

Probabilistic Role Models and the Guarded Fragment

Manfred Jaeger

Institut for Datalogi, Aalborg Universitet, Fredrik Bajers Vej 7E, 9220 Aalborg Ø
jaeger@cs.auc.dk

Abstract

We propose a uniform semantic framework for interpreting probabilistic concept subsumption and probabilistic role quantification through statistical sampling distributions. This general semantic principle serves as the foundation for the development of a probabilistic version of the guarded fragment of first-order logic. A characterization of equivalence in that logic in terms of bisimulations is given.

1 Introduction

Several probabilistic description logics have been proposed over the last ten years [7, 8, 9, 5]. They all focus on replacing the strict subsumption relation $C \sqsubseteq D$ between concepts with a probabilistic subsumption relation

$$P(D | C) = p \quad (1)$$

(“the probability that an object belonging to class C also belongs to class D is p ”). For the specification of concept descriptions C, D a classical, non-probabilistic description logic is used. A language allowing such probabilistic subsumption expressions is then interpreted over standard interpretations I for the underlying description logic, which are augmented by a probability distribution P^I over the domain. This induces probabilities $P^I(C^I)$ for the interpretations of concept descriptions, and hence (1) will be true in the interpretation iff $P^I((C \sqcap D)^I) / P^I(C^I) = p$.

Besides probabilistic subsumption there are several other ways of integrating quantitative, specifically probabilistic, information into a description logic framework. Notably, probabilistic conditions may also be used as concept forming operators. A natural and useful such concept forming operator is obtained by generalizing role quantification $\forall r : C$, and $\exists r : C$ to a probabilistic role quantification

$$P^{=p} r : C. \quad (2)$$

The intended interpretation of the concept description (2) is that in a given interpretation I with domain O it comprises all objects $o \in O$ with the following property: if we randomly select another $o' \in O$ from all objects that are related via r to o , then $o' \in C$ with probability p .

In this paper we will first investigate the semantic foundations of probabilistic description logics that allow expressions of the form (2). We will argue that the adequate way of interpreting such expressions is by interpreting roles as probabilistic transition relations. We will then use this semantic concept to define a probabilistic version of the *guarded fragment* of first-order predicate logic [1]. The resulting *probabilistic guarded fragment* provides a general framework in which many probabilistic description logics can be embedded. In particular, both probabilistic subsumption and probabilistic role quantification are treated in a syntactically and semantically uniform manner.

2 Probabilistic roles

In this section we take a closer look at expressions of the form (2). To make things more graphic, we

use a well-worn example. Consider the expression

$$P^{\geq 0.6} \text{eats} : \text{Fruit} \quad (3)$$

Intuitively, this expression denotes all objects in the domain whose diet consists of at least 60% fruit. Observe that probabilistic concepts of this form are not a replacement for probabilistic subsumption expressions (1), but a means of forming probabilistic concepts. The two can be combined to form statements like

$$P(\text{Herbivore} \mid \text{Animal} \sqcap P^{\geq 0.6} \text{eats} : \text{Fruit}) \geq 0.9 \quad (4)$$

(animals that eat at least 60% fruit are herbivores with probability 0.9).

As a first approach to give semantics to (3), one might just as for probabilistic subsumption statements use a probability distribution P^I on the domain of an interpretation I . This gives us a possible interpretation for the concept description (3) as denoting the set of all objects o for which

$$P^I(\text{Fruit} \mid \text{eats}^I(o, o')) \geq 0.6.$$

To see why this may be inadequate, consider the three statements

$$\text{Herbivore} \sqsubseteq P^{< 0.4} \text{eats_wild}(\text{Fruit}) \quad (5)$$

$$\text{Herbivore} \sqsubseteq P^{> 0.5} \text{eats_zoo}(\text{Fruit}) \quad (6)$$

$$\text{eats_wild} = \text{eats_zoo} \quad (7)$$

The roles `eats_wild` and `eats_zoo` represent what types of food (*banana, kiwi, leaves, insects,...*) are consumed by different types of animals in the wilderness, respectively when kept in a zoo. The probabilistic quantification of these rules specifies the quantitative composition of the diet. E.g. according to (5) fewer than 40% of the meals of a wild antelope will consist of food types in the class `Fruit`. When kept in a zoo, however, this percentage will be greater than 50% (6). In spite of the different quantitative composition, it may still be the case that all food types consumed by an animal in the wilderness will at some time or another also be fed at the zoo (7).

