



A Bayesian Extension of the Description Logic \mathcal{ALC}

Leonard Botha¹ , Thomas Meyer¹ , and Rafael Peñaloza²  

¹ University of Cape Town and CAIR, Cape Town, South Africa
leonardzbotha@gmail.com, tmeyer@cs.uct.ac.za

² University of Milano-Bicocca, Milano, Italy
rafael.penaloz@unibz.it

Abstract. Description logics (DLs) are well-known knowledge representation formalisms focused on the representation of terminological knowledge. A probabilistic extension of a light-weight DL was recently proposed for dealing with certain knowledge occurring in uncertain contexts. In this paper, we continue that line of research by introducing the Bayesian extension \mathcal{BALC} of the DL \mathcal{ALC} . We present a tableau-based procedure for deciding consistency, and adapt it to solve other probabilistic, contextual, and general inferences in this logic. We also show that all these problems remain EXPTIME-complete, the same as reasoning in the underlying classical \mathcal{ALC} .

1 Introduction

Description logics (DLs) [1] are a family of logic-based knowledge representation formalisms designed to describe the terminological knowledge of an application domain. Due to their clear syntax, formal semantics, and the existence of efficient reasoners alongside their expressivity, they have been successfully applied to model several domains, especially from the biomedical sciences. However, in their classical form, these logics are not capable of dealing with uncertainty, which is an unavoidable staple in real-world knowledge. To overcome this limitation, several probabilistic extensions of DLs have been suggested in the literature. The landscape of probabilistic extensions of DLs is too large to be covered in detail in this work. These logics differentiate themselves according to their underlying logical formalism, their interpretation of probabilities, and the kind of uncertainty that they are able to express. For a relevant, although slightly outdated survey, where all these differences are showcased, see [16].

A recently proposed probabilistic DL is the Bayesian extension \mathcal{BEL} of the light-weight \mathcal{EL} . This logic focuses on modelling certain knowledge that holds only in some contexts, together with uncertainty about the current context. One advantage of the formalism underlying \mathcal{BEL} is that it separates the contextual

R. Peñaloza—This work was carried out while the third author was at the Free University of Bozen-Bolzano, Italy.

knowledge, which is *de facto* a classical ontology, from the likelihood of observing this context. As a simple example of the importance of contextual knowledge, consider the knowledge of construction techniques and materials that vary through time. In the context of a modern house, asbestos and modern pipes are not observable, while some classes of houses built during the 1970s do contain both. However, in all contexts we know that asbestos and lead in drinking water have grave health effects. Still, when confronted with a random house, one might not know to which of these contexts it belongs, and by extension whether it is safe to live in. But construction data may be used to derive the probabilities of these contexts.

To allow for complex probabilistic relationships between the contexts, their joint probability distribution is encoded via a Bayesian network (BN) [18]. This logic is closely related to the probabilistic extension of DL-Lite proposed in [12], but uses a less restrictive semantics (for a discussion on the differences between these logics, see [11]). Another similar proposal is Probabilistic Datalog[±] [13], with the difference that uncertainty is represented via a Markov Logic Network, instead of a BN. Since the introduction of \mathcal{BEL} , the main notions behind it have been generalised to arbitrary ontology languages [8]. However, it has also been shown that efficient and complexity-optimal reasoning methods can only be achieved by studying the properties of each underlying ontology language [11].

In this paper, we continue with that line of research and study the Bayesian extension of the propositionally closed DL \mathcal{ALC} . As our main result, we present an algorithm, based on a glass-box modification of the classical tableaux method for reasoning in \mathcal{ALC} , that outputs a description of all the contexts encoding inconsistent knowledge. Using this algorithm, we describe an effective method for deciding consistency of a \mathcal{BALC} knowledge base. We also provide a tight EXPTIME complexity bound for this problem.

This is followed by a study of several crisp and probabilistic variants of the standard DL decision problems; namely, concept satisfiability, subsumption, and instance checking. Interestingly, our work shows that all our problems can be reduced to some basic computations over a context describing inconsistency, and hence are EXPTIME-complete as well. These complexity bounds are not completely surprising, given the high complexity of the classical \mathcal{ALC} . However, our tableaux-based algorithm has the potential to behave better in practical scenarios. This work details and deepens results that have previously been presented in [5, 6].

2 Preliminaries

We start by briefly introducing Bayesian networks and the description logic (DL) \mathcal{ALC} , which form the basis for \mathcal{BALC} .

Bayesian networks (BNs) are graphical models capable of representing the joint probability distribution (JPD) of several discrete random variables in a compact manner [18]. Given a random variable X , $\text{val}(X)$ denotes the set of values that




Fig. 1. A Bayesian network with three Boolean variables.

X can take. For $x \in \text{val}(X)$, $X = x$ is the *valuation* of X taking the value x . This notation is extended to sets of variables in the obvious way. Given a set of random variables V , a *world* ω is a set of valuations containing exactly one valuation for every random variable $X \in V$. A V -*literal* is an ordered pair of the form (X_i, x) , where $X_i \in V$ and $x \in \text{val}(X_i)$. V -literals generalise Boolean literals denoted as x or $\neg x$ for the random variable X . For simplicity, in this paper we will often use the notation X for (X, T) and $\neg X$ for (X, F) . A V -*context* is any set of V -literals. It is *consistent* if it contains at most one literal for each random variable. We will often call V -contexts *primitive contexts*.

A *Bayesian network* is a pair $\mathcal{B} = (G, \Theta)$ where $G = (V, E)$ is a directed acyclic graph (DAG) and Θ is a set of conditional probability distributions for every variable $X \in V$ given its parents $\pi(X)$ on the DAG G ; more precisely, $\Theta = \{P(X = x | \pi(X) = \mathbf{x}') \mid X \in V\}$. \mathcal{B} encodes the JPD of V through the chain rule $P(\mathbf{X} = \mathbf{x}) = \prod_{X_i \in V} P(X_i = x_i \mid \pi(X_i) = \mathbf{x}_j)$.

Figure 1 depicts a BN with three random variables denoting the likelihood of different characteristics of a construction: X stands for a post-1986 building, Y for a renovated building, and Z for the presence of lead pipes. In this case, $P(X, Y, Z) = 0.7 \cdot 0.1 \cdot 0 = 0$; i.e., a renovated post-1986 house has no lead pipes. The *description logic* \mathcal{ALC} is the smallest propositionally closed DL [1, 20]. It is based on *concepts* which correspond to unary predicates of first-order logic, and *roles* corresponding to binary predicates. Formally, given the mutually disjoint sets N_I, N_C , and N_R of *individual*, *concept*, and *role names*, respectively, \mathcal{ALC} concepts are built by the grammar rule $C ::= A \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \exists r.C \mid \forall r.C$, where $A \in N_C$ and $r \in N_R$. \mathcal{ALC} axioms are either *general concept inclusions* (GCIs) of the form $C \sqsubseteq D$, *concept assertions* $C(a)$, or *role assertions* $r(a, b)$ where $a, b \in N_I, r \in N_R$, and C, D are concepts. An *ontology* is a finite set of axioms. We sometimes partition an ontology into the *TBox* \mathcal{T} composed exclusively of GCIs, and the *ABox* \mathcal{A} containing all concept and role assertions.

