Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Concerning Measures in a First-order Logic with Actions and Meta-beliefs *

Daxin Liu[1,2†], Qihui Feng[2†], Vaishak Belle[1], Gerhard Lakemeyer[2],
[1]The University of Edinburgh, UK
[2]RWTH-Aachen University, Germany
{daxin.liu, vbelle}@ed.ac.uk, {feng,lakemeyer}@kbsg.rwth-aachen.de

Abstract

The unification of logic and probability has been seen as a long-standing concern in philosophy and mathematical logic. In this paper, we propose a new general probabilistic modal logic of belief and only-believing in the situation calculus. Our logic can express both continuous and discrete degrees of belief. More importantly, expressing degrees of belief for arbitrary first-order formulas in a dynamic setting is possible for the first time, going well beyond previous proposals where fluents are assumed to be nullary or discrete. We show that our notion of belief retains many of the properties known from the previous related work.

1 Introduction

Starting from Nilsson (1986), there has been a steady body of work on the unification of logic and probability in AI. In so much as logic provides the deductive machinery with certain knowledge, and probability theory serves to capture the quantitative weighting of hypothesis, such unification is of fundamental importance to both computer and cognitive science. Although there were important algebraic efforts by Carnap (1962), Gaifman (1964), and others, it was Nilsson (1986), and later Bacchus (1989) and Halpern (1990), who carefully analyzed the semantics of a logical language that allowed for representing and reasoning about probabilistic assertions.

Since that early work, especially on the differences between the possible-worlds model for capturing subjective knowledge about the domain and the random-worlds model for capturing statistical assertions and random samples (Halpern 1990), there have been extensions for meta-beliefs (Fagin and Halpern 1994), programs (Harel 1984), time (Halpern and Shoham 1991) and processes (Boyen and Koller 1998). In the reasoning about actions community, perhaps the most general model is the work of Bacchus, Halpern, and Levesque (1999)(BHL henceforth). Although there are later works such as probabilistic ASP (de Morais and Finger 2013), probabilistic fluent calculus (Thielscher 2001), and probabilistic dynamic epistemic logic (Kooi 2003), BHL stands out in allowing for a very general definition of belief expressed in first-order logic over a rich theory of actions provided by the situation calculus.

*The title is inspired by the seminal work by Gaifman (1964).
†Corresponding authors; equal contributions

The main advantage of a logical account like BHL is that it allows a specification of beliefs that can be partial or incomplete, in keeping with whatever information is available about the application domain. It does not require specifying a prior distribution over some random variables from which posterior distributions are then calculated, as in Kalman filters, for example (Dean and Wellman 1991). Nor does it require specifying the conditional independences among random variables and how these dependencies change as the result of actions, as in the temporal extensions to Bayesian networks (Pearl 1988). In the BHL model, some logical constraints are imposed on the initial state of belief. These constraints may be compatible with one or very many initial distributions and sets of independence assumptions. All the properties of belief will then follow at a corresponding level of specificity. Put simply: the language allows for multiple initial distributions, possibly infinitely many, over possibly infinitely many actions and sensors, each of which might be characterized by one or more distributions. Formally, The BHL account provides axioms in the situation calculus regarding how the weight associated with a situation changes as the result of acting and sensing. Then belief in a formula is defined in terms of all the accessible situations where the formula is true and summing their weights.

Despite its expressiveness, BHL is still not general enough. For one thing, when the model was introduced in the late 90s, probabilistic robotics was just getting off the ground, so there was not that much immediate interest in modeling continuous distributions. But nowadays, not allowing for continuous distributions in a planning and action language would be considered problematic and limited. In that regard, developments have happened in the last decade:

- In (Belle and Levesque 2013) (BL henceforth), the BHL model was revisited to allow for continuous distributions. The key insight for this generalization was to recognize that because integration over situations is not well-defined, it is possible to integrate over the values of the fluents instead. However, this necessitates knowing which fluents to integrate over, only finitely many nullary fluents were allowed. Thus, there was a loss of expressiveness from a first-order viewpoint for the initial theory.
- In (Belle and Lakemeyer 2017), the logic $\mathcal{DS}$ was introduced, which reconstructs BHL in a modal logic to allow

for simpler semantical machinery to reason about meta-beliefs. (Because the classical situation calculus is defined axiomatically, reasoning about knowledge requires multi-page proofs involving things like Craig's interpolation lemma (Lakemeyer and Levesque 2011)).

- In orthogonal developments, (Liu and Lakemeyer 2021) and (Liu and Feng 2021) proposed the logic $\mathcal{DS}^*$ and $\mathcal{DS}_p$ to provide a semantics for capturing arbitrary distributions in $\mathcal{DS}$ and a definition of regression and progression (respectively) was considered for $\mathcal{DS}$. This extends the results of (Belle and Levesque 2020) that developed regression and progression for BL in allowing for meta-beliefs in the regression and progression operators.
- Finally, in (Belle 2023), BL was revisited in a modal setting, inspired by $\mathcal{DS}$, to introduce the logic $\mathcal{XS}$. But like BL, only finitely many discrete and/or continuous fluent were allowed.

Despite these developments, the question of what a general first-order treatment of beliefs and meta-beliefs in the presence of discrete, continuous, and possibly mixed distributions looks like was still open. It was only recently that (Feng et al. 2023) introduces the logic $\mathcal{OBL}c$ for precisely this generality but only in a static setting. The static setting was involved in itself and required the authors to revisit the notion of measures from first principles. The issue is mainly how to define measures for uncountably many worlds and obtain well-defined beliefs in a first-order setting with quantification and introspection.

In this work, we push the envelope further and investigate how to extend $\mathcal{OBL}c$ to account for actions. This requires some non-trivial treatments. For one thing, one needs to cast the Tarski-like structure of possible worlds to tree-like ones to account for actions (analogous to the tree(s) of situations in the situation calculus (Reiter 2001)). Besides, the most important question is how to incorporate the likelihood of stochastic actions and sensing into beliefs so that well-defined distributions holding initially remain well-defined after the actions have happened. This involves answering the question of how to assign likelihood for uncountable many actions and ensure it is a well-defined probability. To solve these problems, we introduce the notion of upper and lower measures in defining actions' likelihood and the notion of integral over possible worlds in incorporating actions' likelihood into beliefs after actions. When all of this machinery is put together, we get a fully general relational situation calculus with degrees of beliefs and meta-beliefs, over arbitrary distributions, including finite and infinite discrete distributions (e.g., Poisson), continuous (e.g., Gaussian, log normal) and combinations of such discrete and continuous properties (e.g., discrete for some values, and continuous for other ranges). This is the first time such a model has been introduced. It is worth noting that such distributions are considered to be extremely useful in machine learning, and so our approach would bridge the worlds of KR and ML further. Among other things, for illustrative purposes, we discuss a set of novel examples, involving an infinite ladder such that each step gets displaced by actions, and the perturbing of an infinite set of boxes, each with an unknown number of balls, by the addition of an unknown number of new balls. Such examples show the expressive power and flexibility of our new logic, called $\mathcal{PS}$ (for probabilistic situation calculus).

2 The Logic $\mathcal{PS}$

The logic $\mathcal{PS}$ is a first-order (FO) many-sorted modal variant of the situation calculus with modals of (continuous) degrees of belief and actions. For simplicity, we only consider functions with equality ($=$) and omit predicates. There are three sorts: *object*, *action*, and *real number* $\mathbb{R}$.

2.1 The Language

The vocabulary consists of *standard names*, *variables*, and *function symbols*. Standard names can be viewed as a fixed countable domain with the *unique names assumption*. Conventionally, we use n with (sub-)superscripts for standard names, e.g. $n_1, n_2, \dots n', n'', \dots etc.$ We use $a, a' \dots$ for variables of sort action, $x, y \dots$ for sort object or number, and v, v' for any sort.

- *rigid function symbols* of every arity, such as *sonar(x)*, including mathematical functions like $+, \times, e^x$, and a special binary symbol oi;[1]
- *fluent function symbols* of every arity, such as *distanceTo(x)*, *heightOf(y)*, including an unary special symbol l.

Roughly, l specifies the likelihood of actions and oi describes the observational-indistinguishability (alternative choices) among actions.

Besides, connectives $\wedge$, $\neg$, $\forall$ and modal operators $\boldsymbol{B}$, $\boldsymbol{O}$, $[\cdot]$, $\Box$ are used to construct formulas.

Terms (for respective sort) of the language are the least set of expressions such that

1. every standard name and variable is a term;
2. If $t_1, ..., t_k$ are terms and f a k-ary function symbol, then $f(t_1, ..., t_k)$ is a term of the same sort as f.

A term is said to be *rigid*, if and only if it does not contain fluents. *Ground terms* are terms without variables. *Primitive terms* are terms of the form $f(n_1, ..., n_k)$, where f is a function symbol and n_i are object or number standard names. We denote the sets of primitive terms of sort action as $\mathcal{P}_{act}$.

