

# PREGO: An Action Language for Belief-Based Cognitive Robotics in Continuous Domains

Vaishak Belle and Hector Levesque<sup>1</sup>

*This paper is an abridged version of one appearing in the 2014 AAAI Conference on Artificial Intelligence [5].*

**Abstract.** The area of cognitive robotics is often subject to the criticism that the proposals investigated in the literature are too far removed from the kind of continuous uncertainty and noise seen in actual real-world robotics. This paper proposes a new language and an implemented system, called PREGO, based on the situation calculus, that is able to reason effectively about degrees of belief against noisy sensors and effectors in continuous domains. It embodies the representational richness of conventional logic-based action languages, such as context-sensitive successor state axioms, but is still shown to be efficient using a number of empirical evaluations. We believe that PREGO is a powerful framework for exploring real-time reactivity and an interesting bridge between logic and probability for cognitive robotics applications.

## 1 INTRODUCTION

Cognitive robotics, as envisioned in [16], is a high-level control paradigm that attempts to apply knowledge representation (KR) technologies to the reasoning problems faced by an autonomous agent/robot in an incompletely known dynamic world. This is a challenging problem; at its core, it requires a clear understanding of the relationships among the beliefs, perception, and actions of the agent. To that end, sophisticated knowledge-based proposals for reasoning about action and change have been investigated in the literature, as demonstrated in [7, 24, 30, 13], among many others. One major criticism leveled at this line of work, however, is that the theory seems far removed from the kind of continuous uncertainty and noise seen in robotic applications when real sensors and effectors are deployed. The interpreters are often propositional or resort to the closed-world assumption, and at best, only allow limited forms of incomplete information [24, 25]. Rather surprisingly, very little attention has been devoted to the integration of action languages with probability densities or degrees of belief. As far as we know, there has yet to emerge a simple specification language that, (a) has the desirable features of popular action formalisms such as context-dependent successor state axioms [24], (b) allows for expressing discrete and continuous noise in effectors and sensors, and most significantly, (c) is equipped with an effective computational methodology for handling *projection* [24], the reasoning problem at the heart of planning and agent programming. For real-time reactivity, it is also desirable that the reasoning that is needed (under suitable representational stipulations) be as practical as common filtering techniques [32].

<sup>1</sup> Dept. of Computer Science, University of Toronto, Canada; email: {vaishak, hector}@cs.toronto.edu

This paper proposes a system called PREGO as a first step in this direction. Informally, the language of PREGO is built from the following components:

- symbols for the fluents over which one can define a (continuous or discrete) probability distribution;
- action symbols for effectors and sensors, possibly noisy;
- specifications for the preconditions of actions, successor state of fluents, and the results of sensing operations.

In this paper, we study the formal foundations of PREGO, as well as a computational methodology for projecting belief, that is, for computing degrees of belief after any sequence of actions and sensing operations.<sup>2</sup> We show that the language of PREGO can be interpreted as a situation-suppressed dialect of the situation calculus [24], and that the embodied projection mechanism is a special form of goal regression. From a logical point of view, PREGO's handling of continuity and uncertainty goes beyond the capabilities of popular interpreters for KR action languages. From a probability point of view, PREGO allows successor state and sensing axioms that can be arbitrarily complex [24], making it not only significantly more expressive than probabilistic formalisms [6], but also current probabilistic planning formalisms, *e.g.* [27]. To the best of our knowledge, a development of this kind has not been considered in this generality before. PREGO is fully implemented; for empirical evaluations of its behavior on non-trivial projection tasks, see [5].

We structure the paper as follows. We first introduce PREGO and discuss some very simple belief change examples. We then introduce the background logical language of the situation calculus and study PREGO's techniques. Then, related work and conclusions are presented.

## 2 PREGO

The PREGO language is a simple representation language with a LISP-like syntax.<sup>3</sup> A domain in PREGO is modeled as a *basic action theory* (or BAT) made up of the following five expressions (which we will illustrate immediately below):<sup>4</sup>

### 1. (define-fluents *fluent fluent ...*)

A list of all the fluents to be used in the BAT. These can be thought of as probabilistic variables, whose values may or may not be known.

<sup>2</sup> For simplicity, only projection is considered for the language. High-level control structures [17], such as recursion and loops, is left for the future.

<sup>3</sup> The system is realized in the RACKET dialect of the SCHEME family (racket-lang.org). Note that the technical development does not hinge on any feature unique to this programming language.

