

Probabilistic Description Logics for Subjective Uncertainty

Carsten Lutz

Department of Computer Science
University of Bremen
clu@uni-bremen.de

Lutz Schröder

DFKI Bremen and
Dept. of Computer Science, University of Bremen
lutz.schroeder@dfki.de

Abstract

We propose a new family of probabilistic description logics (DLs) that, in contrast to most existing approaches, are derived in a principled way from Halpern's probabilistic first-order logic. The resulting probabilistic DLs have a two-dimensional semantics similar to certain popular combinations of DLs with temporal logic and are well-suited for capturing subjective probabilities. Our main contribution is a detailed study of the complexity of reasoning in the new family of probabilistic DLs, showing that it ranges from PTIME for weak variants based on the lightweight DL $\mathcal{EL}$ to undecidable for some expressive variants based on the DL $\mathcal{ALC}$.

Motivation

Description logics (DLs) are a popular family of knowledge representation formalisms that underlie ontology languages such as the W3C standard OWL. Since traditional DLs are essentially fragments of first-order logic (FOL), they allow only the representation of crisp and definite knowledge, whereas no built-in means are provided to represent uncertainty of any kind. This shortcoming in expressive power has often been criticized as it impedes an adequate modelling of the relevant knowledge in many application areas.

We illustrate the problem in terms of bio-medical applications, where ontologies are being used with particular success. Almost every bio-medical ontology contains uncertain concepts of some sort, although typically modelled in an inappropriate way. Take for example the well-known and widely used medical ontology SNOMED CT (Price and Spackman 2000), which comprises a variety of concepts such as '*Probable tubo-ovarian abscess*', '*Natural death with probable cause suspected*', '*Probable diagnosis*', '*Probably present*', '*Basal cell tumour, uncertain whether benign or malignant*', etc. Similar concepts can be found in other bio-medical ontologies such as GALEN. Since traditional DLs are used, the aspect of uncertainty in these concepts is not reflected in their modelling. For example, nothing is said in SNOMED CT about '*Natural death with probable cause suspected*', except that it is a subclass of '*Natural death*'.

The desire to represent uncertainty in this and other application domains in a proper way has led to various proposals for *probabilistic DLs*, see (Lukasiewicz and Straccia 2008) for a recent survey. These logics differ considerably

regarding the general way in which probabilities are used, in the syntax, in the chosen semantics, and in the intended application. In the present work, we follow a principled approach to defining probabilistic DLs by viewing them as fragments of probabilistic FOL (Halpern 1990) in the same way as traditional DLs are fragments of traditional FOL. As discussed in more detail below, the resulting family of probabilistic DLs has the modelling of *subjective probabilities* as a clear-cut application. We note that a FOL-based approach to probabilistic DLs has already been advocated by Sebastiani (1994), but was never developed in serious detail. To some extent, it can be seen as complementary to the approach taken by Lukasiewicz (2008), who defines a family of probabilistic DLs by viewing them as extensions of probabilistic propositional logics.

In the approach to probabilistic FOL suggested by Halpern (1990), a distinction is made between *statistical probabilities* as formalized by 'Type 1' probabilistic FOL and *subjective probabilities* as formalized by 'Type 2' probabilistic FOL. The statistical view is concerned with probability distributions on the domain, and the interest is typically in conditional probabilities such as $P(\text{LymeDisease} \mid \text{PositiveSerology}) \geq 0.8$, expressing that at least 80% of all patients with positive serological blood tests actually do have Lyme disease. In this context, assertions of facts on the 'instance level' necessarily remain crisp: a particular patient either has or does not have Lyme disease (irrespective of the fact that our knowledge of this may be uncertain), but will never 'have Lyme disease with a probability $0 < p < 1$ '. Contrastingly, the subjective view regards probabilities as degrees of belief. It is thus concerned with probability distributions on a set of possible worlds, each one associated with a standard FOL interpretation that describes a possible view of the world. In this context, probabilistic facts are unproblematic: the statement that patient X has Lyme disease with probability p holds if p is the probability assigned to the set of all worlds in which patient X has Lyme disease. Indeed, uncertain facts of this kind easily arise in situations where diagnosing techniques are not 100% reliable or medical data is obtained from untrusted data sources.

Much of the existing work on probabilistic DLs either does not carefully distinguish between the two kinds of probabilities or is primarily based on the statistical view, see e.g. (Heinsohn 1994; Jaeger 1994; Koller, Levy, and

Pfeffer 1997; Yelland 2000; Ding, Peng, and Pan 2006; Jaeger 2006). Here, we concentrate on the subjective view, which in particular enables us to capture the medical examples given above. We are thus concerned with probabilistic DLs that are fragments of Halpern's Type 2 probabilistic FOL, from which we inherit a standard and easy-to-comprehend semantics, resembling the two-dimensional semantics of many popular temporal DLs (Lutz, Wolter, and Zakharyashev 2008; Gabbay et al. 2003). In particular, our semantics is not entangled with the syntax, as is the case in probabilistic DLs such as (Lukasiewicz 2008; Jaeger 1994).

Our main contribution is an analysis of the computational complexity of knowledge base consistency and instance checking in various members of the new family of probabilistic DLs. In particular, the analysis aims to separate features of our probabilistic DLs that make reasoning difficult from those that do not. We start with a probabilistic logic Prob- $\mathcal{ALC}_c$ that is based on $\mathcal{ALC}$ and allows the application of probabilities only to concepts, but not to roles. This logic provides sufficient expressive power for most applications; we are able to show that reasoning in it is EXPTIME-complete, thus no more difficult than in non-probabilistic $\mathcal{ALC}$. We then prove that this good computational behaviour is retained even if we add explicit probabilistic independence constraints and general linear inequalities between probabilities of concepts.

Next, we extend Prob- $\mathcal{ALC}_c$ with probabilistic roles and show that the resulting DL Prob- $\mathcal{ALC}$ is computationally more problematic: it is 2-EXPTIME-hard even when probability values are restricted to 0 and 1; when independence constraints or linear inequalities are added, reasoning even becomes undecidable. Finally, we study lightweight probabilistic DLs based on the description logic $\mathcal{EL}$. It turns out that the PTIME complexity of non-probabilistic $\mathcal{EL}$ carries over to Prob- $\mathcal{EL}$ only in the case where (i) probability values are restricted to 0 and 1; and (ii) probabilities are only applied to concepts, but not to roles. If (i) is dropped, reasoning becomes EXPTIME-complete, and if (ii) is dropped, it becomes PSPACE-hard.

Longer proofs are deferred to the appendix, which is provided at <http://www.informatik.uni-bremen.de/~clu/papers/>.

The Prob- $\mathcal{ALC}$ Family of Probabilistic DLs

We introduce Prob- $\mathcal{ALC}$, a probabilistic variant of the basic DL $\mathcal{ALC}$. All other probabilistic DLs studied in this paper are either fragments or extensions of Prob- $\mathcal{ALC}$, introduced later as needed. Fix countably infinite sets N_C , N_R , and N_I of *concept names*, *role names*, and *individual names*, respectively. Prob- $\mathcal{ALC}$ concepts are formed according to the syntax rule

$$C ::= A \mid \neg C \mid C \sqcap D \mid \exists r.C \mid P_{\geq n}C \mid \exists P_{*n}r.C$$

where A ranges over N_C , C and D over concepts, r over N_R , $*$ over $\{>, \geq\}$, and n over rational numbers from the interval $[0, 1]$. We use the usual abbreviations $C \sqcup D$ for $\neg(\neg C \sqcap \neg D)$, $\forall r.C$ for $\neg \exists r.\neg C$, $\top$ for $A \sqcap \neg A$, and $\perp$

for $\neg \top$. Moreover, $P_{<n}C$ abbreviates $\neg P_{\geq n}C$, $P_{\leq n}C$ abbreviates $P_{\geq 1-n}\neg C$, and $P_{>n}C$ abbreviates $P_{<1-n}\neg C$.

In DLs, TBoxes are used to formalize an ontology, and ABoxes store instance data. In Prob- $\mathcal{ALC}$, a *TBox* is simply a finite set of *concept inclusions* (CIs) $C \sqsubseteq D$. A (*probabilistic*) *ABox* is an expression formed according to the rule

$$\mathcal{A} ::= C(a) \mid r(a, b) \mid \neg \mathcal{A} \mid \mathcal{A} \wedge \mathcal{A}' \mid P_{\geq n}\mathcal{A}$$

where C , r , and n are as above, $a, b \in N_I$, and $\mathcal{A}, \mathcal{A}'$ range over probabilistic ABoxes. Abbreviations $P_{*n}\mathcal{A}$ for $*$ $\in \{\leq, >, <\}$ are defined as for concepts. A *knowledge base* is a pair $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ with $\mathcal{T}$ a TBox and $\mathcal{A}$ an ABox.

The semantics of standard DLs such as $\mathcal{ALC}$ is based on interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set called the *domain* and $\cdot^{\mathcal{I}}$ is an *interpretation function* that maps each $A \in N_C$ to a subset $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, each $r \in N_R$ to a subset $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and each $a \in N_I$ to an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$. We refer to (Baader et al. 2003) for more details. To provide a semantics for Prob- $\mathcal{ALC}$, we generalize interpretations to probabilistic interpretations, in analogy to Halpern's generalization of FO structures to Type 2 probabilistic FO structures (1990). A *probabilistic interpretation* takes the form

$$\mathcal{I} = (\Delta^{\mathcal{I}}, W, (\mathcal{I}_w)_{w \in W}, \mu)$$

where $\Delta^{\mathcal{I}}$ is the (non-empty) domain, W a non-empty set of *possible worlds*, μ a discrete probability distribution on W , and for each $w \in W$, $\mathcal{I}_w$ is a classical DL interpretation with domain $\Delta^{\mathcal{I}}$ such that $a^{\mathcal{I}_w} = a^{\mathcal{I}_{w'}}$ for all $a \in N_I$ and $w, w' \in W$. Since $a^{\mathcal{I}_w}$ does not depend on w , we write only $a^{\mathcal{I}}$. We usually write $C^{\mathcal{I}, w}$ for $C^{\mathcal{I}_w}$, and likewise for $r^{\mathcal{I}, w}$. For concept names A and role names d , we define the probability

- $p_d^{\mathcal{I}}(A)$ that $d \in \Delta^{\mathcal{I}}$ is an A as $\mu(\{w \in W \mid d \in A^{\mathcal{I}, w}\})$;
- $p_{d,e}^{\mathcal{I}}(r)$ that $d, e \in \Delta^{\mathcal{I}}$ are related by r as $\mu(\{w \in W \mid (d, e) \in r^{\mathcal{I}, w}\})$.

