A Partition-Based First-Order Probabilistic Logic to Represent Interactive Beliefs

Alessandro Panella and Piotr Gmytrasiewicz

University of Illinois at Chicago
Department of Computer Science
Chicago, IL 60607
{apanella,piotr}@cs.uic.edu

Abstract. Being able to compactly represent large state spaces is crucial in solving a vast majority of practical stochastic planning problems. This requirement is even more stringent in the context of multi-agent systems, in which the world to be modeled also includes the mental state of other agents. This leads to a hierarchy of beliefs that results in a continuous, unbounded set of possible interactive states, as in the case of Interactive POMDPs. In this paper, we describe a novel representation for interactive belief hierarchies that combines first-order logic and probability. The semantics of this new formalism is based on recursively partitioning the belief space at each level of the hierarchy; in particular, the partitions of the belief simplex at one level constitute the vertices of the simplex at the next higher level. Since in general a set of probabilistic statements only partially specifies a probability distribution over the space of interest, we adopt the maximum entropy principle in order to convert it to a full specification.

1 Introduction

One of the main problems to be faced in the field of stochastic planning is the curse of dimensionality. Traditional methods based on the enumeration of the state, action, and observation spaces have shown to be unpractical for all but the simplest settings. Factorizing the description of the domain into "features", like in a Bayesian network, has been a prominent direction of research that has led to outstanding results [3, 19]. Yet, in many real-world problems this approach does not suffice, because of the large number of such features. Hence the need to lift the representation from the propositional level to a more abstract level, by exploiting the synergy between first-order logic (FOL) and probability theory, allowing to compactly summarize the regularities of the domain and the interactions between objects. Several applications of this paradigm to MDPs can be found in literature (e.g. [4, 22, 25]). On the other hand, only little work has surfaced that focuses on lifted first-order inference for representing and solving POMDPs [23, 26].

A compact representation of the domain is even more necessary in the context of decision making in partially observable, multi-agent environments, in which an agent needs to model the mental states (beliefs, preferences, and intentions) of other agents, in addition to the physical world. In particular, we focus on the Interactive POMDP (I-POMDP) framework [8], in which the agent maintains a belief about the other agents' *types*, intended as the set of private information involved in their decision making. Each type includes the agent's own $belief$, which is a probability distribution over the state of the world and, recursively, other agents' types. In this paper, we focus on the representation of such interactive beliefs, limiting the discussion to a setting with two agents, i and j. The generalization to scenarios with more than two agents is straightforward. Because of the impossibility of representing infinitely nested beliefs in a finite space, Gmytrasiewicz and Doshi [8] define finitely nested I-POMDPs as a specialization of the infinitely nested ones.

Let us denote as $\Delta(\cdot)$ the regular simplex over the set given as argument. The *interactive state space* at nesting level n for agent i, denoted $IS_{i,n}$, is inductively defined as:

$$\begin{aligned} IS_{i,0} &= S \\ IS_{i,1} &= S \times \Delta(IS_{j,0}) \\ &\vdots \\ IS_{i,n} &= S \times \Delta(IS_{j,n-1}) \end{aligned} \tag{1}$$

One problem of I-POMDPs is that the set of possible beliefs of the other agent is uncountable and unbounded as soon as nesting level 2 is reached [6]. This makes it impossible to even represent interactive beliefs, in that they are not computable functions. For this reason, being able to abstract over the regularities of the interactive state space in order to provide a finite representation is of utmost benefit.

In this paper, we describe for the first time a First-Order Probabilistic Language to express interactive beliefs. The approach is conceived in the context of I-POMDPs, but grows out to be a general representation of nested probability distributions that can be applied in a variety of contexts in multi-agent systems. The main idea is to *recursively partition the belief space into regions, building the belief simplex at the next level over the partitions of the belief simplex at the lower level*, as intuitively depicted in Fig. 1. In this way, we can provide a compact, finite representation of interactive beliefs.

A belief is represented as a set of (probabilistic) sentences, each associated with a probability. These sentences are not required to be non-overlapping, in that the agent should be free to express his belief about any arbitrary statement. As a result of this augmented freedom, a belief base constitutes in general only a set of constraints rather than a full specification of a probability distribution. For this reason, we adopt the well-known *maximum entropy principle* in order to provide a unique distribution associated with the belief base.

The paper is organized as follows. In Sect. 2 we provide a brief survey of the existing related work. Section 3 presents our contribution, describing the probabilistic-logical framework for interactive beliefs in a bottom-up fashion and providing examples to clarify the concepts. Section 4 concludes the paper and hints at directions for future research.

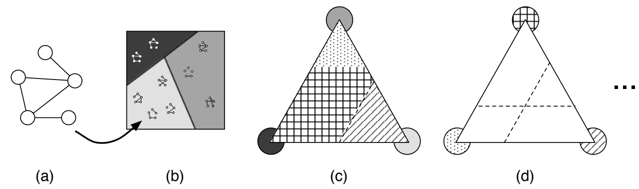

Fig. 1. Qualitative representation of the belief hierarchy. (a) Real state of the world (objects and relations.) (b) Set of possible world states. (c) Agent i's belief simplex over partitions of world states. (d) Agent j's belief simplex over partitions of i's simplex. Note that the partitions at one level are the vertices of the simplex at the next level; this association is represented by matching colors (b-c) and patterns (c-d.)

