Actions, Continuous Distributions and Meta-Beliefs

Vaishak Belle
University of Edinburgh & Alan Turing Institute
United Kingdom
vbelle@ed.ac.uk

ABSTRACT

In this work, we propose a new modal logical language for reasoning about noisy actions and sensors in an epistemic setting. In the reasoning about actions literature, there are only a few frameworks for modelling probabilistic noise, and even less in dealing with continuous probability distributions. In the first model of its kind, we show how a rich theory of actions with beliefs, meta-beliefs and only knowing can be defined over discrete, continuous and mixed discrete-continuous distributions.

KEYWORDS

Logic; Probability; Situation Calculus; Modal Logic

1 INTRODUCTION

The unification of the logic and probability has been seen as a long-standing concern in philosophy and mathematical logic, going back to Carnap [10] and Gaifman [14], at least in terms of early rigorous algebraic studies. In artificial intelligence, starting from Nilsson [31], Bacchus [1] and Halpern [16], a wide range of formalisms encompassing various first-order logical features have been proposed. In so much as a probabilistic underpinning provides the gateway for incorporating probabilistic induction, areas such as statistical relational learning [33] and neuro-symbolic AI [17], are promising candidates for unifying deduction, noisy sensory observations and association-based pattern learning [3, 12, 28, 29].

From a knowledge representation viewpoint, however, especially in the context of reasoning about first-order knowledge over a rich theory of action, probabilistic noise is not extensively studied. Consider, for example, that Reiter's [34] reconsideration of the situation calculus [27] has proven enormously useful for the design of logical agents, essentially paving the way for cognitive robotics [22]. Among other things, it incorporates a simple monotonic solution to the frame problem, leading Reiter to define the notion of regression for basic action theories [43]. The situation calculus, and its counterparts, such as dynamic epistemic logic [42] and the fluent calculus [38], have enjoyed numerous extensions for time, processes, concurrency, exogeneity, reactivity, sensing and knowledge [34].

One criticism leveled at this line of work, and indeed at much of the work in cognitive robotics and reasoning about actions, is that the theory is far removed from the kind of probabilistic uncertainty and noise seen in typical robotic applications [40]. Fortunately, Bacchus, Halpern and Levesque (BHL henceforth) [2] provided a very general account for incorporating degrees of belief, noisy actions and sensors in the situation calculus. It is a surprisingly simple extension to the epistemic situation calculus [36]: instead of a categorical knowledge operator that says whether a formula ϕ is known or not, it permits expressions that quantitatively assesses by how much ϕ is believed. Instead of sensing actions that discard possible worlds based on what was observed, the weights of worlds are adjusted based on how close their fluent values are to the observed value. Most importantly, the main advantage of this logical account 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 [40]. 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 [32]. 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. This is in line with a criticism by McCarthy and Hayes [27] that any formalism forcing us to put numbers on formulas is representationally inadequate.

The simplicity of the BHL model has led to two major classes of extensions. Owing to its limitation to discrete distributions, and the lack of a solution for the projection problem, recent results have demonstrated how it can be extended to continuous distributions [5, 8], as well how a notion of regression and progression [34] can be defined for both the discrete and the continuous model [6, 7].

But perhaps the more major extension is owing to the fact that it is defined axiomatically, as is usual in the situation calculus [34]. Even in the non-probabilistic case of knowledge and actions, an axiomatic definition makes semantic proofs about modalities deeply challenging [20]. The situation is far worse with probabilities: degrees of belief in BHL is defined by summing the weights of situations, but these weights themselves are provided by a successor state axiom that stitches together action executability, unobservable outcomes and likelihoods of actions in one formula. This makes it difficult to unpack, even informally, how degrees of beliefs change over actions and sensing. This motivated a new logical language, the logic $\mathcal{DS}$ [4], which casts the BHL framework in a modal language, allowing a semantical apparatus to reason about actions, beliefs, meta-beliefs (including introspection) and only knowing

[23] in a single logical framework. Extensions to $\mathcal{DS}$ [25, 26] further considered adapting the regression and progression results from the BHL model.

In this work, we argue that such a recasting is more pertinent in the continuous setting. Continuous probabilistic models are widely used in machine learning and robotics [30], and permitting the representation and reasoning of such models within logics of actions would further bridge the knowledge representation and machine learning communities. In the existing situation calculus variant for continuous probabilities (BL henceforth) [8], the tangled account with the weights of situations gets even more severe. While the successor state axiom for a situation's weight (now interpreted as a probabilistic density) incorporates action executability and likelihoods, a definition of belief is provided that accounts for unobservable outcomes, and then integrates over the initial fluent values and the possible outcomes. Owing to the axiomatic foundation, (Lebesgue) integrals need also to be defined axiomatically [8], which not only makes it difficult to study meta-beliefs and their properties, but also to generalize to other accounts of probabilistic measures.

What we attempt here, in other words, is something between $\mathcal{DS}$ and the BL model. We do this by first noting that in the BL model, because the space of situations is not a well-defined space for calculating integrals, they define belief in terms of the values of fluents initially, and then "projecting" how the weights of these situations change after actions. We recast this in a clean modal setting, allowing for a clean specification of how action likelihoods affect the densities of worlds. Integration is used meta-linguistically, and beliefs are allowed to be nested. Overall, our work distinguishes itself from BL in presenting a simple semantic framework for meta-beliefs and only-knowing. On the downside, we will inherit BL's limitation of allowing only finitely many fluents (albeit whose values can come from uncountably infinite sets).

2 THE LOGIC $\mathcal{XS}$

We now define the logic $\mathcal{XS}$ (= *continuous random variables $\vec{X}$ in the situation calculus*). The language is built so as to reason about probabilistic beliefs and meta-beliefs over actions. Recall that we allow fluents to take values from countably infinite and uncountable sets. Likewise, the arguments of sensors and effectors too can values from infinite sets. For simplicity, we will therefore assume the domain of discourse to be consisting of a countably infinite set of *standard names* (the set of objects) together with the set of real numbers $\mathbb{R}$. Quantifiers of the object sort will be understood substitutionally [24]. But we will use variable maps for the reals. (The set of computable reals [41], which includes irrational numbers such as π is assumed to be included in the set of objects so that they can be used in the language.)

2.1 Syntax

Formally, the non-modal fragment of $\mathcal{XS}$ consists of standard first-order logic with = (that is, connectives $\{\wedge, \forall, \neg\}$, syntactic abbreviations $\{\exists, \equiv, \supset\}$) but limited to functions. For simplicity, no predicates are considered. In particular, assume:

- an infinite supply of variables $\{x, y, \ldots, u, v, \ldots\}$;

- rigid function symbols of every arity ≥ 1, *move(x,y)*, *sin(x)* and other arithmetic functions such as + and ×;
- finitely many nullary fluent functions $\leq k$, such as *salary*, *height* and *distance*, often simply denoted $f_1, \ldots, f_k$;
- the following special symbols:
  - a unary functional fluent *poss* to denote the executability of an action;[1] and
  - a unary fluent l that takes an action as its argument and gives the action's likelihood.

