Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR 2020)
Main Track

Changing Beliefs about Domain Dynamics in the Situation Calculus

Toryn Q. Klassen[1,2], **Sheila A. McIlraith**[1,2], **Hector J. Levesque**[1]
[1]Department of Computer Science, University of Toronto, Toronto, Canada
[2]Vector Institute, Toronto, Canada
{toryn, sheila, hector}@cs.toronto.edu

Abstract

Agents change their beliefs about the plausibility of various aspects of domain dynamics—effects of physical actions, results of sensing, and action preconditions—as a consequence of their interactions with the world. In this paper we propose a way to conveniently represent domain dynamics in the situation calculus to support such belief change. Furthermore, we suggest patterns to follow when writing the axioms that describe the effects of actions, and prove how these patterns can control the extent to which observations change the agent's beliefs about action effects. We also discuss the relation of our work to the AGM postulates for belief revision. Finally, we show how beliefs about domain dynamics can be incorporated into a form of regression rewriting to support reasoning.

1 Introduction

In this paper, we present a logical account of belief within the situation calculus (Reiter 2001) in which we show how an agent can change its beliefs about the dynamic properties of actions—effects, preconditions, and sensing—as a consequence of its interactions with the world. For example, the agent may start with some belief about the general effects of an action, but after performing the action and some related sensing, come to realize that the action did not perform as expected, and change its beliefs about the action's effects.

To illustrate the change in beliefs that our account can support, consider an example about picking up and holding objects, where the agent changes its beliefs about how the predicate $\mathsf{H}(x, s)$ (x is held in situation s) changes over time. There's an action $\mathsf{p}(x)$ (the agent tries to pick up x). At one point the agent can believe

$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv a = \mathsf{p}(x) \vee \mathsf{H}(x, s), \tag{1}$$

(i.e., that it's holding an object x if the last action a was trying to pick x up or if it was previously holding x), and then after sensing its failure to pick up a cup c, believe

$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv [a = \mathsf{p}(x) \wedge \neg(\cdots \wedge x = \mathsf{c})] \vee \mathsf{H}(x, s), \tag{2}$$

where the ellipsis stands for an expression identifying when the failure occurred. That is, the agent believes that while it did fail to pick up the cup, that failure was a one-time event. So the agent believes that it will be holding anything it picks up except for that one-time failure. Furthermore, after a second time failing to pick up the cup, the agent believes

$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv (a = \mathsf{p}(x) \wedge x \neq \mathsf{c}) \vee \mathsf{H}(x, s), \tag{3}$$

i.e, that it can only pick up objects other than the cup. Finally, after trying to pick up another object also doesn't result in it being held, the agent believes

$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv (a = \mathsf{p}(x) \wedge \neg \mathsf{S}(x, s)) \vee \mathsf{H}(x, s), \tag{4}$$

i.e., that it can only pick up objects that are not slippery (S). (It assumes the objects it couldn't pick up were slippery.) We will formalize this sequence of belief changes in §5.

We leverage our previous work (Klassen, McIlraith, and Levesque 2018), which represents the beliefs of an agent as being determined by what's true in the most plausible accessible situations (as in (Shapiro et al. 2011)), and where plausibility is measured by counting abnormalities. Beliefs about domain dynamics were not addressed in that work.

Belief revision research following the AGM approach (Alchourrón, Gärdenfors, and Makinson 1985) typically involves revision by arbitrary sentences, but our work does not require that. We are concerned with, instead of directly telling an agent facts about how actions behave, having the agent change its beliefs about dynamics in reaction to observations of the environment (as in our example about picking up objects). It will be up to the axiomatizer to specify the generality of the conclusions the agent should draw (e.g., whether observing a failed attempt to pick up a cup means that that cup can never be picked up, or some broader or narrower conclusion). Indeed, *how to write* such specifications is a major focus of ours (with the downside that the axiomatizer must anticipate these less plausible scenarios). For propositional languages, there has been some work about revising beliefs about domain dynamics (e.g., (Herzig, Perrussel, and Varzinczak 2006; Eiter et al. 2010; Varzinczak 2010; Van Zee et al. 2015)), but they have not usually been concerned with how to specify the generality of conclusions the agent should draw (an exception may be (Eiter et al. 2007; Eiter et al. 2010), which we discuss later).

After some background, in §3 we provide some results on how beliefs about dynamics (effects, sensing, and preconditions) can be determined. We then focus on action effects, suggesting patterns to follow when writing the axioms describing them (§4), and in §5 use these patterns to formalize our example from Equations 1–4. In §6, we provide a result about how (potentially changed) beliefs about action effects can be incorporated into *regression* (Reiter 2001, §4.5), a rewriting technique that can simplify automated reasoning. We discuss related work in §7 before concluding.

2 Preliminaries

2.1 Language and Notation

The situation calculus (McCarthy and Hayes 1969; Reiter 2001) is a language for describing actions and change, with semantics given by (multi-sorted) second-order logic. The sorts are situations, actions, and objects (for convenience, we let the natural numbers be a subsort of objects). We use s as a variable of type situation, a as a variable of type action, x and y as variables of type object, and i and j as variables of type natural number. Predicates are written with an uppercase first letter (e.g., H), and function symbols (including constants) with a lowercase first letter (e.g., sH). For a finite set of formulas Γ, their conjunction can be written as $\bigwedge \Gamma$. We may abbreviate a (possibly empty) sequence of terms $\tau_1, \ldots, \tau_k$ using vector notation as $\vec{\tau}$. A *ground* term does not refer to any variables. We may omit leading universal quantifiers when writing sentences. Also, we use $\forall\phi$ to denote the universal closure of a formula ϕ, i.e., the sentence $\forall\vec{x}.\ \phi$, where $\vec{x}$ is the sequence of all free variables in ϕ.

In the situation calculus, properties that can change (e.g., whether an object is being held) are modelled using *fluents*, predicates (or functions) whose last argument is a situation. For example, $\mathsf{H}(x, s)$ could represent the property of the agent holding x in situation s. We may informally express $\mathsf{H}(x, s)$ by saying that $\mathsf{H}(x)$ is true in s.

Time is modelled as a branching structure: from a situation s, for any action a, $\mathsf{do}(a, s)$ is the future situation that results from performing a in s. We use the abbreviation $\mathsf{do}([a_1, \ldots, a_k], s)$ for $\mathsf{do}(a_k, (\ldots, \mathsf{do}(a_1, s)))$. The notation $s \sqsubseteq s'$ means that s' is the situation resulting from applying zero or more actions in s. The constant s_0 denotes the actual initial situation—the root of the situation tree. There may be other initial situations, that serve the purpose of being epistemic alternatives for the agent. We can use the abbreviation $\mathsf{Init}(s) \overset{\text{def}}{=} \neg\exists a, s'.\ s = \mathsf{do}(a, s')$ to describe initial situations. We also find it convenient to have the function $\mathsf{root}(s)$ (from (Shapiro 2005)) that returns the initial situation preceding s (i.e., the root of s's tree).

2.2 Action Theories

How the actions in any particular domain behave is typically described using some variation of Reiter's basic action theories (Reiter 2001). These include initial state axioms that describe s_0, successor state axioms (SSAs) that describe how actions change the world, preconditions axioms that describe when actions are possible to execute, and (sometimes) sensing axioms that describe how sensors work. Action theories also have unique names axioms for actions.

There also are foundational axioms describing the do function and $\sqsubseteq$. Here, we assume the foundational axioms also describe $\mathsf{root}(s)$ and assert that there exist initial situations where all fluents (not including the special symbols Poss, SF or B, that we will introduce shortly) take all possible combinations of values.

To further describe action theories, we introduce *uniform* formulas (Reiter 2001, Definition 4.4.1). Intuitively, a formula φ is *uniform* in a situation term σ if φ describes only the situation σ. That means, among other details, that φ does not quantify over situations and that σ is the last argument to any fluent appearing in φ. Uniform formulas also cannot refer to the Poss, SF, or B predicates.

Initial state axioms are uniform in s_0. An SSA for a fluent F is a sentence of the form

$$F(\vec{x}, \mathsf{do}(a, s)) \equiv \phi_F(\vec{x}, a, s)$$

where ϕ_F is uniform in s (and all the free variables are implicitly universally quantified). An SSA describes how the value of F in a non-initial situation is determined by the action that just happened and the last situation.