2 Related Work

The integration of first-order logic and probability theory has been an important area of research since the mid and late 80's. Nilsson [18] proposes a probabilistic logic in which a probability value is attached to logic sentences, either propositional or first-order, belonging to a probabilistic knowledge base. He devises a linear problem that, given any query sentence, computes its probability intervals that are consistent with knowledge base. Some years later, Bacchus [2] and Halpern [9] describe how first-order probabibilistic logic comes in two flavors: to express probabilities on possible worlds, such as in the sentence "Tweety flies with probability 0.9," and to express statistical knowledge, as in "90% of birds fly." Subsequent work in probabilistic languages has mostly adopted the first type of semantics. This early approaches provide theoretical basis for the field, but lack practical inference algorithms

The work on probabilistic logic has evolved in what has been recently named Statistical Relational AI (Star AI), that includes a number of different approaches, of which we report a few examples. Koller and Pfeffer [15] define Probabilistic Relational Models (PRMs), borrowing the semantics from relational databases. Like databases, PRMs model complex domains in terms of entities and properties. Moreover, by incorporating a directed causal relations like in Bayesian networks, PRMs allow to express uncertainty over the properties of entities and the relations among them (relational uncertainty.) Markov Logic Networks [20] are collections of first-order logic formulae that are assigned a weight. The atomic formulae appearing in such set constitute the vertices of a Markov Network, whose edges correspond to the logical connectives. The weights determine the potential function assigned to each groups of vertices that compare in the same original formula. In [17] the authors introduce Bayesian Logic (BLOG), a generalization of Bayesian network similar to RPMs that assume an open universe, i.e. the number of objects is unknown and potentially infinite. BLOG models are described by means of a generative semantics that allows to deal with domains of unknown size. Another line of work studies probabilistic logic programs [16] and relational probabilistic conditionals [14]. These approaches adopt the maximum entropy principle to provide semantics to probabilistic knowledge bases.

The study interactive belief systems has been subject of substantial research in the field of game theory and multi-agent systems in general, especially since the introduction of games of incomplete information [12]. Several works [11] study the use of modal logic and its derivations in order to describe players' knowledge about other players' knowledge. Probabilistic extension to modal logic have been proposed [7, 10, 24], and are based on commonly known prior probability distribution on possible worlds and accessibility relation. Aumann [1] describes an approach that embeds knowledge and probabilistic beliefs in the context of interactive epistemology. The work on Interactive POMDPs [8, 5] rejects the common knowledge assumption and proposes a hierarchy of probabilistic beliefs that we take as our starting point, as already described in the introductory section.

3 Probabilistic First-Order Logic to Represent Interactive Beliefs

We begin this section by describing the logic setup. The semantics of first-order logic used here assumes closed universe and unique names. Let $Q = \{q_1, q_2, \ldots, q_{|Q|}\}$ be a *set of predicates*, and let $\rho : Q \rightarrow \mathbb{N}$ be a function that associates each predicate to its *arity*. Given a *domain* (set of constants) $D = \{d_1, d_2, \ldots d_{|D|}\}$, define the *set of ground predicates* G (the Herbrand base), corresponding to each possible instantiation of predicates in the domain, i.e. $G = \{q(d_1, \ldots, d_{\rho(q)}) : q \in Q, d_i \in D \;\; \forall i = 1, \ldots, \rho(q)\}$. An interpretation of Q in domain D is a function $\sigma : G \rightarrow \{T, F\}$ that assigns a truth value to every ground predicate. The set of possible *states of the world* S corresponds to the set of all possible interpretations.

Given a first-order logic sentence ϕ, we denote as $S(\phi)$ the subset of S for which ϕ is true, i.e. the set of models of ϕ (under the usual definition of FOL entailment.) Given a set of FOL sentences Φ, we denote as $S(\Phi)$ the collection of sets of models of the formulae in Φ, i.e. $S(\Phi) = \{S(\phi) : \phi \in \Phi\}$.

3.1 Level 0 Beliefs

In this section, we describe how to represent 0-th level beliefs, i.e. the belief an agent holds about the state of the world. The approach is similar to the one described in [18], that we take as our starting point.

Definition 1 (Level-0 Belief Base). *A Level-0 Belief Base (L0-BB) $\mathcal{B}^{i,0}$ for agent i is a set of pairs of the form $\langle \phi_k, \alpha_k \rangle$, for $k = 1, 2, \ldots, m$, where ϕ_k is a sentence in first-order logic, and α_k a real number between 0 and 1.*

For each pair, α_k intuitively represents i's degree of belief about sentence ϕ_k. In addition to simple pairs, we allow universally quantified expressions of the type $\forall \mathbf{x} \langle \phi(\mathbf{x}), \alpha \rangle$ to appear in the belief base, where $\mathbf{x} = \langle x_1, \ldots, x_l \rangle$ is a tuple of logical variables that are free in FOL formula ϕ. Semantically, this expression is equivalent to the set of pairs resulting from the propositionalization of ϕ, that is:

$$\forall \mathbf{x} \ \langle \phi(\mathbf{x}), \alpha \rangle \equiv \left\{ \langle \textsc{Subst}(\mathbf{x}/\mathbf{d}, \phi), \alpha \rangle : \mathbf{d} \in D^{|\mathbf{x}|} \right\} \ , \tag{2}$$

where $\textsc{Subst}(\mathbf{x}/\mathbf{d}, \phi))$ represent the FOL sentence resulting from ϕ by substituting the tuple of logical variables $\mathbf{x}$ with domain elements $\mathbf{d}$, adopting the notation used in [21]. In the following, we will consider L0-BB's in which all universally quantified pairs have been proposizionalized.

Let $\Phi_{\mathcal{B}} = \{\phi_1, \phi_2, \ldots, \phi_m\}$ be the set of FOL sentences involved in the pairs of a L0-BB $\mathcal{B}^{i,0}$. As mentioned earlier, we do not require the elements of $\Phi_{\mathcal{B}}$ to identify non-overlapping regions of S. Instead, we compute the partitions that are induced by such regions, i.e. every possible overlap, and form a probability distribution on such partitioning. We hence define the set of logical partitions as

$$\Psi_{\mathcal{B}} = \left\{ \bigwedge_{\phi \in \Phi_I} \phi \setminus \bigvee_{\phi \in I \Phi_I^C} \phi : \Phi_I \subseteq \Phi_{\mathcal{B}} \cup \{\top\} \right\} \ , \tag{3}$$

where $\Phi_I^C = (\Phi_{\mathcal{B}} \cup \top) \setminus \Phi_I$. We denote as $\Psi_{\mathcal{B}}(\phi)$ the set of partitions whose union is ϕ:

$$\Psi_{\mathcal{B}}(\phi) = \{\psi \in \Psi_{\mathcal{B}} : S(\psi) \subseteq S(\phi)\} \tag{4}$$

The concept of satisfiability of a belief base is formalized in the following definitions.