If we try to interpret (5)-(7) by a domain distribution P^I , we find that because of (7) for all animals a : $P^I(\{f \mid \text{eats_wild}a\}) = P^I(\{f \mid$

`eats_zoo}a\})`, so that the concepts on the right-hand sides of (5) and (6) are mutually exclusive, and hence the `Herbivore` concept is inconsistent.

The watchful reader will have noticed that the preceding argument depends on our intuitive reading the meaning of (5)-(7), and that somewhat different semantics were attributed to the roles `eats_...` as they appear in (5),(6) on the one hand, and in (7) on the other hand: in the first two expressions, the roles were associated with quantitative information on what animals will typically eat; in the last expression they were taken to provide only the qualitative information of what animals will possibly eat at all. Only under such a dualistic interpretation of the roles can we argue that (5)-(7) should not imply that the concept `Herbivore` is empty. Otherwise we should argue that (7) enables us to substitute e.g. `eats_wild` for `eats_zoo` in (6), which would then necessarily make the right-hand sides of (5) and (6) inconsistent.

It is the goal of the present paper to develop a framework where roles can serve the dual purpose of expressing both what is accessible in principle, and with what frequency different role fillers will be chosen. We may then want to introduce two different roles that coincide with regard to their first, qualitative, interpretation, but not with regard to their second, quantitative, interpretation. As we have seen, this cannot be realized with a semantics based on a single domain probability distribution P^I . Similarly, one can see that intended behavior of the logic cannot be obtained by basing the semantics on a probability distribution over a set of possible interpretations – the second semantic paradigm usually considered for the interpretation of probabilistic logics [6].

Interpreting roles as accessibility relations in the sense of modal logic, it is natural to consider semantics where a qualitative accessibility relation is replaced by probabilistic transition probabilities between states (or possible worlds). This leads to the notion of a probabilistic transition system, as widely studied in the context of distributed and concurrent systems [13, 10]. Such a system consists of a set of states, a set of primitive propositions interpreted at each state, and a set of probabilistic transition relations, each of which

determines for each state a probability distribution over the states of the system. In concurrent systems theory, each transition relation is associated with a possible action the system can perform; they determine the probabilities over possible successor states when that action is taken at a given current state.

In the context of probabilistic description logics, the usual temporal interpretation of probabilistic transitions between states can be replaced with an interpretation as a probabilistic choice of role-fillers among objects. The mathematical structure of the probabilistic model is the same. However, transition probabilities between states are not quite strong enough a semantic concept to support some modeling tasks one might want to solve in a probabilistic description logic. For example, transition probabilities do not determine a probability distribution over possible predecessor states, given that we have just reached a certain state by executing a certain action. Such probabilities, however, are required when we also want to reason e.g. about inverses of probabilistic roles. By a small modification of the model such reasoning can also be supported: instead of assigning to each object and role a conditional distributions over role-fillers, one can define for each role a global distribution over object pairs. For any object o , this global distribution can be conditioned either on o being the first component of an object pair, yielding the distribution of role-fillers for o , or on o being the second argument of an object pair, yielding the distribution of objects that have o as a role-filler.

The interpretation of a role as a probability distribution over pairs of objects induces a coarser, qualitative interpretation as the set of all pairs with nonzero probability. Using these two aspects of a probabilistic role for the interpretation of (5)-(7), we obtain the desired consistency of the `Herbivore` concept.

Interpreting roles by probability distributions over object pairs does not provide an interpretation for a probabilistic concept subsumption (1). Consider, in particular, the case that the conditioning concept description C is a primitive concept A . Then (1) says that an object randomly selected from A will belong to D with probability p . Just as we see roles as random selectors

of pairs of objects, we can view primitive concepts as random selectors of single objects, and interpret each primitive concept by a probability distribution on the domain. Then “randomly selecting from A ” means: sampling according to the domain distribution associated with A , and a formal semantics for $P(C \mid A)$ is obtained. Again we can have the situation that two primitive concepts are identical in their qualitative appearance as to what domain objects can possibly be sampled, but differ with respect to the probabilities. For instance the two concepts `Wild_Mammal` and `Zoo_Mammal` might very well include exactly the same species (assuming, perhaps not realistically, that every species is kept in at least one zoo somewhere). However, the probability of being a giraffe is much higher for a randomly observed `Zoo_Mammal` than for a randomly observed `Wild_Mammal` (which will tend to be a rabbit or squirrel).

We have thus outlined our approach for a very general probabilistic description logic: interpret each role by a probability distribution on pairs of domain objects, and each primitive concept by a probability distribution on the domain. In the following sections we develop this semantic approach. However, rather than introducing a specific probabilistic description logic, we pursue this program in the much more general context of the guarded fragment of first-order logic.