A precondition axiom is a sentence of the form

$$\mathsf{Poss}(\alpha(\vec{x}), s) \equiv \phi_\alpha(\vec{x}, s)$$

where α is an action function symbol (i.e., $\alpha(\vec{x})$ is a term of type action) and $\phi_\alpha(\vec{x}, s)$ is uniform in s. Poss is a special predicate meant to indicate that the action is possible to execute. Similarly, a sensing axiom is a sentence of the form

$$\mathsf{SF}(\alpha(\vec{x}), s) \equiv \phi_\alpha(\vec{x}, s)$$

which is like a precondition axiom except for referring on the LHS to the SF predicate. Intuitively, actions produce binary sensing results, and the SF predicate indicates which results are positive.

Much as traditional modal logics of beliefs determine beliefs using an accessibility relation over possible worlds, in the situation calculus belief can be defined in terms of an accessibility relation $\mathsf{B}(s', s)$ over situations (Scherl and Levesque 2003), saying that s' is accessible from s.[1] Similarly to Scherl and Levesque, we require that B be described by this SSA-like axiom:

$$\begin{aligned}\mathsf{B}(s'', \mathsf{do}(a, s)) \equiv \big[&\exists s'.\ \mathsf{B}(s', s) \wedge (s'' = \mathsf{do}(a, s')) \wedge \\ &\mathsf{Poss}(a, s') \wedge (\mathsf{SF}(a, s') \equiv \mathsf{SF}(a, s))\big]\end{aligned} \tag{5}$$

That is, in any accessible situation the same actions that really occurred have been executed, they had to have been possible, and they produced the same sensing results.

As in some previous papers using the situation calculus, we will make special use of the situation term "now". We will typically describe beliefs or knowledge of the agent with formulas that refer to now, which intuitively will refer to the situation that the agent thinks it's currently in. Given a formula ϕ referring to now, we will write $\phi[s]$ for the formula that is like ϕ but substitutes s for now.

We can define a knowledge operator $\mathbf{Knows}(\phi, s)$ ("ϕ is known in s") to describe what the agent is *certain* of: $\mathbf{Knows}(\phi, s) \overset{\text{def}}{=} \forall s'.\ \mathsf{B}(s', s) \supset \phi[s']$. Note that we are not requiring "knowledge" to be true in reality, just in all accessible situations.

However, we are more interested in *belief*, which we will follow Shapiro et al. (2011) in defining as what's true, not in all accessible situations, but in all the *most plausible* accessible situations. As in our previous work (Klassen, McIlraith, and Levesque 2018), we define one situation as being more plausible than another if the sum of the cardinalities of

[1]Note that the order of arguments to B is the opposite of the convention in modal logics.

the extensions of certain fluents, called *abnormality* fluents, is smaller. We use the notation $s' \leq_{\mathsf{pl}} s''$ to abbreviate a (second-order) formula saying that s' is at least as plausible as s''. A belief operator $\mathbf{Bel}(\phi, s)$ ("ϕ is believed in s") can be defined as an abbreviation for

$$\forall s'. \, [\mathsf{B}(s', s) \wedge \forall s''. \, \mathsf{B}(s'', s) \supset s' \leq_{\mathsf{pl}} s''] \supset \phi[s'].$$

We can think of the agent as using a form of *cardinality-based circumscription* (Liberatore and Schaerf 1997; Sharma and Colomb 1997), a variant of the subset-based minimization originally used in circumscription (McCarthy 1980), in determining its beliefs. In those terms, abnormalities are minimized (w.r.t. cardinality) while all other fluents are allowed to vary.[2] (Note that beliefs can be retracted by shrinking the B relation, unlike with knowledge.)

We will follow the assumption made over most of (Klassen, McIlraith, and Levesque 2018) that abnormality fluents don't change their value over time; i.e., for any abnormality fluent Ab_i, it has this SSA:

$$\mathsf{Ab}_i(\vec{x}_i, \mathsf{do}(a, s)) \equiv \mathsf{Ab}_i(\vec{x}_i, s). \tag{6}$$

We will see that that constraint is not a great limitation, as we can represent implausible events as being predetermined by the initial situation. Note that abnormalities need a situation argument, despite not changing, so that we can associate different plausibility levels with different initial situations.

We will find it convenient to have abnormalities with associated numeric *weights* that determine how much they contribute to the implausibility of a situation. We could do so without changing the formalism by introducing the shorthand $\mathsf{Ab}_i^k(\vec{x}, s) \stackrel{\text{def}}{=} \bigwedge_{j=1}^{k} \mathsf{Ab}_i(j, \vec{x}, s)$. Intuitively, Ab_i^k behaves as an abnormality fluent with weight k should; for it to be true is counted as k abnormalities.

To wrap things up, here's a formal definition of the action theories we'll use, which are essentially the same as those we considered in (Klassen, McIlraith, and Levesque 2018) and so which we give the same name.

Definition 1 (IAAT)**.** An *immutable abnormality action theory (IAAT)* is a set of axioms

$$\Sigma_{\mathsf{found}} \cup \Sigma_{\mathsf{ssa}} \cup \Sigma_{\mathsf{pre}} \cup \Sigma_{\mathsf{sense}} \cup \Sigma_0 \cup \Sigma_{\mathsf{una}} \cup \\ \Big\{ \forall s. \, \mathsf{B}(s, \mathsf{s}_0) \equiv \Big(\mathsf{Init}(s) \wedge \bigwedge_{\phi \in \Sigma_{\mathsf{KB}}} \phi[s] \Big) \Big\}$$

where Σ_{found} is the set of foundational axioms (as we previously described), Σ_{ssa} is the set of successor state axioms (including Equation 5 for B, and axioms for each abnormality fluent in the form of Equation 6), Σ_{pre} is the set of precondition axioms, Σ_{sense} is a set of sensing axioms, Σ_0 is the set of initial state axioms, Σ_{una} is the set of unique names axioms for actions, and Σ_{KB} is a set of axioms (uniform in now) describing what the agent initially believes.

For later use, we will require that the language of an IAAT includes a functional fluent, $\mathsf{history}(s)$, which we will use to store a representation of the sequence of actions that have occurred. To define the history fluent, we assume Σ_0 contains an axiomatization of lists, specifying how concatenation works, and that $\cdot$ is a function symbol for concatenation. We require that Σ_{ssa} contain the following SSA: $\mathsf{history}(\mathsf{do}(a, s)) = \mathsf{history}(s) \cdot a$. Σ_{KB} should contain $\mathsf{history}(\mathsf{now}) = \langle\rangle$, where $\langle\rangle$ denotes the empty list.

Observe that an IAAT constrains the initially accessible situations to be those initial situations where Σ_{KB} is true. We will find it useful to have a symbol to denote the part of an IAAT that describes the domain dynamics:

$$\Sigma_{\mathsf{dyn}} \stackrel{\text{def}}{=} \Sigma_{\mathsf{ssa}} \cup \Sigma_{\mathsf{pre}} \cup \Sigma_{\mathsf{sense}}$$

3 Beliefs about Domain Dynamics

We will be exploring beliefs entailed by our action theories (IAATs) about SSAs, preconditions, and sensing axioms, and how to determine them. Later (§4) we suggest having the descriptions of domain dynamics in the theory refer to abnormalities, so as to describe less plausible ways that the domain might behave. The techniques of this section can then allow us, in some cases, to find beliefs about dynamics that don't refer to abnormalities.

To start, note that given any IAAT Σ, the agent will always believe the SSAs, precondition axioms, and sensing axioms written in it, since they hold at all situations. However, we are more interested in what the agent believes about the domain's dynamics in the situation tree it's on, i.e., in situations following from $\mathsf{root}(\mathsf{now})$. Therefore, we define the notion of an axiom that holds on a (sub)tree, rooted at σ.[3]

Definition 2 (relativized axiom)**.** Let $\phi(s)$ be such that $\forall s. \, \phi(s)$ is an SSA, precondition axiom, or sensing axiom. Then the corresponding *axiom relativized to* σ, where σ is a situation term, is the formula $\forall s. \, (\sigma \sqsubseteq s) \supset \phi(s)$.

Henceforth, when we informally talk about the agent believing an axiom about dynamics, we really mean that it believes the corresponding axiom relativized to $\mathsf{root}(\mathsf{now})$.

Definition 3 ($\Gamma{:}\sigma$)**.** Let Γ be a set of axioms about dynamics. Given a situation term σ, Γ relativized to σ, written $\Gamma{:}\sigma$, is the set of corresponding axioms relativized to σ.