Definition 2 (Satisfiability). *Given a L0-BB $\mathcal{B}^{i,0}$, a probability distribution $p_{i,0}$ over the set of logical partitions $\Psi_{\mathcal{B}}$ is said to* satisfy *$\mathcal{B}^{i,0}$ if, for all $\phi_k \in \Phi_{\mathcal{B}}$, it is true that*

$$\sum_{\psi \in \Psi_{\mathcal{B}}(\phi)} p_i(S(\psi)) = \alpha_k \tag{5}$$

Definition 3 (Consistency). *A L0-BB $\mathcal{B}^{i,0}$ is said to be* consistent *(or* satisfiable*) if there exists a probability distribution $p_{i,0}$ over $\Psi_{\mathcal{B}}$ that satisfies it.*

In general, there exist multiple distributions $p_{i,0}$ satisfying a L0-BB $\mathcal{B}^{i,0}$.[1] This is due to the fact that a L0-BB constitutes in general only a partial specification of a probability distribution over S, as noted by Nilsson [18]. There are different ways to cope with this indeterminacy. A "skeptical" approach is to compute the upper and lower bounds of the probability of each state, and consider such intervals. A more "credulous" solution is to pick one probability distribution among the ones that are consistent with the L0-BB. We follow the latter direction by choosing the maximum entropy (max-ent) distribution [13].

[1] In fact, it can be shown that a L0-BB either has a unique model or admits an uncountably infinite set of models.

Given a L0-BB $\mathcal{B}^{i,0}$, the max-ent probability distribution $p_{i,0}$ over $\Psi_{\mathcal{B}}$ that satisfies $\mathcal{B}^{i,0}$ is given by the solution to the following optimization problem:

$$\max_{p_{i,0}} \Big(- \sum_{\psi \in \Psi_{\mathcal{B}}} p_{i,0}(S(\psi)) \log p_{i,0}(S(\psi)) \Big) \tag{6}$$

subject to

$$\begin{aligned} \textstyle\sum_{\psi \in \Psi_{\mathcal{B}}(\phi)} p_{i,0}(S(\psi)) &= \alpha_k \qquad && \forall\, \phi \in \Phi_{\mathcal{B}} \\ \textstyle\sum_{\psi \in \Psi_{\mathcal{B}}} p_{i,0}(S(\psi)) &= 1 && \\ p_{i,0}(S(\psi)) &\geq 0 && \forall \psi \in \Psi_{\mathcal{B}} \end{aligned} \tag{7}$$

Hence, a L0-BB $\mathcal{B}^{i,0}$ represents the probability space $(S(\Psi_{\mathcal{B}}), 2^{S(\Psi_{\mathcal{B}})}, p_{i,0})$, were $p_{i,0}$ is the max-ent probability distribution just defined. Slightly abusing notation, we will sometimes write $p_{i,0}(\psi)$ rather than $p_{i,0}(S(\psi))$ when this is not source of ambiguity. We clarify the concepts presented in this section by introducing a simple running example.

Example (Grid world) In a world consisting of an $n \times n$ grid, an agent i wants to tag a moving target j. The agent knows his own position, but is uncertain about the target's. The predicate $jPos(x, y)$ indicates the target position, where x and y are integers representing the coordinates of a location on the grid, i.e. $0 \leq x, y < n$. Obviously, the target occupies one and only one position in the grid; this can be expressed by the following FOL sentence:

$$\exists x, y \Big(jPos(x, y) \wedge \neg \exists w, z \big(jPos(x, y) \wedge jPos(w, z) \wedge (x \neq y \vee w \neq z) \big) \Big) \tag{8}$$

Instead of including this fact in the belief base, we assume the agent implicitly knows it, as a limitation on possible worlds. We introduce the auxiliary deterministic predicates $geq(x, k) \equiv x \geq k$ and $leq(x, k) \equiv x \leq k$. Say the agent is interested, in a particular moment in time, to the horizontal location of the target in the grid with respect to the center, i.e. whether the target is in either the left or right half of the map. A plausible L0-BB representing i's belief is:

$$\mathcal{B}^{i,0} = \boxed{\begin{array}{l} \langle \exists x, y (jPos(x, y) \wedge leq(x, \lfloor n/2 \rfloor)),\ 0.8 \rangle \\ \hline \langle \exists x, y (jPos(x, y) \wedge geq(x, \lfloor n/2 \rfloor)), 0.5 \rangle \end{array}} \tag{9}$$

We denote the two FOL sentences in the L0-BB as ϕ_0 and ϕ_1, respectively. The partitioning $\Psi_{\mathcal{B}} = \{\psi_0, \psi_1, \psi_2\}$ induced on the state space is shown in Fig. 2-a. In this case, there is only one probability distribution over $\Psi_{\mathcal{B}}$ that is consistent with the belief base, namely $p_{i,0} = (0.5, 0.3, 0.2)$.

3.2 Level 1 Beliefs

After having described how to represent the beliefs of agent i about the state of the world using first-order logic and probability, we now introduce how to specify agent i's beliefs about agent j's beliefs about S. In order to do so, agent

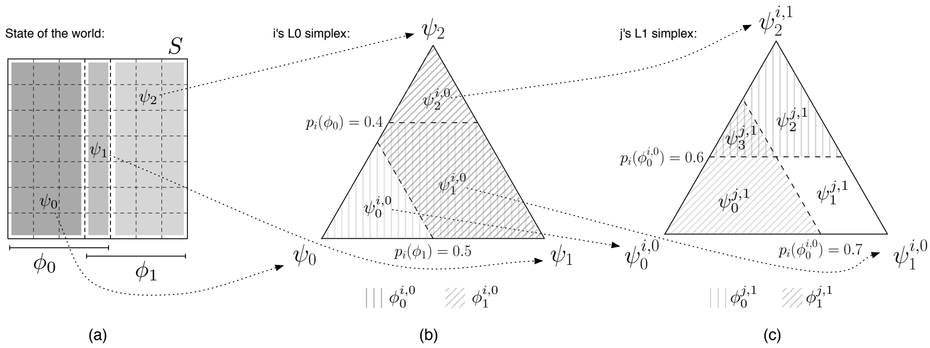

Fig. 2. Example of nested First-Order Belief Base.

i needs a language that allows him to abstract over the space of agent j's 0-th level beliefs, in the same way first-order logic provides abstraction over the set of states of the world in a L0-BB. We call this language *Level-0 First-Order Probabilistic Logic* (L0 FOPL), in that it provides a way to describe j's level 0 beliefs.

