Reasoning about Continuous Uncertainty in the Situation Calculus*

Vaishak Belle and **Hector J. Levesque**
Dept. of Computer Science
University of Toronto
Toronto, Ontario M5S 3H5, Canada
{vaishak, hector}@cs.toronto.edu

Abstract

Among the many approaches for reasoning about degrees of belief in the presence of noisy sensing and acting, the logical account proposed by Bacchus, Halpern, and Levesque is perhaps the most expressive. While their formalism is quite general, it is restricted to fluents whose values are drawn from discrete countable domains, as opposed to the continuous domains seen in many robotic applications. In this paper, we show how this limitation in their approach can be lifted. By dealing seamlessly with both discrete distributions and continuous densities within a rich theory of action, we provide a very general logical specification of how belief should change after acting and sensing in complex noisy domains.

1 Introduction

For many AI applications, and robotics in particular, it is not enough to deal with incomplete knowledge, where some formula ϕ might be unknown. One must also know which of ϕ or $\neg\phi$ is the more likely, and by how much. Perhaps the most general formalism for dealing with *degrees of belief* in formulas, and in particular, with how degrees of belief should evolve in the presence of noisy sensing and acting is the account proposed by Bacchus, Halpern, and Levesque [1999], henceforth BHL. Among its many properties, the BHL model shows precisely how beliefs can be made less certain by acting with noisy effectors, but made more certain by sensing (even when the sensors themselves are noisy).

The main advantage of a logical account like BHL is that it allows a specification of belief that can be partial or incomplete, in keeping with whatever information is available about the application domain. It does not require specifying a prior distribution over some random variables from which posterior distributions are then calculated, as in Kalman filters, for example [Dean and Wellman, 1991]. Nor does it require specifying the conditional independences among random variables and how these dependencies change as the result of actions, as in the temporal extensions to Bayesian networks [Pearl, 1988]. In the BHL model, some logical constraints are imposed on the initial state of belief. These constraints may be compatible with one or very many initial distributions and sets of independence assumptions. All the properties of belief will then follow at a corresponding level of specificity.

Subjective uncertainty is captured in the BHL account using a possible-world model of belief [Kripke, 1963; Hintikka, 1962; Fagin *et al.*, 1995]. In classical possible-world semantics, a formula ϕ is believed to be true when ϕ comes out true in all possible worlds that are deemed accessible. In BHL, the degree of belief in ϕ is defined as a normalized sum over the possible worlds where ϕ is true of some nonnegative *weights* associated with those worlds. (Inaccessible worlds are assigned a weight of zero.) To reason about belief change, the BHL model is then embedded in a rich theory of action and sensing provided by the situation calculus [McCarthy and Hayes, 1969; Reiter, 2001; Scherl and Levesque, 2003]. The BHL account provides axioms in the situation calculus regarding how the weight associated with a possible world changes as the result of acting and sensing. The properties of belief and belief change then emerge as a direct logical consequence of the initial constraints and these changes in weights.

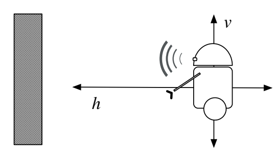

Figure 1: Robot operating in a 2-dimensional world.

To see a very simple example, imagine a robot located at some position on a two-dimensional grid, to the right of a wall parallel to the Y-axis as in Figure 1. Let h be the fluent representing the robot's horizontal distance to the wall. The fluent h would have different values in different possible worlds. In a BHL specification, each of these worlds might be given an initial weight. For example, a uniform distribution might give an equal weight of .1 to ten possible worlds where $h \in \{2, 3, \ldots, 11\}$. The degree of belief in a formula like $(h < 9)$ is then defined as a sum of the weights, and would lead here to a value of .7. The theory of action would then specify how these weights change as the result of acting (such as moving away or towards the wall) and sensing (such as obtaining a reading from a sonar aimed at the wall).

*The authors would like to thank Tim Capes and Steven Shapiro for valuable discussions, Fahiem Bacchus for feedback on a previous version of the paper, and the Natural Sciences and Engineering Research Council of Canada for financial support.

While this model of belief is widely applicable, it does have one serious drawback: it is ultimately based on the addition of weights and is therefore restricted to fluents having discrete countable values. This is in contrast to many robotic applications [Thrun *et al.*, 2005], where event and observation variables are characterized by well-known continuous distributions. There is no way to say in BHL that the initial value of h is any real number drawn from a uniform distribution on the interval $[2, 12]$. One would again expect the belief in $(h < 9)$ to be .7, but instead of being the result of *summing* weights, it must now be the result of *integrating* densities over a suitable space of values, something quite beyond the BHL approach.

The goal of this paper is to show how with minimal additional assumptions this serious limitation of BHL can be lifted. We present a *formal specification* of the degrees of belief in formulas with real-valued fluents (and other fluents too), and how belief changes as the result of acting and sensing. Our account will retain the advantages of BHL but work seamlessly with discrete probability distributions, probability densities, and perhaps most significantly, with difficult combinations of the two. (See Theorem 4 item 4 below.)

The rest of the paper is organized as follows. We review the formal preliminaries, and the BHL model in particular. We then show how the definition of belief in BHL can be reformulated as a different summation, which then provides sufficient foundation for our extension to continuous domains. In the final sections, we discuss related and future work.