Next, we extend $p_d^{\mathcal{I}}(A)$ to complex concepts C and define the extension $C^{\mathcal{I}, w}$ of complex concepts by mutual recursion on C . The definition of $p_d^{\mathcal{I}}(C)$ is exactly as in the base case, with A replaced by C . The extension of complex concepts is defined as follows:

$$\begin{aligned} (\neg C)^{\mathcal{I}, w} &= \{d \in \Delta^{\mathcal{I}} \mid d \notin C^{\mathcal{I}, w}\} \\ (C \sqcap D)^{\mathcal{I}, w} &= \{d \in \Delta^{\mathcal{I}} \mid d \in C^{\mathcal{I}, w} \text{ and } d \in D^{\mathcal{I}, w}\} \\ (\exists r.C)^{\mathcal{I}, w} &= \{d \in \Delta^{\mathcal{I}} \mid \exists e \in C^{\mathcal{I}, w} : (d, e) \in r^{\mathcal{I}, w}\} \\ (P_{\geq n}C)^{\mathcal{I}, w} &= \{d \in \Delta^{\mathcal{I}} \mid p_d^{\mathcal{I}}(C) \geq n\} \\ (\exists P_{*n}r.C)^{\mathcal{I}, w} &= \{d \in \Delta^{\mathcal{I}} \mid \exists e \in C^{\mathcal{I}, w} : p_{d,e}^{\mathcal{I}}(r) * n\} \end{aligned}$$

A probabilistic interpretation $\mathcal{I}$ *satisfies* a concept inclusion $C \sqsubseteq D$ (written $\mathcal{I} \models C \sqsubseteq D$) if $C^{\mathcal{I}, w} \subseteq D^{\mathcal{I}, w}$ for all $w \in W$. It is a *model* of a TBox $\mathcal{T}$ if it satisfies all concept inclusions in $\mathcal{T}$.

To give a semantics to probabilistic ABoxes $\mathcal{A}$, we again use mutual recursion, defining the probability $p^{\mathcal{I}}(\mathcal{A})$ that $\mathcal{A}$ is true as

$$p^{\mathcal{I}}(\mathcal{A}) = \mu(\{w \in W \mid \mathcal{I}, w \models \mathcal{A}\})$$

and defining when a world w of $\mathcal{I}$ satisfies $\mathcal{A}$ (written $\mathcal{I}, w \models \mathcal{A}$) as follows:

$$\begin{aligned} \mathcal{I}, w \models C(a) & \text{ iff } a^{\mathcal{I}} \in C^{\mathcal{I}, w} \\ \mathcal{I}, w \models r(a, b) & \text{ iff } (a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}, w} \\ \mathcal{I}, w \models \neg \mathcal{A} & \text{ iff } \mathcal{I}, w \not\models \mathcal{A} \\ \mathcal{I}, w \models \mathcal{A} \wedge \mathcal{A}' & \text{ iff } \mathcal{I}, w \models \mathcal{A} \wedge \mathcal{I}, w \models \mathcal{A}' \\ \mathcal{I}, w \models P_{\geq n}(\mathcal{A}) & \text{ iff } p^{\mathcal{I}}(\mathcal{A}) \geq n \end{aligned}$$

According to this semantics, the ABox assertions $(P_{*n}C)(a)$ and $P_{*n}(C(a))$ are equivalent, so that we shall not distinguish them in the sequel. We say that $\mathcal{I}$ is a *model* of $\mathcal{A}$ if $\mathcal{I}, w \models \mathcal{A}$ for some w ; it is a *model* of a knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ if it is a model of both $\mathcal{T}$ and $\mathcal{A}$.

We note that the standard translation from $\mathcal{ALC}$ to FO, as e.g. described in (Baader et al. 2003), can be extended to a translation from Prob- $\mathcal{ALC}$ to Type 2 probabilistic FO in a straightforward way. In particular, each concept $P_{\geq n}C$ is translated to the probabilistic FO formula $w(C^{\#}(x)) \geq n$, where $C^{\#}(x)$ is the standard translation of C . The translation of TBoxes and ABoxes is just as simple.

We call a knowledge base $\mathcal{K}$ *consistent* if it has a model. Deciding consistency of knowledge bases is the main reasoning problem studied in this paper. As usual, other standard reasoning problems can be reduced to it. For example, an ABox $\mathcal{A}'$ is called a *consequence* of a knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ if every model of $\mathcal{K}$ is also a model of $\mathcal{A}'$. We can reduce ABox consequence to KB consistency since $\mathcal{A}'$ is a consequence of $\mathcal{K}$ iff the KB $(\mathcal{T}, \mathcal{A} \wedge \neg \mathcal{A}')$ is inconsistent.

Examples

We illustrate how Prob- $\mathcal{ALC}$ can be used to model medical knowledge, revisiting the example classes from SNOMED CT given in the introduction. E.g. ‘*probable tubo-ovarian abscess*’ (Bodenreider, Smith, and Burgun 2004) can be modelled in Prob- $\mathcal{ALC}$ as

$$P_{\geq \alpha} \text{TuboOvarianAbscess}$$

which describes findings that are a tubo-ovarian abscess with probability at least α . We note that this concept is not subsumed by (i.e., does not imply) $\text{TuboOvarianAbscess}$, correctly capturing the fact that a probable tubo-ovarian abscess need not actually be a tubo-ovarian abscess. If we focus on patients instead of on findings, we can distinguish between

$$\exists \text{hasAbnormality}. P_{\geq \alpha} \text{TuboOvarianAbscess},$$

which describes patients who have an abnormality (e.g. a sonographic irregularity) that is a tubo-ovarian abscess with probability at least α , and

$$P_{\geq \alpha} \exists \text{hasDisease}. \text{TuboOvarianAbscess},$$

which describes patients who have a tubo-ovarian abscess with probability at least α (and otherwise do not necessarily have any abnormality). We emphasize that uncertain diagnoses may very well have definite consequences. For example, Lyme disease is typically treated with antibiotics even

when the diagnosis is not entirely certain, due to the combination of the graveness of the disease and the difficulty of diagnosing it with certainty:

$$P_{\geq 0.8} \exists \text{hasDisease}. \text{LymeDisease} \sqsubseteq \exists \text{recommendedTreatment}. \text{Antibiotic}$$

To illustrate probabilistic roles, consider the SNOMED CT concept ‘*natural death with probable cause suspected*’, which expresses two things: on the one hand, there is at least one cause that is considered probable, and on the other hand, no cause is certain. Using probabilistic roles, we can model this as the concept

$$\text{NaturalDeath} \sqcap \exists P_{\geq \alpha} \text{hasCause}. \top \sqcap \neg \exists P_{\geq \beta} \text{hasCause}. \top$$

expressing that there is some (unspecified) phenomenon which is believed to be the cause of death with probability at least α , but nothing is believed to be the cause of death with probability more than β .

As noted in the introduction, uncertainty of instance data is ubiquitous in medicine. Consider, e.g., a scenario where John exhibits fatigue symptoms. Since fatigue is an unspecific symptom of Lyme disease, we may model this by a Prob- $\mathcal{ALC}$ ABox such as

$$\text{Fatigue}(s1) \quad \text{hasSymptom}(\text{John}, s1) \quad P_{\geq 0.1}(C(s1)).$$

where we abbreviate $C = \exists \text{hasCause}. \text{LymeDisease}$. If, however, John additionally exhibits unclear fever, which is also unspecifically related to Lyme disease, one will be more inclined to attribute also John’s fatigue to the suspected case of Lyme disease, so that we update the ABox with assertions such as

$$\begin{aligned} & \text{Fever}(s2) \quad \text{hasSymptom}(\text{John}, s2) \\ & P_{\geq 0.3}(C(s2)) \quad P_{\geq 0.2}(C(s1)) \quad P_{\geq 0.15}(C(s1) \sqcap (C(s2))). \end{aligned}$$

Note that the ABox states that the attributions of the two symptoms to Lyme disease are not independent: their joint probability is higher than the product of the individual probabilities.

Semantic Considerations

We discuss some relevant semantic aspects of Prob- $\mathcal{ALC}$ and related logics and point out several interesting extensions.

Linear inequalities as concepts. In contrast to Prob- $\mathcal{ALC}$, Type 2 probabilistic FO admits the formation of unrestricted linear and polynomial inequalities over probabilities. This inspires the extension Prob- $\mathcal{ALC}^{\text{lineq}}$ (resp. Prob- $\mathcal{ALC}^{\text{polyeq}}$) of Prob- $\mathcal{ALC}$, in which (i) concepts $P_{\geq n}C$ are replaced with linear (resp. polynomial) inequalities $\mathcal{E}$ over expressions $P(C)$, C a concept; and (ii) ABoxes $P_{\geq n}(\mathcal{A})$ are replaced with linear (resp. polynomial) inequalities $\mathcal{E}$ over expressions $P(\mathcal{A})$, $\mathcal{A}$ an ABox. The semantics of a polynomial concept inequality $\mathcal{E}$ is that $\mathcal{E}^{\mathcal{I}, w}$ contains precisely those $d \in \Delta^{\mathcal{I}}$ such that the inequality $\mathcal{E}$ is satisfied when each $P(C)$ in it is replaced by $p_d^{\mathcal{I}}(C)$. The semantics of polynomial ABox inequalities is defined analogously.