Terms are the least set of expressions such that:

- every variable is a term;
- if $t_1, \ldots, t_k$ are terms and f is k-ary function symbol, then $f(t_1, \ldots, t_k)$ is a term.

Well-formed (static) formulas are constructed as usual in first-order logic with equality. They can further be in the context of belief and action modalities.

$\mathcal{XS}$ has two epistemic operators: $\boldsymbol{B}(\alpha : x)$ is to be read as "α is believed with a probability x," where x is a term of the number sort. Next, the modality $\boldsymbol{O}(\alpha_1 : x_1, \ldots, \alpha_k : x_k)$, where x_i is a term of the number sort, is to be read as "all that is believed is: α_1 with probability $x_1, \ldots,$ and α_k with probability x_k." We also use $\boldsymbol{K}\alpha$, to be read as "α is known," as an abbreviation for $\boldsymbol{B}(\alpha : 1)$. We write $\boldsymbol{O}\alpha$, to be read as "α is all that is known," to mean $\boldsymbol{O}(\alpha : 1)$.

$\mathcal{XS}$ has two action modalities $[a]$ and $\square$, in that if α is a formula, then so are $[a]\alpha$ (read: "α holds after a") and $\square\alpha$ (read:"α holds after any sequence of actions.") For $z = a_1 \cdots a_k$, we write $[z]\alpha$ to mean $[a_1]\cdots[a_k]\alpha$. We use *true* to denote truth, which is taken as an abbreviation for, say, $\forall x(x = x)$, and *false* for its negation.

2.2 Semantics

The semantics is given in terms of *possible worlds*. In a dynamic setting, such worlds are defined to interpret not only the current state of affairs, but also how that changes over actions. There are three key complications over non-probabilistic accounts with deterministic acting and sensing [20]:

- we need to be able to specify probabilities over *uncountably many possible* worlds in a well-defined manner;
- to allow for qualitative uncertainty in an inherently quantitative account, beliefs *may not be characterizable* in terms of a single distribution;
- the effects of actions are nondeterministic, and the changes to the state of affairs thereof are (possibly) *not observable* by the agent.

2.2.1 Defining worlds. To begin with, let us assume the elements of $\mathcal{A}$ are of the form $a(c)$, where a is an action symbol, and c is a standard name or a number. Let $\mathcal{Z}$ be all finite sequences of $\mathcal{A}$, including $\langle\rangle$, the empty sequence. Then, a world w is a mapping:

- for every (nullary) fluent f and $z \in \mathcal{Z}$, $w[f, z] \in \mathbb{R}$ (and analogously for fluents of other arities);
- for every k-ary rigid function g, the world maps $\mathbb{R}^k$ to $\mathbb{R}$ that is same for every $z \in \mathcal{Z}$; and

---

[1]This fluent function captures the executability of actions, replacing the usual fluent predicate in the situation calculus [34], and so we require also that for every a, $poss(a) = 0 \vee poss(a) = 1$.

- arithmetic functions such as $sin(x)$, $exp(x)$, and + are interpreted in the usual sense. (This is as in [8].)

Let $\mathcal{W}$ be the set of all such mappings: the set of possible worlds.

2.2.2 Initial distribution. We are now ready to define distributions and epistemic states, like in $\mathcal{DS}$, but for continuous probability spaces. Let us consider a general notion first, not dissimilar to $\mathcal{DS}$, over which we will place further stipulations so as to be integrable. By a *distribution d* we mean a mapping from $\mathcal{W}$ to $\mathbb{R}^{\geq 0}$ (the set of non-negative reals) and an *epistemic state e* is any set of distributions. The idea in $\mathcal{DS}$ is to constrain a d such that it defines a distribution on the worlds in $\mathcal{W}$. Then B is interpreted wrt every $d \in e$.

Let us now consider a d and investigate how it can be used to define a continuous distribution. Moreover, let us focus on the empty sequence $\langle\rangle$, that is, before any actions have occurred.

While there are many notions of continuity in probability theory [9], perhaps the simplest is consider an absolutely continuous distribution $\eta_X(x)$ for a random variable X such that:

$$\Pr[a \leq X \leq b] = \int_a^b \eta_X(x).$$

That is, the probability over an interval is obtained by the Lebesgue integral. As we shall see, the simplicity of our construction means that we will not need to limit ourselves to just integrals, sums will also do, as perhaps would other types of measures [37].

Let us start by observing that given $\mathbb{R}^k$, its Borel sets can be defined as measurable subsets (intervals, Cartesian products of intervals, and so on), leading to a general notion of a probability measure [9, 15]. So given a function $\eta_{X_1,\ldots,X_k}$ that maps $\mathbb{R}^k$ to $\mathbb{R}$, standing for the probability density or mass function, well-understood notion of probability apply in the sense of being able to obtain a well-defined measure on Borel sets of $\mathbb{R}^k$, either finitely many, countably infinitely many or uncountably many. We would obtain a discrete probability distribution with the first two, and a continuous probability distribution with the last.

But our notion of a distribution d is precisely such a function. It maps worlds to reals, where the worlds themselves are elements of $\mathbb{R}^k$. In other words, if d is standing for a probability mass function, there is a number n such that:

$$n = \sum_{\vec{r} \in R^k} \begin{cases} d(w) & \text{if } w[f_1, \langle\rangle] = r_1, \ldots w[f_k, \langle\rangle] = r_k \\ 0 & \text{otherwise} \end{cases}$$

where $R^k \subsetneq \mathbb{R}^k$ is the discrete space corresponding to the values of the fluents. In English: suppose fluents $f_1, \ldots, f_k$ took values from a countably finite set $\left\{ \langle r_1^1, \ldots, r_k^1 \rangle, \ldots, \langle r_1^n, \ldots, r_k^n \rangle \right\}$. There is exactly one world corresponding to each such $\langle r_1^i, \ldots, r_k^i \rangle$. When considering the summation, for every $\vec{r} = \langle r_1^i, \ldots, r_k^i \rangle \in R^k$, we use the d value of the world where the fluents take these values initially, and 0 otherwise. Notice that this is before any actions have happened, because after actions multiple worlds might agree on the values of fluents. (Think, for example, of k boxes and a functional fluent for each box to capture its color: after an action that colors all boxes red, all those fluents would take on the same value.)