<sup>4</sup> For space reasons, we omit action preconditions in this paper.




Figure 1. Robot moving towards a wall.

2. `(define-ini-p-expr expr)`

Here, *expr* is an expression mentioning the fluents from (1) that should evaluate to a number between 0.0 and 1.0. It indicates the probability density given to any initial state in terms of the values of the fluents in that state.

3. `(define-ss-exprs fluent act expr act expr ...)`

This determines the successor-state expressions for the given fluent *fluent*. The *act* parameters are of the form *(name var var ...)* where the *vars* are the arguments of the action and are used in the corresponding *expr*. The idea is that if the action takes place, the new value of the fluent is the value of *expr*. If an action does not appear on the list, the fluent is unchanged by the action.

4. `(define-l-exprs act expr act expr ...)`

The format of *act* and *expr* are as in (3). For each *act*, the *expr* is a numerical expression like in (2) and determines the likelihood of that action. If an action does not appear on the list, it is assumed to have a likelihood of 1.0.

5. `(define-altS act altfn act altfn ...)`

The format of the *act* is as in (3) and (4). The *altfn* is a function of one argument that produces noisy versions of *act* for that argument. If an action does not appear on the list, then it is exact, that is, without a noisy version.

To illustrate these, let us consider the simple scenario depicted in Figure 1, where a robot is moving in a 1-dimensional world towards a wall. Its distance to the wall is given by a fluent *h*. Suppose:

- The robot initially does not know how far it is from the wall, but that the distance satisfies  $2 \leq h \leq 12$ . In other words, the robot believes that the initial value of *h* is drawn from a (continuous) uniform distribution on [2,12].
- The robot has a distance sensor aimed at the wall. The sensor is noisy in that the value *z* read on the sensor differs from the actual value of *h*, but in a reasonable way. We assume that the likelihood of getting *z* is given by a normal distribution whose mean is the true value *h*.
- The robot has an effector for moving exactly *z* units towards or away from the wall, but this motion stops when the wall is reached.

A BAT for this domain is shown in Figure 2. We have a single fluent *h* and two actions: a sensing action *sonar*, and a physical action *fwd*. Note that we can use RACKET arithmetic (e.g. `max` and `-`) and any other function that can be defined in RACKET (e.g. the predefined `UNIFORM` and `GAUSSIAN`). Note also that the *expr* terms are quoted expressions where fluents appear as free variables. (We use backquote and comma to embed the action argument *z* in the expression.)

Once the fluents are defined (here only *h*), the designer provides the initial joint distribution over the fluents using any mathematical function (here, a uniform distribution over one variable).<sup>5</sup> The successor state axiom says that the value of *h* after a *fwd* action is

<sup>5</sup> For simplicity, full joint distributions are assumed; when additional structure is available in the form of conditional independencies, then belief networks, among others, can be used [22].

```
(define-fluents h)
(define-ini-p-expr '(UNIFORM h 2 12))
(define-ss-exprs h
  (fwd z) '(max 0 (- h ,z)))
(define-l-exprs
  (sonar z) '(GAUSSIAN ,z h 4.0))
```

Figure 2. PREGO's input for the simple robot domain.

obtained by subtracting *z* from the previous value or 0, whichever is higher. The likelihood expression says that the sonar readings are expected to be centered on the value of *h* with a standard deviation of 4.0. Since the physical action *fwd* is assumed to be noise-free, `define-altS` is not used in this BAT.

Let us now turn to a noisy version of *fwd*. The idea here is that the agent might attempt a move by 3 units but end up actually moving 3.094 units. 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.

To model this, we imagine a new action symbol *nfwd* with two arguments: the first captures the intended motion amount, and the second captures the actual motion amount. Then, the `ss-exprs` block for the fluent *h* would include:

```
(nfwd x y) '(max 0 (- h ,y))
```

That is, the true value of *h* changes according to the second argument, not the first. Since the robot does not know the actual outcome, all it knows is that `(nfwd 3 z)` occurred for some value of *z*, which is captured using `define-altS`:

```
(define-altS
  (nfwd x y) (lambda (z) '(nfwd ,x ,z)))
```

Finally, to indicate that this noisy move has (say) Gaussian noise, we would include the following in the `l-exprs` block:

```
(nfwd x y) '(GAUSSIAN ,y ,x 1.0)
```

In other words, the actual amount moved is normally distributed around the intended value with a variance of 1.