3 The Probabilistic Guarded Fragment

The guarded fragment was introduced by Andr eka et al. [1] as a natural fragment of first-order predicate logic that, on the one hand, is general enough to allow most modal logics to be interpreted in that fragment, and on the other hand preserves the nice decidability properties of modal logics. The guarded fragment is characterized by the syntactic condition that all quantifiers must be relativized (“guarded”) by an atomic expression that contains all the variables occurring free in the formula quantified over. Furthermore, a relational vocabulary is assumed. In the following we develop the probabilistic guarded fragment by replacing guarded existential and universal quantification with probabilistic quantification relative to

relation-induced probability distributions.

We begin by introducing some notational conventions. We assume a relational vocabulary $S = \{r_1, \dots, r_m\}$ where each r_i is a relation symbol of arity ≥ 1 (unlike the standard guarded fragment, no constant symbols will be allowed). Throughout, $x_1, x_2, \dots, y_1, y_2, \dots$ denote variables and $a_1, a_2, \dots, b_1, b_2, \dots$ denote domain objects. Tuples of variables or objects are written in bold font: $\mathbf{x}, \mathbf{a}, \dots$. The length of a tuple \mathbf{x} is denoted $|\mathbf{x}|$. When $r \in S$, then $r(\mathbf{x}, \mathbf{a})$ denotes an atomic expression consisting of the relation r with arguments from \mathbf{x} and \mathbf{a} . The notation neither implies any particular order of the arguments, nor that all components of \mathbf{x} and \mathbf{a} must actually be used: if e.g. $\mathbf{x} = (x_1, x_2)$, $\mathbf{a} = (a_1)$, then $r(\mathbf{x}, \mathbf{a})$ can be any of the atoms $rx_1x_2a_1, rx_2a_1x_1, rx_1a_1x_1$, etc. Similarly, $\phi(\mathbf{x})$ denotes a formula all of whose free variables are in \mathbf{x} .

Definition 3.1 Let S be a relational vocabulary. The language $\mathcal{L}_{pgf}(S)$ of the probabilistic guarded fragment is defined as follows.

- Relational atoms $r(\mathbf{x})$ ($r \in S$) and equality atoms $x_i = x_j$ are in $\mathcal{L}_{pgf}(S)$.
- $\mathcal{L}_{pgf}(S)$ is closed under boolean operators.
- If $\phi(\mathbf{x}, \mathbf{y}), \psi(\mathbf{x}, \mathbf{y}) \in \mathcal{L}_{pgf}(S)$, $\alpha(\mathbf{x}, \mathbf{y})$ is a positive S -atom such that all free variables of $\phi(\mathbf{x}, \mathbf{y})$ and $\psi(\mathbf{x}, \mathbf{y})$ also occur in $\alpha(\mathbf{x}, \mathbf{y})$, and p a rational number in $[0, 1]$, then

$$P^{\sim p} \mathbf{y} . \alpha(\mathbf{x}, \mathbf{y}) (\phi(\mathbf{x}, \mathbf{y}) \mid \psi(\mathbf{x}, \mathbf{y})) \quad (8)$$

is in $\mathcal{L}_{pgf}(S)$ for $\sim \in \{<, >, =, \geq, \leq\}$.

We here make expressions for conditional probabilities the basic syntactic construct of our language. This is because reasoning with conditional probabilities is key to all interesting probabilistic inference tasks, and the absence of algebraic operators in our language means that we cannot build conditional probability expressions out of unconditional ones. However, we may freely use unconditional expressions $P^{\dots} \dots (\phi(\mathbf{x}, \mathbf{y}))$, which can be seen as abbreviations for conditional expressions with tautological ψ .

Example 3.2 Let `Herbivore`, `Fruit` be unary relation symbols and `eats_wild` binary. Then

$$\chi(x) := P^{<0.4} y . \text{eats_wild}xy(\text{Fruity})$$

is a \mathcal{L}_{pgf} -formula in the free variable x that corresponds to the concept description on the right-hand side of (5), and

$$P^{=1} x . \text{Herbivore}x(\chi(x))$$

is the \mathcal{L}_{pgf} -sentence encoding the terminological axiom (5). The sentence

$$P^{\geq 0.7} x . \text{Herbivore}x(\chi(x))$$

(at least 70% of herbivores eat less than 40% fruit in the wild) combines probabilistic subsumption and probabilistic quantification in the unified form provided by the probabilistic guarded fragment.

