satisfy the same formulas without P_B .

Let us now consider $\mathbf{O}(\Sigma_{\text{dyn}}, \Sigma'_{\text{bel}})_{\text{clink}}$, which leads to a revision by $(B \vee M)$. Since clink has no physical effect, we only examine this revision and omit the progression. The first plausibility level from g consistent with $(B \vee M)$ is $p^* = 1$. Since there is no $w \in g(1) \setminus g(0)$ such that $w \not\models (B \vee M)$, the additional constraint $\neg\delta \Rightarrow \gamma_{p^*}$ is such that:

$$\begin{aligned} \delta &= \neg\gamma_{p^* - 2} \vee (\gamma_{p^*} \wedge \beta) \\ &= \text{FALSE} \vee (\neg P_B \wedge (B \equiv F) \wedge \neg M \wedge (B \vee M)) \\ &= \neg P_B \wedge B \wedge F \wedge \neg M \quad \text{and} \\ \gamma_{p^*} &= \neg P_B \wedge (B \equiv F) \wedge \neg M. \end{aligned}$$

Hence we obtain as the revised set of beliefs:

$$\begin{aligned} \Sigma'_{bel} = \{ & \text{TRUE} \Rightarrow B \vee M, \\ & \neg(B \vee M) \Rightarrow \neg F \wedge \neg M, \\ & F \vee M \Rightarrow F \wedge \neg M, \\ & \neg\delta \Rightarrow \gamma_{p^*}, \\ & P_B \Rightarrow \text{FALSE}, \\ & \neg(B \equiv \text{drop} = \text{drop} \wedge F \vee P_B) \Rightarrow \text{FALSE} \}. \end{aligned}$$

This syntactic revision satisfies the same P_B -free formulas as $f_{\text{drop-clink}}$ from Section 4: we have $g' \models \mathbf{O}(\Sigma_{\text{dyn}}, \Sigma'_{\text{bel}})$ iff $g'(p) = \{w \mid w \models \lambda_p\}$ where:

$$\begin{aligned} \lambda_0 &= \Sigma_{\text{dyn}} \wedge (B \vee M) \wedge (\neg(B \vee M) \supset \neg F \wedge \neg M) \wedge \\ & \quad (F \vee M \supset F \wedge \neg M) \wedge (\neg\delta \supset \gamma_{p^*}) \wedge \\ & \quad \neg P_B \wedge (B \equiv F \vee P_B) \\ &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge B \wedge F \wedge \neg M; \\ \lambda_1 &= \Sigma_{\text{dyn}} \wedge (\neg(B \vee M) \supset \neg F \wedge \neg M) \wedge (\neg\delta \supset \gamma_{p^*}) \wedge \\ & \quad \neg P_B \wedge (B \equiv F \vee P_B) \\ &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge (B \equiv F) \wedge \neg M; \\ \lambda_p &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge (B \equiv F \vee P_B) \\ &= \Sigma_{\text{dyn}} \wedge \neg P_B \wedge (B \equiv F) \quad \text{for all } p \geq 2. \end{aligned}$$

Note that λ_0 is inconsistent with $\neg\delta$ but λ_1 is not, so $\neg\delta \Rightarrow \gamma_{p^*}$ is satisfied at $g'(1)$ and has the effect of asserting $\neg M$.

6 AGM, DP, and NPP Postulates

In this section we relate our framework to the most well known accounts of belief change: AGM [Alchourron *et al.*, 1985; Gärdenfors, 1988], DP [Darwiche and Pearl, 1997], and NPP [Nayak *et al.*, 2003]. We will see that all the AGM postulates and a slight modification of the DP postulates hold, whereas the NPP postulates are not satisfied.

An action r is called a *revision action* when it has no physical effect, that is, $\Sigma_{\text{dyn}} \models \Box[r]F(\vec{x}) \equiv F(\vec{x})$ for all $F \in \mathcal{F}$. The only effect of such a revision action r is the belief revision by $IF(r)$. Since it is equivalent, and to ease the presentation, we consider in the following just the revision instead of progression of \mathbf{O} by a revision action r .

For the rest of this section, let $\Sigma = \mathbf{O}(\Sigma_{\text{dyn}}, \Sigma_{\text{bel}})$ and let β, γ, δ be \mathcal{P} -free fluent sentences. In the following results, we have translated the relevant postulates into our formalism (similarly as in [Shapiro *et al.*, 2011]).

