

Reasoning about Beliefs and Meta-Beliefs by Regression in an Expressive Probabilistic Action Logic

Daxin Liu*, Gerhard Lakemeyer
RWTH Aachen University
{liu, gerhard}@kbsg.rwth-aachen.de

Abstract

In a recent paper Belle and Lakemeyer proposed the logic \mathcal{DS} , a probabilistic extension of a modal variant of the situation calculus with a model of belief based on weighted possible worlds. Among other things, they were able to precisely capture the beliefs of a probabilistic knowledge base in terms of the concept of only-believing. While intuitively appealing, the logic has a number of shortcomings. Perhaps the most severe is the limited expressiveness in that degrees of belief are restricted to constant rational numbers, which makes it impossible to express arbitrary belief distributions. In this paper we will address this and other shortcomings by extending the language and modifying the semantics of belief and only-believing. Among other things, we will show that belief retains many but not all of the properties of \mathcal{DS} . Moreover, it turns out that only-believing arbitrary sentences, including those mentioning belief, is uniquely satisfiable in our logic. For an interesting class of knowledge bases we also show how reasoning about beliefs and meta-beliefs after performing noisy actions and sensing can be reduced to reasoning about the initial beliefs of an agent using a form of regression.

1 Introduction

Because of its expressiveness, first-order logic (FOL) has played a key role in knowledge representation and reasoning since the beginnings of AI [McCarthy, 1963]. Since uncertainty is ubiquitous in many domains like robotics, it seems natural to combine FOL and uncertainty, mixing both categorical beliefs and degrees of belief. Such formalisms are especially useful when reasoning about the beliefs of robots acting and perceiving in uncertain environments.

In a seminal paper, Bacchus, Halpern, and Levesque (BHL) [1999] gave a very rich account of dealing with stochastic actions and noisy sensing as an extension of Reiter’s version of the situation calculus [Raymond, 2001]. A key advantage of BHL’s account is that it allows partial spec-

ifications of (subjective) beliefs, depending on what information is actually available in a particular domain.

Belle and Lakemeyer (BL) [2017] proposed an alternative formulation of BHL’s ideas based on a modal variant of the situation calculus [Lakemeyer and Levesque, 2004], extending earlier work on static probabilistic beliefs [Belle *et al.*, 2016]. In contrast to BHL, which is purely axiomatic, BL’s logic \mathcal{DS} comes equipped with a possible-world semantics together with distributions over worlds. Here a distribution is defined in terms of an assignment of non-negative numbers to the possible worlds. An epistemic state is modeled as a set of distributions and a sentence ϕ is believed with degree r initially, roughly, if for each distribution in the epistemic state the sum of the normalized weights of the worlds that satisfy ϕ is r . Beliefs after performing a sequence of actions are then defined using the notions of *action likelihood* and *observational indistinguishability*, which captures the idea that an agent may not be able to distinguish between certain actions.

The logic has a number of interesting properties such as full introspection of beliefs. It also comes equipped with a notion of *only-believing* to capture the beliefs of a knowledge base (KB) consisting of a set of sentences with degrees of belief attached to them. Despite its appeal \mathcal{DS} also has shortcomings. Perhaps the main one is the lack of expressiveness when it comes to degrees of belief. Currently, the language allows to express degrees of belief only as constants or, more precisely, standard names representing rational numbers. However, it is clearly desirable to express that an agent believes that the values of a random variable are, say, geometrically distributed.¹ Likewise, while an epistemic state in \mathcal{DS} consists of a set of distributions, which in principle would allow an agent to have incomplete information about what distribution is the correct one, this can also not be expressed. Considering a running example drawn from from BL (Figure 1), BL allows a formula like $\mathbf{B}(h = 3 : .5) \wedge \mathbf{B}(h = 4 : .5)$ to express that the robot believes the distance h to the wall is 3 or 4 with equal degree, yet a geometric distribution with expectation 2, like $\forall u. \mathbf{B}(h = u : \mathcal{G}(u, \frac{1}{2}))$, is not expressible in BL. Similarly, it may be desirable to express that all the agent believes is that the distribution of the distance to the wall is geometric with expectation unknown.

¹Since the domain of \mathcal{DS} is countable, only discrete distributions are considered.

*Contact Author




Figure 1: a robot moving towards a wall

Another peculiar aspect of \mathcal{DS} is that the beliefs of an agent depend on properties of the actual world, that is, changing the world can affect an agent’s beliefs even though no (sensing) actions have occurred.

Finally, besides these issues, another important missing piece in BL’s work is the lack of consideration of how to perform projection, that is, reasoning about future beliefs in a dynamic setting.

In this paper, we will address these shortcomings by extending the language and modifying the semantics of \mathcal{DS} . In particular, this allows us to characterize probabilistic KBs (using only-believing), where an agent has incomplete information about belief distributions. Our main technical result is a solution to the projection problem by a novel form of regression. While regression under uncertainty has been considered in the past, we show, for the first time, how to deal with nested beliefs and quantifying into beliefs (quantifying-in).

The rest of the paper is organized as follows. In Section 2, we introduce the syntax and semantics of the new logic called \mathcal{DS}^* . In Section 3, we investigate the properties of the logic. Regression is presented in Section 4. Finally, we discuss related work in Section 5 and conclude.

2 The Logic \mathcal{DS}^*

2.1 Syntax

\mathcal{DS}^* is a first-order modal language with equality and sorts of type *object* and *action*. Implicitly, we assume that *number* is a sub-sort of object and refers to the *computable numbers* \mathbb{C} .² Logical symbols include the usual symbols $\{=, \wedge, \vee, \neg\}$ and abbreviations $\{\exists, \equiv, \supset\}$, as well as variables of every sort. There are function symbols of every sort and arity:

fluent function symbols of every arity, such as $distance(x)$, $nextTo(x, y)$, including two unary special symbols $poss$, l and a special binary symbol oi . Roughly, $poss$ specifies when an action is possible (preconditions), l returns the likelihood of an action and oi describes the *observational indistinguishability* (alternative choices) among actions;

rigid function symbols of every arity, such as $sonar(x)$, including arithmetical functions like $+$, $-$, \times , *etc.*

While fluent functions may vary due to actions, rigid functions are fixed. For simplicity, all action functions are rigid and we do not include predicate symbols in the language.

²We use the computable numbers as they are still enumerable and allow representing distributions mentioning certain real numbers such as Euler’s number e .

The language features a countable set of so-called **standard names** \mathcal{N} , which are isomorphic with a fixed universe of discourse. Roughly speaking, this amounts to having an infinite domain closure axiom together with the unique name assumption. $\mathcal{N} = \mathcal{N}_O \cup \mathcal{N}_A$ where \mathcal{N}_O and \mathcal{N}_A are object standard names and action standard names respectively.

The set of **terms** is the least set which contains the standard names \mathcal{N} , variables of sort object (u, v, x, y, \dots) and action (a, a', \dots) , and if t_1, \dots, t_k are terms, f a k -ary function symbol, then $f(t_1, \dots, t_k)$ is in the set.

