

Probabilistic Reasoning in the Description Logic \mathcal{ALCP} with the Principle of Maximum Entropy

Rafael Peñaloza¹ and Nico Potyka²

¹ Free University of Bozen-Bolzano, Italy

`rafael.penaloz@unibz.it`

² University of Osnabrück, Germany

`npotyka@uni-osnabrueck.de`

Abstract. A central question for knowledge representation is how to encode and handle uncertain knowledge adequately. We introduce the probabilistic description logic \mathcal{ALCP} that is designed for representing context-dependent knowledge, where the actual context taking place is uncertain. \mathcal{ALCP} allows the expression of logical dependencies on the domain and probabilistic dependencies on the possible contexts. In order to draw probabilistic conclusions, we employ the principle of maximum entropy. We provide reasoning algorithms for this logic, and show that it satisfies several desirable properties of probabilistic logics.

1 Introduction

A fundamental element of any intelligent application is storing and manipulating the knowledge from the application domain. Logic-based knowledge representation languages such as description logics (DLs) [1] provide a clear syntax and unambiguous semantics that guarantee the correctness of the results obtained. However, languages based on classical logic are ill-suited for handling the uncertainty inherent to many application domains. To overcome this limitation, various probabilistic logics have been investigated during the last three decades (e.g., [3, 15, 20]). In particular, several probabilistic DLs have been developed [18, 19]. To handle probabilistic knowledge, many approaches require a complete definition of joint probability distributions (JPD) [5, 6, 8, 16, 26]. One approach to avoid a full JPD specification was proposed by Paris [22]: the user gives a partial specification through a set of probabilistic constraints and the partial knowledge is completed by means of the principle of maximum entropy.

In this paper we consider a new probabilistic extension of description logics based on the principle of maximum entropy. In our approach we group different axioms from a knowledge base together into so-called contexts, which are identified by a propositional formula. Intuitively, each context corresponds to a possible situation, in which the associated sub-KB is guaranteed to hold. Uncertainty is associated to the contexts through a set of probabilistic constraints, which are interpreted under the principle of maximum entropy.

To facilitate the understanding of our approach, we focus on the DL \mathcal{ALC} [27] as a prototypical example of a knowledge representation language, and propositional probabilistic constraints as the framework for expressing uncertainty.

As reasoning service we consider subsumption relations between concepts given some partial knowledge of the current context. Since the knowledge in a knowledge base is typically incomplete, one cannot expect to obtain a precise probability for a given consequence. Instead, we compute a belief interval that describes all the probability degrees that can be associated to the consequence without contradiction. The lowest bound of the interval corresponds to a sceptical view, considering only the most fundamental models of the knowledge base. The upper bound, in contrast, reflects the credulous belief in which every context that is not explicitly removed is considered. In the worst-case, we get the trivial interval $[0, 1]$, in the best case, we get a point probability where the upper and lower bounds coincide. In some applications, it might be reasonable to consider only one of these bounds. For instance, if the probability interval that a treatment will cause heavy complications is $[0.01, 0.05]$, we might want to use the upper bound 0.05. In contrast, when the probability interval that a treatment will be successful is $[0.7, 0.9]$, we might be more interested in the lower bound 0.7.

The main contributions of this paper are the following:

- we define the new probabilistic description logic \mathcal{ALCP} that allows for a flexible description of axiomatic dependencies, and its reasoning problems (Section 3);
- we explain in detail how degrees of belief for the subsumption problem can be computed (Section 4); and
- we show that \mathcal{ALCP} satisfies several desirable properties of probabilistic logics (Section 5).³

2 Maximum Entropy

We start by recalling the basic notions of probabilistic propositional logic and the principle of maximum entropy.

Let \mathcal{L} be a propositional language constructed over a finite signature $\text{sig}(\mathcal{L})$, i.e., a set of propositional variables, in the usual way. An \mathcal{L} -interpretation v is a truth assignment of the propositional variables in $\text{sig}(\mathcal{L})$. $\text{Int}(\mathcal{L})$ denotes the set of all \mathcal{L} -interpretations. Satisfaction of a formula $\phi \in \mathcal{L}$ by an \mathcal{L} -interpretation $v \in \text{Int}(\mathcal{L})$ (denoted $v \models \phi$) is defined as usual. A probability distribution over \mathcal{L} is a function $P : \text{Int}(\mathcal{L}) \rightarrow [0, 1]$ where $\sum_{v \in \text{Int}(\mathcal{L})} P(v) = 1$. Probability distributions are extended to arbitrary \mathcal{L} -formulas ϕ by setting $P(\phi) = \sum_{v \models \phi} P(v)$.