The semantics of \mathcal{ALC} is defined by *interpretations*, which are pairs of the form $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where $\Delta^{\mathcal{I}}$ is a non-empty set called the *domain* and $\cdot^{\mathcal{I}}$ is the *interpretation function* that maps every $a \in N_I$ to an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, every $A \in N_C$ to a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and every $r \in N_R$ to a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The interpretation function is extended to arbitrary concepts by defining for any two concepts C, D :

- $(\neg C)^{\mathcal{I}} := \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$,
- $(C \sqcap D)^{\mathcal{I}} := C^{\mathcal{I}} \cap D^{\mathcal{I}}$,
- $(C \sqcup D)^{\mathcal{I}} := C^{\mathcal{I}} \cup D^{\mathcal{I}}$,
- $(\exists r.C)^{\mathcal{I}} := \{\delta \in \Delta^{\mathcal{I}} \mid \exists \eta \in C^{\mathcal{I}}. (\delta, \eta) \in r^{\mathcal{I}}\}$, and
- $(\forall r.C)^{\mathcal{I}} := \{\delta \in \Delta^{\mathcal{I}} \mid \forall \eta \in \Delta^{\mathcal{I}}. (\delta, \eta) \in r^{\mathcal{I}} \Rightarrow \eta \in C^{\mathcal{I}}\}$.

This interpretation *satisfies* the GCI $C \sqsubseteq D$ iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, the concept assertion $C(a)$ iff $a^{\mathcal{I}} \in C^{\mathcal{I}}$ and the role assertion $r(a, b)$ iff $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$. \mathcal{I} is a *model* of the ontology \mathcal{O} iff it satisfies all axioms in \mathcal{O} . An important abbreviation in \mathcal{ALC} is the bottom concept $\perp := A \sqcap \neg A$, where A is any concept name. For any interpretation \mathcal{I} , $\perp^{\mathcal{I}} = \emptyset$.

As a simple example, one can express the notion that water pipes do not contain lead through the GCI $\text{Pipe} \sqsubseteq \forall \text{contains}. \neg \text{Lead}$.

3 \mathcal{BALC}

\mathcal{BALC} is a probabilistic extension of \mathcal{ALC} , in which axioms are considered to hold only in a given (possibly uncertain) context expressed through annotations.

Definition 1 (KB). *Let V be a finite set of discrete random variables. A V -restricted axiom (V -axiom) is an expression of the form α^κ , where α is an \mathcal{ALC} axiom and κ is a V -context. A V -ontology is a finite set of V -axioms. A \mathcal{BALC} knowledge base (KB) over V is a pair $\mathcal{K} = (\mathcal{O}, \mathcal{B})$ where \mathcal{B} is a BN over V , and \mathcal{O} is a V -ontology.*

To define the semantics of \mathcal{BALC} , we extend the notion of an interpretation to also take contexts into account, and interpret the probabilities based on a multiple-world approach. Formally, a V -interpretation is a tuple $\mathcal{V} = (\Delta^{\mathcal{V}}, \cdot^{\mathcal{V}}, v^{\mathcal{V}})$ where $(\Delta^{\mathcal{V}}, \cdot^{\mathcal{V}})$ is an \mathcal{ALC} interpretation and $v^{\mathcal{V}} : V \rightarrow \cup_{X \in V} \text{val}(X)$ is a *valuation function* such that $v^{\mathcal{V}}(X) \in \text{val}(X)$.

Given a valuation function $v^{\mathcal{V}}$, a Bayesian world ω , and a context κ we denote $v^{\mathcal{V}} = \omega$ when $v^{\mathcal{V}}$ assigns to each random variable the same value as it has in ω ; $v^{\mathcal{V}} \models \kappa$ when $v^{\mathcal{V}}(X) = x$ for all $(X, x) \in \kappa$; and $\omega \models \kappa$ when there is $\omega = v^{\mathcal{V}}$ such that $v^{\mathcal{V}} \models \kappa$.

Definition 2 (Model). *The V -interpretation \mathcal{V} is a model of the V -axiom α^κ , ($\mathcal{V} \models \alpha^\kappa$), iff (i) $v^{\mathcal{V}} \not\models \kappa$, or (ii) $(\Delta^{\mathcal{V}}, \cdot^{\mathcal{V}})$ satisfies α . It is a model of the ontology \mathcal{O} iff it is a model of all axioms in \mathcal{O} .*

Notice that \mathcal{BALC} is a generalisation of \mathcal{ALC} . The axiom α^\emptyset holds in all contexts, and hence every \mathcal{ALC} ontology is also a V -ontology from \mathcal{BALC} . In particular, this means that reasoning in \mathcal{BALC} should be at least as hard as doing so in \mathcal{ALC} . For brevity, for the rest of this paper we will abbreviate axioms of the form α^\emptyset simply as α . When it is clear from the context, we will omit the V prefix and refer only to e.g., contexts, GCIs, or ontologies.

V -interpretations focus on only a single world, but KBs have information about the uncertainty of being in one world or another. Probabilistic interpretations combine multiple V -interpretations and the probability distribution from the BN to give information about the uncertainty of some consequences.

Definition 3 (Probabilistic model). A probabilistic interpretation is a pair of the form $\mathcal{P} = (\mathcal{J}, \mathcal{P}_{\mathcal{J}})$, where \mathcal{J} is a finite set of V -interpretations and $\mathcal{P}_{\mathcal{J}}$ is a probability distribution over \mathcal{J} such that $\mathcal{P}_{\mathcal{J}}(\mathcal{V}) > 0$ for all $\mathcal{V} \in \mathcal{J}$. The probabilistic interpretation \mathcal{P} is a model of the axiom α^{κ} ($\mathcal{P} \models \alpha^{\kappa}$) iff every $\mathcal{V} \in \mathcal{J}$ is a model of α^{κ} . \mathcal{P} is a model of the ontology \mathcal{O} iff every $\mathcal{V} \in \mathcal{J}$ is a model of \mathcal{O} .

The distribution $\mathcal{P}_{\mathcal{J}}$ is consistent with the BN \mathcal{B} if for every possible world ω of the variables in V it holds that

$$\sum_{\mathcal{V} \in \mathcal{J}, v^{\mathcal{V}} = \omega} P_{\mathcal{J}}(\mathcal{V}) = P_{\mathcal{B}}(\omega),$$

where $P_{\mathcal{B}}$ is the joint probability distribution defined by the BN \mathcal{B} . The probabilistic interpretation \mathcal{P} is a model of the KB $\mathcal{K} = (\mathcal{O}, \mathcal{B})$ iff it is a (probabilistic) model of \mathcal{O} , and is consistent with \mathcal{B} .

