Reasoning about Probabilities in Unbounded First-Order Dynamical Domains

Citation for published version:
Belle, V & Lakemeyer, G 2017, Reasoning about Probabilities in Unbounded First-Order Dynamical Domains. in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-2017).* IJCAI Inc, pp. 828-836, 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia, 19/08/17. https://doi.org/10.24963/ijcai.2017/115

Digital Object Identifier (DOI):
10.24963/ijcai.2017/115

Published In:
Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-2017)

Reasoning about Probabilities in Unbounded First-Order Dynamical Domains*

Vaishak Belle
University of Edinburgh
United Kingdom
vaishak@ed.ac.uk

Gerhard Lakemeyer†
RWTH Aachen University
Germany
gerhard@cs.rwth-aachen.de

Abstract

When it comes to robotic agents operating in an uncertain world, a major concern in knowledge representation is to better relate high-level logical accounts of belief and action to the low-level probabilistic sensorimotor data. Perhaps the most general formalism for dealing with degrees of belief and, in particular, how such beliefs should evolve in the presence of noisy sensing and acting is the account by Bacchus, Halpern, and Levesque.

In this paper, we reconsider that model of belief, and propose a new logical variant that has much of the expressive power of the original, but goes beyond it in novel ways. In particular, by moving to a semantical account of a modal variant of the situation calculus based on possible worlds with unbounded domains and probabilistic distributions over them, we are able to capture the beliefs of a fully introspective knowledge base with uncertainty by way of an only-believing operator. The paper introduces the new logic and discusses key properties as well as examples that demonstrate how the beliefs of a knowledge base change as a result of noisy actions.

1 Introduction

When it comes to robotic agents operating in an uncertain world, a major concern in knowledge representation is to better relate high-level logical accounts of beliefs and action to the low-level probabilistic sensorimotor data. In these and other applications, it is often not sufficient to say that a formula ϕ is unknown: we may need to say which of ϕ or $\neg\phi$ is more likely, and by how much. Motivated by such concerns, the unification of logic and probability has received much attention in the recent years [Russell, 2015].

Perhaps the most general formalism for dealing with *degrees of belief* in formulas, and in particular, with how that should evolve in the presence of noisy sensing and acting is the first-order logical account by Bacchus, Halpern, and Levesque (BHL) [1999]. 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 domain, making it particularly attractive for general-purpose high-level programming [Lakemeyer and Levesque, 2007].

The BHL account is an extension to Reiter's reworking of the situation calculus [Reiter, 2001]. (For space reasons, we will assume some familiarity with the language.) A modeler is taken to provide the initial beliefs of the agent using a special binary functional fluent p: $p(s', s)$ is understood as the weight accorded to s' when the agent is at s. A basic action theory is taken to include axioms about the preconditions of actions and its effects via successor state axioms as usual, the latter embodying Reiter's monotonic solution to the frame problem. But additionally, a special predicate $oi(a, a')$ denoting *observational indistinguishability* is used to say that executing a may result in a' occurring instead (unbeknownst to the agent), and a special function $l(a, s)$ is used to define the likelihood of the action a at s. A (fixed) successor state axiom for p determines the weight accorded to successors, resulting from oi-related actions:

$$\begin{aligned}
&p(s', do(a, s)) = u \;\equiv \\
&\quad \exists a', s''\ [oi(a, a') \land s' = do(a', s'') \land poss(a', s'') \land \\
&\qquad\qquad u = p(s'', s) \times l(a', s'')] \\
&\lor\ \neg\exists a', s''\ [oi(a, a') \land s' = do(a', s'') \land poss(a', s'') \land u = 0]
\end{aligned}$$

Belief in ϕ, then, is defined in terms of the sum of the weights of accessible situations where ϕ holds. (A simple initial constraint requiring that the belief in TRUE be a number ensures that the summation is well-defined.)

Since the situation calculus is defined axiomatically, no special semantics is needed. However, as argued in [Lakemeyer and Levesque, 2004], when we wish to consider theoretical questions that are not direct entailments of basic action theories, involved arguments based on Tarskian structures or considerable proof-theoretic machinery is needed. In the non-probabilistic epistemic situation calculus alone, elementary questions about knowledge – from $\boldsymbol{K}\alpha \supset (\boldsymbol{K}\beta \lor \boldsymbol{K}\gamma)$, does it follow that $\boldsymbol{K}\alpha \supset \boldsymbol{K}\beta$ or $\boldsymbol{K}\alpha \supset \boldsymbol{K}\gamma$, for example – require multi-page proofs. This situation is clearly much worse if we are arguing about degrees of beliefs. Moreover, when we think about modeling a partially specified domain, our intuition is that a quantitative account subsumes a qualitative specification, and an account with actions subsumes a non-dynamic specification. No such relation has been established for BHL, and we believe deriving such a result would be very involved for the above reasons. Finally, despite being a model

*The authors are grateful to Hector J. Levesque for valuable input on a preliminary draft, and thank the anonymous reviewers for helpful suggestions.

†Also affiliated with the University of Toronto, Canada.

of belief, BHL do not really consider any meta-belief properties: for example, what does *introspection* look like with degrees of belief?

Outside of these three concerns, from a knowledge representation point of view, a key question is this: how do we specify a *knowledge base* in a *probabilistic setting*? What we have in mind is a logical theory from which all of the beliefs and non-beliefs can be inferred.

In this paper, we reconsider the BHL model of belief, and propose a new logical variant that has much of the expressive power of the original. Our starting point is the logic $\mathcal{OBL}$ [Belle *et al.*, 2016] for reasoning about degrees of belief and only knowing, the latter modality allowing a succinct characterization for the beliefs of a knowledge base. The thrust of the semantical apparatus for $\mathcal{OBL}$ is a modal realization of BHL's conceptually simple definition for belief in a first-order setting. Basing our intuitions on the BHL formalism, we now investigate how $\mathcal{OBL}$ can be extended for acting and sensing over noise, yielding our proposal $\mathcal{DS}$. The semantics we propose not only brings to attention key constructions missing in the BHL framework for enabling things like introspection, it is also shown to capture a family of only knowing logics [Levesque and Lakemeyer, 2001; Lakemeyer and Levesque, 2004], roughly the non-probabilistic and non-dynamic counterparts. The resulting picture is a first-order model of belief for both qualitative and quantitative specifications in dynamical systems.

2 The Logic $\mathcal{DS}$

The language is built so as to reason about probabilistic beliefs and meta-beliefs over actions in a first-order setting. Quantification, in particular, is understood substitutionally wrt a fixed countably infinite set of *rigid designators* that exist in all possible worlds [Levesque and Lakemeyer, 2001].

2.1 Syntax

Formally, the non-modal fragment of $\mathcal{DS}$ (= degrees of belief in the situation calculus) consists of standard first-order logic with = (that is, connectives $\{\wedge, \forall, \neg\}$, syntactic abbreviations $\{\exists, \equiv, \supset\}$). In particular, we assume:

- an infinite supply of variables $\{x, y, \ldots, u, v, \ldots\}$;
- rigid function symbols of every arity, such as *obj5* and *move(x,y)*;
- fluent predicates of every arity, such as *Broken(x)*, *NextTo(x,y)*, including the following special symbols:
  - a unary predicate *poss* to denote the executability of an action;
  - a binary predicate *oi* to denote that two actions are indistinguishable from the agent's viewpoint; and
  - a binary predicate l that takes an action as its first argument and the action's likelihood as its second argument.

