

# Logic, Probability and Action: A Situation Calculus Perspective<sup>\*</sup>

Vaishak Belle

University of Edinburgh, UK & Alan Turing Institute, UK  
vai shak@ed . ac . uk

**Abstract.** The unification of logic and probability is a long-standing concern in AI, and more generally, in the philosophy of science. In essence, logic provides an easy way to specify properties that must hold in every possible world, and probability allows us to further quantify the weight and ratio of the worlds that must satisfy a property. To that end, numerous developments have been undertaken, culminating in proposals such as probabilistic relational models. While this progress has been notable, a general-purpose first-order knowledge representation language to reason about probabilities and dynamics, including in continuous settings, is still to emerge. In this paper, we survey recent results pertaining to the integration of logic, probability and actions in the situation calculus, which is arguably one of the oldest and most well-known formalisms. We then explore reduction theorems and programming interfaces for the language. These results are motivated in the context of cognitive robotics (as envisioned by Reiter and his colleagues) for the sake of concreteness. Overall, the advantage of proving results for such a general language is that it becomes possible to adapt them to any special-purpose fragment, including but not limited to popular probabilistic relational models.

## 1 Introduction

The unification of logic and probability is a long-standing concern in AI [72], and more generally, in the philosophy of science [31]. The motivation stems from the observation that (human and agent) knowledge is almost always incomplete. It is then not enough to say that some formula  $\phi$  is unknown. One must also know which of  $\phi$  or  $\neg\phi$  is the more likely, and by how much. On the more pragmatic side, when reasoning about uncertain propositions and statements, it is beneficial to be able to leverage the underlying relational structure. Basically, logic provides an easy way to specify properties that must hold in every possible world, and probability allows us to further quantify the weight and ratio of the worlds that must satisfy a property. For example, the sibling relation is symmetric in every possible world, whereas the influence of smoking among siblings can be considered a statistical property, perhaps only true in 80% of the worlds.

Another argument increasingly made in favor of unifying logic and probability is that perhaps it would help us enable an apparatus analogous to Kahneman's so-called *System 1* versus *System 2* processing in human cognition [43]. That is, we want to

---

<sup>\*</sup> The author was supported by a Royal Society University Research Fellowship.

interface experiential and reactive processing (assumed to be handled by some data-driven probabilistic learning methodology) with cogitative processing (assumed to be handled by a deliberative reasoning methodology).

To that end, numerous developments have been undertaken in AI. Closely following Bayesian networks [70, 52], and particle filters [30, 35], the areas of statistical relational learning and probabilistic relational modeling [28, 36] emerged, and have been very successful. Since the world is rarely static, the application of such proposals to dynamic worlds has also seen many successes, e.g., [68, 80]. However, these closely follow propositional representations, such as Bayesian networks, using logic purely for templating purposes (i.e., syntactic sugar in programming language parlance). So, although the progress has been notable, a general-purpose first-order knowledge representation language to reason about probabilities and dynamics, including in continuous settings, is still to emerge.

In the early days of the field, approaches such as [33] provided a logical language that allowed one to reason about the probabilities of atoms, which could be further combined over logical connectives. That work has inspired numerous extensions for reasoning about dynamics. But this has been primarily in the propositional setting [82, 40, 45], or with discrete probabilistic models [78]. (See [15] for extended discussions.) In this paper, we survey recent results pertaining to the integration of logic, probability and actions in the situation calculus [65, 71]. The situation calculus is one of the oldest and most well-known knowledge representation formalisms. In that regard, the results illustrate that we obtain perhaps the most expressive formalism for reasoning about degrees of belief in the presence of noisy sensing and acting. For that language, we then explore reduction theorems and programming interfaces. Of course, the advantage of proving results for such a general language is that it becomes possible to adapt them to any special-purpose fragment, including but not limited to popular probabilistic relational models.

To make the discussion below concrete, we motivate one possible application of such a language: *cognitive robotics*, as envisioned by Reiter [71] and further discussed in [48]. This is clearly not the only application of a language such as the situation calculus, which has found applications in areas such as service composition, databases, automated planning, decision-theoretic reasoning and multi-agent systems [71, 83].

## 2 Motivation: Cognitive Robotics

The design and control of autonomous agents, such as robots, has been a major concern in artificial intelligence since the early days [65]. Robots can be viewed as systems that need to act purposefully in open-ended environments, and so are required to exhibit everyday commonsensical behavior. For the most part, however, traditional robotics has taken a “bottom-up” approach [79] focusing on low-level sensor-effector feedback. Perhaps the most dominant reason for this is that controllers for physical robots need to address the noise in effectors and sensors, often characterized by continuous probability distributions, which significantly complicates the reasoning and planning problems faced by a robot. While the simplicity of Bayesian statistics, defined over a fixed number of (propositional) random variables, has enabled the successful handling of probabilistic

information in robotics modules, the flip side is that the applicability of contemporary methods is at the mercy of the roboticist's ingenuity. It is also unclear how precisely commonsensical knowledge can be specified using conditional independences between random variables while also accounting for how these dependencies further change as the result of actions.