Following the discussion in section 2, the semantics for \mathcal{L}_{pgf} is straightforward – with one exception: we restrict the semantics to models with at most countable domains. This is because two key concepts we need – conditional distributions and the support of a probability measure – require some non-trivial topological measure theory when applied to uncountable domains.

Definition 3.3 A probabilistic S -structure \mathcal{M} has the form $(M, (P_{\mathbf{r}}^M)_{\mathbf{r} \in S})$, where

- M is a finite or countably infinite domain.
- For $\mathbf{r} \in S$ of arity k , $P_{\mathbf{r}}^M$ is a probability distribution on M^k .

To define the satisfaction relation between models and \mathcal{L}_{pgf} -formulas, we require some additional notation. To simplify matters, the following definitions are given by examples for a ternary relation symbol r .

Let $P_{\mathbf{r}}^M$ be an interpretation of r . Let $a, a' \in M$. Then $P_{\mathbf{r}(a, \cdot, \cdot)}^M$ denotes the distribution on M^2 obtained by conditioning $P_{\mathbf{r}}^M$ on the set $\{(a_1, a_2, a_3) \mid a_1 = a\}$. (Strictly speaking, this conditional is still a distribution on M^3 . However, the first component being fixed, it effectively becomes a distribution on M^2). Similarly, $P_{\mathbf{r}(\cdot, a, a')}^M$ is the conditional of $P_{\mathbf{r}}^M$ on the first component,

given that the second is a and the third is a' . We use $P_{\mathbf{r}(a,\cdot)}^M$ as a generic expression that can stand for fixing arbitrary components of \mathbf{r} to values from \mathbf{a} . If $\mathbf{r}(\mathbf{x}, \mathbf{y}) = r_{x_1}y_1x_2$ and $\mathbf{a} = (a, a')$, then $P_{\mathbf{r}(\mathbf{x}, \mathbf{y})[\mathbf{x}/\mathbf{a}]}^M$ stands for $P_{\mathbf{r}(a,\cdot,a')}^M$.

Definition 3.4 The satisfaction relation between \mathcal{L}_{pgf} -formulas and probabilistic structures under a variable assignment β is defined by

- $\mathcal{M}, \beta \models \mathbf{r}\mathbf{x}$ iff $P_{\mathbf{r}}(\beta(\mathbf{x})) > 0$.

Assuming that the satisfaction relation has been defined for $\phi(\mathbf{x}, \mathbf{y})$, we abbreviate the probability

$$P_{\mathbf{r}(\mathbf{x}, \mathbf{y})[\mathbf{x}/\mathbf{a}]}^M(\{\mathbf{b} \mid \mathcal{M}, [\mathbf{y}/\mathbf{b}] \models \phi(\mathbf{x}, \mathbf{y})\})$$

with

$$P_{\mathbf{r}(\mathbf{x}, \mathbf{y})[\mathbf{x}/\mathbf{a}]}^M(\phi(\mathbf{x}, \mathbf{y})).$$

The semantics for probabilistic quantification then is

- $\mathcal{M}, \beta \models P^{\sim p}\mathbf{y}.\mathbf{r}(\mathbf{x}, \mathbf{y})(\phi(\mathbf{x}, \mathbf{y}) \mid \psi(\mathbf{x}, \mathbf{y}))$
iff $P_{\mathbf{r}(\mathbf{x}, \mathbf{y})[\mathbf{x}/\beta(\mathbf{x})]}^M(\phi(\mathbf{x}, \mathbf{y}) \wedge \psi(\mathbf{x}, \mathbf{y})) \sim p \cdot$
 $P_{\mathbf{r}(\mathbf{x}, \mathbf{y})[\mathbf{x}/\beta(\mathbf{x})]}^M(\psi(\mathbf{x}, \mathbf{y})).$

We can identify formulas in the standard guarded fragments with probabilistic guarded formulas of the forms $P^{=1}\mathbf{y}.\alpha(\mathbf{x}, \mathbf{y})(\phi(\mathbf{x}, \mathbf{y}))$ (universal quantification) and $P^{>0}\mathbf{y}.\alpha(\mathbf{x}, \mathbf{y})(\phi(\mathbf{x}, \mathbf{y}))$ (existential quantification). One can then show that a formula in the standard guarded fragment (without constants) is satisfiable iff its corresponding formula in the probabilistic guarded fragment is satisfiable (for the direction from left to right here the finite-model property of the standard guarded fragment is needed to obtain that our restriction to at most countable structures does not impose any extra limitations on satisfiability). Thus, the probabilistic guarded fragment is a proper extension of the constant-free standard guarded fragment.