Definition 4 (L0 FOPL). *The language $\mathscr{L}^{j,0}$ of j's level 0 probabilistic statements is recursively defined as:*

1. *$P_j(\phi) \bigtriangleup \beta$ is a formula, where ϕ is a formula in first-order logic, $\beta \in [0,1]$, and $\bigtriangleup \in \{<, \leq, =, \geq, >\}$;*
2. *If $\phi^{j,0}$ is a formula, then $\neg\phi^{j,0}$ is a formula;*
3. *If $\phi_1^{j,0}$ and $\phi_2^{j,0}$ are formulae, then $\phi_1^{j,0} \wedge \phi_2^{j,0}$, $\phi_1^{j,0} \vee \phi_2^{j,0}$, $\phi_1^{j,0} \Rightarrow \phi_2^{j,0}$ are formulae;*
4. *If $\phi^{j,0}$ is a formula, then $\exists\mathbf{x}(\phi^{j,0})$ and $\forall\mathbf{x}(\phi^{j,0})$ are formulae, where $\mathbf{x}$ is a subset of the free logical variables of $\phi^{j,0}$;*
5. *A sentence is a formula with no free variables.*

Agent i assigns degrees of belief to some sentences of $\mathscr{L}^{j,0}$. This is represented as a Level 1 Belief Base, defined in the following.

Definition 5 (Level 1 Belief Base). *A* Level 1 Belief Base *(L1-BB) $\mathcal{B}^{i,1}$ for agent i is a collection of pairs $\langle\phi_k^{j,0}, \alpha_k\rangle$, for $k = 1, 2, \ldots, m$, where $\phi_k^{j,0}$ is a sentence of $\mathscr{L}^{j,0}$ and $\alpha_k \in [0,1]$.*

For the sake of presenting the semantics of a L1-BB, we will assume that the quantified statements about j's probabilities (of the type of rule 4 in Definition 4) are expanded into propositional form over the elements of the domain. The basic idea behind the semantics of a L1-BB is that *the partitions over the state of the world correspond to the vertices of the simplex of j's belief about the state of the world.* In turn, agent i maintains a distribution over the partitions of such simplex. This mechanism in intuitively depicted in Fig. 3.

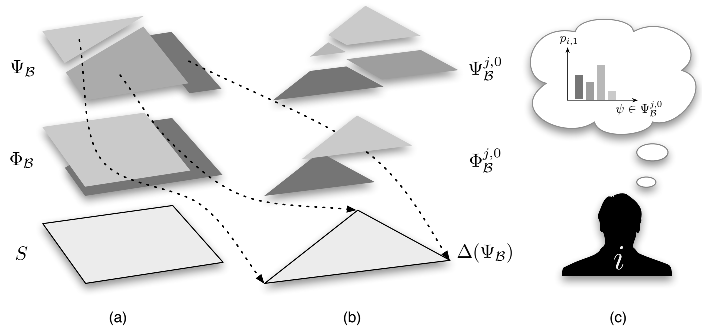

Fig. 3. Intuitive representation of the partition-based belief hierarchy and corresponding notation. (a) States of the world. (b) Agent j's 0-th level beliefs. (c) Agent i maintains a distribution over $\Psi_{\mathcal{B}}^{j,0}$.

We now formalize this process. From the belief base, let us define the set $\Phi_{\mathcal{B}}^{j,0}$ of j's level 0 probabilistic statements about which i holds a degree of belief in $\mathcal{B}^{i,0}$ (Fig 3-a). Formally, we have:

$$\Phi_{\mathcal{B}}^{j,0} = \{\phi^{j,0} : \langle \phi^{j,0}, \alpha \rangle \in \mathcal{B}^{i,1} \text{ for some } \alpha \in [0,1]\} \tag{10}$$

In turn, we also define the set of FOL sentences $\Phi_{\mathcal{B}}$ that appear in some of j's level-0 probabilistic statements. This will allow us to describe the space of j's 0-th level beliefs that agent i considers (Fig 3(b)). Formally, we have:

$$\Phi_{\mathcal{B}} = \{\phi : \phi \circ \phi^{j,0} \text{ for some } \phi^{j,0} \in \Phi_{\mathcal{B}}^{j,0}\} \ , \tag{11}$$

where the circle symbol is read "occurs in" and is recursively defined as:

1. If $\phi^{j,0}$ is of the form $P_j(\phi) \vartriangle \alpha$, then $\phi \circ \phi^{j,0}$;
2. If $\phi \circ \phi_u^{j,0}$, then $\phi \circ (\neg \phi_u^{j,0})$;
3. If $\phi \circ \phi_u^{j,0}$, then $\phi \circ (\phi_u^{j,0} \wedge \phi_v^{j,0})$ and $\phi \circ (\phi_u^{j,0} \vee \phi_v^{j,0})$, for any $\phi_v^{j,0}$.

We denote as $\bigcirc(\phi^{j,0})$ the set of of $\phi \in \Phi_{\mathcal{B}}$ such that $\phi \circ \phi^{j,0}$. The set $\Psi_{\mathcal{B}}$ of partitions of S induced by $\Phi_{\mathcal{B}}$ is defined as in (3), and the set $\Psi_{\mathcal{B}}(\phi)$ as in (4). Let us denote as $\Delta(\Psi_{\mathcal{B}})$ the regular simplex whose vertices are the elements of $\Psi_{\mathcal{B}}$. This simplex is the set of all possible probability distributions that j may hold about the sentences in $\Psi_{\mathcal{B}}$. The semantics of $\mathscr{L}^{j,0}$ is defined on such simplex, as described in the following definition.