2 Preliminaries

The language $\mathcal{L}$ of the situation calculus [McCarthy and Hayes, 1969] is a many-sorted dialect of predicate calculus, with sorts for *actions*, *situations* and *objects* (for everything else, and includes the set of reals $\mathbb{R}$ as a subsort). A situation represents a world history as a sequence of actions. A set of initial situations correspond to the ways the world might be initially. Successor situations are the result of doing actions, where the term $do(a, s)$ denotes the unique situation obtained on doing a in s. The term $do(\alpha, s)$, where α is the sequence $[a_1, \ldots, a_n]$ abbreviates $do(a_n, do(\ldots, do(a_1, s) \ldots))$. Initial situations are defined as 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 here, we use the variable ι to range over initial situations only.

The picture that emerges from the above is a set of trees, each rooted at an initial situation and whose edges are actions. In general, we want the values of predicate and functions to vary from situation to situation. For this purpose, $\mathcal{L}$ includes *fluents* whose last argument is always a situation. Here we assume without loss of generality that all fluents are functional.

Basic action theory

Following [Reiter, 2001], we model dynamic domains in $\mathcal{L}$ by means of a *basic action theory* $\mathcal{D}$, which consists of [1]

1. axioms $\mathcal{D}_0$ that describe what is true in the initial states, including S_0;
2. precondition and successor state axioms that describe the conditions under which actions are executable and the changes to fluents on executing actions respectively;
3. domain-independent *foundational* axioms, the details of which need not concern us here. See [Reiter, 2001].

We obtain many advantages by axiomatizing this way, including a simple solution to the frame problem [Reiter, 2001].

An agent reasons about actions by means of the entailments of $\mathcal{D}$, for which standard Tarskian models suffice. We assume henceforth that models *also* assign the usual interpretations to $=, <, >, 0, 1, +, \times, /, -, e, \pi$ and x^y (exponentials).[2]

Likelihood and degree of belief

The BHL model of belief builds on a treatment of knowledge by Scherl and Levesque [2003]. Here we present a simpler variant based on just two distinguished binary fluents l and p.

The term $l(a, s)$ is intended to denote the likelihood of action a in situation s. For example, suppose $sonar(z)$ is the action of reading the value z from a sensor that measures the distance to the wall, h.[3] We might assume that this action is characterized by a truncated Gaussian error model:[4]

$$\begin{aligned} &l(sonar(z), s) = u \equiv \\ &\quad (z \geq 0 \wedge u = \mathcal{N}(z - h(s); 0, 4)) \vee (z < 0 \wedge u = 0) \end{aligned} \tag{B1}$$

which stipulates that the difference between a nonnegative reading of z and the true value h is normally distributed with a variance of 4 and mean of 0. (A mean of 0 indicates that there is no systematic bias in the reading.) In general, the action theory $\mathcal{D}$ is assumed to contain for each action type A an additional *action likelihood axiom* of the form

4. $l(A(\vec{x}), s) = u \equiv \phi_A(\vec{x}, u, s)$,

where ϕ_A is a formula that characterizes the conditions under which action $A(\vec{x})$ has likelihood u in s. (Actions that have no sensing aspect should be given a likelihood of 1.)

Next, the p fluent determines a probability distribution on situations. The term $p(s', s)$ denotes the relative *weight* accorded to situation s' when the agent happens to be in situation s. The properties of p in initial states, which vary from domain to domain, are specified by axioms as part of $\mathcal{D}_0$. The following nonnegative constraint must be included:

$$\forall \iota, s.\ p(s, \iota) \geq 0 \ \wedge\ (p(s, \iota) > 0 \supset Init(s)) \tag{1}$$

Note that this is a stipulation about initial states ι only. But BHL provide a *successor state axiom* for p, and show that with an appropriate action likelihood axiom, the nonnegative constraint then continues to hold everywhere:

$$\begin{aligned} &p(s', do(a, s)) = u \ \equiv \\ &\qquad \exists s''\ [s' = do(a, s'') \wedge Poss(a, s'') \wedge \\ &\qquad\qquad u = p(s'', s) \times l(a, s'')] \\ &\quad \vee\ \neg\exists s''\ [s' = do(a, s'') \wedge Poss(a, s'') \wedge u = 0] \end{aligned} \tag{2}$$

[1] As usual, free variables in any of these axioms should be understood as universally quantified from the outside.

[2] Alternatively, one could specify axioms for characterizing the field of real numbers in $\mathcal{D}$. Whether or not reals with exponentiation is *first-order* axiomatizable remains a major open question.

[3] Naturally, we assume that the value z being read is not under the agent's control. See BHL for a precise rendering of this nondeterminism in terms of GOLOG operators [Reiter, 2001].

[4] Note that $\mathcal{N}$ is a continuous distribution involving π, e, exponentiation, and so on. Therefore, BHL always consider discrete probability distributions that *approximate* the continuous ones.

That is, the weight of situations s' relative to $do(a, s)$ is the weight of their predecessors s'' times the likelihood of a contingent on the successful execution of a at s''. One consequence of (1) and (2) is that $(p(s', s) > 0)$ will be true only when s' and s share the same history of actions.

