

On the Progression of Situation Calculus Basic Action Theories: Resolving a 10-year-old Conjecture

Stavros Vassos and Hector J. Levesque

Department of Computer Science
University of Toronto
Toronto, ON, M5S 3G4, Canada
{stavros, hector}@cs.toronto.edu

Abstract

In a seminal paper, Lin and Reiter introduced a model-theoretic definition for the progression of the initial knowledge base of a basic action theory. This definition comes with a strong negative result, namely that for certain kinds of action theories, first-order logic is not expressive enough to correctly characterize this form of progression, and second-order axioms are necessary. However, Lin and Reiter also considered an alternative definition for progression which is always first-order definable. They conjectured that this alternative definition is incorrect in the sense that the progressed theory is too weak and may sometimes lose information. This conjecture, and the status of first-order definable progression, has remained open since then. In this paper we present two significant results about this alternative definition of progression. First, we prove the Lin and Reiter conjecture by presenting a case where the progressed theory indeed does lose information. Second, we prove that the alternative definition is nonetheless correct for reasoning about a large class of sentences, including some that quantify over situations. In this case the alternative definition is a preferred option due to its simplicity and the fact that it is always first-order.

Introduction

The situation calculus is a logical language that is specially designed for reasoning about action and change (McCarthy & Hayes 1969). A basic action theory is a logical theory in the situation calculus that describes what holds initially in the world as well as how the world evolves under the effects of actions. An example of a basic action theory is one that captures the dynamics of a board game: part of the theory, the initial knowledge base, describes the initial positions of the pieces on the board, and the rest of the theory characterizes the legal moves of the game and the effects (and non-effects) of performing those moves.

A fundamental problem in reasoning about action and change is to determine whether or not some condition holds after a given sequence of actions has been performed. In other words, we start in an initial situation S_0 , we perform a sequence of actions α taking us to a new situation S_α , and we wish to know if the condition holds in S_α . There are in fact two versions of this problem. The special case where

the condition only refers to S_α is called the (simple) *projection problem* (Reiter 2001). For example, we might want to know if a game piece is at a certain location after move α . The more general case is where the condition may refer to situations in the future of S_α . For example, we might want to know if a game piece can ever get to a certain location after move α . This sort of reasoning, which we will call the *generalized projection problem*, is a prerequisite to other forms of reasoning in dynamic domains such as planning and high-level program execution (Reiter 1993).

The simple projection problem can be solved by *regression* or by *progression* (Lin 2007). Roughly speaking, regression involves taking the condition about S_α and transforming it to an equivalent one about S_0 where we can use the initial knowledge base to answer the question; progression, on the other hand, involves replacing the initial knowledge base in the basic action theory by a new knowledge base that captures the facts that hold in S_α .

For the generalized problem, where the condition may refer to the future of S_α , the case is less clear. There is no result for evaluating such conditions based on regression, and it is not clear if there is a practical definition for progression that is logically correct for this reasoning task.

A model-theoretic definition of progression in the situation calculus that does the trick was first proposed by Lin and Reiter (1997). However, their definition, which we call *LR-progression*, comes with a strong negative result: for certain kinds of basic action theories, first-order logic is not expressive enough and second-order logic is needed. Nonetheless, their result did not preclude the possibility of other forms of progression that could still allow us to solve the generalized problem while remaining first-order definable. In particular, one possible candidate for the new knowledge base is the infinite set of all those first-order sentences about S_α that are entailed by the original basic action theory. We will call this second notion of progression *FO-progression*.

While *FO-progression* clearly captures what holds in S_α , it is not clear that it is sufficient to characterize the future of S_α , even in combination with the rest of the basic action theory. Lin and Reiter conjectured that it was too weak. It has been an open problem whether this conjecture is true or false, rendering unclear also the question whether there can be an alternative to *LR-progression* that solves the generalized problem and is first-order definable.

This paper contains two major results. First of all, we prove the Lin and Reiter conjecture: *FO*-progression is indeed too weak for characterizing the future of S_α . We provide a basic action theory and a sentence about the future of S_α that demonstrate this. This result (Theorem 2) further supports the claim by Lin and Reiter that the progression of unrestricted basic action theories cannot be formalized correctly in first-order logic.

The second result is more positive. *FO*-progression was shown by Lin and Reiter (1997) to be correct for the simple projection problem. Here we prove that it is also correct for a much wider class of sentences including sentences of the form “after α , property ϕ will always be true.” This result (Theorem 4) establishes that *FO*-progression is actually more useful than was originally believed.

