

Reasoning about Discrete and Continuous Noisy Sensors and Effectors in Dynamical Systems

Vaishak Belle^{a,b,*}, Hector J. Levesque^c

^aUniversity of Edinburgh, Edinburgh, United Kingdom.

^bThe Alan Turing Institute, London, United Kingdom.

^cUniversity of Toronto, Toronto, Canada.

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 finite domains, as opposed to the continuous domains seen in many robotic applications. In this work, we show how this limitation in that 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.

Keywords: Knowledge representation, Reasoning about action, Reasoning about knowledge, Reasoning about uncertainty, Cognitive robotics

1. Introduction

On numerous occasions it has been suggested that the formalism [the situation calculus] take uncertainty into account by attaching probabilities to its sentences. We agree that the formalism will eventually have to allow statements about the probabilities of events, but attaching probabilities to all statements has the following objections:

- 1. It is not clear how to attach probabilities to statements containing quantifiers in a way that corresponds to the amount of conviction people have.*
- 2. The information necessary to assign numerical probabilities is not ordinarily available. Therefore, a formalism that required numerical probabilities would be epistemologically inadequate.*

– McCarthy and Hayes [1].

Much of high-level AI research is concerned with the behaviour of some putative agent, such as an autonomous robot, operating in an environment. Broadly speaking, an intelligent agent interacting with a dynamic and incompletely known world grapples with two special sorts of reasoning problems. First, because the world is *dynamic*, it will need to reason about change: how its actions affect the state of the world. Pushing an object on a table, for example, may cause it to fall on the floor, where it will remain unless picked up. Second, because the world is incompletely known, the agent will need to make do with partial specifications about what is true. As a result, the agent will often need to augment what it believes about the world by performing perceptual actions, using sensors of one form or another.

For many AI applications, and robotics in particular, these reasoning problems are more involved. Here, it is not enough to deal with incomplete knowledge, where some formula ϕ might be unknown. One must also know which

*Corresponding author.

Email addresses: vaishak@ed.ac.uk (Vaishak Belle), hector@cs.toronto.edu (Hector J. Levesque)




Figure 1: A simple robot.

of ϕ or $\neg\phi$ is the more *likely*, and by how much. In addition, both the sensors and the effectors that the agent uses to modify its world are often subject to uncertainty in that they are *noisy*.

To see a very simple example, imagine a robot moving towards a wall as shown in Figure 1, and a certain distance h from it. Suppose the robot can move towards and away from the wall, and it is equipped with a distance sensor aimed at the wall. Here, the robot may not know the true value of h but may believe that it takes values from some set, say $\{2, \dots, 11\}$. If the sensor is noisy, a reading of, say, 5 units, does not guarantee that the agent is actually 5 units from the wall, although it should serve to increase the agent's degree of belief in that fact. Analogously, if the robot intends to move by 1 unit and the effector is noisy, it may end up moving by 0.9 units, which the agent does not get to observe. Be that as it may, the robot's degree of belief that it is closer to the wall should increase.

While many proposals have appeared in the literature to address such concerns (*cf.* penultimate section), very few are embedded in a general theory of action whilst supporting features like disjunction and quantification. For example, graphical models such as Bayesian networks can represent and reason about the probabilistic dependencies between random variables, and how that might change over time. However, it lacks first-order features and a rich account of actions. Relational graphical models, including Markov logic networks [2], borrow devices from first-order logic to allow the succinct modelling of relational dependencies, but ultimately they are purely syntactic extensions to graphical models, and do not attempt to address the deeper issues pertaining to the specification of probabilities in the presence of logical connectives and quantifiers. Building on first-order accounts of probabilistic reasoning [3, 4], 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 [5], 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 [6]. 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 [7]. 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 [8, 9, 10]. In classical possible-world semantics, a formula ϕ is believed to be true when ϕ holds 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 [1, 11, 12]. 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.

For example, suppose h is a fluent representing the robot's horizontal distance to the wall in Figure 1. 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 $h \in \{2, 3, \dots, 11\}$. 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). Naturally, the logical language permits weaker specifications, involving disjunctions and quantifiers, and the appropriate behavior would still emerge.

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 finite values. This is in stark contrast to robotics and machine learning applications [13, 14, 15], where event and observation variables are characterized by continuous distributions, or perhaps combinations of discrete and continuous ones. 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.

So, on the one hand, the BHL account and others like it can be seen as general formal theories that attempt to address important philosophical problems such as those raised by McCarthy and Hayes above. But on the other, a serious criticism levelled at this line of work, and indeed at much of the work in reasoning about action, is that the theory is far removed from the kind of probabilistic uncertainty and noise seen in typical robotic applications.

The goal of this work is to show how with minimal additional assumptions this serious limitation of BHL can be lifted.¹ By lifting this limitation, one obtains, for the first time, a logical language for representing real-world robotic specifications without any modifications, but also extend beyond it by means of the logical features of the underlying framework. In particular, 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. More broadly, we believe the model of belief proposed in this work provides the necessary bridge between logic-based reasoning modules, on the one hand, and probabilistic specifications as seen in real-world data-intensive applications, on the other.

The paper is organized as follows. We first review the formal preliminaries, the BHL model in particular, and introduce definitions for modeling continuous probability distributions. 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. We then discuss how this model can be extended for noisy acting, and combinations of discrete and continuous properties. We conclude after discussing related work.

2. A Theory of Action

Our account is formulated in the language of the situation calculus [1], as developed in [11]. The situation calculus is a special-purpose knowledge representation formalism for reasoning about dynamical systems. Informally, the formalism is best understood by arranging the world in terms of three kinds of things: *situations*, *actions* and *objects*. Situations represent “snapshots,” and can be viewed as possible histories. A set of initial situations correspond to the ways the world can be prior to the occurrence of actions. The result of doing an action, then, leads to a successor (non-initial) situation. Naturally, dynamic worlds change the properties of objects, which are captured using predicates and functions whose last argument is always a situation, called *fluents*.

2.1. The Logical Language

Formally, the language \mathcal{L} of the situation calculus is a many-sorted dialect of predicate calculus, with sorts for *actions*, *situations* and *objects* (for everything else). (We do not review standard predicate logic here; see, for example, [17, 18]. We further assume familiarity with the notions of *models*, *structures*, *satisfaction* and *entailment*.) In full length, let \mathcal{L} include:

¹A preliminary version of this work was discussed in [16]. That work was limited to noisy sensors, but assumed deterministic (*i.e.*, noise-free) effectors.

- logical connectives $\neg, \forall, \wedge, =$, with other connectives such as \supset understood for their usual abbreviations;
- an infinite supply of *variables* of each sort;
- an infinite supply of *constant* symbols of the sort object;
- for each $k \geq 1$, *object function* symbols g_1, g_2, \dots of type $(action \cup object)^k \rightarrow object$;
- for each $k \geq 0$, *action function* symbols A_1, A_2, \dots of type $(action \cup object)^k \rightarrow action$;
- a special *situation function* symbol $do: action \times situation \rightarrow situation$;
- a special predicate symbol $Poss: action \times situation$;²
- for each $k \geq 0$, *fluent function* symbols f_1, f_2, \dots of type $(action \cup object)^k \times situation \rightarrow object$;
- a special constant S_0 to represent the actual initial situation.

To reiterate, apart from some syntactic particulars, the logical basis for the situation calculus is the regular (many-sorted) predicate calculus.³ So, terms and well-formed formulas are defined inductively, as usual, respecting sorts. See, for example, [11] for an exposition.

We follow some conventions in the ways we use Latin and Greek alphabets: a for both terms and variables of the action sort (the context would make this clear); s for terms and variables of the situation sort (the context would make this clear); and finally, x, u, v, z, n , and y to range over variables of the object sort. We let ϕ and ψ range over formulas, and Σ over sets of formulas. (These may be further decorated using superscripts or subscripts.)

We sometimes suppress the situation term in a formula ϕ , or use a distinguished variable now , to denote the current situation. Either way, we let $\phi[t]$ denote the formula with the restored situation term t .

We will often use the usual “case” notation with curly braces as a convenient abbreviation for a logical formula:

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

Finally, for convenience, we often introduce formula and term abbreviations that are meant to expand as \mathcal{L} -formulas. For example, we might introduce a new formula A by $A \doteq \phi$, where $\phi \in \mathcal{L}$. Then any expression $E(A)$ containing A is assumed to mean $E(\phi)$. Analogously, if we introduce a new term t by $t = u \doteq \phi(u)$ then any expression $E(t)$ is assumed to mean $\exists u(E(u) \wedge \phi(u))$.

Dynamic worlds are enabled by performing actions, and in the language, this is realized using the *do* operator. That is, the result of doing an action a at situation s is the situation $do(a, s)$. Functional fluents, which take situations as arguments, may then have different values at different situations, thereby capturing *changing properties* of the world. As noted, the constant S_0 is assumed to give the actual initial state of the domain, but the agent may consider others possible that capture the beliefs and ignorance of the agent. In general, we say a situation is an *initial* one when it is a situation without a *predecessor*:

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

The picture that emerges is that situations can be structured as a set of trees, each rooted at an initial situation and whose edges are actions. More conventions: we use ι to range over such initial situations only, and let α denote sequences of action terms or variables, and freely use this with *do*, that is, if $\alpha = [a_1, \dots, a_n]$ then $do(\alpha, s)$ stands for $do(a_n, do(\dots, do(a_1, s) \dots))$.

Domains are modeled in the situation calculus as *axioms*. A set of \mathcal{L} -sentences specify the actions available, what they depend on, and the ways they affect the world. Specifically, these axioms are given in the form of a *basic action theory* [11], reviewed below.

²We will subsequently introduce a few more distinguished predicates when modeling knowledge, sensing and nondeterminism.

³For simplicity, only functional fluents are introduced, and their predicate counterparts are ignored. (Distinguished symbols like *Poss* are an exception.) This is without loss of any generality since predicates can be thought of functions that take one of two values, the first denoting *true* and the other denoting *false*.

2.2. Basic Action Theories

In general, a basic action theory \mathcal{D} is a set of sentences consisting of (free variables understood as universally quantified from the outside):

- an *initial theory* \mathcal{D}_0 that describes what is true initially;
- *precondition axioms*, of the form $Poss(A(\vec{x}, s)) \equiv \beta(\vec{x}, s)$ that describe the conditions under which actions are executable;
- *successor state axioms*, of the form $f(\vec{x}, do(a, s)) = u \equiv \gamma_f(u, \vec{x}, s)$, that describe the changes to fluent values after doing actions;
- domain-agnostic *foundational axioms*, such as a second-order induction axiom for the space of situations and unique name axioms for actions, the details of which need not concern us here [11].

The formulation of successor state axioms, in particular, incorporates Reiter’s monotonic solution to the frame problem [19].

An agent reasons about actions by means of entailments of a basic action theory \mathcal{D} , for which standard first-order (Tarskian) models suffice (although see below). A fundamental task in reasoning about action is that of *projection* [11], where we test which properties hold after actions. Formally, suppose ϕ is a situation-suppressed formula or uses the special symbol *now*. Given a sequence of actions a_1 through a_n , we are often interested in asking whether ϕ holds after these:

$$\mathcal{D} \models \phi[do([a_1, \dots, a_n], S_0)]?$$

This concludes our review of the basic features of the language. In the subsequent sections, we will discuss how the formalism is first extended for knowledge, and then, degrees of belief against discrete probability distributions and beyond. To prepare for that, the rest of the section will introduce three logical constructions that will be used in our work. First, we will define a class of Tarskian structures. Second, we define a logical term standing for summation in the usual mathematical sense, that is to be understood as an abbreviation for a formula involving second-order quantification. Third, we analogously define a logical term standing for integration in the usual mathematical sense.

2.3. \mathbb{R} -interpretations

For our purposes, the notion of entailment will be assumed wrt a class of Tarskian structures that we call \mathbb{R} -interpretations. See [17] for a review of Tarskian structures; we assume some familiarity with the underlying notions. Below, for any \mathcal{L} -term t and \mathcal{L} -interpretation M , we use t^M to mean the domain element that t references. If t has a free variable x and y is any \mathcal{L} -term, then we write t_y^x to mean the \mathcal{L} -term obtained by replacing x in t with y . Finally, for any variable map μ and variable X , we use μ_Y^X to mean a variable map that is exactly like μ except that for the variable X it takes the value Y .

Definition 1. By an \mathbb{R} -interpretation we mean any \mathcal{L} -structure where $=, <, >, 0, 1, +, \times, /, -, e, \pi$, exponentiation and logarithms have their usual interpretations.

(That is, “ $1 + 0 = 1$ ” is true in all \mathbb{R} -interpretations, if “ $x > y$ ” is true then “ $\neg(y > x)$ ” is true, and so on.) So, henceforth, when we write $\Sigma \models \phi$, we mean that in all \mathbb{R} -interpretations where Σ is true, so is ϕ .⁴

Natural numbers can be defined in terms of a predicate by appealing to \mathbb{R} -interpretations. Let

$$Natural(x) \doteq \forall P[(P(0) \wedge \forall x(P(x) \supset P(x + 1))) \supset P(x)].$$

Then:

⁴Alternatively, one could have specified axioms for characterizing the field of real numbers together with Σ . Whether or not reals with exponentiation is *first-order* axiomatizable remains a major open question [20].

Theorem 2. Let M be any \mathbb{R} -interpretation, and c a constant symbol of \mathcal{L} . Then, $M \models \text{Natural}(c)$ iff $c^M \in \mathbb{N}$.

The proof relies on two lemmas. First, we argue that the set of natural numbers satisfies the antecedent $A \doteq P(0) \wedge \forall x(P(x) \supset P(x+1))$:

Lemma 3. For every \mathbb{R} -interpretation M and map μ , $M, \mu_{\mathbb{N}}^P \models A$.

Proof: Since “0” and “+” have their usual interpretations, and $0 \in \mathbb{N}$, $M, \mu_{\mathbb{N}}^P \models P(0)$. Suppose $k \in \mathbb{N}$, and so, $M, \mu_{\mathbb{N}}^P \models P(k)$. Since \mathbb{N} includes the successors of all its elements, $M, \mu_{\mathbb{N}}^P \models P(k+1)$. ■

Next, we argue that every set satisfying the antecedent includes the set of natural numbers:

Lemma 4. For every M and μ as above, and for any set Q , if $M, \mu_Q^P \models A$ then \mathbb{N} is a subset of Q .

Proof: Suppose $M, \mu_Q^P \models A$. We prove the claim by induction on natural numbers. For the base case, we consider $0 \in \mathbb{N}$, and by assumption $M, \mu_Q^P \models P(0)$, and so 0 is in the set Q . For the hypothesis, assume for any $k \in \mathbb{N}$, k is also in Q . That is, $M, \mu_Q^P \models P(k)$. Of course, the successor of k is a natural number and by assumption, $M, \mu_Q^P \models P(k+1)$. Thus, starting from 0, every natural number and its successor is in Q , and so Q must include \mathbb{N} . ■

So, the main claim is as follows:

Proof: Suppose $M \models \text{Natural}(c)$, that is, $M \models \forall P(A \supset P(c))$. Since \mathbb{N} is a subset of the domain by assumption, for any map μ , $M, \mu_{\mathbb{N}}^P \models A \supset P(c)$. By Lemma 3, $M, \mu_{\mathbb{N}}^P \models A$, and so $M, \mu_{\mathbb{N}}^P \models P(c)$. Hence $c^M \in \mathbb{N}$.

Suppose instead $c^M \notin \mathbb{N}$. Suppose for any set Q and any map μ , $M, \mu_Q^P \models A$. By Lemma 4, \mathbb{N} is a subset of Q and so c^M is in Q , that is, $M, \mu_Q^P \models P(c)$. Therefore, $M, \mu_{\mathbb{N}}^P \models A \supset P(c)$. Since this holds for any set Q , $M \models \forall P[A \supset P(c)]$, that is, $M \models \text{Natural}(c)$. ■

2.4. Summation

Here we show how finite summations can be characterized as an abbreviation for an \mathcal{L} -term using second-order quantification.

Let f be any \mathcal{L} -function from \mathbb{N} to \mathbb{R} . Let $\text{SUM}(f, n)$, standing for the sum of the values of f for the argument 1 through n , be defined as an abbreviation:

$$\text{SUM}(f, n) = z \doteq \exists g[g(1) = f(1) \wedge \\ g(n) = z \wedge \\ \forall i(1 \leq i < n \supset g(i+1) = g(i) + f(i+1))].$$

The variable i is understood to be chosen here not to conflict with any of the variables in n and z . The function g from \mathbb{N} to \mathbb{R} is assumed to not conflict with f . (That is, the logical terms are distinct.)