Consider for example the KB $\mathcal{K} = (\mathcal{O}, \mathcal{B})$ where \mathcal{B} is the BN from Fig. 1, and \mathcal{O} contains the axioms

$$\begin{array}{ll} \text{Pipe} \sqsubseteq \forall \text{contains.} \neg \text{Lead}^X & \text{Pipe} \sqsubseteq \forall \text{contains.} \neg \text{Lead}^Y \\ \text{Pipe} \sqsubseteq \exists \text{contains.} \text{Lead}^Z & \text{Water} \sqcap \exists \text{hasAlkalinity.} \text{Low} \sqsubseteq \neg \text{Drinkable}^Z. \end{array}$$

The axioms in the first row express that pipes in post-1986 (context X) and in renovated buildings (context Y) do not contain lead. The axioms in the second row refer exclusively to the context of lead pipes (Z). In this case, our knowledge is that pipes do contain lead, and that water with low alkalinity is not drinkable, as it absorbs the lead from the pipes it travels on. Notice that the first two axioms contradict the third one. This is not a problem because they are required to hold in different contexts. Indeed, notice that any context that makes Z , and either X or Y true has probability 0, and hence can be ignored in the construction of a model.

A *complex context* ϕ is a finite non-empty set of primitive contexts. Note that primitive contexts can be seen as complex ones; e.g., the primitive context κ corresponds to the complex context $\{\kappa\}$. Given a valuation function $v^{\mathcal{V}}$ and a complex context $\phi = \{\alpha_1, \dots, \alpha_n\}$ we say that $v^{\mathcal{V}}$ satisfies ϕ (written as $v^{\mathcal{V}} \models \phi$) iff $v^{\mathcal{V}}$ satisfies at least one $\alpha_i \in \phi$; in particular, if $v^{\mathcal{V}} \models \kappa$ then $v^{\mathcal{V}} \models \{\kappa\}$. Thus, in the following we assume that all contexts are in complex form unless explicitly stated otherwise. Finally we say that ϕ entails ψ ($\phi \models \psi$) iff for all $v^{\mathcal{V}}$ such that $v^{\mathcal{V}} \models \phi$ it follows that $v^{\mathcal{V}} \models \psi$. Or alternatively $\phi \models \psi$ iff for all Bayesian worlds ω such that $\omega \models \phi$ it follows that $\omega \models \psi$.

Given complex contexts $\phi = \{\alpha_1, \dots, \alpha_n\}$ and $\psi = \{\beta_1, \dots, \beta_m\}$ we define the operations

$$\begin{array}{ll} \phi \vee \psi := \phi \cup \psi, & \text{and} \\ \phi \wedge \psi := \bigcup_{\alpha \in \phi, \beta \in \psi} \{\alpha \cup \beta\} = \{\alpha \cup \beta \mid \alpha \in \phi, \beta \in \psi\}. \end{array}$$

These operations generalise propositional disjunction (\vee) and propositional conjunction (\wedge), where disjunction has the property that either one of the two contexts holds and conjunction requires that both hold. It is easy to see that for all worlds ω and complex contexts ϕ, ψ it holds that (i) $\omega \models \phi \vee \psi$ iff $\omega \models \phi$ or $\omega \models \psi$, and (ii) $\omega \models \phi \wedge \psi$ iff $\omega \models \phi$ and $\omega \models \psi$. Two important special complex contexts are top (\top) and bottom (\perp), which are satisfied by all or no world, respectively. If there are n consistent primitive contexts and κ is an inconsistent context, these are defined as $\top := \{\alpha_1, \dots, \alpha_n\}$ and $\perp := \kappa$.

In the next section, we study the problem of consistency of a \mathcal{BALC} KB, and its relation to other reasoning problems.

4 Consistency

The most basic decision problem one can consider is *consistency*. That is, deciding whether a given \mathcal{BALC} KB \mathcal{K} has a probabilistic model or not. To deal with this problem, it is convenient to consider the classical \mathcal{ALC} ontologies that should hold at each specific world. Formally, given the \mathcal{BALC} KB $\mathcal{K} = (\mathcal{O}, \mathcal{B})$ and the world ω , the *restriction* of \mathcal{O} to ω is

$$\mathcal{O}_\omega := \{\alpha \mid \alpha^\kappa \in \mathcal{O}, \omega \models \kappa\}.$$

Recall that a probabilistic model $\mathcal{P} = (\mathcal{J}, \mathcal{P}_\mathcal{J})$ of \mathcal{K} is a class of classical interpretations associated to worlds $(\Delta^\mathcal{V}, \cdot^\mathcal{V}, \omega)$, such that $(\Delta^\mathcal{V}, \cdot^\mathcal{V})$ is a classical model of \mathcal{O}_ω . Moreover, all the interpretations associated with the world ω must add to the probability $P_\mathcal{B}(\omega)$ specified by \mathcal{B} . Using this insight, we obtain the following result.

Theorem 4. *The \mathcal{BALC} KB $\mathcal{K} = (\mathcal{O}, \mathcal{B})$ is consistent iff for every world ω with $P_\mathcal{B}(\omega) > 0$ \mathcal{O}_ω is consistent.*

Based on this result, we can derive a process for deciding consistency that provides a tight complexity bound for this problem.

Corollary 5. *\mathcal{BALC} KB consistency is EXPTIME-complete.*

Proof. There are exponentially many worlds ω . For each of them, we have to check (classical) consistency of \mathcal{O}_ω (in exponential time) and that $P_\mathcal{B}(\omega) > 0$, which is linear in the size of \mathcal{B} . \square

The algorithm described in the proof of this corollary is optimal in terms of *worst-case* complexity, but it also runs in exponential time in the *best case*. Indeed, it enumerates all the (exponentially many) Bayesian worlds. In practice, it is infeasible to use an algorithm that requires exponential time on every instance. For that reason, we present a new algorithm based on the tableau method originally developed for \mathcal{ALC} . To describe this algorithm, we need to introduce some additional notation.