4 Inference

Having defined syntax and semantics of \mathcal{L}_{pgf} , we must address the question of inference in this logic. It is beyond the scope of the present paper to prove either the completeness of some proof system for this logic, or to answer the key question whether \mathcal{L}_{pgf} is decidable. Instead we will

here just propose a Gentzen-style sequent calculus for the logic, and illustrate some inferences it supports. The material in this section must be regarded as preliminary.

We adopt the propositional inference rules and the rules relating to equality from a standard sequent calculus for first-order logic (see e.g. [3]). To simplify matters we here limit ourselves to inference with unconditional probability expressions. A calculus for conditional expressions can be developed along very much the same lines, but will require some additional rules, resp. additional preconditions on some rules. Table 1 shows inference rules relating to probabilistic quantification. The notation P^{\sim} is used for inference rules that refer equally to $P^{>}$ and $P^{=}$. Rules (P1)-(P5) characterize the basic probabilistic properties of the P -quantifier. Rules (I1)-(I3) characterize the interplay of the qualitative and the quantitative aspects of relation symbols. Rules (Q1)-(Q3) are for quantifier introduction.

Example 4.1 Consider the sentence

$$P^{=1}x.\text{Omnivore}x \\ (P^{\leq 0.5}y.\text{feeding}xy(\text{likes}xy)). \quad (9)$$

Here the role $\text{feeding}xy$ represents the relation that animal x is being fed food y at a (public) feeding at the zoo. The probabilities attached to this role may represent the probability that a random visitor to the zoo watches this feeding (i.e. the distribution P_{feeding} is essentially given by the average number of spectators for this feeding). Then (9) says that when we pick some omnivore and watch one of its feedings, then the probability that the animal actually likes the food it is given is less than 0.5. Using the inference rules given above, one can derive the sequent

$$\text{Omnivore}x \rightarrow P^{\leq 0.5}y.\text{feeding}xy(\text{likes}xy) \\ \vdash P^{\leq 0.5}y.\text{feeding}xy(\text{likes}xy \wedge \text{Omnivore}x)$$

(this uses essentially (Q1) and propositional reasoning). The antecedent of this sequent can be derived from (9) using (I2). Since x is not free in (9), we can finally derive using (Q3):

$$P^{\leq 0.5}xy.\text{feeding}xy(\text{likes}xy \wedge \text{Omnivore}x). \quad (10)$$

$\frac{\Gamma \vdash \phi \leftrightarrow \psi \quad \Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi)}{\Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\psi)} \quad (\text{P1})$		
$\frac{\Gamma \vdash P^{\sim 1} \mathbf{y}, \alpha(\mathbf{y}=\mathbf{y})}{\Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi \wedge \alpha)} \quad (\text{P2})$	$\frac{\Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi)}{\Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi \wedge \alpha)} \quad (\text{I1})$	$\frac{\Gamma, \phi \vdash \psi}{\Gamma, P^{> p} \mathbf{x}, \alpha(\phi) \vdash P^{> p} \mathbf{x}, \alpha(\psi)} \quad (\text{Q1})$
$\frac{\Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi)}{\Gamma \vdash P^{> q} \mathbf{y}, \alpha(\phi)} \quad (q < p) \quad (\text{P3})$	$\frac{\Gamma \vdash P^{\sim 1} \mathbf{y}, \alpha(\phi)}{\Gamma \vdash \alpha \rightarrow \phi} \quad (\text{I2})$	$\frac{\Gamma, \phi \vdash \psi \quad \Gamma, \psi \vdash \phi}{\Gamma, P^{\sim p} \mathbf{x}, \alpha(\phi) \vdash P^{\sim p} \mathbf{x}, \alpha(\psi)} \quad (\text{Q2})$
$\frac{\Gamma \vdash \phi \rightarrow \psi \quad \Gamma \vdash P^{\sim q} \mathbf{y}, \alpha(\phi)}{\Gamma \vdash P^{\geq q} \mathbf{y}, \alpha(\psi)} \quad (\text{P4})$	$\frac{\Gamma \vdash \alpha \rightarrow \phi}{\Gamma \vdash P^{\sim 1} \mathbf{y}, \alpha(\phi)} \quad (\text{I3})$	$\frac{\Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi)}{\Gamma \vdash P^{\sim p} \mathbf{x} \mathbf{y}, \alpha(\phi)} \quad (\text{Q3})$
$\frac{\Gamma \vdash \neg(\psi \wedge \phi) \quad \Gamma \vdash P^{\sim p} \mathbf{y}, \alpha(\phi) \quad \Gamma \vdash P^{\sim q} \mathbf{y}, \alpha(\psi)}{\Gamma \vdash P^{\sim p+q} \mathbf{y}, \alpha(\phi \vee \psi)} \quad (\text{P5})$		$\mathbf{x} \text{ not free in } \Gamma$

