



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

How to Progress Beliefs in Continuous Domains

Citation for published version:

Belle, V & Levesque, HJ 2014, How to Progress Beliefs in Continuous Domains. in *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20-24, 2014*. AAAI Press, pp. 438-447.
<<http://www.aaai.org/ocs/index.php/KR/KR14/paper/view/7998>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20-24, 2014

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



How to Progress Beliefs in Continuous Domains

Vaishak Belle and Hector J. Levesque

Dept. of Computer Science
University of Toronto
Toronto, Ontario M5S 3H5, Canada
{vaishak, hector}@cs.toronto.edu

Abstract

When Lin and Reiter introduced the progression of basic action theories in the situation calculus, they were essentially motivated by long-lived robotic agents functioning over thousands of actions. However, their account does not deal with probabilistic uncertainty about the initial situation nor with effector or sensor noise, as often needed in robotic applications. In this paper, we obtain results on how to progress continuous degrees of belief against continuous effector and sensor noise in a semantically correct fashion. Most significantly, and perhaps surprisingly, we identify conditions under which our account is not only as efficient as the filtering mechanisms commonly used in robotics, but considerably more general.

Introduction

Reiter's (2001) reconsideration of the situation calculus (McCarthy and Hayes 1969) has proven enormously useful for the design of logical agents, essentially paving the way for *cognitive robotics* (Lakemeyer and Levesque 2007). Among other things, it incorporates a simple monotonic solution to the frame problem, leading Reiter to define the notion of *regression* for basic action theories (Waldinger 1977). But for long-lived agents like robots, Lin and Reiter (1997) argue that the notion of *progression*, that of continually updating the current view of the state of the world, is perhaps better suited. They show that progression is always second-order definable, and in general, it appears that second-order logic is unavoidable (Vassos and Levesque 2008). However, Lin and Reiter also identify some first-order definable cases by syntactically restricting situation calculus basic action theories, and since then, a number of other special cases have been studied (Liu and Lakemeyer 2009).

While Lin and Reiter intended their work to be used on robots, one criticism leveled at their work, and indeed at much of the work in cognitive robotics, is that the theory is far removed from the kind of continuous uncertainty and noise seen in typical robotic applications. What exactly filtering mechanisms (such as Kalman filters (Thrun, Burgard, and Fox 2005)) have to do with Lin and Reiter's progression has gone unanswered, although it has long been suspected that the two are related.

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This paper remedies this situation. Building on a recent extension to the situation calculus for dealing with continuous uncertainty by Belle and Levesque (2013a) (henceforth BL), we investigate progression in continuous domains. In this paper, we make the following contributions:

1. We introduce a property of basic action theories called *invertibility*, closely related to invertible functions in real analysis (Trench 2003). We identify syntactic restrictions on basic action theories that guarantee invertibility.
2. For our central result, we show a first-order *progression* of continuous degrees of belief against continuous noise in effectors and sensors for action theories that are invertible.
3. Finally, we prove that this account of progression is *efficient*. Perhaps surprisingly, under the additional assumption of context-completeness (Liu and Levesque 2005), it is as efficient as commonly used filtering mechanisms in the robotics literature, while considerably more general.

For the structure of this paper, we begin by introducing the formalism, followed by results along the lines of the above three points and conclude after discussion and related work.

The Situation Calculus

The language \mathcal{L} of the situation calculus is a many-sorted dialect of the predicate calculus, with sorts for *physical actions*, *sensing actions*, *situations* and *objects*, which is a catch-all sort for everything else. A situation represents a history as a sequence of actions. A set of initial situations correspond to the ways the world might be before any action; successor situations are obtained as a result of doing actions. We use α to denote terms of the action sort, and so the term $do(\alpha, s)$ denotes the unique situation obtained on doing α in s . The term $do(\vec{\alpha}, s)$, where $\vec{\alpha}$ is the sequence $\alpha_1 \cdots \alpha_n$, abbreviates $do(\alpha_n, do(\dots, do(\alpha_1, s) \dots))$. Initial situations, then, are those without a predecessor:

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

We let the constant S_0 denote the actual initial situation and use the variable ι to range over initial situations only.

To reason about changing properties, \mathcal{L} is assumed to include *fluents*, whose last argument is always a situation term. Without loss of generality, we restrict ourselves to functional

fluents. In addition, we assume that f_1, \dots, f_k are all the fluents in \mathcal{L} , and that these take no arguments other than a single situation term.¹ Note that this is not a propositional theory in that we allow the *values* of the fluents f_1, \dots, f_k to range over any set. As we are primarily interested in continuous domains, we will make the simplifying assumption that fluents are all real-valued.

Following (Reiter 2001), we call a formula ϕ *uniform* in a situation term σ if σ is the only situation term in ϕ and *rooted* in σ if every situation term in ϕ is of the form $do(\vec{a}, \sigma)$ for some action sequence \vec{a} .² We often suppress the situation argument in a formula ϕ , or use a distinguished variable *now*. Either way, $\phi[\sigma]$ is used to denote the formula with that variable replaced by the situation term σ .

Basic Action Theory

Again following Reiter, we model dynamic domains by means of a set of sentences \mathcal{D} called a *basic action theory*, consisting of:³

- sentences \mathcal{D}_0 describing what is true initially;
- precondition axioms describing the executability of actions using a distinguished symbol *Poss*;⁴
- successor state axioms that describe the change to fluents after actions, of the form $f(do(a, s)) = u \equiv \gamma_f(a, u)[s]$;
- domain-agnostic foundational axioms and the unique name axioms for actions, the details of which need not concern us here (Reiter 2001).

We often lump the components in \mathcal{D} other than \mathcal{D}_0 as Σ . Properties of the formalism follow as entailments of \mathcal{D} , for which standard Tarskian models suffice. We will however be assuming that models assign the usual interpretations to $=, <, >, 0, 1, +, \times, /, -, e, \pi$ and x^y (exponentiation).

The language \mathcal{L} used in this paper is the situation calculus as characterized in (Reiter 2001). Nonetheless, it is very useful to have certain abbreviations that will macro expand to \mathcal{L} -formulas. These abbreviations are used as logical terms, that is, as arguments to functions or predicates. If E is such an expression, its expansion is characterized by a definition of the form $E = u \doteq \phi(u)$, where u is a variable, and $\phi(u)$ is an \mathcal{L} -formula with u free. This should be interpreted as follows. Let $\rho(u)$ be any atomic formula with u free. Then the expression $\rho(E)$ should be understood as standing for the formula $\exists u(\phi(u) \wedge \rho(u))$. To see this in action, consider the definition of the *conditional-if* expression: $u = \text{IF } \exists z \phi \text{ THEN } t_1 \text{ ELSE } t_2$ defined as $(\exists z \phi \supset \forall z(\phi \supset u = t_1)) \wedge (\neg \exists z \phi \supset u = t_2)$. This allows us to write, for example,

$$P(a, \text{IF } Q(f, 3) \text{ THEN } 4 \text{ ELSE } 5, b)$$

¹See BL for a discussion of this particular limitation.