Definition 1 (probabilistic constraints, models). *Given the propositional language \mathcal{L} , a probabilistic constraint (over \mathcal{L}) is an expression of the form*

$$c_0 + \sum_{i=1}^k c_i \cdot \mathbf{p}(\phi_i) \geq 0 \tag{1}$$

where $c_0, c_i \in \mathbb{R}$, and $\phi_i \in \mathcal{L}$, $1 \leq i \leq k$. A probability distribution P over \mathcal{L} is a model of the probabilistic constraint $c_0 + \sum_{i=1}^k c_i \cdot \mathbf{p}(\phi_i) \geq 0$ if and only if

³ The full proofs are available at the technical report [23].

$c_0 + \sum_{i=1}^k c_i \cdot P(\phi_i) \geq 0$. The distribution P is a model of the set of probabilistic constraints \mathcal{R} ($P \models \mathcal{R}$) iff it satisfies all the constraints in \mathcal{R} . The set of all models of \mathcal{R} is denoted by $\text{Mod}(\mathcal{R})$. If $\text{Mod}(\mathcal{R}) \neq \emptyset$, we say that \mathcal{R} is consistent.

Our probabilistic constraints can express the most common types of constraints considered in the literature of probabilistic logics. For instance, probabilistic conditionals $(\psi \mid \phi)[\ell, u]$ are satisfied iff $\ell \cdot P(\phi) \leq P(\psi \wedge \phi) \leq u \cdot P(\phi)$ [17]. That is, the conditional is satisfied iff the conditional probability of ψ given ϕ is between ℓ and u whenever $P(\phi) > 0$. Sometimes $P(\phi) > 0$ is demanded, but strict inequalities are computationally difficult and the semantical differences are negligible in many cases, see [25] for a thorough discussion. These conditions can be expressed in the form (1) as follows

$$\begin{aligned} \mathbf{p}(\psi \wedge \phi) - \ell \cdot \mathbf{p}(\phi) &\geq 0, & \text{and} \\ u \cdot \mathbf{p}(\phi) - \mathbf{p}(\psi \wedge \phi) &\geq 0. \end{aligned}$$

Probabilistic constraints can also express more complex restrictions; for example, we can state that the probability that a bird cannot fly is at most one fourth of the probability that a bird flies through the constraint

$$\frac{1}{4} \mathbf{p}(\text{BIRD} \wedge \text{FLIES}) - \mathbf{p}(\text{BIRD} \wedge \neg \text{FLIES}) \geq 0. \quad (2)$$

To improve readability, we will often rewrite constraints in a more compact manner, using conditionals as in the first example, or e.g. rewriting (2) as $\frac{1}{4} \mathbf{p}(\text{BIRD} \wedge \text{FLIES}) \geq \mathbf{p}(\text{BIRD} \wedge \neg \text{FLIES})$.

In general, consistent sets of probabilistic constraints have infinitely many models, and there is no obvious way to distinguish between them. One well-studied approach for dealing with this diversity is to focus on the model that maximizes the entropy

$$H(P) = - \sum_{v \in \text{Int}(\mathcal{L})} P(v) \cdot \log P(v).$$

From an information-theoretic point of view, the maximum entropy (ME) distribution can be regarded as the most conservative one in the sense that it minimizes the information-theoretic distance (that is, the KL-divergence) to the uniform distribution among all probability distributions that satisfy our constraints. In particular, if there are no restrictions on the probability distributions considered, then the uniform distribution is the ME distribution, see, e.g., [28] for a more detailed discussion of these issues. A complete characterization of maximum entropy for the purpose of uncertain reasoning can be found in [22].

Definition 2 (ME-model). *Let \mathcal{R} be a consistent set of probabilistic constraints. The ME-model $P_{\mathcal{R}}^{\text{ME}}$ of \mathcal{R} is the unique solution of the maximization problem $\arg \max_{P \models \mathcal{R}} H(P)$.*

Existence and uniqueness of $P_{\mathcal{R}}^{\text{ME}}$ follows from the fact that H is strictly concave and continuous, and that the probability distributions that satisfy \mathcal{R} form a

compact and convex set. $P_{\mathcal{R}}^{ME}$ is usually computed by deriving an unconstrained optimization problem by means of the Karush-Kuhn-Tucker conditions. The resulting problem can be solved, for instance, by (quasi-)Newton methods with cost $|Int(\mathcal{L})|^3$, see, e.g., [21] for more details on these techniques.