Situation calculus

The language \mathcal{L} of the situation calculus (McCarthy & Hayes 1969) is first-order with equality and many-sorted, with sorts for actions, situations, and objects (everything else). A *situation* represents a world history as a sequence of actions. The constant S_0 is used to denote the initial situation where no actions have occurred. Sequences of actions are built using the function symbol *do*, such that $do(a, s)$ denotes the successor situation resulting from performing action a in situation s . A *relational fluent* is a predicate whose last argument is a situation, and thus whose value can change from situation to situation. For the scope of this paper we do not allow the language to include *functional fluents* but we note that they can be represented as relational fluents with some extra axioms. Actions need not be executable in all situations, and the predicate $Poss(a, s)$ states that action a is executable in situation s . The language \mathcal{L} also includes the binary predicate symbol \sqsubseteq which provides an ordering on situations. The atom $s \sqsubseteq s'$ means that the action sequence s' can be obtained from the sequence s by performing one or more actions in s . We will typically use the notation $\sigma \sqsubseteq \sigma'$ as a macro for $\sigma \sqsubseteq \sigma' \vee \sigma = \sigma'$.

Often we need to restrict our attention to sentences in \mathcal{L} that refer to a particular situation. For example, the initial knowledge base is a finite set of sentences in \mathcal{L} that do not mention any situation terms except for S_0 . For this purpose, for any situation term σ , we define \mathcal{L}_σ to be the subset of \mathcal{L} that does not mention any other situation terms except for σ , does not mention *Poss*, and where σ is not used by any quantifier (Lin & Reiter 1997). When a formula $\phi(\sigma)$ is in \mathcal{L}_σ we say that it is *uniform in σ* (Reiter 2001). Also, we will use \mathcal{L}^2 to denote the second-order extension of \mathcal{L} that only allows predicate variables that take arguments of sort object. \mathcal{L}_σ^2 then denotes the second-order extension of \mathcal{L}_σ by predicate variables with arguments of sort object.

We will use notation similar to (Gabaldon 2002) and (Reiter 2001) to talk about sequences of actions and situations that are rooted. Let σ be a situation term and δ be a (possibly empty) sequence of action terms $\langle \alpha_1, \dots, \alpha_n \rangle$. We use $do(\delta, \sigma)$ to denote the situation $do(\alpha_n, do(\alpha_{n-1}, \dots do(\alpha_1, \sigma) \dots))$. We say that a situation term σ is *rooted at σ'* iff σ is the term $do(\delta, \sigma')$ for some δ

(in which case, $\sigma' \sqsubseteq \sigma$ clearly holds). Finally, we will use S_α to denote the situation term $do(\alpha, S_0)$.

We will also need to restrict our attention to sentences that refer to σ and the possible futures of σ . We say that τ is *in the future of σ* in ϕ , where ϕ is a rectified sentence in \mathcal{L} , iff

- τ is σ , or
- τ is rooted at some τ' in the future of σ in ϕ , or
- τ is a variable and $\forall \tau'(\tau' \sqsubseteq \tau \supset \beta)$ or $\exists \tau'(\tau' \sqsubseteq \tau \wedge \beta)$ appears in ϕ , where τ' is in the future of σ in ϕ .

We define \mathcal{L}_σ^F as the subset of \mathcal{L} such that for any $\phi \in \mathcal{L}_\sigma^F$ the situation terms in ϕ that appear as arguments of *Poss* or some fluent or the equality predicate are all in the future of σ in ϕ . When a sentence ϕ is in \mathcal{L}_σ^F we say that ϕ is *about the future of σ* .

To see the intuition behind \mathcal{L}_σ^F first note that the sentence $\forall s(S_\alpha \sqsubseteq s \supset \phi(s))$ is in $\mathcal{L}_{S_\alpha}^F$ and expresses that $\phi(s)$ holds in all situations that are rooted at S_α . The recursion allows the sentence $\forall s(S_\alpha \sqsubseteq s \supset \phi(s) \wedge \exists s'(s \sqsubseteq s' \wedge \psi(s')))$ and sentences of this form to be in $\mathcal{L}_{S_\alpha}^F$ as well. In general if a sentence is in \mathcal{L}_σ^F then its truth depends only on situations that are in the future of σ .

Basic action theories

We will be dealing with a specific kind of \mathcal{L} -theory, the so-called basic action theories. The definition that follows is the same as in (Reiter 2001) except that, similarly to (Lake-meyer & Levesque 2004), \mathcal{D}_{ap} consists of a single action precondition axiom for all actions instead of one separate axiom for each action symbol. A *basic action theory* \mathcal{D} has the following form:¹

$$\mathcal{D} = \mathcal{D}_{ap} \cup \mathcal{D}_{ss} \cup \mathcal{D}_{una} \cup \mathcal{D}_{S_0} \cup \mathcal{D}_{fnd}$$

1. \mathcal{D}_{ap} contains a single precondition axiom for all actions of the form $Poss(a, s) \equiv \Pi(a, s)$, where $\Pi(a, s)$ is in \mathcal{L}_s .
2. \mathcal{D}_{ss} is a set of successor state axioms (SSAs), one for each fluent symbol F , of the form $F(\vec{x}, do(a, s)) \equiv \Phi_F(\vec{x}, a, s)$, where $\Phi_F(\vec{x}, a, s)$ is in \mathcal{L}_s . SSAs characterize the conditions under which the fluent has a specific value at situation $do(a, s)$ as a function of situation s .
3. \mathcal{D}_{una} is the set of unique-names axioms for actions: $A(\vec{x}) \neq A'(\vec{y})$, and $A(\vec{x}) = A'(\vec{y}) \supset \vec{x} = \vec{y}$, for each pair of distinct action symbols A and A' .
4. $\mathcal{D}_{S_0} \subseteq \mathcal{L}_{S_0}$ describes the initial situation.
5. \mathcal{D}_{fnd} is a set of domain independent foundational axioms which formally define legal situations and \sqsubseteq .