To illustrate the use of linear inequalities, we make a brief detour to *qualitative* reasoning about probabilities. There are a number of different proposals, of which we consider two. First, Gärdenfors (1975) considers a logic with a binary operator ‘more probable than’. This can be captured in Prob- $\mathcal{ALC}^{\text{lineq}}$ since ‘ C is more probable than D ’ corresponds to the linear concept inequality $P(C) > P(D)$. Second, Herzig (2003) proposes an operator ‘more probably than not’, i.e., the probability of an event is higher than its complement. In Prob- $\mathcal{ALC}^{\text{lineq}}$, this can be expressed as $P(C) > P(\neg C)$. We note that, in general, there is no consensus as to whether subjective probabilities in medical ontologies should be represented quantitatively as in the previous section or qualitatively, e.g. by defining a ‘probable tubo-ovarian abscess’ as

$$P(\text{TuboOvarianAbscess}) > c \cdot P(\neg \text{TuboOvarianAbscess})$$

for some constant c (a generalization of Herzig’s approach). We do believe that it depends on the concrete application which modelling is more appropriate and note that Prob- $\mathcal{ALC}^{\text{lineq}}$ and Prob- $\mathcal{ALC}^{\text{polyeq}}$ support both views.

Polynomial inequalities are strictly more expressive than linear ones. In particular, they capture several kinds of independence constraints, as discussed in the following.

Independence constraints. The semantics of Prob- $\mathcal{ALC}_c$ has no built-in independence of probabilistic events. For example, although one might expect that the ABox

$$P_{\geq 0.5} \text{NearSighted}(\text{John}) \wedge P_{\geq 0.5} \text{Diabetic}(\text{John})$$

has the consequence $P_{\geq 0.25}(\text{NearSighted} \sqcap \text{Diabetic})(\text{John})$ since the two involved events are intuitively independent, it is not hard to see that the semantics of Prob- $\mathcal{ALC}$ does not support such a deduction. Polynomial inequalities make it possible to add independence constraints. For example, we can obtain the desired consequence in the example above when we add the TBox statement $\top \sqsubseteq C$, where C is the polynomial concept inequality

$$P(\text{NearSighted} \sqcap \text{Diabetic}) = P(\text{NearSighted}) \cdot P(\text{Diabetic}).$$

For the use in lower bounds and other negative results, we single out a particular, rather simple form of independence constraint as follows.

We use Prob- $\mathcal{ALC}^{\text{indep}}$ to denote the extension of Prob- $\mathcal{ALC}$ obtained by admitting *independence constraints* of the form $\text{indep}(C, D)$ in the TBox, with C, D concepts. An interpretation $\mathcal{I}$ satisfies $\text{indep}(C, D)$ if for all $d \in \Delta^{\mathcal{I}}$, we have $p_d^{\mathcal{I}}(C) \cdot p_d^{\mathcal{I}}(D) = p_d^{\mathcal{I}}(C \sqcap D)$.

Note that these independence constraints are strictly weaker than polynomial inequalities. For example, the latter also allow to express independence of more than two events. It is of course also possible to consider stronger or different forms of independence constraints that cannot even be expressed by polynomial concept inequalities such as the independence of events $p_d^{\mathcal{I}}(A)$ and $p_e^{\mathcal{I}}(A)$ that involve different domain elements.

Conditional probabilities. In some probabilistic DLs, conditional probabilities play a very central role (Lukasiewicz 2008). We do not include them in Prob- $\mathcal{ALC}$ because in the case of *subjective* probabilities, interesting uses of conditional probabilities are inherently *non-monotonic*, thus in conflict with design goal that Prob- $\mathcal{ALC}$ should be a fragment of (purely monotonic) probabilistic FO.

To give a concrete example, assume that we add TBox expressions of the form $P_{*n}(C|D)$, $* \in \{\geq, \leq\}$ with a standard semantics for conditional probabilities, i.e., $P_{*n}(C|D)$ is satisfied by an interpretation $\mathcal{I}$ whenever for all $d \in \Delta^{\mathcal{I}}$, we have $p_d^{\mathcal{I}}(C \sqcap D) * n \cdot p_d^{\mathcal{I}}(D)$. Now take the TBox

$$\begin{aligned} \top &\sqsubseteq P_{\leq 0.01} \text{HasLymeDisease} \\ \top &\sqsubseteq P_{\geq 0.8} (\text{HasLymeDisease} | \text{HasPositiveSerology}) \end{aligned}$$

expressing the beliefs that (i) everybody has Lyme disease with probability ≤ 0.01 and (ii) if somebody has a positive serology, his probability of having Lyme disease is ≥ 0.8 (contrast this with the different reading under the statistical view of probabilities described in the introduction, where conditional probabilities make perfect sense also in a monotonic logic). If we take an ABox $\{\text{Patient}(\text{John})\}$, we get $P_{\leq 0.01} \text{HasLymeDisease}(\text{John})$ as an expected consequence. If we now learn that John has positive serology and extend our ABox to $\{\text{Patient}(\text{John}), \text{HasPositiveSerology}(\text{John})\}$, we would intuitively expect a *revision* of our belief to $P_{\geq 0.08} \text{HasLymeDisease}(\text{John})$. However, due to the monotonicity of our semantics, the only effect is that the extended ABox is inconsistent with the above TBox.

Note that our choice of a monotonic FO semantics and the non-inclusion of conditional probabilities is deliberate. Indeed, we consider the clear, monotonic semantics of Prob- $\mathcal{ALC}$ to be one of its main features and, following Halpern, believe that it is important to first understand the monotonic aspects of probability before mixing in non-monotonic ones, which are typically much more controversial.

Worlds with probability zero. Note that probabilistic interpretations may contain worlds with probability 0. These worlds represent situations that are infinitely improbable, but not per se impossible (the latter corresponds to a logical inconsistency). The following lemma, which is exploited in the algorithms developed later on, shows that worlds with probability 0 play a rather special role in models of Prob- $\mathcal{ALC}$ knowledge bases. The intuitive reason for this special behaviour is that the probabilistic operators range only over worlds with probability greater than 0.

Lemma 1. *If a Prob- $\mathcal{ALC}^{\text{indep}}$ knowledge base $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is consistent, then it has a model $\mathcal{I}$ with set of worlds W such that for some $w_0 \in W$, we have*

1. $\mu(w_0) = 0$ and $\mathcal{I}, w_0 \models \mathcal{A}$;
2. $\mu(w) > 0$ for all $w \in W \setminus \{w_0\}$.

Prob- $\mathcal{ALC}$ without Probabilistic Roles

We proceed to analyze the computational complexity of Prob- $\mathcal{ALC}$ and its fragments. We start our journey by considering the rather well-behaved fragment Prob- $\mathcal{ALC}_c$ that

is obtained from Prob- $\mathcal{ALC}$ by disallowing the constructor $\exists P_{\geq n} r.C$. In other words, Prob- $\mathcal{ALC}_c$ can speak only about the probability of concepts, but not about that of roles. We believe that Prob- $\mathcal{ALC}_c$ provides sufficient probabilistic expressive power for most applications, which is also illustrated by the examples above.

Our main result on Prob- $\mathcal{ALC}_c$ is that KB consistency is EXPTIME-complete, and thus no more complex than in non-probabilistic $\mathcal{ALC}$. To prove the upper bound, we develop a procedure of the ‘type elimination’ kind that checks for the existence of a decomposed representation of a model that we call a quasimodel. To deal with probabilities, the procedure involves solving systems of linear inequalities. Technically, the reason for the computational well-behavedness is due to the fact that, despite its multi-dimensional semantics, Prob- $\mathcal{ALC}_c$ can essentially be viewed as an independent combination of $\mathcal{ALC}$ and probabilistic propositional logic known as the fusion (Wolter 1998; Schröder and Pattinson 2007).

Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be the knowledge base whose consistency is to be decided. We assume w.l.o.g. that the TBox has the form $\top \sqsubseteq C_{\mathcal{T}}$ and use

- $\text{ccl}(\mathcal{K})$ to denote the *concept closure* of $\mathcal{K}$, i.e., the closure under subconcepts and single negation of $\{C_{\mathcal{T}}\} \cup \{C \mid C(a) \in \mathcal{A}\}$ and
- $\text{acl}(\mathcal{K})$ to denote the *ABox closure* of $\mathcal{K}$, i.e., the closure under single negation and sub-ABoxes of $\mathcal{A} \cup \{C(a) \mid C \in \text{ccl}(\mathcal{K}) \wedge a \in \text{Ind}(\mathcal{A})\}$

An *ABox type* for $\mathcal{K}$ is a pair (t, f) where $t \subseteq \text{acl}(\mathcal{K})$ describes a model of an ABox ‘locally’ at a possible world and $f \in \{0, 1\}$ is a flag that describes whether that world has probability 0 or strictly greater than 0. For brevity, we often identify (t, f) with t and then use f_t to denote f . An ABox type is required to satisfy the following conditions:

1. $C_{\mathcal{T}}(a) \in t$ for all $a \in \text{Ind}(\mathcal{A})$;
2. $(\neg C)(a) \in t$ iff $C(a) \notin t$, for all $(\neg C)(a) \in \text{acl}(\mathcal{K})$;
3. $(C \sqcap D)(a) \in t$ iff $C(a), D(a) \in t$, for all $(C \sqcap D)(a) \in \text{acl}(\mathcal{K})$;
4. $\neg \mathcal{A}' \in t$ iff $\mathcal{A}' \notin t$, for all $\neg \mathcal{A}' \in \text{acl}(\mathcal{K})$;
5. $\mathcal{A}' \wedge \mathcal{A}'' \in t$ iff $\mathcal{A}', \mathcal{A}'' \in t$, for all $\mathcal{A}' \wedge \mathcal{A}'' \in \text{acl}(\mathcal{K})$;
6. if $C(b) \in t$ and $r(a, b) \in t$, then $\exists r.C(a) \in t$, for all $r(a, b), \exists r.C(a) \in \text{acl}(\mathcal{K})$.