Cognitive robotics [48], as envisioned by Reiter and his colleagues [71], follows closely in the footsteps of McCarthy's seminal ideas [64]: it takes the view that understanding the relationships between the beliefs of the agent and the actions at its disposal is key to a commonsensical robot that can operate purposefully in uncertain, dynamic worlds. In particular, it considers the study of knowledge representation and reasoning problems faced by the agent when attempting to answer questions such as [53]:

- *to execute a program, what information does a robot need to have at the outset vs. the information that it can acquire en route by perceptual means?*
- *what does the robot need to know about its environment vs. what need only be known by the designer?*
- *when should a robot use perception to find out if something is true as opposed to reasoning about what it knows was true in the past?*

The goal, in other words, is to develop a theory of high-level control that maps the knowledge, ignorance, intention and desires of the agent to appropriate actions. In this sense, cognitive robotics not only aims to connect to traditional robotics, which already leverages probabilistic reasoning, vision and learning for stochastic control, but also to relate to many other areas of AI, including automated planning, agent-oriented programming, belief-desire-intention architectures, and formal epistemology.

In lieu of this agenda, many sophisticated control methodologies and formal accounts have emerged, summarized in the following section. Unfortunately, despite the richness of these proposals, one criticism leveled at much of the work in cognitive robotics is that the theory is far removed from the kind of continuous uncertainty and noise seen in typical robotic applications. That is, the formal machinery of GOLOG to date does not address the complications due to noise and uncertainty in realistic robotic applications, at least in a way that relates these complications to what the robot believes, and how that changes over actions. The assumptions under which real-time behavior can be expected is also left open. For example, can standard probabilistic projection methodologies, such as Kalman and particle filters, be subsumed as part of a general logical framework?

The results discussed in this article can be viewed as a research agenda that attempts to bridge the gap between knowledge representation advances and robotic systems. By generalizing logic-based knowledge representation languages to reason about discrete and continuous probability distributions in the specification of both the initial beliefs of the agent and the noise in the sensors and effectors, the idea is to contribute to commonsensical and provably correct high-level controllers for agents in noisy worlds.

### 3 Tools of the Trade

To represent the beliefs and the actions, efforts in cognitive robotics would need to rely on a formal language of suitable expressiveness. Reiter's variant of the situation

calculus has perhaps enjoyed the most success among first-order formalisms, although related proposals offer attractive properties of their own.<sup>1</sup> Reiter’s variant was also the language considered in a recent survey on cognitive robotics [48], and so the reported results can easily be put into context.<sup>2</sup>

In this section, we will briefly recap some of the main foundational results discussed in [48]. In a few cases, we report on recent developments expanding on those results.

### 3.1 Language

Intuitively, the language  $\mathcal{L}$  of the situation calculus [65] is a many-sorted dialect of predicate calculus, with sorts for *actions*, *situations* and *objects* (for everything else, and includes the set of reals  $\mathbb{R}$  as a subsort). A situation represents a world history as a sequence of actions. A set of initial situations correspond to the ways the world might be initially. Successor situations are the result of doing actions, where the term  $do(a, s)$  denotes the unique situation obtained on doing  $a$  in  $s$ . The term  $do(\bar{a}, s)$ , where  $\bar{a}$  is the sequence  $[a_1, \dots, a_n]$  abbreviates  $do(a_n, do(\dots, do(a_1, s) \dots))$ . Initial situations are defined as those without a predecessor, and we let the constant  $S_0$  denote the actual initial situation. See [71] for a comprehensive treatment.

The picture that emerges from the above is a set of trees, each rooted at an initial situation and whose edges are actions. In general, we want the values of predicates and functions to vary from situation to situation. For this purpose,  $\mathcal{L}$  includes *fluents* whose last argument is always a situation.

Following [71], dynamic domains in  $\mathcal{L}$  are modeled by means of a *basic action theory*  $\mathcal{D}$ , which consists domain-independent foundational axioms, and a domain-dependent first-order initial theory  $\mathcal{D}_0$  (standing for what is true initially), and domain-dependent precondition and effect axioms, the latter taking the form of so-called successor state axioms that incorporates a monotonic solution to the frame problem [71].

To represent knowledge, and how that changes, one appeals to the possible-worlds approach [34]. The idea is 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. Essentially, situations can be viewed as possible worlds [74]: a special binary fluent  $K$ , taking two situation arguments determines the accessibility relation between worlds. So,  $K(s', s)$  says that when the agent is at  $s$ , he considers  $s'$  possible. *Knowledge*, then, is simply truth at accessible worlds:  $Knows(\phi, s) \doteq \forall s'. K(s', s) \supset \phi[s']$ .

Sensing axioms additionally capture the discovery of the truth values of fluents. For example, to check whether  $f$  is true at  $s$ , we would use:  $SF(sense_{true} f, s) \equiv f(s) = 1$ . A successor state axiom formalizes the incorporation of these sensed values in the agent’s

<sup>1</sup> For example, the fluent calculus [77] offers an intuitive and simple state update mechanism in a first-order setting, and extensions of propositional dynamic logic [41] offer decidable formalisms.