$[t_a]\alpha$ should be read as "α holds after action t_a" and $\Box\alpha$ as "α holds after any sequence of actions". The epistemic expression $\boldsymbol{B}(\alpha\colon r)$ should be read as "α is believed with a degree r". $\boldsymbol{K}\alpha$ means "α is known" and is an abbreviation for $\boldsymbol{B}(\alpha\colon 1)$. $\boldsymbol{O}(\alpha_1\colon r_1, \dots \alpha_k\colon r_k)$ may be read as "all that believed are (the conjunction of) α_i with degree r_i". Similarly, $\boldsymbol{O}\alpha$ means "α is only known" and is an abbreviation for $\boldsymbol{O}(\alpha\colon 1)$. For action sequence $z = a_1 \cdots a_k$, we write $[z]\alpha$ to mean $[a_1]\cdots[a_k]\alpha$. As usual, we treat $\alpha \vee \beta$, $\alpha \supset \beta$, $\alpha \equiv \beta$, and $\exists v.\alpha$ as abbreviations.

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 $\Box$ or $[t_a]$ is called *static*. A formula with no $\boldsymbol{B}$ or $\boldsymbol{O}$ is called *objective*. A formula with no fluent, $\Box$ or $[t_a]$ outside $\boldsymbol{B}$ or $\boldsymbol{O}$ is called *subjective*. A

[1]For simplicity, all action symbols are assumed to be rigid.

formula with no $\boldsymbol{B}$, $\boldsymbol{O}$, $\Box$, $[t_a]$, l, and oi is called a *fluent formula*. A fluent formula without fluent functions is called a *rigid formula*.

2.2 The Semantics

The semantics is given in terms of possible worlds, where a *world* is a tree-like Tarski structure. Formally, we assume a fixed domain of discourse $\mathcal{D} = \mathcal{D}_{obj} \cup \mathcal{D}_{act} \cup \mathbb{R}$ where $\mathcal{D}_{obj}$ is a countable infinite set of objects, $\mathbb{R}$ the real numbers, and $\mathcal{D}_{act}$ an uncountable infinite set of actions. The set of standard names $\mathcal{N} = \mathcal{N}_{obj} \cup \mathcal{N}_{act} \cup \mathcal{N}_{num}$ is a countable subset of $\mathcal{D}$.[2] We require that $\mathcal{N}_{obj} = \mathcal{D}_{obj}$, i.e. object constants are all named. Additionally, we assume elements in $\mathcal{D}_{act}$ are of the form $f_{act}(d_1, d_2, \ldots, d_k)$ where f_{act} is an action function symbol and $d_i \in \mathcal{D}_{obj} \cup \mathbb{R}$. With this, we assume action standard names are just those action primitive terms i.e. $\mathcal{N}_{act} = \mathcal{P}_{act}$. Lastly, we fix $\mathcal{N}_{num}$ to the set of *computable numbers* (Turing 1936), which is a countable subset of $\mathbb{R}$ but still includes important irrational numbers such as π, e. We denote the discrete sub-domain $\mathcal{N}_{obj} \cup \mathcal{N}_{num}$ as $\mathbb{D}$.

Truth of objective formulas Let Z refer to any finite *action sequences*, including the empty sequence $\langle\rangle$, i.e. $Z = (\mathcal{D}_{act})^*$. A *world* is a mapping from all function symbols and Z to functions of the corresponding sorts.

Definition 1. *Formally, a world w satisfies*

1. *for all k-ary object (likewise for action and number) function symbol f_{obj} and action sequence $z \in Z$,* $w[f_{obj}, z] \colon (\mathcal{D})^k \mapsto \mathcal{D}_{obj}$*;*
2. *Rigidity: if f is a rigid function symbol, then for all $(w, z), (w', z')$, $w[f, z] = w'[f, z']$;*
3. *Correctness in math: elementary function symbols (e.g. $+, \times, e^x$) are rigid and interpreted in the usual sense. For example, $w[+, z](1, 1) = 2$ for any w and z.* [3]

Let W be the set of all such worlds. A *variable map* θ maps each variable to an element in $\mathcal{D}$ of the right sort. We write $\theta \sim_v \theta'$ to mean θ and θ' agree except perhaps on variable v. The *denotation* of terms is defined recursively:

Definition 2. *the denotation of a term t under a tuple $\langle w, z, \theta\rangle$ is given as (assuming t, t_i are terms)*

- $\|t\|^z_{w,\theta} = t$ *if* $t \in \mathcal{N}$*;*
- $\|t\|^z_{w,\theta} = \theta(t)$ *if t is a variable;*
- $\|t\|^z_{w,\theta} = w[f, z](\|t_1\|^z_{w,\theta}, ..., \|t_k\|^z_{w,\theta})$ *if t is of the form $f(t_1, ..., t_k)$ where f is a function symbol.*

For simplicity, we write $\|t\|_\theta$ when t is rigid, $\|t\|^z_w$ when t does not contain variables, and $\|t\|$ when t is both rigid and ground. By a *model* we mean a triple $\langle w, z, \theta\rangle$. The truth of objective formulas is then given as:

- $w, z, \theta \models t_1 = t_2$ iff $\|t_1\|^z_{w,\theta}$ and $\|t_2\|^z_{w,\theta}$ are identical;
- $w, z, \theta \models \neg\alpha$ iff $w, z, \theta \not\models \alpha$;
- $w, z, \theta \models \alpha \wedge \beta$ iff $w, z, \theta \models \alpha$ and $w, z, \theta \models \beta$;
- $w, z, \theta \models \forall v.\alpha$ iff $w, z, \theta' \models \alpha$ for all $\theta' \sim_v \theta$;
- $w, z, \theta \models [t_a]\alpha$ iff $w, z \cdot n, \theta \models \alpha$ and $n = \|t_a\|^z_{w,\theta}$;
- $w, z, \theta \models \Box\alpha$ iff $w, z \cdot z', \theta \models \alpha$ for all $z' \in Z$.

Truth of static beliefs To give the semantics of $\boldsymbol{B}$ and $\boldsymbol{O}$, we need the notion of *epistemic state*. We begin with a brief recap of some key concepts in probability theory. A *measure space* is a tuple $\langle X, \mathcal{X}, \mu\rangle$, where X is a set, $\mathcal{X}$ is a σ-algebra on the set X (i.e. a set of subsets containing X and closed under complementation and countable union), and $\mu : \mathcal{X} \mapsto [0, +\infty]$ is a measure. Typical measure spaces include the Lebesgue measure spaces $\langle\mathbb{R}^n, \mathcal{M}, m\rangle$ where $\mathbb{R}^n$ is a n-dimensional Euclidean space, $\mathcal{M}$ is a σ-algebra on $\mathbb{R}^n$, and m is the Lebesgue measure (m respectively corresponds to the length, area, volume, etc of intervals, rectangle, cube, etc). A *probability space* is a special measure space whose measure is normalized, i.e. $\mu(X) = 1$. For probability spaces, usually X is called the *sample space*, $\mathcal{X}$ the *event set*, μ the *probability measure*. A probability space $\langle X, \mathcal{X}, \mu\rangle$ is said to be *complete* if for all $B \in \mathcal{X}$ with $\mu(B) = 0$ and all $A \subseteq B$, it holds $A \in \mathcal{X}$ and $\mu(A) = 0$. Intuitively, completeness means that if an event has zero probability, any subset of it is also an event and has zero probability; likewise, if an event has a probability 1, all its supersets are events and have probability 1. We restrict ourselves to complete probability spaces since each probability space can be uniquely extended to a complete probability space.

An *epistemic state* e is then defined as a set of μ s.t. $\langle W, \mathcal{W}, \mu\rangle$ forms a probability space, where $\mathcal{W}$ is the domain of μ and a σ-algebra on W [4]. Henceforth, we call such μ probability spaces (or distributions). We expand the model with the epistemic state, namely, a model is now a 4-tuple $\langle e, w, z, \theta\rangle$. For objective formulas, truth is given the same meaning as before since the epistemic state e plays no role. Let $\|\alpha\|_{e,\theta} = \{w' \mid e, w', \langle\rangle, \theta \models \alpha\}$. If e contains only a single element μ, we write $\|\alpha\|_{\mu,\theta}$ instead of $\|\alpha\|_{\{\mu\},\theta}$. Let r be a rigid term, the truth for $\boldsymbol{B}$ and $\boldsymbol{O}$ in the static case is given as:

- $e, w, \langle\rangle, \theta \models \boldsymbol{B}(\alpha\colon r)$ iff $\forall\mu$, $\mu \in e$ *implies* $\mu(\|\alpha\|_{\mu,\theta}) = \|r\|_\theta$;
- $e, w, \langle\rangle, \theta \models \boldsymbol{O}(\alpha\colon r)$ iff $\forall\mu$, $\mu \in e$ *iff* $\mu(\|\alpha\|_{\mu,\theta}) = \|r\|_\theta$;

Intuitively, $e, w, \langle\rangle, \theta \models \boldsymbol{B}(\alpha\colon r)$ if for all the probability spaces $\mu \in e$, the set of worlds that satisfy α under μ, i.e. $\|\alpha\|_{\mu,\theta}$, has probability measure $\|r\|_\theta$. Likewise, $e, w, \langle\rangle, \theta \models \boldsymbol{O}(\alpha\colon r)$ iff e is the maximal set of such probability spaces. Clearly, beliefs $\boldsymbol{B}(\alpha\colon r)$ defined in this way are indeed probabilities over possible worlds.