\sqcap -rule	if $(C_1 \sqcap C_2)(x)^\phi \in \mathcal{A}$, and either $C_1(x)^\phi$ or $C_2(x)^\phi$ is \mathcal{A} -insertable then $\mathcal{A}' := (\mathcal{A} \oplus C_1(x)^\phi) \oplus C_2(x)^\phi$
\sqcup -rule	if $(C_1 \sqcup C_2)(x)^\phi \in \mathcal{A}$, and both $C_1(x)^\phi$ and $C_2(x)^\phi$ are \mathcal{A} -insertable then $\mathcal{A}' := \mathcal{A} \oplus C_1(x)^\phi, \mathcal{A}'' := \mathcal{A} \oplus C_2(x)^\phi$
\exists -rule	if $(\exists R.C)(x)^\phi \in \mathcal{A}$, there is no z such that neither $R(x, z)^\phi$ nor $C(z)^\phi$ is \mathcal{A} -insertable, and x is not blocked then $\mathcal{A}' := (\mathcal{A} \oplus R(x, y)^\phi) \oplus C(y)^\phi$, where y is a new individual name
\forall -rule	if $\{(\forall R.C)(x)^\phi, R(x, y)^\psi\} \subseteq \mathcal{A}$, and $C(y)^{\phi \wedge \psi}$ is \mathcal{A} -insertable then $\mathcal{A}' := \mathcal{A} \oplus C(y)^{\phi \wedge \psi}$
\sqsubseteq -rule	if $(C \sqsubseteq D)^\phi \in \mathcal{O}$, x appears in \mathcal{A} , and $(\neg C \sqcup D)(x)^\phi$ is \mathcal{A} -insertable then $\mathcal{A}' := \mathcal{A} \oplus (\neg C \sqcup D)(x)^\phi$

Fig. 2. Expansion rules for constructing $\phi_{\mathcal{K}}^\perp$

We denote the context that describes all worlds ω such that \mathcal{O}_ω is inconsistent as $\phi_{\mathcal{K}}^\perp$. That is $\omega \models \phi_{\mathcal{K}}^\perp$ iff \mathcal{O}_ω is inconsistent. Moreover, $\phi_{\mathcal{B}}$ is a context such that $\omega \models \phi_{\mathcal{B}}$ iff $P(\omega) = 0$. Theorem 4 states that \mathcal{K} is inconsistent whenever there is a world that models both $\phi_{\mathcal{K}}^\perp$ and $\phi_{\mathcal{B}}$. This is formalized in the following result.

Theorem 6. *The KB \mathcal{K} is inconsistent iff $\phi_{\mathcal{K}}^\perp \wedge \neg\phi_{\mathcal{B}}$ is satisfiable.*

To decide consistency, it then suffices to find a method for deriving the contexts $\phi_{\mathcal{K}}^\perp$ and $\phi_{\mathcal{B}}$. For the former, we present a variant of the glass-box approach for so-called axiom pinpointing [4, 15, 17], originally based on the ideas from [2]. This approach modifies the standard tableaux algorithm for \mathcal{ALC} to keep track of the contexts in which the derived elements in the tableau hold. In a nutshell, whenever a rule application requires the use of an axiom from the ontology, this fact is registered as part of a propositional formula. In our case, we need a context, rather than a propositional formula, to take care of the multiple values that the random variables take.

The algorithm starts with the ABox \mathcal{A} from \mathcal{O} . Recall that all the axioms in \mathcal{A} are labeled with a context. The algorithm then creates a set of ABoxes \mathfrak{A} following the rules from Fig. 2. As a pre-requisite for the execution of the algorithm, we assume that all concepts appearing in the ontology are in *negation normal form* (NNF); that is, only concept names can appear in the scope of a negation operator. This assumption is w.l.o.g. because every concept can be transformed into NNF in linear time by applying the De Morgan laws, the duality of the quantifiers, and eliminating double negations. Each rule application chooses an ABox $\mathcal{A} \in \mathfrak{A}$ and replaces it by one or two new ABoxes that expand \mathcal{A} . We explain the details of these rule applications next.

An assertion α^ϕ is \mathcal{A} -insertable to \mathcal{A} iff for all ψ such that $\alpha^\psi \in \mathcal{A}$, $\phi \not\models \psi$. In the expansion rules \oplus is used as shorthand for $\mathcal{A} \oplus \alpha^\phi := (\mathcal{A} \setminus \{\alpha^\psi\}) \cup \{\alpha^{\phi \vee \psi}\}$ if $\alpha^\psi \in \mathcal{A}$ and $\mathcal{A} \cup \{\alpha^\phi\}$ otherwise. The individual x is an *ancestor* of y if there is a chain of role assertions connecting x to y ; x *blocks* y iff x is an ancestor of y and for every $C(y)^\psi \in \mathcal{A}$, there is a ϕ such that $C(x)^\phi \in \mathcal{A}$ and $\psi \models \phi$; y is *blocked* if there is a node that block it.

The algorithm applies the expansion rules until \mathfrak{A} is *saturated*; i.e., until no rule is applicable to any $\mathcal{A} \in \mathfrak{A}$. \mathcal{A} contains a *clash* if $\{A(x)^\phi, \neg A(x)^\psi\} \subseteq \mathcal{A}$ for some individual x and concept name A . We define the context

$$\phi_{\mathcal{A}} := \bigvee_{A(x)^\phi, \neg A(x)^\psi \in \mathcal{A}} \phi \wedge \psi,$$

which intuitively describes all the clashes that appear in \mathcal{A} . When \mathfrak{A} is saturated, we return the context $\phi_{\mathcal{K}}^\perp = \bigwedge_{\mathcal{A} \in \mathfrak{A}} \phi_{\mathcal{A}}$ expressing the need of having clashes in every ABox \mathcal{A} for inconsistency to follow. It is important to notice that the definition of a clash does not impose any constraints on the contexts ϕ and ψ labelling the assertions $A(x)$ and $\neg A(x)$, respectively. Indeed, $A(x)$ and $\neg A(x)$ could hold in contradictory contexts. In that case, the conjunction appearing in $\phi_{\mathcal{A}}$ would not be affected; i.e., this clash will not provide any new information about inconsistency.

Informally, the formula $\phi_{\mathcal{K}}^\perp$ corresponds to the clash formula (or pinpointing formula) for explaining inconsistency of an \mathcal{ALC} ontology [4, 15]. The main differences are that the variables appearing in a context are not necessarily Boolean, but multi-valued, and that the axioms in \mathcal{O} are not labelled with unique variables, but rather with contexts already. Notice that the expansion rules in Fig. 2 generalise the expansion rules for \mathcal{ALC} , but may require new rule applications to guarantee that all possible derivations of a clash are detected. As observed in [3, 4], one has to be careful with termination of the modified method. However, since all the assertions used are unary and binary, the sufficient termination conditions from [4, 19] are satisfied. Hence we obtain the following result.

Theorem 7. *The modified tableau algorithm terminates, and the context $\phi_{\mathcal{K}}^\perp$ is such that for every world ω , $\omega \models \phi_{\mathcal{K}}^\perp$ iff \mathcal{K}_ω is inconsistent.*

We now turn our attention to the computation of the formula $\phi_{\mathcal{B}}$. Recall that in a BN, the joint probability distribution is the product of the conditional probabilities of each variable given its parents. Hence a world ω can only have probability 0 if it evaluates some variable in $X \in V$ and its parents $\pi(X)$ to values x and \mathbf{x} , respectively, such that $P(X = x \mid \pi(X) = \mathbf{x}) = 0$. Thus, to compute $\phi_{\mathcal{B}}$ it suffices to find out the cells in the conditional probability tables in Θ with value 0.