In terms of axioms relativized to $\mathsf{root}(\mathsf{now})$, the agent will still believe the axioms in Σ_{dyn}, i.e., we will always have that $\Sigma \models \forall s. \, \mathbf{Bel}\left(\bigwedge \Sigma_{\mathsf{dyn}}{:}\mathsf{root}(\mathsf{now}), \; s\right)$. However, as the agent changes its beliefs about the abnormality fluents, it may come to believe that various other axioms are equivalent to the original ones (and so also believe them). For example, if Σ includes the SSA

$$\mathsf{H}(x, \mathsf{do}(a, s)) \;\equiv\; (a = \mathsf{p}(x) \wedge \neg \mathsf{Ab}_1(s)) \vee \mathsf{H}(x, s) \tag{7}$$

and the agent comes to believe that Ab_1 is true on the situation tree it's on, then the agent will as a result believe a simpler (relativized) SSA saying that H does not change:

$$\mathsf{H}(x, \mathsf{do}(a, s)) \;\equiv\; (a = \mathsf{p}(x) \wedge \neg \mathsf{True}) \vee \mathsf{H}(x, s)$$

That simplifies to $\mathsf{H}(x, \mathsf{do}(a, s)) \equiv \mathsf{H}(x, s)$.

The following definition will be useful in describing what the agent believes about abnormalities.

[2] The reason for using a cardinality-based approach was technical; see (Klassen, McIlraith, and Levesque 2018, section 5.2).

[3] Such axioms have sometimes been used in action theories with multiple initial situations (Lakemeyer and Levesque 1998).

Definition 4 (Ab account)**.** Suppose we have a language with n abnormality fluents, $\mathsf{Ab}_1, \ldots, \mathsf{Ab}_n$, of possibly differing arities. An *Ab account* ξ is an expression

$$\xi(\mathsf{now}) \stackrel{\text{def}}{=} \bigwedge_{\mathsf{Ab}_i \in R} \forall \vec{x}.\, \mathsf{Ab}_i(\vec{x}, \mathsf{now}) \equiv \xi_i(\vec{x}),$$

where $R \subseteq \{\mathsf{Ab}_1, \ldots, \mathsf{Ab}_n\}$, containing a conjunct corresponding to each Ab_i fluent in R. If Ab_i is an $(m+1)$-ary fluent (where the last of those arguments is the situation) then the expression ξ_i is of the form $\left(\bigvee_{k=1}^{\ell} \bigwedge_{j=1}^{m} x_j = \tau_{jk}\right)$ for some $\ell \geq 0$, where the τ_{jk} are ground terms that do not refer to any situation term. We call R the *range* of ξ.

Intuitively, an Ab account ξ characterizes the extension of each abnormality fluent in its range. Note that if Ab_i is a unary fluent (taking only a situation argument), the expression ξ_i in an Ab account ξ is either True or False. Also, any Ab account requires that there be only finitely many abnormalities, so there can be situations in which no Ab account is true.

Ab accounts are not normally included in action theories, but are things that may be believed or disbelieved by the agent. For example, suppose we're working with a theory including the SSA from Equation 7. If the agent observes that $\mathsf{p}(x)$ fails to make H true of x, then the agent may come to believe the Ab account $(\mathsf{Ab}_1(\mathsf{now}) \equiv \mathsf{True})$. Recall that abnormalities do not change over time, so if Ab_1 is true "now", it was always and will always be true. So, as the next lemma notes, if an agent believes an Ab account holds now, then it believes that account has held and will hold forever.

Lemma 1. *For any IAAT Σ, Ab account ξ, and ground action sequence $\vec{\alpha}$,*

$$\Sigma \models \mathbf{Bel}(\xi(\mathsf{now}) \supset \forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \xi(s), \mathsf{do}(\vec{\alpha}, \mathsf{s}_0))$$

Proof. This follows from abnormalities not changing and the terms in $\xi(s)$ not depending on s. $\square$

The main role to which we put abnormalities is as markers of subjective plausibility. We are typically more interested in the non-abnormality fluents, and what the agent believes about them, i.e., in beliefs about *normal formulas*.

Definition 5 (normal formula)**.** A formula is *normal* if it doesn't refer to any abnormality fluents.

Definition 6 (normalization)**.** Given a formula ϕ and an Ab account ξ, the normalization of ϕ w.r.t. ξ is a formula ϕ' which is like ϕ but, for each Ab_i in the range of ξ, replaces each occurrence of any subformula of the form $\mathsf{Ab}_i(\vec{\tau}, \sigma)$ (where σ is a situation term and $\vec{\tau}$ are other terms) with $\xi_i(\vec{\tau})$.

For example, if ϕ is the SSA from Equation 7,

$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv (a = \mathsf{p}(x) \wedge \neg \mathsf{Ab}_1(x, s)) \vee \mathsf{H}(x, s),$$

and ξ is the Ab account $\mathsf{Ab}_1(x, \mathsf{now}) \equiv (x = \mathsf{c} \vee x = \mathsf{d})$, then the normalization of ϕ with respect to ξ is

$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv (a = \mathsf{p}(x) \wedge \neg(x = \mathsf{c} \vee x = \mathsf{d})) \vee \mathsf{H}(x, s).$$

Note that normalization is defined for any formula ϕ, and if an Ab account ξ includes in its range every abnormality fluent mentioned by ϕ then the result of normalizing ϕ w.r.t. ξ will be a normal formula.

We will see that in some cases the agent will believe the normalizations of certain sentences it believes.

Proposition 1. *Let Σ be an IAAT. Let $\forall s.\, \phi(s)$ be an SSA, precondition axiom, or sensing axiom in Σ. Let $\vec{\alpha}$ be a sequence of ground actions. If there is an Ab account ξ such that $\Sigma \models \mathbf{Bel}(\xi, \mathsf{do}(\vec{\alpha}, \mathsf{s}_0))$ and ϕ' is the normalization of ϕ with respect to ξ, then*

$$\Sigma \models \mathbf{Bel}(\forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \phi'(s), \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)).$$

Proof. Suppose that there is an Ab account ξ such that $\Sigma \models \mathbf{Bel}(\xi, \mathsf{do}(\vec{\alpha}, \mathsf{s}_0))$ and ϕ' is the normalization of ϕ with respect to ξ. By Lemma 1 we have that

$$\Sigma \models \mathbf{Bel}(\forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \xi(s), \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)),$$

and it's easy to see that Σ entails

$$\mathbf{Bel}(\forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \xi(s) \supset [\phi'(s) \equiv \phi(s)], \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)).$$

Therefore, since the agent believes $\forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \phi(s)$ in $\mathsf{do}(\vec{\alpha}, \mathsf{s}_0)$, we get the result. $\square$

Proposition 1 can be applied to show, given particular action theories, that after certain actions the agent believes simpler dynamics axioms than those that were written in its initial knowledge base (we will put it to use in later sections).

A generalization we can make to Proposition 1 is to consider cases where the agent believes a disjunction of Ab accounts (but not necessarily any of the disjuncts). To illustrate why that is useful, consider a scenario where an agent unexpectedly fails to pick up an object and doesn't know if that failure was because the object was red or because the object was fuzzy. Then we might want the agent to believe the disjunction of "I can pick up any non-red object" and "I can pick up any non-fuzzy object". For cases like this, the more general Proposition 2 below is relevant.

Proposition 2. *Let Σ be an IAAT. Let $\forall s.\, \phi(s)$ be an SSA, precondition axiom, or sensing axiom in Σ. Let $\vec{\alpha}$ be a sequence of ground actions. If there are Ab accounts $\xi^1, \ldots, \xi^k$ such that $\Sigma \models \mathbf{Bel}\left(\bigvee_{i=1}^{k} \xi^i, \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)\right)$ and ϕ'_i is the normalization of ϕ with respect to ξ^i for each i, then*

$$\Sigma \models \mathbf{Bel}\left(\bigvee_{i=1}^{k} \forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \phi'_i(s), \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)\right).$$

Proof. Similarly to in the proof of Proposition 1, it's easy to see that for each i, Σ entails

$$\mathbf{Bel}(\forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \xi^i(s) \supset [\phi'_i(s) \equiv \phi(s)], \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)).$$

Therefore, since the agent believes $\forall s \sqsupseteq \mathsf{root}(\mathsf{now}).\, \phi(s)$ in $\mathsf{do}(\vec{\alpha}, \mathsf{s}_0)$, we can get the result (using Lemma 1). $\square$

4 Patterns for SSAs

In this section we consider how the axiomatizer should write SSAs, so that the agent will change its beliefs by the desired amount given new evidence. We suggest patterns to follow, based on a traditional way of writing SSAs in terms of positive and negative effects. Following Reiter (2001, §3.2.7), an SSA for a unary fluent F would be written in the form

$$F(x, \mathsf{do}(a, s)) \equiv \gamma^+(x, a, s) \vee \big(\neg\gamma^-(x, a, s) \wedge F(x, a, s)\big) \qquad (8)$$

where the formula γ^+ describes positive effects on F, i.e., conditions under which F becomes true, and the formula γ^- describes negative effects on F, i.e., conditions under which F becomes false.