<sup>2</sup> There has been considerable debate on why a quantified relational language is crucial for knowledge representation and commonsense reasoning; see references in [57, 26], for example. Moreover, owing to the generality of the underlying language, decidable variants can be developed (e.g., [38, 20]).

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

### 3.2 Reasoning Problems

A fundamental problem underlying almost all applications involving basic action theories is *projection*. Given a sequence of actions  $a_1$  through  $a_n$ , denoted  $\bar{a} = [a_1, \dots, a_n]$ , we are often interested in asking whether  $\phi$  holds after these via entailment:  $\mathcal{D} \models \phi[do(\bar{a}, S_0)]$ ? One of the main results by Reiter is the existence of a reduction operator called *regression* that eliminates the actions:  $\mathcal{D} \models \phi[do(\bar{a}, S_0)]$  iff  $\mathcal{D}_{una} \cup \mathcal{D}_0 \models \mathcal{R}[\phi[do(\bar{a}, S_0)]]$ . Here,  $\mathcal{D}_{una}$  is an axiom that declares that all named actions are unique, and  $\mathcal{R}[\phi[do(\bar{a}, S_0)]]$  mentions only a single situation term,  $S_0$ .

In the worst case, regressed formulas are exponentially long in the length of the action sequence [71], and so it has been argued that for long-lived agents like robots, continually updating the current view of the state of the world, is perhaps better suited. Lin and Reiter [60] proposed a theory of progression that satisfies:  $\mathcal{D} \models \phi[do(\bar{a}, S_0)]$  iff  $\mathcal{D}_{una} \cup \mathcal{P}(\mathcal{D}_0, \bar{a}) \models \phi[S_0]$ . Here  $\mathcal{P}(\mathcal{D}_0, \bar{a})$  is the updated initial theory that denotes the state of the world on doing  $\bar{a}$ . In general, progression requires second-order logic, but many special cases that are definable in first-order logic have since been identified (e.g., [61]).

### 3.3 Closed vs Open Worlds

$\mathcal{D}_0$  is assumed to be any set of first-order formulas, but then computing its entailments, regardless of whether we appeal to regression or progression, would be undecidable. Thus, restricting the theory to be equivalent to a relational database is one possible tractable fragment, but this makes the closed world assumption which is not really desirable for robotics. A second possibility is to assume that at the time of query evaluation, the agent has complete knowledge about the predicates mentioned in the query. This leads to a notion of local completeness [27]. A third possibility is to provide some control over the computational power of the evaluation scheme, leading to a form of *limited reasoning*. First-order fragments such as *proper* and *proper*<sup>+</sup> [56, 47], which correspond to an infinite set of ground literals and clauses respectively, have been shown to work well with projection schemes for restricted classes of action theories [62, 61].

An altogether different and more general strategy for reasoning about incomplete knowledge is to utilize the epistemic situation calculus. A regression theorem was already proved in early work [74], and a progression theorem has been considered in [63]. However, since propositional reasoning in epistemic logic is already intractable [34], results such as the *representation theorem* [57] that shows how epistemic operators can be eliminated under *epistemic closure* (i.e., knowing what one knows as well as what one does not know) needs to be leveraged at least. Alternatively, one could perhaps appeal to limited reasoning in the epistemic setting [46].

### 3.4 High-Level Control

To program agents whose actions are interpreted over a basic action theory, *high-level programming languages* such as GOLOG emerged [54]. These languages contained the usual programming features like sequence, conditional, iteration, recursive procedures, and concurrency but the key difference was that the primitive instruction was an action from a basic action theory. The execution of the program was then understood as  $\mathcal{D} \models Do(\delta, S_0, do(\bar{a}, S_0))$  where  $\delta$  is a GOLOG program, and on starting from  $S_0$ , the program successfully terminates in  $do(\bar{a}, S_0)$ . So, from  $S_0$ , executing the program leads to performing actions  $\bar{a}$ .

As argued in [48], GOLOG programs can range from a fully deterministic instruction  $a_1; \dots; a_n$  to a general search **while**  $\neg\phi$  **do**  $\pi a$ .  $a$ : the former instructs the agent to perform action  $a_1$ , then  $a_2$ , and so on until  $a_n$  in sequence, and the latter instructs to try every possible action (sequence) until the goal is satisfied. It is between these two extremes where GOLOG is most powerful: it enables a partial specification of programs that can perform guided search for sub-goals in the presence of other loopy or conditional plans.

To guide search in the presence of nondeterminism, rewards can be stipulated on situations leading to a decision-theoretic machinery [17]. Alternatively, if the nondeterminism is a result of not knowing the true state of the world, sensing actions can be incorporated during program execution, leading to an *online* semantics for GOLOG execution [73].

## 4 Tools Revisited

In this section, we revisit the results from the previous section and discuss how these have been generalized to account for realistic, continuous, models of noise.