²Formulas that are rooted in σ are about σ and its future. More general formulas about the future of a situation are considered in (Vassos and Levesque 2008), but not here.

³Free variables, here and everywhere else, are assumed to be implicitly quantified from the outside.

⁴For simplicity, we will say no more about *Poss* here, and assume that actions are always executable.

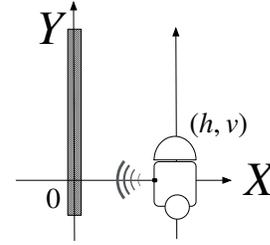


Figure 1: Robot moving parallel to the wall.

meaning $(Q(f, 3) \supset P(a, 4, b)) \wedge (\neg Q(f, 3) \supset P(a, 5, b))$. As a consequence of this convention, we will be able to write Reiter's successor state axioms as equality-expressions, and use the RHS as terms that can be substituted in formulas.

For example, imagine a robot moving in a 2-dimensional plane, parallel to a wall, as shown in Figure 1. Let its coordinate be given by (h, v) , where h and v are real-valued fluents. Consider an action $mv(z)$ that moves the robot away from the X -axis by z units. Its effect on v might be described by:

$$v(do(a, s)) = u \equiv \exists z(a = mv(z) \wedge u = v(s) + z) \vee u = v(s) \wedge \neg \exists z(a = mv(z)).$$

This says that $mv(z)$ is the only action affecting v , thereby incorporating a solution to the frame problem (Reiter 2001). We would now equivalently write this as:

$$v(do(a, s)) = \text{IF } \exists z(a = mv(z)) \text{ THEN } v(s) + z \text{ ELSE } v(s). \quad (1)$$

Thus, successor state axioms are henceforth assumed to be of the form $f(do(a, s)) = E_f(a)[s]$.

We will also use expressions $\int_x t$ and $\sum_x t$ for integration and summation, which are abbreviations defined using second-order logic in BL and (Bacchus, Halpern, and Levesque 1999). Finally, we will often refer to common density functions, such as \mathcal{N} and \mathcal{U} (Gaussian and uniform respectively), which are simply abbreviations for \mathcal{L} -terms built from arithmetic symbols and reals.

Noise and Likelihood

Building on (Bacchus, Halpern, and Levesque 1999), our treatment of noisy acting and sensing will benefit from Reiter's solution to the frame problem. It will involve a new distinguished symbol l for likelihood, where the modeler provides l -axioms in \mathcal{D} , one per action $A(\vec{x})$, of the form:

$$l(A(\vec{x}), s) = \Upsilon_A(\vec{x})[s].$$

Let us illustrate its usage with noisy sensors. (Noisy actions are discussed later.) The concern here is that the value read from a sensor may differ from the true value, but in some reasonable way. Imagine a sonar sensor aimed at the wall, and so its reading would correspond to the true value of fluent h . If noisy, one might assume that the likelihood of a reading z is obtained from a normal curve whose mean is h . We then say it has additive Gaussian noise, and this is modeled as:

$$l(\text{sonar}(z), s) = \mathcal{N}(z; h, 4)[s] \quad (2)$$

which gives the sensor's reading a variance of 4. Note that the l -axioms are more expressive than those typically seen in probabilistic formalisms, and can be context-dependent. For example, to model a sensor with systematic bias at subzero temperatures, we might have

$$l(\text{sonar}(z), s) = \text{IF } \text{temp}(s) > 0 \\ \text{THEN } \mathcal{N}(z; h, 1) \text{ ELSE } \mathcal{N}(z; h + 2, 1).$$

Degrees of Belief

Probabilistic uncertainty about the initial situation and changing degrees of belief are captured using a single distinguished symbol p . The p fluent can be seen as a numeric variant of the K fluent in the Scherl and Levesque scheme (2003). Intuitively, $p(s', s)$ gives us the probabilistic *density* that the agent attributes to s' when at s . In fully discrete domains, p is the *mass* given to s' .

As part of \mathcal{D}_0 , the modeler would provide a characterization of the density attributed by the agent to the initial situations. For example,

$$p(\iota, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; 0, 1)[\iota] \quad (3)$$

says that the agent does not know the initial values of h and v , but thinks of them as drawn independently from a uniform distribution on $[2, 12]$, and from a standard normal. Since p is just like any other fluent, the framework is more expressive than many probabilistic formalisms. For example,

$$\forall \iota (p(\iota, S_0) = \mathcal{U}(h; 2, 3)[\iota]) \\ \vee \\ \forall \iota (p(\iota, S_0) = \mathcal{U}(h; 10, 20)[\iota]) \quad (4)$$

says that the agent takes h to be uniformly distributed on $[2, 3]$ or on $[10, 20]$, but without saying which.

To give p the intended properties, the following non-negative constraint is assumed to be included in \mathcal{D} (Bacchus, Halpern, and Levesque 1999):

$$\forall \iota, s. p(s, \iota) \geq 0 \wedge (p(s, \iota) > 0 \supset \text{Init}(s)). \quad (\text{P1})$$

This says that p values are never negative, and that they are non-zero only for initial situations. While this is a stipulation about initial situations, the following successor state axiom for p ensures that the nonnegative constraint holds everywhere:

$$p(s', \text{do}(a, s)) = \\ \text{IF } \exists s''. s' = \text{do}(a, s'') \wedge \text{Poss}(a, s'') \\ \text{THEN } p(s'', s) \times l(a, s'') \text{ ELSE } 0 \quad (\text{P2})$$

That is, the density of a situation s' relative to $\text{do}(a, s)$ is the density of its predecessor s'' times the likelihood of a contingent on the successful execution of a at s'' . One consequence of (P1) and (P2) is that $(p(s', s) > 0)$ will be true only when s' and s share the same history of actions.

As discussed in BL, to deal with continuity, the following axiom explicating a precise space of initial situations is needed (Baker 1989; Levesque, Pirri, and Reiter 1998):

$$[\forall \vec{x} \exists \iota \bigwedge f_j(\iota) = x_j] \wedge [\forall \iota, \iota'. \bigwedge f_j(\iota) = f_j(\iota') \supset \iota = \iota'] \quad (\text{P3})$$

This states that for every combination of fluents values, there is a unique initial state where the fluents have these values.

With l and p axioms specified in \mathcal{D} , the *degree of belief* in a formula is defined as follows:⁵

Definition 1: Let ϕ be a formula and σ be a situation term, both rooted in *now*. Then $\text{Bel}(\phi, \sigma[S_0])$ is defined as

$$\frac{1}{\eta} \int_{\vec{x}} \text{IF } \exists \iota \bigwedge f_j(\iota) = x_j \wedge \phi[\sigma[\iota]] \\ \text{THEN } p(\sigma[\iota], \sigma[S_0]) \text{ ELSE } 0$$

where η is the *normalization* factor and is the same expression as the numerator but with ϕ replaced by *true*.

This says that belief is the result of integrating the density function p over all possible initial values for the real-valued fluents, and hence by (P3), over all possible initial situations. Note that this definition of belief is restricted to cases where the situation term is rooted in S_0 .