This can then be argued to correspond to summations in the usual mathematical sense as follows:

Theorem 5. Let f be a function symbol of \mathcal{L} from \mathbb{N} to \mathbb{R} , c be a term, and n be a constant symbol of \mathcal{L} . Let M be any \mathbb{R} -interpretation. Then,

$$\text{if } \sum_{i=1}^{n^M} f^M(i) = c^M \text{ then } M \models \text{SUM}(f, n) = c.$$

Proof: We prove by induction on n^M . For the base case, suppose $n^M = 1$. Then, the antecedent would give us $f^M(1)$. As for the consequent, clearly $M \models \text{SUM}(f, n) = f(1)$, and so the base case holds.

Assume the hypothesis holds for n^M , and we prove the case for $(n^M + 1)$. (Note that owing to “1” and “+” having their usual interpretations, $(n+1)^M = n^M + 1$.) So suppose

$$\sum_{i=1}^{n^M+1} f^M(i) = c^M.$$

On expanding the summation expression, we obtain:

$$\sum_{i=1}^{n^M} f^M(i) = c^M - f^M(n^M + 1)$$

By induction hypothesis, $M \models \text{SUM}(f, n) = c - f(n + 1)$, and so, by definition, $M \models \text{SUM}(f, n + 1) = c$. ■

Henceforth, we write:

$$\sum_{i=1}^n t$$

to mean the logical formula $\text{SUM}(t, n)$ for a logical term t . Here, i is assumed to not conflict with any of the variables in n and t .⁵

It is worth noting that this logical formula can be applied to summation expressions where the arguments to the terms are not restricted to natural numbers, but taken from any finite set. For example, suppose H is any finite set of terms $\{h_1, \dots, h_n\}$. We can then use terms such as:

$$\sum_{h \in H} t(h)$$

standing for an abbreviation, similar to $\text{SUM}(t, n)$: let g be a function, and let $g(i) = t(h_i)$. Then clearly the above sum defines the same number as $\sum_{i=1}^n g$. In the sequel, we sometimes sum over a finite set of situations, or a finite vector of values, which is then understood as an abbreviation in this sense.

2.5. Integration

Finally, we characterize integrals as logical terms.⁶ We present this first for a continuous real-valued function of one variable, and then discuss its (straightforward) extension to the many-variable case.

We begin by introducing a notation for limits to positive infinity. For any logical term t and variable x , we introduce a term characterized as follows:

$$\text{LIM}[x, t] = z \doteq \forall u(u > 0 \supset \exists m \forall n(n > m \supset |z - t_n^x| < u).$$

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 . The abbreviation can be argued to correspond to the limit of a function at infinity in the usual sense:

Lemma 6. *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 \mathbb{R} -interpretation of \mathcal{L} . Then we have the following:*

$$\text{If } \lim_{x \rightarrow \infty} g^M(x) = c^M \text{ then } M \models (c = \text{LIM}[x, g]).$$

⁵That is, we use the standard mathematical notation to denote sums (and later: limits and integrals) in two ways that context will disambiguate: first, as the usual mathematical expression, and second, as a well-formed logical formula to be understood as an abbreviation as explained above.

⁶In this article, given a non-negative real-valued function, our notion of an integral of this function is based on the Riemann integral [21], in which case the function is said to be *integrable*. There are limitations to the Riemann integral; for example, the function $f: [0, 1] \rightarrow \mathbb{R}$ where

$$f(x) = \begin{cases} 1 & \text{if } x \text{ is rational} \\ 0 & \text{otherwise} \end{cases}$$

is not integrable in the Riemann account. In the calculus community, generalizations to the Riemann integral, such as the gauge integral [22], have been studied that allow for the integration of such functions. We have chosen to remain within the framework of classical integration, but other accounts may be useful.

The key idea for moving beyond Riemann integrals would be to amend the logical formula standing for the abbreviations we introduce below. (For example, rather than partitioning the domain of a function, partitioning the range would allow us to formalise Lebesgue integrals.)

Proof: Suppose the limit holds. Then, by the standard notion of limits [23], for every real number $\epsilon > 0$, there is a natural number j such that for all natural numbers $k > j$:

$$|g^M(k) - c^M| < \epsilon.$$

Suppose $M \not\models \text{LIM}[x, g] = c$. Then there is some $r > 0$ such that $M, \mu_r^u \not\models \exists m \forall n (n > m \supset |c - g(n)| < u)$. By Theorem 2, the domain of M includes \mathbb{N} , and so let m and n be variables that are mapped to j and k respectively. Then, given that M interprets arithmetic symbols in the usual way, the claim $M, \mu_r^u \not\models (n > m) \supset (|c - g(n)| < u)$ is a contradiction. ■

Henceforth, we write:

$$\lim_{x \rightarrow \infty} t$$

to mean the logical formula $\text{LIM}[x, t]$.

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 :

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

where h stands for $(b - a)/n$. The variables are chosen not to conflict with any of the other variables. We now show:

Lemma 7. *Let g be a function symbol of \mathcal{L} standing for a function from \mathbb{R} to \mathbb{R} , and let a, b, c be constant symbols of \mathcal{L} . Let M be any \mathbb{R} -interpretation of \mathcal{L} . Then we have the following:*

$$\text{If } \int_{a^M}^{b^M} g^M(x) dx = c^M \text{ then } M \models (c = \text{INT}[x, a, b, g]).$$

Proof: Suppose $\int_{a^M}^{b^M} g^M(x) dx = c^M$, that is, g^M is integrable. By definition, then:

$$\lim_{k \rightarrow \infty} h \cdot \sum_{i=1}^k g^M(a^M + h \cdot i) = c^M$$

where $h = (b^M - a^M)/k$. By Lemma 6, we have

$$M \models \lim_{n \rightarrow \infty} h \cdot \sum_{i=1}^n g_{(a+h \cdot i)}^x = c$$

where $h = (b - a)/n$, that is, $M \models \text{INT}[x, a, b, g] = c$. ■

Finally, we define the definite integral of t over all real values of x by the following:

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

The main result for this logical abbreviation is the following:

Theorem 8. *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 \mathbb{R} -interpretation of \mathcal{L} . Then we have the following:*

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

Proof: Suppose $\int_{-\infty}^{\infty} g^M(x) dx = c^M$, that is, g^M is integrable. By definition,

$$\lim_{u \rightarrow \infty} \lim_{v \rightarrow \infty} \int_{-u}^v g^M(x) = c.$$

By Lemma 6 and Lemma 7, we obtain:

$$M \models \lim_{u \rightarrow \infty} \lim_{v \rightarrow \infty} \text{INT}[x, -u, v, g] = c.$$

That is, $M \models \int_x g(x) = c$. ■

The characterization of integrals for a many-variable function f , from \mathbb{R}^k to \mathbb{R} , is then an easy exercise. For variables x_1, \dots, x_k and terms $a_1, \dots, a_k, b_1, \dots, b_k$, and t , we introduce a term $\text{MINT}[x_1, \dots, x_k, a_1, \dots, a_k, b_1, \dots, b_k, t]$ denoting the definite integral of t from $(a_1, \dots, a_k) \in \mathbb{R}^k$ to $(b_1, \dots, b_k) \in \mathbb{R}^k$:

$$\text{MINT}[x_1, \dots, x_k, a_1, \dots, a_k, b_1, \dots, b_k, t] \doteq \lim_{n_1 \rightarrow \infty} \dots \lim_{n_k \rightarrow \infty} h_1 \dots h_k \cdot \sum_{i_1=1}^{n_1} \dots \sum_{i_k=1}^{n_k} t_{(a_1+h_1 \cdot i_1), \dots, (a_k+h_k \cdot i_k)}^{x_1, \dots, x_k}$$

where h_j stands for $(b_j - a_j)/n_j$. The variables are chosen not to conflict with any of the other variables. Finally, we define the definite integral of t over all real values of x_1, \dots, x_k by the following:

$$\int_{x_1, \dots, x_k} t \doteq \lim_{u_1 \rightarrow \infty} \dots \lim_{u_k \rightarrow \infty} \lim_{v_1 \rightarrow \infty} \dots \lim_{v_k \rightarrow \infty} \text{MINT}[x_1, \dots, x_k, -u_1, \dots, -u_k, v_1, \dots, v_k, t].$$

From this we get, as a corollary to Theorem 8:

Corollary 9. *Let g be a function symbol of \mathcal{L} standing for a function from \mathbb{R}^k to \mathbb{R} , and let c be a constant symbol of \mathcal{L} . Let M be any \mathbb{R} -interpretation of \mathcal{L} . Then we have the following:*

$$\text{If } \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} g^M(x_1, \dots, x_k) dx_1 \dots dx_k = c^M \text{ then } M \models (c = \int_{x_1, \dots, x_k} g(x_1, \dots, x_k)).$$

3. A Theory of Knowledge

3.1. Knowledge

An early treatment of *knowledge* in the situation calculus is due to Moore [24]. The classical *possible-world* interpretation for knowledge [8, 9] is based on the notion that there many different ways the world can be, where each world stands for a complete state of affairs. Some of these are considered possible by a putative agent, and they determine what the agent knows and does not know. Moore's observation was that situations can be viewed as possible worlds. (Different from standard modal logics [10], however, worlds are reified as part of the syntax, but this is a minor technicality.) A special binary fluent K , taking two situation arguments, determines the accessibility relation between worlds: $K(s', s)$ says that when the agent is at s , he considers s' possible. (Also different from standard modal logics, we note that the order of the terms in the accessibility relation is reversed.) *Knowledge*, then, is simply truth at accessible worlds:

Definition 10. (Knowledge.) Let ϕ be any situation-suppressed formula. The agent *knowing* ϕ at situation s , written $\text{Knows}(\phi, s)$, is the following abbreviation:

$$\text{Knows}(\phi, s) \doteq \forall s'. K(s', s) \supset \phi[s'].$$

In English: if in every situation s' that is considered possible at s the formula ϕ holds, then the agent knows ϕ at s .⁷

Moore's account has been adapted to the arrangement of dynamic laws via basic action theories by Scherl and Levesque [12]. In this scheme, the initial theory is assumed to specify the agent's initial beliefs. For example, $\text{Knows}(\text{velocity}(\text{obj5}) = 50, S_0)$ expands to $\forall s'. K(s', S_0) \supset \text{velocity}(\text{obj5}, s') = 50$; that is, every accessible world

⁷We do not insist that beliefs are necessarily true in the real world. Perhaps for this reason, "believes" is a more appropriate reading of *Knows*. The two terms are used interchangeably here. Also, see Scherl and Levesque [12] on how features such as *positive* and *negative introspection* can be enabled by constraining the accessibility relation K , in a manner entirely analogous to standard modal logic [10].

initially agrees that the velocity of object 5 is 50, which means that the agent knows that the velocity of object 5 is 50. In addition, $\neg \text{Knows}(\text{velocity}(\text{obj}6) = 60, S_0)$ and $\text{Knows}(f = 0 \vee f = 1, S_0)$ say that initially, the agent does not know that the velocity of object 6 is 60, and knows that the functional fluent f takes a value of either 0 or 1. Equivalently, one specifies an initial constraint on K ; for example:

$$K(t, S_0) \equiv (f = 1 \vee f = 0)[t] \quad (1)$$

characterizes the agent's knowledge about the fluent f taking a value of either 0 or 1.

3.2. Sensing

To specify the behavior of K at non-initial situations, Scherl and Levesque provide a *successor state axiom* for K . This axiom, intuitively, tests whether situations are to remain accessible as actions occur. Without going into full details, assume the provision of domain-specific *sensing* axioms, also part of the basic action theory \mathcal{D} . For example,

$$SF(\text{sensetrue-}f, s) \equiv f(s) = 1.$$

formalizes the sensing outcome for an action to check whether f has value 1 in situation s . Here SF is a special \mathcal{L} -predicate, similar to $Poss$. Then, for s' to be accessible from s , we need the following successor state axiom to be included in \mathcal{D} :

$$K(s', do(a, s)) \equiv \exists s'' [K(s'', s) \wedge s' = do(a, s'') \wedge Poss(a, s'') \wedge (SF(a, s'') \equiv SF(a, s))].$$

This says that if s'' is the predecessor of s' , such that s'' was considered possible at s , then s' would be considered possible from $do(a, s)$ contingent on sensing outcomes.

To appreciate how this axiom works in dynamical domains, assume that for actions without any sensing aspect, such as an action that toggles the value of f , one simply lets SF be vacuously true:

$$SF(\text{toggle-}f, s) \equiv \text{true}.$$

The idea, then, is that the accessibility between situations would not change as physical actions occur. However, as the agent operates in an environment where it senses various properties, those situations that are *incompatible* with the real world regarding the sensed value will be deemed *impossible* after such sensing actions. This leads to a notion of knowledge *expansion*,⁸ as the agent becomes more certain about the true nature of the world. See Figure 2 for an illustration of the axiom: we imagine three situations s , s' and s'' that are epistemically related prior to any sensing (that is, s' and s'' are possible worlds when the agent is at s). These situations disagree on the value of f , and consequently, the agent does *not* know f 's value. After executing a sensing action for the truth of f , however, $do(\text{sensetrue-}f, s'')$ is not epistemically related to $do(\text{sensetrue-}f, s)$. The upshot is that the agent knows the value of f at $do(\text{sensetrue-}f, s)$ and believes that this value is 1.




Figure 2: Situations with accessibility relations after sensing. The numbers inside the circles denote the f values at these situations.

⁸Revising beliefs, where the agent believes ϕ but acquires information to now believe $\neg\phi$, is not dealt with in this work. See [25] for an account.

3.3. Degrees of Belief and Likelihood

The Scherl and Levesque scheme, however, lacks constructs to quantify the agent’s uncertainty. One measure to quantify uncertainty is with *degrees of belief*. What is also lacking in their scheme is the ability to formalize the *probabilistic noise* in effectors and sensors, as seen in many real-world robotics applications [13]. These limitations, to discrete approximations, were addressed by BHL [5].

The BHL scheme builds on Scherl and Levesque’s ideas, especially regarding how accessibility relations between worlds vary as a result of actions. In fact, the reader may observe many parallels between the two extensions. BHL’s remarkably simple proposal consists of introducing two new distinguished fluents, p and l , in addition to K . We present a simpler alternative involving only p and l .

As a simple running example, imagine a robot moving towards a wall, as shown in Figure 1. Its distance to the wall is given by a functional fluent h , and it is assumed to be equipped with a sonar sensor that measures how far the robot is from the wall. In other words, ideally, the sensor’s reading would correspond to the actual value of h .

The p fluent determines a probability distribution on situations, by associating situations with *weights*. More precisely, the term $p(s', s)$ denotes the relative *weight* accorded to situation s' when the agent happens to be in situation s . Of course, p can be seen as a companion to K . As one would for K , the properties of p in initial states, which vary from domain to domain, are specified with axioms as part of \mathcal{D}_0 . For example,

$$p(\iota, S_0) = u \equiv ((h(\iota) = 2 \vee h(\iota) = 3) \wedge u = .5) \vee (h(\iota) \neq 2 \wedge h(\iota) \neq 3 \wedge u = 0). \quad (2)$$

says that those initial situations where h has the integer values 2 or 3 obtain a weight of .5. All other situations, then, obtain 0 weight. We expect, of course, that weights are nonnegative, and that non-initial situations are given a weight of 0 initially. The following nonnegative constraint, also part of \mathcal{D}_0 , ensures this:

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

Note that this is a stipulation about initial situations ι only. But BHL provide a *successor state axiom* for p , to be listed shortly, that ensures that this constraint holds in all situations.

Next, the term $l(a, s)$ is intended to denote the likelihood of action a occurring in situation s . Among other things, l can be used to model noisy sensors. This is perhaps best demonstrated using an example. Imagine a sonar aimed at the wall, which gives a reading for the true value of h . Supposing the sonar’s readings are subject to additive Gaussian noise.⁹ If now a reading of z were observed on the sonar, intuitively, those situations where $h = z$ should be considered more probable than those where $h \neq z$.¹⁰ This occurrence is captured using likelihoods in the formalism. Basically, if $\text{sonar}(z)$ is the sonar sensing action with z being the value read, we specify a likelihood axiom describing its error profile as follows:

$$l(\text{sonar}(z), s) = u \equiv (z \geq 0 \wedge u = \mathcal{N}(z - h(s); \mu, \sigma^2)) \vee (z < 0 \wedge u = 0). \quad (3)$$

This stipulates that the difference between a nonnegative reading of z and the true value h is normally distributed with a variance of σ^2 and mean of μ .¹¹

Clearly, the error profile of various hardware devices is application dependent, and it is this profile that is modeled as shown above using l . Notice, for example, when $\mu = 0$, which indicates that the sensor has no systematic bias, then $l(\text{sonar}(5), s)$ will be higher when $h(s) = 5$ than when $h(s) = 25$. Roughly, then, the idea is that after an observation, the weights on situations would get redistributed based on their compatibility with the observed value.