Theorem 8. *Let $\mathcal{B} = (V, \Theta)$ be a BN, and define*

$$\phi_{\mathcal{B}} := \bigvee_{P(X=x \mid \pi(X)=\mathbf{x})=0} \left((X, x) \wedge \bigwedge_{Y \in \pi(X)} (Y, y) \right).$$

Then for every world ω , $\omega \models \phi_{\mathcal{B}}$ iff $P_{\mathcal{B}}(\omega) = 0$.

Notice that, in general, the context $\phi_{\mathcal{B}}$ can be computed faster than simply enumerating all possible worlds. In particular, if the conditional probability tables in Θ contain no 0-valued cell, then $\phi_{\mathcal{B}} = \perp$; i.e., it is satisfied by no world.

Although consistency is a very important problem to be studied, we are interested also in other reasoning tasks. In particular, we should also take into account the contexts and the probabilities provided by the BN beyond the question of whether they are positive or not. In the next section we study variants of satisfiability and subsumption problems, before turning our attention to instance checking.

5 Satisfiability and Subsumption

In this section, we focus on two problems that depend only on the TBox part of an ontology, and hence assume for the sake of simplicity that the ABox is empty. Thus, we will write a \mathcal{BALC} KB as a pair $(\mathcal{T}, \mathcal{B})$ where \mathcal{T} is a TBox and \mathcal{B} is a BN. We are in general interested in understanding the properties and relationships of concepts.

Given two concepts C, D and a \mathcal{BALC} KB \mathcal{K} , we say that C is *satisfiable* w.r.t. \mathcal{K} iff there exists a probabilistic model $\mathcal{P} = (\mathcal{J}, \mathcal{P}_{\mathcal{J}})$ of \mathcal{K} s.t. $C^{\mathcal{V}} \neq \emptyset$ for all $\mathcal{V} \in \mathcal{J}$. C is *subsumed* by D w.r.t. \mathcal{K} iff for all models $\mathcal{P} = (\mathcal{J}, \mathcal{P}_{\mathcal{J}})$ of \mathcal{K} and all $\mathcal{V} \in \mathcal{J}$ $C^{\mathcal{V}} \subseteq D^{\mathcal{V}}$. It is possible to adapt the well known reductions from the classical case to show that these two problems are EXPTIME-complete.

Theorem 9. *Satisfiability and subsumption w.r.t. \mathcal{BALC} KBs are EXPTIME-complete.*

Proof. Let $\mathcal{K} = (\mathcal{T}, \mathcal{B})$ and C, D be concepts. It is easy to see that C is subsumed by D w.r.t. \mathcal{K} iff $\mathcal{K}' = (\mathcal{T} \cup \{(C \sqcap \neg D)(a)^{\emptyset}\}, \mathcal{B})$ is inconsistent, where a is an arbitrary individual name. Similarly, C is satisfiable iff $\mathcal{K}'' = (\mathcal{T} \cup \{C(a)^{\emptyset}\}, \mathcal{B})$ is consistent. \square

In the following, we study variants of these problems. For a more concise presentation, we will present only the cases for subsumption. Analogous results hold for satisfiability based on the fact that for every \mathcal{ALC} interpretation \mathcal{I} , it holds that $C^{\mathcal{I}} = \emptyset$ iff $C^{\mathcal{I}} \subseteq \perp^{\mathcal{I}}$. First we consider additional information about contexts; afterwards we compute the probability of an entailment, and then the combination of both.

Definition 10 (contextual subsumption). *Let $\mathcal{K} = (\mathcal{T}, \mathcal{B})$ be a \mathcal{BALC} KB, C, D concepts, and κ a context. C is subsumed in context κ w.r.t. \mathcal{K} , denoted as $\mathcal{K} \models (C \sqsubseteq D)^{\kappa}$ if every probabilistic model of \mathcal{K} is also a model of $(C \sqsubseteq D)^{\kappa}$.*

This is the natural extension of entailment to consider also the contexts. In our setting, however, contexts provide a means to express and reason with probabilities.

Definition 11 (subsumption probability). *Let $\mathcal{P} = (\mathcal{J}, \mathcal{P}_{\mathcal{J}})$ be a probabilistic model of the KB \mathcal{K} , κ a context, and C, D two concepts. The probability of $(C \sqsubseteq D)^{\kappa}$ w.r.t. \mathcal{P} is*

$$P_{\mathcal{P}}((C \sqsubseteq D)^{\kappa}) = \sum_{\mathcal{V} \in \mathcal{J}, \mathcal{V} \models (C \sqsubseteq D)^{\kappa}} P_{\mathcal{J}}(\mathcal{V}).$$

The probability of $(C \sqsubseteq D)^\kappa$ w.r.t. \mathcal{K} is

$$P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) = \inf_{\mathcal{P} \models \mathcal{K}} P_{\mathcal{P}}((C \sqsubseteq D)^\kappa).$$

C is positively subsumed by D in κ iff $P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) > 0$; it is p -subsumed iff $P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) \geq p$; it is exactly p -subsumed iff $P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) = p$, and it is almost certainly subsumed iff $P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) = 1$.

That is, the probability of a subsumption in a specific model is the sum of the probabilities of the worlds in which C is subsumed by D in context κ ; notice that this trivially includes all worlds where κ does not hold. In the case where \mathcal{K} is inconsistent we define the probability of all subsumptions as 1 to ensure our definition is consistent with general probability theory (recall that $\inf(\emptyset) = \infty$ in general).

Contextual subsumption is related to subsumption probability in the obvious way. Namely, a KB \mathcal{K} entails a contextual subsumption iff the probability of the subsumption in \mathcal{K} is 1.

Theorem 12. *Given a KB \mathcal{K} , concepts C and D , and a context κ , it holds that:*

$$\mathcal{K} \models (C \sqsubseteq D)^\kappa \text{ iff } P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) = 1.$$

This is convenient as it provides a method of reusing our results from Sect. 4 to compute subsumption probabilities.

Theorem 13. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{B})$ be a consistent KB, C, D two concepts, and κ a context. For the KB $\mathcal{K}' = (\mathcal{T} \cup \{C(a)^\kappa, \neg D(a)^\kappa\}, \mathcal{B})$ it holds that*

$$P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) = \sum_{\omega \models \phi_{\mathcal{K}'}^\perp} P_{\mathcal{B}}(\omega) + 1 - P_{\mathcal{B}}(\kappa).$$

Notice that the formula $\phi_{\mathcal{K}'}^\perp$ requires at most exponential space on the size of \mathcal{T} to be encoded. For each of the exponentially many worlds, computing $P_{\mathcal{B}}(\omega)$ requires polynomial time due to the chain rule. Hence, overall, the computation of the subsumption probabilities requires exponential time. Importantly, this bound does not depend on how $\phi_{\mathcal{K}'}^\perp$ was computed. This provides an exponential upper bound for computing the probability of a subsumption.