Analogously, if d stands for a density function that maps $\mathbb{R}^k$ to $\mathbb{R}$, there is a number n such that:

$$n = \int_{\vec{r} \in \mathbb{R}^k} \left(\begin{cases} d(w) & \text{if } w[f_1, \langle\rangle] = r_1, \ldots w[f_k, \langle\rangle] = r_k \\ 0 & \text{otherwise} \end{cases} \right)$$

This now gives us a generic recipe for dealing with absolutely continuous, discrete, countably infinite and mixed discrete-continuous fluents. We let d map the Cartesian product of the appropriate space of values to $\mathbb{R}$ to capture the probabilistic measure, and sum or integrate as appropriate for that fluent, exactly as one would in standard probability theory [9]. For simplicity of presentation, however, we will use the integration symbol and assume absolutely continuous distributions in the rest of the paper.

All of the above discussion only pertains to the empty sequence, of course, because actions could affect the nature of the distributions. But inspired by [8], we will find a way to define beliefs after actions by recasting it to initial beliefs.

2.2.3 Noisy actions. When a noise-free physical action occurs, it is clear to us (as modelers) but also the agent how the world has changed after the action. Of course, in realistic domains, especially robotic applications, this is not the usual case and a quantitative account of effector noise is needed.

Let us first reflect on what is expected with noisy sensing vs noisy acting. When an agent senses, say a sonar action such as $sonar(z)$, the argument for this action is not chosen by the agent. That is, the world determines what z should be, and based on this reading of z, the agent comes to 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. This is different from noise-free situation calculus [36], where sensing reveals the value of the fluent in the real world.

Noisy acting diverges from that picture even further. The agent intends to do action a, but what actually occurs is a' that is possibly different from a. For example, an agent may want to move 3 units towards a wall, 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 sequences might be possible, which are all indistinguishable in the agent's perspective. Depending on the likelihoods of the action's potential outcomes, some of these successors are considered more probable than others. The agent's belief about what holds then must incorporate these relative likelihoods. Intuitively, if the belief about the position is a bell curve, a noisy action may cause it "flatten" and move the mean (because the robot has possibly moved), and a noisy sensing would cause it "sharpen" slightly, and so on. This is what necessitates a sense-act loop in stochastic domains.

Like with $\mathcal{DS}$, we will need some notational devices for dealing with (noisy) actions. However, there are some significant differences, because in $\mathcal{DS}$, we could simply gather the indistinguishable worlds after actions. However, that would not be useful here because it may not be feasible to provide a compact characterization

of the probability function. For example, an action can transform an absolutely continuous function with a density to a single point with a probability mass. Nonetheless, a characterization of the prior distributions for the initial state is provided by assumption, and we will use it to define probabilities on successor situations, like in [8].

So let us introduce some notation for characterizing the likelihood of a sequence of noisy actions and their executability. We will need to do this precisely because the arguments of the actions are what we will eventually integrate over. So assume noisy actions [8] are of the form $a(x, y)$, where x is the intended argument and y is the actual value of the executed action, and noisy sensing is of the form $a(z)$, where z is the observed value determined by the environment.[2]

First, we extend the application of l to sequences:

Definition 2.1. We define $L : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}^{\geq 0}$ as follows:

- $L(w, \langle \rangle) = 1$ for every $w \in \mathcal{W}$;
- $L(w, z \cdot r) = L(w, z) \times n$ where $w[l(r), z] = n$.

Noise-free physical actions will be axiomatized in a way such that their l-values are 1. For a noisy sensing action $sonar(3)$, its l-value will be determined by how likely it is to observe 3 given that the robot is, say, 4 units away from the wall. For example, if assuming a Gaussian error profile [40] with a standard deviation of 2, the l-value would be $\mathcal{N}(3; 4, 2)$, which is to say the number on the Gaussian curve corresponding to the number 3 on the X-axis. The curve is centered on the number 4 with a spread given by the standard deviation of 2. Analogously, if assuming a Gaussian error profile with a standard deviation of 2, the l-value for the action $move(3, 4)$ would be $\mathcal{N}(4; 3, 2)$, which is to say the number on the Gaussian curve corresponding to the number 4 on the X-axis. We read the action as saying 3 was intended but 4 actually occurred. The curve is centered on the number 3 with a spread given by the standard deviation of 2.

After intending to execute a sequence of actions, the agent needs to also consider those sequences that may be the actual outcomes. For this, we define *action sequence observational indistinguishability* as follows:

Definition 2.2. We define $z \approx z'$:

- $\langle \rangle \approx z'$ iff $z' = \langle \rangle$;
- $z \cdot r \approx z'$ iff $z' = z^* \cdot r$, $z \approx z^*$ and r is a noise-free physical action or a noisy sensor; and
- $z \cdot a(c, c') \approx z'$ iff $z' = z^* \cdot a(c, c'')$ for some c'', and $z \approx z^*$.

Let us unpack this definition. Empty sequences have no indistinguishable alternatives. Noise-free actions have no indistinguishable alternatives, and so if $z \approx z^*$ and r is such an action, $z \cdot r \approx z^* \cdot r$. One might be tempted to lump noisy sensing and noisy actions, but note that with an action like $sonar(3)$, even though the argument is not in control of the agent and is determined by the environment, the reading is *observable* to the agent after performing the action. Thus, noisy sensing actions also have no indistinguishable alternatives.

But for a noisy action $a(c, c')$, any $a(c, c'')$ would mean $z \cdot a(c, c') \approx z \cdot a(c, c'')$.[3] In general, then, given $z \approx z'$, we have $z \cdot a(c, c') \approx z' \cdot a(c, c'')$. By construction, it is easy to see that:

PROPOSITION 2.3. $\approx$ *is an equivalence relation.*

As a matter of notational convenience, given any action sequence z, having n noisy actions, we use $z(c_1, \ldots, c_n)$ to mean that the actual arguments of those actions in z were replaced in corresponding places by $c_1, \ldots, c_n$. For example, given $z = sonar(3) \cdot move(3, 4) \cdot sonar(4) \cdot move(2, 2)$, we let $z(1, 1)$ mean $sonar(3) \cdot move(3, 1) \cdot sonar(4) \cdot move(2, 1)$ and by above, $z \approx z(1, 1)$.

Third, to extend the applicability of *poss* for action sequences, we proceed as follows:

Definition 2.4. Define $Exec(z)$ for any $z \in \mathcal{Z}$ inductively:

- for $z = \langle \rangle$, $Exec(z)$ denotes *true*;
- for $z = a \cdot z'$, $Exec(z)$ denotes $(poss(a) = 1) \wedge [a]Exec(z')$.

2.2.4 Truth. We are finally ready for the semantics. Let us introduce variable maps: such a map v maps real-number variables to $\mathbb{R}$. We write $v' \sim_x v$ to mean v' and v agree on everything except the assignment for variable x.

The denotation of terms (wrt a world w, action sequence z and map v) is defined inductively. If t is a name, then $|t|_{(w, z, v)} = t$. If t is a term of the number sort, then $|t|_{(w, z, v)} = v(t)$. If $t = f(t_1, \ldots, t_m)$, then $|t|_{(w, z, v)} = w[f(r_1, \ldots, r_m), z]$ where $|t_i|_{(w, z, v)} = r_i$.