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 [1], 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). Not only is it embedded in the rich theory of the situation calculus, including the use of Reiter's successor state axioms, it is also a stochastic extension to the categorical epistemic situation calculus. 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 [79]. 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 [70]. 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. (See [15] for extensive discussions.) All the properties of belief will then follow at a corresponding level of specificity.

#### 4.1 Language

The BHL model makes use of two distinguished binary fluents  $p$  and  $l$  [9]. 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  $\mathcal{D}_0$ . The term  $l(a, s)$  is intended to denote the likelihood of action  $a$  in situation  $s$  to capture noisy sensors and effectors. For example, think of a sonar aimed at the wall, which gives a reading for the true value of a fluent  $f$  that corresponds to the distance between the robot and the wall. Supposing the sonar's readings are subject to additive Gaussian noise. If now a reading of  $z$  were observed on the sonar, intuitively, those situations where  $f = z$  should be considered more probable than those where  $f \neq z$ . Then we would have:  $l(\text{sense}_f(z), s) = u \equiv u = \mathcal{N}(z - f(s); 0, 1)$ . Here, a standard normal is assumed, where the mean is 0, and the variance is 1.<sup>3</sup> Analogously, noisy effectors can be modeled using actions with double the arguments:  $l(\text{move}(x, y), s) = u \equiv u = \mathcal{N}(y - x; 0, 1)$ . This says the difference between actual distance moved and the intended amount is normally distributed, corresponding to additive Gaussian noise. Such noise models can also be made context dependent (e.g., specifying the sensor's error profile to be worse for lower temperatures, where the temperature value is situation-dependent). In the case of noisy effectors, the successor state axioms have to be defined to use the second argument, as this is what actually happens at a situation [15].

Analogous to the notion of knowledge, the *degree of belief* in  $\phi$  in situation  $s$  is defined as the weight of accessible worlds where  $\phi$  is true:

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

Here,  $\gamma$  is the normalization factor and corresponds to the numerator but with  $\phi$  replaced by *true*. The change in  $p$  values over actions is specified using a successor state axiom, analogous to the one for  $K$ :  $p(s', do(a, s)) = u \equiv \exists s'' [s' = do(a, s'') \wedge Poss(a, s'') \wedge u = p(s'', s) \times l(a, s'')] \vee \neg \exists s'' [s' = do(a, s'') \wedge Poss(a, s'') \wedge u = 0]$ . This axioms determines how  $l$  affects the  $p$ -value of successor situations.

As the BHL model is defined as a sum over possible worlds, it cannot actually handle Gaussians and other continuous distributions involving  $\pi$ ,  $e$ , exponentiation, and so on. Therefore, BHL always consider discrete probability distributions that *approximate* the continuous ones. However, this limitation was lifted in [15], which shows how  $Bel$  is defined in continuous domains.

#### 4.2 Reasoning Problems

The projection problem in this setting is geared for reasoning about formulas that now mention  $Bel$ . In particular, we might be interested in knowing whether  $\mathcal{D} \models Bel(\phi, do(\bar{a}, S_0)) \geq r$  for a real number  $r$ .

<sup>3</sup> If the specification of the  $p$ -axiom or the  $l$ -axiom includes disjunctions and existential quantifiers, we will then be dealing with uncertainty about distributions. See [14], for example.

One reductive approach would be to translate both  $\mathcal{D}$  and  $\phi$ , which would mention *Bel*, into a predicate logic formula. This approach, however, presents a serious computational problem because belief formulas expand into a large number of sentences, resulting in an enormous search space with initial and successor situations. The other issue with this approach is that sums (and integrals in the continuous case) reduce to complicated second-order formulas.

In [10], it is shown how Reiter's regression operator can be generalized to operate directly on *Bel*-terms. This involves appealing to the likelihood axioms. For example, imagine a robot that is uncertain about its distance  $d$  to the wall, and the prior is a uniform distribution on the interval  $[2, 12]$ . Assume the robot (noise-free) moves away by 2 units and is now interested in the belief about  $d \leq 5$ . Regression would tell the robot that this is equivalent to its initial beliefs about  $d \leq 3$  which here would lead to a value of .1. Imagine then the robot is also equipped with a sonar unit with additive Gaussian noise. After moving away by 2 units, if the sonar were now to provide a reading of 8, then regression would derive that belief about  $d \leq 5$  is equivalent to  $1/\gamma \times \int_2^3 .1 \times \mathcal{N}(6-x; 0, 1) dx$ . Essentially, the posterior belief about  $d \leq 5$  is reformulated as the product of the prior belief about  $d \leq 3$  and the likelihood of  $d \leq 3$  given an observation of 6. That is, observing 8 after moving away by 2 units is equated here to observing 6 initially. (Here,  $\gamma$  is the normalization factor.)