Let $\mathfrak{A}_{\mathcal{K}}$ denote the set of all ABox types for $\mathcal{K}$. Given a set $T \subseteq \mathfrak{A}_{\mathcal{K}}$ and a type $t_0 \in T$, we define a system of linear inequalities $\mathcal{E}(t_0, T)$ over the variables $(x_t)_{(t,1) \in T}$. Intuitively, a solution to this system tells us how to assign ABox types and probabilities to worlds such that we obtain a model that realizes the type t_0 and is coherent regarding the probabilistic operators. The inequalities in $\mathcal{E}(t_0, T)$ are as follows:

1. get the probabilities right:
 - for each $P_{\geq n} \mathcal{A}' \in t_0$: $\sum_{(t,1) \in T \mid \mathcal{A}' \in t} x_t \geq n$
 - for each $\neg P_{\geq n} \mathcal{A}' \in t_0$: $\sum_{(t,1) \in T \mid \mathcal{A}' \in t} x_t < n$

2. guarantee that all types with non-zero probability satisfy the same probabilistic assertions, where $\mathcal{P}(t)$ denotes the set of all expressions in t of the form $P_{\geq n} \mathcal{A}'$:

$$\sum_{(t,1) \in T \mid \mathcal{P}(t) \neq \mathcal{P}(t_0)} x_t = 0$$

3. if $f_{t_0} = 1$, then add the inequality $x_{t_0} > 0$
4. probabilities sum up to one: $\sum_{(t,1) \in T} x_t = 1$

Fix an individual name $a_{\varepsilon} \in \text{Ind}(\mathcal{A})$ that we use to represent ‘anonymous’ (i.e., non-ABox) domain elements. An *element type* for $\mathcal{K}$ is a pair (t, f) where $t \subseteq \{C(a_{\varepsilon}) \mid C \in \text{ccl}(\mathcal{K})\}$ respects Conditions 1 to 3 in the definition of ABox types and $f \in \{0, 1\}$. Since element types refer only to the fixed individual name a_{ε} , we will often drop that name altogether and simply view an element type as a subset of $\text{ccl}(\mathcal{K})$. As with ABox types, we will often confuse (t, f) with t and write f_t to identify f . Let $\mathfrak{T}_{\mathcal{K}}$ denote the set of all element types for $\mathcal{K}$. The systems of inequalities $\mathcal{E}(t_0, T)$ for $T \subseteq \mathfrak{T}_{\mathcal{K}}$ and $t_0 \in T$ are then defined for element types in literally the same way as for ABox types.

A *quasimodel* for $\mathcal{K}$ is a pair $Q = (T, T')$ with $T \subseteq \mathfrak{A}_{\mathcal{K}}$ and $T' \subseteq \mathfrak{T}_{\mathcal{K}}$. A type $t \in T \cup T'$ is *saturated* in Q if for each $\exists r.C(a) \in t$, there exists a $t' \in T'$ with $t' \supseteq \{C\} \cup \{D \mid \forall r.D(a) \in t\}$ and $f_t = f_{t'}$. A type $t \in T$ (resp. $t \in T'$) is *coherent* in Q if the system of inequalities $\mathcal{E}(t, T)$ (resp. $\mathcal{E}(t, T')$) has a nonnegative real solution. Finally, we call Q *proper* for $\mathcal{K}$ if every $t \in T \cup T'$ is saturated and coherent in Q , and there is a $(t, 0) \in T$ with $\mathcal{A} \in t$. The second condition should be clear in view of Lemma 1.

Lemma 2. $\mathcal{K}$ is consistent iff there is a proper quasimodel for $\mathcal{K}$.

Proof. Since the “ $\Rightarrow$ ” direction is uninteresting, we defer it to the appendix and concentrate on “ $\Leftarrow$ ”. Let $Q = (T, T')$ be a proper quasimodel for $\mathcal{K}$. For each $t \in T$ (resp. $t \in T'$), the system of linear inequalities $\mathcal{E}(t, T')$ (resp. $\mathcal{E}(t, T)$) has a solution in the non-negative reals. It is well-known that, thus, it also has a non-negative *rational* solution (Schrijver 1986). For each $t \in T \cup T'$, fix such a solution δ_t . For a rational number r , we use $\text{den}(r)$ to denote the denominator of r . Set

$$c := 1 / \prod_{(t,t') \in (T \times T) \cup (T' \times T'), f_{t'} = 1} \text{den}(\delta_t(x_{t'}))$$

We now construct an interpretation as follows. Choose a set $W = W^+ \uplus \{w_0\}$ such that $|W^+| = 1/c$ and set $\mu(w) = c$ for all $w \in W^+$ and $\mu(w_0) = 0$. For convenience, we set $f(w) = 0$ if $\mu(w) = 0$ and $f(w) = 1$ otherwise.

We highlight some properties that are central for the following construction. Their proof is straightforward based on the observation that for all $(t, t') \in (T \times T) \cup (T' \times T')$ with $f_{t'} = 1$, there is an integer $n \geq 0$ with $c \cdot n = \delta_t(x_{t'})$.

- (a) For each $t_0 \in T$, there is a mapping $\tau : W^+ \rightarrow \{t \mid (t, 1) \in T\}$ such that for each $(t, 1) \in T$, we have $\sum_{w \in W \mid \tau(w) = t} \mu(w) = \delta_{t_0}(x_t)$.

- (b) For each $(t_0, 0) \in T'$, there is a mapping $\tau : W^+ \rightarrow T'$ such that for each $(t, 1) \in T'$, we have $\sum_{w \in W | \tau(w)=t} \mu(w) = \delta_{t_0}(x_t)$.
- (c) For each $(t_0, 1) \in T'$ and $w_0 \in W^+$, there is a mapping $\tau : W^+ \rightarrow T'$ such that $\tau(w_0) = (t_0, 1)$ and for each $(t, 1) \in T'$, we have $\sum_{w \in W | \tau(w)=t} \mu(w) = \delta_{t_0}(x_t)$.

For (c) observe that, by definition of $\mathcal{E}(t, T')$ and since $f_{t_0} = 1$, we have $\delta_{t_0}(x_{t_0}) > 0$, thus we can achieve $\tau(w_0) = t_0$ as required.

The set of domain elements and the interpretation of concept and role names is constructed inductively along with a mapping $\pi_d : W \rightarrow \mathfrak{T}_K$ such that the following invariant is satisfied:

- (*) for all $d \in \Delta^{\mathcal{I}}$ and $w \in W$, we have $(\pi_d(w), f(w)) \in T'$ or there is a $(t, f(w)) \in T$ such that for some $a \in \text{Ind}(\mathcal{A})$, we have $\pi_d(w) = \{C(a_\varepsilon) \mid C(a) \in t\}$.

The details of the construction are as follows:

- There is a $(t_A, 0) \in T$ such that $\mathcal{A} \in t$. To start the construction, we set $\Delta^{\mathcal{I}} = \text{Ind}(\mathcal{A})$. By (a), we find a mapping $\tau : W^+ \rightarrow \{t \mid (t, 1) \in T\}$ such that for each $(t, 1) \in T$, we have $\sum_{w \in W | \tau(w)=t} \mu(w) = \delta_{t_A}(x_t)$. Set $\tau(w_0) := t_A$ and further set
 - $a^{\mathcal{I}} = a$ for each $a \in \mathbf{N}_I$;
 - $A^{\mathcal{I}, w} = \{a \mid A(a) \in \tau(w)\}$ for each concept name A and $w \in W$;
 - $r^{\mathcal{I}, w} = \{(a, b) \mid r(a, b) \in \tau(w)\}$ for each role name r and $w \in W$.

For all $a \in \text{Ind}(\mathcal{A})$ and $w \in W$, set $\pi_a(w) = \{C(a_\varepsilon) \mid C(a) \in \tau(w)\}$. Clearly, (*) is satisfied.

- Repeat the following step indefinitely. Choose $w \in W$, $d \in \Delta^{\mathcal{I}}$, and $\exists r.C \in \pi_d(w)$ such that there is no $e \in \Delta^{\mathcal{I}}$ with $(d, e) \in r^{\mathcal{I}, w}$ and $C \in \pi_e(w)$. By saturatedness and (*), there is a $(t, f(w)) \in T'$ such that $t \supseteq \{C\} \cup \{D \mid \forall r.D \in \pi_d(w)\}$. Add a fresh element e to $\Delta^{\mathcal{I}}$ and distinguish two cases:
 - if $f(w) = 0$, then by (b) there is a mapping $\pi_e : W^+ \rightarrow T'$ such that for each $(t', 1) \in T'$, we have $\sum_{w' \in W | \pi_e(w')=t'} \mu(w') = \delta_t(x_{t'})$; set $\pi_e(w_0) := t$;
 - if $f(w) = 1$, then by (c) there is a mapping $\pi_e : W^+ \rightarrow T'$ such that $\pi_e(w) = t$ and for each $(t', 1) \in T'$, we have $\sum_{w' \in W | \pi_e(w')=t'} \mu(w') = \delta_t(x_{t'})$; set $\pi_e(w_0) := (t, 0)$.
- Set $r^{\mathcal{I}, w} = r^{\mathcal{I}, w} \cup \{(d, e)\}$. For all concept names A and $w' \in W$, set $A^{\mathcal{I}, w'} = A^{\mathcal{I}, w'} \cup \{d\}$ if $A \in \pi_e(w')$.

It is tedious but straightforward to verify that we end up with a model of $\mathcal{K}$. $\square$

Thus, we can decide consistency by checking the existence of a proper quasimodel. This can be done as follows: start with the quasimodel $Q = (\mathfrak{A}_K, \mathfrak{T}_K)$ and then repeatedly delete types from both components that are not saturated or not coherent. If a type $(t, 0)$ with $\mathcal{A} \in t$ survives in the first component after the sets have stabilized, answer ‘satisfiable’. Otherwise, answer ‘unsatisfiable’. It can be proved

that the algorithm decides consistency in Prob- $\mathcal{ALC}_c$ and runs in EXPTIME.

Theorem 3. *Consistency in Prob- $\mathcal{ALC}_c$ is EXPTIME-complete.*

It is interesting to note that the proof of Lemma 2 also establishes a *uniform model property (UMP)*: every consistent Prob- $\mathcal{ALC}_c$ KB has a *uniform model* $\mathcal{I}$, i.e., the probability distribution μ of $\mathcal{I}$ satisfies $\mu(w) = \mu(w')$ for all $w, w' \in W$ with $\mu(w) > 0$ and $\mu(w') > 0$. Our constructions do not yield a *finite model property (FMP)* as the domains of the constructed models may be infinite (whereas the set of worlds is finite). As discussed in the following, this can be fixed by a slight modification which then even yields a *bounded model property (BMP)*. We use the following result about linear programming.