Definition 7 (revisable SSA)**.** We will say that an SSA is a revisable SSA if it is written in the form

$$F(x, \mathsf{do}(a, s)) \equiv \big(\gamma^+(x, a, s) \wedge \neg \textstyle\bigvee_i \epsilon_i(x, a, s)\big) \vee \big(\neg\gamma^-(x, a, s) \wedge F(x, a, s)\big)$$

where γ^+ and γ^- are normal formulas.

Intuitively, each ϵ_i in a revisable SSA describes a less plausible case in which action a fails to make $F(x)$ true. The structure of a revisable SSA could easily be rearranged to instead describe less plausible cases in which F may fail to become false, may become true, or may become false. For reasons of space we'll just consider Definition 7.

What might we want the ϵ_i formulas to look like? We suggest three forms, for dealing with *exceptional objects*, *exceptional classes*, and *one-time exceptions*.

Exceptional objects We may want an agent to conclude from an unexpected observation involving a particular object that actions always affect that object differently. To achieve this, we could make $\epsilon_i(x, a, s)$ take the form $\mathsf{Ab}_j(x, s)$. Intuitively, if the agent comes to believe that $\mathsf{Ab}_j(\mathsf{c}, \mathsf{now})$ is true of a particular object c (e.g., by observing that F does not become true of c when expected), then the agent will conclude that all actions will fail to make F true of c. Note that it's not necessary for the action theory to say anything else about Ab_j for this to work (other than Ab_j's own SSA, specifying that it doesn't change).

Exceptional classes Another sort of generalization that we might want the agent to make on observing an unexpected (non-)effect is that that unexpected behavior will always occur when dealing with objects from a particular class. For example, an agent might conclude from failing to pick up an object that some objects are too slippery to be picked up. To achieve this, we could make $\epsilon_i(x, a, s)$ take the form $[P(x, s) \wedge \mathsf{Ab}_j(s)]$ where P is a fluent. Note that $\mathsf{Ab}_i(s)$ does not take x as an argument, so it being true would mean that *any* objects on the situation tree which s is part of behave abnormally when they have property P.

One-time exceptions We may want an agent to, when observing an unexpected (non-)effect of an action a in situation s, just accept that a had that (non-)effect in s, while not changing its beliefs about how any action will behave in any other situation. This can be viewed as a sort of minimal way of adjusting the agent's beliefs to keep them consistent. We will call such isolated unexpected (non-)effects "one-time exceptions". We could model this by making $\epsilon_i(x, a, s)$ take the form $\mathsf{Ab}_j(\mathsf{history}(s), x, a, s)$. Since the abnormality depends on the sequence of actions $\mathsf{history}(s)$ (from Definition 1), each new unexpected action outcome would require another abnormal atom to be true.

We call a revisable SSA that uses only these three patterns a simple SSA:

Definition 8 (simple SSA)**.** A revisable SSA is a simple SSA if each $\epsilon_i(x, a, s)$ is in one of the following forms (the abnormalities may have associated weights):

1. $\mathsf{Ab}_j(x, s)$ (for exceptional objects),
2. $[P(x, s) \wedge \mathsf{Ab}_j(s)]$ (for an exceptional class),
3. or $\mathsf{Ab}_j(\mathsf{history}(s), x, a, s)$ (for one-time exceptions).

We want to show that simple SSAs behave as desired. To facilitate exposition we introduce the next abbreviation.

Definition 9 ($\vec{\alpha} \rightsquigarrow \phi$)**.** Suppose $\vec{\alpha}$ is a sequence of action terms and ϕ is a formula. Then we define

$$\vec{\alpha} \rightsquigarrow \phi \overset{\text{def}}{=} \mathbf{Bel}\big(\forall(\mathsf{root}(\mathsf{now}) \sqsubseteq s \supset \phi),\ \mathsf{do}(\vec{\alpha}, \mathsf{s}_0)\big)$$

In the case where the length of $\vec{\alpha}$ is 0, we write $\rightsquigarrow \phi$.

That is, $\alpha_1, \ldots, \alpha_k \rightsquigarrow \phi$ is a formula saying that after performing the actions $\alpha_1, \ldots, \alpha_k$ starting from s_0, the agent believes the universal closure of ϕ, where the variable s is restricted to be a successor of $\mathsf{root}(\mathsf{now})$, what the agent thinks is the initial situation. For example,

$$\vec{\alpha} \rightsquigarrow [\mathsf{H}(x, \mathsf{do}(a, s)) \equiv a = \mathsf{p}(x) \vee \mathsf{H}(x, s)]$$

says that after the actions $\vec{\alpha}$, the agent believes that for any situation s which is on the tree rooted at $\mathsf{root}(\mathsf{now})$, and for any object x and action a, the stated relation holds (i.e., x is held after performing a in s just in case $a = \mathsf{p}(x)$ or x was already held in s).

The following proposition illustrates what sorts of normal SSAs an agent may believe when a simple SSA is used in Σ_{ssa}. We'll see a more concrete example in the next section.

Proposition 3. *Suppose Σ is an IAAT with a simple SSA for F, and $\vec{\alpha}$ is a sequence of ground actions. If there is an Ab account ξ such that $\Sigma \models \mathbf{Bel}(\xi, \mathsf{do}(\vec{\alpha}, \mathsf{s}_0))$ and which has in its range all the abnormalities referred to by F's SSA, then*

$$\Sigma \models \vec{\alpha} \rightsquigarrow F(x, \mathsf{do}(a, s)) \equiv \big[\big(\gamma^+(x, a, s) \wedge \neg\phi(x, a, s)\big) \vee \big(\neg\gamma^-(x, a, s) \wedge F(x, s)\big)\big]$$

where ϕ is a (possibly empty) disjunction, containing the following disjuncts, depending on the original simple SSA:

1. *For each $\epsilon_i(x, a, s)$ of the form $\mathsf{Ab}_j(x, s)$, ϕ contains either no corresponding disjunct, or a disjunct of the form $\big[\bigvee_{\tau \in T}(x = \tau)\big]$ for some finite set T of ground terms.*
2. *For each $\epsilon_i(x, a, s)$ of the form $[P(x, s) \wedge \mathsf{Ab}_j(s)]$, ϕ contains either no corresponding disjunct, or $P(x, s)$.*
3. *For each $\epsilon_i(x, a, s)$ of the form $\mathsf{Ab}_j(\mathsf{history}(s), x, a, s)$, ϕ contains either no disjunct, or a disjunct of the form*

$$[\textstyle\bigvee_{\langle\tau_1,\tau_2,\tau_3\rangle \in T} (\mathsf{history}(s) = \tau_1 \wedge x = \tau_2 \wedge a = \tau_3)]$$

for some finite set T of triples of ground terms.

Proof. In the normalization of the original SSA by ξ, the abnormal atoms in each of the $\epsilon_i(x, a, s)$ expressions will get replaced, yielding an SSA as described (for (2), there's some additional simplification needed to remove expressions that include True or False). That that SSA is believed follows from Proposition 1. □

Intuitively, in part (1) of Proposition 3, T is a list of exceptional objects, the result in (2) depends on whether the agent has determined P to be an exceptional class, and in (3), T identifies very specific circumstances for one-time exceptions. Note that a reason we used the history fluent in our one-time exception pattern, rather than just referring to a situation (which also stores a list of actions), is because the right-hand-sides of SSAs are supposed to be uniform formulas, and so cannot refer to equality of situation terms.

The next proposition says that if we write the SSA for F as a simple SSA, then (under some conditions) initially the agent will believe the traditional SSA from Equation 8.