Ground terms are terms without variables. *Primitive terms* are terms of the form $f(n_1, \dots, n_k)$, where f is a function symbol and n_i are object standard names. We denote the sets of primitive terms of sort object and action as \mathcal{P}_O and \mathcal{P}_A , respectively. While object standard names are syntactically like constants, we require that action standard names are all the primitive action terms, i.e. $\mathcal{N}_A = \mathcal{P}_A$. For example, the action $sonar(5)$, where a sonar returns the number 5, is considered a standard action name. For any action term $A(t_1, \dots, t_k)$, A is referred to as the action type.³ Furthermore \mathcal{Z} refers to the set of all finite sequences of action standard names, including the empty sequence $\langle \rangle$. We reserve standard names \top, \perp in \mathcal{N}_O for truth values (to simulate predicates).

The set of **formulas** is the least set which contains atoms, $(t = t')$ where t and t' are terms, and if α and β are formulas, x is a variable, r, r_i are *rigid terms* (terms without fluent functions) of sort number, t_a is of sort action, then $\neg\alpha$, $\alpha \wedge \beta$, $\forall x.\alpha$, $[t_a]\alpha$, $\Box\alpha$, $\mathbf{B}(\alpha : r)$ and $\mathbf{O}(\alpha_1 : r_1, \dots, \alpha_k : r_k)$ are in the set.

$[t_a]\alpha$ should be read as “ α holds after action t_a ” and $\Box\alpha$ as “ α holds after any sequence of actions.” The epistemic expression $\mathbf{B}(\alpha : r)$ should be read as “ α is believed with probability r ”. $\mathbf{K}\alpha$ means “ α is known” and is an abbreviation for $\mathbf{B}(\alpha : 1)$. $\mathbf{O}(\alpha_1 : r_1, \dots, \alpha_k : r_k)$ may be read as “the α_i with a probability r_i are all that is believed”. Similarly, $\mathbf{O}\alpha$ means “ α is only known” and is an abbreviation for $\mathbf{O}(\alpha : 1)$. For action sequence $z = a_1 \dots a_k$, we write $[z]\alpha$ to mean $[a_1] \dots [a_k]\alpha$. α_t^x is the formula obtained by substituting all free occurrences of x in α by t .

A *sentence* is a formula without free variables. We use TRUE as an abbreviation for $\forall x(x = x)$, and FALSE for its negation. A formula with no \Box is called *bounded*; a formula with no \Box or $[t_a]$ is called *static*; a formula with no \mathbf{B} or \mathbf{O} is called *objective*; a formula with no fluent, \Box or $[t_a]$ outside \mathbf{B} or \mathbf{O} is called *subjective*; a formula with no \mathbf{B} , \mathbf{O} , \Box , $[t_a]$, $poss$, l , oi is called a *fluent formula*; a fluent formula without fluent functions is called a *rigid formula*.

2.2 Semantics

The semantics is given in terms of *possible worlds*, which define what is true initially and after any sequence of actions. Compared to non-probabilistic accounts with deterministic actions [Lakemeyer and Levesque, 2004], a number of challenges need to be addressed, including how to specify probabilities over *uncountably* many possible worlds, how to

³An action type is a set of actions that have the same action symbol but bound to different parameters, e.g. $sonar(1)$, $sonar(2)$ are of the same action type

deal with multiple probability distributions entertained by the agent, and how to deal with probabilistic action effects, which may be *unobservable* by the agent.

A **world** w is a mapping from the primitive terms ($\mathcal{P}_O \cup \mathcal{P}_A$) and \mathcal{Z} to \mathcal{N} of the right sort, satisfying:

1. **Rigidity:** If t is a rigid primitive term, then for all $(w, z), (w', z'), w[t, z] = w'[t, z']$;
2. **Arithmetical Correctness:** If f is an arithmetical expression and val is its value in the usual sense, then for all $(w, z), w[f, z] = val$. For example, $w[1 + 1, z] = 2$.

We will sometimes also use piecewise rigid functions such as $f(x) = .3$ if $x \in \{3, 4, 5\}$ and 0 o.w. We assume that these are simply built into the logic as well.⁴

Let \mathcal{W} be the set of all such worlds. Free variables are handled substitutionally by using standard names. We now define the notion of co-referring standard name for ground terms. Given a ground term t , a world w , and action sequence z , we define $|t|_w^z$ (read: the co-referring standard name for t given w, z) recursively by :

1. If $t \in \mathcal{N}$, then $|t|_w^z = t$;
2. $|f(t_1, \dots, t_k)|_w^z = w[f(|t_1|_w^z, \dots, |t_k|_w^z), z]$.

For a rigid term t , we use $|t|$ instead of $|t|_w^z$ for its denotation. We will require that $l(a)$ is of sort number, $poss(a)$ and $oi(a, a')$ only take values \top or \perp , and oi behaves like an equivalence relation (reflexive, symmetric, and transitive). $oi(a, a')$ means a and a' are observationally indistinguishable actions. In the example of Fig. 1, the robot might perform a stochastic action $fwd(x, y)$, where x is its intended forward distance and y is the actual outcome selected by nature. x is observable to the robot while y is not. Then, $oi(fwd(1, 1.1), fwd(1, 0.9))$ says that nature can non-deterministically select 1.1 or 0.9 as a result for the intended value 1.

By a **distribution** d we mean a mapping from \mathcal{W} to $\mathbb{C}^{\geq 0}$ (non-negative computable numbers) and an **epistemic state** e is any set of distributions. By a model, we mean a triple (e, w, z) . In order to prepare for the semantics, we need to extend $l(a), oi(a, a'), poss(a)$ from actions to action sequences:

Definition 1. We define

1. $l^*: \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{C}^{\geq 0}$ as
 $l^*(w, \langle \rangle) = 1$, for every $w \in \mathcal{W}$;
 $l^*(w, z \cdot a) = l^*(w, z) \times n$ where $w[l(a), z] = n$;
2. $z \sim_w z'$ as
 $\langle \rangle \sim_w z'$ iff $z' = \langle \rangle$;
 $z \cdot a \sim_w z'$ iff $z' = z^* \cdot a^*, z \sim_{z^*},$
 $oi(a, a^*), z] = \top$;
3. $exec(z)$ as
for $z = \langle \rangle, exec(z) \equiv \text{TRUE}$;
for $z = z' \cdot a, exec(z) \equiv poss(a) = \top \wedge [a]exec(z')$.

Obviously, \sim_w is an equivalence relation. Since \mathcal{W} is uncountable, to obtain a well-defined sum over uncountably many worlds, the following three conditions are used for evaluating beliefs below:

⁴Alternatively, one could add them as additional assumptions whenever needed.