## Using PREGO

PREGO can be used to reason about what is believed after any sequence of physical or sensing actions. We use


```
(eval-bel expr actions)
```

where *expr* is any Boolean expression (with the fluents as free variables) and *actions* is a list of the actions that occurred. For example, after moving 4 units towards the wall, the robot believes it is as likely as not to be within 3 units:

```
> (eval-bel (< h 3) ((fwd 4)))
0.5
```

Note that PREGO is handling a *mixed distribution* here. After the action, the possibility that  $h = 0$  must be accorded a *weight* of .2 (i.e. from all points where  $h \in [2, 4]$  initially), while points where  $h \in (0, 8]$  retain a *density* of .1.

Likewise, suppose we are interested in the agent’s beliefs after sensing the value 5 on its sonar. Clearly, its beliefs should sharpen around 5, and if the robot obtained a second reading of 5.12, its belief would sharpen further. If we were to plot PREGO’s computed beliefs for the fluent  $h$  after the sequence ((sonar 5) (sonar 5.12)), we would obtain Figure 3. In the following sections, we justify these results and also discuss examples involving the noisy effector.



**Figure 3.** Beliefs about  $h$  initially (solid red), after sensing 5 (magenta markers) and after a second sensor reading (blue squares).

Before ending this section, note that, as mentioned, the language allows arbitrary context-dependent successor-state and likelihood specifications [24]. For example, imagine a binary fluent *wet* which says whether the floor is wet or not. Assume that, initially, it is very likely to be wet:

```
(define-ini-p-expr
  (* (UNIFORM h 2 12)
     (DISCRETE wet #t .8 #f .2)))
```

Then to model a noisy move whose reliability depends on whether or not the floor is wet, we might have:

```
(define-l-exprs
  (nfwd x y) '(GAUSSIAN ,y ,x (if wet 4.0 1.0)))
```

which says that the actual motion is determined by a Gaussian distribution with variance 4 when the floor is wet but with a variance 1 when the floor is dry. What makes the language rich is that there can be physical actions (e.g. mopping up and spilling coffee) that affect this context, as well as sensing actions (e.g. observing a reflection on the floor) that sharpen the agent’s belief about the context.

### 3 LOGICAL FOUNDATIONS

The PREGO language can be interpreted as a dialect of the situation calculus, extended to reason about degrees of belief. We will not go over that language here, except to note:

- it is many-sorted, with sorts for physical actions, sensing actions, situations, and objects (for everything else);
- a set of initial situations correspond to the ways the world might be initially — a constant  $S_0$  denotes the actual one;

- there is a distinguished binary symbol *do* such that  $do(a_1 \cdots a_n, s)$  denotes the situation resulting from performing actions  $a_1$  through  $a_n$  at situation  $s$ .

Fluents capture changing properties about the world. Here we assume that  $f_1, \dots, f_k$  are all the fluents in the language,<sup>6</sup> and that these take no arguments other than a single situation term. Note that this is not a propositional theory in that we allow the *values* of these fluents to range over any set, including the reals  $\mathbb{R}$ .

We following two notational conventions. First, we often suppress the situation argument in a logical expression  $\phi$  or use a distinguished variable *now*, and we let  $\phi[s]$  denote the formula with that variable replaced by  $s$ . Second, we use conditional *if-then-else* expressions in formulas throughout, possibly mentioning quantifiers. We take some liberties with the scope of variables in that  $f = \text{IF } \exists x. \phi \text{ THEN } t_1 \text{ ELSE } t_2$  is to mean  $\exists x [\phi \wedge f = t_1] \vee [(f = t_2) \wedge \neg \exists x. \phi]$ .

### Basic Action Theory

As hinted in PREGO’s presentation, a domain theory  $\mathcal{D}$  is formulated as a BAT [24] which includes sentences  $\mathcal{D}_0$  describing what is true initially, successor state axioms (SSA) of the form  $\forall a, s. f(do(a, s)) = \text{SSA}_f(a)[s]$ , and precondition axioms. For example, the effect of the *fwd* action, from the robot domain above, would be expressed as:<sup>7</sup>

$$h(do(a, s)) = (\text{IF } \exists z(a = fwd(z)) \text{ THEN } \max(0, h - z) \text{ ELSE } h)[s]$$

so as to incorporate Reiter’s solution to the frame problem.

For the task of *projection*, we are interested in the entailments of  $\mathcal{D}$ . Entailment is wrt standard Tarski models, but we assume that the obvious interpretations are assigned to arithmetic symbols, =, and real constants, such as  $\pi$  and  $e$ .