With these two fluents elaborated, the degree of belief can be defined as the total weight of all accessible situations. Let ϕ be a formula with a single free variable of sort situation.[5] Then the degree of belief in ϕ in situation s is defined as an abbreviation by

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \sum_{\{s' : \phi[s']\}} p(s', s), \tag{3}$$

where γ, the normalization factor, is understood throughout as the same expression as the numerator but with ϕ replaced by *true*. Here, for instance, γ is $\sum_{s'} p(s', s)$. Note that we do not have to insist that s' and s share histories since $p(s', s)$ will be 0 otherwise. The BHL paper shows how these summations can be expressed using second-order quantification.

This *Bel* is well defined only when the sum over all those situations s' such that $\phi[s']$ holds is finite. This precludes continuous probability distributions, and for that matter, any domain that suggests an infinite set of situations agreeing on a formula. We now proceed to remove these restrictions.

3 Degree of Belief Reformulated

Prior to treating continuous fluents, our first objective is to reformulate (3), so that instead of summing over situations, we sum over *fluent values*. The next section then shows how this scheme generalizes from summation to *integration*.

First, some notation. We will use two forms of conditional terms as convenient abbreviations in logical formulas. The first is the usual "case" notation with curly braces:

$$z = \begin{cases} t_1 & \text{if } \psi \\ t_2 & \text{otherwise} \end{cases} \doteq (\psi \supset z = t_1) \wedge (\neg\psi \supset z = t_2)$$

The second involves a quantifier and a default value of 0, like in formula (2). If z is a variable, ψ is a formula and t is a term, we use $\langle z.\psi \rightarrow t\rangle$ as a logical term characterized as follows:

$$\begin{aligned} &\langle z.\psi \rightarrow t\rangle = u \doteq \\ &\qquad [(\exists z\psi) \supset \forall z(\psi \supset u = t)] \wedge [(\neg\exists z\psi) \supset u = 0)]. \end{aligned}$$

The notation says that when $\exists z\psi$ is true, the value of the term is t; otherwise, the value is 0. When t uses z (the usual case), this will be most useful if there is a *unique* z that satisfies ψ.

Returning to the task at hand, suppose that there are n fluents $f_1, f_2, \ldots, f_n$ in $\mathcal{L}$ which take no arguments other than the situation argument,[6] and that they take their values from some finite sets. We can rephrase (3) as follows:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \sum_{\vec{x}} \sum_{s'} \begin{cases} p(s', s) & \text{if } \bigwedge f_i(s') = x_i \wedge \phi[s'] \\ 0 & \text{otherwise} \end{cases}$$

In English: for each possible value of the fluents, sum over all possible situations and for each one, if the fluents have those values and ϕ holds, then use the p value, and 0 otherwise.

To arrive at a definition that eschews the summing of situations, we start with the case of initial situations. Let us recall the formalization of the situation calculus presented in [Levesque *et al.*, 1998] for multiple initial situations, which includes an axiom saying there is precisely one initial situation for any possible values of the fluents. For us, this is:

$$[\forall\vec{x}\exists\iota \bigwedge f_i(\iota) = x_i] \wedge [\forall\iota, \iota'. \bigwedge f_i(\iota) = f_i(\iota') \supset \iota = \iota'] \quad (*)$$

Under the assumption $(*)$, we can rewrite (3) for $s = S_0$ as

$$Bel(\phi, S_0) \doteq \frac{1}{\gamma} \sum_{\vec{x}} \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[\iota] \rightarrow p(\iota, S_0)\rangle \tag{4}$$

In fact, the two abbreviations turn out to be equivalent:

Theorem 1: *Let $\mathcal{D}$ be any basic action theory, ϕ any $\mathcal{L}$-formula, and suppose $(*)$ holds. Then the abbreviations for $Bel(\phi, S_0)$ in* (3) *and* (4) *compute the same real number.*

This shows that for S_0, summing over possible worlds can be replaced by summing over fluent values.

But something like (4) will not fly with non-initial situations. For those situations, the assumption that no two agree on all fluent values is untenable. To see why, imagine an action *left*(z) that moves the robot z units to the left (towards the wall) but that the motion stops if the robot hits the wall:

$$\begin{aligned} &h(do(a, s)) = u \equiv \\ &\qquad \neg\exists z(a = left(z)) \wedge u = h(s) \vee \\ &\qquad \exists z(a = left(z) \wedge u = \max(0, h(s) - z)). \end{aligned} \tag{B2}$$

In this case, if we have two initial situations that are identical except that $h = 3$ in one and $h = 4$ in the other, then the two distinct successor situations that result from doing *left*(4) would agree on all fluents (since both would have $h = 0$).