For simplicity, we do not include rigid (non-fluent) predicates or fluent (non-rigid) functions. The terms of the language are the least set of expressions such that:

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

We let $\mathcal{R}$ denote the set of all ground rigid terms. $\mathcal{R}$ is considered to be isomorphic to the domain of discourse. This is similar to [Levesque and Lakemeyer, 2001], where *standard names* are used as the domain of discourse. We also assume that the rationals are a subsort of $\mathcal{R}$, and we reserve the variables u and v to range over the rationals.[1]

If all the variables in an atom are substituted by terms from $\mathcal{R}$, then we call it a *ground* atom. Let $\mathcal{P}$ be the set of ground atoms in $\mathcal{DS}$. $\mathcal{DS}$ has two epistemic operators: $\boldsymbol{B}(\alpha\colon r)$ is to be read as "α is believed with a probability r," where r is a rational number. Next, the modality $\boldsymbol{O}(\alpha_1\colon r_1, \ldots, \alpha_k\colon r_k)$, where α_i does not mention $\{\boldsymbol{O}, \boldsymbol{B}\}$, and r_i is a rational, is to be read as "all that is believed is: α_1 with probability $r_1, \ldots,$ and α_k with probability r_k." We also use $\boldsymbol{K}\alpha$, to be read as "α is known," as an abbreviation for $\boldsymbol{B}(\alpha\colon 1)$. We write $\boldsymbol{O}\alpha$, to be read as "α is all that is known," to mean $\boldsymbol{O}(\alpha\colon 1)$.

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

2.2 Semantics

The semantics is given in terms of *possible worlds.* In a dynamic setting, such worlds are defined to interpret not only the current state of affairs, but also how that changes over actions. There are three key complications over non-probabilistic accounts with deterministic acting and sensing [Lakemeyer and Levesque, 2004]:

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

To begin with, let $\mathcal{Z}$ be all possible sequences of $\mathcal{R}$, including $\langle\rangle$, the empty sequence. Letting $\mathcal{P}$ denote the set of ground atoms as before, let $\mathcal{W}$ be the set of all mappings $\mathcal{P} \times \mathcal{Z} \mapsto \{0, 1\}$, which are the set of all possible worlds.

We will require that at every world $w \in \mathcal{W}$,

- l behaves like a function, that is, for all a, z, there is exactly one rational $n \geq 0$ such that $w[l(a, n), z] = 1$ and for all $n' \neq n$, $w[l(a, n'), z] = 0$;

[1] For simplicity, instead of having variables of the action sort distinct from those of the object sort, we lump both of these together and allow ourselves to use any term as an action or as an object, as in [Lakemeyer and Levesque, 2004]. While this does allow one to construct meaningless atoms, such as *Broken*(*pickup*(*rabbit*)), this is purely for the ease of the technical treatment. A general account can be developed along the lines of [Lakemeyer and Levesque, 2011].

- oi is an equivalence relation (reflexive, symmetric, and transitive) for all z.[2]

By a *distribution* d we mean a mapping from $\mathcal{W}$ to $\mathbb{R}^{\geq 0}$ (the set of non-negative reals) and an *epistemic state* e is any set of distributions.

To prepare for the semantics, we will need four notational devices. First, we extend the application of l to sequences:

Definition 1: We define $l^* : \mathcal{W} \times \mathcal{Z} \mapsto \mathbb{R}^{\geq 0}$ as follows:

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

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

Definition 2: Given any world w, we define $z \sim_w z'$:

- $\langle\rangle \sim_w z'$ iff $z' = \langle\rangle$;
- $z \cdot r \sim_w z'$ iff $z' = z^* \cdot r^*$, $z \sim_w z^*$ and $w[oi(r, r^*), z] = 1$.

Since oi is an equivalence relation, we immediately obtain:

Proposition 3: *$\sim_w$ is an equivalence relation.*

Third, after actions, we will restrict ourselves to *compatible* worlds that agree on observational indistinguishability:[3]

Definition 4: We write $w \approx_{oi} w'$ iff:

- for all a, a', z, $w[oi(a, a'), z] = w'[oi(a, a'), z]$.

Fourth, to extend the applicability of $poss$ for action sequences, we proceed as follows:

Definition 5: Define $exec(z)$ for any $z \in \mathcal{Z}$ inductively:

- for $z = \langle\rangle$, $exec(z)$ denotes TRUE;
- for $z = a \cdot z'$, $exec(z)$ denotes $poss(a) \wedge [a]exec(z')$.

We are finally ready for the semantics where, analogous to $\mathcal{OBL}$, to obtain a well-defined sum over uncountably many worlds, we will introduce some conditions for distributions used for evaluating epistemic operators:

Definition 6: We define NORM, EQ, BND for any d and any set $\mathcal{V} = \{(w_1, z_1), (w_2, z_2), \ldots\}$ as follows:

1. for any $\mathcal{U} \subseteq \mathcal{V}$, $\text{NORM}(d, \mathcal{U}, \mathcal{V}, r)$ iff $\exists b \neq 0$ such that $\text{EQ}(d, \mathcal{U}, b \times r)$ and $\text{EQ}(d, \mathcal{V}, b)$.
2. $\text{EQ}(d, \mathcal{V}, r)$ iff $\text{BND}(d, \mathcal{V}, r)$ and there is no $r' < r$ such that $\text{BND}(d, \mathcal{V}, r')$ holds.
3. $\text{BND}(d, \mathcal{V}, r)$ iff $\neg\exists k, (w_1, z_1), \ldots, (w_k, z_k) \in \mathcal{V}$ such that

$$\sum_{i=1}^{k} d(w_i) \times l^*(w_i, z_i) > r.$$

Intuitively, given $\text{NORM}(d, \mathcal{U}, \mathcal{V}, r)$, r can be seen as the *normalization* of the weights of worlds in $\mathcal{U}$ in relation to the set of worlds $\mathcal{V}$ as accorded by d. Here, $\text{EQ}(d, \mathcal{V}, b)$ expresses that the weight accorded to the worlds in $\mathcal{V}$ is b, and finally $\text{BND}(d, \mathcal{V}, b)$ ensures the weight of worlds in $\mathcal{V}$ is bounded by b. In essence, although $\mathcal{W}$ is uncountable, the conditions BND and EQ admit a well-defined summation of the weights on worlds. See [Belle *et al.*, 2016] for a formal justification of this claim.

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

- $e, w, z \models p$ iff p is an atom and $w[p, z] = 1$;
- $e, w, z \models n_1 = n_2$ iff n_1 and n_2 are identical;
- $e, w, z \models \alpha \wedge \beta$ iff $e, w, z \models \alpha$ and $e, w, z \models \beta$;
- $e, w, z \models \neg\alpha$ iff $e, w, z \not\models \alpha$;
- $e, w, z \models \forall x \alpha$ iff $e, w, z \models \alpha^x_n$ for all $n \in \mathcal{R}$;
- $e, w, z \models [r]\alpha$ iff $e, w, z \cdot r \models \alpha$;
- $e, w, z \models \Box\alpha$ iff $e, w, z \cdot z' \models \alpha$ for all $z' \in \mathcal{Z}$.