### Noise and Degrees of Belief

Our account of probabilistic uncertainty is based on [2, 1], but augmented here to also deal with *continuous* noisy effectors. These extensions to the situation calculus still benefit from Reiter’s solution to the frame problem. They also generalize the Scherl and Levesque [29] proposal, where actions and sensors are noise-free and beliefs are strictly categorical, that is, non-numeric.

Mirroring PREGO’s presentation, our account involves 3 distinguished symbols: *l*, *alt* and *p*.  $\mathcal{D}$  would now contain *l*-axioms of the form  $l(\alpha(\vec{x}), s) = L_\alpha(\vec{x})[s]$ , and *alt*-axioms of the form  $alt(a, u) = a'$ . For example, *nfwd* is modeled by including the appropriate SSA and the following in  $\mathcal{D}$ :<sup>8</sup>

$$l(nfwd(x, y), s) = \mathcal{N}(y; x, 1)[s]. \quad (1)$$

$$alt(nfwd(x, y), z) = nfwd(x, z). \quad (2)$$

Finally, the distinguished symbol *p* can be seen as a *numeric* variant of the accessibility relation in epistemic logics. Intuitively,  $p(s', s)$  is the density that the agent attributes to  $s'$  when at  $s$  [2]. As part of  $\mathcal{D}_0$ ,

<sup>6</sup> Non-logical symbols, such as fluent and action symbols, in the situation calculus are italicized, e.g. fluent  $h$  in PREGO is  $h$  in the language of the situation calculus.

<sup>7</sup> Free variables are implicitly assumed to be quantified from the outside. Note that SSAs are formulated here using conditional expressions, but they macro expand to Reiter’s formulation.

<sup>8</sup> We use  $\mathcal{N}$  and  $\mathcal{U}$  as abbreviations for mathematical formulas defining a Gaussian and uniform density respectively.

the modeler would provide the probability distribution on the agent's initial worlds, using a sentence of the form

$$p(s, S_0) = \text{INIT}(f_1, \dots, f_k)[s].$$

For example,

$$p(s, S_0) = \mathcal{U}(h; 2, 12)[s] \quad (3)$$

embodies the initial uncertainty of our robot from Figure 1.

Given such axioms in  $\mathcal{D}$ , the *degree of belief* in any situation-suppressed  $\phi$  in  $s$  is defined using the abbreviation:

$$\text{Bel}(\phi, s) \doteq \frac{1}{\gamma} \int_{f_1, \dots, f_k} \int_{u_1, \dots, u_n} \text{Density}(\phi, s^*)$$

where the normalization factor  $\gamma$  is the numerator but with  $\phi$  replaced by *true*, and if  $s = do(a_1 \dots a_n, S_0)$  then  $s^* = do(alt(a_1, u_1) \dots alt(a_n, u_n), S_0)$ .

The idea behind *Density* is simple. Starting from  $s$ , the density of  $do(a_1 \dots a_n, s)$  is the  $p$ -value of  $s$  times the likelihoods of each  $a_i$ . By integrating over  $\vec{u}$ ,  $alt(a_i, u_i)$  is taken into account, and so, all possible successors resulting from noisy actions are also considered.<sup>9</sup> For space reasons, we omit the definition; see [2]. We simply note that the belief change mechanism subsumes Bayesian conditioning [22], as used in the robotics literature [32].

## 4 COMPUTING PROJECTION

The projection mechanism seen in `eval-bel` is built on a special form of *goal regression*. In other words, we begin by finding a situation-suppressed expression  $r$  such that

$$\mathcal{D} \models \text{Bel}(\phi, do(a_1 \dots a_n, S_0)) = r[S_0].$$

Because only  $\mathcal{D}_0$  is needed to calculate  $r[S_0]$ , this reduces the calculation of belief after actions and sensing to a calculation in the initial situation in terms of INIT.

Such a regression operator is formulated in [3]. They show how *Bel*-expressions about the future reduce to *Density*-expressions about  $S_0$ . However, their proposal does not deal with noisy effectors. Moreover, their *Density*-expressions expand into formulas that quantify over initial situations. Consequently, considerable logical machinery is needed to further simplify these sentences to a purely numerical formula.

What we propose here is a new treatment that not only generalizes to both noisy acting and sensing, but one that involves only mathematical (as opposed to logical) expressions. Roughly, this is achieved by processing the logical terms in the goal in a *modular* fashion wrt the situation-suppressed RHS (also interpreted as logical terms) of the axioms in  $\mathcal{D}$ . The result is a Boolean expression (about  $S_0$ ), where fluents are free variables, that can be evaluated using any software with numerical integration capabilities. In other words, no logical consequence finding is necessary.