Ergo, we cannot sum over fluent values for non-initial situations unless we are prepared to count some fluent values more than once. However, what Reiter's solution to the frame problem gives us is a way of computing what holds in non-initial situations in terms of what holds in initial ones, which can be used for computing belief at arbitrary successors of S_0:

$$\begin{aligned} &Bel(\phi, do(\alpha, S_0)) \doteq \\ &\qquad \frac{1}{\gamma} \sum_{\vec{x}} \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[do(\alpha, \iota)] \rightarrow \\ &\qquad\qquad p(do(\alpha, \iota), do(\alpha, S_0))\rangle \end{aligned} \tag{5}$$

To say more about how (and why) this definition works, we first note that by (1) and (2), p will be 0 unless its two arguments share the same history. So the s' argument of p in (3) is written as $do(\alpha, \iota)$ in (5). By ranging over all fluent values, we range over all initial ι as before (without ever having to deal with fluent values in non-initial situations). Of course, we test that the ϕ holds and use the p weight in the appropriate non-initial situation. This gives us the following:

Theorem 2: *Let $\mathcal{D}$ be any basic action theory with $(*)$ initially, ϕ any $\mathcal{L}$-formula, and α any sequence of ground actions terms. Then the abbreviations for $Bel(\phi, do(\alpha, S_0))$ in* (3) *and* (5) *compute the same real number.*

[5] The ϕ is usually written either with the situation variable suppressed or with a distinguished variable *now*. Either way, $\phi[t]$ is used to denote the formula with that variable replaced by t.

[6] We will return to this assumption in the penultimate section.

Thus, by incorporating a simple constraint on initial situations, we now have a notion of belief that does not require summing over situations. Our reformulation only applies when we are given an explicit sequence α of actions (including the sensing ones), but this is just what we would expect to be given for the projection problem [Reiter, 2001]. In fact, we can use regression on the ϕ and the p to reduce the belief formula (5) to a formula involving initial situations only.

4 From Weights to Densities

The uncountable nature of continuous domains precludes summing over possible situations. In this section, we present a new formalization of belief in terms of *integrating* over fluent values, and relating that to the space of situations based on the development in the preceding section.

Allowing real-valued fluents implies that there will be uncountably many initial situations. Imagine, for example, that h can now be any nonnegative real number. Then for any nonnegative real x there will be an initial situation where $(h = x)$ is true. Suppose further that $\mathcal{D}_0$ includes:

$$p(\iota, S_0) = \begin{cases} .1 & \text{if } 2 \leq h(\iota) \leq 12 \\ 0 & \text{otherwise} \end{cases}$$

which says that the true value of h initially is drawn from a uniform distribution on [2,12]. Then there are uncountably many situations where p is non-zero initially. So the p fluent now needs to be understood as a *density*, not as a weight. In particular, for any x, we would expect the initial degree of belief in the formula $(h = x)$ to be 0, but in $(h \leq 12)$ to be 1.

But there is more to the story. Interesting subtleties arise with p in non-initial situations. For example, if the robot were to do *left*(4), there would be an uncountable number of situations agreeing on $h = 0$ (namely, those where $2 \leq h \leq 4$ was true initially). In a sense, the point $h = 0$ now has *weight* (the degree of belief in $h = 0$ should be .2), while the other points $h \in (0, 8]$ retain their densities. In effect, we have moved from a density to a *mixed* distribution on h.

One of the advantages of our BHL style approach is that we will not need to specify how to handle changing densities and distributions like this. These will emerge as side-effects. Our proposal for generalizing BHL to the continuous case keeps its simplicity, and consists of:

- a standard situation calculus basic action theory, including $(*)$ to accommodate multiple initial situations;
- action likelihood axioms for the distinguished fluent l, such as the one provided for *sonar* earlier;
- the initial constraint (1) and the successor state axiom (2) for the distinguished fluent p. Note that we now interpret $p(s', s)$ as the *density* of s' when the agent is in s.

Now partition the fluents in $\mathcal{L}$ into two groups, those whose domain is countable, $f_1, \ldots, f_n$, and those whose domain is $\mathbb{R}$, $g_1, \ldots, g_m$. Let us first define the *belief density* of ϕ as:

$$\begin{aligned} &\mathit{Density}(\vec{x}, \vec{y}, \phi, do(\alpha, S_0)) \doteq \\ &\quad \langle\, \iota. \bigwedge g_i(\iota) = x_i \wedge \bigwedge f_j(\iota) = y_j \wedge \phi[do(\alpha, \iota)] \\ &\qquad\qquad \rightarrow\ p(do(\alpha, \iota), do(\alpha, S_0))\,\rangle \end{aligned}$$

Then the *degree of belief* in ϕ is simply an abbreviation for:

$$\mathit{Bel}(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x}} \sum_{\vec{y}} \mathit{Density}(\vec{x}, \vec{y}, \phi, s) \tag{6}$$

That is, the belief in ϕ is obtained by ranging over all possible fluent values, and integrating and summing the densities of situations where ϕ holds. As before, by insisting on an explicit world history, the ι need only range over initial situations, giving us the exact correspondence with fluent values. The normalization factor γ is the numerator but with ϕ replaced by *true*. Integration and infinite sums can be expressed using second-order quantification; see the appendix.

This completes our new definition of belief. We will show that it does the appropriate thing using an example in the next section and its connection to Bayesian conditioning below.

First note, however, that definition (6) will not result in a numeric value when γ is undefined or when $\gamma = 0$, and nor should it. For example, imagine a $\mathcal{D}_0$ with the following:

$$p(\iota, S_0) = \begin{cases} b & \text{if } g(\iota) \text{ is rational} \\ 0 & \text{otherwise} \end{cases}$$

This defines a set whose indicator function is not (Riemann) integrable[7] and no amount of normalization will fix this. Similarly, we get that $\gamma = 0$ after an action whose likelihood is 0, for example, after *sonar*(-3), assuming $\mathcal{D}$ contains (B1).