Proposition 4. *Let Σ be an IAAT. Suppose that the SSA for F in Σ_{ssa} is a simple SSA and that Σ_{KB} (the agent's initial knowledge base, from Definition 1) does not refer to any abnormality fluent. Then*

$$\Sigma \models \rightsquigarrow \big[F(\vec{x}, \mathsf{do}(a, s)) \equiv \gamma^+(\vec{x}, a, s) \vee \big(\neg\gamma^-(\vec{x}, a, s) \wedge F(\vec{x}, a, s)\big)\big]$$

Proof. Since Σ_{KB} does not refer to abnormalities, it's easy to see that there are accessible situations from s_0 in which every abnormality is false. So in s_0 the agent believes the Ab account $\bigwedge_{i=1}^{n} \forall \vec{x}.\ \mathsf{Ab}_i(\vec{x}) \equiv \mathsf{False}$. The normalization of any simple SSA w.r.t. that Ab account is (after some simplification) $F(\vec{x}, \mathsf{do}(a, s)) \equiv \big[\gamma^+(\vec{x}, a, s) \vee \big(\neg\gamma^-(\vec{x}, a, s) \wedge F(\vec{x}, a, s)\big)\big]$. The result follows from Proposition 1. □

While Proposition 3 and Proposition 4 only consider SSAs dealing with less-plausible failures of positive effects, analogous results could be shown for SSAs dealing with other types of less plausible behavior. Note that in some cases it may be possible to more compactly write the SSA by distributing the less plausible conditions throughout it rather than grouping them together as we've done.

5 Example of Changing Beliefs about SSAs

We are now ready to formalize the revision sequence (1–4) described in the example from the introduction. We do so by constructing an IAAT $\Sigma_{Holding}$ with the fluents $\mathsf{H}(x, s)$, saying that x is being held in s, and $\mathsf{S}(x, s)$, that x is slippery in s. The actions are $\mathsf{p}(x)$, the action to (try to) pick up x, and sH, which senses whether anything is held. There are constants c and d (to represent a cup and a dish).

The sensing axioms are

$$\mathsf{SF}(\mathsf{sH}, s) \equiv \exists x.\ \mathsf{H}(x, s) \qquad \mathsf{SF}(\mathsf{p}(x)) \equiv \mathsf{True}$$

Note that picking up does not provide sensing information. All actions are always possible to execute. The SSAs are

$$\mathsf{S}(x, \mathsf{do}(a, s)) \equiv \mathsf{S}(x, s)$$
$$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv [(a = \mathsf{p}(x) \wedge \neg \textstyle\bigvee_i \epsilon_i(a, x, s)) \vee \mathsf{H}(x, s)]$$

where $\bigvee_i \epsilon_i(a, x, s))$ is

$$\mathsf{Ab}_1^2(\mathsf{history}(s), x, a, s) \vee \mathsf{Ab}_2^3(x, s) \vee [\mathsf{S}(x, s) \wedge \mathsf{Ab}_3^4(s)]$$

The disjuncts with lower associated weights (superscripts) are the ones that the agent will tend to find more plausible. So, a one-time exception is more plausible than an exceptional object, which is more plausible than not being able to pick up slippery things (i.e., that slippery objects are an exceptional class). Meanwhile, what's slippery never changes. The initial state axioms include

$$\neg\mathsf{H}(x, \mathsf{s}_0) \qquad \mathsf{Ab}_2^3(x, \mathsf{s}_0)$$

That is, nothing is initially held, and every object is actually abnormal—a consequence of this is that no object can be held in any successor of s_0. So in reality, p actions are ineffectual; they cannot cause anything to become held. The agent's initial knowledge base, Σ_{KB}, is empty.

The four points in Proposition 5 below show how during the action sequence $\mathsf{p}(\mathsf{c})$, sH, $\mathsf{p}(\mathsf{c})$, sH, $\mathsf{p}(\mathsf{d})$, and sH (trying to pick up the cup twice, then trying to pick up the dish, and sensing after each attempt) the agent believes the SSAs from Equations 1–4 from the introduction.

Proposition 5. *Let $\Sigma_{Holding}$ be the IAAT described above. Then it entails each of the following beliefs about SSAs:*

1. $\rightsquigarrow [\mathsf{H}(x, \mathsf{do}(a, s)) \equiv a = \mathsf{p}(x) \vee \mathsf{H}(x, s)]$

2. $\mathsf{p}(\mathsf{c}), \mathsf{sH} \rightsquigarrow$
$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv$
$[a = \mathsf{p}(x) \wedge \neg(\mathsf{history}(s) = \langle\rangle \wedge x = \mathsf{c})] \vee \mathsf{H}(x, s)$

3. $\mathsf{p}(\mathsf{c}), \mathsf{sH}, \mathsf{p}(\mathsf{c}), \mathsf{sH} \rightsquigarrow$
$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv (a = \mathsf{p}(x) \wedge x \neq \mathsf{c}) \vee \mathsf{H}(x, s)$

4. $\mathsf{p}(\mathsf{c}), \mathsf{sH}, \mathsf{p}(\mathsf{c}), \mathsf{sH}, \mathsf{p}(\mathsf{d}), \mathsf{sH} \rightsquigarrow$
$\mathsf{H}(x, \mathsf{do}(a, s)) \equiv (a = \mathsf{p}(x) \wedge \neg\mathsf{S}(x, s)) \vee \mathsf{H}(x, s)$

Proof. We sketch the reason for each entailment. By using Proposition 1, the result can follow from showing which abnormalities the agent believes in the relevant situations. We use the notation $\langle\alpha_1, \ldots, \alpha_k\rangle$ for the term representing the sequence of actions $\alpha_1, \ldots, \alpha_k$.

1. In the initial situation, it's consistent with the agent's knowledge that all abnormalities are false.
2. After the actions $\mathsf{p}(\mathsf{c})$ and sH, the agent knows that executing $\mathsf{p}(\mathsf{c})$ (from a situation with an empty history) failed to cause $\mathsf{H}(\mathsf{c})$. So $\mathsf{Ab}_1^2(\langle\rangle, \mathsf{c}, \mathsf{p}(\mathsf{c})) \vee \mathsf{Ab}_2^3(\mathsf{c}) \vee [\mathsf{S}(\mathsf{c}) \wedge \mathsf{Ab}_3^4]$ must be true at all accessible situations. The most plausible of those are where $\mathsf{Ab}_1^2(\langle\rangle, \mathsf{c}, \mathsf{p}(\mathsf{c}))$ is true and all other abnormalities are false (because Ab_1^2 has the lowest weight). (Note that an $a = \mathsf{p}(\mathsf{c})$ condition could be include in the believed SSA but is redundant.)
3. After $\mathsf{p}(\mathsf{c}), \mathsf{sH}, \mathsf{p}(\mathsf{c}), \mathsf{sH}$, the agent has observed two cases in which picking up c failed. The most plausible accessible situations are those where $\mathsf{Ab}_2^3(\mathsf{c})$ is true and all other abnormalities are false. Note that situations where instead there were two one-time exceptions—where $\mathsf{Ab}_1^2(\langle\rangle, \mathsf{c}, \mathsf{p}(\mathsf{c}))$ and $\mathsf{Ab}_1^2(\langle\mathsf{p}(\mathsf{c}), \mathsf{sH}\rangle, \mathsf{c}, \mathsf{p}(\mathsf{c}))$ are true—are less plausible, as the sum of their weights is 4.

4. After these actions, the agent has seen two failures to pick up c and one to pick up d. The most plausible accessible situations are those where slippery objects can't be picked up (and c and d are slippery), i.e., where $\mathsf{Ab}_3^4 \wedge \mathsf{S}(\mathsf{c}) \wedge \mathsf{S}(\mathsf{d})$ is true, and there are no other abnormalities. □

6 Regression

In this section we suggest a way beliefs about SSAs could be taken advantage of in regression, a syntactic procedure often used in automated reasoning about situation calculus formulas. Reiter (2001) showed that a certain class of formulas, the *regressable* formulas, can be rewritten using regression so as not to refer to any non-initial situations (this can make them easier to prove, since some axioms from a basic action theory will no longer be needed). The essential feature of regression is recursively replacing substitution instances of the left-hand-sides of SSAs with their right-hand-sides.

In regression as it's usually considered, the SSAs used are those the axiomatizer wrote. A novel alternative that our work suggests is to use other SSAs that the agent happens to believe at a given time. A computational advantage might be gained in some cases, because some believed SSAs may lead to much smaller or larger regression rewritings than others. To illustrate, an agent could believe both the SSA $\mathsf{P}(x, \mathsf{do}(a, s)) \equiv (\mathsf{P}(\mathsf{f}(x), s) \wedge \mathsf{P}(\mathsf{g}(x), s))$ and the SSA $\mathsf{P}(x, \mathsf{do}(a, s)) \equiv \mathsf{P}(x, s)$. The first SSA's right-hand-side has twice as many atoms as its left-hand-side, so regressing with it could cause an exponential (in the number of applied actions) blowup, while that doesn't happen using the second SSA. For IAATs, the SSAs given by the axiomatizer will often refer to various implausible conditions, and in many situations the agent will believe simpler SSAs.