Definition 2. We define **BND**, **EQ**, **NORM** for any distribution d and any set $\mathcal{V} = \{(w_1, z_1), (w_2, z_2), \dots\}$ as follows:

1. **BND** (d, \mathcal{V}, n) iff $\neg \exists k, (w_1, z_1), \dots, (w_k, z_k) \in \mathcal{V}$ such that $\sum_{i=1}^k d(w_i) \times l^*(w_i, z_i) > n$.
2. **EQ** (d, \mathcal{V}, n) iff **BND** (d, \mathcal{V}, n) and there is no $n' < n$ such that **BND** (d, \mathcal{V}, n') holds.
3. for any $\mathcal{U} \subseteq \mathcal{V}$, **NORM** $(d, \mathcal{U}, \mathcal{V}, n)$ iff $\exists b \neq 0$ such that **EQ** $(d, \mathcal{U}, b \times n)$ and **EQ** (d, \mathcal{V}, b) .

Intuitively, given **NORM** $(d, \mathcal{U}, \mathcal{V}, n)$, n can be viewed as the *normalized* sum of the weights of worlds in \mathcal{U} wrt d in relation to \mathcal{V} . Here **EQ** (d, \mathcal{V}, b) expresses that the weight of the worlds wrt d in \mathcal{V} is b , and finally **BND** (d, \mathcal{V}, b) ensures the weights of worlds in \mathcal{V} is bounded by b . In essence, even if \mathcal{W} is uncountable, the condition **EQ** and **BND** admit a well-defined summation of the weights on worlds. See [Belle et al., 2016] for a formal justification of this claim.

The truth of sentences in \mathcal{DS}^* is defined as:

- $e, w, z \models t_1 = t_2$ iff $|t_1|_w^z$ and $|t_2|_w^z$ are identical;
- $e, w, z \models \neg \alpha$ iff $e, w, z \not\models \alpha$;
- $e, w, z \models \alpha \wedge \beta$ iff $e, w, z \models \alpha$ and $e, w, z \models \beta$;
- $e, w, z \models \forall x. \alpha$ iff $e, w, z \models \alpha_n^x$ for every standard name n of the right sort;
- $e, w, z \models [t_a]\alpha$ iff $e, w, z \cdot n \models \alpha$ and $n = |t_a|_w^z$;
- $e, w, z \models \Box \alpha$ iff $e, w, z \cdot z' \models \alpha$ for all $z' \in \mathcal{Z}$.

To prepare for the semantics of epistemic operators, let $\mathcal{W}_\alpha^{e, z} = \{(w', z') \mid z' \sim_{w'} z, \text{ and } e, w', \langle \rangle \models [z']\alpha \wedge exec(z')\}$. If $z = \langle \rangle$, we write \mathcal{W}_α^e and if the context is clear, we write \mathcal{W}_α . Intuitively, \mathcal{W}_α is the set of alternatives (world and action sequence pairs) of z that might result in α .

Finally, after actions, we will restrict ourselves to compatible epistemic states that agree on action likelihood:

Definition 3. A distribution d is **compatible** with action sequence z , $d \sim_{comp} z$ iff **NORM** $(d, \mathcal{W}_{\text{TRUE}}^{e, z}, \mathcal{W}_{\text{TRUE}}^{\{d\}, z}, 1)$; A distribution d is **regular** iff $d \sim_{comp} \langle \rangle$.

Essentially, there are two cases where a distribution d may be incompatible with an action sequence z . First, **EQ** $(d, \mathcal{W}_{\text{TRUE}}^{\{d\}, z}, 0)$, i.e. z has zero-likelihood in d . Second, $\neg \text{BND}(d, \mathcal{W}_{\text{TRUE}}^{\{d\}, z}, r)$ for all r . This may happen, for example, if d assigns equal weights to all worlds in \mathcal{W} . Distributions of this sort are called *unmeasurable* [Belle et al., 2016]. Regular distributions are just the measurable ones.

Let (e, w, z) be a model and $e_z = \{d \mid d \in e, d \sim_{comp} z\}$. We are now ready to define the semantics of **B** and **O**:

- $e, w, z \models \mathbf{B}(\alpha : r)$ iff
 $\forall d \in e_z, \text{NORM}(d, \mathcal{W}_\alpha^{\{d\}, z}, \mathcal{W}_{\text{TRUE}}^{\{d\}, z}, n)$ and $n = |r|$;
- $e, w, z \models \mathbf{O}(\alpha_1 : r_1, \dots, \alpha_k : r_k)$ iff $\forall d, d \in e_z$ iff
 $\text{NORM}(d, \mathcal{W}_{\alpha_i}^{\{d\}, z}, \mathcal{W}_{\text{TRUE}}^{\{d\}, z}, n_i)$ for all $1 \leq i \leq k$ where
 $n_i = |r_i|$.

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 set of regular distributions e and w , if $e, w \models \alpha'$ for every

$\alpha' \in \Sigma$, then $e, w \models \alpha$. We say that α is valid ($\models \alpha$) if $\{\} \models \alpha$. Satisfiability is then defined in the usual way. If α is an objective formula, we write $w \models \alpha$ instead of $e, w \models \alpha$. Similarly, we write $e \models \alpha$ instead of $e, w \models \alpha$ if α is subjective.

2.3 Comparison with \mathcal{DS}

The language of \mathcal{DS} differs slightly from \mathcal{DS}^* in that \mathcal{DS} considers fluent predicates instead of fluent functions. Moreover, while in \mathcal{DS} every closed term is a standard name, we follow [Lakemeyer and Levesque, 2011], who use special standard names for objects and actions. The semantic structures in both logics are essentially the same, consisting of worlds, sets of distributions over worlds serving as epistemic states, and action sequences. Our use of rigid mathematical functions, which are not considered in \mathcal{DS} , is similar to the \mathcal{R} -interpretation in [Belle and Levesque, 2018].

Since the semantic rules for \mathbf{B} and \mathbf{O} in both logics are similar, but differ in important details, it is instructive to review their definitions in \mathcal{DS} . The main difference lies in how \mathcal{W}_α is defined and used. \mathcal{DS} uses the notion of compatibility between worlds: $w' \approx_{oi} w$ iff for all a, a' and z , $w'[oi(a, a'), z] = w[oi(a, a'), z]$, that is, w' is compatible with w if they agree on the values of oi . \mathcal{W}_α in \mathcal{DS} is then defined wrt. triples e, w, z : $\mathcal{W}_\alpha^{e, w, z} = \{(w', z') \mid z' \sim_w w', z, w' \approx_{oi} w, \text{ and } e, w', \langle \rangle \models [z']\alpha \wedge \text{exec}(z')\}$. Note that worlds in $\mathcal{W}_\alpha^{e, w, z}$ agree with w wrt. observational indistinguishability. Belief is then defined as $e, w, z \models_{\mathcal{DS}} \mathbf{B}(\alpha : n)$ iff $\forall d \in e$ (not e_z), $\text{NORM}(\mathcal{W}_\alpha^{e, w, z}, \mathcal{W}_{\text{TRUE}}, n)$ for constant n (and similarly for \mathbf{O}).