For the epistemic operators, let $\mathcal{W}^{(e,w,z)}_\alpha = \{(w', z') \mid z' \sim_w z, w' \approx_{oi} w, \text{ and } e, w', \langle\rangle \models [z']\alpha \wedge exec(z')\}$. (We drop the superscript when the context is clear.) That is, these are the pairs of worlds and executable action sequences that agree on oi with w, and where α holds. Then:

- $e, w, z \models \boldsymbol{B}(\alpha : r)$ iff $\forall d \in e$, $\text{NORM}(d, \mathcal{W}_\alpha, \mathcal{W}_{\text{TRUE}}, r)$;
- $e, w, z \models \boldsymbol{O}(\alpha_1 : r_1, \ldots, \alpha_k : r_k)$ iff for all d, $d \in e$ iff $\text{NORM}(d, \mathcal{W}_{\alpha_1}, \mathcal{W}_{\text{TRUE}}, r_1), \ldots, \text{NORM}(d, \mathcal{W}_{\alpha_k}, \mathcal{W}_{\text{TRUE}}, r_k)$.

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

3 Properties

We have the following objectives for this section. First, we would like to argue that the properties of belief proved for $\mathcal{OBL}$ also hold for $\mathcal{DS}$. Second, we would like to study in which sense $\mathcal{DS}$ is downward compatible with $\mathcal{OBL}$, $\mathcal{ES}$ [Lakemeyer and Levesque, 2004] and $\mathcal{OL}$ [Levesque and Lakemeyer, 2001]. Roughly speaking, $\mathcal{DS}$ should be seen as a dynamic extension to $\mathcal{OBL}$, a probabilistic extension to $\mathcal{ES}$, and a probabilistic dynamic extension to $\mathcal{OL}$.

A brief review of the non-dynamic and non-probabilistic counterparts of $\mathcal{DS}$ is found in the appendix. Readers may want to consult this supplementary material prior to perusing the compatibility results.

3.1 Belief and Only Knowing

We begin by listing some properties of $\boldsymbol{B}$, which one would expect from any reasonable probabilistic account of belief and which generalize similar properties found in $\mathcal{OBL}$.

Proposition 7:

1. *If* $\models \Box(\alpha \equiv \beta)$ *then* $\models \Box(\boldsymbol{B}(\alpha : r) \equiv \boldsymbol{B}(\beta : r))$;
2. $\models \Box(\boldsymbol{B}(\alpha \wedge \beta : r) \wedge \boldsymbol{B}(\alpha \wedge \neg\beta : r') \supset \boldsymbol{B}(\alpha : r + r'))$;

[2] We do not think oi being an equivalence relation is controversial, and our sample basic action theories will show that the constraint is also natural and intuitive. As we shall see, it is needed for enabling introspection. Not surprisingly, as BHL are not particularly concerned with meta-beliefs, there is no mention of such a constraint.

[3] Compatibility, it turns out, is needed for enabling introspection, discussions of which also do not appear in BHL.

3. $\models \Box(\boldsymbol{B}(\alpha\colon r) \wedge \boldsymbol{B}(\beta\colon r') \wedge \boldsymbol{B}(\alpha \wedge \beta\colon r'') \supset \boldsymbol{B}(\alpha \vee \beta\colon r + r' - r''))$.

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

Let us now consider the case of believing with certainty, that is, $\boldsymbol{K}\alpha \doteq \boldsymbol{B}(\alpha\colon 1)$. The following properties show that $\boldsymbol{K}$ behaves essentially like the $\boldsymbol{K}$-operator in $\mathcal{ES}$.

Theorem 8:

1. $\models \Box(\boldsymbol{K}\alpha \wedge \boldsymbol{K}(\alpha \supset \beta) \supset \boldsymbol{K}\beta)$;
2. $\models \Box(\forall x \boldsymbol{K}\alpha \supset \boldsymbol{K}\forall x\alpha)$;
3. $\models \Box(\exists x \boldsymbol{K}\alpha \supset \boldsymbol{K}\exists x\alpha)$.

Proof: Here we only prove the first item. Let $e, w, z \models \boldsymbol{K}\alpha \wedge \boldsymbol{K}(\alpha \supset \beta)$ and let $d \in e$. Let w' be any world and z' a sequence of actions such that $z' \sim_w z$, $w \approx_{oi} w'$, $e, w', \langle\rangle \models exec(z')$, and $d(w') \times l^*(w', z') \neq 0$. We will show that $(w', z') \in \mathcal{W}_\beta$. Since $\textsc{Norm}(d, \mathcal{W}_\alpha, \mathcal{W}_{\text{TRUE}}, 1)$ and $\textsc{Norm}(d, \mathcal{W}_{\alpha\supset\beta}, \mathcal{W}_{\text{TRUE}}, 1)$, $(w', z') \in \mathcal{W}_\alpha$ and $(w', z') \in \mathcal{W}_{\alpha\supset\beta}$ (because otherwise Norm would not be 1). Thus $e, w', z' \models \alpha \wedge \alpha \supset \beta$ and, therefore, $e, w', z' \models \beta$. Hence, given the assumptions and the definition of $\mathcal{W}_\beta$, $(w', z') \in \mathcal{W}_\beta$. Since this holds for any world compatible with w and any executable action sequence, $\textsc{Norm}(d, \mathcal{W}_\beta, \mathcal{W}_{\text{TRUE}}, 1)$, that is, $e, w, z \models \boldsymbol{K}\beta$. ∎

We remark that the converse of the Barcan Formula (2.) holds as well. Note, however, that $\Box(\boldsymbol{K}\exists x\alpha \supset \exists x\boldsymbol{K}\alpha)$ is not valid in general, that is, knowing that a formula is true for some individual does not imply that there is a particular individual for which the formula is known to be true.

In the static case of $\mathcal{OBL}$, knowledge is fully introspective (under one proviso). We can show that these properties carry over to the dynamic case as well. The following lemma is needed for the proof.

Lemma 9: *Let e, w, z be given and let $(w', z) \in \mathcal{W}^{(e,w,z)}_{\text{TRUE}}$. Then for all α, $\mathcal{W}^{(e,w,z)}_\alpha = \mathcal{W}^{(e,w',z')}_\alpha$.*

The proof makes use of the fact that w and w' agree on the extension of oi for all action sequences and that $\sim_w$ is an equivalence relation.

Theorem 10:

1. $\models \Box(\boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{K}\boldsymbol{B}(\alpha\colon r))$;
2. $\models \Box(\boldsymbol{K}\text{TRUE} \supset (\neg\boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{K}\neg\boldsymbol{B}(\alpha\colon r)))$.