We will prove (in §6.1) that an agent can use a form of regression, working with any set of SSAs (and precondition axioms and sensing axioms) it believes, to reason about its beliefs. Note that here we apply regression to formulas only *within* belief operators. To regress the whole formula, you would need to additionally apply another form of regression—we call this *full regression* and defer discussing it to §6.2. Also, we leave to future work the important question of how to automatically choose a set of believed SSAs for which regression will be more efficient.

6.1 Regression within Beliefs

Formulas within beliefs typically refer to now. To regress them, we will require them to be "now-regressable", which we define similarly to *regressable* (Reiter 2001, Def. 4.5.1).

Definition 10 (r-regressable)**.** Given a situation term r (e.g., now), a first-order formula ϕ is r-regressable if

- for each term of sort situation mentioned by ϕ, the term has the syntactic form $\mathsf{do}(\vec{\alpha}, r)$
- for each atom of the form $\mathsf{Poss}(\alpha, \sigma)$ or $\mathsf{SF}(\alpha, \sigma)$ mentioned by ϕ, α has the syntactic form $\alpha'(\vec{t})$ where α' is an action function symbol
- ϕ does not have quantification over situations
- ϕ does not mention $\sqsubset$ or compare situations for equality
- ϕ does not mention the B predicate
- ϕ does not mention any functional fluents (this is just for simplicity)

The definition of regression is as follows (based closely on (Reiter 2001, Definition 4.5.3)).

Definition 11. Let $\Delta = \Delta_{\mathsf{ssa}} \cup \Delta_{\mathsf{pre}} \cup \Delta_{\mathsf{sense}}$ be a set of sentences including SSAs, precondition axioms, and sensing axioms for all the fluents and actions. Let ϕ be a now-regressable formula, where WLOG we assume that any variables appearing in ϕ are distinct from those mentioned by Δ. Then the regression of ϕ with respect to Δ is written $\mathcal{R}_1^{\Delta}[\phi]$ and defined case-by-case as follows:

1. ϕ is a situation-independent atom, or a relational fluent atom of the form $F(\vec{\tau}, \mathsf{now})$. Then $\mathcal{R}_1^{\Delta}[\phi] = \phi$.
2. ϕ is a relational fluent atom $F(\vec{\tau}, \mathsf{do}(\alpha, \sigma))$, where the SSA for F in Δ_{ssa} is $F(\vec{x}, \mathsf{do}(a, s)) \equiv \phi_F(\vec{x}, a, s)$. Then $\mathcal{R}_1^{\Delta}[\phi] = \mathcal{R}_1^{\Delta}[\phi_F(\vec{\tau}, \alpha, \sigma)]$.
3. ϕ is an atom of the form $\mathsf{Poss}(\alpha(\vec{\tau}), \sigma)$ or $\mathsf{SF}(\alpha(\vec{\tau}), \sigma)$. In the former case, suppose that the precondition axiom for α in Δ_{pre} is $\mathsf{Poss}(\alpha(\vec{x}), s) \equiv \phi_\alpha(\vec{x}, s)$, and in the latter case, suppose that the sensing axiom for α in Δ_{sense} is $\mathsf{SF}(\alpha(\vec{x}), s) \equiv \phi_\alpha(\vec{x}, s)$. Then $\mathcal{R}_1^{\Delta}[\phi] = \phi_\alpha(\vec{\tau}, s)$.
4. ϕ is a non-atomic formula. Regression is defined inductively as follows: $\mathcal{R}_1^{\Delta}[\neg\phi] = \neg\mathcal{R}_1^{\Delta}[\phi]$, $\mathcal{R}_1^{\Delta}[\phi_1 \wedge \phi_2] = \mathcal{R}_1^{\Delta}[\phi_1] \wedge \mathcal{R}_1^{\Delta}[\phi_2]$, and $\mathcal{R}_1^{\Delta}[\exists x.\, \phi] = \exists x.\, \mathcal{R}_1^{\Delta}[\phi]$.

It can be shown that regressing a now-regressable formula yields a formula uniform in now. The next proposition says that an agent can reason using regression using any set of SSAs that it believes, in the following sense: the agent will believe that any now-regressable formula is equivalent to its regression with respect to those SSAs. (Recall from §2 that $\forall\phi$ is the universal closure of ϕ.)

Proposition 6. *Let $\Delta = \Delta_{\mathsf{ssa}} \cup \Delta_{\mathsf{pre}} \cup \Delta_{\mathsf{sense}}$ be any set of sentences including SSAs, precondition axioms, and sensing axioms for all the fluents and actions. Suppose that σ^* is a ground situation term such that*

$$\Sigma \models \mathbf{Bel}\big(\textstyle\bigwedge \Delta{:}\mathsf{now},\ \sigma^*\big),$$

i.e., the agent in situation σ^* believes that the axioms in Δ apply to future situations. Then for any now*-regressable formula ϕ (which WLOG uses distinct variables from Δ),*

$$\Sigma \models \mathbf{Bel}\big(\forall(\phi \equiv \mathcal{R}_1^{\Delta}[\phi]),\ \sigma^*\big).$$

Proof. Our proof resembles that of the related (Pirri and Reiter 1999, Theorem 2). We assign any now-regressable formula ϕ a triple of numbers, $\mathsf{index}(\phi) = \langle b, d, c\rangle$, where b is 1 if an atom of the form $\mathsf{Poss}(\alpha, \sigma)$ or $\mathsf{SF}(\alpha, \sigma)$ appears in ϕ (and 0 otherwise), d is the greatest depth of nesting of do functions in ϕ, and c is the number of logical connectives/quantifiers in ϕ. The proof is by induction on $\mathsf{index}(\phi)$, with respect to a lexicographic ordering, which we call $\leq_3$.

1. When its index is $\langle 0, 0, 0\rangle$, ϕ is either a situation-independent atom or a relational fluent atom $F(\vec{\tau}, \mathsf{now})$. In either case, $\mathcal{R}_1^{\Delta}[\phi] = \phi$, so the result is trivial.
2. When its index is $\langle 0, d, 0\rangle$ for $d > 0$, ϕ is a relational fluent atom $F(\vec{\tau}, \mathsf{do}(\alpha, \sigma))$. We want to show that Σ entails $\mathbf{Bel}\big(\forall(F(\vec{\tau}, \mathsf{do}(\alpha, \sigma)) \equiv \mathcal{R}_1^{\Delta}[\phi_F(\vec{\tau}, \alpha, \sigma)]), \sigma^*\big)$ where

ϕ_F is from the RHS of the SSA for F in Δ_{ssa}. First, because the agent believes that that SSA applies to now and its successors (and σ is one of those), we get that

$$\Sigma \models \mathbf{Bel}\big(\forall(F(\vec{\tau}, \mathsf{do}(\alpha,\sigma)) \equiv \phi_F(\vec{\tau},\alpha,\sigma)), \sigma^*\big)$$

It can be seen that $\mathsf{index}(\phi_F(\vec{t},\alpha,\sigma)) \leq_3 \langle 0, d-1, c\rangle$ for some c, and since $\langle 0, d-1, c\rangle <_3 \langle 0, d, 0\rangle$, by the inductive hypothesis we get that Σ entails

$$\mathbf{Bel}\big(\forall(\phi_F(\vec{\tau},\alpha,\sigma) \equiv \mathcal{R}_1^{\Delta}[\phi_F(\vec{\tau},\alpha,\sigma)]), \sigma^*\big)$$

Since belief is closed under logical consequence we can put this together with the previous entailment to get the result we want.

3. When its index is $\langle 1, d, 0\rangle$, ϕ is an atom either of the form $\mathsf{Poss}(\alpha(\vec{\tau}),\sigma)$ or $\mathsf{SF}(\alpha(\vec{\tau}),\sigma)$. In either case, the regression of ϕ is $\mathcal{R}_1^{\Delta}[\phi_\alpha(\vec{\tau})]$ where ϕ_α comes from the RHS of a precondition or sensing axiom. It can be seen that $\mathsf{index}(\phi_\alpha(\vec{t},\sigma)) \leq_3 \langle 0, d, c\rangle$ for some c, and $\langle 0, d, c\rangle <_3 \langle 1, d, 0\rangle$. Therefore, this case can be shown similarly to the previous one.
4. When its index is $\langle b, d, c\rangle$ with $c > 0$, ϕ is a non-atomic formula. The result can be seen to follow from the inductive hypothesis and belief being deductively closed. $\square$

6.2 Fully Regressing Formulas

To fully regress formulas containing beliefs (and not just regress formulas within beliefs), we adapt the approach by Schwering and Lakemeyer (2015) from the modal situation calculus. This will *not* subsume the previously described procedure $\mathcal{R}_1$, since for full regression we will not in general be able to make use of axioms that are merely believed. Instead, the relation between the two approaches is complementary; we can (optionally) first use $\mathcal{R}_1$ to make formulas within beliefs uniform in now, and then apply the full regression procedure, which we'll call $\mathcal{R}_2$, to the entire formula. (Also, $\mathcal{R}_1$ is used as a subprocedure by $\mathcal{R}_2$ in a limited way.)