There are, however, legitimate cases of belief that definition (6) does not yet handle. There are two limitations:[8]

Limitation 1: only continuous sensors. We have lifted BHL's successor state axiom for p without any alterations, multiplying $p(s', s)$ by $l(a, s')$ to obtain $p(do(a, s'), do(a, s))$. Since p is a density, l must be one as well. This means that we cannot yet handle a *discrete sensor* that has a finite number of possible readings whose weights sum to 1. (More precisely, the likelihood function cannot give a non-zero result only to a finite number of values.) So although we do not have to approximate Gaussian error models (or any other continuous models) for our sensors as would BHL, a different *weight*-based notion of likelihood is needed for discrete sensors.

Limitation 2: only deterministic actions. When we perform actions that simply change the world, such as *left*, we assume that the probability density is completely transferred from the start to the final situation, which is a form of *imaging* [Lewis, 1976]. The likelihood density of these actions is uniformly 1. So *noisy actions*, which can result in different outcomes with different likelihoods, are not yet handled. Belief would require further summing/integrating over all these outcomes.

Belief change and Bayesian conditioning

A standard model for belief change wrt noisy sensing is *Bayesian conditioning* [Pearl, 1988], which rests on two significant assumptions. First, sensors do not physically change

[7] In the calculus community, the concept of a *gauge integral* has been studied [Swartz, 2001], which is a generalization of the Riemann integral, and allows for the integration of indicator functions such as the one above. We have chosen to remain within the framework of classical integration, but other accounts may be useful.

[8] We do have proposals for handling both of these limitations. The details are left for a longer version of the paper.

the world, and second, conditioning on a random variable g is the same as conditioning on the event of observing g. When a similar set of assumptions are imposed as axioms in an action theory $\mathcal{D}$, we obtain an identical sensor fusion model.

Begin by stipulating that actions are either *physical* or of the *sensing* type [Scherl and Levesque, 2003]. Now, if $obs(z)$ senses the true value of fluent g, then assume the sensor error model to be:

$$l(obs(z), s) = u \;\equiv\; u = Err(z, g(s))$$

where $Err(u_1, u_2)$ is some expression with only two free variables, both numeric. This captures the idea that the error model of a sensor measuring g depends only on the true value of g, and is independent of other factors.[9] Then

Theorem 3: *Suppose $\mathcal{D}$ is as above, ϕ is any $\mathcal{L}$-formula mentioning only g and u is a variable from $\{x_1, \ldots, x_m, y_1, \ldots, y_n\}$:*

$$\mathcal{D} \models Bel(\phi, do(obs(z), S_0)) = \frac{\int_{\vec{x}} \sum_{\vec{y}} Density(\vec{x}, \vec{y}, \phi \wedge g = u, S_0) \times Err(z, u)}{\int_{\vec{x}} \sum_{\vec{y}} Density(\vec{x}, \vec{y}, g = u, S_0) \times Err(z, u)}$$

That is, the *posterior belief* in ϕ is obtained from the prior density and the error likelihoods for all points where ϕ holds given that z is observed, normalized over all points.[10] The usual case for ϕ are formulas such as $b \leq g \leq c$, which is estimated from the prior and error likelihoods for all points in the range $[b, c]$. More generally however, and unlike many probabilistic formalisms, we are able to reason about any logical property ϕ of the random variable g being measured.

5 Example

To reason about the beliefs of our robot, let us build a simple basic action theory $\mathcal{D}$. Suppose the robot's vertical position is given by a fluent v. Now imagine a p as follows:

$$p(\iota, S_0) = \begin{cases} .1 \times \mathcal{N}(v(\iota); 0, 16) & \text{if } 2 \leq h(\iota) \leq 12 \\ 0 & \text{otherwise} \end{cases} \tag{B3}$$

This says that the value of v is normally distributed about the horizontal axis with variance 16, and independently, that the value of h is uniformly distributed between 2 and 12.[11]

We specified the way h changes in (B2). Imagine also:

$$\begin{aligned} v(do(a, s)) = u \equiv \\ &\neg\exists z(a = up(z)) \wedge u = v(s) \vee \\ &\exists z(a = up(z) \wedge u = v(s) + z). \end{aligned} \tag{B4}$$

This says that performing $up(z)$ increments the current value of v by z, but nothing else affects it. Finally, suppose the sonar from (B1) is the only sensor at the robot's disposal.[12] Then:

[9]This is not required in general. We might imagine, for example, that the sonar's accuracy depends on the room temperature, and such accounts are expressible in our formalism.

[10]As mentioned earlier, if the normalizing factor $\gamma = 0$, which corresponds to the case of conditioning on an event that has a 0 probability, Bayes rule is not defined, and neither is the expression here.

[11]Initial beliefs can also be specified for $\mathcal{D}_0$ using *Bel* directly. We remark that a simple distribution is chosen for illustrative purposes. In general, the p specification does not require the variables to be independent, nor does it have to mention all variables.

[12]For a more elaborate example involving multiple competing sensors, see [Belle and Levesque, 2013].