Theorem 9 *The AGM postulates hold:*

1. $\Sigma * \beta$ is deductively closed.
2. $\Sigma * \beta \models \mathbf{B}\beta$.
3. If $\Sigma * \beta \models \mathbf{B}\delta$, then $\Sigma \models \mathbf{B}(\beta \supset \delta)$.
4. If $\Sigma \not\models \mathbf{B}\neg\delta$ and $\Sigma \models \mathbf{B}(\beta \supset \delta)$, then $\Sigma * \beta \models \mathbf{B}\delta$.
5. If $\Sigma \not\models \mathbf{B}\text{FALSE}$ and $\Sigma \not\models \neg\beta$, then $\Sigma * \beta \not\models \mathbf{B}\text{FALSE}$.
6. If $\Sigma \models \beta \equiv \gamma$, then $\Sigma * \beta \equiv \Sigma * \gamma$.
7. If $\Sigma * (\beta \wedge \gamma) \models \mathbf{B}\delta$, then $\Sigma * \beta \models \mathbf{B}(\gamma \supset \delta)$.
8. If $\Sigma * \beta \not\models \mathbf{B}\neg\gamma$ and $\Sigma * \beta \models \mathbf{B}(\gamma \supset \delta)$, then $\Sigma * (\beta \wedge \gamma) \models \mathbf{B}\delta$.

Proof. The proofs are straightforward. Here we only show Postulate 8: suppose $f \models \Sigma$ and the antecedent holds. Then for some $w \in (f * \beta)(0)$, $w \models \gamma$, and for all $w \in (f * \beta)(0)$, $w \models \gamma \supset \delta$. Therefore $(f * \beta \wedge \gamma)(0) \subseteq (f * \beta)(0)$. Since for all $w \in (f * \beta \wedge \gamma)(0)$, $w \models \gamma$, by assumption $w \models \delta$. \square

While our main interest here is not in physical actions, it can be noted that the *KM update postulates* [Katsuno and Mendelzon, 1991] hold with the exception of Postulates 3, 6, and 7 for the reasons noted by [Shapiro *et al.*, 2011].

As for the DP postulates for iterated revision, we need to add the requirement $\beta \not\models \gamma$ to the antecedent of DP2 because we cannot recover from revision by an unsatisfiable formula.

Theorem 10 *The DP postulates hold with a restricted version of the second postulate:*

1. If $\gamma \models \beta$, then $(\Sigma * \beta) * \gamma \models \mathbf{B}\delta$ iff $\Sigma * \gamma \models \mathbf{B}\delta$.
2. If $\beta \models \neg\gamma$ and $\beta \not\models \gamma$, then $(\Sigma * \beta) * \gamma \models \mathbf{B}\delta$ iff $\Sigma * \gamma \models \mathbf{B}\delta$.
3. If $\Sigma * \gamma \models \mathbf{B}\beta$, then $(\Sigma * \beta) * \gamma \models \mathbf{B}\beta$.
4. If $\Sigma * \gamma \not\models \mathbf{B}\neg\beta$, then $(\Sigma * \beta) * \gamma \not\models \mathbf{B}\neg\beta$.

Proof. Again the proofs are reasonably straightforward, so we only show Postulate 3 here: suppose $f \models \Sigma$ and the antecedent holds. If for some $w \in (f * \beta)(0)$, $w \models \gamma$, then $\{w\} \neq ((f * \beta) * \gamma)(0) \subseteq (f * \beta)(0)$ and therefore for all $w \in ((f * \beta) * \gamma)(0)$, $w \models \beta$. Otherwise, for all w with $w \models \gamma$, $w \in f(p) \cup (f * \beta)(0)$ iff $w \in f(p)$. Therefore $(f * \gamma)(0) = ((f * \beta) * \gamma)(0)$. \square

The NPP postulates, however, are not satisfied. This is because natural revision is inconsistent with the third postulate:

3. If $\Sigma \not\models \neg(\beta \wedge \gamma)$ then $(\Sigma * \beta) * \gamma \models \mathbf{B}\delta$ iff $\Sigma * (\beta \wedge \gamma) \models \mathbf{B}\delta$.

Consider our running example: after revising by $(B \vee M)$ and then by $\neg B$, we believe that $\neg M$, whereas after revising by $(B \vee M) \wedge \neg B$ we would believe M .

7 Conclusion

We have developed a logic for reasoning about actions and belief revision. In particular, our approach is able to revise inconsistent sensing information by natural revision. We showed that this formalism is in line with the AGM and DP postulates, but not with the NPP postulates. Most importantly, however, we addressed the belief projection problem by progression: we showed that, if the agent only-believes a conditional knowledge base before an action, then they only-believe another conditional knowledge base after the action.

The next step is to employ this notion of progression in feasible subclasses of the situation calculus such as [Liu and Lakemeyer, 2009]. We then aim to integrate our work with an existing implementation of a limited reasoner about actions and knowledge based on [Lakemeyer and Levesque, 2014].