Proof: Here we only prove the first item. Let $e, w, z \models \boldsymbol{B}(\alpha\colon r)$. We need to show $e, w, z \models \boldsymbol{B}(\boldsymbol{B}(\alpha\colon r)\colon 1)$, that is, for all $d \in e$, $\textsc{Norm}(d, \mathcal{W}_{\boldsymbol{B}(\alpha\colon r)}, \mathcal{W}_{\text{TRUE}}, 1)$ with $\mathcal{W}_{\boldsymbol{B}(\alpha\colon r)} = \{(w', z') \mid w' \approx_{oi} w, z' \sim_w z$, and $e, w', \langle\rangle \models [z']\boldsymbol{B}(\alpha\colon r) \wedge exec(z')\}$. It suffices to show that $\mathcal{W}_{\boldsymbol{B}(\alpha\colon r)} = \mathcal{W}_{\text{TRUE}}$. $\mathcal{W}_{\boldsymbol{B}(\alpha\colon r)} \subseteq \mathcal{W}_{\text{TRUE}}$ holds trivially. Conversely, let $(w', z') \in \mathcal{W}_{\text{TRUE}}$. Then $w' \approx_{oi} w$, $z' \sim_w z$, and $e, w', \langle\rangle \models [z']\text{TRUE} \wedge exec(z')$. It suffices to show that $e, w', \langle\rangle \models [z']\boldsymbol{B}(\alpha\colon r)$, that is, for all $d \in e$, $\textsc{Norm}(d, \mathcal{W}^{(e,w',z')}_\alpha, \mathcal{W}^{(e,w',z')}_{\text{TRUE}}, r)$. By Lemma 9, $\mathcal{W}^{(e,w',z')}_\alpha = \mathcal{W}^{(e,w,z)}_\alpha$ and $\mathcal{W}^{(e,w',z')}_{\text{TRUE}} = \mathcal{W}^{(e,w,z)}_{\text{TRUE}}$. Since, by assumption, for all $d \in e$, $\textsc{Norm}(d, \mathcal{W}^{(e,w,z)}_\alpha, \mathcal{W}^{(e,w,z)}_{\text{TRUE}}, r)$, we are done. ∎

Given the definition of $\boldsymbol{K}$, we immediately obtain

Corollary 11:

1. $\models \Box(\boldsymbol{K}\alpha \supset \boldsymbol{K}\boldsymbol{K}\alpha)$
2. $\models \Box(\boldsymbol{K}\text{TRUE} \supset (\neg\boldsymbol{K}\alpha \supset \boldsymbol{K}\neg\boldsymbol{K}\alpha))$

Let us now briefly turn to the notion of only believing. While it is well defined for arbitrary epistemic states and after any number of actions have occurred, in this paper we will only consider the case of only believing in the initial situation, that is, before any actions have occurred.[4]

By the definition of only believing we immediately obtain:

$$\models \boldsymbol{O}(\alpha_1\colon r_1, \ldots, \alpha_k\colon r\colon k) \supset \bigwedge \boldsymbol{B}(\alpha_i\colon r_i),$$

that is, whatever is only believed is also believed, as it should be. To see how non-beliefs follow from only believing, consider the case of distinct atomic propositions p and q. Then:

$$\models \boldsymbol{O}(p\colon r) \supset \neg\boldsymbol{B}(q\colon r') \text{ for all } r'.$$

This is because for any r' the epistemic state satisfying $\boldsymbol{O}(p\colon r)$ contains distributions d where $\textsc{Norm}(d, \mathcal{W}_q, \mathcal{W}_{\text{TRUE}}, r'')$ for some $r'' \neq r'$.

These properties are all inherited from $\mathcal{OBL}$. In Section 4, we will consider the more interesting case of beliefs evolving from only believing an initial KB after actions have occurred.

3.2 $\mathcal{OBL}$ is part of $\mathcal{DS}$

The main result we prove is:

Theorem 12: *For $\alpha \in \mathcal{OBL} \cap \mathcal{DS}$, $\models_{\mathcal{OBL}} \alpha$ iff $\models_{\mathcal{DS}} \alpha$.*

Here, we understand $\models_X \alpha$ to mean that α is valid in the logic X. Similar to the proof relating $\mathcal{ES}$ and $\mathcal{OL}$ [Lakemeyer and Levesque, 2004], there are two key complications in mapping the static and dynamic logics. First, the domain of quantification $\mathcal{N}$ (nullary function symbols) in $\mathcal{OBL}$ is a proper subset of $\mathcal{R}$. This is addressed by means of a bijection between these domains in the proofs. Second, for every $\mathcal{OBL}$ world w, which is a mapping from ground atoms $\mathcal{P}$ to $\{0, 1\}$, there are infinitely many $\mathcal{DS}$ worlds that agree with w initially. This is addressed by mapping models between the logics so that they agree on the satisfaction of $\mathcal{OBL}$ formulas. However, unlike in $\mathcal{ES}$, a further complication in the probabilistic context is that we will need to think carefully about how to map distributions between the logics so that they remain well-defined.

The proof is long, so we go over the ideas and intermediate lemmas below. First, we will map a $\mathcal{OBL}$ model to a $\mathcal{DS}$ model.

Lemma 13: *Suppose α is as above, and $\circ$ a bijection from $\mathcal{N}$ to $\mathcal{R}$. Then for any $\mathcal{OBL}$ model (e, w), there is a $\mathcal{DS}$ model, which we denote (e°, w°), such that $e, w \models \alpha$ iff $e^\circ, w^\circ \models \alpha^\circ$.*

We construct (e°, w°) as follows:

- Let w° be a $\mathcal{DS}$ world such that for all $\mathcal{OBL}$ atoms p and all z, $w^\circ[p^\circ, z] = 1$ iff $w[p] = 1$. (Moreover, $poss$, oi and l need to be additionally fixed to chosen values, which can be arbitrary; e.g., for all a, a', u, z: $w^\circ[poss(a), z] = 1$, $w^\circ[oi(a, a'), z] = 1$ iff $a = a'$, $w^\circ[l(a, u), z] = 1$ iff $u = 1$.)

[4]To handle only knowing after actions, an account of *progression* is needed [Lin and Reiter, 1997; Lakemeyer and Levesque, 2009], a topic we leave for the future.

- For any $d \in e$, let d° be a $\mathcal{DS}$ distribution such that for all $\mathcal{OBL}$ worlds w, $d^\circ(w^\circ) = d(w)$, and for all $w' \notin \{w^\circ \mid w \text{ is a } \mathcal{OBL} \text{ world}\}$, $d^\circ(w') = 0$.
- Let $e^\circ = \{d^\circ \mid d \in e\}$.

The proof for Lemma 13 is then by an easy induction on α.

In Lemma 13, since each $\mathcal{OBL}$ world maps to infinitely many $\mathcal{DS}$ worlds, mapping distributions is somewhat straightforward because all worlds outside of $\{w^\circ \mid w \text{ is a } \mathcal{OBL} \text{ world}\}$ are set to 0. In the converse direction, it is not obvious how to map distributions in $\mathcal{DS}$ to $\mathcal{OBL}$ distributions, because weights on infinitely many $\mathcal{DS}$ worlds may need to be assigned to a single $\mathcal{OBL}$ world. So, we consider a special class of models where this mapping is obvious.

Definition 14: Suppose w is a $\mathcal{DS}$ world. We call it a *static* world if for all $z \neq \langle\rangle$, for every $\mathcal{DS}$ atom p, $w[p,z] = w[p,\langle\rangle]$.] Let $\mathcal{W}^- \subseteq \mathcal{W}$ be the set of static $\mathcal{DS}$ worlds. We call a distribution d *static* iff for every $w \notin \mathcal{W}^-, d(w) = 0$. We call e *static* if it is a set of static distributions. A model (e, w, z) is called *static* if both e and w are static.

The key consideration here is that there is no loss in terms of entailment for $\mathcal{OBL}$ sentences. Let $\models_{static\ \mathcal{DS}}$ denote the restriction of the satisfaction relation to static models.