Regression

An important computational mechanism for reasoning about actions is regression. A formula ϕ is *regressible* iff the following conditions hold (Reiter 2001):²

1. every situation term in ϕ is rooted at S_0 ;
2. ϕ does not quantify over situations;

¹For the sake of readability we will be omitting the leading universal quantifiers.

²Unlike (Reiter 2001), here unrestricted *Poss* atoms are allowed as a consequence of having a single axiom in \mathcal{D}_{ap} .

3. ϕ does not mention the predicate symbol \sqsubseteq and it does not mention any equality atom built on situation terms.

Reiter introduced a regression operator that eliminates *Poss* atoms in favor of their definitions as given by \mathcal{D}_{ap} , and replaces fluent atoms about $do(\alpha, \sigma)$ by logically equivalent expressions about σ as given by the successor state axioms in \mathcal{D}_{ss} . After repeatedly doing this transformation to a regressable sentence ϕ we get a sentence $\mathcal{R}(\phi)$ in \mathcal{L}_{S_0} such that $\mathcal{D} \models \phi \equiv \mathcal{R}(\phi)$. We omit the definition of the regression operation \mathcal{R} and only state the main theorem as it appears in (Pirri & Reiter 1999):

Theorem 1 (Pirri and Reiter). *Let \mathcal{D} be a basic action theory and ϕ be a regressable sentence of \mathcal{L} . Then $\mathcal{R}(\phi)$ is a sentence in \mathcal{L}_{S_0} . Moreover, $\mathcal{D} \models \phi$ iff $\mathcal{D}_{S_0} \cup \mathcal{D}_{una} \models \mathcal{R}(\phi)$.*

Progression

The progression of \mathcal{D} is a new theory \mathcal{D}' that is able to reason correctly about all situations in the future of S_α . It is typical in the literature to define \mathcal{D}_{S_α} as the progression of \mathcal{D}_{S_0} wrt a ground action α and take \mathcal{D}' to be $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{D}_{S_\alpha}$. In other words, we want to replace \mathcal{D}_{S_0} in \mathcal{D} by a suitable set of sentences \mathcal{D}_{S_α} satisfying the following (Reiter 2001):

- (I) Just as \mathcal{D}_{S_0} is a set of sentences in \mathcal{L}_{S_0} , the sentences of the new knowledge base \mathcal{D}_{S_α} should be uniform in S_α .
- (II) \mathcal{D} and \mathcal{D}' should be equivalent wrt how they describe the situations in the future of S_α .

Whenever \mathcal{D}_{S_α} satisfies these conditions we will say that the progression is *correct*.

Lin and Reiter (1997) gave a model theoretic definition for \mathcal{D}_{S_α} that we call *LR-progression*. Here we review the definition that appears in (Reiter 2001) that is more compact.

Definition 1. Let \mathcal{D} be a basic action theory. A set of sentences \mathcal{D}_{S_α} is an *LR-progression* of \mathcal{D}_{S_0} wrt to ground action α iff the following conditions hold:

1. \mathcal{D}_{S_α} is a set of sentences in $\mathcal{L}_{S_\alpha}^2$;
2. $\mathcal{D} \models (\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{D}_{S_\alpha}$;
3. for every model M' of $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{D}_{S_\alpha}$ there is a model M of \mathcal{D} such that the following conditions hold:
 - (a) M and M' have the same domains;
 - (b) M and M' interpret all non-fluent symbols that do not take any arguments of sort situation identically;
 - (c) for every relational fluent F , and every variable assignment μ ,
$$M, \mu \models S_\alpha \sqsubseteq s \wedge F(\vec{x}, s) \text{ iff } M', \mu \models S_\alpha \sqsubseteq s \wedge F(\vec{x}, s);$$
- (d) for every variable assignment μ ,

$$M, \mu \models S_\alpha \sqsubseteq s \wedge Poss(a, s) \text{ iff } M', \mu \models S_\alpha \sqsubseteq s \wedge Poss(a, s).$$

By the conditions 2 and 3 in the definition it follows that for \mathcal{D} and $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{D}_{S_\alpha}$, any model of one is indistinguishable from some model of the other wrt how they interpret the situations in the future of S_α (Reiter 2001). Therefore *LR-progression* satisfies the condition (II). Moreover the condition 1 says that \mathcal{D}_{S_α} is a set of second-order sentences that

are uniform in S_α . Therefore *LR-progression* also satisfies the condition (I) and thus it is correct.