Definition 6 (L1-BB Probabilistic entailment). *Given a L1-BB $\mathcal{B}^{i,1}$, the probabilistic entailment ($\models$) of a sentence $\phi^{j,0}$ (such that $\bigcirc(\phi_{j,0}) \in \Phi_{\mathcal{B}}$) by a point $p_{j,0} \in \Delta_{\mathcal{B}}(\Psi_{\mathcal{B}})$ is recursively defined as:*

1. *$p_{j,0} \models (P_j(\phi) \vartriangle \beta)$ if and only if $\sum_{\psi \in \Psi_{\mathcal{B}}(\phi)} p_{j,0}(S(\psi)) \vartriangle \beta$;*
2. *$p_{j,0} \models \neg \phi^{j,0}$ if and only if $p_{j,0} \not\models \phi^{j,0}$;*

3. $p_{j,0} \models (\phi_u^{j,0} \wedge \phi_v^{j,0})$ *if and only if* $p_{j,0} \models \phi_u^{j,0}$ *and* $p_{j,0} \models \phi_v^{j,0}$*;*
4. $p_{j,0} \models (\phi_u^{j,0} \vee \phi_v^{j,0})$ *if and only if* $p_{j,0} \models (\neg\phi_u^{j,0} \wedge \neg\phi_v^{j,0})$*;*

If $p_{j,0} \models \phi^{j,0}$, $p_{j,0}$ is said to be a *model* of sentence $\phi^{j,0}$. It is easy to see that the models of a probabilistic sentence $\phi^{j,0}$ are the points in the continuous region of the simplex that is identified by $\phi^{j,0}$. We refer to the set of models of a sentence $\phi^{j,0}$ as $[\Delta(\Psi_{\mathcal{B}})](\phi^{j,0})$, to remark that it corresponds to a subset of the simplex $\Delta(\Psi_{\mathcal{B}})$. To make the notation lighter, we will usually refer to the set of models as $\Delta_{\mathcal{B}}(\phi^{j,0})$. We add the extra symbol $\top^{j,0}$ to the language $\mathscr{L}^{j,0}$ as a shorthand for $P_j(\top) = 1$. Clearly, $\Delta_{\mathcal{B}}(\top^{j,0}) = \Delta(\Psi_{\mathcal{B}})$ (the whole simplex.)

Similarly to (3), we now define the logical probabilistic partitioning induced by $\Phi_{\mathcal{B}}^{j,0}$ as:

$$\Psi_{\mathcal{B}}^{j,0} = \left\{ \bigwedge_{\phi^{j,0} \in \Phi_I^{j,0}} \phi^{j,0} \setminus \bigvee_{\phi^{j,0} \in \Phi_I^{j,0,C}} \phi^{j,0} : \Phi_I^{j,0} \subseteq (\Phi_{\mathcal{B}}^{j,0} \cup \top^{j,0}) \right\}, \tag{12}$$

where $\Phi_I^{j,0,C} = (\Phi_{\mathcal{B}}^{j,0} \cup \top^{j,0}) \setminus \Phi_I^{j,0}$.

The sets $\Delta_{\mathcal{B}}(\psi^{j,0})$ for $\psi^{j,0} \in \Psi_{\mathcal{B}}^{j,0}$ are non-overlapping regions that together cover the simplex $\Delta_{\mathcal{B}}(\Psi_{\mathcal{B}})$ entirely, hence they form a partitioning of such space.

We denote as $\Psi_{\mathcal{B}}^{j,0}(\phi^{j,0})$ the set of elements of $\Psi_{\mathcal{B}}^{j,0}$ that correspond to subsets of $\Delta_{\mathcal{B}}(\phi^{j,0})$, where $\phi^{j,0} \in \Phi_{\mathcal{B}}^{j,0}$. Formally,

$$\Psi_{\mathcal{B}}^{j,0}(\phi^{j,0}) = \left\{\psi^{j,0} \in \mathcal{B}^{j,0} : \Delta_{\mathcal{B}}(\psi^{j,0}) \subseteq \Delta_{\mathcal{B}}(\phi^{j,0})\right\} \tag{13}$$

We now describe the max-ent distribution over the belief simplex induced by a L1-BB. Note that we are defining a distribution over the set of j's distribution about the state of the world, which is a continuous space. Nevertheless, instead of defining a probability directly over the simplex $\Delta(\Psi_{\mathcal{B}})$, we will use the space of partitions of this simplex, namely $\Delta_{\mathcal{B}}(\Psi_{\mathcal{B}}^{j,0})$. This allows the agent to represent his first level interactive belief as a discrete probability distribution, rather than a continuous one. Formally, we define the probability space

$$\left(\Delta_{\mathcal{B}}(\Psi_{\mathcal{B}}^{j,0}), 2^{\Delta_{\mathcal{B}}(\Psi_{\mathcal{B}}^{j,0})}, p_{i,1}\right) \ , \tag{14}$$

where $p_{i,1}$ is the result of the following optimization problem:

$$\max\left(- \sum_{\psi^{j,0} \in \Psi_{\mathcal{B}}^{j,0}} p_{i,1}(\Delta_{\mathcal{B}}(\psi^{j,0})) \log p_{i,1}(\Delta_{\mathcal{B}}(\psi^{j,0}))\right) \tag{15}$$

subject to:

$$\begin{array}{rcll} \sum_{\psi^{j,0} \in \Psi_{\mathcal{B}}^{j,0}(\phi_k^{j,0})} p_{i,1}(\Delta_{\mathcal{B}}(\psi^{j,0})) &=& \alpha_k & \forall\ \phi_k^{j,0} \in \Phi_{\mathcal{B}}^{j,0} \\ p_{i,1}(\Delta_{\mathcal{B}}(\psi^{j,0})) &\geq& 0 & \forall\ \psi^{j,0} \in \Psi_{\mathcal{B}}^{j,0} \\ \sum_{\psi^{j,0} \in \Psi_{\mathcal{B}}^{j,0}} p_{i,1}(\Delta_{\mathcal{B}}(\psi^{j,0})) &=& 1 & \end{array} \tag{16}$$

As before, we will sometimes write $p_{i,1}(\psi^{j,0})$ when this slight abuse of notation does not generate ambiguities.