Theorem 15: *For any $\alpha \in \mathcal{OBL} \cap \mathcal{DS}$, $\models_{\mathcal{DS}} \alpha$ iff $\models_{static\ \mathcal{DS}} \alpha$.*

Lemma 16: *Suppose α is as above and $\circ$ is any bijection from $\mathcal{R}$ to $\mathcal{N}$. Suppose (e, w) is any static model, there is a $\mathcal{OBL}$ model, denoted (e°, w°), such that $e, w \models \alpha$ iff $e^\circ, w^\circ \models \alpha$.*

Here, let w° be a $\mathcal{OBL}$ world such that for every $\mathcal{DS}$ atom p, $w[p,\langle\rangle] = 1$ iff $w^\circ[p^\circ] = 1$. Then,

- For any $d \in e$, let d° be a $\mathcal{OBL}$ distribution such that for all static $\mathcal{DS}$ worlds w, $d^\circ(w^\circ) = d(w)$.
- Let $e^\circ = \{d^\circ \mid d \in e\}$.

By an induction on α, Lemma 16 can be shown. Finally, the proof for Theorem 12 is argued using Lemmas 13 and 16.

3.3 $\mathcal{ES}$ is part of $\mathcal{DS}$

$\mathcal{ES}$ worlds are precisely $\mathcal{DS}$ worlds, and an epistemic state in $\mathcal{ES}$ is defined as any set of worlds.

Truth is defined inductively as usual wrt a model (e, w, z), where, for example:

- $e, w, z \models \boldsymbol{K}\alpha$ iff for all $w' \sim_z w$, if $w' \in e$, $e, w', z \models \alpha$;

where $w' \sim_z w$ is defined inductively on z to posit that, after an action, a world is compatible with the real world provided the action is executable, and they agree on the sensing outcomes of the action, axiomatized using a special fluent *SF*.

When relating $\mathcal{ES}$ to $\mathcal{DS}$, observe that actions are implicitly assumed to be deterministic in $\mathcal{ES}$, which is simple enough to axiomatize using $\Box\forall a, a'\ (oi(a, a') \equiv a = a')$ and $\Box\forall a, u\ (l(a, u) \equiv u = 1)$ that we lump together as Ω. The point of serious divergence, then, is that knowledge in $\mathcal{ES}$ is defined in terms of the sensing results at the real world. Clearly, there is no analogue to this feature in the semantics of $\mathcal{DS}$. So, we restrict ourselves to trivial sensing results $\Gamma = \{\Box\forall a\ (SF(a) \equiv \text{TRUE})\}$. Then, we can show:[5]

Theorem 17: *For any $\alpha \in \mathcal{ES} \cap \mathcal{DS}$, $\Gamma \models_{\mathcal{ES}} \alpha$ iff $\Omega \models_{\mathcal{DS}} \Box\boldsymbol{K}\text{TRUE} \supset \alpha$.*

The proof rests on two lemmas, the first of which maps an $\mathcal{ES}$ model (e, w, z) to a $\mathcal{DS}$ model (e^*, w, z) where

$$e^* = \{d \mid \text{NORM}(d, \mathcal{U}, \mathcal{V}, 1)\}$$

where $\mathcal{U} = \{(w, \langle\rangle) \mid w \in e\}$ and $\mathcal{V} = \{(w, \langle\rangle) \mid w \in \mathcal{W}\}$, and the second of which maps a $\mathcal{DS}$ model (e, w, z) to a $\mathcal{ES}$ model (e^*, w, z) where $e^* = \{w' \mid \exists d \in e, d(w') > 0\}$.

3.4 $\mathcal{OL}$ is part of $\mathcal{DS}$

It was shown in [Belle *et al.*, 2016] that $\mathcal{OL}$ is part of $\mathcal{OBL}$:

Theorem 18: *For $\alpha \in \mathcal{OL} \cap \mathcal{OBL}$, $\models_{\mathcal{OL}} \alpha$ iff $\models_{\mathcal{OBL}} \boldsymbol{K}\text{TRUE} \supset \alpha$.*

By means of Theorem 12, we obtain:

Corollary 19: *For $\alpha \in \mathcal{OL} \cap \mathcal{DS}$, $\models_{\mathcal{OL}} \alpha$ iff $\models_{\mathcal{DS}} \boldsymbol{K}\text{TRUE} \supset \alpha$.*

4 Basic Action Theories

Let us now consider the equivalent of basic action theories of the situation calculus. Since there is no explicit notion of situations in $\mathcal{DS}$ and the uniqueness of names is built into the semantics, our basic action theories do not require foundational axioms. Our basic action theories will include the usual ingredients on the executability of actions and their effects as in $\mathcal{ES}$, but will additionally include axioms about observational indistinguishability and action likelihoods.

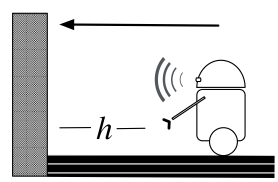

Figure 1: robot moving towards a wall

The first-order nature of the logic allows us to capture domains quite beyond standard probabilistic formalisms – see [Bacchus *et al.*, 1999; Belle and Levesque, 2015], for example – but to illustrate the features of the language, we will focus on two variants using a simple 1-dimensional robot. Imagine a robot to the right of a wall as in Figure 1. Let h be the fluent representing the distance to the wall. A noisy effector $move(x, y)$ brings the robot closer to the wall, with the understanding that x is how much the robot intends to move, but y is what actually happens. (A negative argument is understood as moving *away* from the wall.) A noisy sensor $sonar(z)$ provides a reading z that is taken as an estimate for h. The idea is that repeated sensing increases the robot's confidence about h.

Throughout, we have the following in common:

[5]We suspect it is possible to extend the result to non-trivial *SF* values, but only by means of an explicit axiomatization of the sensing outcomes as seen in the real world, similar to INDIGOLOG [Sardina *et al.*, 2004], a discussion of which we leave for an extended version of the paper.

- We assume the sensor action is always executable, and the move action is executable provided the argument does not cause the robot the hit the wall:[6]

$$\Box poss(a) \equiv \exists x, y, u(a = move(x, y) \land h(u) \land y \leq u) \lor \exists z(a = sonar(z) \land \text{TRUE}).$$

- Similar to BHL, noisy actions are retrofitted in successor state axioms by postulating that the effects of an action are against the actual argument:

$$\Box[a]h(u) \equiv \exists x, y(a = move(x, y) \land h(u + y)) \lor \forall x, y\ (a \neq move(x, y) \land h(u)).$$

- Clearly, all instances of *move* with an intended argument of x are observationally indistinguishable from each other. A sonar action is only observationally indistinguishable to itself, as the agent sees the actual argument corresponding to the reading on the sensor:

$$\Box oi(a, a') \equiv \exists x, y, z(a = move(x, y) \land a' = move(x, z)) \lor (a = sonar(z) \land a' = a).$$

- We assume the noise model of a sensor to depend on the true value of h, and that for the effector to depend on how much the actual argument diverges from the intended argument:

$$\Box l(a, u) \equiv \exists x, y, z(a = sonar(z) \land u = \Theta(h, z, .8, .1)) \lor (a = move(x, y) \land u = \Theta(x, y, .6, .2)) \lor (a \neq sonar(z) \land a \neq move(x, y) \land u = 0)$$

where,

$$\Theta(u, v, c, d) = \begin{cases} c & \text{if } u = v \\ d & \text{if } |u - v| = 1 \\ 0 & \text{otherwise} \end{cases}$$

So, observing a value z on the sonar means that it is very likely the true value, and at most it is off by a unit.[7] Analogously, for the move action.