[2] Even if the domain is uncountable, we can only assign standard names to a countable subset of it.

[3] Sometimes, one might also wish to use the predicate '$<$' (similarly for '$\leq$') in formulas, this can be done by assuming a rigid function $lessthan$ which takes values from $\{0, 1\}$, additionally, for all worlds w and real number x, y, $w[lessthan](x, y) = 1$ iff $x < y$.

[4] Allowing multiple distributions rather than a single distribution in the epistemic state can avoid the problem that *de re* knowledge about degrees of belief, i.e. formulas such as $\exists x.\boldsymbol{K}(\boldsymbol{B}(\phi\colon x))$, is valid (Gabaldon and Lakemeyer 2007).

Truth of beliefs after stochastic actions Unlike the static case, defining beliefs after actions is more involved in a stochastic domain. There are at least two questions: 1) how to specify the non-deterministic effects of stochastic actions; 2) how to incorporate the likelihood of stochastic actions or sensing into beliefs.

The first question can be solved in the same way as in BHL. Namely, instead of saying a stochastic action has non-deterministic effects, we view the stochastic action as a set of ground actions (mutual *alternatives*) that are observationally indistinguishable to the agent and each has a deterministic effect. This is done via the special rigid function symbol oi.[5] For example, $oi(fwd(1,1), fwd(1,0))$ says that the agent cannot distinguish a successful forward $fwd(1,1)$ and a failure one $fwd(1,0)$.

The second question is more involved. To begin with, one would wish to express that some effects are more likely than others. This can be achieved by oi together with a special fluent $l(a)$ to specify the likelihood of an action a. However, the problem is how to ensure l is indeed a probability among potentially uncountably many action alternatives for stochastic actions. Recall the probability space consists of three ingredients: sample space, event set, and probability measures, it is challenging to use the single fluent l to express such a rich structure. Our solution is as follows:

1. there are finitely many action types, which are function symbols; Note that this does not imply a finite number of actions, as the parameters of actions can vary over an infinite domain, even the uncountable domain $\mathbb{R}$.
2. stochastic actions have at most 3 parameters, namely, they are of the form $sa(x, y_d, y_c)$ where x is the *controllable* and *observable* while y_d and y_c are the *uncontrollable* and *unobservable* to the agent. Additionally, y_d ranges over the discrete domain $\mathbb{D}$ while y_c ranges over the continuous domain $\mathbb{R}$. A formalism that allows more parameters in actions such as $sa(x_1, x_2, \ldots x_k, y_d, y_c)$ or $sa(x, y_d, y_{c_1}, y_{c_2}, \ldots, y_{c_k})$ is possible. For the simplicity of presentation, we restrict to three parameters. Nevertheless, actions may contain less than 3 parameters, like the action $fwd(x, y_c)$ does not contain the uncontrollable discrete parameter y_d;
3. the alternative relations oi only exist among actions of the same action symbols and with the same set of controllable parameters x.

Formally, we expand the model with an equivalence relation (or alternative relation) o among actions in $\mathcal{D}_{act}$. Moreover, for formulas involving oi, we assign truth as:

- $e, w, z, \theta, o \models oi(t_a, t'_a) = 1$ iff $\langle \|t_a\|^z_{w,\theta}, \|t'_a\|^z_{w,\theta}\rangle \in o$

For simplicity, we only consider the particular alternative relation given by:[6] $o_0 := \{\langle a, a'\rangle | \bigvee_i \exists x, y_c, y_d, y'_c, y'_d . a = sa_i(x, y_c, y_d)$ and $a' = sa_i(x, y'_c, y'_d)\}$. In the rest of the paper, we write $e, w, z, \theta \models \alpha$ instead of $e, w, z, \theta, o_0 \models \alpha$.

[5]For simplicity, we set oi to be rigid. Allowing oi to vary would cause counter-intuitive result, see (Liu and Lakemeyer 2021)

[6]This amounts to put the following sentence Σ_{oi} as a background theory: $\forall a, a'.\Box(oi(a, a') = 1 \equiv \bigvee_i \exists x, y_c, y_d, y'_c, y'_d . a = sa_i(x, y_c, y_d) \wedge a' = sa_i(x, y'_c, y'_d))$.

Now, we show how to ensure l is indeed a probability among potentially uncountably many action alternatives for stochastic actions.

Definition 3. *A world w is called proper if and only if*

- *for every stochastic action $a = sa(x, y_d, y_c)$ and action sequence z,*

$$\sum_{y_d \in \mathbb{D}} \int_{y_c \in \mathbb{R}} w[l, z](sa(x, y_d, y_c)) dy_c = 1. \qquad (1)$$

Implicitly, Eq. (1) requires that the fluent l has to be a Lebesgue integrable function in terms of parameters y_c. Essentially, a world w is proper if and only if the fluent l indeed specifies a probability distribution over alternatives (defined by o_0) for stochastic actions. This is better illustrated by the theorem below.

In the rest of the paper, whenever we mention a world, we mean a proper one. Meanwhile, by an epistemic state, we mean a set μ s.t. $\langle W_p, \mathcal{W}_p, \mu\rangle$ forms a probability space where W_p is the set of proper worlds, $\mathcal{W}_p$ is the domain of μ and also a σ-algebra on W_p, and μ is a probability measure.

Given a stochastic action $a = sa(x, y_d, y_c)$ and a set of its alternative $A' \subseteq A_a := \{a' | \langle a, a'\rangle \in o_0\}$, let

$$\mathbb{D}_{A'} = \{y'_d | \exists y'_c . \langle a, sa(x, y'_d, y'_c)\rangle \in A'\}$$
$$\mathbb{R}_{y_d, A'} = \{y'_c | \langle a, sa(x, y_d, y'_c)\rangle \in A'\}.$$

Namely, $\mathbb{D}_{A'}$ is the set of all possible discrete values in A', and $\mathbb{R}_{y_d, A'}$ is the set of all possible real values when fixing y_d in A'.

Theorem 1. *Given a proper world w and a stochastic action $a = sa(x, y_d, y_c)$, let $\mathcal{A}_a = \{A' \subseteq A_a \,|\, \forall y_d \in \mathbb{D}, \mathbb{R}_{y_d, A'}$ is Lebesgue measurable$\}$, for every action sequence z, defining $\nu^{w,z} \colon \mathcal{A}_a \mapsto \mathbb{R}^{\geq 0}$ as*

$$\nu^{w,z}(A') = \sum_{y'_d \in \mathbb{D}} \int_{y'_c \in \mathbb{R}_{y_d, A'}} w[l, z](sa(x, y'_d, y'_c)) dy'_c \quad (2)$$

then $\langle A_a, \mathcal{A}_a, \nu^{w,z}\rangle$ forms a probability space.

The proof is based on two facts: 1) the Lebesgue measurable sets form a σ-algebra, hence $\mathcal{A}_a$ is a σ-algebra on A_a; 2) by Def. 3 and using Lebesgue integral, one can show $\nu^{w,z}$ is a probability measure on $\mathcal{A}_a$. We call a set of actions $A' \subseteq A_a$ *measurable* if $A \in \mathcal{A}_a$.

The probability defined above can be extended from a single action to sequences in the following way.

Definition 4. *We define:*

1. *$l^* \colon W_p \times Z \rightarrow \mathbb{R}^{\geq 0}$ as*
   (a) *$l^*(w, \langle\rangle) = 1$, for every $w \in W_p$;*
   (b) *$l^*(w, z \cdot a) = l^*(w, z) \times w[l, z](a)$.*
2. *$z \sim z'$ as*
   (a) *$\langle\rangle \sim z'$ iff $z' = \langle\rangle$;*
   (b) *$z \cdot a \sim z'$ iff $z' = z^* \cdot a^*, z \sim z^*$, and $\langle a, a^*\rangle \in o_0$*

Let $\bigodot$ be the concatenation of actions. Given stochastic action sequence $z = \bigodot_{i=1}^k sa_i(x_i, y_{d_i}, y_{c_i})$, we denote

- the set of alternatives of z as $Z_z = \{z \,|\, z \sim z'\}$;

- the set of all possible discrete values of $\vec{y}'_d$ in the subset Z' of z's alternatives as $\mathbb{D}_{Z'} = \{\vec{y}'_d \mid \exists \vec{y}'_c. \bigodot_{i=1}^{k} a_i(x_i, y'_{d_i}, y'_{c_i}) \in Z' \subseteq Z_z\}$ where $\vec{y}'_d = [y'_{d_1}; \dots; y'_{d_k}]$ and $\vec{y}'_c = [y'_{c_1}; \dots; y'_{c_k}]$;
- the set of all possible real values of $\vec{y}'_c$ in Z' given $\vec{y}'_d$ as $\mathbb{R}_{\vec{y}'_d, Z'} = \{\vec{y}'_c \mid \bigodot_{i=1}^{k} sa_i(x_i, y'_{d_i}, y'_{c_i}) \in Z' \subseteq Z_z\}$.