Truth in $\mathcal{XS}$ is defined wrt (e, w, z, v) as follows:

(1) $e, w, z, v \models (f = t)$ iff $w[f, z] = |t|_{(w, z, v)}$;
(2) $e, w, z, v \models (t_1 = t_2)$ iff $|t_1|_{(w, z, v)}$ and $|t_2|_{(w, z, v)}$ are identical;
(3) $e, w, z, v \models \alpha \wedge \beta$ iff $e, w, z, v \models \alpha$ and $e, w, z, v \models \beta$;
(4) $e, w, z, v \models \neg \alpha$ iff $e, w, z, v \not\models \alpha$;
(5) $e, w, z, v \models \forall x \alpha$ iff $e, w, z, v \models \alpha_n^x$ for all names n, where x is a variable of the object sort;
(6) $e, w, z, v \models \forall x \alpha$ iff $e, w, z, v' \models \alpha$ for all $v' \sim_x v$, where x is a variable of the number sort;
(7) $e, w, z, v \models [t]\alpha$ iff $e, w, z \cdot n, v \models \alpha$ where $n = |t|_{(w, z, v)}$;
(8) $e, w, z, v \models \Box \alpha$ iff $e, w, z \cdot z', v \models \alpha$ for all $z' \in \mathcal{Z}$.

Hereafter, to simplify the presentation, we often write $e, w, z \models \alpha$ to mean $e, w, z, v \models \alpha$ for all maps v.

2.2.5 Epistemic operators. For the epistemic operators, we need to bridge the intuition of defining a probability using d for the empty sequence with actions. This is made all the more complicated by the fact that actions are noisy. To work this out, let us start with what we understand so far: consider a semantics for the degree of belief of a modality-free formula ϕ wrt $\langle \rangle$ and a singleton $e = \{d\}$. Based on what was discussed before, we might arrive at:

$$e, w, \langle \rangle, v \models B(\phi : n) \text{ iff } n = \int_{\mathbb{R}^k} \begin{cases} d(w^*) & \text{if } \psi \\ 0 & \text{otherwise} \end{cases}$$

---

[2]Actions with more arguments are permitted, we are restricting it this way for presentation purposes.

[3] Note that in $\mathcal{DS}$, there was a distinguished predicate for capturing the indistinguishability of actions, which could also be axiomatized in the basic action theory. Our approach is less general, but simpler: we assume that for every noisy action $a(x, y)$, all and only the actions $a(x, z)$ are alternatives. (We expect the likelihood axiom to determine which of these are more likely than the other.) Ultimately, because we define integration in terms of the parameters of the actions, this restriction is unavoidable, so we drop the use of this distinguished predicate.

What might ψ be? We understand from before that we choose worlds corresponding to every vector of values from $\mathbb{R}^k$, and so to obtain the probability of ϕ, we would need to only consider those tuples where it holds. Formally, ψ is

$$w^*[f_1, \langle \rangle] = r_1, \ldots, w^*[f_k, \langle \rangle] = r_k, (e, w^*, \langle \rangle, v \models \phi).$$

That is, for every vector of values from $\mathbb{R}^k$, consider the w^* where the fluents take these values initially, and test whether ϕ is true at the model with w^* as the real world. If it is, use the d-value, otherwise ignore.

But what is the relation between d and e? An epistemic state is a set of distributions, so for arbitrary epistemic states, we would require that for every $d \in e$:

$$n = \int_{\mathbb{R}^k} \begin{cases} d(w^*) & \text{if } \psi \\ 0 & \text{otherwise} \end{cases}$$

Notice that the epistemic state in the ψ-condition is e and not some $d \in e$, and this allows for introspection.

To now handle actions, let us reflect on how belief changes. With noise-free actions, the values of fluents change at worlds, so all we would need to do is:

$$e, w, a, v \models \boldsymbol{B}(\phi : n) \text{ iff for all } d \in e, \left(\int_{\mathbb{R}^k} \begin{cases} d(w^*) & \text{if } \psi \\ 0 & \text{otherwise} \end{cases} \right) = n$$

where ψ is

$$w^*[f_i, \langle \rangle] = r_i \text{ for every } i, \ (e, w^*, a, v \models \phi).$$

That is, we integrate over $\mathbb{R}^k$, and for every vector of real values, if the world corresponding to those values satisfies ϕ after a, then we use its d-value, else we ignore it. This can be extended to any sequence of noise-free actions.

With noisy sensing, the account is very much the same except that we need to adjust the weight of the world based on the value read on the sensor.[4] That is,

$$e, w, a, v \models \boldsymbol{B}(\phi : n) \text{ iff for all } d \in e, \left(\int_{\mathbb{R}^k} \begin{cases} d(w^*) \times \theta & \text{if } \psi \\ 0 & \text{o.w.} \end{cases} \right) = n$$

where o.w. = otherwise, and ψ is

$$w^*[f_i, \langle \rangle] = r_i \text{ for every } i, \ w^*[l(a), \langle \rangle] = \theta, \ (e, w^*, a, v \models \phi).$$

Recall that we have L precisely for obtaining the likelihood of a sequence, so we may use that instead of referring to θ.

We now consider the case of noisy actions. The likelihood of an always continuous noisy action is given by a density function. Intuitively, when an action with argument x is intended, uncountably many y are possible and we would need to integrate over these. The density of a world is now further adjusted for each such outcome. This means that we need to introduce a new integration symbol

for every noisy action. Formally, given a noisy action $a(c, c')$, we have:[5]

$$e, w, a(c, c'), v \models \boldsymbol{B}(\phi : n) \text{ iff for every } d \in e, n = \eta$$

where η is:

$$\int_{\mathbb{R}} \int_{\mathbb{R}^k} \begin{cases} d(w^*) \times \mathrm{L}(w^*, a(c, c'')) & \text{if } \psi \\ 0 & \text{o.w.} \end{cases}$$

where ψ is

$$w^*[f_i, \langle \rangle] = r_i \text{ for every } i, \ (e, w^*, a(c, c''), v \models \phi)$$

In English, given a real number c'' for the actual outcome, and real numbers $r_1, \ldots, r_k$ for the fluent values, we first identify the world w^* where the fluents take on the latter values. Because $a(c, c'')$ is indistinguishable from $a(c, c')$, if ϕ is true after $a(c, c'')$, then we use the d-value of the world. Its weight is further adjusted using the likelihood of this alternate action. So for every world, we consider every possible outcome, and depending on whether ϕ is true, we use the density term (product of the initial density term and the likelihood). This now allows us to provide the general definition of belief:

(8) $e, w, z, v \models \boldsymbol{B}(\phi : n)$ iff for every $d \in e, n = \eta$

where η is:

$$\int_{\mathbb{R}^{k+m}} \begin{cases} d(w^*) \times \mathrm{L}(w^*, z') & \text{if } \psi \\ 0 & \text{otherwise} \end{cases}$$

where ψ is:

$$w^*[f_i, \langle \rangle] = r_i \text{ for every } i, \ z' = z(c_1, \ldots, c_m), (e, w^*, z', v) \models \phi.$$

In English: for k real numbers, we identify the world w^* where the fluents take these values. Given m noisy actions in z, for the m real numbers $\langle c_1, \ldots, c_m \rangle$, we consider $z' = z(c_1, \ldots, c_m)$. Because z' is indistinguishable from z, provided ϕ is true after z' at w^* means that we can use its d-value. It is adjusted as per the likelihood of z'.