Corollary 14. *The probability of a subsumption w.r.t. a KB can be computed in exponential time on the size of the KB.*

Obviously, an exponential-time upper bound for computing the exact probability of a subsumption relation immediately yields an EXPTIME upper bound for deciding the other problems introduced in Definition 11. All these problems are also generalisations of the subsumption problem in \mathcal{ALC} . More precisely, given an \mathcal{ALC} TBox \mathcal{T} , we can create the \mathcal{BALC} KB $\mathcal{K} = (\mathcal{T}', \mathcal{B})$ where \mathcal{T}' contains all the axioms in \mathcal{T} labelled with the context x and \mathcal{B} contains only one Boolean node x that holds with probability 1. Given two concepts C, D $\mathcal{T} \models C \sqsubseteq D$ iff

C is almost certainly subsumed by D in context x . Since subsumption in \mathcal{ALC} is already EXPTIME-hard, we get that all these problems are EXPTIME-complete.

In practice, however, it may be too expensive to compute the exact probability when we are only interested in determining lower bounds, or the possibility of observing an entailment; for instance, when considering positive subsumption. Notice that, according to our semantics, a contextual GCI $(C \sqsubseteq D)^\kappa$ will hold in any world ω such that $\omega \not\models \kappa$. Thus, if the probability of this world is positive ($P_{\mathcal{B}}(\omega) > 0$), we can immediately guarantee that $P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) > 0$. Thus, positive subsumption can be decided without any ontological reasoning for any context that is not almost certain. In all other cases, the problem can still be reduced to inconsistency.

Theorem 15. *The concept C is positively subsumed by D in context κ w.r.t. $\mathcal{K} = (\mathcal{T}, \mathcal{B})$ iff $\mathcal{K}' = (\mathcal{T} \cup \{C(a)^\kappa, \neg D(a)^\kappa\}, \mathcal{B})$ is inconsistent or $P_{\mathcal{B}}(\kappa) < 1$.*

Assuming that the KB \mathcal{K} is consistent, the inconsistency in the KB \mathcal{K}' from this theorem can only arise from the inclusion of the two assertions which are required to hold in context κ . If the context κ is not known before hand, it is also possible to leverage the inconsistency decision process, which is the most expensive part of this method. Let $\mathcal{K}_\emptyset := (\mathcal{T} \cup \{C(a)^\emptyset, \neg D(a)^\emptyset\}, \mathcal{B})$. That is, we extend \mathcal{K} with assertions negating the subsumption relation, which should hold in all contexts. From Theorem 7 we conclude that $\phi_{\mathcal{K}_\emptyset}^\perp$ encodes all the contexts in which $C \sqsubseteq D$ must hold. Notice that the computation of $\phi_{\mathcal{K}_\emptyset}^\perp$ does not depend on the context κ but can be used to decide positive subsumption for any context.

Corollary 16. *The concept C is positively subsumed by D in context κ w.r.t. $\mathcal{K} = (\mathcal{T}, \mathcal{B})$ iff κ entails $\phi_{\mathcal{K}_\emptyset}^\perp$ or $P_{\mathcal{B}}(\kappa) < 1$.*

Considering the probabilities of contextual subsumption relations may lead to unexpected results arising from the contextual semantics. Indeed, it always holds (see Theorem 15) that $P_{\mathcal{K}}((C \sqsubseteq D)^\kappa) \geq 1 - P_{\mathcal{B}}(\kappa)$. In other words, the probability of a subsumption in a very unlikely context will always be very high, regardless of the KB and concepts used. In some cases, it may be more meaningful to consider a *conditional* probability under the assumption that the context κ holds.

Definition 17 (conditional subsumption). *Let $\mathcal{P} = (\mathcal{J}, P_{\mathcal{J}})$ be a probabilistic model of the KB \mathcal{K} , κ, λ two contexts with $P_{\mathcal{B}}(\lambda) > 0$, and C, D two concepts. The conditional probability of $(C \sqsubseteq D)^\kappa$ given λ w.r.t. \mathcal{P} is*

$$P_{\mathcal{P}}((C \sqsubseteq D)^\kappa \mid \lambda) = \frac{\sum_{\mathcal{V} \in \mathcal{J}, \mathcal{V} \models \lambda, \mathcal{V} \models (C \sqsubseteq D)^\kappa} P_{\mathcal{J}}(\mathcal{V})}{P_{\mathcal{B}}(\lambda)}.$$

The conditional probability of $(C \sqsubseteq D)^\kappa$ given λ w.r.t. \mathcal{K} is

$$P_{\mathcal{K}}((C \sqsubseteq D)^\kappa \mid \lambda) = \inf_{\mathcal{P} \models \mathcal{K}} P_{\mathcal{P}}((C \sqsubseteq D)^\kappa \mid \lambda).$$

This definition follows the same principles of conditioning in probability theory, but extended to the open world interpretation provided by our model-based semantics. Notice that, in addition to the scaling factor $P_{\mathcal{B}}(\lambda)$ in the denominator, the nominator is also differentiated from Definition 11 by considering only the worlds that satisfy the context λ already.

Consider the numerator in the definition of conditional probabilities. Notice that it is pretty similar to Definition 11, except that the sum restricts to only the worlds that satisfy the context λ . Thus, the numerator can be obtained from the contextual probability in the context $\kappa \wedge \lambda$ excluding the worlds that violate λ . More formally, we have that

$$\begin{aligned} \sum_{\mathcal{V} \in \mathcal{J}, v^{\mathcal{V}} \models \lambda, \mathcal{V} \models (C \sqsubseteq D)^{\kappa}} P_{\mathcal{J}}(\mathcal{V}) &= \sum_{v^{\mathcal{V}} \models \lambda, v^{\mathcal{V}} \not\models \kappa} P_{\mathcal{J}}(\mathcal{V}) + \sum_{v^{\mathcal{V}} \models \lambda \wedge \kappa, \mathcal{V} \models C \sqsubseteq D} P_{\mathcal{J}}(\mathcal{V}) \\ &= P_{\mathcal{B}}(\lambda) - P_{\mathcal{B}}(\lambda \wedge \kappa) + P_{\mathcal{P}}((C \sqsubseteq D)^{\lambda \wedge \kappa}) - 1 + P_{\mathcal{B}}(\lambda \wedge \kappa) \\ &= P_{\mathcal{P}}((C \sqsubseteq D)^{\lambda \wedge \kappa}) + P_{\mathcal{B}}(\lambda) - 1. \end{aligned}$$

Thus we get the following result.

Theorem 18. $P_{\mathcal{K}}((C \sqsubseteq D)^{\kappa} \mid \lambda) = \frac{P_{\mathcal{P}}((C \sqsubseteq D)^{\lambda \wedge \kappa}) + P_{\mathcal{B}}(\lambda) - 1}{P_{\mathcal{B}}(\lambda)}$.