To summarize, all that was needed to reason about degrees of belief and continuity in the situation calculus were the following components:

- \mathcal{D}_0 as usual (possibly mentioning p);
- the action components and foundational axioms as before, including the l -axioms, and p -constraints (P1), (P2) and (P3) (lumped together as Σ);
- abbreviations for *Bel*, integrals, and summation.

It is easy to see that this framework generalizes the categorical knowledge treatment of (Scherl and Levesque 2003). It can also be shown that the belief change mechanism exhibited by the framework subsumes Bayesian conditioning (Pearl 1988). Finally, the scheme is also applicable to discrete domains: for a discrete fluent f , one would simply replace $\int_{\vec{x}}$ with \sum_x ranging over the possible values of f .

Noisy actions

The account of belief in BL left noisy actions for future work. The idea behind noisy actions is that an agent might attempt a physical move of 3 units, say, but as a result of the limited accuracy of effectors, actually move 3.094 units. Thus, unlike sensors, where the reading is nondeterministic, *observable*, but does not affect fluents, the outcome of noisy actions is nondeterministic, *unobservable* and changes fluent properties. Of course, when attempting to move 3 units, the agent knows that an actual move by 3.094 units is much more likely than 30.94 units.

We suggest a simple specification of noisy actions where instead of an action like $mv(x)$, we use an action $mv'(x, y)$ where x is the intended motion (3 units) known to the agent and y is the actual motion (perhaps 3.094 units) unknown to the agent. The successor state axiom for h would be written to reflect the fact that its value is changed by the y value. Unlike noise-free actions where the likelihood is always 1, noisy actions would have non-trivial l -axioms like this one:

$$l(mv'(x, y), s) = \mathcal{N}(y; x, 1 + \text{wet})[s].$$

⁵This is a minor reworking of the definition given in BL, but as it turns out, it has the advantage of working seamlessly with noisy actions as well.

This is a context-dependent axiom that says that the actual value moved is normally distributed around the intended value, but with a variance that depends on the *wet* fluent (a measure of how wet the floor is).

To capture the fact that the agent need not know the actual amount moved, we use *fluents arguments* to the action. For example, to find out the degree of belief in ϕ after a noisy move action, instead of asking for the value of

$$Bel(\phi, do(mv'(3, 3.094), S_0)),$$

we instead ask for the value of something like

$$Bel(\phi, do(mv'(3, g(S_0)), S_0)),$$

where g is a fluent whose value is not known, that stands for the actual amount moved. The fluent g here represents nature's choice of a value, and as it is unaffected by any action, its successor state axiom should be $g(do(a, s)) = g(s)$. The fluent will have different values in different initial states, of course, and so the likelihood of the corresponding mv' action will also vary, as will the density given to the successor situation by (P2). The definition of belief will then integrate over all possible values of g as appropriate.

In general, a fluent can be used like this to handle any argument to an action whose value is not known to the agent. We believe this account is much simpler than the version of noisy actions given in (Bacchus, Halpern, and Levesque 1999) which relied on Golog.

Invertible Action Theories

What we are after is an account of progression for belief in the presence of continuous noise. However, subtleties arise with the p fluent even in simple domains. For example, imagine a noise-free action $towall(z)$ that moves the robot towards the wall but stopping when the robot hits the wall:

$$h(do(a, s)) = \text{IF } \exists z(a = towall(z)) \\ \text{THEN } \max(0, h(s) - z) \text{ ELSE } h(s). \quad (5)$$

If the robot were to begin with (3) and perform the action $towall(4)$, beliefs about the new value of h become much more complex. Roughly, those points where $h \in [2, 4]$ initially are mapped to a single point $h = 0$ that should then obtain a probability *mass* of .2, while the other points retain their initial *density* of .1. In effect, a probability density on h is transformed into a mixed density / distribution (and the (P3) assumption no longer holds). BL deal with this issue using *regression*: beliefs are regressed to the initial state where (P3) does hold, and all the actual belief calculations can be done in the initial state (Belle and Levesque 2013b).

In this paper, we develop a logical theory of progression for basic action theories where such mixed distributions do not arise. In this section, we provide a simple definition to support this, and then discuss general syntactic restrictions that satisfy this requirement.

Definition 2: Given a basic action theory \mathcal{D} , a fluent f is said to be *invertible* if there is an expression $H_f(a)$ uniform in *now* such that $\mathcal{D} \models f(s) = H_f(a)[do(a, s)]$. We say that \mathcal{D} is invertible if every fluent in the theory is invertible.

Intuitively, a fluent is invertible when we can find a dual formulation of its successor state axiom, that is, where we can characterize the *predecessor* value of a fluent in terms of its current value.

There are three syntactic conditions on a basic action theory \mathcal{D} that are sufficient to guarantee its invertibility:

- i. There is an ordering on fluents such that all the fluents that appear in $E_f(a)$ other than f are earlier in the ordering.
- ii. Any situation term in $E_f(a)$ appears as an argument to one of the fluents.
- iii. The mapping from the value of $f(s)$ to the value of $f(do(a, s))$ given by $E_f(a)$ is bijective (Trench 2003). (This is understood in the usual set-theoretic sense.)

Before considering some examples, here is the result:

Theorem 3: *If a basic action theory satisfies (i), (ii) and (iii) above, then it is invertible.*

Proof sketch: The proof is by induction on the ordering given by (i). By (iii), we can take $f(do(a, s)) = E_f(a)[s]$ and solve for $f(s)$, obtaining an equation $f(s) = H$, where H mentions $f(do(a, s))$ and possibly other fluents $f'(s)$ that appear earlier in the ordering. By induction, each $f'(s)$ in H can be replaced by $H_{f'}(a)[do(a, s)]$. By (ii), the result will then be uniform in $do(a, s)$, and thus we obtain a formula $H_f(a)$ where $\mathcal{D} \models f(s) = H_f(a)[do(a, s)]$ as desired. ■

Example 4: Consider (1). This trivially satisfies (i) and (ii). The mapping from $v(s)$ to $v(do(a, s))$ is bijective and so (iii) is satisfied also. (In general, any $E_f(a)$ that is restricted to addition or multiplication by constants will satisfy (iii).) So the fluent is invertible and we have $v(s) = H_v(a)[do(a, s)]$, where $H_v(a) = \text{IF } \exists z(a = mv(z)) \text{ THEN } v - z \text{ ELSE } v$.

Example 5: Consider (5). Here the mapping is not bijective because of the *max* function and the fluent h is not invertible. If $h(do(a, s)) = 0$ where $\alpha = towall(4)$, then the value of $h(s)$ cannot be determined and can be anything less than 4.

Example 6: Consider a successor state axiom like this:

$$v(do(a, s)) = \text{IF } \exists z(a = mv(z)) \text{ THEN } (v(s))^z \text{ ELSE } v(s).$$

For $\alpha = mv(2)$, we obtain a squaring function, which is not bijective. Indeed, from $v(do(\alpha, s)) = 9$, one cannot determine whether $v(s)$ was -3 or 3. So the fluent is not invertible.