LR-progression comes with a strong negative result, namely that if we restrict \mathcal{D}_{S_α} to be first-order then an *LR-progression* does not always exist (Lin & Reiter 1997). Nonetheless, there is an alternative definition according to which a first-order \mathcal{D}_{S_α} always exists. The idea is to let \mathcal{D}_{S_α} be the infinite set of first-order entailments of \mathcal{D} in \mathcal{L}_{S_α} (Pednault 1987). We call this second notion of progression *FO-progression* and to avoid confusion we will be using \mathcal{F}_{S_α} to refer to it. We introduce the following definition:³

Definition 2. Let \mathcal{D} be a basic action theory and \mathcal{F}_{S_α} be a set of sentences in \mathcal{L}_{S_α} . \mathcal{F}_{S_α} is an *FO-progression* of \mathcal{D} wrt to ground action α iff for all ϕ in \mathcal{L}_{S_α} , $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \phi$ iff $\mathcal{D} \models \phi$.

It is clear that any *FO-progression* satisfies the condition (I). It has been open though whether it also satisfies the condition (II) since it was first formulated as a problem in (Lin & Reiter 1997). In fact, following intuitions and results in (Peppas, Foo, & Williams 1995) Lin and Reiter conjectured that there is a counter example that shows that *FO-progression* does not always satisfy the property (II). Here we state the conjecture in an equivalent way using the terminology that we introduced in this paper.

Conjecture 1 (Lin and Reiter). *There is a basic action theory \mathcal{D} , a ground action α , and a sentence ϕ in $\mathcal{L}_{S_\alpha}^F$ such that $\mathcal{D} \models \phi$ but $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \not\models \phi$, where \mathcal{D}_{S_0} is the initial knowledge base of \mathcal{D} , and \mathcal{F}_{S_α} is an *FO-progression* of \mathcal{D}_{S_0} wrt α .*

FO-progression is not correct for $\mathcal{L}_{S_\alpha}^F$

In this section we give a proof of Conjecture 1 thus resolving the open question whether *FO-progression* is correct. The proof is based on the notion of *unnamed objects* that we will be defining shortly. We will present a basic action theory \mathcal{D}_1 , a ground action A , and a sentence $\phi^* \in \mathcal{L}_{S_\alpha}^F$, for which Conjecture 1 holds. We start by presenting \mathcal{D}_1 and the intuitions behind its definition.

Definition 3. Let \mathcal{L}_1 be the situation calculus language that consists of a binary fluent symbol F , two constant action symbols A, B , a constant object symbol 0 , and a unary function symbol n that takes an argument of sort object. Let \mathcal{D}_1 be the basic action theory of \mathcal{L}_1 that is defined as follows.

- \mathcal{D}_{ap} consists of the sentence $Poss(a, s) \equiv true$.
- \mathcal{D}_{ss} consists of the following sentence:

$$F(x, do(a, s)) \equiv a = A \wedge x = 0 \vee a = B \wedge \neg F(x, s) \wedge \exists y(x = n(y) \wedge F(y, s)). \quad (1)$$

- \mathcal{D}_{una} consists of the sentence $A \neq B$.
- \mathcal{D}_{S_0} consists of the following sentences:

$$\forall a(a = A \vee a = B) \quad (2)$$

³Unlike (Lin & Reiter 1997) and in order to be consistent with the idea of replacing \mathcal{D}_{S_0} with a new set, we insist that it is not \mathcal{F}_{S_α} but $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha}$ that entails the same set of ϕ in \mathcal{L}_{S_α} as \mathcal{D} .

$$\forall x(x \neq 0 \equiv \exists y n(y) = x) \quad (3)$$

$$\forall x \forall y (n(x) = n(y) \supset x = y) \quad (4)$$

$$F(0, S_0) \wedge \forall x (F(x, S_0) \supset F(n(x), S_0)) \quad (5)$$

$$\exists x \neg F(x, S_0) \quad (6)$$

- \mathcal{D}_{ind} is the domain independent foundational axioms.

\mathcal{D}_1 was carefully defined so that all of its models satisfy two properties that we will take advantage in the sequel. Before we state the properties we need to introduce some notation. Observe that each of the ground terms of sort object in \mathcal{L}_1 has the form $n^k(0)$, i.e. it is constructed by a finite number of applications of the function n to the constant 0.

Definition 4. Let GT be the set of all the ground terms of sort object in \mathcal{L} and M be an \mathcal{L} -structure. For every q in the object domain of M we will say that q is *named* iff there is a term $t \in GT$ such that t is interpreted as q , and *unnamed* otherwise. Also, we will say that M is a *term structure* iff all the elements of the object domain of M are named.

The first property of the models of \mathcal{D}_1 is due to \mathcal{D}_{S_0} which can only be satisfied in models that have unnamed objects.