We lump these axioms together as Σ.

Example 20: We begin by assuming a standard discrete uniform distribution on h, say, on the range $\{2, 3, 4\}$. In particular, we will consider the entailments of the following theory:

$$\boldsymbol{O}(h(2)\colon 1/3; h(3)\colon 1/3; h(4)\colon 1/3; \Sigma\colon 1). \tag{1}$$

The following are entailments of (1):

1. $\boldsymbol{B}(h(5)\colon 0)$

   By means of defining a probability distribution over 3 possible values for h, other values are impossible.
2. $[sonar(2)]\boldsymbol{B}(h(4)\colon 0) \land [sonar(2)]\boldsymbol{B}(h(2)\colon 8/9)$

   Obtaining a reading of 2 on the sensor means that being 4 units away is no longer possible, whereas the agent's confidence in being 2 units away from the wall increases.
3. $[sonar(2)][move(1, 1)]\boldsymbol{B}(h(2)\colon 8/45)$

   Consider that, in case of an exact move, the probability of being 2 units away would have been 0. See Figure 2 for the degrees of belief in other values of h.

[6]Free variables are assumed to be implicitly quantified from the outside.

[7]It is also possible to handle discrete approximations of Gaussians, for example.

***h* values**	2	3	4	5
initially	1/3	1/3	1/3	0
after *sonar(2)*	8/9	1/9	0	0
after *move(1,1)*	8/45	5/9	11/45	1/45

Figure 2: distribution on h values in Example 20

Example 21: We now assume the same dynamic axioms as above (given by Σ), but imagine the agent to only have qualitative uncertainty about h's value. Consider:[8]

$$\boldsymbol{O}(\exists u\ (h(u) \land u > 1) \land \Sigma\colon 1) \tag{2}$$

So, the agent considers infinitely many h values possible.

The following are entailed by (2):

1. $\boldsymbol{B}(h(1)\colon 0) \land \neg\boldsymbol{B}(h(4)\colon 0)$

   Values for $h > 1$ cannot be ruled out. Here, only knowing allows us to infer non-beliefs.
2. $[sonar(2)]\boldsymbol{B}(h(4)\colon 0)$

   The nature of the likelihood axioms for the sensor is such that obtaining a reading of 2 eliminates infinitely many possibilities.

The following is not entailed by (2):

3. $[sonar(2)]\exists u, v\ \boldsymbol{B}(h(2)\colon u) \land \boldsymbol{B}(h(3)\colon v) \land u > v$

   A sensor reading of 2 means that the robot is either 2 or 3 units away from the wall, as the prior on being 1 unit away is 0. Despite the reading favoring the case for the robot being two units away from the wall, qualitative uncertainty about h's value means there are distributions where $h(2)$ has a low or even 0 prior probability, and therefore, it does not follow that the degree of belief in $h(2)$ necessarily trumps that in $h(3)$.

5 Related Work

Reasoning about probabilities is widely studied in the logical literature – see, for example, [Gaifman, 1964; Nilsson, 1986; Halpern, 2003]; we briefly survey the major camps below. At the outset, we remark that a key feature of our work is only knowing,[9] which has not been considered for probabilistic specifications, except for $\mathcal{OBL}$ that we build on.

The inspiration for our work, and perhaps the one closest in spirit, is the work of BHL on degrees of belief in the situation calculus. It is an axiomatic proposal with a conceptually attractive definition of belief in a first-order setting. The thrust of our work is in providing a semantic basis for that proposal, and as we noted, without constructions like $\sim_w$ and $\approx_{oi}$, meta-beliefs do not work right away in the BHL framework. BHL also do not consider only knowing.

In less restrictive settings than full first-order $\mathcal{OL}$, Gabaldon and Lakemeyer [2007] consider a logic of only knowing and probability by meta-linguistically enforcing finitely

[8]For expressions such as $u > 1$, it is implicitly assumed that we are quantifying over the rationals, which can be accomplished by adding sorts to the language. This is left out of the semantics for simplicity.

[9]Only knowing is related to notions such as *minimal knowledge* [Halpern and Moses, 1984] and *total knowledge* [Pratt-Hartmann, 2000]. See [Levesque and Lakemeyer, 2001] for discussions.

many equivalence classes for possible worlds. Consequently, quantification also ranges over a finite set. In a game theory context, Halpern and Pass [2009] have considered a (propositional) version of only knowing with probability distributions.

Reasoning about knowledge and probability has appeared in a number of works prior to BHL, of course, in computer science [Nilsson, 1986; Fagin and Halpern, 1994], game theory [Monderer and Samet, 1989; Heifetz and Mongin, 2001], among others [Halpern, 2003]. Properties discussed in this paper, such as introspection and additivity, are also well studied [Aumann, 1999]. Notably, the work of Fagin and Halpern [1994] can be seen to be at the heart of BHL (and our work). The Fagin-Halpern scheme is a general one formulated for Kripke frames [Fagin *et al.*, 1995], but it is propositional. We also consider the simple case where a set of global distributions apply to $\mathcal{W}$ as seen at every world; in theirs, the probability spaces can differ arbitrarily across the worlds. The Fagin-Halpern scheme shares some similarity with probabilistic logics for programs [Kozen, 1981] and variants thereof [Halpern and Tuttle, 1993; Van Benthem *et al.*, 2009].

There are many previous first-order accounts of probabilities, such as [Bacchus, 1990] and [Halpern, 1990]; see [Ognjanovic and Raškovic, 2000] for a comprehensive list. Limited versions of probabilistic logics have also become popular in the machine learning literature [Poole, 2003; Domingos and Webb, 2012; Getoor and Taskar, 2007], with things like a finite domain assumption built-in. In these latter formalisms, the logical syntax is mostly used to succinctly represent large probabilistic graphical models with many interacting random variables over classes and hierarchies, and as such, the knowledge base is assumed to be equivalent to a single joint distribution over these variables. First-order accounts such as [Halpern, 1990] allow logical connectives over beliefs, as a result of which, like in BHL and $\mathcal{DS}$, beliefs may not correspond to any single distribution. However, our emphasis on a first-order theory of actions and only knowing distinguishes us from much of this work.

6 Conclusions

We proposed a first-order modal logic of subjective probabilities, noisy acting and sensing, and explored its formal properties. The language allows us to express BHL-style basic action theories, albeit in a crisp semantical framework, with introspection and only knowing. Bridging logical representations and probabilistic data is becoming a central problem in cognitive robotics and much of artificial intelligence; so, having a general specification language for analyzing systems and high-level programs is vital.

For the future, we think the following directions seem promising. From a reasoning viewpoint, it would be interesting to consider how to implement a possibly restricted (perhaps, even propositional) version of $\mathcal{DS}$ as a reasoning service. For starters, this would undoubtedly need an automated way to reason about actions, so that queries after actions can be reduced to purely static ones. Our sense is that earlier results for the BHL scheme [Belle and Levesque, 2013b; 2014] can be reformulated for $\mathcal{DS}$ with some effort.