Progression too could potentially be addressed by expanding formulas involving *Bel*-terms, but it is far from clear what precisely this would look like. In particular, given initial beliefs about fluents (such as the one about  $d$  earlier), we intuit that a progression account would inform us how this distribution changed. For example, on moving away from the wall by 2 units, we would now expect  $d$  to be uniformly distributed on the interval  $[4, 14]$ . However, this leads to a complication: because if the robot had instead moved towards the wall by 4 units, then those points where  $d \in [2, 4]$  initially are mapped to a single point  $d = 0$  that should then obtain a probability mass of .2, while the other points retain their initial density of .1. In [11], it is shown that for a certain class of basic action theories called *invertible theories*, such complications are avoidable, and moreover, the progressed database can be specified by means of simple syntactic manipulations.

### 4.3 Closed vs Open Worlds

The closed vs open world discussion does not seem immediately interesting here, because, after all, the language is clearly open in the sense of not knowing the values of fluents, and according a distribution to these values. However, consider that the closed-world assumption was also motivated previously by computational concerns. In that regard, the above regression and progression results already studied special cases involving conjugate distributions [18], such as Gaussians which admit attractive analytical simplifications. For example, efficient Kalman filters [79] often make the assumption that the initial prior and the noise models are Gaussians, in which case the posterior would also be a Gaussian. In [12], it is further shown that when the initial belief is a Bayesian network, by way of regression, projection can be handled effectively by sampling. (That is, once the formula is regressed, the network is sampled and the samples are evaluated against the regressed formula.)



In the context of probabilistic specifications, the notion of “open”-ness can perhaps be interpreted differently. We can take this to mean that we do not know the distribution of the random variables, or even that the set of random variables is not known in advance. As argued earlier, this is precisely the motivation for the BHL scheme, and a recent modal reformulation of BHL illustrates the properties of such a language in detail [7]. A detailed demonstration of how such specifications would work in the context of *robot localization* was given in [14].

The question of how to effectively compute beliefs in such rich settings is not clear, however. We remark that various static frameworks have emerged for handling imprecision or uncertainty in probabilistic specifications [66, 24, 58]. For example, when we have finitely many random variables but there is uncertainty about the underlying distribution, credal representations are of interest [24], and under certain conditions, they can be learned and reasoned with in an efficient manner [58]. On the other hand, when we have infinitely many random variables (but with a single underlying distribution), proposals such as [72] and [2] are of interest, the latter being a weighted representation inspired by *proper*<sup>+</sup> knowledge bases. Extending these to allow uncertainty about the underlying distribution may also be possible. Despite being static, by means of regression or progression, perhaps such open knowledge bases can be exploited for cognitive robotics applications, but that remains to be seen.

#### 4.4 High-Level Control

A high-level programming language that deals with noise has to reason about two kinds of complications. First, when a noisy physical or sensing action in the program is performed, we must condition the next instruction on how the belief has changed as a result of that action. Second, because sensing actions in the language are of the form *sense*( $z$ ) that expects an input  $z$ , an *offline* execution would simulate possible values for  $z$  whereas an *online* execution would expect an external source to provide  $z$  (e.g., reading off the value of a sonar). We also would not want the designer to be needlessly encumbered by the error profiles of the various effectors and sensors, so she has to be encouraged to program around *sense-act* loops; that is, every action sequence should be accompanied with a suitable number of sensing readings so that the agent is “confident” (i.e., the distribution of the fluent in question is *narrow*) before performing more actions. In [13], such a desiderata was realized to yield a stochastic version of knowledge-based programming [71]. Primitive instructions are *dummy* versions of noisy actions and sensors; e.g., *move*( $x,y$ ) is simply *move*( $x$ ) and *sonar*( $z$ ) is simply *sonar*. The idea then is that the modeler simply uses these dummy versions as she would with noise-free actions, but the execution semantics incorporates the change in belief. It is further shown that program execution can be realized by means of a *particle filtering* [79] strategy: weighted samples are drawn from the initial beliefs, which correspond to initial situations, and on performing actions, fluent values at these situations are updated by means of the successor state axioms. The degree of belief in  $\phi$  corresponds to summing up the weights of samples where  $\phi$  is true.

Such an approach can be contrasted with notable probabilistic relational modelling proposals such as [68]: the difference mainly pertains to three sources of generality.

First, a language like the situation calculus allows knowledge bases to be arbitrary quantificational theories, and BHL further allows uncertainty about the distributions defined for these theories. Second, the situation calculus, and by extension, GOLOG and the paradigm in [13] allows us to reason about non-terminating and unbounded behavior [23]. Third, since an explicit belief state is allowed, it becomes possible to provide a systematic and generic treatment for multiple agents [44, 6].

On the issue of tractable reasoning, an interesting observation is that because these programs require reasoning with an explicit belief state [34], one might wonder whether the programs can be “compiled” to a *reactive plan*, possibly with loops, where the next action to be performed depends only on the sensing information received in the current state. This relates knowledge-based programming to *generalized planning* [55, 76], and of course, the advantage is also that numerous strategies have been identified to synthesize such loopy, reactive plans. Such plans are also shown to be sufficient for goal achievability [59]; however, knowledge-based programs are known to be exponentially more succinct than loopy, reactive plans [49]. In [42], a generic algorithmic framework was proposed to synthesize such plans in noise-free environments. How the correctness of such plans should be generalized to noisy environments was considered in [8, 3]. The algorithmic synthesis problem was then considered in [81].