3 The Probabilistic Description Logic \mathcal{ALCP}

\mathcal{ALCP} is a probabilistic extension of the classical description logic \mathcal{ALC} capable of expressing complex logical and probabilistic relations. As with classical DLs, the main building blocks in \mathcal{ALCP} are *concepts*. Syntactically, \mathcal{ALCP} concepts are constructed exactly as \mathcal{ALC} concepts. Given two disjoint sets \mathbf{N}_C of *concept names* and \mathbf{N}_R of *role names*, \mathcal{ALCP} concepts are built using the grammar rule $C ::= A \mid \neg C \mid C \sqcap C \mid \exists r.C$, where $A \in \mathbf{N}_C$ and $r \in \mathbf{N}_R$. Note that we can derive disjunction, universal quantification and subsumption from these rules by using logical equivalences like $C_1 \sqcup C_2 \equiv \neg(\neg C_1 \sqcap \neg C_2)$. The knowledge of the application domain is expressed through a finite set of axioms that restrict the way the different concepts and roles may be interpreted. To express both probabilistic and logical relationships, each axiom is annotated with a formula from \mathcal{L} that intuitively expresses the context in which this axiom holds.

Definition 3 (KB). *An \mathcal{L} -restricted general concept inclusion (\mathcal{L} -GCI) is of the form $\langle C \sqsubseteq D : \kappa \rangle$ where C, D are \mathcal{ALCP} concepts and κ is an \mathcal{L} -formula. An \mathcal{L} -TBox is a finite set of \mathcal{L} -GCIs. An \mathcal{ALCP} knowledge base (KB) over \mathcal{L} is a pair $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ where \mathcal{R} is a set of probabilistic constraints and \mathcal{T} is an \mathcal{L} -TBox.*

Example 4. Consider an application modeling beliefs about bacterial and viral infections using the concept names *strep* (streptococcal infection), *bac* (bacterial infection), *vir* (viral infection), *inf* (infection), and *ab* (antibiotic); and the role names *sf* (suffers from), and *suc* (successful treatment); and the propositional variables *RES* (antibiotic resistance), and *H* (heavy use of antibiotics by patient). Define the \mathcal{L} -TBox \mathcal{T}_{exa} containing the \mathcal{L} -GCIs

$$\begin{aligned} \langle \exists sf.bac \sqsubseteq \exists suc.ab : \neg RES \wedge \neg H \rangle, \quad \langle \exists sf.vir \sqsubseteq \neg \exists suc.ab : \top \rangle, \quad \langle strep \sqsubseteq bac : \top \rangle, \\ \langle \exists sf.bac \sqsubseteq \neg \exists suc.ab : RES \rangle, \quad \langle bac \sqsubseteq inf : \top \rangle, \quad \langle vir \sqsubseteq inf : \top \rangle, \end{aligned}$$

where \top is any \mathcal{L} -tautology. For example, the first \mathcal{L} -GCI states that a bacterial infection can be treated successfully with antibiotics if no antibiotic resistance is present and there was no heavy use of antibiotics; the second one states that viral infections can never be treated with antibiotics successfully. Consider additionally the set \mathcal{R} containing the probabilistic constraints containing

$$(RES)[0.05], \quad (RES \mid H)[0.8].$$

That is, the probability of an antibiotic resistance is 5% if no further information is given. If the patient used antibiotics heavily, the probability increases to 80%.

Notice that the probabilistic constraints, and hence the representation of the uncertainty in the knowledge, refer only to the propositional formulas that label the \mathcal{L} -GCIs. In \mathcal{ALCP} , the uncertainty of the knowledge is handled through these propositional formulas as explained next.

A possible world interprets both the axiom language (i.e., the concept and role names) and the context language (the propositional variables). Intuitively, it describes a possible context (\mathcal{L} -interpretation) together with the relationships between concepts in that situation (\mathcal{ALC} -interpretation).

Definition 5 (possible world). A possible world is a triple $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}, v^{\mathcal{I}})$ where $\Delta^{\mathcal{I}}$ is a non-empty set (called the domain), $v^{\mathcal{I}}$ is an \mathcal{L} -interpretation, and $\cdot^{\mathcal{I}}$ is an interpretation function that maps every concept name A to a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and every role name r to a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$.

The interpretation function $\cdot^{\mathcal{I}}$ is extended to complex concepts as usual in DLs by letting $(\neg C)^{\mathcal{I}} := \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$; $(\exists r.C)^{\mathcal{I}} := \{d \in \Delta^{\mathcal{I}} \mid \exists e \in \Delta^{\mathcal{I}}.(d, e) \in r^{\mathcal{I}}, e \in C^{\mathcal{I}}\}$; and $(C \sqcap D)^{\mathcal{I}} := C^{\mathcal{I}} \cap D^{\mathcal{I}}$. A possible world is a model of an \mathcal{L} -GCI iff it satisfies the description logic constraint of the axiom whenever it satisfies the context.