The first thing to note is that our semantics of \mathbf{B} (and \mathbf{O}) is independent of the real word, while \mathcal{DS} is not. In particular, the truth of $\mathbf{B}(\alpha : n)$ depends on both e and w because of the definition of $\mathcal{W}_\alpha^{e, w, z}$. As a result, even if e is kept fixed, once w changes, beliefs can vary. For example, consider a fluent *lightOn*, which indicates whether a light is on or off. In \mathcal{DS} , using functions as in \mathcal{DS}^* instead of predicates, we obtain:

Example 1. Let Σ_{oi} be $\Box oi(a, a') = \top \equiv \text{lightOn} = \perp \vee \text{lightOn} = \top \wedge a = a'$.⁵ Let w, w' be two worlds such that $w \models_{\mathcal{DS}} \text{lightOn} = \top$ and $w' \models_{\mathcal{DS}} \text{lightOn} = \perp$. Additionally, suppose $e, w \models_{\mathcal{DS}} \mathbf{K}\Sigma_{oi}$ and $e, w' \models_{\mathcal{DS}} \mathbf{K}\Sigma_{oi}$. Then it follows that $e, w \models_{\mathcal{DS}} \mathbf{K}(\text{lightOn} = \top)$ and $e, w' \models_{\mathcal{DS}} \mathbf{K}(\text{lightOn} = \perp)$.

In English: Σ_{oi} says actions are mutually distinguishable if the light is on. There are two worlds w, w' where the light is on and off, respectively. The agent with epistemic state e knows Σ_{oi} in both worlds. Then it follows that in world w the agent knows that the light is on while in w' the agent knows that the light is off. This seems counter-intuitive, especially since the agent has not performed any actions.

This is the main reason why in \mathcal{DS}^* the definition of \mathcal{W}_α does not refer to the real world. We remark for the same reason we require the r in $\mathbf{B}(\alpha : r)$ to be rigid and hence world-independent (and similarly for \mathbf{O}). But note that r can oth-

⁵Free variables are implicitly universal quantified outside. The modality \Box has lower syntactic precedence than the connectives, and $[\cdot]$ has the highest priority.

erwise be an arbitrary mathematical function, allowing us to express things like distributions as in $\forall u. \mathbf{B}(h = u : \mathcal{G}(u, \frac{1}{2}))$.

Lastly, notice that in the definition of \mathbf{B} and \mathbf{O} , we use $\mathcal{W}_\alpha^{\{d\}, z}$ instead of $\mathcal{W}_\alpha^{e, z}$ (or $\mathcal{W}_\alpha^{e, w, z}$ as in \mathcal{DS}^*). While this makes no difference when it comes to believing objective sentences, it does impact beliefs about beliefs in important ways, in particular introspection, as discussed below. Perhaps the greatest benefit of this change is that only-believing a sentence becomes uniquely satisfiable in \mathcal{DS}^* . Among other things, this allows us to model agents which are unsure about their beliefs. For example, if p is an atomic sentence, then $\mathbf{O}(\mathbf{B}(p : .3) \vee \mathbf{B}(p : .4))$ is satisfied by the set of all regular distributions where the degree of belief in p is either .3 or .4. Yet, in \mathcal{DS} the above sentence is unsatisfiable.

3 Properties of Belief and Only-believing

The new logic shares many properties with \mathcal{DS} like additivity, and the Barcan formula. Formally, we have:

Proposition 1. for any formula α, β and rigid terms r, r', r'' :

1. If $\models \Box(\alpha = \beta)$ then $\models (\mathbf{B}(\alpha : r) \equiv \mathbf{B}(\beta : r))$;
2. $\models \Box(\mathbf{B}(\alpha \wedge \beta : r) \wedge \mathbf{B}(\alpha \wedge \neg\beta : r') \supset \mathbf{B}(\alpha : r + r'))$;
3. $\models \Box(\mathbf{B}(\alpha : r) \wedge \mathbf{B}(\beta : r') \wedge \mathbf{B}(\alpha \wedge \beta : r'') \supset \mathbf{B}(\alpha \vee \beta : r + r' - r''))$;
4. $\models \Box(\mathbf{K}\alpha \wedge \mathbf{K}(\alpha \supset \beta) \supset \mathbf{K}\beta)$;
5. $\models \Box(\forall x \mathbf{K}\alpha \supset \mathbf{K}\forall x\alpha)$;
6. $\models \Box(\exists x \mathbf{K}\alpha \supset \mathbf{K}\exists x\alpha)$.

Introspection, which is valid in \mathcal{DS} , is not valid in \mathcal{DS}^* in general. However, as we will see below, it does hold under certain assumptions.

3.1 Introspection

In the static case, for positive introspection, we have:

Theorem 1. For any α and r , $\models \mathbf{B}(\alpha : r) \supset \mathbf{K}(\mathbf{B}(\alpha : r))$.

Negative introspection like $\neg\mathbf{B}(\alpha : r) \supset \mathbf{K}(\neg\mathbf{B}(\alpha : r))$ does not hold. For example, let $\{d_1\} \models \mathbf{B}(\alpha : r)$ and $\{d_2\} \not\models \mathbf{B}(\alpha : r)$, then for $e = \{d_1, d_2\}$, we have $e \models \neg\mathbf{B}(\alpha : r)$. Yet, $e \not\models \mathbf{K}\neg\mathbf{B}(\alpha : r)$ due to $\mathcal{W}_{\neg\mathbf{B}(\alpha : r)}^{\{d_1\}} = \emptyset$ and $\text{NORM}(d_1, \mathcal{W}_{\neg\mathbf{B}(\alpha : r)}^{\{d_1\}}, \mathcal{W}_{\text{TRUE}}^{\{d_1\}}, 0)$.

Besides, $\Box(\mathbf{B}(\alpha : r) \supset \mathbf{K}(\mathbf{B}(\alpha : r)))$ is not valid in general. This is because we allow worlds with different observational indistinguishability in \mathcal{W}_α , and as pointed out in \mathcal{DS} , the consistency of observational indistinguishability between worlds and actions in \mathcal{W}_α (\approx_{oi}) are crucial for introspection. However, when oi is rigid, positive introspection holds.

$$\Box oi(a, a') = \top \equiv \psi \quad (1)$$

where ψ is a rigid formula. Then we have:

Theorem 2. Let Σ_{oi} be as above. Then for any α and r , $\mathbf{K}(\Sigma_{oi}) \models \Box\mathbf{B}(\alpha : r) \supset \mathbf{K}(\mathbf{B}(\alpha : r))$

Let us now consider a case where negative introspection also holds. By a belief distribution we mean the joint distribution of a finite set of random variables. Formally, assuming a finite set of nullary fluents $\vec{h} = \{h_1, \dots, h_m\}$,⁶ a belief distribution \mathbf{B}_{init}^f of \vec{h} is a formula of the form $\forall \vec{u}. \mathbf{B}(\vec{h} = \vec{u} : f(\vec{u}))$, where \vec{u} is a set of variables, $\vec{h} = \vec{u}$ stands for $\bigwedge h_i = u_i$, f is a rigid function of sort number with free variables \vec{u} . Taking our robot example, a belief distribution might be \mathbf{B}_{init}^f (for the single fluent h), where f is given by

$$f(u) = \begin{cases} \frac{1}{3} & u \in \{3, 4, 5\} \\ 0 & o.w. \end{cases} \quad (2)$$

Such belief distributions will play an important role later.