Acknowledgments

We thank the anonymous reviewers for helpful comments and pointers to related work. This work was partially supported by the DAAD Go8 project 56266625 and by the DFG research unit FOR 1513. The first author was supported by a DAAD doctoral scholarship.

References

- [Alchourron *et al.*, 1985] Carlos E. Alchourron, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [Bacchus *et al.*, 1999] Fahiem Bacchus, Joseph Y. Halpern, and Hector J. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1–2):171–208, 1999.
- [Booth and Meyer, 2006] Richard Booth and Thomas Meyer. Admissible and restrained revision. *Journal of Artificial Intelligence Research*, 26:127–151, 2006.
- [Boutilier, 1993] Craig Boutilier. Revision sequences and nested conditionals. In *Proc. IJCAI*, pages 519–525, 1993.
- [Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1):1–29, 1997.
- [Delgrande and Levesque, 2012] James P. Delgrande and Hector J. Levesque. Belief revision with sensing and fallible actions. In *Proc. KR*, pages 148–157, 2012.
- [Demolombe and Pozos Parra, 2006] Robert Demolombe and Maria del Pilar Pozos Parra. Belief revision in the situation calculus without plausibility levels. In *Foundations of Intelligent Systems*, LNCS, pages 504–513. 2006.
- [Fang and Liu, 2013] Liangda Fang and Yongmei Liu. Multiagent knowledge and belief change in the situation calculus. In *Proc. AAAI*, pages 304–312, 2013.
- [Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in flux: Modeling the dynamics of epistemic states*. The MIT press, 1988.
- [Gelfond and Lifschitz, 1993] Michael Gelfond and Vladimir Lifschitz. Representing action and change by logic programs. *The Journal of Logic Programming*, 17(2):301–321, 1993.
- [Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.
- [Kowalski and Sergot, 1989] Robert Kowalski and Marek Sergot. A logic-based calculus of events. In *Foundations of knowledge base management*, pages 23–55. 1989.
- [Lakemeyer and Levesque, 2004] Gerhard Lakemeyer and Hector J. Levesque. Situations, si! situation terms, no! In *Proc. KR*, pages 516–526, 2004.
- [Lakemeyer and Levesque, 2009] Gerhard Lakemeyer and Hector J. Levesque. A semantical account of progression in the presence of defaults. In *Proc. IJCAI*, pages 842–847, 2009.
- [Lakemeyer and Levesque, 2011] Gerhard Lakemeyer and Hector J. Levesque. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence*, 175(1):142–164, 2011.
- [Lakemeyer and Levesque, 2014] Gerhard Lakemeyer and Hector J. Levesque. Decidable reasoning in a fragment of the epistemic situation calculus. In *Proc. KR*, pages 468–477, 2014.
- [Levesque and Lakemeyer, 2001] Hector J. Levesque and Gerhard Lakemeyer. *The Logic of Knowledge Bases*. MIT Press, 2001.
- [Lewis, 1973] David Lewis. *Counterfactuals*. John Wiley & Sons, 1973.
- [Lin and Reiter, 1997] Fangzhen Lin and Ray Reiter. How to progress a database. *Artificial Intelligence*, 92(1):131–167, 1997.
- [Liu and Lakemeyer, 2009] Yongmei Liu and Gerhard Lakemeyer. On first-order definability and computability of progression for local-effect actions and beyond. In *Proc. IJCAI*, pages 860–866, 2009.
- [McCarthy, 1963] John McCarthy. Situations, Actions, and Causal Laws. Technical Report AI Memo 2, AI Lab, Stanford University, July 1963.
- [Nayak *et al.*, 2003] Abhaya C. Nayak, Maurice Pagnucco, and Pavlos Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146:193–228, 2003.
- [Pearl, 1990] Judea Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proc. TARK*, pages 121–135, 1990.
- [Reiter, 2001] Raymond Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, 2001.
- [Scherl and Levesque, 2003] Richard Scherl and Hector J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1–2):1–39, 2003.
- [Schwering and Lakemeyer, 2014] Christoph Schwering and Gerhard Lakemeyer. A semantic account of iterated belief revision in the situation calculus. In *Proc. ECAI*, pages 801–806, 2014.
- [Schwering and Lakemeyer, 2015] Christoph Schwering and Gerhard Lakemeyer. Projection in the epistemic situation calculus with belief conditionals. In *Proc. AAAI*, pages 1583–1589, 2015.
- [Shapiro *et al.*, 2011] Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.
- [Spohn, 1988] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, pages 105–134. 1988.
- [Thielscher, 1999] Michael Thielscher. From situation calculus to fluent calculus: State update axioms as a solution to the inferential frame problem. *Artificial Intelligence*, 111(1):277–299, 1999.
- [van Benthem, 2007] Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.