The semantics for only knowing then is:

(9) $e, w, z, v \models \boldsymbol{O}(\phi_1 : n_1, \ldots, \phi_k : n_k)$ iff for all $d, d \in e$ iff $n_i = \eta_i$

where η_i is the same expression as η above except in using ϕ_i.

For any sentence α, we write $e, w \models \alpha$ instead of $e, w, \langle \rangle \models \alpha$. When Σ is a set of sentences and α is a sentence, we write $\Sigma \models \alpha$ (read: "Σ logically entails α") to mean that for every e and w, if $e, w \models \alpha'$ for every $\alpha' \in \Sigma$, then $e, w \models \alpha$. Finally, we write $\models \alpha$ (read: "α is valid") to mean $\{\} \models \alpha$.

We introduce some syntactic sugar for $\boldsymbol{B}$. Our modal operator for belief is of the form $\boldsymbol{B}(\phi : x)$, but we often write $\boldsymbol{B}\phi = x$. This

---

[4]With continuous distributions, a number of involved issues may arise in the context of conditioning. (Conditioning corresponds to an observation so can be captured using sensing actions.) Given a continuous fluent, a sensing action might suggest a single point or a discrete set of points to be likely and all others unlikely, for example. It is also possible to condition on null sets, leading to the Borel–Kolmogorov paradox [9, 15]. We will not have much to say about such concerns beyond what is already discussed in the literature, and simply assume that the d-value of worlds as well as its adjustment after actions always integrates to 1.

[5]Related to the conditioning remark, there is an important consideration on measurability [9] more generally that we are glossing over in our treatment. Our set of worlds is defined over $\mathbb{R}^k$, and we simply assumed that the initial beliefs are (Lebesgue) integrable, as is the product of the initial density with the likelihood of actions. These assumptions are perhaps acceptable in the context of most robotics applications, where we expect initial beliefs and likelihood axioms to involve Gaussian, uniform and other classical distributions. But there are examples of sets that are not Lebesgue integrable (such as the Vitali set [19]): this would mean formulas expressing such sets would not be satisfiable in our logic. This is an issue also with prior accounts formulated using Lebesgue integrals, such as [8]. We ignore this issue in its entirety for now. A general account would have to address the problem of defining integrals over arbitrary probability spaces, both initially and after actions. Finally, notice that by assuming beliefs are always integrable, we have that $\boldsymbol{K}$ *true* is valid, as is $\Box \boldsymbol{K}$ *true*. (Recall that we use the abbreviation $\boldsymbol{K}\alpha \doteq \boldsymbol{B}\alpha = 1$.)

is extended for arithmetic inequalities $\circ \in \{\neq, \leq, \geq, <, >\}$ in an obvious manner. That is,

($8_\circ$) $e, w, z, v \models \boldsymbol{B}\phi \circ n$ iff for every $d \in e, \eta \circ n$,

where η is before, given in the semantic definition (8).

3 PROPERTIES

We can show that reasonable properties regarding $\boldsymbol{B}$ and $\boldsymbol{O}$ as considered for $\mathcal{DS}$ also hold in $\mathcal{XS}$. They essentially follow the same style of argumentation (which is desirable), but the proofs are different because the semantic structures and the definition of $\boldsymbol{B}$ is entirely different. For what follows, it will be useful to introduce a bit of additional notation for the RHS of "iff" in the notion of truth for $\boldsymbol{B}$ and $\boldsymbol{O}$. In particular, here is a slightly more notationally precise version of item (8) on the definition of truth:

(8) $e, w, z, v \models \boldsymbol{B}(\phi: n)$ iff for every $d \in e, n = \eta^d(\phi, z)$

where $\eta^d(\phi, z)$ is:

$$\int_{\mathbb{R}^{k+m}} \begin{cases} d(w^*) \times \mathrm{L}(w^*, z') & \text{if } \psi(\vec{r}, \vec{c}, \phi, z) \\ 0 & \text{otherwise} \end{cases}$$

where $\psi(\vec{r}, \vec{c}, \phi, z)$ is:

$w^*[f_i, \langle\rangle] = r_i$ for every i, $z' = z(c_1, \ldots, c_m), (e, w^*, z', v) \models \phi$.

We often drop the superscript on η when the context is obvious.

Let us start with how degrees of belief operate over logical equivalence and connectives:

PROPOSITION 3.1. *The following can be shown in* $\mathcal{XS}$:
(1) *If* $\models \Box(\alpha \equiv \beta)$ *then* $\models \Box(\boldsymbol{B}(\alpha: r) \equiv \boldsymbol{B}(\beta: r))$;
(2) $\models \Box(\boldsymbol{B}(\alpha \wedge \beta: r) \wedge \boldsymbol{B}(\alpha \wedge \neg\beta: r') \supset \boldsymbol{B}(\alpha: r + r'))$;
(3) $\models \Box(\boldsymbol{B}(\alpha: r) \wedge \boldsymbol{B}(\beta: r') \wedge \boldsymbol{B}(\alpha \wedge \beta: r'')$
$\qquad\qquad \supset \boldsymbol{B}(\alpha \vee \beta: r + r' - r''))$.

Note the $\Box$ in front of formulas, which indicates that the properties hold after any sequence of actions.

PROOF. For item 1, consider any (e, w, z). Suppose $e, w, z \models \boldsymbol{B}\alpha = r$. By definition, for every $d \in e, r = \eta(\alpha, z)$. But $\alpha \equiv \beta$ by assumption, and so when we check $e, w^*, z' \models \alpha$ in $\psi(\vec{r}, \vec{c}, \alpha, z)$, it is also the case that $e, w^*, z' \models \beta$. This must mean $e, w, z \models \boldsymbol{B}\beta = r$. The argument is identical for the other direction.

For item 2, consider any e, w, z such that the left hand side of the implication holds. This means that for every $d \in e, \eta(\alpha \wedge \beta, z) = r$ and $\eta(\alpha \wedge \neg\beta, z) = r'$. It is a bit tedious, but by using the semantical definition, we note the following: When expanding the η expressions, we see that after z, certain posterior worlds where $\alpha \wedge \beta$ holds integrate to r, and others where $\alpha \wedge \neg\beta$ hold integrate to r'. But these are distinct, and so the posterior worlds where α holds is the union. Hence we get $e, w, z \models \boldsymbol{B}(\alpha) = r + r'$. Item 3 is analogous. □

Let us now consider the relationship between degrees of belief and meta-knowledge. We have:

PROPOSITION 3.2. (1) $\models \Box(\boldsymbol{B}(\alpha: r) \supset \boldsymbol{K}\boldsymbol{B}(\alpha: r))$; and
(2) $\models \Box(\neg\boldsymbol{B}(\alpha: r) \supset \boldsymbol{K}\neg\boldsymbol{B}(\alpha: r))$.