**Definition** Given a BAT  $\mathcal{D}$ , a situation-suppressed expression  $e$ , and an action sequence  $\sigma$ , we define  $\mathcal{R}[e, \sigma]$  as a situation-suppressed formula  $e'$  as follows:

1. If  $e$  is a fluent:
  - if  $\sigma = \epsilon$  (is empty), then  $e' = e$ ;

<sup>9</sup> For discrete fluents and discrete noisy effectors, one would replace  $\int_f$  and  $\int_u$  by  $\sum_f$  and  $\sum_u$  respectively. Let us also remark that both summations and integrals can be defined as terms in the logical language, as shown in [2].

- if  $\sigma = \sigma' \cdot a$  then  $e' = \mathcal{R}(\text{SSA}_e(a), \sigma')$ .
2. If  $e$  is a number, constant or variable, then  $e' = e$ .
  3. If  $e$  is  $\text{Bel}(\phi, now)$  then

$$e' = \frac{1}{\gamma} \int_{\mathcal{F}} \text{INIT} \times \mathcal{G}[\phi, \sigma]$$

where  $\mathcal{G}$  is an operator for obtaining a (mathematical) expression from the belief argument  $\phi$ , defined below.

4. Else  $e$  is  $(e_1 \circ e_2 \circ \dots \circ e_n)$  and

$$e' = (\mathcal{R}[e_1, \sigma] \circ \mathcal{R}[e_2, \sigma] \circ \dots \circ \mathcal{R}[e_n, \sigma])$$

where  $\circ$  is any mathematical operator over expressions, such as  $\neg, \wedge, =, +, \text{IF}, \mathcal{N}$ , etc.

As in [24], fluents are simplified one action at a time using appropriately instantiated SSAs. The main novelty here is how *Bel* is regressed using INIT, the RHS of the  $p$ -axiom in  $\mathcal{D}_0$ , with the argument  $\phi$  handled separately, and how  $\mathcal{R}$  works over arbitrary mathematical functions in a modular manner; e.g.  $\mathcal{R}[\mathcal{N}(t_1; t_2, t_3), \sigma]$  would give us  $\mathcal{N}(\mathcal{R}[t_1, \sigma]; \mathcal{R}[t_2, \sigma], \mathcal{R}[t_3, \sigma])$ . Now,  $\mathcal{G}$ :

**Definition** Let  $\mathcal{D}$  and  $\sigma$  be as above. Given any situation-suppressed fluent formula  $\phi$ , we define  $\mathcal{G}[\phi, \sigma]$  to be a situation-suppressed expression  $e$  as follows:

1. If  $\sigma = \epsilon$ , then  $e = \text{IF } \phi \text{ THEN } 1 \text{ ELSE } 0$ .
2. Else,  $\sigma = \sigma' \cdot \alpha(t)$ ; let  $\alpha(t') = alt(\alpha(t), u)$  and

$$e = \int_u \mathcal{R}[\mathcal{L}_\alpha(t'), \sigma'] \times \mathcal{G}[\mathcal{R}[\phi, \alpha(t')], \sigma'].$$

Essentially,  $\mathcal{G}$  integrates over all possible outcomes for a noisy action using *alt*. The likelihood of these outcomes is determined using  $\mathcal{L}_\alpha$ , the RHS of the  $l$ -axioms.

For our main theorem,  $\mathcal{R}$  is shown to have this property:

**Theorem 1** Let  $\mathcal{D}$  and  $\sigma$  be as above, and let  $e$  be any situation-suppressed expression. Then

$$\mathcal{D} \models e[do(\sigma, S_0)] = (\mathcal{R}[e, \sigma])[S_0].$$

When  $e$  is a belief formula, this allows us to reduce the belief calculation to the initial situation, as desired:

**Corollary 2** Let  $\mathcal{D}, \phi$  and  $\sigma$  be as above. Then

$$\mathcal{D} \models \text{Bel}(\phi, do(\sigma, S_0)) = (\mathcal{R}[\text{Bel}(\phi, now), \sigma])[S_0].$$

## From Specification to PREGO

$\mathcal{R}$  is a reduction operator that takes as input any Boolean expression (where fluents are free variables) and outputs a new one, leading to a surprisingly straightforward implementation that exactly follows Definitions 4 and 4. We demonstrate this using a BAT. Let  $\mathcal{D}$  include the situation calculus counterparts for our robot domain, i.e. (1), (2), (3) and:

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

$$h(do(a, s)) = (\text{IF } \exists x, y (a = \text{nfwd}(x, y)) \text{ THEN } \max(0, h - y) \text{ ELSE } h)[s] \quad (5)$$

Assume now that the robot *senses* 5 initially on the sonar, and then it attempts a noisy move of 2 units *away* from the wall (by providing

a negative argument to  $nfwd$ ). In reality, assume a move of 2.1 units occurs. (As we shall see, the second argument of  $nfwd$  determines the change to the  $h$  fluent, but does not affect beliefs about  $h$ , and so it can be any arbitrary number.) Suppose we are further interested in the robot’s beliefs about  $\phi = (h \leq 7)$ .  $\mathcal{R}$  works as follows:

$$\begin{aligned} & \mathcal{R}[Bel(\phi, now), do(sonar(5) \cdot nfwd(-2, -2.1), S_0)] \\ &= \frac{1}{\gamma} \int_h \mathcal{U}(h; 2, 12) \times \mathcal{G}[\phi, sonar(5) \cdot nfwd(-2, -2.1)] \\ &= \frac{1}{\gamma} \int_h \mathcal{U}(h; 2, 12) \times \int_u \left( \mathcal{R}[L_{nfwd}(-2, u), sonar(5)] \times \right. \\ & \quad \left. \mathcal{G}[\mathcal{R}[\phi, nfwd(-2, u)], sonar(5)] \right) \\ &= \frac{1}{\gamma} \int_h \mathcal{U}(h; 2, 12) \times \int_u \mathcal{N}(u; -2, 1) \times \mathcal{G}[\psi, sonar(5)] \\ &= \frac{1}{\gamma} \int_h \mathcal{U}(h; 2, 12) \times \int_u \mathcal{N}(u; -2, 1) \times \mathcal{N}(5; h, 4) \times \mathcal{G}[\psi, \epsilon] \end{aligned}$$

where  $\psi = (\mathcal{R}[\phi, nfwd(-2, u)]) = (\max(0, h - u) \leq 7)$ , and  $\mathcal{G}[\psi, \epsilon] = \text{IF } \psi \text{ THEN } 1 \text{ ELSE } 0$ .

In the `PREGO` system, the supporting regression can be examined using a second function, `regr-bel`, as follows:

```
> (regr-bel (<= h 7) ((sonar 5) (nfwd -2 -2.1)))
' (/
  (INTEGRATE (h u)
    (* (UNIFORM h 2 12) (GAUSSIAN u -2 1.0)
      (GAUSSIAN 5 h 4.0)
      (if (<= (max 0 (- h u)) 7) 1.0 0.0)))
  (INTEGRATE (h w)
    (* (UNIFORM h 2 12) (GAUSSIAN w -2 1.0)
      (GAUSSIAN 5 h 4.0))))
```

The answer is a quotient of two integrals (or summations in the discrete case), where the denominator is the normalization factor. The integrand of the numerator is a product of four terms: the first coming from `INIT`, the next two coming from the noisy acting and sensing, and the final due to the regression of the argument of `Bel`. What `eval-bel` then does is to evaluate these integrals numerically using Monte Carlo sampling [21]:<sup>10</sup>

```
> (eval-bel (<= h 7) ((sonar 5) (nfwd -2 -2.1)))
0.47595449413426844
```

## 5 RELATED WORK

The `PREGO` framework is based on the situation calculus. While this language is quite expressive, popular interpreters either make the closed-world assumption, or only allow certain kinds of disjunctive knowledge [24, 9, 8]. In other words: no degrees of belief. The notable exception to this is the MDP-inspired `DTGLOG` and related proposals [24]. Although a full comparison is difficult since `DTGLOG` implements a particular planning methodology while `PREGO` is just a specification language, there are significant differences: `DTGLOG` is for fully observable domains and only supports discrete probabilities. There has been recent work on POMDP extensions [28, 36]. However, they assume discrete noise in effectors, and make other strong structural assumptions, such as context-free (`STRIPS`-style) actions. In

<sup>10</sup> For standard distributions such as  $\mathcal{N}$  and  $\mathcal{U}$ , points can be generated for  $f_1, \dots, f_k$  (i.e. the fluents) and  $u_1, \dots, u_n$  (i.e. the new integration variables introduced for noisy actions) using `INIT` and the  $l$ -axioms respectively. These points are then tested for the regression of the argument to `Bel` (e.g.  $\mathcal{G}[\psi, \epsilon]$  above).

our view, context-dependent SSAs are one of the reasons to consider using a language like the situation calculus in the first place.