Proposition 4 (Fagin, Halpern, and Megiddo 1990). *If a system of r linear inequalities with integer coefficients with length ℓ has a non-negative solution, then it has a non-negative solution with at most r entries positive, and where the size of each member of the solution is $\mathcal{O}(r\ell + r \log(r))$.*

Here, the *length* of an integer denotes the number of bits of its binary representation and the *size* of a rational number is the sum of the lengths of the binary representations of the numerator and denominator. A careful analysis shows that the number of worlds in the constructed models is bounded by $2^{2^{\mathcal{O}(K)}}$. It is then simple to obtain the same bound on the number of DL elements: in the model construction, we can stop adding new domain elements as soon as in every world, there is an element of every type. We thus obtain the following combination of a UMP and a BMP.

Corollary 5. *Every consistent Prob- $\mathcal{ALC}_c$ KB has a uniform model in which the number of worlds and domain elements is bounded by $2^{2^{\mathcal{O}(K)}}$.*

We have no proof that a super-polynomial number of worlds can actually be enforced (and neither a super-exponential number of domain elements). It is straightforward to extend all constructions in this section to Prob- $\mathcal{ALC}_c^{\text{lineq}}$.

Polynomial Inequalities

The EXPTIME upper bound can be adapted to the extension Prob- $\mathcal{ALC}_c^{\text{polyeq}}$ of Prob- $\mathcal{ALC}_c$ with polynomial inequalities. This extension is particularly useful because of Prob- $\mathcal{ALC}_c^{\text{polyeq}}$ ’s ability to express independence constraints.

The main idea of the upper bound adaptation is to generalize the systems of inequalities $\mathcal{E}(t_0, T)$ for ABox types and $\mathcal{E}(t_0, T')$ for element types to take into account the additional expressive power of polynomial inequalities. More specifically, Point 1 of the definition of $\mathcal{E}(t_0, T)$ is replaced with the following:

- add each polynomial ABox inequality $\mathcal{E} \in t_0$ to $\mathcal{E}(t_0, T)$ after replacing each expression $P(\mathcal{A})$ that occurs in $\mathcal{E}$ with $\sum_{(t,1) \in T | \mathcal{A} \in t} x_t$;
- for each $\mathcal{E}(a) \in t_0$, add the polynomial concept inequality $\mathcal{E}$ to $\mathcal{E}(t_0, T')$ after replacing each expression $P(C)$ that occurs in $\mathcal{E}$ with $\sum_{(t,1) \in T | C(a) \in t} x_t$.

In Point 2 of the definition of $\mathcal{E}(t_0, T)$, $\mathcal{P}(t)$ now denotes the set of all polynomial ABox inequalities in t plus all elements $\mathcal{E}(a) \in t$ with $\mathcal{E}$ a polynomial concept inequality. Denote the resulting systems by $\mathcal{E}^*(t_0, T)$. The generalization of $\mathcal{E}(t_0, T')$ into $\mathcal{E}^*(t_0, T')$ is analogous.

Clearly, $\mathcal{E}^*(t_0, T)$ is a system of polynomial inequalities instead of linear ones. Consequently, the complexity of solving such systems goes up from PTIME to PSPACE, which is currently the best known upper bound (Canny 1988) (the best known lower bound is NP). Thus, our EXPTIME upper bound for consistency becomes an EXPSPACE upper bound. To push it back down to EXPTIME, we first make the following observation, which is originally due to (Fagin, Halpern, and Megiddo 1990).

Lemma 6. *If a system $\mathcal{E}^*(t_0, T)$ has a solution, then it has a solution with at most $|\mathcal{K}|^2$ variables non-zero, and likewise for systems $\mathcal{E}^*(t_0, T')$.*

Proof. Let δ be a solution for $\mathcal{E}^*(t_0, T)$. For each $\mathcal{A} \in \text{acl}(\mathcal{K})$, set $\delta(\mathcal{A}) = \sum_{t \in T | \mathcal{A} \in t} \delta(x_t)$. The system $\hat{\mathcal{E}}(t_0, T)$ is obtained from $\mathcal{E}^*(t_0, T)$ by replacing the inequalities derived from polynomial (ABox and concept) inequalities with

$$\sum_{(t,1) \in T | \mathcal{A} \in t} x_t = \delta(\mathcal{A}) \quad \text{for all } \mathcal{A} \in \text{acl}(\mathcal{K})$$

Since $\mathcal{E}^*(t_0, T)$ is satisfiable, so is $\hat{\mathcal{E}}(t_0, T)$. Since $\hat{\mathcal{E}}(t_0, T)$ consists of $\mathcal{O}(|\mathcal{K}|^2)$ linear inequalities, Proposition 4 yields the existence of a solution with at most $\mathcal{O}(|\mathcal{K}|^2)$ non-zeros. By construction of $\hat{\mathcal{E}}(t_0, T)$, this solution is also a solution for $\mathcal{E}^*(t_0, T)$. $\square$

Instead of checking a system $\mathcal{E}^*(t_0, T)$ directly for solvability, Lemma 6 allows us to consider all systems $\mathcal{E}^*(t_0, T^*)$ with $T^* \subseteq T$ such that $|T^*| \leq |\mathcal{K}|^2$ and $t_0 \in T^*$. If any of these systems has a solution δ , then we obtain a solution for $\mathcal{E}^*(t_0, T)$ by setting $\delta(x_t) = 0$ for all $t \in T \setminus T^*$. Conversely, solvability of $\mathcal{E}^*(t_0, T)$ and Lemma 6 implies the existence of a solution δ for $\mathcal{E}^*(t_0, T)$ and a subset $T^* \subseteq T$ such that $|T^*| \leq |\mathcal{K}|^2$ and $\delta(x_t) = 0$ for all $t \in T \setminus T^*$. Clearly, δ is also a solution for $\mathcal{E}^*(t_0, T^*)$. Each system $\mathcal{E}^*(t_0, T^*)$ is of size polynomial in $|\mathcal{K}|$ and there are exponentially many such systems. Thus, the solvability of $\mathcal{E}^*(t_0, T)$ can be checked in EXPTIME.

A second problem that arises is the loss of the UMP, which means that the model construction underlying Lemma 2, which assumes a rational solution and produces uniform models, no longer works.

Lemma 7. *Prob- $\mathcal{ALC}_c^{\text{indep}}$ does not enjoy the UMP.*

Proof. Consider the knowledge base $\mathcal{K} = (T, \{C(a)\})$ with

$$\begin{aligned} \mathcal{T} &= \{\text{indep}(A, B)\} \\ C &= P_{=\frac{1}{8}}(A \sqcap B) \sqcap P_{=\frac{1}{8}}(\neg A \sqcap \neg B) \end{aligned}$$

It is not hard to see that $\mathcal{K}$ is consistent: simply take the

probabilistic interpretation $\mathcal{I}$ with

$$\begin{aligned} \Delta^{\mathcal{I}} &= \{d\} \\ W &= \{w_S \mid S \subseteq \{A, B\}\} \\ \mu(w_{\{A,B\}}) &= \mu(w_{\emptyset}) = \frac{1}{8} \\ \mu(w_{\{A\}}) &= \frac{3}{8} - \frac{1}{\sqrt{8}} \\ \mu(w_{\{B\}}) &= \frac{3}{8} + \frac{1}{\sqrt{8}} \\ X^{\mathcal{I}, w_S} &= \begin{cases} \{d\} & \text{if } X \in S \\ \emptyset & \text{otherwise} \end{cases} \quad \text{for } X \in \{A, B\} \end{aligned}$$

It is easy to verify that $\mathcal{I}$ is a (non-uniform) model of $\mathcal{K}$. In the appendix, we show that *all* models $\mathcal{I}$ of $\mathcal{K}$ satisfy

$$p_d^{\mathcal{I}}(A) = \frac{1}{2} \pm \frac{1}{\sqrt{8}}$$

for every $d \in \Delta^{\mathcal{I}}$. Since these numbers are irrational and every world in a uniform model has rational probability, there is no uniform model of $\mathcal{K}$. $\square$

We now describe a more subtle model construction that replaces the one used in the proof of Lemma 2 and works also for real-valued solutions. Of course, it constructs non-uniform models. Proper quasimodels for Prob- $\mathcal{ALC}_c^{\text{polyeq}}$ are defined exactly as for Prob- $\mathcal{ALC}_c$, except that the generalized systems $\mathcal{E}^*(t, T)$ and $\mathcal{E}^*(t, T')$ are used.

Lemma 8. *$\mathcal{K}$ is consistent iff there is a proper quasimodel for $\mathcal{K}$.*

Proof. Since the “ $\Leftarrow$ ” direction is straightforward, we only consider the “ $\Rightarrow$ ” direction in the proof of Lemma 2. In particular, we show how to define W and μ such that the Properties (a) to (c) still hold. For each $t \in T$ (resp. $t \in T'$), let δ_t be a non-negative solution of $\mathcal{E}(t, T)$ (resp. $\mathcal{E}(t, T')$). Let $\frac{1}{u} \in (0, 1)$ be a rational number such that $\frac{1}{u} < \frac{1}{2} \cdot \delta_t(t')$ for all $(t, t') \in (T \times T) \cup (T' \times T')$ with $f_{t'} = 1$ and $\delta_t(t') > 0$. Define a set

$$S_\varepsilon = \{\delta_t(t') \bmod u \mid (t, t') \in (T \times T) \cup (T' \times T'), f_{t'} = 1\}.$$

and let S be the smallest set of real numbers from the interval $[0, \frac{1}{u}]$ that contains 0 and $\frac{1}{u}$ and is closed under addition and subtraction of values from S_ε . Let $s_0, \dots, s_k$ be an enumeration of the elements of S in ascending order. It is not hard to see that S is finite. Define the set of worlds

$$W := \{0\} \cup \{(x, y) \mid 1 \leq x \leq u \wedge 1 \leq y \leq k\}$$

and set $W^+ = W \setminus \{0\}$. Intuitively, we can think of worlds as partitioning the interval $[0, 1]$. Actually, there are two layers of partitions: the x -component uniformly partitions the interval $[0, 1]$ into subintervals of length $\frac{1}{u}$. The y -component then partitions these subintervals in a non-uniform way. Put $\mu(0) = 0$ and $\mu(x, y) = s_y - s_{y-1}$. We now show that (a) to (c) hold:

- For (a), let $t_0 \in T$. We define a mapping $\tau : W^+ \rightarrow \{t \mid (t, 1) \in T\}$ as follows. Let $t_1, \dots, t_k$ be an enumeration of those elements $t \in T$ such that $\delta_{t_0}(t) > 0$. Intuitively, we want to select intervals for $t_1, \dots, t_k$ such that $[0, 1]$ is covered and the interval for t_1 meets that of t_2 meets that of t_3 etc. For $1 < i < k$, set

- $b_i = \sum_{1 \leq p < i} \delta_{t_0}(t_p)$ (the beginning point of the interval for t_i);
- $e_i = \sum_{1 \leq p \leq i} \delta_{t_0}(t_p)$ (the end point of the interval for t_i);
- $r_1 = \delta_{t_0}(t_1) \bmod \frac{1}{u}$
- $r_i = \delta_{t_0}(t_i) \bmod \frac{1}{u} - a_i$; (the ‘remainder’ that t_i leaves in the last ‘ x -interval’ that it overlaps with);
- $a_i = \frac{1}{u} - d_{i-1}$; (the length of the t_i -part of the first ‘ x -interval’ that t_i overlaps with);
- $a_k = \frac{1}{u} - d_{k-1}$.