Lemma 1. *No model of \mathcal{D}_1 is a term structure.*

Proof Sketch: The intuition is that the sentence (5) is satisfied only in a structure M where for all named objects q , $M, \mu_q^x \models F(x, S_0)$, while the sentence (6) is satisfied only in a structure M that has an element q' in the object domain such that $M, \mu_{q'}^x \not\models F(x, S_0)$. Therefore $\{(5), (6)\}$ can only be satisfied in a structure that has an unnamed object. ■

The second property is stated in the following lemma:

Lemma 2. *Let M be a model of \mathcal{D}_1 . For every action sequence δ , $M, \mu_q^x \models F(x, do(\delta, S_A))$ iff q is the denotation of $n^k(0)$, where k is*

- the number 0, if the last action in δ is A ;
- the number of B actions that appear after the last occurrence of action A in $\langle A, \delta \rangle$, otherwise.

Proof Sketch: Consider the sentence (1), the successor state axiom for F . First note that in S_A , F is false for all the elements of the object domain except for the denotation of 0. In $do(B, S_A)$ then, F is false for all the elements of the object domain except for the denotation of $n(0)$. This is because $\neg F(x, S_A) \wedge \exists y (x = n(y) \wedge F(y, S_A))$ is true only for $x = n(0)$ and $y = 0$. The formal proof is done by induction on the length of δ using a similar argument and the fact that the sentences (3), (4) in \mathcal{D}_{S_0} ensure that y and $n(y)$ are different objects. ■

We now present the sentence $\phi^* \in \mathcal{L}_{S_A}^F$ that we will be using to prove the conjecture.

Definition 5. Let $\phi^* \text{ be } \exists x \forall s (S_A \sqsubseteq s \supset \neg F(x, s))$.

First we show that by the two properties of \mathcal{D}_1 that we identified earlier the following lemma holds for \mathcal{D}_1 and ϕ^* :

Lemma 3. $\mathcal{D}_1 \models \phi^*$.

Proof Sketch: Consider a model M of \mathcal{D}_1 . By Lemma 2 it follows that for every situation in the future of S_A there can only be named objects for which F is true. By Lemma 1 it follows that there exists at least one unnamed object in the domain. Therefore there is an x such that $F(x, s)$ can never be true in any situation in the future of S_A , which implies that $M \models \exists x \forall s (S_A \sqsubseteq s \supset \neg F(x, s))$. ■

Now we will proceed to show that $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1 \not\models \phi^*$, where \mathcal{F}_1 is an *FO-progression* of \mathcal{D}_{S_0} wrt A .

Definition 6. Let \mathcal{F}_1 be the set $\{\forall x (x = 0 \equiv F(x, S_A)), (2), (3), (4)\}$.

It is not difficult to show that, unlike \mathcal{D}_1 , $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1$ has a term model, in particular one that has the natural numbers as the domain for objects and interprets the constant symbol 0 as the number 0 and the function symbol n as the successor function.

Lemma 4. *There is a model of $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1$ that is a term structure.*

The important point is that even though $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1$ fails to capture a property that \mathcal{D}_1 has, namely that \mathcal{D}_1 is not satisfied in any term structure, the next lemma shows that \mathcal{F}_1 is in fact an *FO-progression* wrt to A .

Lemma 5. \mathcal{F}_1 is an *FO-progression* of the initial knowledge base of \mathcal{D}_1 wrt to ground action A .

The reason is that $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1$ and \mathcal{D}_1 entail the same set of sentences in \mathcal{L}_{S_A} . The formal proof is long and tedious and involves model-theoretic techniques for constructing elementarily equivalent structures, such as the use of the upward Lowenheim-Skolem theorem.

We now show the last lemma we need in order to prove Conjecture 1.

Lemma 6. $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1 \not\models \phi^*$.

Proof Sketch: Consider the term model M of $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1$ that we sketched for Lemma 4, where 0 is interpreted as the number zero and $n^k(0)$ is interpreted as $k \in \mathbb{N}$. Note that the property about F that is proven in Lemma 2 also holds for all the models of $(\mathcal{D}_1 - \mathcal{D}_{S_0}) \cup \mathcal{F}_1$. It follows that for every x in the object domain there is a sequence of actions after which $F(x, s)$ becomes true, which implies that $M \models \forall x \exists s (S_A \sqsubseteq s \wedge F(x, s))$ or equivalently that $M \models \neg \phi^*$. ■

The next theorem establishes that the conjecture by Lin and Reiter is indeed true and thus closes the corresponding open question about the correctness of *FO-progression*.

Theorem 2. *Conjecture 1 holds.*

Proof. By Lemma 3, Lemma 5, and Lemma 6. ■

FO-progression is correct for $\mathcal{L}_{S_\alpha}^O$

In the previous section we showed that in the general case *FO-progression* is not correct. In this section we show that it is nonetheless correct for addressing certain non-trivial reasoning problems. First, we review the result by Lin and Reiter (1997) that *FO-progression* is correct for addressing the (simple) projection problem.