Schwering and Lakemeyer's approach used *conditional beliefs*. Intuitively, a conditional belief in ψ given ϕ, which we will write as $\mathbf{Con}(\psi \mid \phi, s)$, means that in the most plausible accessible situations from s where ϕ is true, ψ is also true. This can be defined as an abbreviation using $\leq_{\mathsf{pl}}$. Belief can be related to conditional belief in the usual way, i.e., $\mathbf{Bel}(\phi, s)$ could equivalently be defined as $\mathbf{Con}(\phi \mid \mathsf{True}, s)$. When fully regressing formulas containing beliefs, we will assume that any expression of the form $\mathbf{Bel}(\phi,\sigma)$ has been replaced with $\mathbf{Con}(\phi \mid \mathsf{True}, \sigma)$.

We can use the regression operator $\mathcal{R}_1$ that we previously defined *within* conditional beliefs, though we need that the agent be *certain* of the axioms Δ used (as opposed to just believing them like in Proposition 6), because in the most plausible accessible situations where the conditional's antecedent is true, merely believed axioms may not hold.

Lemma 2. *Let Σ be an IAAT and $\Delta = \Delta_{\mathsf{ssa}} \cup \Delta_{\mathsf{pre}} \cup \Delta_{\mathsf{sense}}$ a set of axioms such that $\Sigma \models \mathbf{Knows}\left(\bigwedge \Delta{:}\mathsf{now}, \mathsf{s}_0\right)$. Then for any* now*-regressable formulas ψ_1 and ψ_2 using distinct variables from Δ,*

$$\Sigma \models \forall[\mathbf{Con}(\psi_2 \mid \psi_1, \mathsf{s}_0) \equiv \mathbf{Con}(\mathcal{R}_1^{\Delta}[\psi_2] \mid \mathcal{R}_1^{\Delta}[\psi_1], \mathsf{s}_0)].$$

Proof. The key is to note that it would suffice to show that

$$\Sigma \models \forall[\mathbf{Knows}((\psi_1 \equiv \mathcal{R}_1^{\Delta}[\psi_1]) \wedge (\psi_2 \equiv \mathcal{R}_1^{\Delta}[\psi_2]), \mathsf{s}_0)].$$

This is because that would mean that the most plausible accessible situations where ψ_1 is true are exactly the most plausible accessible situations where $\mathcal{R}_1^{\Delta}[\psi_1]$ is true, and whether ψ_2 is true at those situations is equivalent to whether $\mathcal{R}_1^{\Delta}[\psi_2]$ is true at those situations. The proof is similar to that of Proposition 6 but substitutes **Knows** for **Bel**. $\square$

In Lemma 2, we considered knowledge only in s_0 because that's all we'll need for the role that $\mathcal{R}_1$ plays within the broader procedure $\mathcal{R}_2$ that we're going to define. Next we need to establish how conditional beliefs in a situation are related to the previous situation. Schwering and Lakemeyer (2015, Theorem 5) described this, and we adapt their result.

Lemma 3. *For any IAAT Σ and* now*-regressable formulas ψ_1 and ψ_2, $\Sigma \models \forall a, s.\ \mathbf{Con}(\psi_2 \mid \psi_1, \mathsf{do}(a,s)) \equiv \beta(a,s)$, where $\beta(a,s)$ abbreviates the following:*

$$\begin{aligned}&\big[\mathsf{SF}(a,s) \wedge \mathbf{Con}(\psi_2[\mathsf{do}(a,\mathsf{now})] \mid \\ &\qquad \mathsf{SF}(a,\mathsf{now}) \wedge \mathsf{Poss}(a,\mathsf{now}) \wedge \psi_1[\mathsf{do}(a,\mathsf{now})], s)\big] \vee \\ &\big[\neg\mathsf{SF}(a,s) \wedge \mathbf{Con}(\psi_2[\mathsf{do}(a,\mathsf{now})] \mid \\ &\qquad \neg\mathsf{SF}(a,\mathsf{now}) \wedge \mathsf{Poss}(a,\mathsf{now}) \wedge \psi_1[\mathsf{do}(a,\mathsf{now})], s)\big].\end{aligned}$$

Proof. This can be seen to follow using the SSA for B. $\square$

A class of formula that can be fully regressed can be defined. The *fully regressable formulas* have a number of restrictions, including (among less notable ones), that any term of sort situation mentioned by ϕ (outside beliefs) has the syntactic form $\mathsf{do}(\vec{\alpha}, \mathsf{s}_0)$, and for any expression $\mathbf{Con}(\psi_2 \mid \psi_1, \sigma)$ appearing in ϕ, ψ_1 and ψ_2 are now-regressable.

We can now describe the regression procedure, which is much like that from Schwering and Lakemeyer (2015).

Definition 12. Let $\Gamma = \Gamma_{\mathsf{ssa}} \cup \Gamma_{\mathsf{pre}} \cup \Gamma_{\mathsf{sense}}$ be a set of sentences including SSAs, precondition axioms, and sensing axioms for all the fluents and actions. The (full) regression of ϕ with respect to Γ, where ϕ is fully-regressable (and uses distinct variables from Γ), is written $\mathcal{R}_2^{\Gamma}[\phi]$ and defined case-by-case as follows:

1. ϕ is a formula of the form $\mathbf{Con}(\psi_2 \mid \psi_1, \mathsf{do}(\alpha,\sigma))$. Let $\beta(a,s)$ be as in Lemma 3. Then
$$\mathcal{R}_2^{\Gamma}[\mathbf{Con}(\psi_2 \mid \psi_1, \mathsf{do}(\alpha,\sigma))] = \mathcal{R}_2^{\Gamma}[\beta(\alpha,\sigma)]$$
2. ϕ is a formula of the form $\mathbf{Con}(\psi_2 \mid \psi_1, \mathsf{s}_0)$. Then
$$\mathcal{R}_2^{\Gamma}[\mathbf{Con}(\psi_2 \mid \psi_1, \mathsf{s}_0)] = \mathbf{Con}(\mathcal{R}_1^{\Gamma}[\psi_2] \mid \mathcal{R}_1^{\Gamma}[\psi_1], \mathsf{s}_0)$$
where $\mathcal{R}_1$ is the regression operator from Definition 11.
3. If ϕ is a formula of any other form, it is regressed by $\mathcal{R}_2$ analogously to $\mathcal{R}_1$ in Definition 11.

Proposition 7. *Suppose that Σ is an IAAT. For any fully-regressable formula ϕ (not sharing variables with Σ_{dyn}), $\Sigma \models \forall(\phi \equiv \mathcal{R}_2^{\Sigma_{\mathsf{dyn}}}[\phi])$.*

Proof sketch. This can be proved by induction. The correctness of case (1) can be shown using Lemma 3. For case (2), the result follows from Lemma 2. $\square$

It can be shown that the result of full regression (on a fully-regressable formula) will be a formula where all the situation terms outside of conditional beliefs are s_0, and all the ones inside are now. Further details on full regression, including on how not all axioms from an IAAT are needed to entail a fully regressed sentence, will be included in the first author's PhD dissertation.

7 Discussion and Related Work

The AGM postulates for revision (Alchourrón, Gärdenfors, and Makinson 1985) are widely used for describing rational belief change. Our framework, in modelling belief in much the same way as Shapiro et al. (2011), inherits their results about the relation to AGM. Briefly, Shapiro et al. defined revision by ϕ in terms of a "revision action" that sensed that ϕ was true, and showed that this satisfied a slightly modified version of the AGM postulates. Their postulates refer only to beliefs about formulas that describe the current situation, so don't directly describe how beliefs about dynamics should change. However, in our framework, changes in beliefs about dynamics result from changing beliefs about abnormalities, so there is still some relevance.