Example 7: Consider this successor state axiom:

$$v(do(a, s)) = \text{IF } \exists z(a = rapid(z)) \wedge wax(s) \leq .3 \\ \text{THEN } v(s) + (fuel(s))^z \text{ ELSE } v(s).$$

Suppose further:

$$fuel(do(a, s)) = \text{IF } \exists z(a = fill(z)) \text{ THEN } z \cdot fuel(s) \text{ ELSE } fuel(s). \\ wax(do(a, s)) = \text{IF } a = clean \text{ THEN } wax(s)/10 \text{ ELSE } wax(s).$$

Here, the fuel tank and the wax on the floor determine the distance moved, and one can clean the floor wax and fill the tank. This theory is invertible, and $H_v(a)$ is given by

$$H_v(a) = \text{IF } \exists z(a = rapid(z)) \wedge wax \leq .3 \\ \text{THEN } v - fuel^z \text{ ELSE } v.$$

That is, because $rapid(z)$ does not affect wax and $fuel$, we simply invert the successor state axiom for v and relativize everything to $do(a, s)$. If (say) the action $rapid(z)$ also affected $fuel$, by the ordering in (i), we would first obtain the H -expression for $fuel$ and use it in the H -expression for v .

Finally, note that the bijection property does not prevent us from using non-bijective functions, such as squares, in the successor state axiom of v , provided that these only apply to the other fluents. (The remaining fluents essentially behave as constants at any given situation.) In our experience, many commonly occurring successor state axioms are invertible.

Progression

In this section we propose a definition of progression that applies to any invertible basic action theory. Note that the definition of invertibility imposes no constraint on \mathcal{D}_0 . So the definition in this section is general in that the p may appear in \mathcal{D}_0 in an unrestricted way, such as in (4). Given such a theory $\mathcal{D}_0 \cup \Sigma$ and a ground action α , we define a transformation \mathcal{D}'_0 such that $\mathcal{D}'_0 \cup \Sigma$ agrees with $\mathcal{D}_0 \cup \Sigma$ on the future of α . Then, in the next section, we will consider how \mathcal{D}'_0 grows as a result of this progression.

Classical Progression

Let us first consider the simpler case of progression for a \mathcal{D}_0 that does not mention the p fluent (and the quantification over initial situations that comes with it), and so where \mathcal{D}_0 is uniform in S_0 . In this case, because we are assuming a finite set of nullary fluents, any basic action theory can be shown to be *local-effect* (Liu and Lakemeyer 2009), where progression is first-order definable. The new theory is computed by appealing to the notion of *forgetting* (Lin and Reiter 1994). If the basic action theory is invertible, however, the progression can *also* be defined in another way. Let \mathcal{D}'_0 be \mathcal{D}_0 but with any $f(S_0)$ term in it replaced by $H_f(\alpha)[S_0]$.

Theorem 8 *Let $\mathcal{D}_0 \cup \Sigma$ be any invertible basic action theory not mentioning p and α any ground action. Then for any \mathcal{L} -formula ϕ rooted in now*

$$\mathcal{D}_0 \cup \Sigma \models \phi[do(\alpha, S_0)] \text{ iff } \mathcal{D}'_0 \cup \Sigma \models \phi[S_0].$$

Example 9: Consider (1), and the $H_v(a)$ from Example 4. Suppose $\mathcal{D}_0 = \{v(S_0) > 10\}$. Then:

$$\begin{aligned} \mathcal{D}'_0 &= (H_v(mv(3))[S_0]) > 10 \\ &= (\text{IF } \exists z(mv(z) = mv(3)) \\ &\quad \text{THEN } v(S_0) - z \text{ ELSE } v(S_0)) > 10 \\ &= (v(S_0) - 3 > 10) \\ &= (v(S_0) > 13). \end{aligned}$$

Therefore, as expected, the progression of $v(S_0) > 10$ wrt a noise-free motion of 3 units is $v(S_0) > 13$. (The unique name axiom and arithmetic are used in the simplification.)

Progressing Degrees of Belief

There are two main complications when progressing beliefs wrt noisy sensors and actions. First, the p fluent will have to take the likelihood of the action α into account. Second, \mathcal{D}_0 need not be uniform in S_0 , since p typically requires quantification over initial situations (as in (3), for example). This leads to the following definition:

Definition 10: Let $\mathcal{D}_0 \cup \Sigma$ be an invertible basic action theory and α be a ground action of the form $A(\vec{t})$ where \vec{t} is uniform in now.⁶ Then $Pro(\mathcal{D}_0, \alpha)$ is defined as \mathcal{D}_0 with the following substitutions:

- $p(\iota, S_0)$ is replaced by $\frac{p(\iota, S_0)}{\Upsilon_A(\vec{t})[\iota]}$;
- every other fluent term $f(u)$ is replaced by $H_f(\alpha)[u]$.

Here, $\Upsilon_A(\vec{x})$ refers to the RHS of the likelihood axiom for $A(\vec{x})$.

The main result of this paper is the correctness of this definition of progression:

Theorem 11: *Under the conditions of the definition above, let $\mathcal{D}'_0 = Pro(\mathcal{D}_0, \alpha)$. Suppose that $\mathcal{D}_0 \models (\Upsilon_A(\vec{t}) \neq 0)[S_0]$. Then for any \mathcal{L} -formula ϕ rooted in now,*

$$\mathcal{D}_0 \cup \Sigma \models \phi[do(\alpha, S_0)] \text{ iff } \mathcal{D}'_0 \cup \Sigma \models \phi[S_0].$$

Proof sketch: The proof is based on ideas in (Lin and Reiter 1997). To show that $\mathcal{D}'_0 \cup \Sigma$ is the progression of $\mathcal{D}_0 \cup \Sigma$, we need to show that for any model M , M is a model of $\mathcal{D}'_0 \cup \Sigma$ iff there is a model M' of $\mathcal{D}_0 \cup \Sigma$ such that for any ϕ , $M \models \phi[S_0]$ iff $M' \models \phi[do(\alpha, S_0)]$. We then use the fact that for any situation term σ , $f(\sigma)$ and $f(do(\alpha, \sigma))$ are related precisely to each other in terms of $E_f(\alpha)$ and $H_f(\alpha)$ to satisfy this model-theoretic property. ■

This theorem gives us the desired property for *Bel* (which is defined in terms of p) as a corollary:

Corollary 12: *Suppose $\mathcal{D}_0, \Sigma, \mathcal{D}'_0, \phi$, and α are as above. Then for all real numbers n :*

$$\begin{aligned} \mathcal{D}_0 \cup \Sigma \models Bel(\phi, do(\alpha, S_0)) = n \\ \text{iff } \mathcal{D}'_0 \cup \Sigma \models Bel(\phi, S_0) = n. \end{aligned}$$

Thus the degree of belief in ϕ after a physical or sensing action is equal to the initial belief in ϕ in a progressed theory.

We now present some examples, considering, in turn, noise-free actions, noisy sensing and finally noisy actions.