Now set $\tau(0) = 0$ and $\tau(x, y) = t_i$, $1 \leq i \leq k$, whenever one of the following holds:

- $b_i \leq \frac{x-1}{u}$ and $\frac{x}{u} \leq e_i$;
- $\frac{x-1}{u} + s_y \leq e_i$ and $\frac{x}{u} > b_{i+1}$,
- $\frac{x-1}{u} < e_{i+1}$ and $\frac{x-1}{u} + s_y > b_i$.

It can be verified that for all $t \in T$, we have $\sum_{w \in W | \tau(w)=t} \mu(w) = \delta_{t_0}(x_t)$ as required.

- (b) can be proved similarly to (a);
- For (c), let $(t_0, 1) \in T'$ and $(x, y) \in W^+$. Fix a mapping $\tau : W^+ \rightarrow \{t \mid (t, 1) \in T'\}$ exactly as in (a). To achieve that $\tau(x, y) = t_0$, select an $x' \in \{1, \dots, u\}$ such that $\tau(x, y) = t_0$ for all $y \in \{1, \dots, k\}$. Such an x' exists by construction of τ and since $\frac{1}{u} < \frac{1}{2} \cdot \delta_{t_0}(t_0)$. Now swap the values of (x, y) and (x', y) for all $y \in \{1, \dots, k\}$. It is easy to see that, as before the swaps, we have $\sum_{w \in W | \tau(w)=t} \mu(w) = \delta_{t_0}(x_t)$. $\square$

We have thus established the following result.

Theorem 9. *Consistency in Prob- $\mathcal{ALC}_c^{\text{polyeq}}$ is EXPTIME-complete.*

Prob- $\mathcal{ALC}$ with Probabilistic Roles

Reasoning in full Prob- $\mathcal{ALC}$ is much more challenging than in Prob- $\mathcal{ALC}_c$. This is unsurprising given the similarity of Prob- $\mathcal{ALC}$ and the modal DL $S5_{\mathcal{ALC}}$ together with the well-known difficulties of dealing with global roles in $S5_{\mathcal{ALC}}$ (which correspond to probabilistic roles in Prob- $\mathcal{ALC}$) (Gabbay et al. 2003; Artale, Lutz, and Toman 2007). In fact, we can use the connection to $S5_{\mathcal{ALC}}$ to carry over some interesting first observations. For example, we can adapt a well-known construction from $S5_{\mathcal{ALC}}$ to show that Prob- $\mathcal{ALC}$ does not enjoy the *finite model property (FMP)*: the KB $\mathcal{K} = (\mathcal{T}, \{P_{>0} \neg A(a)\})$ with

$$\mathcal{T} = \{\neg A \sqsubseteq P_{>0} A, A \sqsubseteq \exists P_{\geq 1} r. \neg A, \neg A \sqsubseteq \forall P_{\geq 1} r. \neg A\}$$

is consistent, but all models of $\mathcal{K}$ comprise both an infinite set of worlds and an infinite domain. It is also straightforward to find a polytime reduction of $S5_{\mathcal{ALC}}$ with global roles to Prob- $\mathcal{ALC}$, which yields the following lower bound.

Theorem 10. *Consistency in Prob- $\mathcal{ALC}$ is 2-EXPTIME-hard, even for ABoxes of the form $A(a)$.*

In the case of $S5_{\mathcal{ALC}}$, the 2-EXPTIME lower bound is tight. For Prob- $\mathcal{ALC}$, we do not even know whether consistency is decidable. What we *do* know, however, is that consistency in Prob- $\mathcal{ALC}$ becomes undecidable once we add independence constraints or admit linear inequalities as concepts. This can be proved by a reduction of the halting problem of two-register machines.

Theorem 11. *Consistency in Prob- $\mathcal{ALC}^{\text{indep}}$ and Prob- $\mathcal{ALC}^{\text{lineq}}$ is undecidable.*

To identify a decidable variant of Prob- $\mathcal{ALC}$ that does include probabilistic roles, we restrict the probability constants that can occur to 0 and 1, i.e., we only admit the probabilistic constructors $P_{>0}C$, $P_{=1}C$, $\exists P_{>0}r.C$, and $\exists P_{=1}r.C$. The resulting logic is called Prob- $\mathcal{ALC}^{01}$. This kind of restriction has previously been used to regain decidability of undecidable probabilistic logics, see for example (Brázdil et al. 2008). It may appear that decidability and a 2-EXPTIME upper bound for Prob- $\mathcal{ALC}^{01}$ can be obtained by a straightforward reduction to $S5_{\mathcal{ALC}}$. Somewhat surprisingly, however, probabilistic ABoxes and in particular the presence of worlds with probability 0 pose serious difficulties for such a reduction; essentially, Prob- $\mathcal{ALC}^{01}$ has higher expressive power than $S5_{\mathcal{ALC}}$ in that it allows for a distinction between logically impossible and ‘infinitely improbable’ events. Therefore, the following is proved from first principles in the appendix, making use of quasimodels similar to those used in (Artale, Lutz, and Toman 2007) for $S5_{\mathcal{ALC}}$.

Theorem 12. *Consistency in Prob- $\mathcal{ALC}^{01}$ is 2-EXPTIME-complete.*

Prob- $\mathcal{EL}$ Lower Bounds

The $\mathcal{EL}$ family of DLs is a popular family of lightweight ontology languages. Introduced in (Baader, Brandt, and Lutz 2005), the design of the $\mathcal{EL}$ family of DLs aims at tractability of standard reasoning problems such as consistency while still providing sufficient expressive power for formulating ontologies. In fact, members of the $\mathcal{EL}$ family are used in important and well-known ontologies such as SNOMED CT and underlie the OWL 2 EL profile of the OWL 2 ontology language. We consider probabilistic variants of the basic member $\mathcal{EL}$ of the $\mathcal{EL}$ -family.

In (non-probabilistic) $\mathcal{EL}$, concepts are formed according to the syntax rule

$$C ::= \top \mid A \mid C \sqcap D \mid \exists r.C$$

where A ranges over N_C , C and D over concepts and r over N_R . $\mathcal{EL}$ ABoxes are simply conjunctions of assertions $C(a)$ and $r(a, b)$. Since $\mathcal{EL}$ and the probabilistic extensions considered in what follows do not contain negation, deciding consistency of knowledge bases is uninteresting: *every* KB is consistent. Instead, we consider *instance checking* which is the problem to decide, given a KB $\mathcal{K}$, an individual name a , and a concept C , whether $\mathcal{I}, w \models C(a)$ for all models $\mathcal{I}$ of $\mathcal{K}$ and all $w \in \Delta^{\mathcal{I}}$ (written $\mathcal{K} \models C(a)$). Note that, in probabilistic extensions of $\mathcal{EL}$, C can be of the form $P_{>n}C'$, and thus instance checking can be used to find out whether

the probability of a concept assertion $C(a)$ exceeds a given threshold.

Unfortunately, it turns out that naive extensions of $\mathcal{EL}$ with probabilistic constructors are typically intractable. This is due to the fact that, as observed in (Baader, Brandt, and Lutz 2005) and exploited in many hardness proofs since, every *non-convex* extension of basic $\mathcal{EL}$ is typically as hard as the corresponding variant of $\mathcal{ALC}$. Here, a logic is convex if for all KBs $\mathcal{K}$ and concepts $D_1, \dots, D_n$ such that $\mathcal{K} \models D_1 \sqcup \dots \sqcup D_n(a)$, we have $\mathcal{K} \models D_i(a)$ for some i . Already the extension of $\mathcal{EL}$ with the concept constructor $P_{>n}$ turns out to be non-convex. To see this, choose $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ with $\mathcal{T} = \emptyset$, $\mathcal{A} = \{C(a)\}$, and

$$\begin{aligned} C &= P_{>0.4}A_1 \sqcap P_{>0.4}A_2 \sqcap P_{>0.4}A_3 \\ D_1 &= P_{>0}(A_1 \sqcap A_2) \\ D_2 &= P_{>0}(A_1 \sqcap A_3) \\ D_3 &= P_{>0}(A_2 \sqcap A_3) \end{aligned}$$

In a similar way, non-convexity can be shown for the extension of $\mathcal{EL}$ with any of the constructors $P_{\geq n}$, $\exists P_{>n}r$, $\top$, and $\exists P_{\geq n}r$. Since the latter two cases may be a bit less obvious, we explicitly give a counterexample for convexity: for $\mathcal{EL}$ extended with $\exists P_{>0}r$, choose $\mathcal{A} = \{C(a)\}$ and

$$\begin{aligned} \mathcal{T} &= \{ \exists r_1.\top \sqcap \exists r_2.\top \sqsubseteq \exists P_{>0}s_1.\top \\ &\quad \exists r_1.\top \sqcap \exists r_3.\top \sqsubseteq \exists P_{>0}s_1.\top \\ &\quad \exists r_2.\top \sqcap \exists r_3.\top \sqsubseteq \exists P_{>0}s_2.\top \} \\ C &= \exists P_{>0.4}r_1.\top \sqcap \exists P_{>0.4}r_2.\top \sqcap \exists P_{>0.4}r_3.\top \\ D_1 &= \exists P_{>0}s_1.\top \\ D_2 &= \exists P_{>0}s_2.\top \end{aligned}$$

In each case, we can employ a standard form of reduction to show EXPTIME-hardness, see (Baader, Brandt, and Lutz 2005). The upper bound stems from Theorem 3.