Compared to the original AGM postulates, a difference of Shapiro et al.'s version was that because revision was defined in terms of sensing, it was not possible to revise by invalid formulas. Furthermore, if the agent's beliefs become inconsistent, they stay inconsistent (because the accessibility relation can only contract over time), which violates what is sometimes called the "triviality" postulate (however, under some assumptions the agent's beliefs will not become inconsistent). Note that we can still model unreliable sensing results by, for example, instead of sensing whether $P(\mathsf{now})$ is true, sensing whether $[P(\mathsf{now}) \vee \mathsf{Ab}_i(\mathsf{history}(\mathsf{now}), \mathsf{now})]$ is true, which allows for the possibility of a "false positive" result if the relevant abnormality is true.

Existing works on belief revision in the situation calculus have supported having SSAs describing conditional effects and the agent revising its beliefs about when those conditions hold. For instance, Schwering, Lakemeyer, and Pagnucco (2017, §4.2) gave an example where there is an SSA saying that dropping fragile objects breaks them, and the agent revises its beliefs about whether a particular object is fragile. However, the effect of such revisions on what SSAs the agent believes was not discussed (and so neither was regression with SSAs that the agent believes but were not written by the axiomatizer).

Eiter et al. (2007; 2010) describe how a preference order can be defined on propositional transition diagrams by valuing a diagram as the weighted sum of the "query" formulas it entails (Eiter et al. 2007, §4.2). The queries are written in a propositional temporal-logic-like language. It appears this approach could describe preferences on how general of effects action have. However, unlike our work theirs is in a propositional setting without sensing actions.

Delgrande and Levesque (2013) considered actions which could fail (and non-deterministic actions more generally). Their formalization (in the situation calculus) was rather different from ours, as the failure of an action was represented by the agent "intending" to execute one action but actually executing another. Fang and Liu (2013) similarly had an approach, in a multi-agent setting, where agents could be uncertain about what actions had occurred. These works did not discuss having the agent generalize from past failures to reach new conclusions about future action behavior.

A limitation of our approach is that the generalizations the agent can draw from observations have to be specified in advance, as opposed to being determined by some general inductive principles. In contrast, research in *inductive logic programming (ILP)* (Muggleton and de Raedt 1994; De Raedt 2017) has dealt with the problem of inducing general first-order rules given examples. ILP has been applied to learning event calculus theories (Moyle and Muggleton 1997; Katzouris, Artikis, and Paliouras 2019), and also to learning action models in the field of *relational reinforcement learning* (Walker et al. 2007; Rodrigues et al. 2010). On the other hand, we have focused on providing a way for the axiomatizer to precisely and explicitly control the plausibility assigned to different possible dynamics.

Working within the event calculus, Mueller (2006, Chapter 12) used abnormality predicates within descriptions of the environment dynamics, so as to model phenomena like default effects and default events. That was not combined with explicitly modelling belief or belief revision, though.

Britz and Varzinczak (2018) distinguish in an example between two reasons a light might fail to turn on, "either because the light bulb is blown (the current situation is abnormal) or because an overcharge resulted from switching the light (the action behaves abnormally)." In our framework, we would represent both cases as abnormal situations (with the latter using an abnormality fluent that also takes as arguments the action and history, so as to treat overcharges like "one-time exceptions").

8 Conclusion

We have shown how changes of beliefs about SSAs, precondition axioms, and sensing axioms can be modelled using action theories that assign plausibility to situations by counting abnormalities. We described several patterns for writing SSAs that refer to abnormalities, to allow for more or less general changes of belief in response to unexpected observations. We have also shown how beliefs about domain dynamics can be incorporated in regression.

A limitation of our approach is that beliefs about domain dynamics are only changed in response to observations of the present state, as opposed to in response to being given arbitrary facts about dynamics, such as you might read in a physics textbook or a fantasy story. We've also assumed that the agent always knows what actions have occurred. However, it would be natural to revise an agent's beliefs about what actions have occurred (perhaps the reason it's not now holding the cup is that someone else took it, for example). Our approach should be compatible with more general epistemic accessibility relations, such as Shapiro and Pagnucco's (2004), which allowed for exogenous actions that the agent is unaware of. Finally, while we've described our approach in classical logic, it could easily be adapted to a version of $\mathcal{ES}$ (Lakemeyer and Levesque 2011), a modal variant of the situation calculus.

Acknowledgements

The authors gratefully acknowledge funding from the Natural Sciences and Engineering Research Council of Canada and the Canada CIFAR AI Chairs program (Vector Institute).

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2):510–530.

Britz, K., and Varzinczak, I. 2018. From KLM-style conditionals to defeasible modalities, and back. *Journal of Applied Non-Classical Logics* 28(1):92–121.

De Raedt, L. 2017. Inductive logic programming. In Sammut, C., and Webb, G. I., eds., *Encyclopedia of Machine Learning and Data Mining*. Springer US. 648–656.

Delgrande, J. P., and Levesque, H. J. 2013. A formal account of nondeterministic and failed actions. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 861–868.

Eiter, T.; Erdem, E.; Fink, M.; and Senko, J. 2007. Comparing action descriptions based on semantic preferences. *Annals of Mathematics and Artificial Intelligence* 50(3):273–304.

Eiter, T.; Erdem, E.; Fink, M.; and Senko, J. 2010. Updating action domain descriptions. *Artificial Intelligence* 174(15):1172–1221.

Fang, L., and Liu, Y. 2013. Multiagent knowledge and belief change in the situation calculus. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 304–312.

Herzig, A.; Perrussel, L.; and Varzinczak, I. J. 2006. Elaborating domain descriptions. In *ECAI 2006, 17th European Conference on Artificial Intelligence*, 397–401.

Katzouris, N.; Artikis, A.; and Paliouras, G. 2019. Parallel online event calculus learning for complex event recognition. *Future Generation Computer Systems* 94:468–478.

Klassen, T. Q.; McIlraith, S. A.; and Levesque, H. J. 2018. Specifying plausibility levels for iterated belief change in the situation calculus. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, 257–266.

Lakemeyer, G., and Levesque, H. J. 1998. AOL: a logic of acting, sensing, knowing, and only knowing. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 316–327.

Lakemeyer, G., and Levesque, H. J. 2011. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence* 175(1):142–164.

Liberatore, P., and Schaerf, M. 1997. Reducing belief revision to circumscription (and vice versa). *Artificial Intelligence* 93(1):261–296.

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

McCarthy, J. 1980. Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence* 13(1–2):27–39.

Moyle, S., and Muggleton, S. 1997. Learning programs in the event calculus. In *Inductive Logic Programming (ILP-97)*, 205–212.

Mueller, E. T. 2006. *Commonsense Reasoning*. Morgan Kaufmann Publishers.

Muggleton, S., and de Raedt, L. 1994. Inductive logic programming: Theory and methods. *The Journal of Logic Programming* 19–20:629–679.

Pirri, F., and Reiter, R. 1999. Some contributions to the metatheory of the situation calculus. *Journal of the ACM* 46(3):325–361.

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

Rodrigues, C.; Gerard, P.; Rouveirol, C.; and Soldano, H. 2010. Incremental learning of relational action rules. In *2010 Ninth International Conference on Machine Learning and Applications (ICMLA)*, 451–458.

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

Schwering, C., and Lakemeyer, G. 2015. Projection in the epistemic situation calculus with belief conditionals. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 1583–1589.

Schwering, C.; Lakemeyer, G.; and Pagnucco, M. 2017. Belief revision and projection in the epistemic situation calculus. *Artificial Intelligence* 251:62–97.

Shapiro, S., and Pagnucco, M. 2004. Iterated belief change and exogeneous actions in the situation calculus. In *Proceedings of the 16th Eureopean Conference on Artificial Intelligence, ECAI'2004*, 878–882.

Shapiro, S.; Pagnucco, M.; Lespérance, Y.; and Levesque, H. J. 2011. Iterated belief change in the situation calculus. *Artificial Intelligence* 175(1):165–192.

Shapiro, S. 2005. *Specifying and Verifying Multiagent Systems Using the Cognitive Agents Specification Language (CASL)*. Ph.D. Dissertation, University of Toronto.

Sharma, N., and Colomb, R. 1997. Towards an integrated characterisation of model-based diagnosis and configuration through circumscription policies. Technical Report 364, Department of Computer Science, University of Queensland.

Van Zee, M.; Doder, D.; Dastani, M.; and Van Der Torre, L. 2015. AGM revision of beliefs about action and time. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, 3250–3256.

Varzinczak, I. J. 2010. On action theory change. *Journal of Artificial Intelligence Research* 37:189–246.

Walker, T.; Torrey, L.; Shavlik, J. W.; and Maclin, R. 2007. Building relational world models for reinforcement learning. In *Inductive Logic Programming, 17th International Conference, ILP 2007*, 280–291.