Example 13: Let us consider an action theory with a vertical action $mv(z)$, a sensing action $sonar(z)$ and two horizontal actions: *towards* moves the robot halfway towards the wall and *away* moves the robot halfway away from the wall. Formally, let $\mathcal{D}_0 \cup \Sigma$ be an action theory where \mathcal{D}_0 contains just (3), and Σ includes

- foundational axioms and (P1)-(P3) as usual;
- a l -axiom for $sonar(z)$, namely (2);
- l -axioms for the other actions, which are noise-free, and so these simply equal 1;
- a successor state axiom for v , namely (1);
- the following successor state axiom for h :

$$\begin{aligned} h(do(a, s)) = \\ \text{IF } a = \textit{away} \text{ THEN } 3/2 \cdot h(s) \\ \text{ELSE IF } a = \textit{towards} \text{ THEN } 1/2 \cdot h(s) \\ \text{ELSE } h(s). \end{aligned}$$

⁶In the most common case (like noise-free or sensing actions), the arguments to the action would simply be a vector of constants.

We noted that (5) does not satisfy our invertibility property. This variant, however, is invertible. The H -expression for v was derived in Example 4. The H -expression for h is:

$$H_h(a) = \begin{array}{l} \text{IF } a = \textit{away} \text{ THEN } 2/3 \cdot h \\ \text{ELSE IF } a = \textit{towards} \text{ THEN } 2 \cdot h \\ \text{ELSE } h. \end{array}$$

We now consider the progression of \mathcal{D}_0 wrt the action *away*. First, the instantiated H -expressions would simplify to:

- $H_h(\textit{away}) = 2/3 \cdot h$;
- $H_v(\textit{away}) = v$.

Next, since *away* is noise-free, we have $\Upsilon_{\textit{away}} = 1$. Putting this together, we obtain $\mathcal{D}'_0 = \textit{Pro}(\mathcal{D}_0, \textit{away})$ as:

$$\begin{aligned} p(s, S_0) &= \mathcal{U}(2/3 \cdot h; 2, 12) \times \mathcal{N}(v; 0, 1) [s] \\ &= \mathcal{U}(h; 3, 18) \times \mathcal{N}(v; 0, 1) [s] \end{aligned}$$

That is, the new p is one where h is uniformly distributed on $[3, 18]$ and v is independently drawn from a standard normal distribution (as before). This leads to a *shorter* and *wider* density function, as depicted in Figure 2. Here are three simple properties to contrast the original *vs.* the progressed:

- $\mathcal{D}_0 \cup \Sigma \models \textit{Bel}(h \geq 9, S_0) = .3$.

The *Bel* term expands as:

$$\frac{1}{\eta} \int_x \int_y \text{IF } \exists \iota (h = x \wedge v = y \wedge h \geq 9) [\iota] \text{ THEN } \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; 0, 1) [\iota] \text{ ELSE } 0$$

which simplifies to the integration of a density function:

$$\begin{aligned} &\frac{1}{\eta} \int_x \int_y \begin{cases} .1 \times \mathcal{N}(y; 0, 1) & \text{if } x \in [2, 12], x \geq 9 \\ 0 & \text{otherwise} \end{cases} \\ &= \frac{1}{\eta} \int_x \int_y \begin{cases} .1 \times \mathcal{N}(y; 0, 1) & \text{if } x \in [9, 12] \\ 0 & \text{otherwise} \end{cases} \\ &= .3. \end{aligned}$$

Only those situations where $h \in [2, 12]$ initially are given non-zero p values and by the formula in the *Bel*-term, only those where $h \geq 9$ are to be considered.

- $\mathcal{D}_0 \cup \Sigma \models \textit{Bel}(h \geq 9, \textit{do}(\textit{away}, S_0)) = .6$.

For any initial situation ι , $h[\textit{do}(\textit{away}, \iota)] \geq 9$ only when $h[\iota] \geq 6$, which is given an initial belief of .6.

- $\mathcal{D}'_0 \cup \Sigma \models \textit{Bel}(h \geq 9, S_0) = .6$.

Basically, *Bel* simplifies to an expression of the form:

$$\frac{1}{\eta} \int_x \int_y \begin{cases} 1/15 \cdot \mathcal{N}(y; 0, 1) & \text{if } x \in [3, 18], x \geq 9 \\ 0 & \text{otherwise} \end{cases}$$

giving us .6.

Example 14: Let $\mathcal{D}_0 \cup \Sigma$ be exactly as above, and consider its progression wrt *towards*. It is easy to verify that for instantiated H -expressions we get:

- $H_h(\textit{towards}) = 2 \cdot h$;
- $H_v(\textit{towards}) = v$;

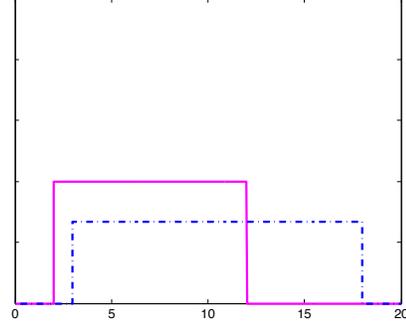


Figure 2: Belief change about h : initially (solid magenta) and after moving away (dotted blue).

Here too, because *towards* is noise-free, $\Upsilon_{\textit{towards}}$ is simply 1, which is to say the $\mathcal{D}'_0 = \textit{Pro}(\mathcal{D}_0, \textit{towards})$ is defined as:

$$\begin{aligned} p(s, S_0) &= \mathcal{U}(2 \times h; 2, 12) \times \mathcal{N}(v; 0, 1) [s] \\ &= \mathcal{U}(h; 1, 6) \times \mathcal{N}(v; 0, 1) [s]. \end{aligned}$$

The new distribution on h is *narrower* and *taller*, as shown in Figure 3. Here we might contrast \mathcal{D}_0 and \mathcal{D}'_0 as follows:

- $\mathcal{D}_0 \cup \Sigma \models \textit{Bel}(h \in [2, 3], S_0) = .1$.
- $\mathcal{D}'_0 \cup \Sigma \models \textit{Bel}(h \in [2, 3], S_0) = .2$.

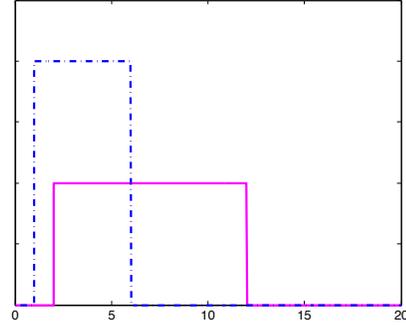


Figure 3: Belief change about h : initially (solid magenta) and moving towards the wall (dotted blue).

Example 15: Let $\mathcal{D}_0 \cup \Sigma$ be as in the previous examples. Consider its progression wrt the action *sonar(5)*. Sensing actions do not affect fluents, so for H -expressions we have:

- $H_h(\textit{sonar}(5)) = h$;
- $H_v(\textit{sonar}(5)) = v$.