Table 1: Inference rules

This is a quite different statement than (9): it says that when we randomly pick some feeding event in the zoo, then the probability that we will watch an omnivore eating food it likes is less than 0.5. Unlike (9) this statement can also be true simply due to the fact that we are unlikely to select an omnivore feeding to watch. Note, too, that if the first quantifier in (9) had stated just a lower bound $P^{>0.9} \dots$, then no non-trivial upper bound $P^{\leq q} \dots$ could have been derived for the sentence (10), as then it would be possible that the distribution P_{feeding} is in its first argument highly biased towards the few omnivores that actually like most of their food.

5 Bisimulations

In this section we provide a semantic characterization of logical equivalence in \mathcal{L}_{pgf} . For predicate logics, characterizations of this kind are typically given in a game-theoretic or algebraic setting (the Ehrenfeucht-Fraïssé method, see e.g. [4]). In the context of modal logics, closely related methods are called *bisimulations* [10, 1]. Here we use the terminology of bisimulations, but the techniques we use are in fact very close to the Fraïssé style algebraic method.

Definition 5.1 Let \mathcal{M}, \mathcal{N} be two probabilistic S -structures. A *partial isomorphism* is a bijection $f : \mathbf{a} \mapsto \mathbf{b}$ between a finite subset $\mathbf{a} \subseteq M$ and a finite subset $\mathbf{b} \subseteq N$, such that for all $\mathbf{r} \in S$ and all $\tilde{\mathbf{a}} \subseteq \mathbf{a}$: $\mathcal{M} \models \mathbf{r}\tilde{\mathbf{a}}$ iff $\mathcal{N} \models \mathbf{r}f(\tilde{\mathbf{a}})$.

The following is the crucial concept needed to extend the standard techniques for predicate logics to the probabilistic setting.

Definition 5.2 Let F be a set of partial isomorphisms between \mathcal{M} and \mathcal{N} . Let \sim_F be the

symmetric and transitive closure of the relation $\{(\mathbf{a}, \mathbf{b}) \mid \exists f \in F : f(\mathbf{a}) = \mathbf{b}\}$. For $\mathbf{a} \in M^k$ let $[\mathbf{a}]^F$ be the equivalence class in M^k of \mathbf{a} with respect to \sim_F . Thus, $\mathbf{a}' \in [\mathbf{a}]^F$ iff there exists a finite sequence $\mathbf{a} \xrightarrow{f_1} \mathbf{b}_1 \xleftarrow{f_2} \mathbf{a}_2 \xrightarrow{f_3} \dots \mathbf{b}_m \xleftarrow{f_{m+1}} \mathbf{a}'$ with $f_1, \dots, f_{m+1} \in F$. Finally, define for \mathbf{a}, \mathbf{a}' :

$$[\mathbf{a}']_{\mathbf{a}}^F := \{\mathbf{a}'' \mid \mathbf{a}\mathbf{a}'' \in [\mathbf{a}\mathbf{a}']^F\} \quad (11)$$

(the section of the equivalence class $[\mathbf{a}\mathbf{a}']$ along \mathbf{a}).

For the following definition, $\text{dom}(f)$ denotes the domain of the partial isomorphism f . A tuple \mathbf{a} , resp. \mathbf{b} is called *guarded*, if $\mathcal{M} \models \mathbf{r}\mathbf{a}$, resp. $\mathcal{N} \models \mathbf{r}\mathbf{b}$, for some $\mathbf{r} \in S$.

Definition 5.3 A k -bisimulation between \mathcal{M} and \mathcal{N} is a sequence F_0, F_1, \dots, F_k of non-empty sets of partial isomorphisms, such that

- (*forth*) for all $f : \mathbf{a} \mapsto \mathbf{b} \in F_i$ and for all guarded \mathbf{a}' there exists $g \in F_{i-1}$ with $\mathbf{a}' \in \text{dom}(g)$, such that for $\mathbf{a}^* := \mathbf{a}' \cap \mathbf{a}$, $\mathbf{a}^{**} := \mathbf{a}' \setminus \mathbf{a}$, and $\mathbf{b}^* := f(\mathbf{a}^*)$:

$$g(\mathbf{a}^*) = \mathbf{b}^* \quad (12)$$

and for all $\mathbf{r} \in S$

$$P_{\mathbf{r}(\mathbf{a}^*, \cdot)}^M([\mathbf{a}^{**}]_{\mathbf{a}^*}^{F_{i-1}}) = P_{\mathbf{r}(\mathbf{b}^*, \cdot)}^N([\mathbf{b}^{**}]_{\mathbf{b}^*}^{F_{i-1}}) \quad (13)$$