Definition 6 (model of TBox). A possible world $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}, v^{\mathcal{I}})$ is a model of the \mathcal{L} -GCI $\langle C \sqsubseteq D : \kappa \rangle$ ($\mathcal{I} \models \langle C \sqsubseteq D : \kappa \rangle$) iff (i) $v^{\mathcal{I}} \not\models \kappa$, or (ii) $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. It is a model of the \mathcal{L} -TBox \mathcal{T} iff it is a model of all \mathcal{L} -GCIs in \mathcal{T} .

The classical DL \mathcal{ALC} is a special case of \mathcal{ALCP} where all the axioms are annotated with an \mathcal{L} -tautology \top . To preserve the syntax of classical DLs, we denote such \mathcal{L} -GCIs as $C \sqsubseteq D$ instead of $\langle C \sqsubseteq D : \top \rangle$. In this case, the condition (i) from Definition 6 cannot be satisfied, and hence a model is required to satisfy $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all \mathcal{L} -GCIs $C \sqsubseteq D$ in the TBox. For a deeper introduction to classical \mathcal{ALC} , see [1].

According to our semantics, we only demand that the \mathcal{L} -GCIs are satisfied in some specific contexts. Thus, it is often useful to focus on the classical \mathcal{ALC} TBox that contains the knowledge that holds in a particular situation. For a KB $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ and $v \in \text{Int}(\mathcal{L})$, the v -restricted TBox is the \mathcal{ALC} TBox

$$\mathcal{T}_v := \{C \sqsubseteq D \mid \langle C \sqsubseteq D : \kappa \rangle \in \mathcal{T}, v \models \kappa\}.$$

The possible world \mathcal{I} satisfies \mathcal{T}_v ($\mathcal{I} \models \mathcal{T}_v$) if for all \mathcal{L} -GCIs $C \sqsubseteq D \in \mathcal{T}_v$ it holds that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. In the following, we will often consider *subsumption* and *strong non-subsumption* between concepts w.r.t. a restricted TBox. We say that C is *subsumed* by D w.r.t. \mathcal{T}_v ($\mathcal{T}_v \models C \sqsubseteq D$) if for every $\mathcal{I} \models \mathcal{T}_v$ it holds that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. Dually, C is *strongly non-subsumed* by D w.r.t. \mathcal{T}_v ($\mathcal{T}_v \models C \not\sqsubseteq D$) if for every $\mathcal{I} \models \mathcal{T}_v$, $C^{\mathcal{I}} \not\subseteq D^{\mathcal{I}}$ holds. Notice that strong non-subsumption requires that the inclusion between axioms does not hold in *any* possible world satisfying \mathcal{T}_v . Hence, this condition is more strict than just negating the subsumption relation.

We now describe how the probabilistic constraints are handled in our logic. An \mathcal{ALCP} -interpretation consists of a finite set of possible worlds and a probability function over these worlds.

Definition 7 (ALCP-interpretation). An ALCP-interpretation is a pair of the form $\mathcal{P} = (\mathfrak{I}, P_{\mathfrak{I}})$, where \mathfrak{I} is a non-empty, finite set of possible worlds and $P_{\mathfrak{I}}$ is a probability distribution over \mathfrak{I} .

Each ALCP-interpretation induces a probability distribution over \mathcal{L} . The probability of a context can be obtained by adding the probabilities of all possible worlds in which this context holds.

Definition 8 (distribution induced by \mathcal{P}). Let $\mathcal{P} = (\mathfrak{I}, P_{\mathfrak{I}})$ be an ALCP-interpretation. The probability distribution $P^{\mathcal{P}} : \text{Int}(\mathcal{L}) \rightarrow [0, 1]$ induced by \mathcal{P} is defined by $P^{\mathcal{P}}(v) := \sum_{\mathcal{I} \in \mathfrak{I}|_v} P_{\mathfrak{I}}(\mathcal{I})$, where $\mathfrak{I}|_v = \{(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}, v^{\mathcal{I}}) \in \mathfrak{I} \mid v^{\mathcal{I}} = v\}$.

As usual, reasoning is restricted to interpretations that satisfy the restrictions imposed by the knowledge base. In our case, we have to demand that the interpretation is consistent with both the classical and the probabilistic part of our knowledge base. That is, we consider only those possible worlds that satisfy both the terminological knowledge (\mathcal{T}) and the probabilistic constraints (\mathcal{R}).

Definition 9 (model). Let $\mathcal{P} = (\mathfrak{I}, P_{\mathfrak{I}})$ be an ALCP-interpretation. \mathcal{P} is consistent with the TBox \mathcal{T} if every $\mathcal{I} \in \mathfrak{I}$ is a model of \mathcal{T} . \mathcal{P} is consistent with the set of probabilistic constraints \mathcal{R} iff $P^{\mathcal{P}} \models \mathcal{R}$. The ALCP-interpretation \mathcal{P} is a model of the KB $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ iff it is consistent with both \mathcal{T} and \mathcal{R} . As usual, a KB is consistent iff it has a model.