Lemma 7. Let \mathcal{D} be a basic action theory and \mathcal{F}_{S_α} be an FO-progression of \mathcal{D}_{S_0} wrt to ground action α . Then, for any sentence $\phi(s) \in \mathcal{L}_s$ and any situation term σ that is rooted at S_α , $\mathcal{D} \models \phi(\sigma)$ iff $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \phi(\sigma)$.

This result can be extended using the properties of regression. First, we define the set \mathcal{L}_σ^R which is a generalization of the set of regressable sentences.

Definition 7. A formula ϕ is in \mathcal{L}_σ^R iff the following conditions hold:

- every term of sort situation mentioned in ϕ is rooted at σ ;
- ϕ does not quantify over situations;
- ϕ does not mention the predicate symbol \sqsubseteq and it does not mention any equality atom built on situation terms.

$\mathcal{L}_{S_0}^R$ is exactly the set of regressable sentences while \mathcal{L}_σ^R is the subset of it that can also be regressed down to σ . For example $F(do(A, S_0)) \wedge G(do(B, S_0))$ is in $\mathcal{L}_{S_0}^R$ but not in $\mathcal{L}_{S_A}^R$, while $F(do(A, S_0)) \wedge G(do(A, S_0))$ is in both.

We introduce a generalized regression operator \mathcal{R}_σ for formulas in \mathcal{L}_σ^R . This operator works exactly the same as $\mathcal{R}(\phi)$ regressing atoms according to the precondition and successor state axioms in \mathcal{D} , except that it only does so until a sentence uniform in σ is obtained. Like Theorem 1, two similar results can be obtained for \mathcal{R}_σ .

Corollary 1. Let \mathcal{D} be a basic action theory and ϕ be a sentence in \mathcal{L}_σ^R . Then, $\mathcal{R}_\sigma(\phi)$ is a sentence in \mathcal{L}_σ such that $\mathcal{D} \models \phi$ iff $\mathcal{D} \models \mathcal{R}_\sigma(\phi)$.

Corollary 2. Let \mathcal{D} be a basic action theory, \mathcal{F}_{S_α} be a set of sentences in \mathcal{L}_{S_α} , and ϕ be a sentence in $\mathcal{L}_{S_\alpha}^R$, where α is a ground action. Then, $\mathcal{R}_{S_\alpha}(\phi)$ is a sentence in \mathcal{L}_{S_α} such that $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \phi$ iff $\mathcal{F}_{S_\alpha} \cup \mathcal{D}_{una} \models \mathcal{R}_{S_\alpha}(\phi)$.

It is easy then to extend Lemma 7 and show that an FO-progression is correct not only for reasoning about sentences in \mathcal{L}_σ but also for any sentence in $\mathcal{L}_{S_\alpha}^R$.

Lemma 8. Let \mathcal{D} be a basic action theory, \mathcal{F}_{S_α} be an FO-progression of \mathcal{D}_{S_0} wrt to ground action α , and ϕ be a sentence in $\mathcal{L}_{S_\alpha}^R$. Then, $\mathcal{D} \models \phi$ iff $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \phi$.

Proof. By Corollary 1, $\mathcal{R}_{S_\alpha}(\phi)$ is in \mathcal{L}_{S_α} and $\mathcal{D} \models \phi$ iff $\mathcal{D} \models \mathcal{R}_{S_\alpha}(\phi)$. By Lemma 7, this holds iff $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \mathcal{R}_{S_\alpha}(\phi)$. By Corollary 2 and since \mathcal{D} and $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha}$ share the same \mathcal{D}_{ss} , this holds iff $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \phi$. ■

We will now show that an FO-progression is correct for a much wider class of sentences that may also quantify over situations. We extend \mathcal{L}_σ^R as follows.

Definition 8. Let σ be a situation term. \mathcal{L}_σ^Q is the smallest set such that the following conditions hold:

1. if $\phi(s) \in \mathcal{L}_s^R$ then $\phi(\sigma)$ and $\forall s(\sigma \sqsubseteq s \supset \phi(s))$ are in \mathcal{L}_σ^Q ;
2. if $\phi, \psi \in \mathcal{L}_\sigma^Q$ then so are $\neg\phi, \phi \wedge \psi$.

\mathcal{L}_σ^Q is the subset of \mathcal{L}_σ^R that restricts the quantifiers for situation variables to appear only in sub-formulas of the form $\forall s(\sigma \sqsubseteq s \supset \phi(s))$, where $\phi(s)$ does not have free variables other than s . Consider, for instance, the set $\mathcal{L}_{S_\alpha}^Q$. The sentence ϕ^* , that we used to show that FO-progression is not correct for \mathcal{L}_σ^R , is an example of a sentence not in $\mathcal{L}_{S_\alpha}^Q$.

Nonetheless $\mathcal{L}_{S_\alpha}^Q$ is quite large and includes many interesting cases, such as sentences expressing state invariants of the form “after the execution of α it is ensured that $\phi(s)$ will always hold” or “after the execution of α there is no way to achieve $\phi(s)$ ”, as well as boolean combinations of those.

Before we proceed to proving the main result of the section we need the following theorem.