⁹Note that Gaussians are continuous distributions involving π , e , exponentiation, and so on. Therefore, BHL always consider discrete probability distributions that *approximate* the continuous ones.

¹⁰As usual, the reading observed is not in the control of the agent. Here, we assume that the value is given to us, and in that sense, the language is geared for projection (cf. Section 5.2). For example, we might be interested in the beliefs of the agent after obtaining a specific sequence of readings on the sonar. Integrating this language with an *online* framework that obtains such readings from an external source is addressed in [26].

¹¹We understand $\mathcal{N}(u; \mu, \sigma^2)$ as an abbreviation for the mathematical expression $e^{-\frac{(u-\mu)^2}{2\sigma^2}} / \sqrt{2 \cdot \pi \cdot \sigma^2}$.

One may contrast such likelihood specifications to (trivial) ones for deterministic physical actions,¹² such as an action $move(z)$ of moving towards the wall by precisely z units. For such actions, we simply write

$$l(move(z), s) = u \equiv u = 1$$

in which case the p value of s is the same as that for $do(move(z), s)$. Thus, this is a form of *imaging* [27], where the weights of worlds are simply “transferred” to their successors.

Formally, we add *action likelihood* axioms to \mathcal{D} :

Definition 11. Action likelihood axioms for each action type A are sentences of the form:

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

Here $\psi_A(\vec{x}, u, s)$ is any formula characterizing the conditions under which action $A(\vec{x})$ has likelihood u in s .

In general, likelihood axioms can depend on any number of features of the world besides the fluent that the sensor is measuring. For example, imagine that the sonar’s accuracy depends on the room temperature. We could then specify an error profile as follows:

$$\begin{aligned} l(sonar(z), s) = u \equiv & \\ & (z \geq 0 \wedge temp(s) \geq 0 \wedge u = \mathcal{N}(z - h(s); \mu, 1)) \vee \\ & (z \geq 0 \wedge temp(s) < 0 \wedge u = \mathcal{N}(z - h(s); \mu, 16)) \vee \\ & (z < 0 \wedge u = 0). \end{aligned} \quad (4)$$

That is, the sonar’s accuracy worsens severely when the temperature drops below 0, as seen by the larger variance.

Having introduced the new fluents, we are now ready to provide the successor state axiom for p , which is analogous to the one for K :

$$\begin{aligned} p(s', do(a, s)) = u \equiv & \\ & \exists s'' [s' = do(a, s'') \wedge Poss(a, s'') \wedge \\ & \quad u = p(s'', s) \times l(a, s'')] \\ \vee \neg \exists s'' [s' = do(a, s'') \wedge Poss(a, s'')] \wedge u = 0. \end{aligned} \quad (P2)$$

This says that 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'' .




Figure 3: Situations with accessibility relations after noisy sensing. The numbers inside the circles denote the h values at these situations. Dotted circles denote a lower weight in relation to their epistemic alternatives.

¹²Noisy actions will also involve non-trivial likelihood axioms. Their treatment, however, is deferred to a subsequent section.

To see an application of this axiom using the specifications (2) and (3), consider two situations s and s' associated with the same weight initially, as shown in Figure 3. These situations have h values 2 and 3 respectively. Suppose the robot obtains a reading of 2 on the sensor. Given a sensor with mean $\mu = 0$ and variance $\sigma^2 = 1$, the likelihood axiom is such that the weight of the successor of s is higher than that of s' because the h value at s coincides with the sensor reading. More precisely, the weight for the successor of s is given by the prior weight .5 multiplied by the likelihood factor $\mathcal{N}(z - h(s); \mu, \sigma^2) = \mathcal{N}(2 - 2; 0, 1) = \mathcal{N}(0, 0, 1)$. The weight for the successor of s' is obtained analogously.

Interestingly, by means of the above axiom, if predecessors are not *epistemically* related, which is another way of saying that $p(s'', s)$ is 0, then their successors will also not be. Similarly, when a is not executable at s'' , its successor is no longer accessible from $do(a, s)$. One other 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. This is because (P1) insists that initial situations are only epistemically related to other initial ones, and (P2) respects this relation over actions.

We are now prepared to define the *degree of belief* in a formula ϕ at a situation s , written $Bel(\phi, s)$. (Henceforth, whenever a formula ϕ appears in the context of Bel , we assume that it is either situation-suppressed, or it only mentions the situation term *now*.) Intuitively, this is simply the *weight of accessible situations*. Formally:

Definition 12. (Degrees of belief.) Suppose ϕ is any situation-suppressed \mathcal{L} -formula. Then the *degree of belief* in ϕ is an abbreviation for:

$$Bel(\phi, s) \doteq \frac{1}{\gamma_{\{s':\phi[s']\}}} \sum p(s', s),$$

where γ , the normalization factor, is understood (throughout) as the same expression as the numerator but with ϕ replaced by *true*. So, here, 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, as discussed above. The summation term in this logical formula is not a new logical symbol, but simply an abbreviation for a second-order formula, by way of Section 2.4.

3.4. Discussion

Let us conclude this section by remarking that since p is a fluent, the syntax of basic action theories allows us to express probabilistic knowledge in a very general way, quite beyond standard probabilistic formalisms [28, 7]. For example, to represent categorical uncertainty like (1), we would do:

$$p(t, S_0) = \begin{cases} 1 & \text{if } (h = 2 \vee h = 3)[t] \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

In English: all initial situations where the value of h is either 2 or 3 are considered epistemically possible, and accorded a weight of 1. All other initial situations are accorded a weight of 0. Observe, however, that this p -specification does not say which value of h is more likely. Thus, unlike standard probabilistic frameworks, we do not assume that it is always possible to find a single probability distribution for the robot to use.

It is also possible to handle partial specifications. Let us contrast (2) with the following initial axiom for p :

$$p(t, S_0) = \begin{cases} .1 & \text{if } h(t) \in \{1, 2, \dots, 10\} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Then, letting a basic action theory include the sentence:

$$(2) \vee (6)$$

means that the robot believes h is uniformly distributed on $\{2, 3\}$ or on $\{1, \dots, 10\}$ without being able to say which. To reiterate, we do not assume that it is always possible to find a single probability distribution for the robot to use.

Of course, a much weaker specification is possible by replacing one of these probabilistic alternatives with categorical ones. For example, suppose the basic action theory were to instead include:

$$(5) \vee (6).$$

In this case, the agent believes that h may take a value from $\{2, 3\}$ or is uniformly distributed on $\{1, 2, \dots, 10\}$ without being able to say which.

In sum, the framework allows one to freely combine categorical and probabilistic specifications, leading to a very general model of belief. For simplicity of presentation, we consider a fairly simple set of specifications in this article, and refer readers to [29] for more involved ones.

4. Belief Reformulated

The definition for degrees of belief, as given by BHL, is intuitive and simple. It is closely fashioned after the semantics for belief in modal probability logics [30], where the probabilities of formulas is calculated from the weights of possible worlds satisfying the formula. Unfortunately, this definition is not easily amenable to generalizations. Notice, for example, that Bel is well-defined only when the sum over all those situations s' such that $\phi[s']$ holds is finite. This immediately precludes domains that involve an infinite set of situations agreeing on a formula. Moreover, the definition does not have an obvious analogue for continuous probability distributions. Observe that such an analogue would involve *integrating* over the space of situations, which makes little sense. Indeed, it is not certain what the space of situations would look like in general, but even if this was fixed, how such a thing can be tinkered with so as to obtain an appropriate notion of *integration* is far from obvious.

Therefore, what we propose is to shift the calculating of probabilities from situations to *fluent values*, that is, to the well-understood domain of numbers. The current section is an exploration of this idea. What we will show in this section is that Definition 12 can be reformulated as a summation over numeric indices. That will allow, among other things, a seamless generalization from summation to integration, which is to be the topic of the next section.

To prepare for that, in addition to the usual case notation used, for example, in (6), we will introduce another kind of conditional term for convenience. This involves a quantifier and a default value of 0, like in formula (P2). 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:

$$\langle z.\psi \rightarrow t \rangle = u \doteq [(\exists z\psi) \supset \forall z(\psi \supset u = t)] \wedge [(\neg\exists z\psi) \supset u = 0].$$

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, we will now need a way to enumerate the primitive fluent terms of the language. Intuitively, these correspond to the probabilistic variables in the language. Perhaps the simplest way is to assume there are n fluents f_1, f_2, \dots, f_n in \mathcal{L} which take no arguments other than the situation argument,¹³ and that they take their values from some finite sets. We can then rephrase Definition 12 as follows:

Definition 13. Suppose ϕ is as before. Let $Bel(\phi, s)$ be an abbreviation for:

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

where γ is the numerator but with ϕ replaced by *true*, as usual.

(For readability, we often drop the index variables in sums and connectives when the context makes it clear: in this case, i ranges over the set $\{1, \dots, n\}$, that is, the indices of the fluents in \mathcal{L} .) Definition 13 suggests that for each possible value of the fluents, we are to sum over all possible situations and for each one, if the fluents have those

¹³Essentially, functional fluents in \mathcal{L} are assumed to not take any object arguments. More generally, if we assume that the arguments of k -ary fluents are drawn from finite sets, an analogous enumeration of ground functional fluent terms is possible.

Understandably, from the point of view of situation calculus basic action theories, where fluents are also usually allowed to take arguments from *any* set, including infinite ones, this is a limitation. But in probabilistic terms, this would correspond to having a joint probability distribution over infinitely many, perhaps uncountably many, random variables. We know of no general logical account of this sort, and we have as yet no good ideas about how to deal with it. It remains to be seen whether ideas from probability theory on high dimensions [31, 32] and infinite-dimensional probabilistic graphical models [33] can be leveraged for our purposes.

values and ϕ holds, then the p value is to be used, and 0 otherwise. Roughly speaking, if one were to group situations satisfying $\bigwedge f_i(s) = x_i$ into sets for every possible vector \vec{x} , the union of these sets would give the space of situations. Our claim about the relationship between the two abbreviations can be made precise as follows:

Theorem 14. *Let \mathcal{D} be any basic action theory and ϕ any \mathcal{L} -formula. Then the abbreviations for $Bel(\phi, s)$ from Definition 12 and 13 define the same number.*

Proof: For the proof, we focus solely on the numerators of the two abbreviations. That is,

$$\sum_{\{s':\phi[s']\}} p(s', s) \quad (\dagger)$$

on the one hand, and

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

on the other. We show that these expressions define the same number. With this, the case for the denominators follows trivially, since *true* is a special case of ϕ . Then, the claim is proven.

Let S be a set such that $s' \in S$ iff $p(s', s) > 0$. (That is, for any ground situation term s , this is the set of all ground situation terms s' such that $p(s', s) > 0$.) Then, let $T \subseteq S$ be the set such that $s' \in T$ iff $\phi[s']$. Intuitively, T is the set of all situations that are epistemically related to s , but where ϕ holds. It is easy to see that

$$(\dagger) = \sum_{s' \in T} p(s', s).$$

Suppose now f_i ranges over $\{c_{1i}, \dots, c_{ki}\}$, and so by extension, suppose $\vec{f} = \langle f_1, \dots, f_n \rangle$ ranges over $\{\vec{c}_1^i, \vec{c}_2^i, \dots, \vec{c}_k^i\}$. For any \vec{c}_j^i in that set, let $T_j \subseteq T$ be a set such that $s' \in T_j$ iff $\bigwedge f_i(s') = c_{ji}$. That is, T_j identifies those situations from T where fluents satisfy the vector of values \vec{c}_j^i . Observe that

$$(\ddagger) = \sum_{j=1}^k \sum_{s' \in T_j} p(s', s).$$

But, of course, $T = \bigcup T_j$. Therefore,

$$\sum_{j=1}^k \sum_{s' \in T_j} p(s', s) = \sum_{s' \in T} p(s', s).$$

Therefore, (\dagger) and (\ddagger) define the same number. ■

Be that as it may, Definition 13 still involves summations over situations. To arrive at a definition that eschews the summing of situations, we start with the case of initial situations. In this matter, we will be insisting on a precise space of initial situations. For this, let us recall the axiomatization of the situation calculus presented in [34] for multiple initial situations, which includes a sentence saying there is precisely one initial situation for *any* possible vector of fluent values. This can be written as follows:

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

(Recall that i ranges over the indices of the fluents in \mathcal{L} , that is, $\{1, \dots, n\}$.) Under the assumption (P3), we can rewrite Definition 12 for $s = S_0$ as

$$Bel(\phi, S_0) \doteq \frac{1}{\gamma} \sum_{\vec{x}} \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[\iota] \rightarrow p(\iota, S_0) \rangle \quad (\text{B0})$$

The two abbreviations, in fact, are equivalent:

Theorem 15. *Let \mathcal{D} be any basic action theory, ϕ any \mathcal{L} -formula, and suppose \mathcal{D}_0 includes (P3). Then the abbreviations for $Bel(\phi, S_0)$ in Definition 12 and (B0) define the same number.*

Proof: As in Theorem 14, we focus on the numerators for the two abbreviations. That is,

$$\sum_{\{s':\phi[s']\}} p(s', S_0) \quad (\dagger)$$

on the one hand, and

$$\sum_{\vec{x}} \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[\iota] \rightarrow p(\iota, S_0) \rangle \quad (\ddagger)$$

on the other. We show that these expressions define the same number. The denominators represent a special case, and so the claim will follow.

Let S be the set of initial situations. Suppose f_i ranges over $\{c_i, c'_i, \dots, c''_i\}$. By way of (P3), for any vector of values $\langle c_1, \dots, c_n \rangle$ for the vector of fluents $\langle f_1, \dots, f_n \rangle$, there exists a (unique) situation $s \in S$ such that $\bigwedge f_i(s) = c_i$. Let $T \subseteq S$ be such that $s \in T$ iff $\phi[s]$. It is easy to see that

$$(\ddagger) = \sum_{s' \in T} p(s', S_0).$$

Since (P1) ensures that $p(s', S_0) > 0$ only if s' is an initial situation, we get that

$$(\dagger) = \sum_{s' \in T} p(s', S_0).$$

Therefore (\dagger) and (\ddagger) define the same number. ■

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

Unfortunately, (B0) is only geared for initial situations. For non-initial situations, the assumption that no two agree on all fluent values is untenable. To see why, imagine an action $move(z)$ that moves the robot z units to the left (towards the wall) but that the motion stops if the robot hits the wall. The successor state axiom for fluent h , then, might be like this:

$$\begin{aligned} h(do(a, s)) = u &\equiv \\ &\neg \exists z (a = move(z) \wedge u = h(s) \vee \\ &\exists z (a = move(z) \wedge u = \max(0, h(s) - z)). \end{aligned} \quad (7)$$

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 $move(4)$ 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.

It turns out there is a simple way to circumvent this issue by appealing to Reiter's solution to the frame problem. Indeed, Reiter's solution gives us 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 . More precisely,

Definition 16. (Degrees of belief (reformulated).) Let ϕ be any \mathcal{L} -formula. Given any sequence of ground action terms $\alpha = [a_1, \dots, a_k]$, let

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \sum_{\vec{x}} P(\vec{x}, \phi, s)$$

where if $s = do(\alpha, S_0)$ then

$$P(\vec{x}, \phi, do(\alpha, S_0)) \doteq \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[do(\alpha, \iota)] \rightarrow p(do(\alpha, \iota), do(\alpha, S_0)) \rangle.$$

(As before, i ranges over the indices of the fluents in \mathcal{L} .) To say more about how (and why) this definition works, we first note that by (P1) and (P2), p will be 0 unless its two arguments share the same history. So the s' argument of p in Definition 12 is *expanded* and written as $do(\alpha, \iota)$ in Definition 16. By ranging over all fluent values, we range over all initial ι as before, but 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. In particular, owing to p 's successor state axiom (P2), the weight for non-initial situations accounts for the likelihood of actions executed in the history. We now establish the following result:

Theorem 17. *Let \mathcal{D} be any basic action theory with (P3) initially, ϕ any \mathcal{L} -formula, and α any sequence of ground actions terms. Then the abbreviations for $\text{Bel}(\phi, \text{do}(\alpha, S_0))$ in Definition 12 and Definition 16 define the same number.*

Proof: As in Theorem 14, we will focus on the numerators of the two abbreviations. That is,

$$\sum_{\{s':\phi[s']\}} p(s', \text{do}(\alpha, S_0)) \quad (\dagger)$$

on the one hand, and

$$\sum_{\vec{x}} \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[\text{do}(\alpha, \iota)] \rightarrow p(\text{do}(\alpha, \iota), \text{do}(\alpha, S_0)) \rangle. \quad (\ddagger)$$

on the other. We show that these expressions define the same number. The denominators represent a special case, and so the claim will follow.