Example (Grid world, cont'd) In the grid world example introduced before, suppose the moving target j is an agent on its own, maintaining a probability distribution over S and over i's belief about S. We assume j knows that i is concerned about his horizontal position w.r.t. the center. A possible L1-BB is:

$$B^{j,1} = \begin{array}{|c|}\hline \langle P_i(\phi_0) \geq 0.4, 0.4\rangle \\ \hline \langle P_i(\phi_1) > 0.5, 0.7\rangle \\ \hline \end{array} \tag{17}$$

where ϕ_0 and ϕ_1 are the same as in (9). We identify the set

$$\Phi_{\mathcal{B}}^{i,0} = \{P_i(\phi_0) \geq 0.5, P_i(\phi_1) < 0.4\} = \{\phi_0^{i,0}, \phi_1^{i,0}\} \tag{18}$$

The probabilistic sentences $\phi_0^{i,0}$ and $\phi_1^{i,0}$ induce three partitions on i's L0 belief simplex, as shown in Fig. 2-b. Again, there is only one consistent probability distribution over $\Phi_{\mathcal{B}}^{i,0}$, namely $p_{j,1} = (0.3, 0.1, 0.6)$.

3.3 Level n Beliefs

In this section, we follow the same steps as for the L1-BB and generalize the approach to any level of nesting. Intuitively, we need to represent agent i's degree of belief about agent j's beliefs about i's beliefs about... and so on, down to level 0 beliefs about the state of the world. In order to do so, we build a logical-probabilistic framework that allows the definition of a probability distribution over the other agent's $(n-1)$-th level beliefs in a recursive fashion. The result will be that, by partitioning the belief simplices at each level of the hierarchy, we can provide a finite representation of the interactive beliefs at any level of nesting. The intuition behind this process is that *the partitions of the belief simplex at any level $n-1$ corresponds to the vertices of the simplex at level n*. Since we need to specify the beliefs over some simplex at level $n-1$, we begin the description by formally defining the language of *Level $n-1$ First-Order Probabilistic Logic* (L$(n-1)$ FOPL) for agent j.

Definition 7 (L$(n-1)$ FOPL). *Given a set of predicate symbols Q and a domain D, the language $\mathscr{L}^{j,n-1}$ of j's level $n-1$ probabilistic statements is recursively defined as:*

1. *$P_j(\phi^{i,n-2}) \triangle \beta$ is a formula, where $\phi^{i,n-2}$ is a formula of language $\mathscr{L}^{i,n-2}$, $\beta \in [0,1]$, and $\triangle \in \{<, \leq, =, \geq, >\}$;*
2. *If $\phi^{j,n-1}$ is a formula, then $\neg\phi^{j,n-1}$ is a formula;*
3. *If $\phi_1^{j,n-1}$ and $\phi_2^{j,n-1}$ are formulae, then $\phi_1^{j,n-1} \wedge \phi_2^{j,n-1}$, $\phi_1^{j,n-1} \vee \phi_2^{j,n-1}$, $\phi_1^{j,n-1} \Rightarrow \phi_2^{j,n-1}$ are formulae;*
4. *If $\phi^{j,n-1}$ is a formula, then $\exists\mathbf{x}(\phi^{j,n-1})$ and $\forall\mathbf{x}(\phi^{j,n-1})$ are formulae, where $\mathbf{x}$ is a subset of the free logical variables of $\phi^{j,n-1}$;*
5. *A sentence is a formula with no free variables.*

For convenience, we use the symbol $\top^{j,n-1}$ as a shorthand for $P_j(\top^{i,n-2}) = 1$. We represent i's beliefs about some sentences of the language $\mathscr{L}^{j,n-1}$ as a Level n Belief Base.

Definition 8 (Level n Belief Base). *A* Level n Belief Base (Ln-BB) $\mathcal{B}_i^n$ *for agent i is a collection of pairs $\langle \phi_k^{j,n-1}, \alpha_k \rangle$, for $k = 1, 2, \ldots, m$, where $\phi^{j,n-1}$ is a sentence of $\mathscr{L}^{j,n-1}$ and $\alpha_k \in [0,1]$.*

In order to be able to define the probabilistic semantics of a Ln-BB, we propositionalize each quantified probabilistic statement down the nesting hierarchy.

A Ln-BB represents agent i's degree of belief about a number of probabilistic statements regarding j's beliefs. These are the level $n-1$ probabilistic statements that appear in the tuples of the Ln-BB. Formally, we define the set of such statements as:

$$\Phi_{\mathcal{B}}^{j,n-1} = \{\phi^{j,n-1} : \langle \phi^{j,n-1}, \alpha \rangle \in \mathcal{B}^{i,n} \text{ for some } \alpha \in [0,1]\} \tag{19}$$