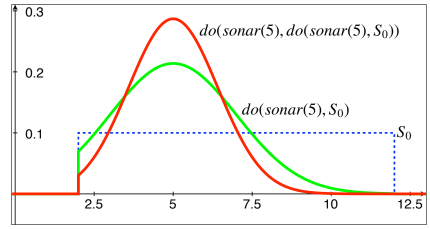

Figure 2: Belief density change for h at S_0 (in blue), after sensing 5 (in green), and after sensing 5 twice (in red).

Theorem 4: *Let $\mathcal{D}$ contain the union of* (1), (2), ($*$) *and* (B1)–(B4). *Then the following are logical entailments of $\mathcal{D}$:*

1. $Bel([h = 3 \vee h = 4 \vee h = 7], S_0) = 0.$
   Although we are integrating a density function $q(x_1, x_2)$ over all real values, $q(x_1, x_2) = 0$ unless $x_1 \in \{3, 4, 7\}$.

2. $Bel(h \leq 9, S_0) = .7.$
   Here we are integrating a function that is 0 except when $2 \leq x_1 \leq 9$. So this is $\int_{\mathbb{R}} \int_2^9 .1 \times \mathcal{N}(x_2; 0, 16)\, \mathrm{d}x_1\, \mathrm{d}x_2 = .7$.

3. $Bel(h > 7v, S_0) \approx .6.$
   Beliefs about any mathematical expression involving the random variables, even when that does not correspond to well known density functions, are entailed.

4. $Bel(h = 0, do(left(4), S_0)) = .2.$
   Here a *continuous* distribution evolves into a *mixed* one. By (B2), $h = 0$ holds after the action iff $h \leq 4$ held before. So this results in $\int_{\mathbb{R}} \int_2^4 .1 \times \mathcal{N}(x_2; 0, 16)\, \mathrm{d}x_1\, \mathrm{d}x_2 = .2$.

5. $Bel(h \leq 3, do(left(4), S_0)) = .5.$
   Bel's definition is amenable to a set of h values, where one value has a weight of .2, and all the other real values have a uniformly distributed density of .1.

6. $Bel([\exists a, s.\ now = do(a, s) \wedge h(s) > 1], do(left(4), S_0)) = 1.$
   It is possible to refer to earlier or later situations using *now* as the current situation. This says that after moving, there is full belief that $(h > 1)$ held before the action.

7. $Bel(h = 4, do([left(4), left(-4)], S_0)) = .2$
   $Bel(h = 4, do([left(-4), left(4)], S_0)) = 0.$
   The point $h = 4$ has 0 weight initially (like in item 1). Moving leftwards *first* means many points "collapse", and so this point (now having h value 0) gets .2 weight which is retained on moving away. But not vice versa.

8. $Bel(-1 \leq v \leq 1, do(left(4), S_0)) =$
   $Bel(-1 \leq v \leq 1, S_0) = \int_{-1}^{1} \mathcal{N}(x_2; 0, 16)\mathrm{d}x_2.$
   Owing to the solution to the frame problem, belief in v is unaffected by a lateral motion. For $v \in [-1, 1]$, it is the area between $[-1, 1]$ bounded by the specified Gaussian.

9. $Bel(-1 \leq v \leq 7, do(up(2.5), S_0)) = Bel(-3.5 \leq v \leq 4.5, S_0).$
   After the action $up(2.5)$, the Gaussian for v's value has its mean "shifted" by 2.5 because the density associated with $v = x_2$ initially is now associated with $v = x_2 + 2.5$.

10. $Bel(h \leq 9, do(sonar(5), S_0)) \approx .97.$
    $Bel(h \leq 9, do([sonar(5), sonar(5)], S_0)) \approx .99.$

Compared to item 2, belief in $h \leq 9$ is sharpened by obtaining a reading of 5 on the sonar, and sharpened to almost certainty on a second reading of 5. This is because the p function, according to (2), incorporates the likelihood of each *sonar*(5) action. Starting with the density function in item 2, each sensor reading multiplies the expression to be integrated by $\mathcal{N}(5 - x_1; 0, 4)$, as given by (B1). These changing densities are shown in Figure 2.

6 Related Work and Discussions

Belief update via sensor information has been a fundamental concern of autonomous agent formalisms. On the one hand, we have probabilistic formalisms such as Bayesian networks [Pearl, 1988; Lerner *et al.*, 2002], and Kalman and particle filters [Dean and Wellman, 1991; Fox *et al.*, 2003]. These have difficulties handling strict uncertainty. Moreover, since rich models of actions are rarely incorporated, shifting conditional dependencies and distributions are hard to address in a general way. While there are graphical formalisms with an account of actions, such as [Darwiche and Goldszmidt, 1994; Hajishirzi and Amir, 2010], they too have difficulties handling strict uncertainty and quantification. To the best of our knowledge, no existing probabilistic formalism handles changes in state variables like those considered here.

Logical formalisms, on the other hand, such as [Fagin and Halpern, 1994; Bacchus, 1990], provide means to specify properties about the domain together with probabilities about propositions. Related to these are relational probabilistic models [Ng and Subrahmanian, 1992; Milch *et al.*, 2005] and Markov logics [Richardson and Domingos, 2006; Domingos *et al.*, 2006; Tran and Davis, 2008; Choi *et al.*, 2010]. Here too explicit actions are seldom addressed.