Theorem 13. *In $\mathcal{EL}$ extended with any of $P_{>n}C$, $P_{\geq n}C$, $\exists P_{>n}r$, $\top$, and $\exists P_{\geq n}r$, instance checking is EXPTIME-hard. In the former two cases, it is EXPTIME-complete.*

In the latter two cases, we do not even know decidability. However, non-convexity implies that these logics can only be decidable if the corresponding version of Prob- $\mathcal{ALC}$ is.

The counterexamples used above suggest that, in order to attain convexity, we have to restrict ourselves to the probability values 0 and 1. We will see later that this indeed suffices to achieve tractability when probabilistic roles are disallowed entirely. Here we show that, despite guaranteeing convexity, the mentioned restriction is still not sufficient for tractability if probabilistic roles are present. The proof is by a reduction of the word problem of deterministic polynomially space-bounded Turing machines.

Theorem 14. *Instance checking in $\mathcal{EL}$ extended with $P_{>0}C$ and at least one of $\exists P_{=1}r.C$ and $\exists P_{>0}r.C$ is PSPACE-hard.*

Proof (sketch). We concentrate on Prob- $\mathcal{EL}_{r=1}^{01}$. The proof is by reduction of the word problem of deterministic, polynomially space-bounded Turing machines. Let $M = (Q, \Sigma, \Gamma, \delta, q_0, q_{acc}, q_{rej})$ be such a machine, $x \in \Sigma$

an input of length n , and $m = p(n)$ the space bound of M on x . Our aim is to construct in polynomial time a TBox $\mathcal{T}$ and concept C_0 such that $\mathcal{K} = (\mathcal{T}, \emptyset) \models C_0(a)$ iff M accepts x . The basic idea is that each model $\mathcal{I}$ of $\mathcal{T}$ will take the form of an infinite r -chain of probability 1, i.e., there are $d_0, d_1, \dots \in \Delta^{\mathcal{I}}$ such that $p_{d_i, d_{i+1}}^{\mathcal{I}}(r) = 1$ for all $i \geq 0$. For every d_i , there will be a world w such that the concept memberships of d represent the initial configuration of M on x . When going *backwards* in the chain but staying in the world w , the concept memberships evolve according to the computation of M on x . Since this holds for all d_i (i.e., at arbitrary distance from d_0), it follows that for each configuration c that is encountered during the computation, there is a world w where the concept memberships of d_0 represent c . It is then easy to use C_0 to check whether any of these configurations is accepting. $\square$

The best known upper bound is the 2-EXPTIME one from Theorem 12.

Prob- $\mathcal{EL}$ Upper Bound

We consider the probabilistic DL Prob- $\mathcal{EL}_c^{01}$ that extends $\mathcal{EL}$ with the probability constructors $P_{>0}C$ and $P_{=1}C$. Our main result is that instance checking can be decided in PTIME and is thus no more difficult than in $\mathcal{EL}$. In Prob- $\mathcal{EL}_c^{01}$, an ABox is a set of assertions $C(a)$, $r(a, b)$, $P_{>0}r(a, b)$, and $P_{=1}r(a, b)$. Thus, we do allow probabilistic roles, but only in the ABox.

To prove the PTIME upper bound, we may w.l.o.g. restrict our attention to instance checking problems of the form $\mathcal{K} = (\mathcal{T}, \mathcal{A}) \models A(a)$, where A is a concept *name* that occurs in $\mathcal{T}$. We assume that Prob- $\mathcal{EL}_c^{01}$ -TBoxes are in the following *normal form*. A *basic concept* is either $\top$, a concept name, or a concept of the form $P_{*n}A$ with A a concept name. For a TBox to be in normal form, we require that every concept inclusion is of one of the forms

$$X_1 \sqcap \dots \sqcap X_n \sqsubseteq X, \quad \exists r.X \sqsubseteq A, \quad X \sqsubseteq \exists r.A$$

where the X_i and X denote basic concepts and A denotes a concept names. It is standard to show that, by introducing fresh concept names, every TBox $\mathcal{T}$ can be converted in polynomial time into a TBox $\mathcal{T}'$ in normal form such that $(\mathcal{T}, \mathcal{A}) \models A(a)$ iff $(\mathcal{T}', \mathcal{A}) \models A(a)$ for all ABoxes $\mathcal{A}$, $a \in \text{Ind}(\mathcal{A})$, and all concept names A that occur in $\mathcal{T}$ (Baader, Brandt, and Lutz 2005). We also assume that ABoxes are in *normal form*, i.e. in all assertions $C(a)$, C is a concept *name*. As for TBoxes, this can be achieved by introducing fresh concept names.

Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a Prob- $\mathcal{EL}_c^{01}$ -KB with $\mathcal{T}$ and $\mathcal{A}$ in normal form, and let $A_0(a_0)$ be an assertion for which we want to decide whether $\mathcal{K} \models A_0(a_0)$. We use $N_C^{\mathcal{K}}$ (resp. $N_R^{\mathcal{K}}$) to denote the set of concept names (resp. role names) that occur in $\mathcal{K}$; $C^{\mathcal{K}}$ (resp. $\mathcal{P}_0^{\mathcal{K}}, \mathcal{P}_1^{\mathcal{K}}$) to denote the set of basic concepts (resp. concepts $P_{>0}A$, $P_{=1}A$) that occur (possibly as a subconcept) in $\mathcal{K}$; and $\mathcal{R}_0^{\mathcal{K}}$ to denote the set of assertions $P_{>0}r(a, b)$ in $\mathcal{A}$.

Our algorithm is of the same kind as the one presented in (Baader, Brandt, and Lutz 2005), i.e., it builds a representation of a ‘least’ model of $\mathcal{K}$ (in the sense of Horn logic). The

set of worlds of this model is $V := \{0, \varepsilon, 1\} \cup \mathcal{P}_0^\mathcal{K} \cup \mathcal{R}_0^\mathcal{K}$, where 0 is the only world with probability 0 and all other worlds have uniform probability $1/|V \setminus \{0\}|$. Intuitively, the worlds in $\mathcal{P}_0$ and $\mathcal{R}_0^\mathcal{K}$ serve as witnesses for the eponymous concepts $P_{>0}A$ and assertions $P_{>0}r(a, b)$; the world 1 is used to collect concepts that a domain element has to satisfy with probability 1; and ε serves as a general witness that is used to deal with existential restrictions. We represent the model as a quasimodel, defined as follows.

An *ABox quasistate* $\hat{Q}$ for $\mathcal{K}$ is a mapping that associates with each $a \in \text{Ind}(\mathcal{A})$ and each $v \in V$ a subset $\hat{Q}(a, v) \subseteq \mathcal{C}^\mathcal{K}$. An *element quasistate* Q for $\mathcal{K}$ is a mapping that associates to each $v \in V$ a subset $Q(v) \subseteq \mathcal{C}^\mathcal{K}$. To treat ABox quasistates and element quasistates uniformly, we fix an individual name a_ε and for each element quasistate Q and $v \in V$, use $Q(a_\varepsilon, v)$ to denote $Q(v)$ whenever convenient. A *quasimodel* $\mathfrak{M}$ for $\mathcal{K}$ is a pair $(\hat{Q}, f)$ with $\hat{Q}$ an ABox quasistate for $\mathcal{K}$ and f a mapping that assigns to each $(A, i) \in \Omega := \mathcal{N}_\mathcal{C}^\mathcal{K} \times \{0, \varepsilon\}$ an element quasistate $f_{A,i}$ for $\mathcal{K}$. Intuitively, the quasistate $f_{A,i}$ describes a domain element that serves as a witness for existential restrictions $\exists r.A$ satisfied by some domain element in world 0 if $i = 0$ and in a world from $V \setminus \{0\}$ if $i = \varepsilon$.

The algorithm starts with the quasimodel $(\hat{Q}, f)$, where

- $\hat{Q}(a, 0) = \{\top\} \cup \{A \mid A(a) \in \mathcal{A}\}$ for all $a \in \text{Ind}(\mathcal{A})$;
- $\hat{Q}(a, v) = \{\top\}$ for all $a \in \text{Ind}(\mathcal{A})$ and $v \in V \setminus \{0\}$;
- for every $(A, 0) \in \Omega$, $f_{A,0}$ is defined by setting $f_{A,0}(0) = \{\top, A\}$ and $f_{A,0}(v) = \{\top\}$ for all $v \in V \setminus \{0\}$;
- for every $(A, \varepsilon) \in \Omega$, $f_{A,\varepsilon}$ is defined by setting $f_{A,\varepsilon}(\varepsilon) = \{\top, A\}$ and $f_{A,\varepsilon}(v) = \{\top\}$ for all $v \in V \setminus \{\varepsilon\}$.

This quasimodel is then extended by applying the completion rules shown in Figure 1 until no more rules apply. In the figure, Q ranges over $\hat{Q}$ and over the quasistates $(f_{A,i})_{(A,i) \in \Omega}$ (relying upon our notation of $f_{A,i}(a_\varepsilon v)$ as a synonym for $f_{A,i}(v)$), and v ranges over V unless otherwise specified. The function $\gamma : V \rightarrow \{0, \varepsilon\}$ is defined by setting $\gamma(0) = 0$ and $\gamma(v) = \varepsilon$ for all $v \in V \setminus \{0\}$. Note that in rule **R6**, we use $f_{A,\varepsilon}(\varepsilon)$ as a witness for the existential restriction $\exists r.A$ satisfied by a in v , for all $v \in V \setminus \{0\}$ instead of only for $v = \varepsilon$. This seemingly incorrect approach is actually necessary to ensure correctness, and has to be compensated by a careful model construction.

Lemma 15. *For all $A_0 \in \mathcal{N}_\mathcal{C}$ and $a \in \text{Ind}(\mathcal{A})$, $\mathcal{K} \models A_0(a_0)$ iff $A_0 \in \hat{Q}(a_0, 0)$.*

Our proof also establishes a combined UMP and BMP (c.f. Corollary 5), with exact bounds $\mathcal{O}(|\mathcal{K}|)$ on the number of worlds and $\mathcal{O}(|\mathcal{K}|^2)$ on the number of domain elements.