Lemma 1. *Let \mathbf{B}_{init}^f be a joint belief distribution and α a subjective static formula whose fluents are all in \mathbf{B}_{init}^f , then $\mathbf{B}_{init}^f \models \alpha$ or $\mathbf{B}_{init}^f \models \neg\alpha$.*

The proof is by induction on the size of formula α . The lemma says, roughly, that \mathbf{B}_{init}^f determines completely what the agent believes and does not believe.

Theorem 3. *For any α, r , let \mathbf{B}_{init}^f be a belief distribution of fluents in α and Σ_{oi} as before:*

- $\mathbf{B}_{init}^f \models \neg\mathbf{B}(\alpha : r) \supset \mathbf{K}(\neg\mathbf{B}(\alpha : r))$;
- $\mathbf{B}_{init}^f \wedge \mathbf{K}(\Sigma_{oi}) \models \square\neg\mathbf{B}(\alpha : r) \supset \mathbf{K}(\neg\mathbf{B}(\alpha : r))$.

The first part holds as a direct consequence of Lemma 1: once an epistemic state satisfies \mathbf{B}_{init}^f and $\neg\mathbf{B}(\alpha : r)$, then all distributions considered must satisfy $\neg\mathbf{B}(\alpha : r)$, hence $\mathbf{K}(\neg\mathbf{B}(\alpha : r))$ is satisfied. The proof of the second is rather similar to its counter-part of the positive introspection.

3.2 Only-believing

Let us now turn to the notion of only-believing, which is intended to capture the beliefs of an agent's probabilistic KB.

Proposition 2. *Let α_i be arbitrary sentences and r_i arbitrary rigid terms of sort number, we have*

1. $\models \mathbf{O}(\alpha_1 : r_1; \dots \alpha_k : r_k) \supset \bigwedge \mathbf{B}(\alpha_i : r_i)$;
2. $\models \mathbf{O}(\alpha : r) \supset \neg\mathbf{B}(h(\vec{n}) = m : r')$ for all r, r' , and α , where \vec{n} and m are std. names and h is a fluent not in α ;
3. $\models \mathbf{O}(\alpha_1 : r_1; \alpha_2 : r_2) \equiv \mathbf{O}(\mathbf{B}(\alpha_1 : r_1) \wedge \mathbf{B}(\alpha_2 : r_2))$.

The second part says that the agent has no beliefs about things not mentioned in the KB. Note that this is not true if \mathbf{O} is replaced by \mathbf{B} . The third part says only-believing reduces to only-believing a single subjective sentence with certainty.

Theorem 4 (Unique Model Theorem). *For arbitrary α_i and r_i , there is a unique set of regular distributions e such that $e \models \mathbf{O}(\alpha_1 : r_1; \dots \alpha_k : r_k)$.*

⁶As discussed in [Belle and Levesque, 2018], allowing fluents with arguments would result in joint distribution over infinitely many random variables, which is generally problematic in probability theory.

Proof. Let $e = \{d \mid \text{NORM}(d, \mathcal{W}_{\alpha_i}^{\{d\}}, \mathcal{W}_{\text{TRUE}}^{\{d\}}, n_i), n_i = |r_i|, 1 \leq i \leq k\}$. Clearly $e \models \mathbf{O}(\alpha_1 : r_1; \dots \alpha_k : r_k)$, and there cannot be any other e with that property. \square

While only-believing is always uniquely satisfied by an e , it may be the case that e is empty. For example, $\emptyset \models \mathbf{O}(p \wedge \neg p : .6)$ and $\emptyset \models \mathbf{O}(\mathbf{O}(p))$. For the latter this is because no single distribution only-believes p .

Corollary 1. $\models \mathbf{O}(\alpha : r) \supset \mathbf{B}(\beta : r')$ or $\models \mathbf{O}(\alpha : r) \supset \neg\mathbf{B}(\beta : r')$ for arbitrary α, β and all r, r' .

We remark that in \mathcal{DS} the unique-model property for \mathbf{O} only holds for objective sentences. In the case of non-objective sentences, there may be multiple epistemic states that only-believe them or none at all. For example and as mentioned earlier, $\mathbf{O}(\mathbf{B}(p : .3) \vee \mathbf{B}(p : .4))$ is not satisfiable in \mathcal{DS} . Such feature makes \mathcal{DS}^* much more expressive than \mathcal{DS} and many other probabilistic formalism. For example, $\mathbf{O}(\exists x. \forall u. \mathbf{B}(h = u : \mathcal{G}(u, x)))$ says the agent only-believes that the distance is geometricly distributed.

4 Reasoning by Regression

An important reasoning task for the agent is to determine what holds after actions. Regression is a powerful technique to solve this problem by transforming a query about the future to a query about the initial state [Raymond, 2001]. In this section, we first introduce the notion of basic action theories (BATs) that are required for regression. Thereafter, we explore how regression works in the setting of beliefs.

4.1 Basic Action Theories

BATs were first introduced by Reiter [2001] to encode the effects and preconditions of actions. In a dynamic uncertain setting, BATs additionally have to specify the observational indistinguishability among actions and likelihoods of actions. Formally, given finite action types A_1, \dots, A_k , $\Sigma = \Sigma_{pre} \cup \Sigma_{post} \cup \Sigma_{oi} \cup \Sigma_l$ consists of the union of these sets:

- Σ_{pre} : A set consisting of a single axiom of the form $\square_{poss}(a) = \top \equiv \bigvee_i \exists \vec{x}_i. a = A_i(\vec{x}_i) \wedge \pi_i$ to characterize action preconditions.
- Σ_{post} : A set of successor state axioms, one for each fluent, of the form $\square[a]h = u \equiv \bigvee_i \exists \vec{x}_i. a = A_i(\vec{x}_i) \wedge \gamma_i \vee \bigwedge_i \forall \vec{x}_i. a \neq A_i(\vec{x}_i) \wedge h = u$ to characterize action effects, also providing a solution to the frame problem [Raymond, 2001].
- Σ_l : A single axiom of the form $\square l(a) = q \equiv \bigvee_i \exists \vec{x}_i. a = A_i(\vec{x}_i) \wedge \phi_i$ to capture action likelihoods.

Here π_i, γ_i, ϕ_i are fluent formulas with free variables among \vec{x}_i . γ_i (respectively, ϕ_i) has free variable u (q) and are functional in u (q). Moreover, Σ_{oi} consists of a single axiom as in Eq. 1, which forces oi to be rigid (Σ_{oi} is also in form of disjunctions wrt action types A_1, \dots, A_k).