Here *sonar(z)* is noisy, and we have $\Upsilon_{\textit{sonar}(5)} = \mathcal{N}(5; h, 4)$. This means that the progression $\mathcal{D}'_0 = \textit{Pro}(\mathcal{D}_0, \textit{sonar}(5))$ is

$$\frac{p(\iota, S_0)}{\mathcal{N}(5; h, 4)[\iota]} = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; 0, 1)[\iota],$$

which simplifies to the following:

$$p(\iota, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; 0, 1) \times \mathcal{N}(5; h, 4) [\iota].$$

As can be noted in Figure 4, the robot's belief about h 's true value around 5 has sharpened. Consider, for example, that:

- $\mathcal{D}_0 \cup \Sigma \models \text{Bel}(h \leq 9, S_0) = .7$.
- $\mathcal{D}'_0 \cup \Sigma \models \text{Bel}(h \leq 9, S_0) \approx .97$.

If we were to progress \mathcal{D}'_0 further wrt a second sensing action, say $\text{sonar}(5.9)$, we would obtain the following:

$$p(t, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; 0, 1) \times \mathcal{N}(5; h, 4) \times \mathcal{N}(5.9; h, 4) [t].$$

As can be seen in Figure 4, the robot's belief about h would sharpen significantly after this second sensing action. If we let $\mathcal{D}''_0 = \text{Pro}(\mathcal{D}'_0, \text{sonar}(5.9))$ then:

- $\mathcal{D}''_0 \cup \Sigma \models \text{Bel}(h \leq 9, S_0) \approx .99$.

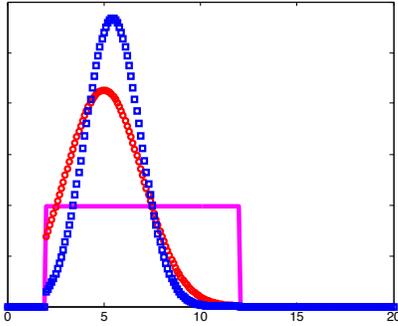


Figure 4: Belief change about h : initially (solid magenta), after sensing 5 (red circles), and after sensing twice (blue squares).

Example 16: Let \mathcal{D}_0 be the conjunction of (3) and $v(S_0) = 3$. Let Σ be the union of:

- (P1)-(P3) and domain-independent foundational axioms;
- a successor state axiom for h as above;
- a noisy move action mv' with the following l -axiom:

$$l(mv'(x, y), s) = \mathcal{N}(y; x, 2)$$

- a successor state axiom for v using this noisy move:

$$v(\text{do}(a, s)) = \text{IF } \exists x, y (a = mv'(x, y)) \text{ THEN } v(s) + y \text{ ELSE } v(s).$$

(Recall that for a noisy move $mv'(x, y)$, x is the intended motion and y is the actual motion.) This is inverted using the same idea as in Example 4.

Consider the progression of $\mathcal{D}_0 \cup \Sigma$ wrt $mv'(2, g(S_0))$. The simplified H -expressions are as follows:

- $H_h(mv'(2, g)) = h$;
- $H_v(mv'(2, g)) = v - g$.

By definition, occurrences of $v(u)$ in \mathcal{D}_0 are to be replaced by $H_v(mv'(2, g))[u]$. Also, $\Upsilon_{mv'}(2, z) = \mathcal{N}(z; 2, 2)$. Therefore, $\mathcal{D}'_0 = \text{Pro}(\mathcal{D}_0, mv'(2, g(S_0)))$ is defined to be

$$(v - g)[S_0] = 3 \wedge \left(\frac{p(t, S_0)}{\mathcal{N}(g; 2, 2)[t]} = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v - g; 0, 1)[t] \right)$$

This simplifies to the conjunction of these two sentences:

- $v(S_0) = 3 + g(S_0)$; and
- $p(t, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; g, 1) \times \mathcal{N}(g; 2, 2) [t]$.

Thus the noisy action has had two effects: the actual position has shifted by $g(S_0)$ units, and the belief about the position has also shifted by an amount g drawn from a normal distribution centered around 2. This leads to a *shifted* and *wider* curve seen in Figure 5. As expected, the agent is considerably less confident about its position after a noisy move. Here, for example, are the degrees of belief about being located within 1 unit of the best estimate (that is, the mean):

- $\mathcal{D}_0 \cup \Sigma \models \text{Bel}(v \in [-1, 1], S_0) \approx .68$.
- $\mathcal{D}'_0 \cup \Sigma \models \text{Bel}(v \in [1, 3], S_0) \approx .34$.

Basically, Bel expands to an expression of the form

$$\frac{1}{\eta} \int_{x,y,z} \begin{cases} .1\mathcal{N}(y; z, 1) \cdot \mathcal{N}(z; 2, 2) & \text{if } x \in [2, 12], y \in [1, 3] \\ 0 & \text{otherwise} \end{cases}$$

where η is

$$\int_{x,y,z} \begin{cases} .1\mathcal{N}(y; z, 1) \cdot \mathcal{N}(z; 2, 2) & \text{if } x \in [2, 12] \\ 0 & \text{otherwise} \end{cases}$$

leading to .34.

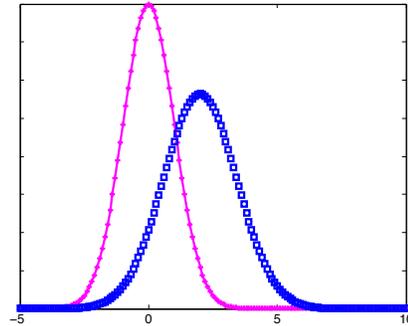


Figure 5: Belief change about v : initially (solid magenta) and after a noisy move of 2 units (blue squares).

Constant Growth

As seen in the examples, the result of progression is a theory \mathcal{D}'_0 which is essentially obtained by adding H -expressions, and so, its size is linear in the size of the *action theory*:

Theorem 17: *Suppose $\mathcal{D} = \mathcal{D}_0 \cup \Sigma$ is any invertible basic action theory. After the iterative progression of $\mathcal{D}_0 \cup \Sigma$ wrt a sequence σ , the size of the new initial theory is $O(|\mathcal{D}| \times |\sigma|)$.*

Therefore, progression is both computable and efficient in the sense of (Liu and Lakemeyer 2009). But for realistic robotic applications, even this may not be enough, especially over millions of actions. Consider, for example, that to calculate a degree of belief it will be necessary to integrate the numerical expression for p . What we turn to in this section is a special case that would guarantee that over any length of action sequences, the size of the progressed theory does not change beyond a constant factor. It will use the notion of *context-completeness* (Liu and Levesque 2005) and a few simplification rules.

Definition 18 Suppose $F \subseteq \{f_1, \dots, f_k\}$ is any set of fluents, and $\mathcal{D}_0 \cup \Sigma$ is any invertible basic action theory. We say that \mathcal{D}_0 is *complete* wrt F if for any $\phi \in \mathcal{L}_F$, either $\mathcal{D}_0 \models \phi$ or $\mathcal{D}_0 \models \neg\phi$, where \mathcal{L}_F is the sublanguage of \mathcal{L} restricted to the fluents in F .

Definition 19: An invertible basic action theory $\mathcal{D}_0 \cup \Sigma$ is said to be *context-complete* iff

- for every fluent f , \mathcal{D}_0 is complete wrt every fluent other than f appearing in the successor state axiom of f ;
- \mathcal{D}_0 is complete wrt every fluent appearing in a conditional expression in the likelihood axioms.