In particular, this means that also conditional probabilities can be computed through contextual probabilities, with a small overhead of computing the probability (on the BN \mathcal{B}) of the conditioning context λ .

As in the contextual case, if one is only interested in knowing that the subsumption is possible (that is, that it has a positive probability), then one can exploit the complex context describing the inconsistent contexts which, as mentioned before, can be precompiled to obtain the contexts in which a subsumption relation must be satisfied. However, in this case, entailment between contexts is not sufficient; one must still compute the probability of the contextual subsumption.

Corollary 19. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{B})$ be a consistent KB, C, D concepts, and κ, λ contexts s.t. $P_{\mathcal{B}}(\lambda) > 0$. $P_{\mathcal{K}}((C \sqsubseteq D)^{\kappa} \mid \lambda) > 0$ iff $P_{\mathcal{P}}((C \sqsubseteq D)^{\lambda \wedge \kappa}) > 1 - P_{\mathcal{B}}(\lambda)$.*

In other words, $P((C \sqsubseteq D)^{\kappa} \mid \lambda) > 0$ iff C is p -subsumed by D in $\kappa \wedge \lambda$ with $p = 1 - P_{\mathcal{B}}(\lambda)$.

Analogously to Definitions 10, 11, and 17, it is possible to define the notions of consistency of a concept C to hold in only some contexts, and based on it, the (potentially conditional) probability of such a contextual consistency problem. As mentioned already, it is well known that for every \mathcal{ALC} interpretation \mathcal{I} it holds that $C^{\mathcal{I}} = \emptyset$ iff $C^{\mathcal{I}} \subseteq \perp^{\mathcal{I}}$. Hence, all these problems can be solved through a direct reduction to their related subsumption problem.

We now turn our attention to the problem of instance checking. In this problem, the ABox also plays a role. Hence, we consider once again ontologies \mathcal{O} that can have in addition to GCIs, concept and role assertions.

6 Instance Checking

We consider a probabilistic extension to the classical instance checking problem. In \mathcal{BALC} we call this problem probabilistic instance checking and we define both a decision problem and probability calculation for it next.

Given a KB \mathcal{K} and a context κ , the individual name a is an *instance* of the concept C in κ w.r.t. \mathcal{K} , written $\mathcal{K} \models C(a)^\kappa$, iff for all probabilistic models $\mathcal{P} = (\mathcal{J}, P_{\mathcal{J}})$ of \mathcal{K} and for all $\mathcal{V} \in \mathcal{J}$ it holds that $\mathcal{V} \models C(a)^\kappa$. That is, if every interpretation in \mathcal{P} satisfies the assertion $C(a)^\kappa$. Note that as before, instance checking in \mathcal{ALC} is a special case of this definition, that can be obtained by considering a BN with only one variable that is true with probability 1. Notice that, contrary to the case of satisfiability studied at the end of the last section, it is not possible to reduce instance checking to subsumption since an instance may be caused by ABox assertions. However, it may be reduced to consistency.

Theorem 20. *Given $a \in N_I$, a concept C , a context κ , and a KB $\mathcal{K} = (\mathcal{O}, \mathcal{B})$, $\mathcal{K} \models C(a)^\kappa$ iff the KB $\mathcal{K}' = (\mathcal{O} \cup \{(-C(a))^\kappa\}, \mathcal{B})$ is inconsistent.*

In particular, this means that instance checking is at most as hard as deciding consistency. As mentioned already, it is also at least as hard as instance checking in the classical \mathcal{ALC} . Hence we get the following result.

Lemma 21. *Instance checking in a \mathcal{BALC} KB is EXPTIME-complete.*

Let us now consider the probabilistic entailments related to instance checking.

Definition 22 (instance probability). *The probability of an instance in a probabilistic model $\mathcal{P} = (\mathcal{J}, P_{\mathcal{J}})$ of the KB \mathcal{K} is*

$$P_{\mathcal{P}}(C(x)^\kappa) = \sum_{\mathcal{V} \in \mathcal{J}, \mathcal{V} \models C(x)^\kappa} P_{\mathcal{J}}(\mathcal{V}).$$

The instance probability w.r.t. a KB \mathcal{K} is

$$P_{\mathcal{K}}(C(x)^\kappa) = \inf_{\mathcal{P} \models \mathcal{K}} P_{\mathcal{P}}(C(x)^\kappa).$$

The conditional probability of an instance in a particular probabilistic model $\mathcal{P} = (\mathcal{J}, P_{\mathcal{J}})$ is

$$P_{\mathcal{P}}(C(x)^\kappa \mid \lambda) = \frac{\sum_{\mathcal{V} \in \mathcal{J}, \mathcal{V} \models \lambda, \mathcal{V} \models C(x)^\kappa} P_{\mathcal{J}}(\mathcal{V})}{P_{\mathcal{B}}(\lambda)},$$

The probability of the conditional instance in \mathcal{K} is:

$$P_{\mathcal{K}}(C(x)^\kappa \mid \lambda) = \inf_{\mathcal{P} \models \mathcal{K}} P_{\mathcal{P}}(C(x)^\kappa \mid \lambda)$$

The probability of all instance checks for an inconsistent KB is always 1 to keep our definitions consistent with probability theory.

As we did for subsumption, we can exploit the reasoning techniques for deciding inconsistency of a \mathcal{BALC} KB to find out the contextual and conditional probabilities of an instance. Moreover, the method can be further optimised in the cases where we are only interested in probabilistic bounds. In particular, we can adapt Theorem 18 to this case.

Theorem 23. $P_{\mathcal{K}}(C(x)^{\kappa} \mid \lambda) = \frac{P_{\mathcal{K}}(C(x)^{\kappa \wedge \lambda}) + P_{\mathcal{B}}(\lambda) - 1}{P(\lambda)}$.

7 Conclusions

We have presented a new probabilistic extension of the DL \mathcal{ALC} based on the ideas of Bayesian ontology languages, in which certain knowledge is dependent on the uncertain context where it holds. Our work extends the results on \mathcal{BEL} [9, 10] to a propositionally closed ontology language. The main notions follow the basic ideas of Bayesian ontology languages [11]; however, by focusing on a specific logic, we are able to produce a tableaux-based decision algorithm for KB consistency, in contrast to the generic black-box algorithms proposed in the literature. Our algorithm extends the classical tableau algorithm for \mathcal{ALC} with techniques originally developed for axiom pinpointing. The main differences are the use of multi-valued variables in the definition of the contexts, and the possibility of having complex contexts (not only unique variables) labeling individual axioms. In general, we have shown that adding context-based uncertainty to an ontology does not increase the complexity of reasoning in this logic: all (probabilistic) reasoning problems can still be solved in exponential time.