Example 2. *In our robot example, besides the stochastic action $fwd(x, y)$, assume the robot can perform a noisy sensing sonar(z) to locate itself, where z denotes the observed value*

of sonar. Then the following is a possible BAT:

$$\begin{aligned} \Box poss(a) &= \top \equiv \text{TRUE} \\ \Box [a]h = u &\equiv \exists x, y. a = fwd(x, y) \wedge u = h - y \\ &\quad \vee \forall x, y. a \neq fwd(x, y) \wedge h = u \\ \Box oi(a, a') &= \top \equiv \exists x, y, z. a = fwd(x, y) \wedge a' = fwd(x, z) \\ &\quad \vee a = sonar(z) \wedge a' = a \\ \Box l(a) = q &\equiv \exists x, y, z. a = fwd(x, y) \wedge q = \theta(x, y, .1, .8) \\ &\quad \vee a = sonar(z) \wedge q = \theta(z, h, .2, .6), \text{ where} \end{aligned}$$

$$\theta(x, y, m, n) = \begin{cases} m & |x - y| = 1 \\ n & x = y \\ 0 & o.w. \end{cases}$$

In English: actions are always executable; distance h can only be affected by $fwd(x, y)$ and the value is determined by value y , not the intended value x ; two actions are observationally indistinguishable if and only if they are both forward actions with the same intended value or they are identical sensing action; action likelihood is specified by θ .

4.2 Regressing Objective Formulas

The regression of objective formulas is similar to [Claßen, 2013]. A formula where each quantifier has a distinct variable is called *rectified*. Let \mathcal{T}_Σ be the set of action terms that can be generated by the finite action types of BAT Σ , i.e. $\mathcal{T}_\Sigma = \cup_i \{A_i(\vec{r}) \mid \vec{r} \text{ are standard names or variables}\}$. In the rest of the paper, action terms and action sequences refer to terms in \mathcal{T}_Σ and sequences in $(\mathcal{T}_\Sigma)^*$.

An objective formula is called *regressible* if it is bounded, rectified and every function fluent has the form $h = r'$, where r' does not mention any fluent. In addition every *poss* and *oi* is of the form $poss(t) = \top$, $oi(t, t') = \top$ s.t. $t, t' \in \mathcal{T}_\Sigma$.⁷ Apart from being bounded, all other syntactic restrictions on regressible formulas are only introduced for convenience in order to simplify the following definition. Unbounded formulas such as $\Box \alpha$ are not regressible because \Box represents an infinite number of action sequences.

We define $\mathfrak{R}[\alpha]$, the regression of α wrt Σ , to be $\mathfrak{R}[\langle \rangle, \alpha]$, with $\mathfrak{R}[z, \alpha]$ for any action sequence $z \in (\mathcal{T}_\Sigma)^*$ defined as:

- $\mathfrak{R}[z, \alpha] = \alpha$ if α is rigid;
- $\mathfrak{R}[z, \neg \alpha] = \neg \mathfrak{R}[z, \alpha]$, and similarly for \wedge, \vee ;
- $\mathfrak{R}[z, [A_i(\vec{r})]\alpha] = \mathfrak{R}[z \cdot A_i(\vec{r}), \alpha]$;
- $\mathfrak{R}[z, poss(A_i(\vec{r})) = \top] = \mathfrak{R}[z, (\pi_i)_{\vec{r}}^{\vec{u}_i}]$, π_i is from Σ_{pre} 's RHS, and similarly for l and oi ;
- $\mathfrak{R}[z, h = r']$ is defined inductively:
 1. $\mathfrak{R}[\langle \rangle, h = r'] = (h = r')$;
 2. $\mathfrak{R}[z \cdot A_i(\vec{r}), h = r] = \mathfrak{R}[z, (\gamma_i)_{\vec{r}, r'}^{\vec{u}_i, u}]$, where γ_i is from Σ_{post} 's RHS.

Note that $\mathfrak{R}[\alpha]$ always results in a fluent formula if α is regressible.

⁷It is easy to see that every formula can be rewritten this way, e.g. $\models h = h' \equiv \exists x. h = x \wedge h' = x$, for fluents h, h' .

Theorem 5 (Regression Theorem). *Let \mathcal{D} be a set of fluent sentences (describing the initial state), Σ a BAT, α a regressible sentence. Then $\mathcal{D} \cup \Sigma \models \alpha$ iff $\mathcal{D} \models \mathfrak{R}[\alpha]$.*

The proof follows the same argument as the proof a corresponding theorem in [Claßen, 2013]. As an illustration, consider the BAT Σ of the robot example. $\mathfrak{R}[[fwd(1, 1)]h = 10] = (h = 9)$. So to determine whether $[fwd(1, 1)]h = 10$ is entailed by some $\mathcal{D} \cup \Sigma$, we only need to check whether \mathcal{D} entails $h = 9$.

4.3 Regressing Beliefs

To reason about future beliefs, besides basic action theories, we need to specify what is believed initially. By the notion of a belief distribution, our initial KB is expressed via a single sentence of the form $O(B_{init}^f \wedge \Sigma)$, where B_{init}^f is a belief distribution of fluents mentioned by BAT Σ .

The regression operator now takes a single belief distribution B_{init}^f as a parameter, namely $\mathfrak{R}[B_{init}^f, z, \alpha]$.

A formula is called *regressible* wrt B_{init}^f if all fluents are mentioned in B_{init}^f , it does not mention O , and otherwise satisfies the constraints for objective regressible formulas.

Taking the B_{init}^f of the robot example (see Eq. 2), $[fwd(1, 1)]K(h = 1)$ and $[sonar(2)]K([fwd(1, 1)]B(h = 3 : \frac{1}{2}))$ are regressible, while $\Box K(h = 1)$ and $B(h = h' : .6)$, are not, assuming h, h' are fluents. If the context is clear, we leave out B_{init}^f . Given a BAT Σ and $z = t_1 \cdot t_2 \cdot \dots \cdot t_k$ where $t_i \in \mathcal{T}_\Sigma$, if B_{init}^f is *compatible* with z, Σ (Def. 4 below), the regression of belief is defined as:

$$\mathfrak{R}[z, B(\alpha : r)] = (r = \frac{1}{\eta} \sum_{\vec{u}} f(\vec{u}) \sum_{z' \in \Psi} \mathfrak{T}[z', \vec{u}] \mathfrak{J}(\vec{u}, \mathfrak{R}[z', \alpha]))$$

where η is the normalizer, which has the same form as the numerator but replacing α with TRUE , $\sum_{z' \in \Psi}$ is an abbreviation of $\sum_{t'_1 \in \{t'_1 \mid \models \psi_1\}} \dots \sum_{t'_k \in \{t'_k \mid \models \psi_k\}}$, and ψ_i is the RHS of Σ_{oi} instantiated by t_i and t'_i . $z' = [t'_1] \dots [t'_k]$. \vec{u} ranges over $(\mathcal{N}_O)^m$ where m is the number of fluents and t'_i ranges over \mathcal{N}_A .⁸

$\mathfrak{T}[z, \vec{u}]$ is recursively defined as:

1. $\mathfrak{T}[\langle \rangle, \vec{u}] = 1$;
2. $\mathfrak{T}[z \cdot t, \vec{u}] = LH(z, t, \vec{u}) \times \mathfrak{T}[z, \vec{u}]$,

Where $\mathfrak{J}(\vec{u}, \alpha) = IF \alpha_{\vec{u}}^{\vec{h}} THEN 1 ELSE 0, LH(z, t, \vec{u}) = \lambda(\mathfrak{R}[z, l(t) = q_t]_{\vec{u}}^{\vec{h}})$ and λ is a function that extracts the value of q_t after regression. Essentially $LH(z, t, \vec{u})$ returns the likelihoods of action t after action sequence z , $\mathfrak{T}[z, \vec{u}]$ returns the likelihoods of z when the initial state satisfies $\vec{h} = \vec{u}$, and $\mathfrak{J}(\vec{u}, \alpha)$ indicates whether α follows from the initial state $\vec{h} = \vec{u}$. For space reasons, we leave out the details of how to specify λ , but the ideas are essentially the same as in [Belle and Levesque, 2020]. See the regression ex. below.