## 5 Related Work and Discussions

There are many threads of research in AI, automated planning and robotics that are close in spirit to what is reported here. For example, belief update via the incorporation of sensor information has been considered in probabilistic formalisms such as Bayesian networks [70, 52], Kalman and particle filters [79]. But these have difficulties handling strict uncertainty. Moreover, since rich models of actions are rarely incorporated, shifting conditional dependencies and distributions are hard to address in a general way. While there are graphical formalisms with an account of actions, such as [25, 39], 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 possible in the BHL scheme. Related to these are relational probabilistic models [67, 66, 32, 21]. Although limited accounts for dynamic domains are common here [50, 69], explicit actions are seldom addressed in a general way. We refer interested readers to discussions in [15], where differences are also drawn to prior developments in reasoning about actions, including stochastic but non-epistemic GOLOG dialects [37].

Arguably, many of the linguistic restrictions of such frameworks is often motivated by computational considerations. So what is to be gained by a general approach? This question is especially significant when we take into account that numerous “hybrid” approaches have emerged over the years that provide a bridge between a high-level language and a low-level operation [19, 51]. Our sense is that while these and other approaches are noteworthy, and are extended in a modular manner to keep things tractable and workable on an actual physical robot, it still leaves a lot at the mercy of the robotist’s ingenuity. For example, extending an image recognition algorithm to reason about a structured world is indeed possible, but it is more likely than not that this ontology is also useful for a number of other components, such as the robot’s grasping arm; more-

over, changes to one must mean changes to all. Abstracting a complex behavior module of a robot is a painstaking effort: often the robot’s modules are written in different programming languages with varying levels of abstraction, and to reduce these interactions to *atoms* in the high-level language would require considerable know-how of the system. Moreover, although a roboticist can abstract probabilistic sensors in terms of high-level categorical ones, there is loss in detail, as it is not clear at the outset which aspect of the sensor data is being approximated and by how much. Thus, all of these “bottom-up” approaches ultimately challenge the claim that the underlying theory is a genuine characterization of the agent.

In service of that, the contributions reported in this work attempt to express all the (inner and outer) workings of a robot in a *single mathematical language*: a mathematical language that can capture rich structure as well as natively reason about the probabilistic uncertainty plaguing a robot; a mathematical language that can reason with all available information, some of which may be probabilistic, and some categorical; a mathematical language that can reason about the physical world at different levels of abstraction, in terms of objects, atoms, and whatever else physicists determine best describes our view of the world. Undoubtedly, given this glaring expressiveness, the agenda will raise significant challenges for the applicability of the proposal in contemporary robots, but our view is that, it will also engender novel extensions to existing algorithms to cope with the expressiveness. Identifying tractable fragments, for example, will engender novel theoretical work. As already discussed, many proposals from the statistical relational learning community are very promising in this regard, and are making steady progress towards the overall ambition. (But as discussed, they fall short in terms of being able to reason about non-terminating behavior, arbitrary first-order quantification, among other things, and so identifying richer fragments is a worthwhile direction.) It is also worth remarking that the tractability of reasoning (and planning) has been the primary focus of much of the research in knowledge representation. The broader question of how to learn models has received lesser attention, and this is precisely where statistical relational learning and related paradigms will prove useful [4]. (It would be especially interesting to consider relational learning with neural modules [29].) Indeed, in addition to approaches such as [22, 75], there have been a number of advances recently on learning dynamic representations (e.g., [68]), which might provide fertile ground to lift such ideas for cognitive robotics. Computability results for qualitative learning in dynamic epistemic logic has been studied in [16]. Recently, proper<sup>+</sup> knowledge bases were shown to be polynomial-time learnable for querying tasks [5]. Ultimately, learning may provide a means to coherently arrive at action descriptions at different levels of granularity from data [26]. In the long term, the science of building a robot, which currently is more of an art, can perhaps be approached systematically. More significantly, through the agenda of cognitive robotics, we might gain deep insights on how commonsense knowledge and actions interact for general-purpose, open-ended robots. In that regard, the integration of logic, probability and actions will play a key role.