Similar to Theorem 1, we have:

Theorem 2. *Given a proper world w and a sequence of stochastic actions $z = a_1 \cdots a_k$ where $a_i = sa_i(x_i, y_{d_i}, y_{c_i})$, let $\mathcal{Z}_z = \{Z' \subseteq Z_z \mid \forall \vec{y}'_d \in \mathbb{D}^{|\vec{y}_d|}, \mathbb{R}_{\vec{y}'_d, Z'}$ is Lebesgue measurable$\}$, defining $\nu^w_{l^*}\colon \mathcal{Z}_z \mapsto \mathbb{R}^{\geq 0}$ as:*

$$\nu^w_{l^*}(Z') = \textstyle\sum_{\vec{y}'_d \in \mathbb{D}_{Z'}} \int_{\vec{y}'_c \in \mathbb{R}_{\vec{y}'_d, Z'}} l^*(w, \bigodot_{i=1}^{k} a'_i) d\vec{y}'_c$$

where a'_i are actions obtained from a_i by replacing y_{d_i}, y_{c_i} with y'_{d_i}, y'_{c_i}, then $\langle Z_z, \mathcal{Z}_z, \nu^w_{l^*}\rangle$ forms a probability space.

The proof is similar to the proof of Theorem 1. We call a set of action sequences $Z' \subseteq Z_z$ *measurable* if $Z' \in \mathcal{Z}$.

A remark is that the definition of $\nu^{w,z}$ in Theorem 1 requires an action sequence z as a parameter while the definition of $\nu^w_{l^*}$ does not. In fact, $\nu^w_{l^*}$ is implicitly defined wrt the empty sequence $\langle\rangle$.

Now, the task remains to incorporate the well-defined likelihood of actions into the model of belief. For that, we need two additional notations: *the likelihood for an unmeasurable set of action sequences* and the *integral over possible worlds*.

Given a proper world w, even if only considering simple propositions like $distanceToWall = 9$, there might be an unmeasurable set of stochastic action sequences z (mutually alternative) such that $w, z \models distanceToWall = 6$ holds. Hence, inevitably, one needs to assign likelihoods for unmeasurable sets of action sequences. To do so, we borrow ideas from inner and outer Lebesgue measures. Formally,

Definition 5. *Given a proper world w and a sequence of stochastic action z, let $\nu^w_{l^*}$ be as in Theorem 4, for any $Z' \subset Z_z$, we define $\nu^{w+}_{l^*}$ and $\nu^{w-}_{l^*}$ as:*

$$\nu^{w+}_{l^*}(Z') = \inf\{\nu^w_{l^*}(Z'') \mid Z'' \text{ measurable, } Z' \subseteq Z''\}$$
$$\nu^{w-}_{l^*}(Z') = \sup\{\nu^w_{l^*}(Z'') \mid Z^* \text{ measurable, } Z'' \subseteq Z'\}.$$

$\nu^{w+}_{l^*}(Z')$ is called the *outer measure* of Z' while $\nu^{w-}_{l^*}(Z')$ is called the *inner measure* of Z'. Fig. 1 provides a conceptual illustration of $\nu^{w+}_{l^*}$ and $\nu^{w-}_{l^*}$. Dashed dots in the grid represent the measurable sets and the arrow on the left represents the direction of inclusion relations among sets. The set Z' (gray dot) is unmeasurable and its inner measure equals the measure of Z^{in} (black dot) while its outer measure equals the minimal of measures between Z_1^{out} and Z_2^{out} (hollow dots). A remark is that $\nu^{w+}_{l^*}(Z') = \nu^{w-}_{l^*}(Z')$ iff Z' is measurable.

The last ingredient is the notion of *integral over possible worlds*. Let $f_X(x) = 1\ if\ x \in X\ else\ 0$ be a *characteristic function* of the set X. Given a probability space $\langle X, \mathcal{X}, \mu\rangle$, the integral of a characteristic function of a measurable set $X' \subseteq X$ is defined in the standard way as: $\int_{x\in X} f_{X'}(x) d\mu = \mu(X')$. A *simple function* $g(x) : X \mapsto \mathbb{R}$ over a probability space $\langle X, \mathcal{X}, \mu\rangle$ is a finite linear combination of characteristic functions of measurable sets in that probability space. Namely, $g(x) = \sum_i b_i f_{X_i}(x)$ for some characteristic functions $f_{X_i}(x)$ and X_i are measurable ($b_i \in \mathbb{R}$). The integral of simple function $g(x)$ is defined as $\int_{x\in X} g(x) d\mu = \sum_i b_i \mu(X_i)$. Given an epistemic state e, let $\mu \in e$, we call a function $f : W_p \mapsto \mathbb{R}^{\geq 0}$ measurable on μ iff for all $b \in \mathbb{R}^{\geq 0}$, $\{w \in W_p \mid f(w) < b\} \in \mathcal{W}_p$.

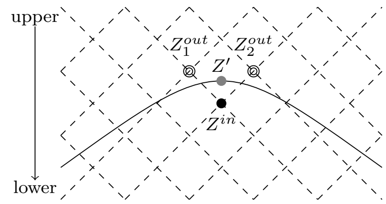

Figure 1: Illustration of $\nu^{w+}_{l^*}(Z')$ and $\nu^{w-}_{l^*}(Z')$.

Theorem 3. *Given a probability space μ, let $f : W_p \mapsto \mathbb{R}^{\geq 0}$ be a measurable function on $\mathcal{W}_p$ where $\mathcal{W}_p$ is the domain of μ. Then there exists a sequence of non-negative simple functions $\{g_0, g_1, \dots g_n, \dots\}$ s.t. $\forall w \in W_p$*

$$g_n(w) \leq g_{n+1}(w) \text{ and } \lim_{n\to\infty} g_n(w) = f(w)$$

The theorem suggests that any measurable function can be approximated arbitrarily close by using a series of simple functions. Hence, we exploit the limit of the integral over such simple functions to define the integral over measurable functions. This is in the same spirit as the definition of Lebesgue integral (Stein and Shakarchi 2009). [7]

Definition 6 (Integral over Possible Worlds)**.** *The integral of a measurable function f on μ is defined as*

$$\int_w f d\mu = \lim_{n\to\infty} \int_w g_n d\mu$$

where g_n are as in Theorem 3.

For any model $\langle e, w, z, \theta\rangle$ and formula α, let $Z^{\alpha}_{e,w,z,\theta} = \{z' \mid z' \sim z, e, w, z', \theta \models \alpha\}$. If α is a sentence, we write $Z^{\alpha}_{e,w,z}$. If the context of w and z is clear, we write Z^{α}_{e}. If e contains only a single element μ, we write Z^{α}_{μ} instead of $Z^{\alpha}_{\{\mu\}}$. Clearly, $Z^{\text{TRUE}}_{e,w,z,\theta} = Z_z$ for all e, w, z, θ.

Now we are ready to define the truth conditions of $\boldsymbol{B}$ and $\boldsymbol{O}$ after stochastic actions:

- $e, w, z, \theta \models \boldsymbol{B}(\alpha\colon r)$ iff $\forall \mu, \mu \in e$ *implies* $\int_w \nu^{w+}_{l^*}(Z^{\alpha}_{\mu,w,z,\theta}) d\mu = \int_w \nu^{w-}_{l^*}(Z^{\alpha}_{\mu,w,z,\theta}) d\mu = \|r\|_\theta$;
- $e, w, z, \theta \models \boldsymbol{O}(\alpha\colon r)$ iff $\forall \mu, \mu \in e$ *iff* $\int_w \nu^{w+}_{l^*}(Z^{\alpha}_{\mu,w,z,\theta}) d\mu = \int_w \nu^{w-}_{l^*}(Z^{\alpha}_{\mu,w,z,\theta}) d\mu = \|r\|_\theta$;

[7] The key to the proof is how to construct such $g_n(w)$. W.l.o.g. assume $f(w) \in [0, 1]$. For any n, let $f_{W_{n,i}}(w)$ be the characteristic function of the set $W_{n,i} = \{w \mid i/2^n \leq f(w) < (i+1)/2^n\}$. Then it is not hard to show $g_n(w) = \sum_{i=1}^{2^n} \frac{i-1}{2^n} f_{W_{n,i}}(w)$ is the desired simple function.