Definition 4. *Belief distribution B_{init}^f is compatible with z, Σ iff $\models \sum_{\vec{u}} f(\vec{u}) \times \sum_{z' \in \Psi} (\mathfrak{T}[z', \vec{u}]) \neq 0$.*

⁸Note that summation is treated here like any other rigid mathematical function for simplicity. It can be defined using second-order logic as shown in [Belle and Levesque, 2018], for example.

Basically, f is *compatible* with z, Σ if and only if z has non-zero likelihood under f and Σ . In case \mathbf{B}_{init}^f is not compatible with z, Σ , we define: $\mathfrak{R}[z, \mathbf{B}(\alpha : r)] \equiv \text{TRUE}$

Theorem 6. Let $\mathcal{D}, \Sigma, \mathbf{B}_{init}^f$ be as before, then for all regressable formulas α ,

$$\mathcal{D} \cup \Sigma \cup \mathcal{O}(\mathbf{B}_{init}^f \wedge \Sigma) \models \alpha \text{ iff } \mathcal{D} \models \mathfrak{R}[\alpha].$$

Example 3. Let Σ be as in Ex. 2 and \mathbf{B}_{init}^f as in Eq. 2. Then $\mathcal{O}(\mathbf{B}_{init}^f \wedge \Sigma) \models$

$$[fwd(1, 2)]\mathbf{B}(h = 2 \wedge [sonar(3)]\exists x.\mathbf{B}(h = x : \frac{5}{8}) : \frac{3}{10})$$

$$\mathfrak{R}[fwd(1, 2), \mathbf{B}(h = 2 \wedge [sonar(3)]\exists x.\mathbf{B}(h = x : \frac{5}{8}) : \frac{3}{10})]$$

$$\equiv \frac{3}{10} = \frac{1}{\eta} \sum_{u,y} f(u) \left(\begin{array}{c} \mathfrak{I}[fwd(1, y), u] \times \\ \mathfrak{I}(u, \mathfrak{R}[fwd(1, y), h = 2 \\ \wedge [sonar(3)]\exists x.\mathbf{B}(h = x : \frac{5}{8})]) \end{array} \right)$$

$$\equiv \frac{3}{10} = \frac{1}{\eta} \sum_{u,y} f(u) \left(\begin{array}{c} \theta(1, y, .1, .8) \times \\ \mathfrak{I}(u, h - y = 2) \end{array} \right)$$

$$\equiv \frac{3}{10} = \frac{1}{\eta} \times .3 \equiv \text{TRUE}$$

The first equivalence is because Σ_{oi} declares an observational indistinguishability among $fwd(x, y)$ for all y .

Note that $LH(\langle \cdot \rangle, fwd(1, y), u) = \theta(1, y, .1, .8)$ and $LH(fwd(1, y), sonar(3), u) = \theta(3, u - y, .2, .6)$:

$$\mathfrak{R}[fwd(1, y), l(sonar(3)) = q_t]$$

$$\equiv \mathfrak{R}[fwd(1, y), q_t = \theta(3, h, .2, .6)]$$

$$\equiv \mathfrak{R}[fwd(1, y), \exists u'.h = u' \wedge q_t = \theta(3, u', .2, .6)]$$

$$\equiv \exists u'.h = u' + y \wedge q_t = \theta(3, u', .2, .6)$$

Therefore $\lambda(\exists u'.u = u' + y \wedge q_t = \theta(3, u', .2, .6)) = \theta(3, u - y, .2, .6)$. The second equivalence is because $\mathfrak{R}[fwd(1, y), [sonar(3)]\exists x.\mathbf{B}(h = x : \frac{5}{8})]$ (denoted as $\mathfrak{R}[\mathbf{IN}]$) returns TRUE in the same vein as the outer regression:

$$\mathfrak{R}[\mathbf{IN}] \equiv \exists x.\mathfrak{R}[fwd(1, y) \cdot sonar(3), \mathbf{B}(h = x : \frac{5}{8})]$$

$$\equiv \exists x.\frac{5}{8} = \frac{1}{\eta'} \sum_{u,y'} f(u) \left(\begin{array}{c} \theta(1, y', .1, .8)\theta(3, u - y', .2, .6) \\ \mathfrak{I}(u, h - y' = x) \end{array} \right)$$

$$\equiv \left(\begin{array}{c} \exists x.x = 2 \wedge \frac{5}{8}\eta' = .06 \vee \\ x = 3 \wedge \frac{5}{8}\eta' = .2 \vee x = 4 \wedge \frac{5}{8}\eta' = .06 \vee \\ x \neq 2 \wedge x \neq 3 \wedge x \neq 4 \wedge \frac{5}{8}\eta' = 0 \end{array} \right)$$

$$\equiv \exists x.x = 3 \equiv \text{TRUE}$$

Here, the second line holds because $\mathfrak{I}(u, h - y' = x)$ is non-zero only if $u - y' = x$ and $\theta(3, u - y', .2, .6)$ is non-zero only if $u - y'$ is among $\{2, 3, 4\}$. The third line is due to $\eta' = .32$ by computation. Hence $\models \mathfrak{R}[\mathbf{IN}]$ and moreover $\models \mathfrak{R}[fwd(1, 2), \mathbf{B}(h = 2 \wedge [sonar(3)]\exists x.\mathbf{B}(h = x : \frac{5}{8}) : \frac{3}{10})]$. By Theorem 6, $\mathcal{O}(\mathbf{B}_{init}^f \wedge \Sigma) \models [fwd(1, 2)]\mathbf{B}(h = 2 \wedge [sonar(3)]\exists x.\mathbf{B}(h = x : \frac{5}{8}) : \frac{3}{10})$.

5 Related Work

Our logic builds on \mathcal{DS} [Vaishak and Gerhard, 2017], a modal logic of degrees of belief based on a modal variant of the situation calculus [Lakemeyer and Levesque, 2011]. In particular, our logic allows to express belief distributions, while the

degree of belief in \mathcal{DS} is restricted to numbers. BHL define *Bel* via a special fluent p and belief distributions can be expressed axiomatically, but they leave the question of how to represent an agent's KB open.