Notice that ALCP-KBs can express both, logical and probabilistic dependencies between axioms. For instance, two \mathcal{L} -GCIs $\langle C_1 \sqsubseteq D_1 : \kappa_1 \rangle$ and $\langle C_2 \sqsubseteq D_2 : \kappa_2 \rangle$ where $\kappa_1 \Rightarrow \kappa_2$ express that whenever the first \mathcal{L} -GCI is satisfied, the second one must also hold. Similarly, the probabilistic dependencies between axioms are expressed via the probabilistic constraints of the labeling formulas.

We are interested in computing degrees of belief for subsumption relations between concepts. We define the conditional probability of a subsumption relation given a context with respect to a given ALCP-interpretation following the usual notions of conditioning.

Definition 10 (probability of subsumption). Let C, D be concepts, κ a context and \mathcal{P} an ALCP-interpretation. The conditional probability of $C \sqsubseteq D$ given κ with respect to \mathcal{P} is

$$Pr_{\mathcal{P}}(C \sqsubseteq D \mid \kappa) := \frac{\sum_{\mathcal{I} \in \mathfrak{I}, \mathcal{I} \models \kappa, \mathcal{I} \models C \sqsubseteq D} P_{\mathfrak{I}}(\mathcal{I})}{\sum_{\mathcal{I} \in \mathfrak{I}, \mathcal{I} \models \kappa} P_{\mathfrak{I}}(\mathcal{I})}. \quad (3)$$

Notice that the denominator in (3) can be rewritten as

$$\sum_{\mathcal{I} \in \mathfrak{I}, \mathcal{I} \models \kappa} P_{\mathfrak{I}}(\mathcal{I}) = \sum_{v \models \kappa} \sum_{\mathcal{I} \in \mathfrak{I}|_v} P_{\mathfrak{I}}(\mathcal{I}) = \sum_{v \models \kappa} P^{\mathcal{P}}(v) = P^{\mathcal{P}}(\kappa).$$

As usual, the conditional probability is only well-defined when $P^{\mathcal{P}}(\kappa) > 0$.

Recall that the set of probabilistic constraints \mathcal{R} may be satisfied by an infinite class of probability distributions. In the spirit of maximum entropy reasoning, we consider only the most conservative ones in the sense that they induce the ME-model $P_{\mathcal{R}}^{ME}$ of \mathcal{R} .

Definition 11 (ME- \mathcal{ALCP} -model). An \mathcal{ALCP} -model \mathcal{P} of \mathcal{K} is called an ME- \mathcal{ALCP} -model of \mathcal{K} iff $P^{\mathcal{P}} = P_{\mathcal{R}}^{ME}$. The set of all ME- \mathcal{ALCP} -models of \mathcal{K} is denoted by $Mod_{ME}(\mathcal{K})$. \mathcal{K} is called ME-consistent iff $Mod_{ME}(\mathcal{K}) \neq \emptyset$.

Note that ME-consistency is a strictly stronger notion of consistency. ME-consistent knowledge bases are always consistent, but the converse does not necessarily hold if the classical TBox obtained from \mathcal{T} by restricting to a context is inconsistent as we show in the following example.

Example 12. Let $\text{sig}(\mathcal{L}) = \{x\}$ and $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ be the KB with $\mathcal{R} = \emptyset$ and $\mathcal{T} = \{A \sqcup \neg A \sqsubseteq A \sqcap \neg A : x\}$. Since $A \sqcup \neg A \sqsubseteq A \sqcap \neg A$ is contradictorial, each \mathcal{ALCP} -model of \mathcal{K} must satisfy $\neg x$. There certainly are such models, but in each such model \mathcal{P} , $P^{\mathcal{P}}(x) = 0$. However, since $\mathcal{R} = \emptyset$, we have $P_{\mathcal{R}}^{ME}(x) = 0.5$ and hence \mathcal{K} has no ME-model.

ME-inconsistency rules out some undesired cases in which the whole knowledge base is consistent, but the TBox restricted to some context is inconsistent. The following theorem gives a simple characterization of ME-consistency: to verify ME-consistency of a KB, it suffices to check consistency of the TBoxes induced by the \mathcal{L} -interpretations that have positive probability with respect to $P_{\mathcal{R}}^{ME}$. By the properties of the ME distribution, these are the interpretations that are not explicitly restricted to have zero probability through \mathcal{R} .

Theorem 13. The KB $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ is ME-consistent iff for every $v \in \text{Int}(\mathcal{L})$ such that $P_{\mathcal{R}}^{ME}(v) > 0$, \mathcal{T}_v is consistent.