Action logics share the motivation of the work here. Recent proposals, for example [Van Benthem *et al.*, 2009], treat sensor fusion. However, these and related frameworks [Halpern and Tuttle, 1993; Kushmerick *et al.*, 1995] are propositional. Proposals based on the situation and fluent calculi are first-order [Bacchus *et al.*, 1999; Poole, 1998; Boutilier *et al.*, 2000; Mateus *et al.*, 2001; Shapiro, 2005; Gabaldon and Lakemeyer, 2007; Fritz and McIlraith, 2009; Belle and Lakemeyer, 2011; Thielscher, 2001], but none of them deal with continuous sensor noise, and nor do the extensions for continuous processes [Herrmann and Thielscher, 1996]. We are also not aware of any logical approach for uncertainty, dynamical or not, that deals with the integration of continuous variables within the language.

Before wrapping up, let us remark that one of the strong limitations of our work from the point of view of situation calculus basic action theories is that our functional fluents take no argument other than the situation term. While we do allow the *values* of the fluents to range over any set (including the reals), fluents are also usually allowed to take *arguments* from any set, including infinite ones. In probabilistic terms, this would correspond to having a joint probability distribution over infinitely many, perhaps uncountably many, random variables. We know of no existing work of this sort, and we have as yet no good ideas about how to deal with it.

7 Conclusions

Robotic applications have to deal with numerous sources of uncertainty, the main culprit being sensor noise. Probabilistic error models have proven to be powerful in state estimation, allowing the beliefs of a robot to be strengthened over time. But to use these models, the modeler is left with the difficult task of deciding how the domain is to be captured in terms of random variables, and shifting conditional independences and distributions. In the BHL model, one simply provides a specification of some initial beliefs, characterizes the physical laws of the domain, and suitable posterior beliefs are entailed. The applicability of BHL to real-world robotics was limited, however, by its inability to handle continuous distributions, a limitation we lift in this paper. By recasting the assessment of belief in terms of fluent values, we now seamlessly combine the situation calculus with discrete probability distributions, densities and combinations of the two. We demonstrated that distributions evolve appropriately after actions, emerging as a side-effect of the general specification. Nondeterminism in acting and sensing is treated in a longer version of the paper.

Regarding future work, on the more computational side, we have noted that the definition of belief seems amenable to regression. (See [Belle and Levesque, 201x] for preliminary results.) Assuming we are given priors and likelihoods drawn from tractable continuous distributions [Box and Tiao, 1973], one promising line of research would be to see if posteriors can be calculated efficiently by regressing belief after action to a formula concerning the initial state only.

Appendix: Integrals in Logic

Logical formulas can be used to characterize a variety of sorts of integrals. Here we show the simplest possible case: the definite integral from $-\infty$ to ∞ of a continuous real-valued function of one variable. Other complications, including functions with multiple variables, are treated in an extended version of the paper.

We begin by introducing a notation for limits to positive infinity. For any logical term t and variable x, we let $\lim_{x\to\infty} t$ stand for a term characterized by the following:

$$\lim_{x\to\infty} t = z \;\doteq\; \forall u(u > 0 \supset \exists m \forall n(n > m \supset \left|z - t_n^x\right| < u)).$$

The variables u, m, and n are understood to be chosen here not to conflict with any of the variables in x, t, and z.

Next, for any variable x and terms a, b, and t, we introduce a term $\text{INT}[x, a, b, t]$ denoting the definite integral of t over x from a to b:

$$\text{INT}[x, a, b, t] \;\doteq\; \lim_{n\to\infty} h \cdot \sum_{i=1}^{n} t^x_{(a+h\cdot i)}$$

where h stands for $(b - a)/n$. The variable n is chosen not to conflict with any of the other variables. The summation term is a finite sum (for each n) as defined in the BHL paper. Finally, we define the definite integral of t over all real values of x by the following:

$$\int_x t \;\doteq\; \lim_{u\to\infty} \lim_{v\to\infty} \text{INT}[x, -u, v, t].$$

The main result for this logical abbreviation is the following:

Theorem 5: *Let g be a function symbol of $\mathcal{L}$ standing for a function from $\mathbb{R}$ to $\mathbb{R}$, and let c be a constant symbol of $\mathcal{L}$. Let M be any logical interpretation of $\mathcal{L}$ such that the function g^M is continuous everywhere. Then we have the following:*

$$\textit{If} \int_{-\infty}^{\infty} g^M(x)\,.\,dx \;=\; c^M \;\; \textit{then} \;\; M \models (c = \int_x g(x)).$$

References

[Bacchus *et al.*, 1999] 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.

[Bacchus, 1990] F. Bacchus. *Representing and Reasoning with Probabilistic Knowledge*. MIT Press, 1990.

[Belle and Lakemeyer, 2011] V. Belle and G. Lakemeyer. A semantical account of progression in the presence of uncertainty. In *Proc. AAAI*, pages 165–170, 2011.

[Belle and Levesque, 2013] V. Belle and H. J. Levesque. Robot location estimation in the situation calculus. In *Symposium on Logical Formalizations of Commonsense Reasoning*, 2013.