Henceforth, we write Z^α or Z^α_μ for $Z^\alpha_{\mu,w,z,\theta}$. The semantics above imposes two implicit constraints:

1. the integral $\int_w \nu^{w+}_{l^*}(Z^\alpha)d\mu$ and $\int_w \nu^{w-}_{l^*}(Z^\alpha)d\mu$ requires that functions $\nu^{w+}_{l^*}(Z^\alpha)$ and $\nu^{w-}_{l^*}(Z^\alpha)$ have to be measurable;
2. the integrals of outer and inner measure of Z^α are equal implies that the set $\{w \mid Z^\alpha_{\{\mu\},w,z,\theta}$ is unmeasurable$\}$ has zero-measure (i.e. zero-probability) under μ. This is crucial to ensure that $\boldsymbol{B}$ is indeed a well-defined probability (see properties later).

Intuitively, given μ, $\nu^{w+}_{l^*}(Z^\alpha)$ (resp. for $\nu^{w-}_{l^*}$) represents the superior (inferior) probability of α after z in the world w. Furthermore, the integrals, i.e. $\int_w \nu^{w+}_{l^*}(Z^\alpha)d\mu$ and $\int_w \nu^{w-}_{l^*}(Z^\alpha)d\mu$, are just the probability of all worlds where α holds after z.

Truth of belief after sensing actions The story for sensing actions is a bit different. To begin with, sensing outcomes are all observable to the agent hence sensing does not have alternatives other than themselves, yet one would still wish to express that some outcomes are more likely than others. Besides, after observing some sensing outcomes, some unrealistic epistemic states should be ruled out.

To address the first problem, in complying with our previous assumption on stochastic actions, we assume that: 1) sensing does not have the unobservable parameters y_c or y_d but only the observable parameters x, i.e. of the form $sen(x)$; 2) the alternative relations o_0 has to be expanded. That is to say, instead of the alternative relations o_0, we consider the alternative relation o_1 now where $o_1 = o_0 \cup \{\langle a, a\rangle | \bigvee_i \exists x.a = sen_i(x)\}$. This also means that the satisfaction relation $\models$ is regarding o_1 instead of o_0 henceforth. Besides, we redefine the concept of *proper world* to cover sensing. Formally,

Definition 7. *we call a world w proper if and only if*

- *for every stochastic action $a = sa(x, y_d, y_c)$ and action sequence z, w satisfies the constraint in Equation (1);*
- *for every action sequence $z = a_1 \cdots a_k$ where $a_i = sa_i(x_i, y_{d_i}, y_{c_i})$ or $a_i = sen_i(x_i)$, $l^*(w, z)$ is a measurable function in terms of $\vec{y}_c$ on $\mathbb{R}^{|\vec{y}_c|}$ where $\vec{y}_c = [y_{c_1}; \ldots; y_{c_k}]$ is the set of all uncontrollable continuous parameters in z.*

Essentially, the second condition above is to ensure that the function $l^*(w, z)$ (z may contain sensing) like the one in Theorem 2 is integrable. With such a definition, we have a theorem that is similar to Theorem 2. Let the notion Z_z, $\mathbb{D}_{Z'}$, and $\mathbb{R}_{\vec{y}'_d,Z'}$ be defined exactly the same as before but z could contain sensing of the form $sen_i(x)$, then we have:

Theorem 4. *Given a proper world w and an action sequence $z = a_1 \cdots a_k$ where $a_i = sa_i(x_i, y_{d_i}, y_{c_i})$ or $a_i = sen_i(x_i)$, let $\mathcal{Z}_z = \{Z' \subseteq Z_z \mid \forall \vec{y}'_d \in \mathbb{D}^{|\vec{y}_d|}, \mathbb{R}_{\vec{y}'_d,Z'}$ is Lebesgue measurable$\}$, defining $\nu^w_{l^*}: \mathcal{Z}_z \mapsto \mathbb{R}^{\geq 0}$ as:*

$$\nu^w_{l^*}(Z') = \sum_{\vec{y}'_d \in \mathbb{D}_{Z'}} \int_{\vec{y}'_c \in \mathbb{R}_{\vec{y}'_d,Z'}} l^*(w, \textstyle\bigodot_{i=1}^k a'_i) d\vec{y}'_c,$$

where a'_i are actions obtained from a_i by replacing y_{d_i}, y_{c_i} with y'_{d_i}, y'_{c_i}, then $\langle Z_z, \mathcal{Z}_z, \nu^w_{l^*}\rangle$ forms a measure space.

The task remains to update the agent's epistemic state after sensing. For this, we define the *compatibility* between action sequence and epistemic states. Formally, given μ and z, we require that there exists a value η^μ_z (or simply η) s.t. $\eta = \int_w \mu^w_{l^*}(Z^{\text{TRUE}})d\mu$. Intuitively, η represents the likelihood of actions sequence z under μ.

We call μ *compatible* with action sequence z iff $0 < \eta^\mu_z < \infty$. Namely, the sequence z has a positive finite likelihood in μ. Given $z \in Z$, we denote $e_z = \{\mu \mid \mu \in e \text{ and } 0 < \eta^\mu_z < \infty\}$ as the compatible subset of e wrt z. For any set $Z' \subset Z_z$, the inner measure $\nu^{w-}_{l^*}(Z')$ and outer measure $\nu^{w+}_{l^*}(Z')$ can be defined similarly as in Def. 5. Now we can define the truth conditions of $\boldsymbol{B}$ and $\boldsymbol{O}$ after any action sequences:

- $e, w, z, \theta \models \boldsymbol{B}(\alpha\colon r)$ iff $\forall \mu$, $\mu \in e_z$ *implies* $\frac{1}{\eta}\int_w \nu^{w+}_{l^*}(Z^\alpha_\mu)d\mu = \frac{1}{\eta}\int_w \nu^{w-}_{l^*}(Z^\alpha_\mu)d\mu = \|r\|_\theta$.
- $e, w, z, \theta \models \boldsymbol{O}(\alpha\colon r)$iff $\forall \mu$, $\mu \in e_z$ *iff* $\frac{1}{\eta}\int_w \nu^{w+}_{l^*}(Z^\alpha_\mu)d\mu = \frac{1}{\eta}\int_w \nu^{w-}_{l^*}(Z^\alpha_\mu)d\mu = \|r\|_\theta$.

For a sentence α, we write $e, w \models \alpha$ to mean $e, w, \langle\rangle, \theta \models \alpha$ for all variable maps θ. 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 $e, w \models \alpha$. We say that α is valid ($\models \alpha$) if $\{\} \models \alpha$. Satisfiability is then defined in the usual way (wrt o_1). 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.

3 Properties

Now that we have introduced the semantics, let us examine the properties of our logic. Despite the increasing expressiveness, our notions of belief and only-believing retain many of the properties known from its static predecessor, the logic $\mathcal{OBL}c$, and its discrete predecessor, the logic $\mathcal{DS}$.

3.1 Additivity and Equivalence

We show that the notion of belief $\boldsymbol{B}$ satisfies the properties of probability as in $\mathcal{DS}$. [8]

- If $\models \Box\alpha \equiv \beta$ then $\models \Box\boldsymbol{B}(\alpha : r) \equiv \boldsymbol{B}(\beta : r)$
- $\models \Box\boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{B}(\neg\alpha\colon 1 - r)$
  This means $\boldsymbol{B}$ satisfies the complement law of probability. The proof is based on the fact that for all w, $\nu^{w+}_{l^*}(Z^{\neg\alpha}) = \mu^w_{l^*}(Z^{\text{TRUE}}) - \nu^{w-}_{l^*}(Z^\alpha)$, which can be derived from the definition of $\nu^{w+}_{l^*}$ and $\nu^{w-}_{l^*}$.
- $\models \Box\boldsymbol{B}(\alpha \wedge \beta : r) \wedge \boldsymbol{B}(\alpha \wedge \neg\beta : r') \supset \boldsymbol{B}(\alpha : r + r')$
  This can be proved in the same spirit as above with the fact that (likewise for $\nu^{w-}_{l^*}$) $\nu^{w+}_{l^*}(Z^\alpha) = \nu^{w+}_{l^*}(Z^{\alpha\wedge\beta}) + \nu^{w+}_{l^*}(Z^{\alpha\wedge\neg\beta})$.
- $\models \Box\boldsymbol{B}(\alpha : r) \wedge \boldsymbol{B}(\beta : r') \wedge \boldsymbol{B}(\alpha \wedge \beta : r'') \supset \boldsymbol{B}(\alpha \vee \beta : r + r' - r'')$

[8] Conventionally, free variables are implicitly universally quantified outside. The modality $\Box$ has lower syntactic precedence than the connectives, and $[\cdot]$ has the highest priority.

The above properties are stemming from the fact that $\boldsymbol{B}$ is indeed a probability over the possible world.

3.2 Knowledge

Now, let us turn to the properties of knowledge $\boldsymbol{K}$. Recall that $\boldsymbol{K}\alpha$ is an abbreviation for $\boldsymbol{B}(\alpha\colon 1)$. We show that our $\boldsymbol{K}$ enjoys many properties of its qualitative counterpart in the logic $\mathcal{ES}$ (Lakemeyer and Levesque 2011) and its quantitative yet discrete counterpart in the logic $\mathcal{DS}$ (Belle and Lakemeyer 2017), including the universal and existential versions of the Barcan formula.