For the rest of this paper we consider only ME-consistent KBs. Hence, whenever we speak of a KB \mathcal{K} , we implicitly assume that \mathcal{K} has at least one ME-model.

We are interested in computing the probability of a subsumption relation w.r.t. a given KB \mathcal{K} . Notice that, although we consider only one probability distribution $P_{\mathcal{R}}^{ME}$, there can still exist many different ME-models of \mathcal{K} , which yield different probabilities for the same subsumption relation. One way to handle this is to consider the smallest and largest probabilities that can be consistently associated to this relation. We call them the *sceptical* and the *credulous* degrees of belief, respectively.

Definition 14 (degree of belief). Let C, D be \mathcal{ALCP} concepts, κ a context, and $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ an \mathcal{ALCP} KB. The sceptical degree of belief of $C \sqsubseteq D$ given κ w.r.t. \mathcal{K} is

$$\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa) := \inf_{\mathcal{P} \in Mod_{ME}(\mathcal{K})} Pr_{\mathcal{P}}(C \sqsubseteq D \mid \kappa).$$

The credulous degree of belief of $C \sqsubseteq D$ given κ w.r.t. \mathcal{K} is

$$\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa) := \sup_{\mathcal{P} \in Mod_{ME}(\mathcal{K})} Pr_{\mathcal{P}}(C \sqsubseteq D \mid \kappa).$$

Example 15. Consider \mathcal{K}_{exa} from Example 4. If we ask for the degrees of belief that a patient who suffers from an infection can be successfully treated with

antibiotics, we obtain

$$\begin{aligned}\mathcal{B}_{\mathcal{K}_{\text{exa}}}^s(\exists\text{sf.inf} \sqsubseteq \exists\text{suc.ab} \mid \top) &= 0, \\ \mathcal{B}_{\mathcal{K}_{\text{exa}}}^c(\exists\text{sf.inf} \sqsubseteq \exists\text{suc.ab} \mid \top) &= 1.\end{aligned}$$

These bounds are not very informative, but they are perfectly justified by our knowledge base since we do not know anything about the effectiveness of antibiotics with respect to infections in general. However, for a patient who suffers from a streptococcal infection we get

$$\begin{aligned}\mathcal{B}_{\mathcal{K}_{\text{exa}}}^s(\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab} \mid \top) &= 0.9405, \\ \mathcal{B}_{\mathcal{K}_{\text{exa}}}^c(\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab} \mid \top) &= 0.95.\end{aligned}$$

If we know that this patient used antibiotics heavily in the past, then there is nothing in our knowledge base that guarantees the existence of a successful treatment. Hence, the degrees of belief become

$$\begin{aligned}\mathcal{B}_{\mathcal{K}_{\text{exa}}}^s(\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab} \mid \text{H}) &= 0 \\ \mathcal{B}_{\mathcal{K}_{\text{exa}}}^c(\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab} \mid \text{H}) &= 0.2.\end{aligned}$$

Our definition of the sceptical degree of belief raises a philosophical question: should there be no difference between the degree of belief 0 and an infinitely small degree of belief? A dual question arises for the credulous degree of belief and the probability 1. However, as we show in the next section, the sceptical and credulous degrees of belief actually correspond to minimum and maximum rather than to infimum and supremum (see Corollary 20) so that these questions become vacuous. From the following theorem we can conclude that every intermediate degree can also be obtained by some model of the KB.

Theorem 16 (Intermediate Value Theorem). *Let $p_1 < p_2$ and \mathcal{P}_1 and \mathcal{P}_2 be two ME-ALCP-models of the KB $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ such that $\text{Pr}_{\mathcal{P}_1}(C \sqsubseteq D \mid \kappa) = p_1$ and $\text{Pr}_{\mathcal{P}_2}(C \sqsubseteq D \mid \kappa) = p_2$. Then for each p between p_1 and p_2 there exists an ME-ALCP-model \mathcal{P} of \mathcal{K} such that $\text{Pr}_{\mathcal{P}}(C \sqsubseteq D \mid \kappa) = p$*

As we will show in Corollary 20, both the sceptical degree $\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa)$ and the credulous degree $\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa)$ are in fact witnessed by some ME-models. Therefore it is meaningful to consider the whole interval of beliefs between $\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa)$ and $\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa)$.

Definition 17 (belief interval). *Let C, D be ALCP concepts, $\kappa \in \mathcal{L}$ a context and $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ a ALCP KB. The belief interval for $C \sqsubseteq D$ w.r.t. \mathcal{K} given κ is*

$$\mathcal{B}_{\mathcal{K}}(C \sqsubseteq D \mid \kappa) := [\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa), \mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa)].$$

4 Computing Beliefs

In this section we show how to compute the belief interval. The first theorem states that the sceptical degree of belief for a subsumption relation can be

computed by adding the probabilities of those \mathcal{L} -interpretations w that entail this subsumption in the corresponding restricted TBox \mathcal{T}_w .