[Belle and Levesque, 201x] V. Belle and H. J. Levesque. Reasoning about probabilities in dynamic systems using goal regression. In preparation, 201x.

[Boutilier *et al.*, 2000] C. Boutilier, R. Reiter, M. Soutchanski, and S. Thrun. Decision-theoretic, high-level agent programming in the situation calculus. In *Proc. AAAI*, pages 355–362, 2000.

[Box and Tiao, 1973] G. E. P. Box and G. C. Tiao. *Bayesian inference in statistical analysis*. Addison-Wesley, 1973.

[Choi *et al.*, 2010] J. Choi, E. Amir, and D. J. Hill. Lifted inference for relational continuous models. In *Proc. UAI*, pages 126–134, 2010.

[Darwiche and Goldszmidt, 1994] A. Darwiche and M. Goldszmidt. Action networks: A framework for reasoning about actions and change under uncertainty. In *Proc. UAI*, pages 136–144, 1994.

[Dean and Wellman, 1991] T. Dean and M. Wellman. *Planning and control*. Morgan Kaufmann Publishers Inc., 1991.

[Domingos *et al.*, 2006] P. Domingos, S. Kok, H. Poon, M. Richardson, and P. Singla. Unifying logical and statistical AI. In *Proc. AAAI*, pages 2–7, 2006.

[Fagin and Halpern, 1994] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.

[Fagin *et al.*, 1995] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.

[Fox *et al.*, 2003] D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *Pervasive Computing, IEEE*, 2(3):24–33, 2003.

[Fritz and McIlraith, 2009] C. Fritz and S. A. McIlraith. Computing robust plans in continuous domains. In *Proc. ICAPS*, pages 346–349, 2009.

[Gabaldon and Lakemeyer, 2007] A. Gabaldon and G. Lakemeyer. ESP: A logic of only-knowing, noisy sensing and acting. In *Proc. AAAI*, pages 974–979, 2007.

[Hajishirzi and Amir, 2010] H. Hajishirzi and E. Amir. Reasoning about deterministic actions with probabilistic prior and application to stochastic filtering. In *Proc. KR*, 2010.

[Halpern and Tuttle, 1993] J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *J. ACM*, 40:917–960, 1993.

[Herrmann and Thielscher, 1996] Christoph S. Herrmann and Michael Thielscher. Reasoning about continuous processes. In *AAAI/IAAI, Vol. 1*, pages 639–644, 1996.

[Hintikka, 1962] Jaakko Hintikka. *Knowledge and belief: an introduction to the logic of the two notions*. Cornell University Press, 1962.

[Kripke, 1963] Saul Kripke. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94, 1963.

[Kushmerick *et al.*, 1995] N. Kushmerick, S. Hanks, and D.S. Weld. An algorithm for probabilistic planning. *Artificial Intelligence*, 76(1):239–286, 1995.

[Lerner *et al.*, 2002] U. Lerner, B. Moses, M. Scott, S. McIlraith, and D. Koller. Monitoring a complex physical system using a hybrid dynamic bayes net. In *Proc. UAI*, pages 301–310, 2002.

[Levesque *et al.*, 1998] H. J. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Electron. Trans. Artif. Intell.*, 2:159–178, 1998.

[Lewis, 1976] D. Lewis. Probabilities of conditionals and conditional probabilities. *The philosophical review*, pages 297–315, 1976.

[Mateus *et al.*, 2001] 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.

[McCarthy and Hayes, 1969] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, pages 463–502, 1969.

[Milch *et al.*, 2005] B. Milch, B. Marthi, S. J. Russell, D. Sontag, D. L. Ong, and A. Kolobov. BLOG: Probabilistic models with unknown objects. In *Proc. IJCAI*, pages 1352–1359, 2005.

[Ng and Subrahmanian, 1992] R. Ng and V.S. Subrahmanian. Probabilistic logic programming. *Information and Computation*, 101(2):150–201, 1992.

[Pearl, 1988] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.

[Poole, 1998] D. Poole. Decision theory, the situation calculus and conditional plans. *Electron. Trans. Artif. Intell.*, 2:105–158, 1998.

[Reiter, 2001] R. Reiter. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT Press, 2001.

[Richardson and Domingos, 2006] M. Richardson and P. Domingos. Markov logic networks. *Machine learning*, 62(1):107–136, 2006.

[Scherl and Levesque, 2003] R. B. Scherl and H. J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.

[Shapiro, 2005] S. Shapiro. Belief change with noisy sensing and introspection. In *NRAC Workshop*, pages 84–89, 2005.

[Swartz, 2001] C. Swartz. *Introduction to gauge integrals*. World Scientific Publishing Company Incorporated, 2001.

[Thielscher, 2001] M. Thielscher. Planning with noisy actions (preliminary report). In *Proc. Australian Joint Conference on Artificial Intelligence*, pages 27–45, 2001.

[Thrun *et al.*, 2005] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.

[Tran and Davis, 2008] S. D. Tran and L. S. Davis. Event modeling and recognition using markov logic networks. In *Proc. ECCV*, pages 610–623, 2008.

[Van Benthem *et al.*, 2009] J. Van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96, 2009.