- (*back*) for all $f : \mathbf{a} \mapsto \mathbf{b} \in F_i$ and for all guarded \mathbf{b}' there exists $g \in F_{i-1}$ with $\mathbf{b}' \in \text{dom}(g^{-1})$, such that for $\mathbf{b}^* := \mathbf{b}' \cap \mathbf{b}$, $\mathbf{b}^{**} := \mathbf{b}' \setminus \mathbf{b}$, and $\mathbf{a}^* := f^{-1}(\mathbf{b}^*)$ (12) and (13) hold.

A bisimulation is an infinite sequence F_0, F_1, \dots with the *back* and *forth* property.

Following a common strategy, we will derive a stratified characterization of logical equivalence

in \mathcal{L}_{pgf} by classifying formulas according to their *nesting depth* [12]. Nesting depth is a variation of quantifier rank: the application of any probability quantifier $P^{\sim q} \mathbf{y}.\alpha$ increments the nesting depth of a formula by one, irrespective of the number of variables in \mathbf{y} . We denote with \mathcal{L}_{pgf}^k the set of formulas with nesting depth at most k .

For $\mathbf{a} \subseteq M, \mathbf{b} \subseteq N$ with $|\mathbf{a}| = |\mathbf{b}|$ we write $\mathcal{M}[\mathbf{a}] \equiv_{\mathcal{L}_{pgf}^k} \mathcal{N}[\mathbf{b}]$ iff for all $\phi(\mathbf{x}) \in \mathcal{L}_{pgf}^k$: $\mathcal{M} \models \phi(\mathbf{a})$ iff $\mathcal{N} \models \phi(\mathbf{b})$.

Theorem 5.4 The following are equivalent for finite \mathcal{M}, \mathcal{N} :

- (i) $\mathcal{M}[\mathbf{a}] \equiv_{\mathcal{L}_{pgf}^k} \mathcal{N}[\mathbf{b}]$
- (ii) There exists a k -bisimulation $\mathbf{F} : \mathcal{M} \rightarrow \mathcal{N}$ with $f(\mathbf{a}) = \mathbf{b}$ for some $f \in F_k$.

If \mathcal{M} or \mathcal{N} is infinite, then the implication (ii) \rightarrow (i) holds.

Due to space limitations the proof is omitted here. For the direction (i) \rightarrow (ii) it heavily relies on the finiteness of \mathcal{M}, \mathcal{N} . However, it is an open problem whether this implication might actually also hold for infinite structures.




Figure 1: Bisimilar unary structures

Example 5.5 Figure 1 shows two finite structures for a vocabulary consisting of unary relations r and s . The distributions P_r, P_s are specified in the format $P_r(a); P_s(a)$ at each domain element a . The structures differ in the size of their (qualitative) interpretations of r, s . However, both structures agree on the probabilities with which an object sampled from one relation will belong to the other relation as well (for instance, with probability 0.8 is an object sampled from r also in s). This means that the two structures are bisimilar: let F be the set of all partial isomorphisms

f with $dom(f)$ of size one. The dashed arrow in Figure 1 indicates one f in F . We show that the constant sequence $F_i := F$ has the *forth* property. For this let $f : a \mapsto b \in F$. Since the vocabulary is unary, and all domain elements in M belong to at least one relation, guarded tuples \mathbf{a}' are just all singletons a' . For $a' = a$ one obtains that $g = f$ is as required for the *forth* property ((12) trivially holds, and (13) is vacuous because $\mathbf{a}^{**} = \emptyset$). If $a' \neq a$ then there exists some b' such that $g : a' \mapsto b' \in F$ (possibly $b' = b$). Now (12) is vacuous, because $\mathbf{a}^* = \emptyset$. Condition (13) becomes

$$P_{r(\cdot)}^M([a']^F) = P_{r(\cdot)}^N([b']^F),$$

(and similarly for s). Since here an equivalence class $[a]^F$ just contains all a' that agree with a with respect to membership in r, s , these conditions are satisfied by \mathcal{M}, \mathcal{N} .

The preceding example can be easily generalized to provide a simple and complete characterization of bisimilar unary structures. It can be interpreted as showing that for unary vocabularies \mathcal{L}_{pgf} essentially becomes a propositional probability logic in the style of Nilsson [11].