Note that this is analogous to the set $\Phi_{\mathcal{B}}^{j,0}$ defined in (10) for a L1-BB. As in the previous subsection, we now need to provide a notion of probabilistic entailment for the language $\mathscr{L}^{j,n-1}$. We consider the regular simplex whose vertices are i's $(n-2)$-th level probabilistic statements that appear in the belief base. Formally, we define the set $\Phi_{\mathcal{B}}^{i,n-2}$ of sentences of $\mathscr{L}^{i,n-2}$ that occur in some element of $\Phi_{\mathcal{B}}^{j,n-1}$. Mathematically,

$$\Phi_{\mathcal{B}}^{i,n-2} = \{\phi^{i,n-2} : \phi^{i,n-2} \circ \phi^{j,n-1} \text{ for some } \phi^{j,n-1}\} \ , \tag{20}$$

where again we use the $\circ$ operator introduced in the previous section, generalized to $\mathscr{L}^{j,n-1}$ (we omit the full definition for conciseness.) We now consider the set of logical partitions induced by $\Phi_{\mathcal{B}}^{i,n-2}$:

$$\Psi_{\mathcal{B}}^{i,n-2} = \left\{ \bigwedge_{\phi^{i,n-2} \in \Phi_I^{i,n-2}} \phi^{i,n-2} \setminus \bigvee_{\phi^{i,n-2} \in \Phi_I^{i,n-2,C}} \phi_h^{i,n-2} \right. \\ \left. : \Phi_I^{i,n-2} \subseteq (\Phi_{\mathcal{B}}^{i,n-2} \cup \top^{i,n-2}) \right\} \ , \tag{21}$$

where $\Phi_I^{i,n-2,C} = (\Phi_{\mathcal{B}}^{i,n-2} \cup \top^{i,n-2}) \setminus \Phi_I^{i,n-2}$.

The regular simplex that has the elements of $\Psi_{\mathcal{B}}^{i,n-2}$ as vertices, denoted $\Delta(\Psi_{\mathcal{B}}^{i,n-2})$, is the space of j's $n-1$ level probability distributions. Hence, it represents the set over which the Ln-BB of agent i induces a max-ent distribution (remember that we are defining probability distributions over probability distributions over...) Instead of considering a distribution over this continuous space, we consider a distribution over partitions of such space. To do so, we first need the notion of probabilistic entailment for this case, that is a straightforward generalization of the level 1 entailment defined in the previous section, and is not reported here for brevity.

A distribution $p_{j,n-1}$ that entails a sentence $\phi^{j,n-1}$ is said to be a *model* of $\phi^{j,n-1}$. The set of models of $\phi^{j,n-1}$ is denoted as $[\Delta(\Psi_{\mathcal{B}}^{i,n-2})](\phi^{j,n-1})$, and is usually abbreviated as $\Delta_{\mathcal{B}}(\phi^{j,n-1})$.

Each element of $\Phi_{\mathcal{B}}^{j,n-1}$ corresponds therefore to a region of the simplex $\Delta(\Psi_{\mathcal{B}}^{i,n-2})$. In order to obtain the max-ent probability distribution encoded in the Ln-BB we need to define the probability space given by the partitions induced by the sentences $\phi^{j,n-1} \in \Phi_{\mathcal{B}}^{j,n-1}$. The set of logical partitions induced by $\Phi_{\mathcal{B}}^{j,n-1}$ is defined as in (21), by substituting $(n-2)$ with $(n-1)$, and i with j. We do not report the complete definition for brevity.

As before, we also define the set of logical partitions $\Psi_{\mathcal{B}}^{j,n-1}(\phi^{j,n-1})$ whose union is the set $\Delta_{\mathcal{B}}(\phi^{j,n-1})$. At this point, we are ready to introduce the max-ent probability distribution over j's belief partitions given by the Ln-BB. To this sake, we define the probability space

$$\left(\Delta_{\mathcal{B}}(\Psi_{\mathcal{B}}^{j,n-1}), 2^{\Delta_{\mathcal{B}}(\Psi_{\mathcal{B}}^{j,n-1})}, p_{i,n}\right) \ , \tag{22}$$

where $p_{i,n}$ is the solution to the following optimization problem:

$$\max_{p_{i,n}} \left(- \sum_{\psi^{j,n-1} \in \Psi_{\mathcal{B}}^{j,n-1}} p_{i,n}(\psi^{j,n-1}) \log p_{i,n}(\psi^{j,n-1}) \right) \tag{23}$$

subject to the constraints:

$$\begin{aligned} \sum_{\substack{\psi^{j,n-1} \in \\ \Psi_{\mathcal{B}}^{j,n-1}(\phi_k^{j,n-1})}} p_{i,n}(\psi^{j,n-1}) &= \alpha_k \qquad \forall\ \phi_k^{j,n-1} \in \Phi_{\mathcal{B}}^{j,n-1} \\ \sum_{\psi^{j,n-1} \in \Psi_{\mathcal{B}}} p_{i,n}(\psi^{j,n-1}) &= \ 1 \\ p_{i,n}(\psi^{j,n-1}) &\geq \ 0 \qquad \forall\ \psi^{j,n-1} \in \Psi_{\mathcal{B}}^{j,n-1} \end{aligned} \tag{24}$$

Above, we use the abbreviated notation $p_{i,n}(\psi^{j,n-1})$ instead of $p_{i,n}(\Delta_{\mathcal{B}}(\psi^{j,n-1}))$.

Example (Grid world, cont'd) We borrow the structure of the grid world example seen for the L1-BB. Assume now that agent i models the target j as a rational agent, who is in turn maintaining a belief over i's beliefs. The L2-BB for agent i we consider is:

$$B^{i,2} = \begin{array}{|c|}\hline \langle P_j(P_i(\phi_0) \geq 0.4) < 0.4,\, 0.2\rangle \\ \hline \langle P_j(P_i(\phi_1) > 0.5) < 0.7,\, 0.6\rangle \\ \hline \end{array} \tag{25}$$

We consider the set:

$$\Phi_{\mathcal{B}}^{j,1} = \{P_j(\phi_0^{i,0}) < 0.4, P_j(\phi_1^{i,0}) < 0.7\} = \{\phi_0^{j,1}, \phi_1^{j,1}\} \ , \tag{26}$$

where $\phi_0^{i,0}$ and $\phi_1^{i,0}$ are defined in (18). This set induces four non-empty partitions on j's level 1 belief simplex, as shown in Fig. 2-c. We now compute the max-ent distribution $p_{i,2}$ over the four partitions:

$$\max_{p_{i,2}} \left(- \sum_{k=0}^{3} p_{i,2}(\psi_k^{j,1}) \log p_{i,2}(\psi_k^{j,1}) \right), \quad \text{s.t.} \tag{27}$$

$$\begin{bmatrix} 0\,0\,1\,1 \\ 1\,0\,0\,1 \\ 1\,1\,1\,1 \end{bmatrix} \begin{bmatrix} p_{i,2}(\psi_0^{j,1}) \\ p_{i,2}(\psi_1^{j,1}) \\ p_{i,2}(\psi_2^{j,1}) \\ p_{i,2}(\psi_3^{j,1}) \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.6 \\ 1 \end{bmatrix} \tag{28}$$

The resulting probability distribution is: $p_{i,2} = (0.48, 0.32, 0.08, 0.12)$.