## References

1. F. Bacchus, J. Y. Halpern, and H. J. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1–2):171 – 208, 1999.

2. V. Belle. Weighted model counting with function symbols. In *UAI*, 2017.
3. V. Belle. On plans with loops and noise. In *AAMAS*, 2018.
4. V. Belle. Symbolic logic meets machine learning: A brief survey in infinite domains, 2020.
5. V. Belle and B. Juba. Implicitly learning to reason in first-order logic. *NeurIPS*, 2019.
6. V. Belle and G. Lakemeyer. Multiagent only knowing in dynamic systems. *Journal of Artificial Intelligence Research*, 49, 2014.
7. V. Belle and G. Lakemeyer. Reasoning about probabilities in unbounded first-order dynamical domains. In *IJCAI*, 2017.
8. V. Belle and H. Levesque. Foundations for generalized planning in unbounded stochastic domains. In *KR*, 2016.
9. V. Belle and H. J. Levesque. Reasoning about continuous uncertainty in the situation calculus. In *Proc. IJCAI*, 2013.
10. V. Belle and H. J. Levesque. Reasoning about probabilities in dynamic systems using goal regression. In *Proc. UAI*, 2013.
11. V. Belle and H. J. Levesque. How to progress beliefs in continuous domains. In *KR*, 2014.
12. V. Belle and H. J. Levesque. PREGO: An Action Language for Belief-Based Cognitive Robotics in Continuous Domains. In *Proc. AAAI*, 2014.
13. V. Belle and H. J. Levesque. Allegro: Belief-based programming in stochastic dynamical domains. In *IJCAI*, 2015.
14. V. Belle and H. J. Levesque. A logical theory of localization. In *Studia Logica*. 2015.
15. V. Belle and H. J. Levesque. Reasoning about discrete and continuous noisy sensors and effectors in dynamical systems. *Artificial Intelligence*, 262:189–221, 2018.
16. T. Bolander and N. Gierasimczuk. Learning actions models: Qualitative approach. In *International Workshop on Logic, Rationality and Interaction*, pages 40–52. Springer, 2015.
17. C. Boutilier, R. Reiter, M. Soutchanski, and S. Thrun. Decision-theoretic, high-level agent programming in the situation calculus. In *Proc. AAAI*, pages 355–362, 2000.
18. G. E. P. Box and G. C. Tiao. *Bayesian inference in statistical analysis*. Addison-Wesley, 1973.
19. W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artif. Intell.*, 114(1-2):3–55, 1999.
20. D. Calvanese, G. De Giacomo, M. Montali, and F. Patrizi. First-order  $\mu$ -calculus over generic transition systems and applications to the situation calculus. *Inf. Comput.*, 259(3):328–347, 2018.
21. J. Choi, E. Amir, and D. J. Hill. Lifted inference for relational continuous models. In *Proc. UAI*, pages 126–134, 2010.
22. J. Choi, A. Guzman-Rivera, and E. Amir. Lifted relational kalman filtering. In *Proc. IJCAI*, pages 2092–2099, 2011.
23. J. Claßen and G. Lakemeyer. A logic for non-terminating golog programs. In *KR*, pages 589–599, 2008.
24. F. G. Cozman. Credal networks. *Artificial Intelligence*, 120(2):199 – 233, 2000.
25. A. Darwiche and M. Goldszmidt. Action networks: A framework for reasoning about actions and change under uncertainty. In *Proc. UAI*, pages 136–144, 1994.
26. E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103, 2015.
27. G. De Giacomo and H. J. Levesque. Projection using regression and sensors. In *IJCAI*, 1999.
28. L. De Raedt and K. Kersting. Statistical relational learning. In *Encyclopedia of Machine Learning*, pages 916–924. Springer, 2011.
29. L. De Raedt, R. Manhaeve, S. Dumancic, T. Demeester, and A. Kimmig. Neuro-symbolic=neural+ logical+ probabilistic. In *NeSy'19 @ IJCAI*, 2019.

30. T. Dean and M. Wellman. *Planning and control*. Morgan Kaufmann Publishers Inc., 1991.
31. L. Demey, B. Kooi, and J. Sack. Logic and probability. 2013.
32. P. Domingos, S. Kok, H. Poon, M. Richardson, and P. Singla. Unifying logical and statistical AI. In *Proc. AAAI*, pages 2–7, 2006.
33. R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.
34. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
35. D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *Pervasive Computing, IEEE*, 2(3):24–33, 2003.
36. L. Getoor and B. Taskar. Introduction to statistical relational learning (adaptive computation and machine learning). 2007.
37. H. Grosskreutz and G. Lakemeyer. ccgolog – a logical language dealing with continuous change. *Logic Journal of the IGPL*, 11(2):179–221, 2003.
38. Y. Gu and M. Soutchanski. A description logic based situation calculus. *Ann. Math. Artif. Intell.*, 58(1-2):3–83, 2010.
39. H. Hajishirzi and E. Amir. Reasoning about deterministic actions with probabilistic prior and application to stochastic filtering. In *Proc. KR*, 2010.
40. J. Y. Halpern and M. R. Tuttle. Knowledge, probability, and adversaries. *J. ACM*, 40:917–960, 1993.
41. A. Herzig, J. Lang, D. Longin, and T. Polacsek. A logic for planning under partial observability. In *Proc. AAAI/IAAI*, pages 768–773, 2000.
42. Y. Hu and G. De Giacomo. A generic technique for synthesizing bounded finite-state controllers. In *ICAPS*, 2013.
43. D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
44. R. F. Kelly and A. R. Pearce. Complex epistemic modalities in the situation calculus. In *KR*, 2008.
45. N. Kushmerick, S. Hanks, and D. Weld. An algorithm for probabilistic planning. *Artificial Intelligence*, 76(1):239–286, 1995.
46. G. Lakemeyer and Y. Lespérance. Efficient reasoning in multiagent epistemic logics. In *Proc. ECAI*, pages 498–503, 2012.
47. G. Lakemeyer and H. J. Levesque. Evaluation-based reasoning with disjunctive information in first-order knowledge bases. In *Proc. KR*, pages 73–81, 2002.
48. G. Lakemeyer and H. J. Levesque. Cognitive robotics. In *Handbook of Knowledge Representation*, pages 869–886. Elsevier, 2007.
49. J. Lang and B. Zanuttini. Probabilistic knowledge-based programs. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
50. T. Lang, M. Toussaint, and K. Kersting. Exploration in relational domains for model-based reinforcement learning. *Journal of Machine Learning Research (JMLR)*, 13(Dec):3691–3734, 2012.
51. S. Lemaignan, R. Ros, L. Mösenlechner, R. Alami, and M. Beetz. Oro, a knowledge management platform for cognitive architectures in robotics. In *IROS*, 2010.
52. U. Lerner, B. Moses, M. Scott, S. McIlraith, and D. Koller. Monitoring a complex physical system using a hybrid dynamic bayes net. In *Proc. UAI*, pages 301–310, 2002.
53. H. Levesque and R. Reiter. High-level robotic control: Beyond planning. Position paper at AAAI Spring Symposium on Integrating Robotics Research, 1998.
54. H. Levesque, R. Reiter, Y. Lespérance, F. Lin, and R. Scherl. Golog: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31:59–84, 1997.
55. H. J. Levesque. What is planning in the presence of sensing? In *Proc. AAAI/IAAI*, pages 1139–1146, 1996.