Theorem 18. *Let $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ be a KB, C, D two concepts, and κ a context such that $P_{\mathcal{R}}^{ME}(\kappa) > 0$. Then*

$$\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa) = \frac{\sum_{w \in \text{Int}(\mathcal{L}), \mathcal{T}_w \models C \sqsubseteq D, w \models \kappa} P_{\mathcal{R}}^{ME}(w)}{P_{\mathcal{R}}^{ME}(\kappa)}.$$

Dually, the credulous degree of belief for a subsumption relation can be computed by removing all the situations in which this relation cannot possibly hold.

Theorem 19. *Let $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ be a KB, C, D two concepts, and κ a context with $P_{\mathcal{R}}^{ME}(\kappa) > 0$. Then*

$$\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa) = 1 - \frac{\sum_{w \in \text{Int}(\mathcal{L}), \mathcal{T}_w \models C \not\sqsubseteq D, w \models \kappa} P_{\mathcal{R}}^{ME}(w)}{P_{\mathcal{R}}^{ME}(\kappa)}.$$

To prove these theorems, one can build two models of the KB \mathcal{K} , \mathcal{P} and \mathcal{Q} such that $Pr_{\mathcal{P}}(C \sqsubseteq D \mid \kappa)$ and $Pr_{\mathcal{Q}}(C \sqsubseteq D \mid \kappa)$ are those degrees expressed by Theorems 18 and 19, respectively. As a byproduct of these proofs, we obtain that the infimum and supremum that define the sceptical and the credulous degrees of belief actually correspond to minimum and maximum taken by some ME-models, yielding the following corollary.

Corollary 20. *Let \mathcal{K} be an \mathcal{ALCP} KB, C, D be two concepts, and κ be a context. There exist two ME-models \mathcal{P}, \mathcal{Q} of \mathcal{K} with $\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa) = Pr_{\mathcal{P}}(C \sqsubseteq D \mid \kappa)$ and $\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa) = Pr_{\mathcal{Q}}(C \sqsubseteq D \mid \kappa)$.*

The direct consequence of Theorems 18 and 19 is that if we want to compute the belief interval for $C \sqsubseteq D$ given some context, it suffices to identify all \mathcal{L} -interpretations whose induced (classical) TBoxes entail the subsumption relation $C \sqsubseteq D$ (for the sceptical belief) or the strong non-subsumption $C \not\sqsubseteq D$ (for credulous belief). Recall that every set of propositional interpretations can be represented by a propositional formula. This motivates the following definition.

Definition 21 (consequence formula). *An \mathcal{L} -formula ϕ is a consequence formula for $C \sqsubseteq D$ (respectively $C \not\sqsubseteq D$) w.r.t. the \mathcal{L} -TBox \mathcal{T} if for every $w \in \text{Int}(\mathcal{L})$ it holds that $w \models \phi$ iff $\mathcal{T}_w \models C \sqsubseteq D$ (respectively $\mathcal{T}_w \models C \not\sqsubseteq D$).*

If we are able to compute these consequence formulas, then the computation of the belief interval can be reduced to the evaluation of the probability of these formulas w.r.t. the ME-distribution satisfying \mathcal{R} .

Theorem 22. *Let $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ be an \mathcal{ALCP} KB, ϕ and ψ be consequence formulas for $C \sqsubseteq D$ and $C \not\sqsubseteq D$ w.r.t. \mathcal{T} , respectively, and κ a context. Then $\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa) = P_{\mathcal{R}}^{ME}(\phi \mid \kappa)$ and $\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa) = 1 - P_{\mathcal{R}}^{ME}(\psi \mid \kappa)$.*

Algorithm 1 Computing degrees of belief

Input: KB $\mathcal{K} = (\mathcal{R}, \mathcal{T})$, concepts C, D , context κ
Output: Belief degrees $(\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D | \kappa), \mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D | \kappa))$
 $\ell_s \leftarrow 0; \ell_c \leftarrow 0$
for all $v \in \text{Int}(\mathcal{L})$ **do**
 if $v \models \kappa$ **then**
 if $\mathcal{T}_v \models C \sqsubseteq D$ **then**
 $\ell_s \leftarrow \ell_s + P_{\mathcal{R}}^{ME}(v)$
 else if $\mathcal{T}_v \models C \not\sqsubseteq D$ **then**
 $\ell_c \leftarrow \ell_c + P_{\mathcal{R}}^{ME}(v)$
return $(\ell_s / P_{\mathcal{R}}^{ME}(\kappa), 1 - \ell_c / P_{\mathcal{R}}^{ME}(\kappa))$