PROOF. We show the first item, and the second is analogous. Consider any model (e, w, z) such that $e, w, z \models \boldsymbol{B}\alpha = r$. By assumption, for every $d \in e, r = \eta(\alpha, z)$. But by assumption, we observe that for

every such $d, \eta(\boldsymbol{B}\alpha = r, z) = 1$. This can be checked by expanding the definition of $\boldsymbol{B}$ and noting that one of the conditions for using the d-value of a world w^* is when $e, w^*, z \models \boldsymbol{B}\alpha = r$. Since this is true at every w^* by assumption (the real world is, in fact, irrelevant), we integrate and consider the density of every $\vec{r} \in \mathbb{R}^k$, and so we get the probability of 1. □

Following this style of argument, it is not hard to see that modus ponens and weak $S5$ properties [11] holds for $\boldsymbol{K}$:

PROPOSITION 3.3. *The following are valid in* $\mathcal{XS}$:
(1) $\models \Box(\boldsymbol{K}\alpha \wedge \boldsymbol{K}(\alpha \supset \beta) \supset \boldsymbol{K}\beta)$;
(2) $\models \Box(\boldsymbol{K}\alpha \supset \boldsymbol{K}\boldsymbol{K}\alpha)$; *and*
(3) $\models \Box(\neg\boldsymbol{K}\alpha \supset \boldsymbol{K}\neg\boldsymbol{K}\alpha)$.

Finally, the benefit of using only knowing is that is provides the means to succinctly define what is known as well as what is not known. Consider the following properties:

PROPOSITION 3.4. *The following are valid in* $\mathcal{XS}$:
(1) $\boldsymbol{O}(\alpha: r) \models \boldsymbol{B}(\alpha: r)$.
(2) $\boldsymbol{O}\alpha \models \boldsymbol{K}\alpha$.
(3) *Suppose* $\{\alpha, \beta\}$ *do not mention modalities, and* $\alpha \not\models \beta$. *Then* $\boldsymbol{O}\alpha \models \neg\boldsymbol{K}\beta$.

PROOF. For item 1, suppose $e, w, z \models \boldsymbol{O}(\alpha: r)$. This means that every d such that $\eta(\alpha, z) = r$ is in e, and for every $d \in e, \eta(\alpha, z) = r$. Clearly then $e, w, z \models \boldsymbol{B}(\alpha: r)$ as it only uses the second condition. Item 2 is a special case of item 1 with $r = 1$.

For item 3, suppose $e, w, z \models \boldsymbol{O}\alpha$. Consider that for every $d \in e, \eta(\alpha, z) = 1$. Since $\alpha \not\models \beta$, it cannot be that $\eta(\beta, z) = 1$. In other words, there may be worlds where β is true but because $\alpha \not\models \beta$, their d-values would not considered in the integration when obtaining $\eta(\beta, z)$. Although it is possible that worlds where $\neg\beta$ are given a d-value of 0 by some distributions, there must be distributions that accord a non-zero density to worlds where $\neg\beta$. By definition of only-knowing, these latter distributions must also be in e, and in them, $\eta(\beta, z) < 1$. Therefore, $e, w, z \not\models \boldsymbol{K}\beta$. □

There are perhaps many other properties one could explore in the context of knowledge, belief and actions [4, 24], but the above illustrate the reasonableness of the semantic definition for many of the ones considered for $\mathcal{DS}$. Let us now explore the use of $\mathcal{XS}$ for basic action theories.

4 BASIC ACTION THEORY

Using the example from [8], imagine a robot facing the wall, as in Figure 1. Let h be the fluent representing the robot's horizontal distance to the wall. The fluent h would have different values in different possible worlds. In a discrete setting, the set of worlds where h take on a particular value might be given a discrete probability, whereas in the continuous case, they might be given a density. For example, the following initial theory:

$$\forall x(0 < x \leq 10 \supset \boldsymbol{B}(h = x) = 0.1)$$

might be understood as ascribing a probability of 0.1 to $h \in \{1, \ldots, 10\}$ in a discrete setting. In a continuous setting, we might use:

$$\forall x \boldsymbol{B}(h \leq x) = \int_{-\infty}^{x} \mathcal{U}(x; 0, 10). \tag{1}$$

Here, $\mathcal{U}(x; a, b)$, standing for a continuous uniform distribution, returns a density of $1/(b - a)$ when $x \in [a, b]$, and 0 otherwise. In our case, we have a density of 0.1 to $h \in [0, 10]$. In other words, the worlds where $h = 2$ are all collectively accorded a probability of 0.1 in the discrete setting. In a continuous setting, by construction, there is only one world where $h = 2$, and its probability is 0. If we were, however, interested in the probability of $h \leq 2$, we would integrate the density accorded to every world where $h \in [0, 2]$, and so we would get:

$$\int_0^2 0.1 = 0.2.$$

By extension, we have a belief of 0.5 for $h \leq 5$, 0.7 for $h \leq 7$, and 0 for $h \geq 11$.

As discussed previously, the logic also permits uncertainty about distributions. For example, a sentence of the following form:

$$\forall x(0 \leq x \leq 10 \supset \boldsymbol{B}(h \leq x) \neq 0)$$

says that any distribution that accords a non-zero probability density (or mass, depending on whether x is discrete) to the range $0 \leq x \leq 10$ is permitted. So an epistemic state satisfying this sentence will include infinitely many distributions, including ones where:

- the density is 0.1 for every $x \in [0, 10]$;
- the density is 0.05 for $x \in [0, 5]$ and 0.15 for $x \in [5, 10]$;
- the density is 0.15 for every $x \in [3, 8]$ and 0.05 for every $x \in [0, 10]$ but $x \notin [3, 8]$; and so on.

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). In particular, let the action $move(x, y)$ mean the robot intends to move by x units by y is the actual outcome, and $sonar(u)$ to say that a reading of u was obtained on the robot's sonar sensor. Using the specification from [8], we have the following axioms:[6]

- Noisy actions are retrofitted in successor state axioms by postulating that the effects are based on the actual argument:

$$\Box[a]h = u \equiv \exists x, y(a = move(x, y) \wedge u = h - y) \vee$$
$$\forall x, y(a \neq move(x, y) \wedge u = h).$$

  This says that doing a move action means that the subsequent value of h is its current value reduced by the actual outcome y, otherwise it is the same as the current value. This is the equivalent of Reiter's monotonic solution to the frame problem in the $\mathcal{DS}$ family of logics [4, 20].