From the representational viewpoint, the most immediate extension would be to allow both discrete and continuous probabilities, similar to [Belle and Levesque, 2013a]. This would allow the logic to capture the sorts of realistic sensor models seen in practice [Thrun *et al.*, 2005]. But even when limited to discrete probabilistic variables, there is an important limitation with the current syntax. Our belief operators were inherited from $\mathcal{OBL}$, which take the form $\boldsymbol{B}(\alpha, r)$, where r is a rational. But suppose we are interested in characterizing discrete probabilistic variables taking values from infinite sets. A simple case is a *geometric distribution*: imagine that the probabilistic variable h from our examples can take values from $\mathbb{N}$, and the probability of taking the value $u \in \mathbb{N}$ is $(1-\theta)^u \times \theta$ for a given $\theta \in [0, 1]$. For example, if $\theta = .5$, then the value 0 is accorded a probability of .5, the value 1 a probability of .25, and so on. Although we could use an infinite theory such as $\{\boldsymbol{B}(h(0)\colon .5), \boldsymbol{B}(h(1)\colon .25), \ldots\}$, it would be nice to be able to represent the distribution as a formula, perhaps in the context of $\boldsymbol{O}$. At first glance, this would require the second argument of $\boldsymbol{B}$ to possibly taking on variables of the rational sort rather than just rational constants, but how to design such a language and specify a semantics for it is an open question.

Appendix

We briefly review the non-dynamic and non-probabilistic counterparts of $\mathcal{DS}$: the logics $\mathcal{OL}$ [Levesque and Lakemeyer, 2001], $\mathcal{OBL}$ [Belle *et al.*, 2016], and $\mathcal{ES}$ [Lakemeyer and Levesque, 2004].

The Logic $\mathcal{OL}$

The non-modal fragment of $\mathcal{OL}$ consists of standard first-order logic with = (that is, connectives $\{\wedge, \forall, \neg\}$ and a countably infinite set of *standard names* $\mathcal{N}$, which, as usual, includes the set of rationals as a subsort. (Unlike $\mathcal{DS}$, where we needed to consider k-ary rigid function symbols for building action terms, $\mathcal{OL}$ can be limited to nullary rigid terms, making the account simpler.) As before, we let $\mathcal{P}$ be the set of ground atoms. Finally, as one would surmise, $\mathcal{OL}$ has two epistemic operators: $\boldsymbol{K}$ and $\boldsymbol{O}$.

The semantics of $\mathcal{OL}$, given using possible worlds, is defined over the set of mappings from the ground atoms in $\mathcal{P}$ to $\{0, 1\}$. By an *epistemic state* e, we mean any set of worlds. By a *model*, we mean a pair (e, w).

Given $\alpha \in \mathcal{OL}$ and a model (e, w), the definition of truth is defined inductively. We omit the logical connectives, and discuss the case for atoms and modalities:[10]

- $e, w \models p$ iff $w[p] = 1$;
- $e, w \models \boldsymbol{K}\alpha$ iff for all $w' \in e$, $e, w' \models \alpha$;
- $e, w \models \boldsymbol{O}\alpha$ iff for all w', $w' \in e$ iff $e, w' \models \alpha$.

[10]Readers familiar with the generic approach of accessibility relations between worlds in Kripke frames [Hughes and Cresswell, 1972] will realize that our notion of an epistemic state offers a simplified account, essentially corresponding to *weak S5*.

That is, knowing α amounts to α being true in all worlds in e, and only-knowing α amounts to e being precisely those worlds where α holds. A definition for entailment is given as we would for $\mathcal{DS}$.

The $\boldsymbol{K}$ modality in $\mathcal{OL}$ exhibits the usual properties of introspection, as well as the universal and existential versions of the Barcan formula; for example:

- $\models \boldsymbol{K}\alpha \supset \boldsymbol{K}\boldsymbol{K}\alpha$;
- $\models \exists x\boldsymbol{K}\alpha \supset \boldsymbol{K}\exists x\alpha$ but $\not\models \boldsymbol{K}\exists x\alpha \supset \exists x\boldsymbol{K}\alpha$.

The new modality here in comparison to classical epistemic logic is $\boldsymbol{O}$, and in that regard, only knowing implies knowing, but also not believing what does not logically follow from the knowledge base:

- $\models \boldsymbol{O}\alpha \supset \boldsymbol{K}\alpha$
- $\models \boldsymbol{O}\phi \supset \boldsymbol{K}\psi$ iff $\models \phi \supset \psi$
- $\models \boldsymbol{O}p \supset \neg\boldsymbol{K}q$ for distinct atoms p and q.

Clearly, the last property is not true in classical epistemic logic, because knowing p does not preclude knowing $p \wedge q$.

The Logic $\mathcal{OBL}$

From a syntax point of view, $\mathcal{OBL}$ is essentially identical to $\mathcal{OL}$ except for including modalities $\boldsymbol{B}(\alpha\colon r)$ and $\boldsymbol{O}(\alpha_1\colon r_1, \ldots, \alpha_k\colon r_k)$, with $\boldsymbol{K}\alpha$ and $\boldsymbol{O}\alpha$ serving as abbreviations as in $\mathcal{DS}$.

Given the set of ground atoms $\mathcal{P}$ and the set of worlds $\mathcal{W}$, which are mappings from $\mathcal{P}$ to $\{0, 1\}$, a *distribution* d is understood as a function from $\mathcal{W}$ to the set of non-negative reals $\mathbb{R}^{\geq 0}$. An *epistemic state* e is defined as a set of distributions. By a *model*, we mean a pair (e, w).

For the epistemic operators to be well-defined, we define conditions $\textsc{Bnd}, \textsc{Eq}$ and $\textsc{Norm}$ as in $\mathcal{DS}$, but these conditions now take the signature $(d, \mathcal{V}, r)$ where $\mathcal{V} \subseteq \mathcal{W}$. (That is, we are not considering world and action pairs, since there are no actions in the language.) For example: $\textsc{Bnd}(d, \mathcal{V}, r)$ iff there is no $k, w_1, \ldots, w_k \in \mathcal{V}$ such that

$$\sum_{i=1}^{k} d(w_i) \ > r.$$

With that, the semantic rules can be defined using (once again, omitting the connectives):

- $e, w \models p$ iff $w[p] = 1$;
- $e, w \models \boldsymbol{B}(\alpha\colon r)$ iff
  for all $d \in e$, $\textsc{Norm}(d, \{w' \mid e, w' \models \alpha\}, r)$;
- $e, w \models \boldsymbol{O}(\alpha_1\colon r_1, \ldots, \alpha_k\colon r_k)$ iff
  for all d, $d \in e$ iff $\textsc{Norm}(d, \{w' \mid e, w' \models \alpha_1\}, r_1)$,
  $\ldots, \textsc{Norm}(d, \{w' \mid e, w' \models \alpha_k\}, r_k)$.

Incidentally, $\neg\boldsymbol{K}\textsc{true}$ is satisfiable, which means that while positive introspection holds unconditionally, negative introspection needs to be predicated on believing $\textsc{true}$:[11]

[11]Since $\boldsymbol{B}(\textsc{true}\colon 1)$ is not valid, the authors in [Belle *et al.*, 2016] consider the notion of *measurable* epistemic states, which are precisely those that satisfy $\boldsymbol{B}(\textsc{true}\colon 1)$. All of this carries over to $\mathcal{DS}$, but for space reasons we will not pursue the topic here.