While the situation calculus has received a lot of attention, there are, of course, other action languages; e.g. see [31, 33, 14] for treatments on probabilities in other formalisms. They are, however, limited to discrete probabilities. Thus, we differ from these and many others [23, 18, 12, 11] in being able to address continuity in a general way, except for [2] that we build on. See [1, 2] for more discussions. In addition, `PREGO` is seen to be more expressive than current probabilistic planning languages [15, 35, 27]. Other continuous models for planning, such as [19], are procedural rather than declarative, and do not support contextual SSAs. As argued in [1], the same representational limitations also apply to most probabilistic formalisms, such as Kalman Filters and Dynamic Bayesian Networks [6].<sup>11</sup> Finally, recent work on relational probabilistic languages [26, 20] and probabilistic programming [10] feature sophisticated stochastic models, but do not handle actions.

## 6 CONCLUSIONS

This paper proposed a new declarative representation language, and studied a formal and computational account for projection with degrees of beliefs. The language allows for discrete and continuous fluents, noisy actions and noisy sensors. It incorporates important features of action languages, thereby providing an interesting bridge between realistic robotic concerns, on the one hand, and logic-based representation languages, on the other. In [5], we also discuss empirical studies that demonstrate why we feel that the `PREGO` system is powerful enough to explore real-time reactivity in cognitive robotics applications. To the best of our knowledge, no other proposal of this generality has been investigated.

There are two main avenues for the future. First, progression. The regression system in `PREGO` maintains the initial state, and is appropriate for planning. For some applications, it is desirable to periodically update the system [34]. Following [4], we would like to explore progression in `PREGO` and perhaps evaluate that against particle filters [32]. Second, in a companion paper, we intend to consider the equivalent of `GOLOG`’s program structures for `PREGO`.

## REFERENCES

- [1] F. Bacchus, J. Y. Halpern, and H. J. Levesque, ‘Reasoning about noisy sensors and effectors in the situation calculus’, *Artificial Intelligence*, **111**(1–2), 171 – 208, (1999).
- [2] V. Belle and H. J. Levesque, ‘Reasoning about continuous uncertainty in the situation calculus’, in *Proc. IJCAI*, (2013).
- [3] V. Belle and H. J. Levesque, ‘Reasoning about probabilities in dynamic systems using goal regression’, in *Proc. UAI*, (2013).
- [4] V. Belle and H. J. Levesque, ‘How to progress beliefs in continuous domains’, in *Proc. KR*, (2014).
- [5] V. Belle and H. J. Levesque, ‘PREGO: An Action Language for Belief-Based Cognitive Robotics in Continuous Domains’, in *AAAI*, (2014).
- [6] X. Boyen and D. Koller, ‘Tractable inference for complex stochastic processes’, in *Proc. UAI*, pp. 33–42, (1998).
- [7] G. De Giacomo, H.J. Levesque, and S. Sardina, ‘Incremental execution of guarded theories’, *ACM Transactions on Computational Logic*, **2**(4), 495–525, (2001).
- [8] Yi Fan, Minghui Cai, Naiqi Li, and Yongmei Liu, ‘A first-order interpreter for knowledge-based golog with sensing based on exact progression and limited reasoning’, in *Proc. AAAI*, (2012).

<sup>11</sup> As shown in [3], when the `BAT` is restricted to normally-distributed fluents and effectors, regression can be shown to yield expressions identical to a Kalman filter. However, Kalman filters only maintain the current world state, and so they correspond to a kind of progression [34].