56. H. J. Levesque. A completeness result for reasoning with incomplete first-order knowledge bases. In *Proc. KR*, pages 14–23, 1998.
57. H. J. Levesque and G. Lakemeyer. *The logic of knowledge bases*. The MIT Press, 2001.
58. A. Levray and V. Belle. Learning tractable credal networks. In *AKBC*, 2020.
59. F. Lin and H. J. Levesque. What robots can do: Robot programs and effective achievability. *Artif. Intell.*, 101(1-2):201–226, 1998.
60. F. Lin and R. Reiter. How to progress a database. *Artificial Intelligence*, 92(1-2):131–167, 1997.
61. Y. Liu and G. Lakemeyer. On first-order definability and computability of progression for local-effect actions and beyond. In *Proc. IJCAI*, pages 860–866, 2009.
62. Y. Liu and H. Levesque. Tractable reasoning with incomplete first-order knowledge in dynamic systems with context-dependent actions. In *Proc. IJCAI*, pages 522–527, 2005.
63. Y. Liu and X. Wen. On the progression of knowledge in the situation calculus. In *IJCAI*, 2011.
64. J. McCarthy. Programs with common sense. In *Semantic Information Processing*, pages 403–418. MIT Press, 1968.
65. J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, pages 463–502, 1969.
66. B. Milch, B. Marthi, S. J. Russell, D. Sontag, D. L. Ong, and A. Kolobov. BLOG: Probabilistic models with unknown objects. In *Proc. IJCAI*, pages 1352–1359, 2005.
67. R. Ng and V. Subrahmanian. Probabilistic logic programming. *Information and Computation*, 101(2):150–201, 1992.
68. D. Nitti. *Hybrid Probabilistic Logic Programming*. PhD thesis, KU Leuven, 2016.
69. D. Nitti, V. Belle, and L. D. Raedt. Planning in discrete and continuous markov decision processes by probabilistic programming. In *ECML*, 2015.
70. J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.
71. R. Reiter. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT Press, 2001.
72. S. J. Russell. Unifying logic and probability. *Commun. ACM*, 58(7):88–97, 2015.
73. S. Sardina, G. De Giacomo, Y. Lespérance, and H. J. Levesque. On the semantics of deliberation in indigolog—from theory to implementation. *Annals of Mathematics and Artificial Intelligence*, 41(2-4):259–299, 2004.
74. R. B. Scherl and H. J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.
75. A. Shirazi and E. Amir. First-order logical filtering. In *Proc. IJCAI*, pages 589–595, 2005.
76. S. Srivastava. *Foundations and Applications of Generalized Planning*. PhD thesis, Department of Computer Science, University of Massachusetts Amherst, 2010.
77. M. Thielscher. From situation calculus to fluent calculus: state update axioms as a solution to the inferential frame problem. *Artificial Intelligence*, 111(1-2):277–299, 1999.
78. M. Thielscher. Planning with noisy actions (preliminary report). In *Proc. Australian Joint Conference on Artificial Intelligence*, pages 27–45, 2001.
79. S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
80. S. D. Tran and L. S. Davis. Event modeling and recognition using markov logic networks. In *Proc. ECCV*, pages 610–623, 2008.
81. L. Treszkai and V. Belle. A correctness result for synthesizing plans with loops in stochastic domains. *International Journal of Approximate Reasoning*, 2020.
82. J. Van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96, 2009.
83. F. Van Harmelen, V. Lifschitz, and B. Porter. *Handbook of knowledge representation*. Elsevier, 2008.