Let S be the set of initial situations, as determined by (P3). Let $S^* = \{\text{do}(\alpha, s') : s' \in S\}$. Let $T \subseteq S^*$ such that $s'' \in T$ iff $\phi[s'']$. It is easy to see that

$$(\ddagger) = \sum_{s' \in T} p(s', \text{do}(\alpha, S_0)).$$

On the other hand, by (P1) and (P2), $p(s, \text{do}(\alpha, S_0)) = 0$ for all $s \notin S^*$. This means that

$$(\dagger) = \sum_{s' \in T} p(s', \text{do}(\alpha, S_0)).$$

Therefore (\dagger) and (\ddagger) define the same number. ■

Thus, by incorporating a simple constraint on initial situations, we now have a notion of belief that does not require summing over situations.

Readers may notice that 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* [11], where we are interested in inferring whether a formula holds after an action sequence. In fact, we can use regression on the ϕ and the p to reduce the belief formula from Definition 16 to a formula involving initial situations only. See [35] for work in this direction.

5. From Weights to Densities

The framework presented so far is *fully discrete*, which is to say that fluents, sensors and effectors are characterized by finite values and finite outcomes. Belief in ϕ , in particular, is the summing over a finite set of situations where ϕ holds. We now generalize this framework. We structure our work by first focusing on *fully continuous domains*, which is to say that fluents, sensors and effectors are characterized by values and outcomes ranging over \mathbb{R} . This section, in particular, explores the very first installment: effectors are assumed to be deterministic, but *sensors* have continuous noisy error profiles. The next section, then, allows both effectors and sensors to have continuous noisy profiles. Further generalizations are deferred to Section 7.

Let us begin by observing that 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. This, in particular, is made possible by the developments in the preceding section.

Allowing real-valued fluents implies that there will be uncountably many initial situations. Imagine, for example, the scenario from Figure 1, and that the fluent 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:

$$p(\iota, S_0) = \begin{cases} .1 & \text{if } 2 \leq h(\iota) \leq 12 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

which says that the true value of h initially is drawn from a uniform distribution on the interval $[2, 12]$. 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. (That is, we now interpret $p(s', s)$ as the *density* of s' when the agent is in s .) 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.

When actions enter the picture, even if deterministic, there is more to be said. Numerous subtleties arise with p in non-initial situations. For example, if the robot were to do a *move*(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*, and the degree of belief in $h = 0$ should be .2. On the other hand, the other points $h \in (0, 8]$ should retain their densities. That is, belief in $h \leq 2$ should be .4 but belief in $h = 2$ should still be 0. In effect, we have moved from a continuous to a *mixed* distribution on h . Of course, a subsequent rightward motion will retain this mixed density. For example, if the robot were to now move away by 4 units, the belief in $h = 4$ would then be .2.

To address the concern of belief change in continuous domains, we now present a generalization to BHL. One of the advantages of our approach is that we will not need to specify how to handle changing densities and distributions like the ones above. These will emerge as side-effects, that is, shifting density changes will be *entailed* by the action theory.

For our formulation of belief, we first observe that we have fluents f_1, \dots, f_n in \mathcal{L} as before, that take no argument other than the situation term but which now take their values from \mathbb{R} . Then:

Definition 18. (Degrees of belief (continuous noisy sensors).) Let ϕ be any situation-suppressed \mathcal{L} -formula, and $\alpha = [a_1, \dots, a_k]$ any ground sequence of action terms. The *degree of belief* in ϕ at s is an abbreviation:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x}} P(\vec{x}, \phi, s)$$

where, as in Definition 16, if $s = do(\alpha, S_0)$ then

$$P(\vec{x}, \phi, do(\alpha, S_0)) \doteq \langle \iota. \bigwedge f_i(\iota) = x_i \wedge \phi[do(\alpha, \iota)] \rightarrow p(do(\alpha, \iota), do(\alpha, S_0)) \rangle.$$

That is, the belief in ϕ is obtained by ranging over all possible fluent values, and integrating the densities of situations where ϕ holds. If we were to compare the above definition to Definition 16, we see that we have simply shifted from summing over finite domains to integrating over reals. In fact, we could read P as the (*unnormalized*) density associated with a situation satisfying ϕ . As discussed, by insisting on an explicit world history, the ι need only range over initial situations, giving us the exact correspondence with fluent values.

This completes our new definition of belief. To summarize, our extension to the BHL scheme is defined using a few convenient abbreviations, such as for *Bel* and mathematical integration, and where an action theory consists of:

1. \mathcal{D}_0 (with (P1)) as usual, but now also including (P3);
2. precondition axioms as usual;
3. successor state axioms, including one for p , namely (P2), as usual;
4. foundational domain-independent axioms as usual; and
5. action likelihood axioms, one for each action type.

Note that, apart from (P3) and *Bel*'s new abbreviation, we carry over precisely the same components as would BHL. By and large, the extension, thus, retains the simplicity of their proposal, and comes with minor additions. We will show that it has reasonable properties using an example and its connection to Bayesian conditioning below.

In the sequel, we assume, without explicitly mentioning so, that basic action theories include the sentences (P1), (P2) and (P3).

5.1. Bayesian Conditioning

We now explicate the relationship between our definition for *Bel* and *Bayesian conditioning* [7]. Bayesian conditioning is a standard model for belief change wrt noisy sensing [13] and it rests on two significant assumptions. First, sensors do not physically change the world, and second, conditioning on a random variable f is the same as conditioning on the event of observing f .

In general, in the language of the situation calculus, there need not be a distinction between sensing actions and physical actions. In that case, the agent's beliefs are affected by the sensed value as well as any other physical changes that the action might enable to adequately capture the "total evidence" requirement of Bayesian conditioning.

The second assumption expects that sensors only depend on the true value for the fluent. For example, in the formulation of (3) the sonar's error profile is determined solely by h . But to suggest that the error profile might depend on other factors about the environment, as formulated by (4) for example, goes beyond this simplified view. In fact, here, the agent also learns about the room temperature, apart from sensing the value of h .

Thus, our theory of action admits a view of dynamical systems far richer than the standard setting where Bayesian conditioning is applied. Be that as it may, when a similar set of assumptions are imposed as axioms in an action theory, we obtain a sensor fusion model identical to Bayesian conditioning. This connection was demonstrated in BHL for the discrete case. We prove the property formally for continuous variables below.

We begin by stipulating that actions are either of the *physical* type or of the *sensing* type [12], the latter being the kind that do not change the value of any fluent, that is, such actions do not appear in the successor state axioms for any fluent. Now, if $obs(z)$ senses the true value of fluent f , then assume the sensor error model to be:

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

where $Err(u_1, u_2)$ is some expression with only two free variables, both numeric. This captures the notion established above: the error model of a sensor measuring f depends only on the true value of f , and is independent of other factors. Finally, for simplicity, assume $obs(z)$ is always executable:

$$Poss(obs(z), s) \equiv true.$$

Then we obtain:

Theorem 19. *Suppose \mathcal{D} is any basic action theory with likelihood and precondition axioms for $obs(z)$ as above, ϕ is any \mathcal{L} -formula mentioning only f , and $u \in \{x_1, \dots, x_n\}$ that f takes a value from. Then we obtain:*

$$\mathcal{D} \models Bel(\phi, do(obs(z), S_0)) = \frac{\int_{\vec{x}} [P(\vec{x}, \phi \wedge f = u, S_0) \times Err(z, u)]}{\int_{\vec{x}} [P(\vec{x}, f = u, S_0) \times Err(z, u)]}$$

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. The proof for the theorem is as follows.

Proof: Without loss of generality, assume f takes the value x , and the remaining fluents are f_1, \dots, f_n that take values from x_1, \dots, x_n . Let a denote $obs(z)$. From Definition 18, $Bel(\phi, do(obs(z), S_0))$ is an abbreviation for:

$$\frac{1}{\gamma} \int_{\vec{x}} \frac{P(\vec{x}, \phi, do(a, S_0))}{\gamma} = \tag{a}$$

$$\frac{1}{\gamma} \int_{\vec{x}} \langle \iota.f(\iota) = x \wedge \bigwedge_{i=1}^n f_i(\iota) = x_i \wedge \phi[do(a, \iota)] \rightarrow p(do(a, \iota), do(a, S_0)) \rangle = \tag{b}$$

$$\frac{1}{\gamma} \int_{\vec{x}} \langle \iota.f(\iota) = x \wedge \bigwedge f_i(\iota) = x_i \wedge \phi[\iota] \rightarrow p(do(a, \iota), do(a, S_0)) \rangle = \tag{c}$$

$$\frac{1}{\gamma} \int_{\vec{x}} \langle \iota.f(\iota) = x \wedge \bigwedge f_i(\iota) = x_i \wedge (\phi \wedge f = x)[\iota] \rightarrow \underline{p(do(a, \iota), do(a, S_0))} \rangle = \tag{d}$$

$$\frac{1}{\gamma} \int_{\vec{x}} \langle \iota.f(\iota) = x \wedge \bigwedge f_i(\iota) = x_i \wedge (\phi \wedge f = x)[\iota] \rightarrow \underline{p(\iota, S_0) \times Err(z, x)} \rangle = \tag{e}$$

$$\frac{1}{\gamma} \int_{\vec{x}} \frac{Err(z, x) \times \langle \iota.f(\iota) = x \wedge \bigwedge f_i(\iota) = x_i \wedge (\phi \wedge f = x)[\iota] \rightarrow p(\iota, S_0) \rangle}{\gamma} = \tag{f}$$

$$\frac{1}{\gamma} \int_{\vec{x}} Err(z, x) \times P(\vec{x}, \phi \wedge f = x, S_0).$$

The arguments underline those parts of the sentences that are being reduced. Step (a) expands P . In step (b), owing to the fact that sensing actions do not change fluent values (by our first assumption tailored to Bayesian conditioning), $\phi[do(a, \iota)]$ is equivalent to $\phi[\iota]$. In step (c), we observe that what appears to the left of the right-arrow in step (b) is equivalent to one where $\phi[\iota]$ is replaced by $(\phi \wedge f = x)[\iota]$. (The main reason for introducing the formula $f = x$ is to allow us to identify those situations where f takes the value x which are all to be multiplied by the error likelihood for observing z when the true value is x .) In step (d), we use (P2) and the fact that a is always executable to replace $p(do(a, \iota), do(a, S_0))$ by $p(\iota, S_0) \times l(a, \iota) = p(\iota, S_0) \times Err(z, x)$. Now note that the term appearing in the context of an integral is suggesting that if there is an initial situation where a particular condition holds, then a certain value is returned, and otherwise 0 is returned. This allows us to place the term $Err(z, x)$ on the outside in step (f), giving us the required numerator that appears in the claim. The expansion of the denominator is analogous with $true$ being a special case for ϕ , and so we are done. ■

The usual case for posteriors are formulas such as $b \leq f \leq c$, which is estimated from the prior and error likelihoods for all points in the range $[b, c]$, as demonstrated by the following consequence:

Corollary 20. *Suppose \mathcal{D} is any basic action theory with likelihood and precondition axioms for $obs(z)$ as above, f is any \mathcal{L} -fluent, and u is a variable from $\vec{x} = \langle x_1, \dots, x_m \rangle$ that f takes a value from. Then we obtain:*

$$\mathcal{D} \models Bel(a \leq f \leq b, do(obs(z), S_0)) = \frac{\int_{\vec{x}} [P(\vec{x}, f = u \wedge a \leq u \leq b, S_0) \times Err(z, u)]}{\int_{\vec{x}} [P(\vec{x}, f = u, S_0) \times Err(z, u)]}$$

More generally however, and unlike many probabilistic formalisms, we are able to reason about any logical property ϕ of the random variable f being measured.

5.2. Example

Using an example, we demonstrate the formalism and Theorem 19 in particular. To reason about the beliefs of our robot, let us build a simple basic action theory \mathcal{D} . We extend the setting from Figure 1 to a 2-dimensional grid, as shown in Figure 4. As before, let h be the fluent denoting its horizontal position (that is, its distance to the wall), and let the robot's vertical position be given by a fluent v . The components of \mathcal{D} are as below.




Figure 4: A robot in a 2-dimensional grid.

- Imagine a p of the form:

$$p(\iota, S_0) = \begin{cases} .1 \times \mathcal{N}(v(\iota); 0, 16) & \text{if } 2 \leq h(\iota) \leq 12 \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

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.

Note also that initial beliefs can be specified for \mathcal{D}_0 using Bel directly. For example, to express that the true value of h is believed to be uniformly distributed on the interval $[2, 12]$ we might equivalently include the following theory in \mathcal{D}_0 :

$$\{Bel(2 \leq h \leq 12, S_0) = .1, Bel(h \leq 2 \vee h \geq 12, S_0) = 0\},$$

and analogously for the fluent v .

For this example, a simple distribution has been chosen for illustrative purposes. In general, recall from Section 3.4 that the p -specification does not require the variables to be independent, nor does it have to mention all variables.

- For simplicity, let us assume that actions are always executable, *i.e.*, that \mathcal{D} includes

$$Poss(a, s) \equiv true \quad (10)$$

for all actions a . For this example, we assume three action types: action $move(z)$ that moves the robot z units towards the wall, action $up(z)$ that moves the robot z units away from the horizontal axis, and action $sonar(z)$ that gives a reading of z for the distance between the robot and the wall.

- The successor state axiom for h is as in (7), and the one for v is as follows:

$$v(do(a, s)) = u \equiv \neg \exists z (a = up(z) \wedge u = v(s) \vee \exists z (a = up(z) \wedge u = v(s) + z). \quad (11)$$

- For the sensor device, suppose its error model is given as follows:

$$l(sonar(z), s) = u \equiv (z \geq 0 \wedge u = \mathcal{N}(z - h(s); 0, 4)) \vee (z < 0 \wedge u = 0). \quad (12)$$

The error model says that for nonnegative z readings, the difference between the reading and the true value is normally distributed with mean 0 (which indicates that there is no systematic bias) and variance 4.¹⁴

For the remaining (physical) actions, we let

$$l(move(z), s) = 1, \quad l(up(z), s) = 1 \quad (13)$$

since they are assumed to be deterministic for this section.

Then we obtain:

Theorem 21. *Let \mathcal{D} be a basic action theory that is the union of $\{(9), (10), (7), (11), (12), (13)\}$. Then the following are logical entailments of \mathcal{D} :*

1. $Bel([h = 3 \vee h = 4 \vee h = 7], S_0) = 0$.

To see how this follows, let us begin by expanding $Bel([h = 3 \vee h = 4 \vee h = 7], S_0)$:

$$\frac{1}{\gamma} \int_{\vec{x}} P(\vec{x}, h = 3 \vee h = 4 \vee h = 5, S_0). \quad (a)$$

For the rest of the section, let h take its value from x_1 and v take its value from x_2 . By means of (P3), there is exactly one situation for any set of values for x_1 and x_2 . The P term for any such situation, however, is 0 unless $h = 3 \vee h = 4 \vee h = 5$ holds at the situation. Thus, (a) basically simplifies to:

$$\frac{1}{\gamma} \int_{\vec{x}} \begin{cases} 1 \times \mathcal{N}(x_2; 0, 16) & \text{if } x_1 \in \{3, 4, 5\} \\ 0 & \text{otherwise} \end{cases} = 0.$$

In effect, although we are integrating a function $\delta(x_1, x_2)$ over all real values, $\delta(x_1, x_2) = 0$ unless $x_1 \in \{3, 4, 7\}$.

¹⁴For a more elaborate example involving multiple competing sensors and systematic bias, see [36].

2. $Bel(h \leq 9, S_0) = .7$.