Theorem 16. *Instance checking in $\text{Prob-}\mathcal{EL}_\mathcal{C}^{\text{LOI}}$ can be decided in PTIME.*

Related Work

There is a large number of proposals for probabilistic DLs that differ widely in many fundamental aspects. To start

R1	If $X_1, \dots, X_n \in Q(a, v)$, $X_1 \sqcap \dots \sqcap X_n \sqsubseteq X \in \mathcal{T}$, and $X \notin Q(a, v)$ then $Q(a, v) := Q(a, v) \cup \{X\}$
R2	If $P_{>0}A \in Q(a, v)$ and $A \notin Q(a, P_{>0}A)$ then $Q(a, P_{>0}A) := Q(a, P_{>0}A) \cup \{A\}$
R3	If $P_{=1}A \in Q(a, v)$ and $A \notin Q(a, v)$ with $v \neq 0$ then $Q(a, v) := Q(a, v) \cup \{A\}$
R4	If $A \in Q(a, v)$ with $v \neq 0$, $P_{>0}A \in \mathcal{P}_0^\mathcal{K}$, and $P_{>0}A \notin Q(a, v')$ then $Q(a, v') := Q(a, v') \cup \{P_{>0}A\}$
R5	If $A \in Q(a, 1)$, $P_{=1}A \in \mathcal{P}_1^\mathcal{K}$, and $P_{=1}A \notin Q(a, v)$ then $Q(a, v) := Q(a, v) \cup \{P_{=1}A\}$
R6	If $X \in Q(a, v)$, $X \sqsubseteq \exists r.A \in \mathcal{T}$, $Y \in f_{A,\gamma(v)}(\gamma(v))$, $\exists r.Y \sqsubseteq B \in \mathcal{T}$, and $B \notin Q(a, v)$ then $Q(a, v) := Q(a, v) \cup \{B\}$
R7	If $r(a, b) \in \mathcal{A}$, $X \in \hat{Q}(b, 0)$, $\exists r.X \sqsubseteq A \in \mathcal{T}$, and $A \notin \hat{Q}(a, 0)$ then $Q(a, 0) := Q(a, 0) \cup \{A\}$
R8	If $P_{>0}r(a, b) \in \mathcal{A}$, $X \in \hat{Q}(b, P_{>0}r(a, b))$, $\exists r.X \sqsubseteq A \in \mathcal{T}$, and $A \notin \hat{Q}(a, P_{>0}r(a, b))$ then $\hat{Q}(a, P_{>0}r(a, b)) := \hat{Q}(a, P_{>0}r(a, b)) \cup \{A\}$
R9	If $P_{=1}r(a, b) \in \mathcal{A}$, $X \in \hat{Q}(b, v)$ with $v \neq 0$, $\exists r.X \sqsubseteq A \in \mathcal{T}$, and $A \notin \hat{Q}(a, v)$ then $\hat{Q}(a, v) := \hat{Q}(a, v) \cup \{A\}$

Figure 1: Completion rules

with, several proposals assume that *concrete* probability distributions are specified, typically in terms of a Bayes net (Koller, Levy, and Pfeffer 1997; Yelland 2000; Ding, Peng, and Pan 2006; da Costa and Laskey 2006). These proposals often address statistical probabilities rather than subjective ones. In contrast, we address subjective probabilities and our semantics is ‘free’ in the sense that no concrete distributions are fixed. This approach was pioneered by Fagin, Halpern, and Megiddo (Fagin, Halpern, and Megiddo 1990) and has first been integrated with DLs by Heinsohn (Heinsohn 1994) and Jaeger (Jaeger 1994). The main feature of the latter two proposals are conditional probabilities in the TBox, expressing statistical probabilities. The semantics is then loosely related to, yet different from Halpern’s Type 1 probabilistic FO. Jaeger additionally admits probabilistic ABox statements, which address subjective probabilities.

A more recent approach to probabilistic DLs is due to Lukasiewicz (2008), who proposes a probabilistic extension of expressive DLs such as *SHIQ*. The resulting DL *P-SHIQ* encompasses both statistical and subjective features, and also addresses the non-monotonic aspects of probabilistic knowledge using a semantics based on Lehmann’s lexicographic entailment. The differences to our framework are manifold. First, *Prob-ALC* does not include non-monotonic features, which on the one-hand precludes features such as subjective conditional probabilities, but on the other hand enables a transparent semantics. In contrast, the semantics of *P-SHIQ* is often felt to be rather difficult to understand precisely because of its non-monotonic features and use of lexicographic entailment (Klinov, Parsia, and Sattler 2009). Second, *P-SHIQ* allows probabilistic TBox state-

ments, which is not foreseen in the current version of Prob- $\mathcal{ALC}$ (but could be added). Conversely, Prob- $\mathcal{ALC}$ allows probabilistic concepts and roles, which are not present in P- $\mathcal{SHIQ}$. And third, it is discussed in (Klinov, Parsia, and Sattler 2009) that P- $\mathcal{SHIQ}$ restricts ABoxes in a number of ways, e.g. separating ‘classical’ and ‘probabilistic individuals’ and disallowing probabilistic role assertions. None of these restrictions is required in Prob- $\mathcal{ALC}$.

Conclusion

We have proposed a new family of probabilistic DLs that are derived in a straightforward way from Halpern’s Type 2 probabilistic FOL. We have also provided a first but substantial analysis of the complexity of reasoning in these logics. We conjecture that the most well-behaved variants Prob- $\mathcal{ALC}_c$ and Prob- $\mathcal{EL}_c^{01}$ can be extended with many standard DL constructors without losing their good computational properties. In addition to Type 2 logics for reasoning about subjective probabilities, Halpern also suggests Type 1 logics for reasoning about statistical probabilities. Alas, probabilistic DLs derived from Type 1 logic in the spirit of the present work are rather inexpressive and not overly interesting; details can be found in the appendix. An issue for future research is to investigate propositional TBoxes statements of the form $P_{*n}(C \sqsubseteq D)$. In contrast to the probabilistic constructors in this paper, which express uncertainty about concrete facts, such TBoxes can be used to express uncertainty about general knowledge.

References

- Artale, A.; Lutz, C.; and Toman, D. 2007. A description logic of change. In *International Joint Conference on Artificial Intelligence, IJCAI-07*, 218–223.
- Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook*. Cambridge University Press.
- Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the $\mathcal{EL}$ envelope. In *International Joint Conference on Artificial Intelligence, IJCAI-05*, 364–369.
- Bodenreider, O.; Smith, B.; and Burgun, A. 2004. The ontology-epistemology divide: A case study in medical terminology. In *Formal Ontology in Information Systems, FOIS 04*, 185–195. IOS Press.
- Brázdil, T.; Forejt, V.; Kretínský, J.; and Kucera, A. 2008. The satisfiability problem for probabilistic ctl. In *LICS*, 391–402. IEEE Computer Society.
- Canny, J. F. 1988. Some algebraic and geometric computations in PSPACE. In *Symposium on Theory of Computing, STOC 1988*, 460–467. ACM.
- da Costa, P. C. G., and Laskey, K. B. 2006. PR-OWL: A framework for probabilistic ontologies. In *Formal Ontology in Information Systems, FOIS 06*, volume 150 of *Frontiers in Artificial Intelligence and Applications*, 237–249. IOS Press.
- Ding, Z.; Peng, Y.; and Pan, R. 2006. BayesOWL: Uncertainty modeling in semantic web ontologies. In *Soft Computing in Ontologies and Semantic Web*, number 204 in *Studies in Fuzziness and Soft Computing*, 3–29. Springer.
- Fagin, R.; Halpern, J.; and Megiddo, N. 1990. A logic for reasoning about probabilities. *Inform. Comput.* 87:78–128.
- Gabbay, D. M.; Kurucz, Á.; Wolter, F.; and Zakharyashev, M. 2003. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier.
- Halpern, J. Y. 1990. An analysis of first-order logics of probability. *Artif. Intell.* 46:311–350.
- Heinsohn, J. 1994. Probabilistic description logics. In *Uncertainty in Artificial Intelligence, UAI 94*, 311–318. Morgan Kaufmann.
- Jaeger, M. 1994. Probabilistic reasoning in terminological logics. In *Principles of Knowledge Representation and Reasoning, KR 94*. Morgan Kaufmann. 305–316.
- Jaeger, M. 2006. Probabilistic role models and the guarded fragment. *Int. J. Uncertainty, Fuzziness and Knowledge-Based Systems* 14:43–60.
- Klinov, P.; Parsia, B.; and Sattler, U. 2009. On correspondences between probabilistic first-order and description logics. In *Description Logics*, volume 477 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Koller, D.; Levy, A.; and Pfeffer, A. 1997. P-CLASSIC: A tractable probabilistic description logic. In *AAAI 97*, 390–397. AAAI Press.
- Lukasiewicz, T., and Straccia, U. 2008. Managing uncertainty and vagueness in description logics for the semantic web. *J. Web Sem.* 6(4):291–308.
- Lukasiewicz, T. 2008. Expressive probabilistic description logics. *Artif. Intell.* 172:852–883.
- Lutz, C.; Wolter, F.; and Zakharyashev, M. 2008. Temporal description logics: A survey. In *International Symposium on Temporal Representation and Reasoning, TIME 2008*, 3–14. IEEE Computer Society.
- Price, C., and Spackman, K. 2000. SNOMED Clinical Terms. *Br. J. Healthcare Comput. Inf. Mgmt.* 17:27–31.
- Schrijver, A. 1986. *Theory of linear and integer programming*. Wiley Interscience.
- Schröder, L., and Pattinson, D. 2007. Modular algorithms for heterogeneous modal logics. In *Automata, Languages and Programming, ICALP 07*, volume 4596 of *LNCS*, 459–471. Springer.
- Sebastiani, F. 1994. A probabilistic terminological logic for modelling information retrieval. In *ACM-SIGIR Conference on Research and Development in Information Retrieval, SIGIR 1994*, 122–130. Springer.
- Wolter, F. 1998. Fusions of modal logics revisited. In *Advances in modal logic*, volume 1 of *CSLI Lect. Notes*. CSLI. 361–379.
- Yelland, P. M. 2000. An alternative combination of Bayesian networks and description logics. In *Principles of Knowledge Representation and Reasoning, KR 00*, 225–233. Morgan Kaufmann.