Theorem 3. Let \mathcal{D} be a basic action theory, \mathcal{F}_{S_α} be an FO-progression of \mathcal{D}_{S_0} wrt ground action α , and \mathcal{C} be a set of sentences in \mathcal{L} . Let \mathcal{D}' be $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha}$ and assume that the following holds for $\mathcal{D}, \mathcal{D}', \mathcal{C}$.

Let M be a model of \mathcal{D} and M' be a model of \mathcal{D}' . If it holds that for all $\phi \in \mathcal{L}_{S_\alpha}$, $M \models \phi$ iff $M' \models \phi$, then it also holds that for all $\phi \in \mathcal{C}$, $M \models \phi$ iff $M' \models \phi$.

Then, for all $\phi \in \mathcal{C}$, $\mathcal{D} \models \phi$ iff $\mathcal{D}' \models \phi$.

This theorem specifies a method for proving that an FO-progression is correct for a class of sentences \mathcal{C} . Essentially it reduces the question about *entailment* (the two theories entail the same set of sentences in \mathcal{C} , provided they entail the same set of sentences in \mathcal{L}_{S_α}) to a simpler question about *satisfaction* (any two models of the theories satisfy the same set of sentences in \mathcal{C} , provided they satisfy the same set of sentences in \mathcal{L}_{S_α}). The proof is long and relies on the foundational result that $\mathcal{D} - \mathcal{D}_{fnd}$ is equivalent to \mathcal{D} wrt the entailment of sentences uniform in some σ (Lin & Reiter 1997) and the Compactness Theorem of first-order logic.

So, in order to show that an FO-progression is correct for $\mathcal{L}_{S_\alpha}^Q$ it suffices to prove the following lemma:

Lemma 9. Let \mathcal{D} be a basic action theory and \mathcal{F}_{S_α} be an FO-progression of \mathcal{D}_{S_0} wrt ground action α . Let M be a model of \mathcal{D} and M' be a model of $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha}$ such that for all $\phi \in \mathcal{L}_{S_\alpha}$, $M \models \phi$ iff $M' \models \phi$. Then, for all $\phi \in \mathcal{L}_{S_\alpha}^Q$, $M \models \phi$ iff $M' \models \phi$.

Proof. By induction on the construction of $\phi \in \mathcal{L}_{S_\alpha}^Q$. The only interesting part is the base of the induction where we have two cases: i) ϕ is in $\mathcal{L}_{S_\alpha}^R$ or ii) ϕ is $\forall s(S_\alpha \sqsubseteq s \supset \psi(s))$, where $\psi(s)$ is in \mathcal{L}_s^R . Case i) follows from Lemma 8. For case ii) we will use a trick to deal with the quantification over situations to reduce it to case i). We prove the (\Rightarrow) direction by contradiction and the other one follows similarly.

Let $M \models \forall s(S_\alpha \sqsubseteq s \supset \psi(s))$ where $\psi(s)$ is in \mathcal{L}_s^R and suppose that $M' \not\models \forall s(S_\alpha \sqsubseteq s \supset \psi(s))$. It follows that there is an element q of the situation domain such that $M', \mu_q^s \models S_\alpha \sqsubseteq s \wedge \neg\psi(s)$. Since M' satisfies the foundational axioms \mathcal{D}_{fnd} , this element q is reachable from the denotation of S_α by a finite number of applications of the function do . In particular let e_1, \dots, e_n be elements of the action domain such that $do^{M'}((e_1, \dots, e_n), S_\alpha^{M'}) = q$. It follows that $M' \models \gamma$, where γ is the following sentence:

$$\exists a_1 \dots \exists a_n \neg\psi(do(\langle a_1, \dots, a_n \rangle, S_\alpha)).$$

By the hypothesis $\psi(s)$ is in \mathcal{L}_s^R and so γ is in $\mathcal{L}_{S_\alpha}^R$. By case i) it follows that $M \models \gamma$. Since M satisfies the foundational axioms \mathcal{D}_{fnd} , it follows that $M \models \exists s(S_\alpha \sqsubseteq s \wedge \neg\psi(s))$ or equivalently $M \not\models \forall s(S_\alpha \sqsubseteq s \supset \psi(s))$ which is a contradiction. Thus our assumption is wrong and $M' \models \forall s(S_\alpha \sqsubseteq s \supset \psi(s))$. ■

Now we are ready to state and prove the main theorem of the section, namely that *FO*-progression is correct for $\mathcal{L}_{S_\alpha}^Q$.