We might contrast this with the previous property in that for any given value for x_1 and x_2 , the P term is 0 when $x_1 > 9$. When $x_1 \leq 9$, however, the p value for the situation is obtained from the specification given by (9). That is, we have:

$$\begin{aligned} \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} \text{Initial specification given by (9)} & \text{if } h \leq 9 \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } \exists t. h(t) = x_1, x_1 \in [2, 12] \text{ and } h(t) \leq 9 \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_{\mathbb{R}} \int_2^9 .1 \times \mathcal{N}(x_2; 0, 16) dx_1 dx_2 & \end{aligned}$$

The numerator evaluates to .7, and the denominator to 1.

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. To evaluate this one, for example, observe that we have

$$\begin{aligned} \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } x_1 \in [2, 12] \text{ and } x_1 > 7x_2 \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_2^{12} \int_{-\infty}^{x_1/7} .1 \times \mathcal{N}(x_2; 0, 16) dx_1. & \end{aligned}$$

4. $Bel(h = 0, do(move(4), S_0)) = .2$.

Here a *continuous* distribution evolves into a *mixed* distribution. This results from $Bel(h = 0, do(move(4), S_0))$ first expanding as:

$$\frac{1}{\gamma} \int_{\vec{x}} P(\vec{x}, h = 0, do(move(4), S_0)) \quad (a)$$

The P term, then, simplifies to:

$$\langle t. \bigwedge f(t) = x \wedge (h = 0)[do(move(4), t)] \rightarrow p(t, S_0) \rangle \quad (b)$$

That is, since $move(z)$ has no error component, $l(move(z), s) = 1$ for any s in accordance with \mathcal{D} . Therefore, $p(do(a, t), do(a, S_0)) = p(t, S_0)$. Now (b) says that for every possible value for h and v , if there is an initial situation where $h = 0$ holds after moving leftwards, then its p value is to be considered. Note that for any initial situation s where $h(s) \in [2, 4]$, we get $h(do(move(4), s)) = 0$ by (7). This leaves us with:

$$\begin{aligned} \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } \exists t. h(t) = x_1, x_1 \in [2, 12], h(t) \in [2, 4] \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } x_1 \in [2, 4] \\ 0 & \text{otherwise} \end{cases} & \end{aligned} \quad (c)$$

We can show that $\gamma = 1$, which means (c) = .2. This change in beliefs is shown in Figure 5.

5. $Bel(h \leq 3, do(move(4), S_0)) = .5$.

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.

$$6. \text{Bel}([\exists a, s. \text{now} = \text{do}(a, s) \wedge h(s) > 1], \text{do}(\text{move}(4), S_0)) = 1.$$

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.

$$7. \text{Bel}(h = 4, \text{do}([\text{move}(4), \text{move}(-4)], S_0)) = .2$$

$$\text{Bel}(h = 4, \text{do}([\text{move}(-4), \text{move}(4)], S_0)) = 0.$$

The point $h = 4$ has 0 weight initially, as shown in item 1. Roughly, if the agent were to move leftwards first then many points would “collapse”, as shown in item 4. The point would then obtain a h value of 0, and have a weight of .2. The weight is then retained on moving away by 4 units, where the point once again gets h value 4. On the other hand, if this entire phenomena were reversed then none of these features are observed because the collapsing does not occur and the entire space remains fully continuous.

$$8. \text{Bel}(-1 \leq v \leq 1, \text{do}(\text{move}(4), S_0)) = \text{Bel}(-1 \leq v \leq 1, S_0) = \int_{-1}^1 \mathcal{N}(x_2; 0, 16) dx_2.$$

Owing to Reiter’s solution to the frame problem, belief in v is unaffected by a lateral motion. That is, a leftwards motion does not change v in accordance with (11). As per (9), the initial belief in $v \in [-1, 1]$ is the area between $[-1, 1]$ bounded by the specified Gaussian.

$$9. \text{Bel}(v \leq 7, \text{do}(\text{up}(2.5), S_0)) = \text{Bel}(v \leq 4.5, S_0).$$

After the action $\text{up}(2.5)$, 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 $v = x_2 + 2.5$. Intuitively, we have:

$$\begin{aligned} \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } \exists t. v(t) = x_2 \text{ and } (v \leq 7)[\text{do}(\text{up}(2.5), t)] \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } \exists t. v(t) = x_2 \text{ and } v(t) \leq 4.5 \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_{\vec{x}} \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } x_2 \leq 4.5 \\ 0 & \text{otherwise} \end{cases} & \end{aligned}$$

$$10. \text{Bel}(h \leq 9, \text{do}(\text{sonar}(5), S_0)) \approx .97.$$

$$\text{Bel}(h \leq 9, \text{do}([\text{sonar}(5), \text{sonar}(5)], S_0)) \approx .99.$$

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 (P2), incorporates the likelihood of each $\text{sonar}(5)$ action. More precisely, the belief term in the first entailment simplifies to:

$$\frac{1}{\gamma} \int_{\vec{x}} \langle t. h(t) = x_1 \wedge v(t) = x_2 \wedge h(t) \leq 9 \rightarrow p(t, S_0) \times \mathcal{N}(5 - x_1; 0, 4) \rangle \quad (a)$$

Note that we have replaced $(h \leq 9)[\text{do}(\text{sonar}(5), S_0)]$ by $(h \leq 9)[t]$ since $\text{sonar}(z)$ does not affect h . From (a), we get

$$\frac{1}{\gamma} \int_{\vec{x}} \mathcal{N}(5 - x_1; 0, 4) \times \langle t. h(t) = x_1 \wedge v(t) = x_2 \wedge h(t) \leq 9 \rightarrow p(t, S_0) \rangle \quad (b)$$

We know from (9) that those initial situations where $h \leq 2$ have p values 0. Therefore, from (b), we get:

$$\begin{aligned} \frac{1}{\gamma} \int_{\vec{x}} \mathcal{N}(5 - x_1; 0, 4) \times \begin{cases} .1 \times \mathcal{N}(x_2; 0, 16) & \text{if } x_1 \in [2, 9] \\ 0 & \text{otherwise} \end{cases} &= \\ \frac{1}{\gamma} \int_{\mathbb{R}} \int_2^9 \mathcal{N}(5 - x_1; 0, 4) \times .1 \times \mathcal{N}(x_2; 0, 16) dx_1 dx_2 & \end{aligned}$$

After a second reading of 5 from the sonar, the expansion for belief is analogous, except that the function to be integrated gets multiplied by a second $\mathcal{N}(5 - x_1; 0, 4)$ term. It is then not hard to see that belief sharpens significantly with this multiplicand. The agent’s changing densities are shown in Figure 6.




Figure 5: Belief update for h after physical actions. Initial belief at S_0 (in solid red) and after a leftward move of 4 units (in blue with point markers).




Figure 6: Belief change for h at S_0 (in solid red), after sensing 5 (in green with circular markers), and after sensing 5 twice (in blue with square markers).

6. Noisy Acting

In the presentation so far, we assumed physical actions to be deterministic. By this we mean that when a physical action a occurs, it is clear to us (as modelers) but also the agent how the world has changed on a . Of course, in realistic domains, especially robotic applications, this is *not* the usual case. In this section, in a domain that has continuous fluents, we show how our current account of belief can be extended to reason with sensors as well as effectors that are noisy.

In line with the rest of this work, effector noise is given a quantitative account. Let us first reflect on what is expected with noisy acting. When an agent senses, as in the case of $sonar(z)$, the argument for this action is not chosen by the agent. That is, the world decides what z should be, and based on this reading of z , the agent comes to certain conclusions about its own state. The noise factor, then, simply addresses the phenomena that the number z returned may differ from the true value of whatever fluent the sensor is measuring.

Noisy acting diverges from that picture in the following sense. The agent *intends* to do action a , but what actually occurs is a' that is possibly different from a . For example, the agent may want to move 3 units, but, unbeknownst to the agent, it may move by 3.042 units. The agent, of course, does not observe this outcome. Nevertheless, provided the agent has an account of its effector's inaccuracies, it is reasonable for the agent to believe that it is in fact closer to the wall, even if it may not be able to precisely tell by how much. Intuitively, the result of a nondeterministic action is that any number of successor situations might be obtained, which are all indistinguishable in the agent's perspective (until sensing is performed). Depending on the likelihoods of the action's potential outcomes, some of these successor situations are considered more probable than others. The agent's belief about what holds then must incorporate these relative likelihoods. So, in our view, nondeterminism is really an *epistemic* notion.

6.1. Noisy Action Types

Following [5], perhaps the simplest extension to make all this precise is to assume that deterministic actions such as $move(x)$ now have companion action types $move(x, y)$ in \mathcal{L} . The intuition is that x represents the *nominal* value, which is the number of units that the agent intends to move, while y represents the actual distance moved. The actual value of y in any ground action, of course, is not observable for the agent. This simple idea will need three adjustments to our account:

1. successor state axioms need to be built using these new action types;
2. the formalism must allow the modeler to formalize that certain outcomes are more likely than others, that is, noisy actions may be associated with a probabilistic account of the various outcomes; and
3. the notion of belief must incorporate the nominal value, the range of possible outcomes and their likelihoods.

First, we address successor state axioms. These are now specified as usual, but using the second argument, which is the actual outcome, rather than the nominal value, which is ignored. For example, for the fluent h , instead of (7), we will now have:

$$h(do(a, s)) = u \equiv \exists x, y [a = move(x, y) \wedge u = \max(0, h - y)] \vee \neg \exists x, y [a = move(x, y) \wedge u = h(s)]. \quad (14)$$

The reason for this modification is obvious. If y is the actual outcome then the fluent change should be contingent on this value rather than what was intended. It is important to note that no adjustment to the existing (Reiter's) solution to the frame problem is necessary.

6.2. The Golog approach

The foremost issue now is to use the above idea to allow for more than one possible successor situation. Clearly, we do not want the agent to control the actual outcome in general. So the approach taken by BHL is to think of picking the second argument as a *nondeterministic Golog* program [11]. Briefly, *Golog* is an agent programming proposal where one is allowed to formulate complex actions that denote sequential and nondeterministic executions of actions, among others, and is essentially a basic action theory. Given the action $move(x, y)$, for example, the *Golog* program $MOVE(x)$ might stand for the abbreviation $\pi y. move(x, y)$, which corresponds to a ground action $move(x, n)$

where n is chosen nondeterministically. For our purposes, we would then imagine that the agent executes *Golog* programs.

There are some advantages to this approach: namely, we only have to look at the logical entailments, including ones mentioning *Bel*, of such *Golog* programs. Since traces of these programs account for many potential outcomes, *Bel* does the right thing and accommodates all of these when considering knowledge. But the disadvantage is that the resulting formal specification turns out to be unnecessarily complex, at least as far as projection is concerned.

For projection tasks, we show that we can settle on a simpler alternative, one that does not appeal to *Golog*. Like BHL, we assume that the world is deterministic, where the result of doing a ground action leads to a distinct successor. Roughly, the intuition then is that when a noisy action is performed, the various outcomes of the action as well as the potential successor situations that are obtained wrt these are treated at the level of belief.

6.3. Alternate Action Axioms

Inspired by [37], our approach is based on the introduction of a distinguished predicate *Alt*. The idea is this: if $Alt(a, a', \vec{z})$ holds for ground action $a = A(\vec{c})$ then we understand this to mean that the agent believes that any instance of $a' = A(\vec{z})$ might have been executed instead of a . Here, \vec{z} denotes the range of the arguments for potential outcomes.

To see how that gets used with the new action types such as $move(x, y)$, consider the ground action $move(3, 3.1)$. So, the agent intends to move by 3 units but what has actually occurred is a move by 3.1 units. Since the agent does not observe the latter argument, from its perspective, what occurred could have been a move by 2.9 units, but also perhaps (although less likely) a move by 9 units. Thus, the ground actions $move(3, 2.9)$ and $move(3, 9)$ are *Alt*-related to $move(3, 3.1)$. (The likelihoods for these may vary, of course.) In logical terms, we might have an axiom of the following form in the background theory:

$$Alt(move(x, y), a', z) \equiv a' = move(x, z). \quad (15)$$

This to be read as saying that $move(3, z)$ for every $z \in \mathbb{R}$ are alternatives to $move(3, 3.1)$. If we required that z is only a certain range from 3, for example, we might have:

$$Alt(move(x, y), a', z) \equiv a' = move(x, z) \wedge |x - z| \leq c$$

where c bounds the magnitude of the maximal possible error. On the other hand, for actions such as $sonar(z)$, which do not have any alternatives, we simply write:

$$Alt(sonar(x), a', z) \equiv a' = sonar(z) \wedge z = x. \quad (16)$$

This says that $sonar(x)$ is only *Alt*-related to itself.

With this simple technical device, one can now additionally constrain the likelihood of various outcomes using l . For example:

$$l(move(x, y), s) = u \equiv u = \mathcal{N}(y - x; \mu, \sigma^2) \quad (17)$$

says that the difference between nominal value and the actual value is normally distributed with mean μ and variance σ^2 . This essentially corresponds to the standard additive Gaussian noise model in robotics [13].

To see an example of how, say, (15) and (17) work together with the successor state axiom (P2) for p , consider three situations s, s' and s'' associated with the same density, as shown in Figure 7. Suppose their h values are 6, 6.1 and 5.9 respectively. After attempting to move 3 units, the action $move(3, z)$ for any $z \in \mathbb{R}$ may have occurred. So, for each of the three situations, we explore successors from different values for z . Assume the motion effector is defined by a mean $\mu = 0$ and variance $\sigma^2 = 1$. Then, the p -value of the situation $do(move(3, 5.7), s')$, for example, is obtained from the p -value for s' multiplied by the likelihood of $move(3, 5.7)$, which is $\mathcal{N}(5.7 - 3; 0, 1) = \mathcal{N}(2.7; 0, 1)$. Thus, the successor situation $do(move(3, 5.7), s')$ is much less likely than the successor situation $do(move(3, 3), s)$, as should be the case.

In general, we define *alternate actions* axioms that are to be a part of the basic action theory henceforth:

Definition 22. Let $A(\vec{x}, \vec{y})$ be any action. *Alternate actions axioms* are sentences of the form:

$$Alt(A(\vec{x}, \vec{y}), a', \vec{z}) \equiv a' = A(\vec{x}, \vec{z}) \wedge \psi(\vec{x}, \vec{y}, \vec{z})$$

where ψ is a formula that characterizes the relationship between the nominal and true values.

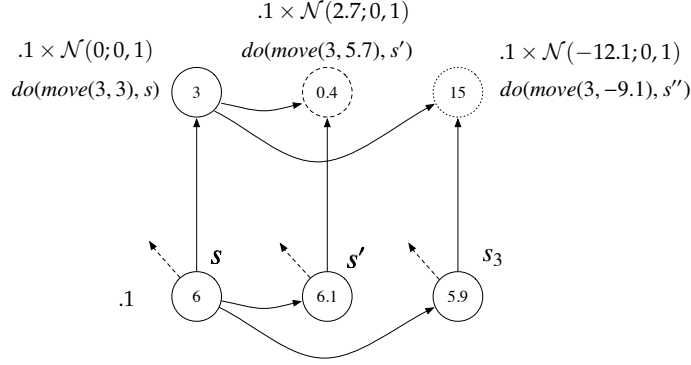


Figure 7: Situations with accessibility relations after noisy acting. The numbers inside the circles denote the h values at these situations. Dotted circles denote a lower density in relation to their epistemic alternatives.

The one limitation with this definition is that only actions of the same *type*, *i.e.*, built from the same function symbol, are alternatives to each other. This does not allow, for example, situations where the agent intends a physical move, but instead unlocks the door. Nevertheless, this definition is not unreasonable because noisy actions in robotic applications typically involve additive noise [13]. Moreover, this limitation only assists us in arriving at a simple and familiar definition for belief. A more involved definition would allow for other variants.

6.4. A Definition for Belief

We have thus far successfully augmented successor state axioms and extended the formalism for modeling noisy actions. The final question, then, is how can the outcomes of a noisy action, and their likelihoods, be accounted for? Indeed, a formula might not only be true as a result of the actions intended, but also as a result of those that were not.

Consider the simple case of deterministic actions, where the density associated with s is simply transferred to $do(a, s)$. This is an instance of Lewis's *imaging* [27]. In contrast, if a and a' are *Alt*-related, then the result of doing a at s would lead to successor situations $do(a, s)$ and $do(a', s)$. Moreover, unlike noisy sensors, a and a' may affect fluent values in different ways, which is certainly the case with $move(3, 3.074)$ and $move(3, 3)$ on the fluent h . Thus, the idea then is that when reasoning about the agent's beliefs about ϕ , one would need to integrate over the densities of all those potential successors where ϕ would hold.