Theorems 6, 7 and 8 yield an effective decision method \mathcal{BALC} KB consistency, through the computation and handling of two complex contexts. Notice that $\phi_{\mathcal{B}}$ can be computed in linear time on the size of \mathcal{B} , and satisfiability of the context $\phi_{\mathcal{K}}^{\perp} \wedge \phi_{\mathcal{B}}$ can be checked in non-deterministic polynomial time in the size of this context. However, the tableau algorithm for computing $\phi_{\mathcal{K}}^{\perp}$ is not optimal w.r.t. worst-case complexity. In fact, in the worst case it requires double exponential time, although the formula itself is only of exponential size. The benefit of this method, as in the classical case, is that it provides a better behaviour in the average case. To further improve the efficiency of our approach, one can think of adapting the methods from [21] to construct a compact representation of the context—akin to a binary decision diagram (BDD) [7, 14] for multi-valued variables—allowing for efficient weighted model counting. A task for future work is to exploit these data structures for practical development of our methods.

Recall that the most expensive part of our approach is the computation of the context $\phi_{\mathcal{K}}^{\perp}$. By slightly modifying the KB, we have shown that one computation of this context suffices to solve different problems of interest; in particular, contextual and conditional entailments—being subsumption, satisfiability, or instance checking—can be solved using $\phi_{\mathcal{K}'}^{\perp}$, for an adequately constructed \mathcal{K}' , regardless of the contexts under consideration.

An important next step will be to implement the methods described here, and compare the efficiency of our system to other probabilistic DL reasoners based on similar semantics. In particular, we would like to compare against the tools from [21]. Even though this latter system is also based on an extension of the tableaux algorithm for DLs, and use multiple-world semantics closely related to ours, a direct comparison would be unfair. Indeed, [21] makes use of stronger independence assumptions than ours. However, a well-designed experiment can shed light on the advantages and disadvantages of each method.

Another interesting problem for future work is to extend the query language beyond instance queries. To maintain some efficiency, this may require some additional restrictions on the language or the probabilistic structure. A more detailed study of this issue is needed.

References

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications, 2nd edn. Cambridge University Press, Cambridge (2007)
2. Baader, F., Hollunder, B.: Embedding defaults into terminological knowledge representation formalisms. *J. Autom. Reasoning* **14**(1), 149–180 (1995). <https://doi.org/10.1007/BF00883932>
3. Baader, F., Peñaloza, R.: Axiom pinpointing in general tableaux. In: Olivetti, N. (ed.) TABLEAUX 2007. LNCS (LNAI), vol. 4548, pp. 11–27. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-73099-6_4
4. Baader, F., Peñaloza, R.: Axiom pinpointing in general tableaux. *J. Logic Comput.* **20**(1), 5–34 (2010). <https://doi.org/10.1093/logcom/exn058>. Special issue: Tableaux and analytic proof methods
5. Botha, L.: The Bayesian description logic ALC. Master’s thesis, University of Cape Town, South Africa (2018)
6. Botha, L., Meyer, T., Peñaloza, R.: The Bayesian description logic BALC. In: Ortiz, M., Schneider, T. (eds.) Proceedings of the 31st International Workshop on Description Logics (DL 2018), CEUR Workshop Proceedings, vol. 2211. CEUR-WS.org (2018). <http://ceur-ws.org/Vol-2211/paper-09.pdf>
7. Brace, K.S., Rudell, R.L., Bryant, R.E.: Efficient implementation of a BDD package. In: Proceedings of the 27th ACM/IEEE Design Automation Conference, DAC 1990, pp. 40–45. ACM, New York (1990). <https://doi.org/10.1145/123186.123222>
8. Ceylan, İ.İ.: Query answering in probabilistic data and knowledge bases. Ph.D. thesis, Dresden University of Technology, Germany (2018)
9. Ceylan, İ.İ., Peñaloza, R.: The Bayesian description logic \mathcal{BEL} . In: Demri, S., Kapur, D., Weidenbach, C. (eds.) IJCAR 2014. LNCS (LNAI), vol. 8562, pp. 480–494. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08587-6_37
10. Ceylan, İ.İ., Peñaloza, R.: Tight complexity bounds for reasoning in the description logic \mathcal{BEL} . In: Fermé, E., Leite, J. (eds.) JELIA 2014. LNCS (LNAI), vol. 8761, pp. 77–91. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11558-0_6
11. Ceylan, İ.İ., Peñaloza, R.: The Bayesian ontology language BEL. *J. Autom. Reasoning* **58**(1), 67–95 (2017). <https://doi.org/10.1007/s10817-016-9386-0>

12. d'Amato, C., Fanizzi, N., Lukasiewicz, T.: Tractable reasoning with bayesian description logics. In: Greco, S., Lukasiewicz, T. (eds.) SUM 2008. LNCS (LNAI), vol. 5291, pp. 146–159. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87993-0_13
13. Gottlob, G., Lukasiewicz, T., Simari, G.I.: Answering threshold queries in probabilistic datalog+/- ontologies. In: Benferhat, S., Grant, J. (eds.) SUM 2011. LNCS (LNAI), vol. 6929, pp. 401–414. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23963-2_31
14. Lee, C.Y.: Representation of switching circuits by binary-decision programs. *Bell Syst. Tech. J.* **38**, 985–999 (1959)
15. Lee, K., Meyer, T.A., Pan, J.Z., Booth, R.: Computing maximally satisfiable terminologies for the description logic ALC with cyclic definitions. In: Parsia, B., Sattler, U., Toman, D. (eds.) Proceedings of the 2006 International Workshop on Description Logics (DL 2006), CEUR Workshop Proceedings, vol. 189. CEUR-WS.org (2006). http://ceur-ws.org/Vol-189/submission_29.pdf
16. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the semantic web. *J. Web Semant.* **6**(4), 291–308 (2008). <https://doi.org/10.1016/j.websem.2008.04.001>
17. Meyer, T.A., Lee, K., Booth, R., Pan, J.Z.: Finding maximally satisfiable terminologies for the description logic ALC. In: Proceedings of the Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, pp. 269–274. AAAI Press (2006). <http://www.aaai.org/Library/AAAI/2006/aaai06-043.php>
18. Pearl, J.: Bayesian networks: a model of self-activated memory for evidential reasoning. In: Proceedings of Cognitive Science Society (CSS-7), pp. 329–334 (1985)
19. Peñaloza, R.: Axiom-pinpointing in description logics and beyond. Ph.D. thesis, Dresden University of Technology, Germany (2009)
20. Schmidt-Schauß, M., Smolka, G.: Attributive concept descriptions with complements. *Artif. Intell.* **48**(1), 1–26 (1991). [https://doi.org/10.1016/0004-3702\(91\)90078-X](https://doi.org/10.1016/0004-3702(91)90078-X)
21. Zese, R., Bellodi, E., Riguzzi, F., Cota, G., Lamma, E.: Tableau reasoning for description logics and its extension to probabilities. *Ann. Math. Artif. Intell.* **82**(1–3), 101–130 (2018). <https://doi.org/10.1007/s10472-016-9529-3>