Theorem 4. *Let \mathcal{D} be a basic action theory, \mathcal{F}_{S_α} be an *FO*-progression of \mathcal{D}_{S_0} wrt to ground action α , and ϕ be a sentence in $\mathcal{L}_{S_\alpha}^Q$. Then $\mathcal{D} \models \phi$ iff $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha} \models \phi$.*

Proof. By Theorem 3 and Lemma 9. ■

Related and Future Work

Other people have looked into definitions for the progression of basic action theories under different assumptions. Liu and Levesque (2005) study the special case where the domain of discourse is fixed to a countable set of named objects, Claßen and Lakemeyer (2006) focus on the \mathcal{ES} variant of the situation calculus, and Thielscher (1999) defines a dual representation for the basic action theories based on state update axioms that explicitly define the direct effects of each action. In order for the progression to remain in first-order a special form is assumed for the structure of the basic action theory. This is in contrast to our line of work where we identify a class of sentences for which a first-order progression of an unrestricted theory is correct. A similar but much weaker result is due to Shirazi and Amir (2005). Shirazi and Amir prove that for those cases that progression is first-order definable, their variant of progression is correct for answering queries uniform in some situation term.

With respect to the proof for Conjecture 1, the notion of unnamed objects is also used in a different way in (Lin & Reiter 1997) to show that a first-order *LR*-progression does not always exist. Also, Theorem 4 shares intuitions with a result that appears in (Savelli 2006): whenever a basic action theory entails that there exists a situation satisfying a condition, at least one such situation must be found within a *predetermined distance* from the initial situation. The proof of this result relies on two ideas that we also used, namely the trick we used in the proof for Lemma 9 to deal with the quantification over situations, and the use of the Compactness Theorem for the proof of Theorem 3. The main difference is that we have separated the use of each of the ideas in such way that Theorem 3 and a different trick can be used to prove a result about progression under different assumptions. Finally, we note that Theorem 4 and the result in (Savelli 2006) also imply corresponding results about regression that we intend to investigate in future work.

Our future work also focuses on the following. First, we intend to investigate when a finite *FO*-progression can be found. Note that an *FO*-progression is not necessarily an infinite set. The case of \mathcal{F}_1 for the basic action \mathcal{D}_1 is such an example. Second, we showed that an *FO*-progression is not correct unless we restrict our attention to a subset of the sentences about the future of S_α , namely the set $\mathcal{L}_{S_\alpha}^Q$. Nonetheless, depending on what kind of a basic action theory we consider, there are cases where *FO*-progression is correct for all the sentences about the future of S_α . As a special case note that when \mathcal{D}_{S_0} is empty, *FO*-progression is correct.⁴ We want to identify practical cases where *FO*-progression can safely be used for all reasoning tasks.

⁴Let \mathcal{F}_{S_α} be the empty set. \mathcal{D} and $(\mathcal{D} - \mathcal{D}_{S_0}) \cup \mathcal{F}_{S_\alpha}$ coincide.

Conclusions

In this paper we presented two significant results about the progression of basic action theories. First, we proved a conjecture by Lin and Reiter and showed that an alternative definition for a progressed theory loses information. Second, we proved that this alternative definition is nonetheless correct for reasoning about a large class of sentences, including some that quantify over situations. Moreover, we provided a general method for proving the correctness of the alternative definition that can be used under different assumptions. We conclude that, under practical conditions, the alternative definition is a preferred option due to its simplicity and the fact that it is always first-order.

References

- Claßen, J., and Lakemeyer, G. 2006. A semantics for ADL as progression in the situation calculus. In *Proc. NRAC-06*.
- Gabalton, A. 2002. Non-markovian control in the situation calculus. In *Proc. AAAI-02*, 519–524.
- Lakemeyer, G., and Levesque, H. J. 2004. Situations, si! situation terms, no! In *Proceedings of KR-2004*.
- Lin, F., and Reiter, R. 1997. How to progress a database. *Artificial Intelligence* 92(1-2):131–167.
- Lin, F. 2007. Chapter: Situation calculus. In Lifschitz, V.; van Harmelen, F.; and Porter, F., eds., *Handbook of Knowledge Representation*. Elsevier.
- Liu, Y., and Levesque, H. J. 2005. Tractable reasoning with incomplete first-order knowledge in dynamic systems with context-dependent actions. In *Proc. IJCAI-05*.
- McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence* 4:463–502.
- Pednault, E. P. D. 1987. *Toward a mathematical theory of plan synthesis*. Ph.D. Dissertation, Stanford University, Stanford, CA, USA.
- Peppas, P.; Foo, N.; and Williams, M. 1995. On the expressibility of propositions. *Logique et Analyse* 139–149:251–272.
- Pirri, F., and Reiter, R. 1999. Some contributions to the metatheory of the situation calculus. *Journal of the ACM* 46(3):261–325.
- Reiter, R. 1993. Proving properties of states in the situation calculus. *Artificial Intelligence* 64(2):337–351.
- Reiter, R. 2001. *Knowledge in Action. Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press.
- Savelli, F. 2006. Existential assertions and quantum levels on the tree of the situation calculus. *Artificial Intelligence* 170(6):643–652.
- Shirazi, A., and Amir, E. 2005. First-order logical filtering. In *Proc. of IJCAI-05*, 589–595.
- Thielscher, M. 1999. From situation calculus to fluent calculus: State update axioms as a solution to the inferential frame problem. *Artificial Intelligence* 111(1-2):277–299.