To make this precise, let us first consider the result of doing a single action a at S_0 . The *degree of belief* in ϕ after doing a is now an abbreviation for:

$$Bel(\phi, do(a, S_0)) \doteq \frac{1}{\gamma} \int_{\vec{x}} \int_z P(\vec{x}, z, \phi, do(a, S_0))$$

where

$$P(\vec{x}, z, \phi, do(a, S_0)) \doteq \langle \iota, b. \bigwedge f_i(\iota) = x_i \wedge Alt(a, b, z) \wedge \phi[do(b, \iota)] \rightarrow p(do(b, \iota), do(a, S_0)) \rangle.$$

(As before, the i ranges over the indices of the fluents in \mathcal{L} , that is, $\{1, \dots, n\}$.) The intuition is this. Recall that by integrating over \vec{x} , all possible initial situations are considered by $f_i(\iota) = x_i$. Analogously, by integrating over z , all possible action outcomes are considered by $Alt(a, b, z)$. Supposing $a = A(x, y)$, for each outcome $b = A(x, z)$,¹⁵ we test whether ϕ holds at the resulting situation $do(b, \iota)$ as before, and use its p -value. Here, this p -value is given by $p(do(b, \iota), do(a, S_0))$, where the first argument is the successor of interest $do(b, \iota)$ and the second is the real world $do(a, S_0)$.

The generalization, then, for a sequence of actions is as follows:

¹⁵For ease of presentation, we assume that the nominal and the actual arguments involve a single variable.

Definition 23. (Degrees of belief (continuous noisy effectors and sensors.)) Suppose ϕ is any \mathcal{L} -formula. Then the *degree of belief* in ϕ at s , written $Bel(\phi, s)$, is defined as an abbreviation:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x}} \int_{\vec{z}} P(\vec{x}, \vec{z}, \phi, s)$$

where, if $s = do([a_1, \dots, a_k], S_0)$, then

$$P(\vec{x}, \vec{z}, \phi, s) \doteq \langle \iota, b_1, \dots, b_k. \wedge f_i(\iota) = x_i \wedge \wedge Alt(a_j, b_j, z_j) \wedge \phi[do([b_1, \dots, b_k], \iota)] \rightarrow p(do([b_1, \dots, b_k], \iota), do([a_1, \dots, a_k], S_0)) \rangle.$$

(Here, i ranges over $\{1, \dots, n\}$ as before, and j ranges over the indices of the ground actions $\{1, \dots, k\}$.) That is, given any sequence, for all possible \vec{z} values, we consider alternate sequences of ground action terms and integrate the densities of successor situations that satisfy ϕ , using the appropriate p -value.

6.5. Example

Let us now build a simple example with noisy actions. Consider the robot scenario in Figure 1. Suppose the basic action theory \mathcal{D} includes the foundational axioms, and the following components.

The initial theory \mathcal{D}_0 includes the following p specification:

$$p(\iota, S_0) = \begin{cases} .5 & \text{if } 10 \leq h(\iota) \leq 12 \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

For simplicity, let h be the only fluent in the domain, and assume that actions are always executable. The successor state axiom for the fluent h is (14). For p , it is the usual one, *viz.* (P2).

We imagine two actions in this domain, one of which is the noisy move $move(x, y)$ and a sonar sensing action $sonar(z)$. For the alternate actions axioms, let us use (15) and (16).

Finally, we specify the likelihood axioms. Let the sonar's error profile be

$$l(sonar(z), s) = u \equiv (z \geq 0 \wedge u = \mathcal{N}(z - h(s); 0, .25)) \vee (z < 0 \wedge u = 0). \quad (19)$$

Readers may note that this sonar is more accurate than the one characterized by (12), as it has a smaller variance. Regarding the likelihood axiom for $move(x, y)$, let that be:

$$l(move(x, y), s) = u \equiv u = \mathcal{N}(y - x; 0, 1). \quad (20)$$

This completes the specification of \mathcal{D} .

Theorem 24. *The following are entailments of \mathcal{D} .*

1. $Bel(h \geq 11, do(move(-2, -2.01), S_0)) \approx .95$

We first observe that for calculating the degrees of belief, we have to consider all those successors of initial situations wrt $move(-2, z)$ for every z , where ϕ holds. By (P2), the p value for such situations is the initial p value times the likelihood of $move(-2, z)$, which is $\mathcal{N}(z + 2; 0, 1)$ by (20). Therefore, we get

$$\begin{aligned} & \frac{1}{\gamma} \int_x \int_z \begin{cases} .5 \times \mathcal{N}(z + 2; 0, 1) & \text{if } \exists \iota. x \in [10, 12], h(\iota) = x \text{ and } (h \geq 11)[do(move(-2, z), \iota)] \\ 0 & \text{otherwise} \end{cases} = \\ & \frac{1}{\gamma} \int_x \int_z \begin{cases} .5 \times \mathcal{N}(z + 2; 0, 1) & \text{if } \exists \iota. x \in [10, 12], h(\iota) = x \text{ and } h(\iota) - z \geq 11 \\ 0 & \text{otherwise} \end{cases} = \\ & \frac{1}{\gamma} \int_{-\infty}^{\infty} \int_{10}^{12} \begin{cases} .5 \times \mathcal{N}(z + 2; 0, 1) & \text{if } x \in [10, 12] \text{ and } x - z \geq 11 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$




Figure 8: Belief change for h at S_0 (in solid red), a noisy move away from the wall (in green with circular markers), and after a second noisy move towards the wall (in blue with square markers).

It is not hard to see that had the action been deterministic, the degree of belief in $h \geq 11$ after moving away by 2 units should have been precisely 1. In Figure 8, we see the effect of this move, where the range of h values with non-zero densities extends considerably more than 2 units.

2. $Bel(h \geq 10, do([move(-2, -2.01), move(2, 2.9)], S_0)) \approx .74$

The argument proceeds in a manner identical to the previous demonstration. The density function is further multiplied by a factor of $\mathcal{N}(u - 2; 0, 1)$, from (20) and (P2). More precisely, we have

$$\frac{1}{\gamma} \int_x \int_z \int_u \begin{cases} .5 \times \mathcal{N}(z + 2; 0, 1) \times \mathcal{N}(u - 2; 0, 1) & \text{if } \exists t. x \in [10, 12], h(t) = x \text{ and } h(t) - z - u \geq 10 \\ 0 & \text{otherwise} \end{cases} =$$

$$\frac{1}{\gamma} \int_z \int_u \int_{10}^{12} \begin{cases} .5 \times \mathcal{N}(z + 2; 0, 1) \times \mathcal{N}(u - 2; 0, 1) & \text{if } x \in [10, 12] \text{ and } x - z - u \geq 10 \\ 0 & \text{otherwise} \end{cases}$$

If the action were deterministic, yet again the degree of belief about $h \geq 10$ would be 1 after the intended actions. That is, the robot moved away by 2 units and then moved towards the wall by another 2 units, which means that h 's current value should have been precisely what the initial value was.

See Figure 8 for the resulting density change. Intuitively, the resulting density changes as effectuated by the moves degrades the agent's confidence considerably. In Figure 8, for example, we see that in contrast to a single noisy move, the range of h values considered possible has extended further, leading to a wide curve.

3. $Bel(h \geq 11, do([move(-2, -2.01), sonar(11.5)], S_0)) \approx .94$

This demonstrates the result of a sensing action after a noisy move. Using arguments analogous to those in the previous item, it is not hard to see that we have:

$$\frac{1}{\gamma} \int_x \int_z \begin{cases} .5 \times \mathcal{N}(z + 2; 0, 1) \times \mathcal{N}(x - z - 11.5; 0, .25) & \text{if } \exists t. x \in [10, 12], h(t) = x \text{ and } h(t) - z \geq 11 \\ 0 & \text{otherwise} \end{cases}$$

4. $Bel(h \geq 11, do([move(-2, -2.01), sonar(11.5), sonar(12.6)], S_0)) \approx .99$

In this case, two successive readings around 12 strengthens the agent's belief about $h \geq 11$. The density function




Figure 9: Belief change for h at S_0 (in solid red), after a noisy move (in green with circular markers), after sensing once (in blue with square markers), and after sensing twice (in solid magenta).

is multiplied by $\mathcal{N}(x - z - 12.6; 0, .25)$ because of (P2) and (12) as follows:

$$\frac{1}{\gamma} \int_x \int_z \begin{cases} \delta \times \mathcal{N}(x - z - 11.5; 0, .25) \times \mathcal{N}(x - z - 12.6; 0, .25) & \text{if } \exists t. x \in [10, 12], h(t) = x, h(t) - z \geq 11 \\ 0 & \text{otherwise} \end{cases} =$$

$$\frac{1}{\gamma} \int_z \int_x \begin{cases} \delta \times \mathcal{N}(x - z - 11.5; 0, .25) \times \mathcal{N}(x - z - 12.6; 0, .25) & \text{if } x \in [10, 12] \text{ and } x - z \geq 11 \\ 0 & \text{otherwise} \end{cases}$$

where $\delta = .5 \times \mathcal{N}(z + 2; 0, 1)$, and γ is

$$\int_z \int_x \begin{cases} \delta \times \mathcal{N}(x - z - 11.5; 0, .25) \times \mathcal{N}(x - z - 12.6; 0, .25) & \text{if } x \in [10, 12] \\ 0 & \text{otherwise} \end{cases}$$

In Figure 9, the agent's increasing confidence is shown as a result of these sensing actions. Note that even though the sensors are noisy, the agent's belief about h 's true value sharpens because the sensor is a fairly accurate one.

7. Generalization

Many real-world problems have both continuous and discrete components (sensors, fluents, and/or effectors). Not surprisingly, discrete sensors can be easily modeled in the current scheme, as they only affect the p -values. Regarding fluents and effectors, it turns out that accommodating the more general case is an easy exercise, where an integration symbol in Bel corresponding to a continuous fluent or action argument is replaced by a summation symbol.

To clarify, we proceed as follows. We begin with an example for discrete sensors, introduce a general definition for Bel in the above sense, and finally conclude with an example that demonstrates this general setting.

7.1. Example

We understand a discrete sensor to mean a sensing action that is characterized by a finite number of possible observations. Thus, these observations would be associated with a probability rather than a density. Imagine the robot scenario from Figure 1. Suppose that instead of a sensor that returns a number indicating the distance to the wall, the robot is equipped with a crude binary version. This latter sensor simply indicates whether the robot is close or far from the wall.

Formally, suppose there is a sensing action $sensewall(z)$ where $z \in \{close, far\}$. A noise-free model for the closeness sensor might be as follows:

$$l(sensewall(z), s) = u \equiv (z = close \wedge ((h(s) \leq 1 \wedge u = 1) \vee (h(s) > 1 \wedge u = 0))) \vee \\ (z = far \wedge ((h(s) > 1 \wedge u = 1) \vee (h(s) \leq 1 \wedge u = 0))).$$

For the more interesting case of a noisy sensor, assume the following error profile:

$$l(sensewall(z), s) = u \equiv (z = close \wedge h(s) \leq 3 \wedge u = 2/3) \vee \\ (z = close \wedge h(s) > 3 \wedge u = 1/3) \vee \\ (z = far \wedge h(s) \geq 4 \wedge u = 4/5) \vee \\ (z = far \wedge h(s) < 4 \wedge u = 1/5). \quad (21)$$

Notice that the behavior of the sensor differs for the two values. So, for $z = close$ the sensor considers $h \leq 3$ as a measure of closeness, and has an accuracy of $2/3$. For $z = far$, however, the sensor takes $h \geq 4$ to be a measure of being far, and has an accuracy $4/5$.

We now build a simple action theory using this sensor. For simplicity, assume noise-free physical actions and a single fluent, h . Let \mathcal{D} be the union of the p specification (8), the successor state axiom (7), the above likelihood model (21), in addition to (P1), (P2) and (P3). For the actions $A(z)$ in \mathcal{D} , including the sensor, let

$$Alt(A(z), a', u) \equiv a' = A(u) \wedge u = z.$$

We now analyze belief change on the application of this sensor.

Theorem 25. *The following are entailments of \mathcal{D} :*

1. $Bel(h \leq 4, do(sensewall(close), S_0)) = 3/11$

By (P2), after $a = sensewall(close)$ the term $p(do(a, s'), do(a, s))$ is obtained from $p(s', s) \times Err(close, h(s'))$, where $Err(u_1, u_2)$ denotes the error profile (21) of the closeness sensor. To compute the belief term, let us first resolve the normalization factor γ . It is not hard to see that γ evaluates to

$$\int_x \begin{cases} .1 \times 2/3 & \text{if } \exists t. h(t) = x, x \in [2, 12] \text{ and } h(t) \leq 3 \\ .1 \times 1/3 & \text{if } \exists t. h(t) = x, x \in [2, 12] \text{ and } h(t) > 3 \\ 0 & \text{otherwise} \end{cases}$$

Since \mathcal{D}_0 assigns a non-zero density to only those situations where $h \in [2, 12]$ one obtains the above normalization factor. Moreover, after doing $sensewall(close)$, the density of those situations where $h \leq 3$ is multiplied by a factor of $2/3$, while the density of the remaining situations is multiplied by a factor of $1/3$.

Formulating the numerator is analogous, except that only those situations where $h \leq 4$ are to be considered. To be precise, the degree of belief in $h \leq 4$ after the sensing action is:

$$\frac{1}{\gamma} \int_x \begin{cases} .1 \times 2/3 & \text{if } x \in [2, 12], x \leq 3, x \leq 4 \\ .1 \times 1/3 & \text{if } x \in [2, 12], x > 3, x \leq 4 \\ 0 & \text{otherwise} \end{cases} = \\ \frac{1}{\gamma} \int_x \begin{cases} .1 \times 2/3 & \text{if } x \in [2, 3] \\ .1 \times 1/3 & \text{if } x \in (3, 4] \\ 0 & \text{otherwise} \end{cases}$$

This then amounts to $3/11$.

2. $Bel(h \leq 4, do([move(1), sensewall(close)])) = 5/12$

We proceed in a manner analogous to the previous item. Using (P2) and (21), we obtain:

$$\frac{1}{\gamma} \int_x \begin{cases} .1 \times 2/3 & \text{if } \exists t. h(t) = x, h(\text{do}(\text{move}(1), t)) \leq 3 \text{ and } h(\text{do}(\text{move}(1), t)) \leq 4 \\ .1 \times 1/3 & \text{if } \exists t. h(t) = x, h(\text{do}(\text{move}(1), t)) > 3 \text{ and } h(\text{do}(\text{move}(1), t)) \leq 4 \\ 0 & \text{otherwise} \end{cases} =$$

$$\frac{1}{\gamma} \int_x \begin{cases} .1 \times 2/3 & \text{if } x \in [2, 12], x - 1 \leq 3 \text{ and } x - 1 \leq 4 \\ .1 \times 1/3 & \text{if } x \in [2, 12], x - 1 > 3 \text{ and } x - 1 \leq 4 \\ 0 & \text{otherwise} \end{cases}$$

The numerator amounts to 5/3 and the normalization factor γ is 4, leading to a degree of belief of 5/12.