- $\models \Box\boldsymbol{K}\alpha \supset \boldsymbol{K}(\alpha \vee \beta)$
  The proof is based on the fact that for all w, $Z^{\alpha} \subseteq Z^{\alpha\vee\beta}$.
- $\models \Box\boldsymbol{K}\alpha \wedge \boldsymbol{K}\beta \supset \boldsymbol{K}(\alpha \wedge \beta)$

  Proof. Suppose that $e, z \models \boldsymbol{K}\alpha \wedge \boldsymbol{K}\beta$. By the complement law $e, z \models \boldsymbol{B}(\neg\alpha\colon 0) \wedge \boldsymbol{B}(\neg\beta\colon 0)$. Additionally, $\models \Box\boldsymbol{K}\alpha \supset \boldsymbol{K}(\alpha \vee \beta)$, which means $e, z \models \boldsymbol{B}(\neg\alpha \wedge \neg\beta\colon 0)$. According to the addition law, $\models \boldsymbol{B}(\neg\alpha\colon 0) \wedge \boldsymbol{B}(\neg\beta\colon 0) \wedge \boldsymbol{B}(\neg\alpha \wedge \neg\beta\colon 0) \supset \boldsymbol{B}(\neg\alpha \vee \neg\beta\colon 0)$. Thus, $e, z \models \boldsymbol{K}(\alpha \wedge \beta)$. □

- $\models \Box\boldsymbol{K}\alpha \wedge \boldsymbol{K}(\alpha \supset \beta) \supset \boldsymbol{K}\beta$

  Proof. Suppose $e, z \models \boldsymbol{K}\alpha \wedge \boldsymbol{K}(\alpha \supset \beta)$. It suffice to prove $e, z \models \boldsymbol{K}((\alpha \vee \beta) \wedge (\neg\alpha \vee \beta))$. By the first property of knowledge, $e, z \models \boldsymbol{K}(\alpha \vee \beta)$. $\boldsymbol{K}(\neg\alpha \vee \beta)$ is equivalent to $\boldsymbol{K}(\alpha \supset \beta)$. Based on the second property of knowledge, $e, z \models \boldsymbol{K}\beta$. □

- $\models \Box\exists x.\boldsymbol{K}\alpha \supset \boldsymbol{K}\exists x.\alpha$

  Proof. Henceforth, we write Z^{α}_{θ} for $Z^{\alpha}_{\{\mu\},w,z,\theta}$.
  Suppose $e, z, \theta \models \exists x.\boldsymbol{K}\alpha$. By semantics, $\exists\theta'$ s.t. $\theta' \sim_x \theta$, $e, z, \theta' \models \boldsymbol{K}\alpha$. Namely, for all $\mu \in e_z$,
  $\eta = \int_w \nu^{w+}_{l^*}(Z^{\alpha}_{\theta'})d\mu = \int_w \nu^{w-}_{l^*}(Z^{\alpha}_{\theta'})d\mu$. Meanwhile, $Z^{\exists x.\alpha}_{\theta} = \bigcup_{\{\theta'' \mid \theta'' \sim_x \theta\}} Z^{\alpha}_{\theta''}$. Therefore, $Z^{\alpha}_{\theta'} \subseteq Z^{\exists x.\alpha}_{\theta}$. Hence (likewise for $\nu^{w-}_{l^*}$)
  $$\eta = \int_w \nu^{w+}_{l^*}(Z^{\alpha}_{\theta'})d\mu \leq \int_w \nu^{w+}_{l^*}(Z^{\exists x.\alpha}_{\theta})d\mu$$
  $$\leq \int_w \nu^{w+}_{l^*}(Z^{\text{TRUE}}_{\theta})d\mu = \eta.$$
  Thus, $\eta = \int_w \nu^{w+}_{l^*}(Z^{\exists x.\alpha}_{\theta})d\mu = \int_w \nu^{w-}_{l^*}(Z^{\exists x.\alpha}_{\theta})d\mu$. □

The converse of the above formula does not hold: knowing that α holds for someone does not imply knowing that individual. For the universal version, it holds

- $\models \Box\boldsymbol{K}\forall x.\alpha \supset \forall x.\boldsymbol{K}\alpha$
  This can be proved in a similar spirit as above and based on the fact that $Z^{\forall x.\alpha} = \bigcap_{\theta} Z^{\alpha}_{\theta} \subseteq Z^{\alpha}_{\theta}$.

The converse of the above formula, i.e. $\Box\forall x.\boldsymbol{K}\alpha \supset \boldsymbol{K}\forall x.\alpha$, is not valid as already shown in (Feng et al. 2023) for the static cases. The reason is that in a probability space with uncountably many samples, there could be uncountably many distinct events where each of them has probability 1 (those satisfying $\boldsymbol{K}\alpha$), any countable intersection of them has probability 1, yet an uncountable intersection does not have probability 1.

Besides Barcan formula, we have the following properties in terms of introspection:

- $\models \Box\boldsymbol{K}\alpha \supset \boldsymbol{K}\boldsymbol{K}\alpha$
- $\models \Box\boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{K}\boldsymbol{B}(\alpha\colon r)$

That is, the agent has knowledge about what is known or believed. However, negative introspection is not valid as in (Feng et al. 2023), i.e. $\not\models \Box\neg\boldsymbol{B}(\alpha\colon r) \supset \boldsymbol{K}\neg\boldsymbol{B}(\alpha\colon r)$. This is because there exists $\mu \in e$ that satisfies $\neg\boldsymbol{B}(\alpha\colon r)$ does not imply for all $\mu \in e$ satisfies $\neg\boldsymbol{B}(\alpha\colon r)$, which is exactly the condition of $e \models \boldsymbol{K}\neg\boldsymbol{B}(\alpha\colon r)$.

3.3 Only-Believing

Lastly, we examine the properties of $\boldsymbol{O}$ here. Only-knowing (or only-believing) captures the intuition that the beliefs and non-beliefs of an agent are precisely those that follow from its knowledge base. Hence it is useful to characterize a knowledge base.

To begin with, the unique model theorem holds for only-believing as in the work (Liu and Lakemeyer 2021).

Theorem 5 (Uniqe Model Theorem)**.** *For any sentence α and rigid ground term r, there is a unique epistemic state e such that $e \models \boldsymbol{O}(\alpha\colon r)$.*

Proof. By the semantics, we have $e \models \boldsymbol{O}(\alpha\colon r)$ iff $e = \{\mu \mid \mu(\|\alpha\|_{\mu}) = \|r\|\}$. Clearly, there is only one such e. □

Besides, as in the logic $\mathcal{OL}$ (Levesque and Lakemeyer 2001), only-knowing implies knowing and not knowing about what is not entailed by the knowledge base. Below, let α be an arbitrary sentence and ϕ, ψ objective sentences:

- $\models \boldsymbol{O}(\alpha\colon r) \supset \boldsymbol{B}(\alpha\colon r)$
- $\models \boldsymbol{O}\phi \supset \boldsymbol{K}\psi$ iff $\phi \models \psi$
- $\models \boldsymbol{O}(\alpha\colon r) \supset \neg\boldsymbol{B}(h(\vec{n}) = m\colon r')$ for all r', where $\vec{n}, m$ are standard names and h is a fluent not in α.

E.g. let $\Sigma := \boldsymbol{O}(\mathit{weightOf}(A) \leq 50\colon 0.9)$, then we have $\Sigma \models \boldsymbol{B}(\mathit{weightOf}(A) \leq 50\colon 0.9)$ and $\Sigma \models \forall x.\neg\boldsymbol{B}(1.7 \leq \mathit{heightOf}(A)\colon x)$. That is, only-believing the person A's weight is less than 50kg with a degree 0.9 entails believing the person's weight is less than 50kg and also, not believing A's height is greater than 1.7m with any degree.

4 Reasoning about Actions and Beliefs

In this section, we show the expressiveness and flexibility of our logic through examples.

4.1 Expressing Belief Distributions

To reason about beliefs, one needs to specify what is believed initially. Let us consider the robot moving example in Fig. 2 from (Belle and Lakemeyer 2017). In the (discrete) logic $\mathcal{DS}$ and its extensions, a nullary fluent h is used to indicate the robot's horizontal distance toward the wall. Moreover, the robot's uncertain belief about h is specified via a formula of the form $\forall x.\boldsymbol{B}(h = x\colon f(x))$. Namely, the

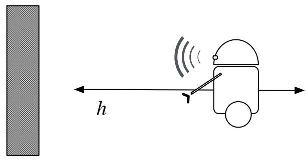

Figure 2: A robot moving towards a wall.

rigid mathematical function $f(x)$ expresses a *belief distribution* of the random variable h where $\sum_x f(x) = 1$. This can be easily extended for multiple but finite nullary fluents.