That is, there is sufficient information in \mathcal{D}_0 to simplify all the conditionals appearing in the context formulas of the successor state axioms and the likelihood axioms.

STRIPS actions are trivially context-complete, and so are Reiter’s *context-free* successor state axioms where only rigid symbols appear in the RHS (Reiter 2001). In Example 7, if \mathcal{D}_0 is complete wrt the fluents *fuel* and *wax*, then the theory would be context-complete. Note that \mathcal{D}_0 does not need to be complete wrt the fluent v in that example, and this is precisely why they are interesting. Indeed, both (5) and (1) are also context-complete because, by definition, E_f may mention f , and (say) use its previous value.⁷ The reader may further verify that all the density change examples developed in the paper are context-complete.

As a result of context-completeness, the body of the successor state axiom for a fluent f would simplify to an expression involving only f . As a consequence, H -expressions simplify to expressions of the form:

$$f(s) = \theta(f(\text{do}(a, s)))$$

where $\theta(x)$ is some rigid \mathcal{L} -term, possibly an arithmetic expression, with a single free variable x . Analogously, if $A(\vec{t})$ is any ground action, then,

$$\Upsilon_A(\vec{t}) = \delta(\vec{t}, \vec{f}).$$

where $\delta(\vec{x}, \vec{y})$ is some rigid \mathcal{L} -term, possibly an arithmetic expression, with free variables $\vec{x} \cup \vec{y}$. So in a progression, every fluent f would be replaced by an arithmetic expression using only f , resulting in a theory that is linear in the size of the *initial theory* \mathcal{D}_0 only:

Theorem 20: *Suppose $\mathcal{D}_0 \cup \Sigma$ is any invertible basic action theory that is also context-complete. After the iterative progression of $\mathcal{D}_0 \cup \Sigma$ wrt a sequence σ , the size of the new initial theory is $O(|\mathcal{D}_0| \times |\sigma|)$.*

This can make a substantial difference in the size of the expression for p . Nevertheless, after n actions, we might still end up with an expression consisting of a product of n terms. We now simplify this computation further by appealing to the analytical tractability of *conjugate distributions* (Box and Tiao 1973; Bertsekas and Tsitsiklis 2008), such as Gaussians, where the product of two Gaussians is definable as another Gaussian. We show how this helps with an example.

⁷Strictly speaking, our notion of context-completeness is inspired by, but not the same as the one in (Liu and Levesque 2005). This pertains to the clarification just made: we allow successor state axioms to use the fluent’s previous value.

Example 21: Let $\mathcal{D}_0 \cup \Sigma$ be as in Example 16. We noted its progression wrt $mv'(2, g(S_0))$ includes:

$$p(\iota, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; g, 1) \times \mathcal{N}(g; 2, 2) [\iota].$$

If we progress this sentence further wrt a second noisy action $mv'(3, g'(S_0))$, we would obtain:

$$p(\iota, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v - g'; g, 1) \times \mathcal{N}(g; 2, 2) \times \mathcal{N}(g'; 3, 2) [\iota] \quad (6)$$

By letting g^* be a new fluent symbol such that $g^* = g + g'$, it follows that (6) is equivalent to the following:

$$p(\iota, S_0) = \mathcal{U}(h; 2, 12) \times \mathcal{N}(v; g^*, 1) \times \mathcal{N}(g^*; 5, 4) [\iota].$$

Analogously, after any number n of noisy moves, we would only need to integrate over 3 variables vs. $n + 2$ in (6).

Theorem 22: *Suppose $\mathcal{D}_0 \cup \Sigma$ is any invertible action theory that is context-complete and the likelihood axioms are normal distributions. After the iterative progression of $\mathcal{D}_0 \cup \Sigma$ wrt a sequence σ , the size of the new initial theory is $O(|\mathcal{D}_0|)$.*

That is, the size of the progression does not grow beyond a constant factor. This result is easily adapted to other conjugate distributions. Essentially, like other *filtering* techniques (Thrun, Burgard, and Fox 2005), such as Kalman filters, we conclude that progression can be made very efficient. But unlike some of these other techniques, we are not limited to Gaussians. More significantly, unlike all other probabilistic formalisms, we can allow arbitrary context formulas, and under the reasonable assumption of context-completeness, still remain efficient.

Related Work and Discussion

This work builds on Lin and Reiter’s (1997) notion of progression. Other advances on progression have been made since then (Liu and Lakemeyer 2009; Vassos 2009), mainly by appealing to the notion of *forgetting* (Lin and Reiter 1994). We were motivated by concerns about continuity, and this led to the notion of invertible theories. These theories allowed us to perform first-order progression by inverting successor state axioms in a way that, as far as we know, has not been investigated before. Although we restricted \mathcal{L} to nullary real-valued fluents, we suspect that invertibility and its connection to progression may apply more generally. This is left for future investigations.

The progression of categorical knowledge against noise-free effectors and sensors is considered in (Liu and Wen 2011; Lakemeyer and Levesque 2009). The progression of discrete degrees of belief wrt context-completeness is considered in (Belle and Lakemeyer 2011). In the fluent calculus (Thielscher 2001), a dual form of successor state axioms is used, leading naturally to a form of progression. However, continuity is not considered in any of these.

The form of progression considered here follows Lin and Reiter and differs from weaker forms including the one proposed by Liu and Levesque (2005), and the notion of *logical filtering* (Shirazi and Amir 2005; Hajishirzi and Amir 2010),

which is a form of (approximate) progression.⁸ Interestingly, logical filtering is inspired by Kalman filters (Thrun, Burgard, and Fox 2005), although the precise connection is not considered. In situation calculus terminology, Kalman filters and its variants are derived using *strongly* context-free (Reiter 2001) noisy actions and sensors, with additive Gaussian noise, over normally distributed fluents. See, for example, (Belle and Levesque 2013b). Indeed, as discussed in BL, the framework used here is significantly more general than probabilistic formalisms used in the robotics and uncertainty literature, including Dynamic Bayesian Networks (Boyan and Koller 1998; Darwiche and Goldszmidt 1994) and hybrid control structures (McIlraith et al. 2000), among others.

See BL for a discussion of how the formalism used here also differs from other logical accounts for reasoning about uncertainty, such as probabilistic logics (Fagin and Halpern 1994), Markov logics (Richardson and Domingos 2006; Choi, Guzman-Rivera, and Amir 2011), probabilistic planning languages (Kushmerick, Hanks, and Weld 1995; Younes and Littman 2004; Sanner 2011), dynamic logics (Van Benthem, Gerbrandy, and Kooi 2009), and previous first-order proposals (Mateus et al. 2001; Thielscher 2001; Martin and Thielscher 2009; Fritz and McIlraith 2009; Poole 1998). In essence, none of these address continuous uncertainty and continuous noise in a general way.