7.2. A General Definition for Belief

The main idea is to simply allow the range of some fluents to be taken from finite sets. So, reconsider Definition 18, where the fluents and the action arguments were assumed to take values from \mathbb{R} . Then, suppose fluents f_1, \dots, f_n takes values from \mathbb{R} , and fluents g_1, \dots, g_m take values from finite sets. Intuitively, f_i is to be seen as a continuous probabilistic variable, and g_i as a discrete probabilistic variable.¹⁶ Analogously, suppose a_1, \dots, a_o are action types such that $Alt(a_j, b, z)$ holds for $z \in \mathbb{R}$, and suppose d_1, \dots, d_l are action types such that $Alt(d_{\bar{j}}, r, u)$ holds for u taken from a finite set. Intuitively, a_j is to be seen as a noisy action characterized by a continuous probability distribution, and $d_{\bar{j}}$ as a noisy action characterized by a discrete probability distribution. Then:

Definition 26. (Degrees of belief (discrete and continuous fluents, sensors and effectors)). Suppose ϕ is any \mathcal{L} -formula. The *degree of belief* in ϕ at s , written $Bel(\phi, s)$, is an abbreviation:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x}} \sum_{\vec{y}} \int_{\vec{z}} \sum_{\vec{u}} P(\vec{x} \cdot \vec{y} \cdot \vec{z} \cdot \vec{u}, \phi, s)$$

where if $s = \text{do}(\alpha, S_0)$ for $\alpha = [a_1, \dots, d_1, \dots, a_o, \dots, d_l]$, then

$$P(\vec{x} \cdot \vec{y} \cdot \vec{z} \cdot \vec{u}, \phi, s) \doteq$$

$$\langle t, b_1, \dots, b_o, r_1, \dots, r_l. \wedge f_i(t) = x_i \wedge \wedge g_{\bar{i}}(t) = y_{\bar{i}} \wedge$$

$$\wedge Alt(a_j, b_j, z_j) \wedge \wedge Alt(d_{\bar{j}}, r_{\bar{j}}, u_{\bar{j}}) \wedge \phi[\text{do}([b_1, \dots, r_l], t)]$$

$$\rightarrow p(\text{do}([b_1, \dots, r_l], t), \text{do}([a_1, \dots, d_l], S_0)) \rangle.$$

Here, \cdot is for the concatenation of variables, i ranges over the indices of the continuous fluents, \bar{i} over the indices of the discrete fluents, j over the indices of the continuous actions, and \bar{j} over the indices of the discrete actions.

Naturally, by means of (P3), there is an initial situation for every possible real number for f_1, \dots, f_n and for every possible vector of values for g_1, \dots, g_m . The intuition is as before. That is, owing to a bijection between situations and the vector of fluent values, for any given value for $x_1, \dots, x_n, y_1, \dots, y_m$, there is a unique initial situation where f_i has the value x_i and $g_{\bar{i}}$ has the value $y_{\bar{i}}$. The only difference to what we had earlier is that instead of just integrating over possible values for \vec{x} , of course, we integrate over values for \vec{x} and sum over possible values for \vec{y} while using the p -values of successor situations where ϕ holds. When a noisy action occurs, the space of possible successor situations is determined by $Alt(a_j, b, z)$ for a noisy action with a continuous probabilistic model, and the space is determined by $Alt(d_{\bar{j}}, r, u)$ otherwise.

¹⁶Of course, discrete probabilistic variables can also take values from infinite sets, and by way of Section 2.5, limits can be used to define the sum of an infinite sequence of terms, that is, a series. For simplicity of presentation, however, we assume discrete fluents and action arguments take values from finite sets.

7.3. Example

To demonstrate the more intricate case of discrete and continuous fluents, and discrete and continuous action arguments, we consider an example of a robot moving towards the wall with a window, as shown in Figure 10. Like with Figure 1, let a fluent h denote the distance to the wall. Let w be a fluent that captures the status of the window in the sense of whether it is opened or closed, with $w = 1$ meaning that it is open and $w = 0$ meaning that it is closed.

To change the status of the window, we imagine a noisy effector $setwin(x, y)$ where the agent intends on setting w to x , but it is, in fact, set to y . We formally account for this effector using:

$$Alt(setwin(x, y), a', z, s) \equiv a' = setwin(x, z) \quad (22)$$

$$l(setwin(x, y), s) = u \equiv ((x = 0 \vee x = 1) \wedge x = y \wedge u = .75) \vee ((x = 0 \vee x = 1) \wedge (y = 0 \vee y = 1) \wedge |x - y| = 1 \wedge u = .25) \vee ((\neg(x = 0 \vee x = 1) \vee \neg(y = 0 \vee y = 1)) \wedge u = 0). \quad (23)$$

This can be seen as saying that $\{x, y\}$ take values $\{0, 1\}$, and the likelihood of x and y agreeing is $.75$, and that of disagreeing is $.25$. So, it is very likely that the action succeeds.

From this, we then provide the following successor state axiom for w :

$$w(do(a, s)) = u \equiv \exists x, y [setwin(x, y) \wedge u = y] \vee \neg \exists x, y [setwin(x, y) \wedge u = w(s)]. \quad (24)$$

For simplicity, assume that the action $move(z)$ moves the robot towards the wall by z units (that is, it is deterministic), for which we use the successor state axiom (7).

Analogously, imagine a noisy sensor $seewin(z)$ that provides a reading of x for the status of the window, but with the following error profile:

$$l(seewin(z), s) = u \equiv (z = 1 \wedge w(s) = 1 \wedge u = .8) \vee (z = 1 \wedge w(s) = 0 \wedge u = .2) \vee (z = 0 \wedge w(s) = 1 \wedge u = .3) \vee (z = 0 \wedge w(s) = 0 \wedge u = .7). \quad (25)$$

That is, when returning 1, the sensor gives the correct reading with a probability of $.8$, but when returning 0, the sensor gives the correct reading with a probability of $.7$. Alternate action axioms are specified as usual for sensors:

$$Alt(seewin(x), a', z, s) \equiv a' = seewin(z) \wedge z = x. \quad (26)$$




Figure 10: Robot moving towards a wall that has a window.

To finalize the example, let \mathcal{D} be a union of (7), (22), (23), (24), (26), (25), together with (P1), (P2), (P3), and the following initial axiom for p :

$$p(t, S_0) = \begin{cases} .5 \times .6 & \text{if } 10 \leq h(t) \leq 12 \text{ and } w(t) = 1 \\ .5 \times .4 & \text{if } 10 \leq h(t) \leq 12 \text{ and } w(t) = 0 \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

That is, h and w are independent, h is uniformly distributed on $[10, 12]$ and the window being open has a probability of .6.

Theorem 27. *The following are entailed by \mathcal{D} :*

1. $Bel(w = 0, S_0) = .4$

We are to compute the following term, which is easily shown to be equal to .4:

$$\frac{1}{\gamma} \int_x \sum_y \begin{cases} .5 \times .6 & \exists t. x \in [10, 12], h(t) = x, w(t) = y, y = 1 \text{ and } w(t) = 0 \\ .5 \times .4 & \exists t. x \in [10, 12], h(t) = x, w(t) = y, y = 0 \text{ and } w(t) = 0 \\ 0 & \text{otherwise} \end{cases}$$

2. $Bel(w = 0, do(move(1), S_0)) = .4$

Owing to Reiter's solution to the frame problem, belief in the closed window does not change on moving laterally. If we were to expand the belief term, we could get a logical term equivalent to the one in the previous item.

3. $Bel(w = 0, do(setwin(0, 1), S_0)) = .75$

After attempting to close the window, the agent's belief about the window being closed is .75. Not surprisingly, since the action sets the final value of w (as opposed to toggle its value), the degree of belief precisely corresponds to the probability of the action getting successfully executed in (23). Expanding the belief term, we first obtain:

$$\frac{1}{\gamma} \int_x \sum_y \sum_z \begin{cases} .5 \times .6 \times \delta(0, z) & \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 1, (w = 0)[do(setwin(0, z), t)] \\ .5 \times .4 \times \delta(0, z) & \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 0, (w = 0)[do(setwin(0, z), t)] \\ 0 & \text{otherwise} \end{cases}$$

Here $\delta(0, z)$ is the probability assigned to the substitution of z in $setwin(0, z)$ via (23). Simplifying the above, we get:

$$\frac{1}{\gamma} \int_x \sum_y \sum_z \begin{cases} .5 \times .6 \times .75 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 1, z = 0, (w = 0)[do(setwin(0, z), t)] \\ .5 \times .4 \times .75 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 0, z = 0, (w = 0)[do(setwin(0, z), t)] \\ 0 & \text{otherwise} \end{cases} =$$

$$\frac{1}{\gamma} \int_x \sum_y \sum_z \begin{cases} .5 \times .6 \times .75 & \text{if } x \in [10, 12], y = 1, z = 0 \\ .5 \times .4 \times .75 & \text{if } x \in [10, 12], y = 0, z = 0 \\ 0 & \text{otherwise} \end{cases}$$

Essentially, there are only two ways that the window can be closed after doing $setwin(0, z)$. Either the window is closed initially and $z = 0$, or the window is open and again $z = 0$. This leads to .75.

4. $Bel(w = 0, do([setwin(0, 1), seewin(0)], S_0)) = .875$

After observing that the window is closed from its sensor, the robot's belief about the window being closed increases. Proceeding in a manner analogous to above, after simplifications, we get:

$$\frac{1}{\gamma} \int_x \sum_y \sum_z \begin{cases} .5 \times .6 \times .75 \times .7 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 1, z = 0, (w = 0)[do(\alpha, t)] \\ .5 \times .4 \times .75 \times .7 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 0, z = 0, (w = 0)[do(\alpha, t)] \\ 0 & \text{otherwise} \end{cases}$$

where $\alpha = [\text{setwin}(0, z), \text{seewin}(0)]$. The numerator simplifies to $.75 \times .7$. It can be shown that γ is:

$$\int_x \sum_y \sum_z \begin{cases} .5 \times .6 \times .75 \times .7 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 1, z = 0, (\text{true})[do(\alpha, t)] \\ .5 \times .4 \times .75 \times .7 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 0, z = 0, (\text{true})[do(\alpha, t)] \\ .5 \times .6 \times .25 \times .3 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 1, z = 1, (\text{true})[do(\alpha, t)] \\ .5 \times .4 \times .25 \times .3 & \text{if } \exists t. h(t) = x, x \in [10, 12], w(t) = y, y = 0, z = 1, (\text{true})[do(\alpha, t)] \\ 0 & \text{otherwise} \end{cases}$$

which leads to $.75 \times .7 + .25 \times .3$. Thus, the belief in the window being closed is $.875$.

7.4. Summary

This completes the specification for reasoning about beliefs with discrete and continuous fluents against noisy effectors and sensors. To summarize, the generalization of the BHL scheme to arbitrary domains is defined using convenient abbreviations for *Bel*, sums and integrals, and where an action theory consists of:

- \mathcal{D}_0 describing the initial state, which also includes (P1) and (P3);
- successor state axioms as before, including a fixed one for p , namely (P2);
- alternate actions axioms for noisy effectors;
- likelihood axioms for noisy sensors and noisy effectors; and
- precondition axioms and foundational axioms as before.

It is perhaps also worth noting that our specification extends BHL in a minimal way. They do not need a fixed space of situations, which we do by (P3), but this is very reasonable [34]. Likelihood axioms are specified for us in much the same manner as they would. Their treatment of noisy actions and sensors is slightly different in expecting the modeler to provide *indistinguishability* axioms. Briefly, these axioms speculate the set of actions that are observationally indistinguishable to the agent. Roughly, then, this serves the same purpose as our alternate actions axioms. However, their approach requires the successor state axiom for p to include notions of indistinguishability. But perhaps most significantly, as we mentioned earlier, their approach needs to appeal to *Golog* to reason about belief change after noisy effectors, which we do not.

To conclude our technical treatment, let us attempt to simplify Definition 26. We will first abuse notation and use a single symbol to denote either sums or integrals, with the understanding that they expand appropriately for a given term t . For the sake of the discussion, let us use $\int_x t$ to mean the integration of the function $t(x)$ from negative to positive infinity when the function takes values from \mathbb{R} , and the sum of terms otherwise. Under this notational convention, let us suppose f_1, \dots, f_n are all the fluents in the language, some of which take values from \mathbb{R} and others take values from finite sets. Suppose a_1, \dots, a_l are all the action types in the language such that for some action types $Alt(a, b, z)$ holds for $z \in \mathbb{R}$, and for the remaining action types, $Alt(a, b, z)$ holds for z taking values from finite sets. Then, we obtain a proposal like Definition 18:

Definition 28. (Degrees of belief (simplified and general).) Suppose ϕ is any \mathcal{L} -formula, and let $\int_x t$ denote integration or summation over term t as appropriate. Then the *degree of belief* in ϕ at s , written $Bel(\phi, s)$, is defined as an abbreviation:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x} \cdot \vec{z}} P(\vec{x}, \vec{z}, \phi, s)$$

where, if $s = do(\alpha, S_0)$ for $\alpha = [a_1, \dots, a_k]$ and suppose $\beta = [b_1, \dots, b_k]$, then

$$P(\vec{x}, \vec{z}, \phi, s) \doteq \langle \iota, b_1, \dots, b_k. \wedge f_i(t) = x_i \wedge \wedge Alt(a_j, b_j, z_j) \wedge \phi[do(\beta, \iota)] \rightarrow p(do(\beta, \iota), do(\alpha, S_0)) \rangle.$$

That is, i ranges over the indices of the fluents, and j over the indices of the ground action terms.

8. Related Work

This article focused on degrees of belief in a first-order dynamic setting. In particular, the existing scheme of BHL was generalized to handle both discrete and continuous probability distributions while retaining all the advantages. Related efforts on belief update via sensor information can be broadly classified into two camps: the literature on probabilistic formalisms, and those that extend logical languages. We discuss them in turn. At the outset, we remark that the focus of our work is on developing a general framework, and not on computational considerations, efficiency or otherwise.

From the perspective of probabilistic modeling, graphical models [28], such as Bayesian networks [7], are important formalisms for dealing with probabilistic uncertainty in general, and the uncertainty that would arise from noisy sensors in particular. Mainly, when random variables are defined by a set of dependencies, the density function can be compactly factorized using these formalisms. The significance of such formalisms is computational, with reasoning methods, such as filtering, being a fundamental component of contemporary robotics and machine learning technologies [6, 38, 13]. On the representation side, however, these formalisms have difficulties handling strict uncertainty, as would arise from connectives such as disjunctions. (Proposals such as credal networks [39] allow for certain types of partial specifications, but still do not offer the generality of arbitrary logical constraints.) 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 model frameworks with an account of actions, such as [40, 41], 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.

This inherent limitation of probabilistic formalisms led to a number of influential proposals on combining logical and probabilistic specifications [42]. (The synthesis of deductive reasoning and the probability calculus has a long and distinguished history [43, 44] that we do not review here; see [45] and references therein.) The works of Bacchus and Halpern [30, 3], for example, provide the means to specify properties about the domain together with probabilities about propositions; see [46] for a recent list on first-order accounts of probability. But these do not explicitly address reasoning about actions. As we noted, treating actions in a general way requires, among other things, addressing the frame problem, reasoning about what happened in the past and projecting the future, handling contextual effects, as well as appropriate semantical machinery. We piggybacked on the powerful situation calculus framework, and extended that theory for reasoning about continuous uncertainty.

In a similar vein, from a modal logical perspective, the interaction between categorical knowledge, on the one hand, and degrees of belief, on the other, is further discussed in [30, 47, 48]. While these are essentially propositional, there are first-order variants [49]. Actions are not explicitly addressed, however.

Recently in AI, limited versions of probabilistic logics have been discussed, in the form of relational graphical models, Markov logic networks, probabilistic databases and probabilistic programming [50, 51, 52, 53, 54, 55, 56, 33, 57]. Some have been further extended for continuous probability distributions and temporal reasoning [58, 59, 60, 61]. Overall, these limit the first-order expressiveness of the language, do not treat actions in a general way, and do not handle strict uncertainty. Admittedly, syntactical restrictions in these frameworks are by design, in the interest of tractability (or at least decidability) wrt inference, as they have origins in the richer probabilistic logical languages mentioned above [30, 3]. From the point of view of a general-purpose representation language, however, they are lacking in the kinds of features that we emphasise here.

From the perspective of dynamical systems, closest in spirit to our work here are knowledge representation languages for reasoning about action and knowledge, which we refer to as action logics. The situation calculus [1, 11], which has been the sole focus of this paper, is one such language. There are others, of course, such as the event calculus [62], dynamic logic [63, 64], the fluent calculus [65], and formalisms based on the stable model semantics [66].

In the situation calculus, a monotonic solution to the frame problem was provided by Reiter [19]. The situation calculus was extended to reason about knowledge whilst incorporating this solution in [12], and to reason about noisy effectors and sensors by BHL.¹⁷ Other action logics have enjoyed similar extensions. For example, [68] proposes an

¹⁷Throughout, we operated under the setting of knowledge expansion, that is, observations are assumed to resolve the agent's uncertainty and never contradict what is believed. The topic of *belief revision* lifts this assumption [67], but it is not considered here. See [25] for an account of belief revision in the situation calculus.

extension to dynamic logic for reasoning about degrees of belief and noisy actions, and [69] provides a computational framework for probabilistic reasoning using the stable model semantics, but they are propositional. In [70, 71], the fluent calculus was extended for probabilistic beliefs and noisy actions. None of these admit continuous probability distributions.