In our logic, we could do the same for discrete beliefs, but to express continuous belief distributions, we exploit a similar formula: a (potentially continuous) belief distribution over a finite set of nullary fluents $\vec{h} = \{h_1, h_2, \dots h_k\}$ is a formula of the form $\forall \vec{x}.\boldsymbol{B}(\bigwedge_i h_i \leq x_i \colon f(\vec{x}))$ where f is a rigid mathematical function satisfying $f(-\infty) = 0$ and $f(+\infty) = 1$. [9] We write $\boldsymbol{B}^f$ for short. Intuitively, $f(\vec{x})$ represents the *joint cumulative distribution function* of random variables $\vec{h}$. For example, the formula $\forall x.\boldsymbol{B}(h \leq x \colon \mathcal{U}(x; 10, 11))$ expresses the robot believes its distance is distributed uniformly among the interval $[10, 11]$, where

$$\mathcal{U}(x; 10, 11) = \begin{cases} 0 & x < 10 \\ x - 10 & x \in [10, 11) \\ 1 & 11 \leq x \end{cases}$$

Note that $\boldsymbol{B}^f$ could also be a discrete or mixture distribution. The belief distributions $\forall x.\boldsymbol{B}(h \leq x \colon f'(x))$ and $\forall x.\boldsymbol{B}(h \leq x \colon f''(x))$ below respectively express such distributions.

$$f'(x) = \begin{cases} 0 & x < 1 \\ 0.5 & x \in [1, 2) \\ 1 & 2 \leq x \end{cases} \quad f''(x) = \begin{cases} 0 & x \leq 0 \\ 0.5 & x = 0 \\ x + 0.5 & x \in (0, 0.5] \\ 1 & 1 \leq x \end{cases}$$

The latter is rather useful since the robot's belief after actions is usually a mixture distribution even if it starts with a pure continuous one. Imagine that the robot cannot get across the wall and its initial belief distortion is given as $\mathcal{U}(x; 10, 11)$, then after a deterministic action $fwd(10.5)$, it will end up in $\boldsymbol{B}^{f''}$. This is because some probability mass will accumulate at the point $h = 0$.

We comment that expressing such mixture distribution is impossible in the logic $\mathcal{DS}$ since it is discrete nor in the BL framework. The reason is that their notion of belief extends the BHL formalism and is defined in terms of normalized sums or integral of weights of situations. Since it is unclear how to integrate over situations, they map situations to fluent values which are purely continuous, then belief is given in terms of integral over the continuous values of fluents that are pure continuous or pure discrete.

4.2 Basic Action Theories

To infer beliefs after actions, one needs to specify a theory of action besides the initial beliefs. In our logic, the basic action theory includes a set of successor state axioms, one for each fluent, incorporating Reiter's (2001) solution to the frame problem, and a set of likelihood axioms, one for each action symbol [10].

The following is a possible basic action theory for the robot moving example

- the successor state axiom Σ_{post} for fluent h as:

$$\begin{aligned} \square[a]h = v \equiv & \exists x, y_c.a = fwd(x, y_c) \wedge v = h - y \\ & \vee \forall x, y_c.a \neq fwd(x, y_c) \wedge v = h \end{aligned}$$

That is, the robot's new location is always determined by the actual outcome y_c of the action $fwd(x, y_c)$;

- the likelihood axioms Σ_l for stochastic action $fwd(x, y_c)$ and noisy sensing $sonar(x)$

$$\begin{aligned} \square l(fwd(x, y_c)) &= \mathcal{N}(y_c; x, 1) \\ \square l(sonar(x)) &= \mathcal{N}(x; h, 0.25) \end{aligned}$$

Namely, always, the likelihood of $fwd(x, y_c)$ is given by a Gaussian density for the outcome y_c centered on the intended value x with a spread of 1. Likewise, the sensing model is given by a Gaussian density for the value read x centered on the true value h (the fluent that it is measuring) with a spread of 0.25. Let $\Sigma_{dyn} = \Sigma_{post} \cup \Sigma_l$ and Σ be as:

$$\forall x.\boldsymbol{B}(h \leq x \colon \mathcal{U}(x; 10, 11)) \wedge \boldsymbol{K}\Sigma_{dyn}$$

then Σ entails the following:

- $[fwd(2, 2.1)]\boldsymbol{B}(h \leq 9 \colon 0.6844)$

Intuitively, when the robot intends to forward 2 units, even if nature selects 2.1 as an outcome, the robot has no observation of this and considers $fwd(2, y_c)$ possible for all y_c with y_c distributed as $\mathcal{N}(y_c; 2, 1)$. The posterior of $h \leq 9$ amounts to the product of the prior of h and the probability of y_c subject to $h - y \leq 9$, this is because Σ_{post} says the locations after action $fwd(x, y_c)$ is given by $h - y_c$ where h is the initial location. Hence the posterior amounts to

$$\int_{\mathbb{R}} \int_{\mathbb{R}} \left(\begin{cases} \mathcal{N}(y; 2, 1) & \psi \\ 0 & \text{otherwise} \end{cases} \right) dy_c dx = 0.6844 \quad (3)$$

here ψ is the conjunction of three conditions: 1) $h = x$; 2) $x \in [10, 11]$ since h distributes uniformly in [10,11]; 3) $x - 9 \leq y_c$. Hence, $\Sigma \models [fwd(2, 2.1)]\boldsymbol{B}(h \leq 9 \colon 0.6844)$. Formally,

Proof. Suppose $e \models \Sigma$. For all $\mu \in e$. By definition, $Z^{h \leq 9}_{\mu, w, fwd(2,2.1)} = \{z' | z' \sim fwd(2, 2.1), w, z' \models h \leq 9)\} = \{a' | \exists y_c.a' = fwd(2, y_c), \|h\|^{\langle\rangle}_w - y_c \leq 9\}$. Hence $\mathbb{R}_{Z^{h \leq 9}} = \{y_c \mid \|h\|^{\langle\rangle}_w - 9 \leq y_c\}$. By Def. of $\nu^{w+}_{l^*}$,

$$\begin{aligned} \nu^{w+}_{l^*}(Z^{h \leq 9}) &= \int_{y_c \in \mathbb{R}_{Z^{h \leq 9}}} l^*(w, fwd(2, y_c)) dy_c \\ &= \int_{\mathbb{R}} \left(\begin{cases} \mathcal{N}(y_c; 2, 1) & \|h\|^{\langle\rangle}_w - y_c \leq 9 \\ 0 & \text{otherwise} \end{cases} \right) dy_c. \end{aligned}$$

[9] Here $f(+\infty) = 1$ should be understand as $\lim_{x \to \infty} f(x) = 1$ and $\lim_{x \to \infty} f(x) = y$ can be defined as $\lim_{x \to \infty} f(x) = y := \forall u.(u > 0 \supset \exists m.\forall v.(v > m) \supset |y - f(v)| < u)$.

[10] For simplicity, we do not include precondition axioms for actions and assume all actions are executable yet the usual "impossible actions" will have a zero likelihood.

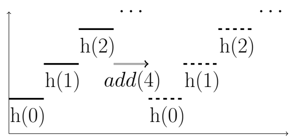

Figure 3: Infinitely many random variables $h(x)$.

Consequently $\int_w \nu_{l^*}^{w+}(Z^{h\leq 9})d\mu$ amounts to

$$\int_w \int_{\mathbb{R}} \left(\begin{cases} \mathcal{N}(y_c; 2, 1) & \|h\|_w^{\langle\rangle} - y_c \leq 9 \\ 0 & \text{otherwise} \end{cases} \right) dy_c d\mu \quad (4)$$

which is essentially the same as Eq.(3) and equal to 0.6844. (we omit the proof for space reasons) The same result holds for $\nu_{l^*}^{w-}$, hence, $e \models [fwd(2, 2.1)]\boldsymbol{B}(h \leq 11\colon 0.6844)$. □

- $[fwd(2, 2.1) \cdot sonar(8)]\boldsymbol{B}(h \leq 9\colon 0.9778)$

After the sonar reads a value less than 9, i.e. $sonar(8)$, the robot's belief of being less than 9 is enhanced. After the action sequence $fwd(2, 2.1) \cdot sonar(8)$, the robot considers the set of sequences $fwd(2, y_c) \cdot sonar(8)$ possible for all y_c. Moreover, the likelihood of such a sequence is given by $\mathcal{N}(y_c; 2, 1) \times \mathcal{N}(8; h, 0.25)$, the posterior of $h \leq 9$ is

$$\int_{\mathbb{R}} \int_{\mathbb{R}} \left(\begin{cases} \mathcal{N}(y_c; 2, 1)\mathcal{N}(8; h, 0.25) & \psi' \\ 0 & \text{otherwise} \end{cases} \right) dy_c dx \quad (5)$$

where ψ' is the conjunction of: 1) $x = h + y_c$; 2) $x \in [10, 11]$; 3) $x - 9 \leq y_c$. Condition 1) is because h here represents the location after the sequence $fwd(2, y_c) \cdot sonar(8)$, hence, its initial values would be $h + y_c$ as sensing does not change the world. After normalized by η (η can be calculated in the same way without Condition 3)), the result is 0.9778. Thus, $\Sigma \models [fwd(2, 2.1)]\boldsymbol{B}(h \leq 9\colon 0.9778)$.