- We let the noisy effector and sensor have Gaussian error profiles, as is standard in probabilistic robotics [40]:

$$\Box l(move(x, y)) = \mathcal{N}(y; x, 1), \Box l(sonar(z)) = \mathcal{N}(z; h, 0.25)$$

  The first sentence says that the likelihood of $move(x, y)$ is given by a Gaussian density for the number y centered on the true value x with a spread of 1. Likewise, the sensing model is given by a Gaussian density for the value read z centered on the true value h (the fluent that it is measuring) with a spread of 0.25.

- For simplicity, let actions be always executable:

$$\Box poss(a) = 1 \equiv a = move(x, y) \vee a = sonar(z).$$

_____
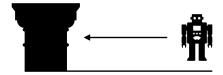[6]Free variables are implicitly quantified from the outside.

Figure 1: A robot moving towards a wall.

Note that successor state and likelihood axioms are stipulated to hold for every action sequence, the rough equivalent of quantifying over situations as needed in basic action theories [21].

Lumping the axioms involving actions as Σ_{dyn}, let us consider entailments wrt the following background theory Σ:[7]

$$(h = 11) \wedge \Sigma_{dyn} \wedge \boldsymbol{B}(\Sigma_{dyn} : 1)$$

together with (1), but using $\mathcal{U}(x; 10, 12)$. Here, the robot is actually 11 units from the wall in the real world, but believes it might be anything from 10 to 12 units from the wall.

The following are entailments of Σ (plotted in Figure 2):

- $[move(-2, -2.32)](h = 13.32) \wedge \boldsymbol{B}(h \geq 11) \leq 0.97$.
  If the robot intends to move away by 2 units, it does not control the actual distance it moves away by, nor does it get to observe it after the move happens. Note that had the move action been noise-free, the degree of belief in $h \geq 11$ would have been one. So the overall outcome is that the robot believes it away from the wall but the distribution has "spread." Note that the entailment also captures the effect on the true distance, as a result of the actual value being included in the background theory Σ. However, to reason about belief, we do not require any mention of h's value in the real world. It is purely for illustration purposes. Of course, the action $move(-2, 3.32)$ is also possible, in which case the robot might have actually moved towards the wall. To unwrap how the degree of belief is obtained, consider that as per Σ, there is only a single d that ascribes a probability density of 0.5 to every $h \in [10, 12]$. That is, let $(e, w, \langle \rangle)$ be any model satisfying Σ. To reason about beliefs after action, using the definition of truth, we have to obtain the following expression:

$$\int_{\mathbb{R}} \int_{\mathbb{R}} \begin{cases} 0.5 \times \mathcal{N}(z; -2, 1) & \text{if } \psi \\ 0 & \text{otherwise} \end{cases}$$

  where ψ is the following three conditions: (a) $x \in [10, 12]$; (b) $w^*[h, \langle \rangle] = x$; and (c) $(e, w^*, move(-2, z)) \models h \geq 11$. So only those worlds where h initially in $[10, 12]$ possibly have non-zero densities. (So whatever value h takes in the real world w is not relevant.) Next, in these worlds, if the move of $move(-2, z)$ leads to $h \geq 11$, then its prior density of 0.5 is now multiplied by the density ascribed to z given that the intended move was -2. (So whatever distance the robot moved in the real world w is also not relevant.) Using the successor state axiom, we would, in fact, see that in any w^* where $h = x$ initially and $(x - z) \geq 11$ is accorded such a density. (Recall that we assign higher likelihoods to z taking

_____
[7]The example below also works with only knowing, but since we do not reason about negative beliefs, we limit ourselves to $\boldsymbol{B}$ in Σ for simplicity.

on negative values, especially around -2.) This expression is integrated over all values of x and z to obtain about 0.95, hence the entailed inequality.

- $[move(-2, -2.32) \cdot sonar(11.5) \cdot sonar(12.6)](h = 13.32) \wedge B(h \geq 11) \geq 0.98$.

  Here, two consecutive sensing actions means that the robot is fairly certain (but not absolutely certain) about the spread of h. For any $(e, w, \langle\rangle)$ satisfying Σ, we expand belief after the actions to the following expression:

  $$\int_{\mathbb{R}} \int_{\mathbb{R}} \begin{cases} \theta & \text{if } \psi \\ 0 & \text{otherwise} \end{cases}$$

  where ψ is identical to above, except that the condition (c) is:

  $$e, w^*, move(-2, z) \cdot sonar(11.5) \cdot sonar(12.6) \models h \geq 11.$$

  That is, for any w^* such that its initial h value is in $[10, 12]$, we use the density only provided $h \geq 11$ is satisfied at $(e, w^*, move(-2, z) \cdot sonar(11.5) \cdot sonar(12.6))$. But using the successor state axiom and the fact that sensing actions do not affect the values of fluents, we use the density provided $h = x, x \in [10, 12]$ and $x - z \geq 11$.

  So what is the density term θ? Recall from the above example, that after a noisy action, its the prior multiplied by the density accorded to the actual argument of the action being z, that is: $0.5 \times \mathcal{N}(z; -2, 1)$. After two sensing actions, this term needs to account for likelihood of the observed value given the true value. So θ is: $0.5 \times \mathcal{N}(z; -2, 1) \times \mathcal{N}(11.5; h, 0.25) \times \mathcal{N}(12.6; h, 0.25)$. So the closer the observed values (11.5 followed by 12.6) are to h in w^*, the "higher" the density term. Not surprisingly, it is the probabilistic analogue to exact sensing [36] where worlds that disagree with the sensed value are simply discarded. Here worlds that disagree are not discarded but weighted less (proportional to the difference in values) than worlds that agree.

- $[move(-2, -2.32) \cdot move(2, 1.9)]B(h \geq 10) \leq 0.8$.

  Moving away by 2 units and moving back by 2 units means that the degree of belief in $h \geq 10$ is not quite one, owing to the two noisy moves. That is, had the move action been noise-free, the robot would be back where it had started, which means the robot would know that $h \geq 10$. The entailment expands as before, except that we have three integrals: one to range over values of h, one to range over the second argument of the first action $move(-2, z)$, and one to range over the second argument of the second action $move(2, z')$.

5 RELATED WORK & CONCLUSIONS

Given the interest in unifying logic and probability, there are an extensive list of related work – we refer interested readers to discussions in [8] – but very few that are closely related. Probabilistic models, Kalman filters, decision theoretic and probabilistic planning languages are either not logics (in allowing for arbitrary connectives and quantifiers) or not general models of actions [18, 35]. Relational probabilistic models [33] offer some logical features (such as clausal reasoning), but not embedded in a general model of action, in allowing to reason about unbounded sequences of actions.
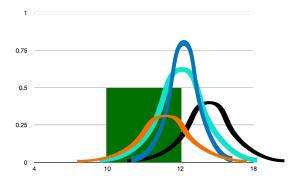

Figure 2: The prior as a green rectangle. The posterior after a single noisy move away from the wall as a black bell curve, a subsequent sensing action as a teal bell curve, and a second sensing action as a navy blue bell curve. The posterior after a noisy move away from the wall and a further noisy move towards the wall given as an orange bell curve.