The situation calculus has also been extended for uncertainty modeling in other directions. For example, [72] consider discrete noisy actions over complete knowledge, that is, no degrees of belief. In later work, [73] treat continuous random variables as meta-linguistic functions, and so their semantics is not provided in the language of the situation calculus. This seems sufficient for representing things like products of probabilistic densities, but it is not an epistemic account in a logical or probabilistic sense. A final prominent extension to the situation calculus for uncertainty is the embedding of decision-theoretic planning in *Golog* [74, 75]. Here, actions are allowed to be nondeterministic, but the assumption is that the actual state of the world is fully observable. (It essentially corresponds to a fully observable Markov decision process [76].) In this sense, the picture is a special case of the BHL framework. It is also not developed as a model of belief. While this line of work has been extended to a partially observable setting [77], the latter extension is also not developed as a model of belief. Perhaps most significantly, neither of these support continuous probability distributions, nor strict uncertainty at the level of probabilistic beliefs.

It is worth noting that real-valued fluents in action logics turn out to be useful for modeling resources and time. See, for example, [78, 79, 80]. These are, in a sense, complementary to an account of belief.

Outside the logical literature, there are a variety of formalisms for modeling noisy dynamical systems. Of these, partially observable Markov decision processes (over discrete and continuous random variables) are perhaps the most dominant [13]. They can be seen as belonging to the literature on probabilistic planning languages [81, 82]. Recent probabilistic planning languages [83], moreover, combine continuous Bayesian networks and classical planning languages. Planning languages, generally speaking, only admit a limited set of logical connectives, constrain the language for specifying dynamic laws (that is, they limit the syntax of the successor state axioms), and do not handle strict uncertainty.

In sum, to our knowledge, our proposal is the first of its kind to handle degrees of belief, noisy sensors and effectors over discrete and continuous probability distributions in a general way. The proposal allows for partial and incomplete specifications, and the properties of belief will then follow at a corresponding level of specificity. Moreover, to our knowledge, no other logical formalism for uncertainty deals with the integration of continuous variables within the language.

9. Conclusions

Many real-world applications, such as robotics, have to deal with numerous sources of uncertainty, the main culprit being sensor noise. Probabilistic error models have proven to be essential 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 was limited, however, by its inability to handle continuous distributions, a limitation we lift in this article. By recasting the assessment of belief in terms of fluent values, we now seamlessly combine the situation calculus with discrete probability distributions, densities and difficult combinations of the two. We demonstrated that distributions evolve appropriately after actions, emerging as a side-effect of the general specification. Our formal framework was then shown to easily accommodate the interaction between discrete and continuous fluents, discrete and continuous noise models, and logical connectives. At a specification level, the framework provides the necessary bridge between logic-based action formalisms and probabilistic ones.

Armed with this general specification language, we are in a position to investigate specialized reasoning machinery. To give a few examples, in [35, 84], we identified general projection techniques, where we transform a property of belief after a sequence of (noisy) actions and observations to what is believed initially. In later work [85], we provided an efficient implementation of a projection technique under some reasonable assumptions, in service of enabling richer domain axiomatizations for robotics applications. Finally, a version of *Golog* was recently embedded in our model of belief [26], in the style of knowledge-based programming [10, 86].

A major criticism leveled at much of the work in cognitive robotics [87], and logic-based knowledge representation more generally is that the languages are far removed from the kind of uncertainty and noise seen in machine learning and robotics applications. A formal language such as the one considered in this article addresses this concern. It also shows the advantages of appealing to logical machinery: firstly, in admitting natural, rich and intuitive physical laws; secondly, in allowing belief specifications that can exploit the full power of logical connectives, thereby going considerably beyond standard probabilistic formalisms; and thirdly, in alleviating the burden of determining how these probabilistic beliefs are affected in dynamical systems. In the long term, we hope it takes steps towards a general-purpose epistemologically-adequate representation language as envisioned by McCarthy and Hayes.

References

- [1] J. McCarthy, P. J. Hayes, Some philosophical problems from the standpoint of artificial intelligence, in: *Machine Intelligence*, 1969, pp. 463–502.
- [2] M. Richardson, P. Domingos, Markov logic networks, *Machine learning* 62 (1) (2006) 107–136.
- [3] F. Bacchus, *Representing and Reasoning with Probabilistic Knowledge*, MIT Press, 1990.
- [4] J. Halpern, An analysis of first-order logics of probability, *Artificial Intelligence* 46 (3) (1990) 311–350.
- [5] F. Bacchus, J. Y. Halpern, H. J. Levesque, Reasoning about noisy sensors and effectors in the situation calculus, *Artificial Intelligence* 111 (1–2) (1999) 171 – 208.
- [6] T. Dean, M. Wellman, *Planning and control*, Morgan Kaufmann Publishers Inc., 1991.
- [7] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*, Morgan Kaufmann, 1988.
- [8] S. Kripke, Semantical considerations on modal logic, *Acta Philosophica Fennica* 16 (1963) 83–94.
- [9] J. Hintikka, *Knowledge and belief: an introduction to the logic of the two notions*, Cornell University Press, 1962.
- [10] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, *Reasoning About Knowledge*, MIT Press, 1995.
- [11] R. Reiter, *Knowledge in action: logical foundations for specifying and implementing dynamical systems*, MIT Press, 2001.
- [12] R. B. Scherl, H. J. Levesque, Knowledge, action, and the frame problem, *Artificial Intelligence* 144 (1-2) (2003) 1–39.
- [13] S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics*, MIT Press, 2005.
- [14] Z. Zamani, S. Sanner, P. Poupart, K. Kersting, Symbolic dynamic programming for continuous state and observation POMDPs., in: *NIPS*, 2012, pp. 1403–1411.
- [15] K. Murphy, *Machine learning: a probabilistic perspective*, The MIT Press, 2012.
- [16] V. Belle, H. J. Levesque, Reasoning about continuous uncertainty in the situation calculus, in: *Proc. IJCAI*, 2013.
- [17] H. Enderton, *A mathematical introduction to logic*, Academic press New York, 1972.
- [18] R. Smullyan, *First-order logic*, Dover Publications, 1995.
- [19] R. Reiter, The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression, in: *Artificial intelligence and mathematical theory of computation: Papers in honor of John McCarthy*, Academic Press, 1991, pp. 359–380.
- [20] D. Marker, Model theory and exponentiation, *Notices of the AMS* 43 (7) (1996) 753–759.
- [21] W. F. Trench, *Introduction to real analysis*, Prentice Hall, 2003.
- [22] C. Swartz, *Introduction to gauge integrals*, World Scientific Publishing Company Incorporated, 2001.
- [23] H. J. Keisler, *Elementary calculus: An infinitesimal approach*, Courier Corporation, 2012.
- [24] R. C. Moore, A Formal Theory of Knowledge and Action, in: *Formal Theories of the Commonsense World*, Ablex, Norwood, NJ, 1985, pp. 319–358.
- [25] S. Shapiro, M. Pagnucco, Y. Lespérance, H. J. Levesque, Iterated belief change in the situation calculus, *Artif. Intell.* 175 (1) (2011) 165–192.
- [26] V. Belle, H. J. Levesque, Allegro: Belief-based programming in stochastic dynamical domains, in: *IJCAI*, 2015.
- [27] D. Lewis, Probabilities of conditionals and conditional probabilities, *The philosophical review* (1976) 297–315.
- [28] D. Koller, N. Friedman, *Probabilistic Graphical Models - Principles and Techniques.*, MIT Press, 2009.
- [29] V. Belle, H. J. Levesque, Robot location estimation in the situation calculus, *J. Applied Logic* 13 (4) (2015) 397–413.
- [30] R. Fagin, J. Y. Halpern, Reasoning about knowledge and probability, *J. ACM* 41 (2) (1994) 340–367.
- [31] P. Billingsley, *Probability and Measure*, 3rd Edition, Wiley-Interscience, 1995.
- [32] G. Da Prato, *An Introduction to Infinite-Dimensional Analysis*, Universitext, Springer, 2006.
- [33] P. Singla, P. M. Domingos, Markov logic in infinite domains, in: *UAI*, 2007, pp. 368–375.
- [34] H. J. Levesque, F. Pirri, R. Reiter, Foundations for the situation calculus, *Electron. Trans. Artif. Intell.* 2 (1998) 159–178.
- [35] V. Belle, H. J. Levesque, Reasoning about probabilities in dynamic systems using goal regression, in: *Proc. UAI*, 2013.
- [36] V. Belle, H. J. Levesque, Robot location estimation in the situation calculus, in: *Symposium on Logical Formalizations of Commonsense Reasoning*, 2013.
- [37] J. P. Delgrande, H. J. Levesque, Belief revision with sensing and fallible actions, in: *Proc. KR*, 2012.
- [38] D. Fox, J. Hightower, L. Liao, D. Schulz, G. Borriello, Bayesian filtering for location estimation, *Pervasive Computing*, *IEEE* 2 (3) (2003) 24–33.
- [39] F. G. Cozman, Credal networks, *Artificial Intelligence* 120 (2) (2000) 199 – 233.
- [40] A. Darwiche, M. Goldszmidt, Action networks: A framework for reasoning about actions and change under uncertainty, in: *Proc. UAI*, 1994, pp. 136–144.
- [41] H. Hajishirzi, E. Amir, Reasoning about deterministic actions with probabilistic prior and application to stochastic filtering, in: *Proc. KR*, 2010.
- [42] N. Nilsson, Probabilistic logic, *Artificial intelligence* 28 (1) (1986) 71–87.

- [43] R. Carnap, Logical foundations of probability, Routledge and Kegan Paul London, 1951.
- [44] H. Gaifman, Concerning measures in first order calculi, *Israel Journal of Mathematics* 2 (1) (1964) 1–18.
- [45] J. Y. Halpern, Reasoning about Uncertainty, MIT Press, 2003.
- [46] Z. Ognjanovic, M. Raškovic, Some first-order probability logics, *Theoretical Computer Science* 247 (1–2) (2000) 191 – 212.
- [47] D. Monderer, D. Samet, Approximating common knowledge with common beliefs, *Games and Economic Behavior* 1 (2) (1989) 170–190.
- [48] A. Heifetz, P. Mongin, Probability logic for type spaces, *Games and Economic Behavior* 35 (1-2) (2001) 31 – 53.
- [49] V. Belle, G. Lakemeyer, H. J. Levesque, A first-order logic of probability and only knowing in unbounded domains, in: Proc. AAAI, 2016.
- [50] D. Koller, A. Pfeffer, Probabilistic frame-based systems, in: AAAI/IAAI, 1998, pp. 580–587.
- [51] L. Getoor, N. Friedman, D. Koller, B. Taskar, Learning probabilistic models of relational structure, in: ICML, 2001, pp. 170–177.
- [52] S. J. Russell, Unifying logic and probability, *Commun. ACM* 58 (7) (2015) 88–97.
- [53] D. Poole, First-order probabilistic inference, in: Proc. IJCAI, 2003, pp. 985–991.
- [54] P. Domingos, W. A. Webb, A tractable first-order probabilistic logic, in: Proc. AAAI, 2012.
- [55] L. D. Raedt, A. Kimmig, H. Toivonen, Problog: A probabilistic prolog and its application in link discovery, in: Proc. IJCAI, 2007, pp. 2462–2467.
- [56] D. Suciu, D. Olteanu, C. Ré, C. Koch, Probabilistic databases, *Synthesis Lectures on Data Management* 3 (2) (2011) 1–180.
- [57] B. Milch, B. Marthi, S. J. Russell, D. Sontag, D. L. Ong, A. Kolobov, BLOG: Probabilistic models with unknown objects, in: Proc. IJCAI, 2005, pp. 1352–1359.
- [58] J. Choi, A. Guzman-Rivera, E. Amir, Lifted relational kalman filtering, in: Proc. IJCAI, 2011, pp. 2092–2099.
- [59] D. Nitti, T. D. Laet, L. D. Raedt, A particle filter for hybrid relational domains, in: IROS, 2013, pp. 2764–2771.
- [60] C. Anderson, P. Domingos, D. Weld, Relational markov models and their application to adaptive web navigation, in: Proc. SIGKDD, ACM, 2002, pp. 143–152.
- [61] V. Belle, A. Passerini, G. Van den Broeck, Probabilistic inference in hybrid domains by weighted model integration, in: IJCAI, 2015.
- [62] R. Kowalski, M. Sergot, A logic-based calculus of events, *New Generation Computing* 4 (1986) 67–95.
- [63] H. van Ditmarsch, W. van der Hoek, B. Kooi, *Dynamic Epistemic Logic*, 1st Edition, Springer Publishing Company, Incorporated, 2007.
- [64] H. P. van Ditmarsch, A. Herzig, T. D. Lima, From situation calculus to dynamic epistemic logic, *J. Log. Comput.* 21 (2) (2011) 179–204.
- [65] M. Thielscher, Reasoning robots: the art and science of programming robotic agents, *Applied logic series*, Springer, Dordrecht, 2005.
- [66] M. Gelfond, V. Lifschitz, Representing action and change by logic programs, *The Journal of Logic Programming* 17 (2-4) (1993) 301–321.
- [67] C. E. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, *Journal of Symbolic Logic* 50 (1985) 510–530.
- [68] B. Kooi, Probabilistic dynamic epistemic logic, *Journal of Logic, Language and Information* 12 (4) (2003) 381–408.
- [69] C. Baral, M. Gelfond, J. N. Rushton, Probabilistic reasoning with answer sets, *TPLP* 9 (1) (2009) 57–144.
- [70] M. Thielscher, Planning with noisy actions (preliminary report), in: Proc. Australian Joint Conference on Artificial Intelligence, 2001, pp. 27–45.
- [71] Y. Martin, M. Thielscher, Integrating reasoning about actions and Bayesian networks, in: International Conference on Agents and Artificial Intelligence, Valencia, Spain, 2009.
- [72] P. Mateus, A. Pacheco, J. Pinto, A. Sernadas, C. Sernadas, Probabilistic situation calculus, *Annals of Math. and Artif. Intell.* 32 (1-4) (2001) 393–431.
- [73] C. Fritz, S. A. McIlraith, Computing robust plans in continuous domains, in: Proc. ICAPS, 2009, pp. 346–349.
- [74] C. Boutilier, R. Reiter, M. Soutchanski, S. Thrun, Decision-theoretic, high-level agent programming in the situation calculus, in: Proc. AAAI, 2000, pp. 355–362.
- [75] C. Boutilier, R. Reiter, B. Price, Symbolic dynamic programming for first-order MDPs, in: Proc. IJCAI, 2001, pp. 690–697.
- [76] C. Boutilier, T. Dean, S. Hanks, Decision-theoretic planning: Structural assumptions and computational leverage, *Journal of Artificial Intelligence Research* 11 (1) (1999) 94.
- [77] S. Sanner, K. Kersting, Symbolic dynamic programming for first-order pomdps, in: Proc. AAAI, 2010, pp. 1140–1146.
- [78] H. Grosskreutz, G. Lakemeyer, ccgolog – a logical language dealing with continuous change, *Logic Journal of the IGPL* 11 (2) (2003) 179–221.
- [79] C. S. Herrmann, M. Thielscher, Reasoning about continuous processes, in: AAAI/IAAI, Vol. 1, 1996, pp. 639–644.
- [80] M. Fox, D. Long, Modelling mixed discrete-continuous domains for planning, *J. Artif. Intell. Res. (JAIR)* 27 (2006) 235–297.
- [81] N. Kushmerick, S. Hanks, D. Weld, An algorithm for probabilistic planning, *Artificial Intelligence* 76 (1) (1995) 239–286.
- [82] H. Younes, M. Littman, PPDDL 1. 0: An extension to pddl for expressing planning domains with probabilistic effects, Tech. rep., Carnegie Mellon University (2004).
- [83] S. Sanner, Relational dynamic influence diagram language (rddl): Language description, Tech. rep., Australian National University (2011).
- [84] V. Belle, H. J. Levesque, How to progress beliefs in continuous domains, in: Proc. KR, 2014.
- [85] V. Belle, H. J. Levesque, PREGO: An Action Language for Belief-Based Cognitive Robotics in Continuous Domains, in: Proc. AAAI, 2014.
- [86] R. Reiter, On knowledge-based programming with sensing in the situation calculus, *ACM Trans. Comput. Log.* 2 (4) (2001) 433–457.
- [87] G. Lakemeyer, H. J. Levesque, Cognitive robotics, in: *Handbook of Knowledge Representation*, Elsevier, 2007, pp. 869–886.