4.3 Beyond Nullary Fluent

Another appealing point of our logic is that when it comes to expressing degrees of belief, fluents do not have to be nullary, which goes beyond previous proposals such as BHL, $\mathcal{DS}$, and $\mathcal{XS}$ and their variants.

As a concrete example, imagine there are infinitely many random variables $h(x)$ with $x \in \mathbb{N}^{\geq 0}$ distributed uniformly among the interval $[x, x + 1]$ which looks like an infinite ladder in Fig. 3. Supposing deterministic action $increase(b)$ will increase the values of $h(x)$ by b, then $h(x)$ should be distributed uniformly among the interval $[x + b, x + b + 1]$ after the action $increase(b)$. Formally,

Example 1. *let* Σ_{dyn} *be as:*

$$\Box[a]h(x) = v \equiv \exists x'.a = increase(x') \wedge v = h(x) + x' \vee \forall x'.a \neq increase(x') \wedge v = h(x)$$

$\Box l(increase(x')) = 1$. *Then, we have that*

- $\models \forall x, y.\boldsymbol{B}(h(x) \leq y\colon \mathcal{U}(y, x, x + 1)) \wedge \boldsymbol{K}\Sigma_{dyn} \supset [increase(4)]\forall x, y.\boldsymbol{B}(h(x) \leq y\colon \mathcal{U}(y, x + 4, x + 5))$

We comment that the result does not have to limit to uniform distributions and deterministic actions. A more complex example could be as follows: suppose there are infinitely many boxes and each has infinitely many balls, yet the number of balls $numOfBall(x)$ ($x \in \mathbb{N}^{\geq 0}$) is distributed as Poisson with an expectation of 6. Now the robot observes some balls are added to the boxes. It cannot tell the exact number of balls that are added to each box, yet the number of new balls for each box is another Poisson distribution with an expectation of 3. As a result, the robot would believe that $numOfBall(x)$ is distributed as Poisson with an expectation of 9. This can be formulated as:

Example 2. *let* Σ_{dyn} *be as:*

$$\Box[a]numOfBall(x) = v \equiv \exists y.a = add(y) \wedge v = numOfBall(x) + y \vee \forall y.a \neq add(y) \wedge v = h(x)$$

$$\Box l(add(y)) = \begin{cases} 3^y e^{-3}/y! & y \in \mathbb{N}^{\geq 0} \\ 0 & \textit{otherwise} \end{cases}$$

Then we have:

- $\models \forall x, y.\boldsymbol{B}(h(x) = y\colon 6^y e^{-6}/y!) \wedge \boldsymbol{K}\Sigma_{dyn} \supset [add(4)]\forall x, y.\boldsymbol{B}(h(x) = y\colon 9^y e^{-9}/y!)$

To sum up, our logic is capable of capturing interesting results such as the sum of two Poisson is another Poisson, the sum of two Gaussian is another Gaussian, and the product of two log-norms is another log-norm, making it extremely expressive.

5 Related Work & Conclusion

We have reviewed the most related work in the introduction. $\mathcal{PS}$ could be viewed as a continuous extension to $\mathcal{DS}$ and a dynamic extension to $\mathcal{OBL}c$. Less relevantly, the axiomatic approach BL (Belle and Levesque 2013; Belle and Levesque 2018) can also express degrees of belief in a dynamic setting (yet only for the nullary fragments). Recently, the BL framework is recast into the modal logic $\mathcal{XS}$ (Belle 2023). BL is inspired by the well-known work BHL (Bacchus, Halpern, and Levesque 1999) which integrates previous work in reasoning about knowledge and probability (Nilsson 1986; Bacchus 1989; Halpern 1990).

There are also works on limited forms of probabilistic logic (see (Belle and Levesque 2018) for a discussion), such as Bayesian networks (Pearl 1988), relational graphical models (De Raedt, Kimmig, and Toivonen 2007), and probabilistic databases (Suciu et al. 2011). They emphasize more on the reasoning side (limiting expressiveness to obtain tractability in reasoning) while our work focuses more on the representation side.

In conclusion, we propose the logic $\mathcal{PS}$, a modal logic of (continuous) degree of belief and actions. Our logic is rich in expressiveness as specifying believing arbitrary first-order formulas is possible. Our logic also has reasonable properties.

For future work, developing an account for progression is promising (Liu and Feng 2021). Moreover, investigating fragments where reasoning is tractable is also possible. Besides, it is interesting to see how the logic fits in the context of epistemic programming (Belle and Levesque 2015; Liu and Lakemeyer 2022).

Acknowledgments

This work has been supported by the Royal Society University Research Fellowship, the UKRI Strategic Priorities Fund to the UKRI Research Node on Trustworthy Autonomous Systems Governance and Regulation (EP/V026607/1, 2020-2024), the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) RTG 2236/2 'UnRAVeL', and the EU ICT-48 2020 project TAILOR (No. 952215).

References

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

Bacchus, F. I. 1989. Representing and reasoning with probabilistic knowledge. *Artificial Intelligence Series*.

Belle, V., and Lakemeyer, G. 2017. Reasoning about probabilities in unbounded first-order dynamical domains. In *IJCAI*, 828–836.

Belle, V., and Levesque, H. 2013. Reasoning about continuous uncertainty in the situation calculus. In *Twenty-Third International Joint Conference on Artificial Intelligence*. Citeseer.

Belle, V., and Levesque, H. 2015. Allegro: Belief-based programming in stochastic dynamical domains. In *International Joint Conference on Artificial Intelligence (IJCAI)*, volume 2015, 2762–2769.

Belle, V., and Levesque, H. J. 2018. Reasoning about discrete and continuous noisy sensors and effectors in dynamical systems. *Artificial Intelligence* 262:189–221.

Belle, V., and Levesque, H. J. 2020. Regression and progression in stochastic domains. *Artificial Intelligence* 281:103247.

Belle, V. 2023. Actions, continuous distributions and meta-beliefs. In *The 22nd International Conference on Autonomous Agents and Multiagent Systems, 2023*.

Boyen, X., and Koller, D. 1998. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, 33–42.

Carnap, R. 1962. *Logical foundations of probability*, volume 2. Citeseer.

de Morais, E. M., and Finger, M. 2013. Probabilistic answer set programming. In *2013 Brazilian Conference on Intelligent Systems*, 150–156. IEEE.

De Raedt, L.; Kimmig, A.; and Toivonen, H. 2007. Problog: A probabilistic prolog and its application in link discovery. In *IJCAI*, volume 7, 2462–2467. Hyderabad.

Dean, T., and Wellman, M. 1991. *Planning and Control*. The Morgan Kaufmann Series in Representation & Reasoning.

Fagin, R., and Halpern, J. Y. 1994. Reasoning about knowledge and probability. *Journal of the ACM (JACM)* 41(2):340–367.

Feng, Q.; Liu, D.; Belle, V.; and Lakemeyer, G. 2023. A logic of only-believing over arbitrary probability distributions. In *The 22nd International Conference on Autonomous Agents and Multiagent Systems, 2023*.

Gabaldon, A., and Lakemeyer, G. 2007. Esp: A logic of only-knowing, noisy sensing and acting. In *Proceedings of the national conference on artificial intelligence*, volume 22, 974. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

Gaifman, H. 1964. Concerning measures in first order calculi. *Israel journal of mathematics* 2:1–18.

Halpern, J. Y., and Shoham, Y. 1991. A propositional modal logic of time intervals. *Journal of the ACM (JACM)* 38(4):935–962.

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

Harel, D. 1984. *Dynamic logic*. Springer.

Kooi, B. P. 2003. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information* 12:381–408.

Lakemeyer, G., and Levesque, H. J. 2011. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence* 175(1):142–164.

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

Liu, D., and Feng, Q. 2021. On the progression of belief. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 18, 465–474.

Liu, D., and Lakemeyer, G. 2021. Reasoning about beliefs and meta-beliefs by regression in an expressive probabilistic action logic. In *IJCAI*, 1951–1958.

Liu, D., and Lakemeyer, G. 2022. On the verification of belief programs. *arXiv preprint arXiv:2204.12562v3*.

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

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann.

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

Stein, E. M., and Shakarchi, R. 2009. *Real analysis: measure theory, integration, and Hilbert spaces*. Princeton University Press.

Suciu, D.; Olteanu, D.; Ré, C.; and Koch, C. 2011. Probabilistic databases. *Synthesis lectures on data management* 3(2):1–180.

Thielscher, M. 2001. Planning with noisy actions (preliminary report). In *AI 2001: Advances in Artificial Intelligence: 14th Australian Joint Conference on Artificial Intelligence Adelaide, Australia, December 10–14, 2001 Proceedings 14*, 495–506. Springer.

Turing, A. 1936. On computable numbers, with an application to the entscheidungsproblem. *J. of Math* 58:345–363.