The concept of only-believing to capture the beliefs of a KB was first introduced by Levesque [1990], from which the version in \mathcal{DS} derives. Only-believing without full introspection was studied in [Humberstone, 1987].

Reasoning about knowledge and probability has been addressed in many domains: in computer science [Nilsson, 1986; Fagin and Halpern, 1994], in game theory [Monderer and Samet, 1989; Aviad and Philippe, 2001] among others [Halpern, 2003]. The work of [Fagin and Halpern, 1994] is at the heart of BHL and this work. Most recently, [Belle and Levesque, 2018] reconsiders the BHL framework and extends it to continuous domains.

In terms of belief regression, our ideas are inspired by [Belle and Levesque, 2014b], an extension of [Belle and Levesque, 2013a; Belle and Levesque, 2013b] where regression is proposed for the BHL framework. What distinguishes us perhaps most is that our regression operator works for meta-beliefs as well.

The main challenge of regression in stochastic domains is how to address noisy sensor and effectors. This has been studied in standard probabilistic frameworks [Fox *et al.*, 2003; Thrun *et al.*, 2005; Boyen and Koller, 1998], and in symbolic programming [Boutilier *et al.*, 2001; Boutilier *et al.*, 2000].

6 Conclusions

We proposed a logic of belief for the specification of very expressive KBs in uncertain dynamic domains. Our main result is a novel form of regression which applies to beliefs about beliefs, including quantifying-in. We remark that, while our result assumes a single initial belief distribution, we believe that it can be extended to multiple distributions.

For future work, projection by progression, that is, by modifying the initial beliefs, is an interesting avenue. Existing work on progression like [Belle and Lakemeyer, 2011; Belle and Levesque, 2014a] may be a good starting point. Another direction would be to use the logic for the verification of belief programs [Belle and Levesque, 2015].

Acknowledgments

This work has been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) RTG 2236 'UnRAVeL' and by the EU ICT-48 2020 project TAILOR (No. 952215).

References

- [Aviad and Philippe, 2001] Heifetz Aviad and Mongin Philippe. Probability logic for type spaces. *Games and economic behavior*, 35(1-2):31–53, 2001.
- [Belle and Lakemeyer, 2011] Vaishak Belle and Gerhard Lakemeyer. A semantical account of progression in the presence of uncertainty. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*. AAAI Press, 2011.

- [Belle and Levesque, 2013a] Vaishak Belle and Hector Levesque. Reasoning about continuous uncertainty in the situation calculus. In *Twenty-Third International Joint Conference on Artificial Intelligence*. IJCAI/AAAI, 2013.
- [Belle and Levesque, 2013b] Vaishak Belle and Hector Levesque. Reasoning about probabilities in dynamic systems using goal regression. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*. AUA Press, 2013.
- [Belle and Levesque, 2014a] Vaishak Belle and Hector Levesque. How to progress beliefs in continuous domains. In *Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*. AAAI Press, 2014.
- [Belle and Levesque, 2014b] Vaishak Belle and Hector Levesque. PREGO: an action language for belief-based cognitive robotics in continuous domains. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 989–995. AAAI Press, 2014.
- [Belle and Levesque, 2015] Vaishak Belle and Hector Levesque. ALLEGRO: belief-based programming in stochastic dynamical domains. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 2762–2769. AAAI Press, 2015.
- [Belle and Levesque, 2018] Vaishak Belle and Hector J Levesque. Reasoning about discrete and continuous noisy sensors and effectors in dynamical systems. *Artificial Intelligence*, 262:189–221, 2018.
- [Belle and Levesque, 2020] Vaishak Belle and Hector J Levesque. Regression and progression in stochastic domains. *Artificial Intelligence*, 281:103247, 2020.
- [Belle *et al.*, 2016] Vaishak Belle, Gerhard Lakemeyer, and Hector Levesque. A first-order logic of probability and only knowing in unbounded domains. In *Thirtieth AAAI Conference on Artificial Intelligence*, pages 893–899. AAAI Press, 2016.
- [Boutilier *et al.*, 2000] Craig Boutilier, Raymond Reiter, Mikhail Soutchanski, Sebastian Thrun, et al. Decision-theoretic, high-level agent programming in the situation calculus. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 355–362. AAAI Press / The MIT Press, 2000.
- [Boutilier *et al.*, 2001] Craig Boutilier, Raymond Reiter, and Bob Price. Symbolic dynamic programming for first-order mdps. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 690–700. Morgan Kaufmann, 2001.
- [Boyen and Koller, 1998] Xavier Boyen and Daphne Koller. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 33–42. Morgan Kaufmann, 1998.
- [Claßen, 2013] Jens Claßen. *Planning and verification in the agent language Golog*. PhD thesis, Hochschulbibliothek der Rheinisch-Westfälischen Technischen Hochschule Aachen, 2013.
- [Fagin and Halpern, 1994] Ronald Fagin and Joseph Y Halpern. Reasoning about knowledge and probability. *Journal of the ACM (JACM)*, 41(2):340–367, 1994.
- [Fahiem *et al.*, 1999] Bacchus Fahiem, Halpern Joseph Y, and Levesque Hector J. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2):171–208, 1999.
- [Fox *et al.*, 2003] Dieter Fox, Jeffrey Hightower, Lin Liao, Dirk Schulz, and Gaetano Borriello. Bayesian filtering for location estimation. *IEEE pervasive computing*, 2(3):24–33, 2003.
- [Halpern, 2003] Joseph Y Halpern. *Reasoning about uncertainty*. MIT Press, 2003.
- [Humberstone, 1987] I. L. Humberstone. The modal logic of ‘all and only’. *Notre Dame J. Formal Log.*, 28(2):177–188, 1987.
- [Lakemeyer and Levesque, 2004] Gerhard Lakemeyer and Hector J Levesque. Situations, si! situation terms, no! In *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004)*, pages 516–526. AAAI Press, 2004.
- [Lakemeyer and Levesque, 2011] Gerhard Lakemeyer and Hector J Levesque. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence*, 175(1):142–164, 2011.
- [Levesque, 1990] Hector J. Levesque. All I know: A study in autoepistemic logic. *Artif. Intell.*, 42(2-3):263–309, 1990.
- [McCarthy, 1963] John McCarthy. Situations, actions, and causal laws. Technical report, Dept. of computer science, Stanford Univ, 1963.
- [Monderer and Samet, 1989] Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, 1989.
- [Nilsson, 1986] Nils J Nilsson. Probabilistic logic. *Artificial intelligence*, 28(1):71–87, 1986.
- [Raymond, 2001] Reiter Raymond. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT press, 2001.
- [Thrun *et al.*, 2005] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.
- [Vaishak and Gerhard, 2017] Belle Vaishak and Lakemeyer Gerhard. Reasoning about probabilities in unbounded first-order dynamical domains. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pages 828–836. ijcai.org, 2017.