- [9] A. Finzi, F. Pirri, and R. Reiter, ‘Open world planning in the situation calculus’, in *Proc. AAAI*, pp. 754–760, (2000).
- [10] Noah D. Goodman, Vikash K. Mansinghka, Daniel M. Roy, Keith Bonawitz, and Joshua B. Tenenbaum, ‘Church: a language for generative models’, in *UAI*, pp. 220–229. AUAI Press, (2008).
- [11] Henrik Grosskreutz and Gerhard Lakemeyer, ‘ccgolog – a logical language dealing with continuous change’, *Logic Journal of the IGPL*, **11**(2), 179–221, (2003).
- [12] Henrik Grosskreutz and Gerhard Lakemeyer, ‘Probabilistic complex actions in golog’, *Fundam. Inform.*, **57**(2-4), 167–192, (2003).
- [13] Andreas Herzig, Jerome Lang, and Pierre Marquis, ‘Action representation and partially observable planning using epistemic logic’, in *Proc. IJCAI*, pp. 1067–1072, (2003).
- [14] Luca Iocchi, Thomas Lukasiewicz, Daniele Nardi, and Riccardo Rosati, ‘Reasoning about actions with sensing under qualitative and probabilistic uncertainty’, *ACM Transactions on Computational Logic*, **10**, 5:1–5:41, (2009).
- [15] N. Kushmerick, S. Hanks, and D.S. Weld, ‘An algorithm for probabilistic planning’, *Artificial Intelligence*, **76**(1), 239–286, (1995).
- [16] H. Levesque and R. Reiter. High-level robotic control: Beyond planning. Position paper at AAAI Spring Symposium on Integrating Robotics Research, 1998.
- [17] H. Levesque, R. Reiter, Y. Lespérance, F. Lin, and R. Scherl, ‘Golog: A logic programming language for dynamic domains’, *Journal of Logic Programming*, **31**, 59–84, (1997).
- [18] P. Mateus, A. Pacheco, J. Pinto, A. Sernadas, and C. Sernadas, ‘Probabilistic situation calculus’, *Annals of Math. and Artif. Intell.*, **32**(1-4), 393–431, (2001).
- [19] Nicolas Meuleau, Emmanuel Benazera, Ronen I. Brafman, Eric A. Hansen, and Mausam, ‘A heuristic search approach to planning with continuous resources in stochastic domains’, *J. Artif. Intell. Res. (JAIR)*, **34**, 27–59, (2009).
- [20] B. Milch, B. Bonet, S. J. Russell, D. Sontag, D. L. Ong, and A. Kolobov, ‘BLOG: Probabilistic models with unknown objects’, in *Proc. IJCAI*, pp. 1352–1359, (2005).
- [21] Kevin P Murphy, *Machine learning: a probabilistic perspective*, The MIT Press, 2012.
- [22] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*, Morgan Kaufmann, 1988.
- [23] D. Poole, ‘Decision theory, the situation calculus and conditional plans’, *Electron. Trans. Artif. Intell.*, **2**, 105–158, (1998).
- [24] R. Reiter, *Knowledge in action: logical foundations for specifying and implementing dynamical systems*, MIT Press, 2001.
- [25] Raymond Reiter, ‘On knowledge-based programming with sensing in the situation calculus’, *ACM Trans. Comput. Log.*, **2**(4), 433–457, (2001).
- [26] M. Richardson and P. Domingos, ‘Markov logic networks’, *Machine learning*, **62**(1), 107–136, (2006).
- [27] S. Sanner, ‘Relational dynamic influence diagram language (rddl): Language description’, Technical report, Australian National University, (2011).
- [28] S. Sanner and K. Kersting, ‘Symbolic dynamic programming for first-order pomdps’, in *Proc. AAAI*, pp. 1140–1146, (2010).
- [29] R. B. Scherl and H. J. Levesque, ‘Knowledge, action, and the frame problem’, *Artificial Intelligence*, **144**(1-2), 1–39, (2003).
- [30] T.C. Son and C. Baral, ‘Formalizing sensing actions—a transition function based approach’, *Artificial Intelligence*, **125**(1-2), 19–91, (2001).
- [31] M. Thielscher, ‘Planning with noisy actions (preliminary report)’, in *Proc. Australian Joint Conference on Artificial Intelligence*, pp. 27–45, (2001).
- [32] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, MIT Press, 2005.
- [33] J. Van Benthem, J. Gerbrandy, and B. Kooi, ‘Dynamic update with probabilities’, *Studia Logica*, **93**(1), 67–96, (2009).
- [34] S. Vassos and H. Levesque, ‘On the Progression of Situation Calculus Basic Action Theories: Resolving a 10-year-old Conjecture’, in *Proc. AAAI*, pp. 1004–1009, (2008).
- [35] H. Younes and M. Littman, ‘PPDDL 1. 0: An extension to pddl for expressing planning domains with probabilistic effects’, Technical report, Carnegie Mellon University, (2004).
- [36] Z. Zamani, S. Sanner, P. Poupart, and K. Kersting, ‘Symbolic dynamic programming for continuous state and observation POMDPs.’, in *NIPS*, pp. 1403–1411, (2012).