Figure 2: Bisimilar structures with binary r

Example 5.6 Figure 2 shows two structures \mathcal{M}, \mathcal{N} for a single binary relation r . The solid arrows in the figure indicate the pairs with nonzero probability P_r . Irrespective of the concrete probability values, the two structures are bisimilar. To see this, let F be the set of all partial isomorphisms f with $dom(f)$ of size two. The dashed arrows indicate one such $f : (a_1, a_2) \mapsto (b_1, b_2)$. Again, let $F_i = F$ for all i . For f and a tuple $\mathbf{a}' = (a'_1, a'_2)$ in the *forth* property there exist three possibilities: either 0, 1 or 2 components of \mathbf{a}' can be equal to one of a_1 or a_2 . We consider the case of one component being equal, e.g. $\mathbf{a}'(a_1, a_4)$. Then $g : (a_1, a_4) \mapsto (b_1, b_2) \in F$, and (12) is satisfied. In \mathcal{M} all pairs of neighboring elements are equivalent with respect to \sim_F .

In particular, $[a_4]_{a_1}^F$ consists of the two neighbors $\{a_2, a_4\}$ of a_1 . In \mathcal{N} $[b_2]_{b_1}^F$ consists of the single element b_2 . In both cases, the probability of reaching these sets from a_1 , resp. b_1 is equal to one, i.e. (13) holds. The other cases for \mathbf{a}' are similar.

In both of our examples, the bisimulation sequence was a constant F . It is easy to see that for finite bisimilar structures the F_i can always be chosen to be the same for all i .

6 Conclusion

In this paper we have introduced a probabilistic version of the guarded fragment, and shown how, in particular, it can be used as a probabilistic extension of description logics that permits probabilistic subsumption and probabilistic role quantification in a uniform framework.

The key design choice in the logic presented here is to view each k -ary relation as a random selector of k -tuples from the domain. It is not always natural to assume that there should be a one-to-one correspondence between different sampling distributions that should be distinguished for a domain, and the relations used to describe qualitative dependencies between objects in the domain. However, the logic \mathcal{L}_{pgf} allows, for instance, to use some relation symbols exclusively as standard logical atoms, and others exclusively as guards in probability quantifiers. In this way one can effectively undo the assumed duality of relation symbols, and reason with arbitrary sampling distributions and arbitrary qualitative relations.

In introducing our semantics in section 2 we have emphasized its basis in the semantics of probabilistic transition systems. However, it is also quite closely related to some other probabilistic frameworks. First, being essentially given through probabilities on (products of) the domain of a structure, our semantics can be viewed as a variant of statistical probability semantics [6, 2] with multiple sampling distributions. In the context of probabilistic description logics, Koller et al. [9] also have proposed a framework where probabilistic selection of role fillers can be represented by different domain distributions. Apart from this similarity, however, the system in [9] is

entirely different from the one presented here, as it is a Bayesian network based representation language for completely specified probability distributions, not a logic-based representation language for (partial) probabilistic knowledge.

Key problem for future work is the decidability question for the probabilistic guarded fragment.

References

- [1] H. Andr eka, I. N emeti, and J. van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 1998.
- [2] F. Bacchus. *Representing and Reasoning With Probabilistic Knowledge*. MIT Press, 1990.
- [3] J. Barwise. An introduction to first-order logic. In J. Barwise, editor, *Handbook of Mathematical Logic*, pages 5–46. North Holland, 1977.
- [4] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer Verlag, second edition, 1999.
- [5] R. Giugno and T. Lukasiewicz. P-SHOQ(D): A probabilistic extension of SHOQ(D) for probabilistic ontologies in the semantic web. In *Proceedings of JELIA'02*, pages 86–97, 2002.
- [6] J.Y. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46:311–350, 1990.
- [7] J. Heinsohn. Probabilistic description logics. In *Proceedings of UAI-94*, pages 311–318, 1994.
- [8] M. Jaeger. Probabilistic reasoning in terminological logics. In *Proceedings of KR94*, pages 305–316, 1994.
- [9] D. Koller, A. Levy, and A. Pfeffer. P-classic: A tractable probabilistic description logic. In *Proceedings of AAAI-97*, pages 390–397, 1997.
- [10] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 1991.
- [11] N. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–88, 1986.
- [12] M. Otto. Modal and guarded characterisation theorems over finite transition systems. To appear in *Annals of Pure and Applied Logic*.
- [13] M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proceedings of the 26th IEEE Symposium on Foundations of Computer Science*, pages 327–338, 1985.