4 Conclusion and Future Work

In this paper, we have contribute a novel theoretical framework to compactly represent a hierarchy of interactive beliefs exploiting first-order logic and probability theory. The main idea is to partition the belief simplex at each level of the hierarchy, and let the simplex at level $n-1$ constitute the vertices of the simplex at level n. We have shown that, by recursively partitioning the belief simplices, the representation of the interactive space is finite, thus overcoming the unboundedness of the space of distributions that is typical of standard, enumeration-based representations.

There are several directions for future research. First, we will develop a feasible implementation of our proposed theoretical system and will evaluate the computational costs, both of exact and approximate inference techniques. In particular, we intend to study the ties between our approach and existing first-order probabilistic systems, such as Relational Probabilistic Models and Markov Logic Networks, and possibly extend them towards the interactive beliefs semantics presented in this paper.

Second, we want to embed this novel representation of interactive beliefs in decision making algorithms. One possible application is to extend the work of Sanner and Kersting [23] on First-Order POMDPs to interactive settings. In particular, we believe that our partition-based interactive belief system is suitable to be embedded in decision making frameworks such as Interactive POMDPs. In fact, the optimal value function for (I-)POMDPs divides the belief simplex in partitions corresponding to the optimal policy for each such region. Hence, we want to explore the use of interactive first-order belief bases to recursively represent the relevant belief partitions of the other agents.

References

1. Aumann, R.J.: Interactive epistemology II: Probability. International Journal of Game Theory 28(3), 301–314 (1999)
2. Bacchus, F.: Representing and reasoning with probabilistic knowledge: a logical approach to probabilities. MIT Press, Cambridge, MA, USA (1990)
3. Boutilier, C., Poole, D.: Computing optimal policies for partially observable decision processes using compact representations. In: Proceedings of the National Conference on Artificial Intelligence. pp. 1168–1175 (1996)
4. Boutilier, C., Reiter, R., Price, B.: Symbolic Dynamic Programming for First-Order MDPs. In: IJCAI 01. pp. 690–700. Morgan Kaufmann (2001)

5. Doshi, P., Gmytrasiewicz, P.J.: Monte Carlo sampling methods for approximating interactive POMDPs. J. Artif. Int. Res. 34, 297–337 (March 2009)
6. Doshi, P., Perez, D.: Generalized point based value iteration for interactive POMDPs. In: Proceedings of the 23rd national conference on Artificial Intelligence - Volume 1. pp. 63–68. AAAI Press (2008)
7. Fagin, R., Halpern, J.Y.: Reasoning about knowledge and probability. J. ACM 41, 340–367 (March 1994)
8. Gmytrasiewicz, P.J., Doshi, P.: A framework for sequential planning in multi-agent settings. Journal of Artificial Intelligence Research 24, 24–49 (2005)
9. Halpern, J.Y.: An analysis of first-order logics of probability. In: Proceedings of the 11th international joint conference on Artificial intelligence - Volume 2. pp. 1375–1381. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1989)
10. Halpern, J.Y.: Reasoning about Uncertainty. The MIT Press (2003)
11. Halpern, J.Y., Moses, Y.: Knowledge and common knowledge in a distributed environment. J. ACM 37, 549–587 (July 1990)
12. Harsanyi, J.C.: Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic Model. Management Science 14(3), 159–182 (1967)
13. Jaynes, E.T.: Information Theory and Statistical Mechanics. Physical Review Online Archive (Prola) 106(4), 620–630 (May 1957)
14. Kern-Isberner, G., Thimm, M.: Novel semantical approaches to relational probabilistic conditionals. In: Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning (2010)
15. Koller, D., Pfeffer, A.: Probabilistic frame-based systems. In: Proceedings of the 15th National Conference on Artificial Intelligence (AAAI). pp. 580–587 (1998)
16. Lukasiewicz, T., Kern-Isberner, G.: Probabilistic logic programming under maximum entropy. In: In Proc. ECSQARU-99, LNCS 1638. pp. 279–292. Springer (1999)
17. Milch, B., Marthi, B., Russell, S.J., Sontag, D., Ong, D.L., Kolobov, A.: Blog: Probabilistic models with unknown objects. In: Probabilistic, Logical and Relational Learning (2005)
18. Nilsson, N.J.: Probabilistic logic. Artif. Intell. 28, 71–88 (February 1986)
19. Poupart, P.: Exploiting structure to efficiently solve large scale Partially Observable Markov Decision Processes. Ph.D. thesis, University of Toronto, Toronto, Ont., Canada (2005), aAINR02727
20. Richardson, M., Domingos, P.: Markov logic networks. Mach. Learn. 62, 107–136 (February 2006)
21. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. Pearson Education, 3rd edn. (2010)
22. Sanner, S., Boutilier, C.: Practical solution techniques for first-order MDPs. Artif. Intell. 173, 748–788 (April 2009)
23. Sanner, S., Kersting, K.: Symbolic dynamic programming for first-order POMDPs. In: AAAI 10. pp. 1140–1146 (2010)
24. Shirazi, A., Amir, E.: Factored models for probabilistic modal logic. In: Proceedings of the 23rd national conference on Artificial intelligence - Volume 1. pp. 541–547. AAAI Press (2008)
25. Wang, C., Joshi, S., Khardon, R.: First order decision diagrams for relational MDPs. J. Artif. Int. Res. 31, 431–472 (March 2008)
26. Wang, C., Khardon, R.: Relational partially observable MDPs. In: AAAI 10. pp. 1153–1157 (2010)