- $\models \boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{K}\boldsymbol{B}(\alpha\colon r)$ for arbitrary r;
- so, $\models \boldsymbol{K}\alpha \supset \boldsymbol{K}\boldsymbol{K}\alpha$ is then a special case;
- $\boldsymbol{K}\textsc{true} \models \neg\boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{K}\neg\boldsymbol{B}(\alpha\colon r)$ for arbitrary r.

Moreover, $\boldsymbol{B}$ admits additivity and equivalence properties; for example:

- $\models \boldsymbol{B}(\alpha \wedge \beta\colon r) \wedge \boldsymbol{B}(\alpha \wedge \neg\beta\colon r') \supset \boldsymbol{B}(\alpha\colon r + r')$;
- if $\models \alpha \equiv \beta$, then $\models \boldsymbol{B}(\alpha\colon r) \equiv \boldsymbol{B}(\beta\colon r)$.

Finally, regarding $\boldsymbol{O}$, the properties already discussed in the context of $\mathcal{OL}$ also hold in $\mathcal{OBL}$.

The Logic $\mathcal{ES}$

The logic $\mathcal{ES}$ semantically reconstructs the situation calculus in a logic of only knowing. To allow for action terms, as we have in $\mathcal{DS}$, one assumes k-ary rigid function terms. The set of atoms $\mathcal{P}$ are then obtained by applying all ground terms $\mathcal{R}$ to the predicates in the language.

Worlds in $\mathcal{ES}$ are no longer static entities that understand the (current) state of affairs. That is, a world maps $\mathcal{P}$ and $\mathcal{Z}$ to $\{0, 1\}$, where the set of action sequences $\mathcal{Z}$ is defined as we have for $\mathcal{DS}$. By a *model*, we mean a triple (e, w, z) where $z \in \mathcal{Z}$.

To account for how knowledge changes after (noise-free) sensing, one defines $w' \sim_z w$, which is to be read as saying "w' and w agree on the sensing for z", as follows:

- if $z = \langle\rangle$, $w' \sim_z w$ for every w';
- $w' \sim_{z\cdot a} w$ iff $w' \sim_z w$, $w[poss(a), z] = 1$ and $w'[SF(a), z] = w[SF(a), z]$.

The truth rules are (omitting connectives, as usual):

- $e, w, z \models p$ iff p is an atom and $w[p, z] = 1$;
- $e, w, z \models \boldsymbol{K}\alpha$ iff for all $w' \sim_z w$, if $w' \in e$, $e, w', z \models \alpha$;
- $e, w, z \models \boldsymbol{O}\alpha$ iff for all $w' \sim_z w$, $w' \in e$ iff $e, w', z \models \alpha$.

Properties of knowledge and only-knowing take on the same form as in $\mathcal{OL}$, with the addition that introspection holds for all action sequences; for example:

- $\models \square(\boldsymbol{K}\alpha \supset \boldsymbol{K}\boldsymbol{K}\alpha)$;
- $\models \square(\boldsymbol{O}\alpha \supset \boldsymbol{K}\alpha)$.

Action theories for $\mathcal{ES}$ are similar to $\mathcal{DS}$, except that acting and sensing is assumed to be noise-free. Among other things, there are no likelihood and observational indistinguishability axioms, and noise-free sensing is axiomatized by means of a distinguished predicate SF. In the context of moving towards a wall, for example, we might have:

$$\square SF(a) \equiv \exists u(a = isClose \wedge (h(u) \wedge u \leq 5) \vee a \neq isClose)$$

which says that SF is true on executing $isClose$ only when the robot is within 5 units of the wall, and is false otherwise. Semantically, then, given worlds w' and w, we would have $w' \sim_{isClose} w$ only when the robot is within 5 units from the wall in both worlds.

References

[Aumann, 1999] R. J. Aumann. Interactive epistemology II: probability. *Int. J. Game Theory*, 28(3):301–314, 1999.

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

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

[Belle and Levesque, 2013a] V. Belle and H. J. Levesque. Reasoning about continuous uncertainty in the situation calculus. In *Proc. IJCAI*, 2013.

[Belle and Levesque, 2013b] V. Belle and H. J. Levesque. Reasoning about probabilities in dynamic systems using goal regression. In *Proc. UAI*, 2013.

[Belle and Levesque, 2014] V. Belle and H. J. Levesque. How to progress beliefs in continuous domains. In *Proc. KR*, 2014.

[Belle and Levesque, 2015] V. Belle and H. J. Levesque. A logical theory of localization. In *Studia Logica*. 2015.

[Belle *et al.*, 2016] V. Belle, G. Lakemeyer, and H. J. Levesque. A first-order logic of probability and only knowing in unbounded domains. In *Proc. AAAI*, 2016.

[Domingos and Webb, 2012] P. Domingos and W. A. Webb. A tractable first-order probabilistic logic. In *AAAI*, 2012.

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

[Fagin *et al.*, 1995] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.

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

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

[Getoor and Taskar, 2007] L. Getoor and B. Taskar, editors. *An Introduction to Statistical Relational Learning*. MIT Press, 2007.

[Halpern and Moses, 1984] J. Y. Halpern and Y. Moses. Towards a theory of knowledge and ignorance: Preliminary report. In *Proc. NMR*, pages 125–143, 1984.

[Halpern and Pass, 2009] J. Y. Halpern and R. Pass. A logical characterization of iterated admissibility. In *Proc. TARK*, pages 146–155, 2009.

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

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

[Halpern, 2003] J. Y. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.

[Heifetz and Mongin, 2001] A. Heifetz and P. Mongin. Probability logic for type spaces. *Games and Economic Behavior*, 35(1-2):31 – 53, 2001.

[Hughes and Cresswell, 1972] G. E. Hughes and M. J. Cresswell. *An introduction to modal logic*. Methuen London, 1972.

[Kozen, 1981] D. Kozen. Semantics of probabilistic programs. *Journal of Computer and System Sciences*, 22(3):328 – 350, 1981.

[Lakemeyer and Levesque, 2004] G. Lakemeyer and Hector J. Levesque. Situations, Si! situation terms, No! In *Proc. KR*, pages 516–526, 2004.

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

[Lakemeyer and Levesque, 2009] G. Lakemeyer and H. J. Levesque. A semantical account of progression in the presence of defaults. In *IJCAI*, pages 842–847, 2009.

[Lakemeyer and Levesque, 2011] G. Lakemeyer and H. J. Levesque. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence*, 175:142–164, 2011.

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

[Lin and Reiter, 1997] F. Lin and R. Reiter. How to progress a database. *Artificial Intelligence*, 92(1-2):131–167, 1997.

[Monderer and Samet, 1989] D. Monderer and D. Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, June 1989.

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

[Ognjanovic and Rašković, 2000] Z. Ognjanovic and M. Rašković. Some first-order probability logics. *Theoretical Computer Science*, 247(1–2):191 – 212, 2000.

[Poole, 2003] D. Poole. First-order probabilistic inference. In *Proc. IJCAI*, pages 985–991, 2003.

[Pratt-Hartmann, 2000] I. Pratt-Hartmann. Total knowledge. In *Proc. AAAI*, pages 423–428, 2000.

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

[Russell, 2015] S. J. Russell. Unifying logic and probability. *Commun. ACM*, 58(7):88–97, 2015.

[Sardina *et al.*, 2004] 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.

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

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