The closest ones, therefore, are from the knowledge representation literature. Of the ones permitting probabilities, proposals are either propositional [13] or limited to discrete distributions [39]. From an expressiveness viewpoint BHL and $\mathcal{DS}$ are the most general, and from a continuous viewpoint, BL is the closest, which we recast in a modal language with meta-beliefs and only knowing here.

There is some low-hanging fruit: resolving projection, via, say regression or progression [6, 7], for example. But reworking these should be readily possible given the closeness of the semantical definition of B to BL. Arguably, the most interesting direction for the future would be allow for infinitely many fluent terms.

REFERENCES

[1] F. Bacchus. 1990. *Representing and Reasoning with Probabilistic Knowledge*. MIT Press.
[2] F. Bacchus, J. Y. Halpern, and H. J. Levesque. 1999. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence* 111, 1–2 (1999), 171 – 208.
[3] V. Belle. 2021. Logic Meets Learning: From Aristotle to Neural Networks. In *Neuro-Symbolic Artificial Intelligence: The State of the Art*. IOS Press, 78–102.
[4] V. Belle and G. Lakemeyer. 2017. Reasoning about Probabilities in Unbounded First-Order Dynamical Domains. In *IJCAI*.
[5] V. Belle and H. J. Levesque. 2013. Reasoning about Continuous Uncertainty in the Situation Calculus. In *Proc. IJCAI*.
[6] V. Belle and H. J. Levesque. 2013. Reasoning about Probabilities in Dynamic Systems using Goal Regression. In *Proc. UAI*.
[7] V. Belle and H. J. Levesque. 2014. How to progress beliefs in continuous domains. In *Proc. KR*.
[8] V. Belle and H. J. Levesque. 2018. Reasoning about discrete and continuous noisy sensors and effectors in dynamical systems. *Artificial Intelligence* 262 (2018), 189–221.
[9] P. Billingsley. 1995. *Probability and Measure* (3 ed.). Wiley-Interscience.
[10] R. Carnap. 1951. *Logical foundations of probability*. Routledge and Kegan Paul London.

[11] B.F. Chellas. 1980. *Modal logic.* Cambridge University Press.

[12] B. de Finetti. 1972. *Probability, Induction, and Statistics.* New York: John Wiley.

[13] Y. A Feldman. 1986. A decidable propositional dynamic logic with explicit probabilities. *Inf. Control* 63 (November 1986), 11–38. Issue 1-2. https://doi.org/10.1016/S0019-9958(84)80039-X

[14] H. Gaifman. 1964. Concerning measures in first order calculi. *Israel Journal of Mathematics* 2, 1 (1964), 1–18.

[15] P.R. Halmos. 1950. Measure theory. *Van Nostrad Reinhold Company* (1950).

[16] J.Y. Halpern. 1990. An analysis of first-order logics of probability. *Artificial Intelligence* 46, 3 (1990), 311–350.

[17] P. Hitzler and M. Saker. 2022. Neuro-Symbolic Artificial Intelligence: The State of the Art. (2022).

[18] L> Kaelbling, M. Littman, and A. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101, 1–2 (1998), 99 – 134. https://doi.org/10.1016/S0004-3702(98)00023-X

[19] A. Kharazishvili. 2011. Measurability properties of Vitali sets. *The American Mathematical Monthly* 118, 8 (2011), 693–703.

[20] G. Lakemeyer and H. Levesque. 2004. Situations, Si! Situation terms, No!. In *Proc. KR.* 516–526.

[21] G. Lakemeyer and H. Levesque. 2005. Semantics for a useful fragment of the situation calculus. In *IJCAI.* 490–496.

[22] G. Lakemeyer and H. J. Levesque. 2007. Cognitive robotics. In *Handbook of Knowledge Representation.* Elsevier, 869–886.

[23] H. J. Levesque. 1990. All I know: a study in autoepistemic logic. *Artificial Intelligence* 42, 2-3 (1990), 263–309.

[24] H. J. Levesque and G. Lakemeyer. 2001. *The logic of knowledge bases.* The MIT Press.

[25] D. Liu and Q. Feng. 2021. On the Progression of Belief. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning.* 465–474. https://doi.org/10.24963/kr.2021/44

[26] D. Liu and G. Lakemeyer. 2021. Reasoning about Beliefs and Meta-Beliefs by Regression in an Expressive Probabilistic Action Logic. IJCAI.

[27] J. McCarthy and P. J. Hayes. 1969. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In *Machine Intelligence.* 463–502.

[28] J.R. Milton. 2011. Induction before Hume. In *Inductive Logic*, Stephan Hartmann Dov M. Gabbay and John Woods (Eds.). Handbook of the History of Logic, Vol. 10. North-Holland, 1 – 41. https://doi.org/10.1016/B978-0-444-52936-7.50001-X

[29] S. Muggleton, L. De Raedt, D. Poole, I. Bratko, P. Flach, K. Inoue, and A. Srinivasan. 2012. ILP turns 20. *Machine learning* 86, 1 (2012), 3–23.

[30] K. Murphy. 2012. *Machine learning: a probabilistic perspective.* The MIT Press.

[31] N. J. Nilsson. 1986. Probabilistic logic. *Artificial intelligence* 28, 1 (1986), 71–87.

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

[33] L. De Raedt, K. Kersting, S. Natarajan, and D. Poole. 2016. Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis lectures on artificial intelligence and machine learning* 10, 2 (2016), 1–189.

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

[35] S. Sanner. 2011. *Relational Dynamic Influence Diagram Language (RDDL): Language Description.* Technical Report. Australian National University.

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

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

[38] M. Thielscher. 1999. From situation calculus to fluent calculus: state update axioms as a solution to the inferential frame problem. *Artificial Intelligence* 111, 1-2 (1999), 277–299. https://doi.org/10.1016/S0004-3702(99)00033-8

[39] M. Thielscher. 2001. Planning with Noisy Actions (Preliminary Report). In *Proc. Australian Joint Conference on Artificial Intelligence.* 27–45.

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

[41] A. M. Turing. 1936. On computable numbers, with an application to the entscheidungsproblem. *J. of Math* 58 (1936), 345–363.

[42] H. P. van Ditmarsch, A. Herzig, and T. De Lima. 2011. From Situation Calculus to Dynamic Epistemic Logic. *J. Log. Comput.* 21, 2 (2011), 179–204.

[43] R. Waldinger. 1977. Achieving several goals simultaneously. In *Machine Intelligence.* Vol. 8. 94–136.