Conclusions

Lin and Reiter developed the notion of progression, with long-lived agents in mind. However, their account does not deal with probabilistic uncertainty nor with noise, as seen in real-world robotic applications. In the work here, we consider semantically correct progression in the presence of continuity. By first identifying what we called invertible basic action theories, we obtained a new way of computing progression. Under the additional restriction of context-completeness, progression is very efficient. Most significantly, by working within a richer language, we have obtained progression machinery that, to the best of our knowledge, has not been discussed elsewhere, and goes beyond existing techniques. The unrestricted nature of the specification of the p fluent, for example, which we inherit from (Bacchus, Halpern, and Levesque 1999), allows for agents whose beliefs are not determined by a unique distribution.

One major topic for future investigation is an implementation. But apart from simply progressing a theory, we would also like to support the numerical calculation of degrees of belief. For this, we may need to appeal to numerical integration techniques such as Monte Carlo sampling (Murphy 2012), and so we would also like to study formal constraints on \mathcal{D}_0 which guarantee that sampling can be done in an effective manner.

⁸In (Shirazi and Amir 2005; Hajishirzi and Amir 2010), the notion of a *permuting action* is introduced for computing their form of progression, which bears some similarity to invertible fluents. However, as mentioned above, our work on continuity led itself to invertibility. Neither continuous uncertainty nor continuous noise is considered in (Shirazi and Amir 2005; Hajishirzi and Amir 2010).

Acknowledgements

The authors are grateful to the Natural Sciences and Engineering Research Council of Canada for financial support. The authors also thank the reviewers for their feedback.

References

- Bacchus, F.; Halpern, J. Y.; and Levesque, H. J. 1999. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence* 111(1–2):171 – 208.
- Baker, A. B. 1989. A simple solution to the yale shooting problem. In *Proc. KR*, 11–20.
- Belle, V., and Lakemeyer, G. 2011. A semantical account of progression in the presence of uncertainty. In *Proc. AAAI*, 165–170.
- Belle, V., and Levesque, H. J. 2013a. Reasoning about continuous uncertainty in the situation calculus. In *Proc. IJCAI*.
- Belle, V., and Levesque, H. J. 2013b. Reasoning about probabilities in dynamic systems using goal regression. In *Proc. UAI*.
- Bertsekas, D., and Tsitsiklis, J. 2008. *Introduction to probability*. Athena Scientific, 2nd edition edition.
- Box, G. E. P., and Tiao, G. C. 1973. *Bayesian inference in statistical analysis*. Addison-Wesley.
- Boyan, X., and Koller, D. 1998. Tractable inference for complex stochastic processes. In *Proc. UAI*, 33–42.
- Choi, J.; Guzman-Rivera, A.; and Amir, E. 2011. Lifted relational kalman filtering. In *Proc. IJCAI*, 2092–2099.
- Darwiche, A., and Goldszmidt, M. 1994. Action networks: A framework for reasoning about actions and change under uncertainty. In *Proc. UAI*, 136–144.
- Fagin, R., and Halpern, J. Y. 1994. Reasoning about knowledge and probability. *J. ACM* 41(2):340–367.
- Fritz, C., and McIlraith, S. A. 2009. Computing robust plans in continuous domains. In *Proc. ICAPS*, 346–349.
- Hajishirzi, H., and Amir, E. 2010. Reasoning about deterministic actions with probabilistic prior and application to stochastic filtering. In *Proc. KR*.
- Kushmerick, N.; Hanks, S.; and Weld, D. 1995. An algorithm for probabilistic planning. *Artificial Intelligence* 76(1):239–286.
- Lakemeyer, G., and Levesque, H. J. 2007. Cognitive robotics. In *Handbook of Knowledge Representation*. Elsevier. 869–886.
- Lakemeyer, G., and Levesque, H. 2009. A semantical account of progression in the presence of defaults. In *Conceptual Modeling: Foundations and Applications*. Springer. 82–98.
- Levesque, H. J.; Pirri, F.; and Reiter, R. 1998. Foundations for the situation calculus. *Electron. Trans. Artif. Intell.* 2:159–178.
- Lin, F., and Reiter, R. 1994. Forget it. In *Working Notes of AAAI Fall Symposium on Relevance*, 154–159.
- Lin, F., and Reiter, R. 1997. How to progress a database. *Artificial Intelligence* 92(1-2):131–167.

- Liu, Y., and Lakemeyer, G. 2009. On first-order definability and computability of progression for local-effect actions and beyond. In *Proc. IJCAI*, 860–866.
- Liu, Y., and Levesque, H. 2005. Tractable reasoning with incomplete first-order knowledge in dynamic systems with context-dependent actions. In *Proc. IJCAI*, 522–527.
- Liu, Y., and Wen, X. 2011. On the progression of knowledge in the situation calculus. In *Proc. IJCAI*, 976–982.
- Martin, Y., and Thielscher, M. 2009. Integrating reasoning about actions and Bayesian networks. In *International Conference on Agents and Artificial Intelligence (ICAART)*.
- Mateus, P.; Pacheco, A.; Pinto, J.; Sernadas, A.; and Sernadas, C. 2001. Probabilistic situation calculus. *Annals of Math. and Artif. Intell.* 32(1-4):393–431.
- McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, 463–502.
- McIlraith, S.; Biswas, G.; Clancy, D.; and Gupta, V. 2000. Hybrid systems diagnosis. In *Proc. of Workshop on Hybrid Systems: Computation and Control (HSCC 2000)*, LNCS, 282–295.
- Murphy, K. P. 2012. *Machine learning: a probabilistic perspective*. The MIT Press.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.
- Poole, D. 1998. Decision theory, the situation calculus and conditional plans. *Electron. Trans. Artif. Intell.* 2:105–158.
- Reiter, R. 2001. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT Press.
- Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1):107–136.
- Sanner, S. 2011. Relational dynamic influence diagram language (rddl): Language description. Technical report, Australian National University.
- Scherl, R. B., and Levesque, H. J. 2003. Knowledge, action, and the frame problem. *Artificial Intelligence* 144(1-2):1–39.
- Shirazi, A., and Amir, E. 2005. First-order logical filtering. In *Proc. IJCAI*, 589–595.
- Thielscher, M. 2001. Planning with noisy actions (preliminary report). In *Proc. Australian Joint Conference on Artificial Intelligence*, 27–45.
- Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic Robotics*. MIT Press.
- Trench, W. F. 2003. *Introduction to real analysis*. Prentice Hall.
- Van Benthem, J.; Gerbrandy, J.; and Kooi, B. 2009. Dynamic update with probabilities. *Studia Logica* 93(1):67–96.
- Vassos, S., and Levesque, H. 2008. On the Progression of Situation Calculus Basic Action Theories: Resolving a 10-year-old Conjecture. In *Proc. AAAI*, 1004–1009.
- Vassos, S. 2009. *A Reasoning Module for Long-Lived Cognitive Agents*. Ph.D. Dissertation, University of Toronto, Toronto, Canada.
- Waldinger, R. 1977. Achieving several goals simultaneously. In *Machine Intelligence*, volume 8. 94–136.
- Younes, H., and Littman, M. 2004. PPDDL 1. 0: An extension to pddl for expressing planning domains with probabilistic effects. Technical report, Carnegie Mellon University.