Example 23. In our running example, one can see that a consequence formula for $\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab}$ is $\neg\text{RES} \wedge \neg\text{H}$. Indeed, in order to deduce this consequence it is necessary to satisfy the first axiom of \mathcal{T}_{exa} , which is only guaranteed in the context $\neg\text{RES} \wedge \neg\text{H}$. Similarly, RES is a consequence formula for $\exists\text{sf.strep} \not\sqsubseteq \exists\text{suc.ab}$. Knowing both the consequence formulas and the ME-model, we can deduce

$$\begin{aligned} \mathcal{B}_{\mathcal{K}_{\text{exa}}}^s(\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab} \mid \top) &= P_{\mathcal{R}}^{ME}(\neg\text{RES} \wedge \neg\text{H}) = 0.9405, \text{ and} \\ \mathcal{B}_{\mathcal{K}_{\text{exa}}}^c(\exists\text{sf.strep} \sqsubseteq \exists\text{suc.ab} \mid \text{H}) &= 1 - P_{\mathcal{R}}^{ME}(\text{RES} \mid \text{H}) = 0.2. \end{aligned}$$

In particular, Theorem 22 implies that the belief interval can be computed in two phases. The first phase uses purely logical reasoning to compute the consequence formulas, while the second phase applies probabilistic inferences to compute the degrees of belief from these formulas. We now briefly explain how the consequence formulas can be computed.

Notice first that subsumption and non-subsumption are monotonic consequences in the sense of [2]; that is, if an \mathcal{ALC} TBox \mathcal{T} entails the subsumption $C \sqsubseteq D$, then every superset of \mathcal{T} also entails this consequence. Similarly, adding more axioms to a TBox entailing $C \not\sqsubseteq D$ does not remove this entailment. Moreover, the set of all \mathcal{L} -formulas (modulo logical equivalence) forms a distributive lattice ordered by generality, in which \mathcal{L} -interpretations are all the join prime elements. Thus, the consequence formulas from Definition 21 are in fact the so-called *boundaries* from [2]. Hence, they can be computed using any of the known boundary computation approaches.

Assuming that the number of contexts is small in comparison to the size of the TBox, it is better to compute the degrees of belief through a more direct approach following Theorems 18 and 19. In order to compute $\mathcal{B}_{\mathcal{K}}^s(C \sqsubseteq D \mid \kappa)$ and $\mathcal{B}_{\mathcal{K}}^c(C \sqsubseteq D \mid \kappa)$, it suffices to enumerate all interpretations $v \in \text{Int}(\mathcal{L})$ and check whether $\mathcal{T}_v \models C \sqsubseteq D$ or $\mathcal{T}_v \models C \not\sqsubseteq D$, and $v \models \kappa$, or not (see Algorithm 1). This approach requires $2^{|\text{sig}(\mathcal{L})|}$ calls to a standard \mathcal{ALC} reasoner, and each of these calls runs in exponential time on $|\mathcal{T}|$ [9]. Notice that this algorithm has an *any-time* behaviour: it is possible to stop its execution at any moment and obtain an approximation of the belief interval. Moreover, the longer the algorithm runs, the better this approximation becomes. Thus, this method is adequate for a

system where finding good approximations efficiently may be more important than computing the precise answers.

5 Properties

We now investigate some properties of probabilistic logics [22]. First we show that \mathcal{ALCP} is *language and representation invariant*. Invariance is meant with respect to logical objects. Language invariance means that just extending the language without changing the knowledge base should not affect reasoning results. Representation invariance means that equivalent knowledge bases should yield equal inference results. Notice that different notions of *representation dependence* exist in the literature. For instance, in [11] a very different notion is considered, where the language and the knowledge base are changed simultaneously. This case is not covered by our notion of representation invariance. \mathcal{ALCP} also satisfies an *independence* property; i.e., reasoning results about a part of the language are not changed, when we add knowledge about an independent part of the language. Finally, \mathcal{ALCP} is *continuous* in the sense that minor changes in the probabilistic knowledge expressed by a knowledge base cannot induce major changes in the reasoning results.

Theorem 24 (Representation invariance). *Let $\mathcal{K}_i = (\mathcal{R}_i, \mathcal{T}_i)$, $i \in \{1, 2\}$, be two KBs such that $\text{Mod}(\mathcal{R}_1) = \text{Mod}(\mathcal{R}_2)$ and $\text{Mod}(\mathcal{T}_1) = \text{Mod}(\mathcal{T}_2)$. Then for all concepts C, D and contexts $\kappa \in \mathcal{L}$, $\mathcal{B}_{\mathcal{K}_1}(C \sqsubseteq D \mid \kappa) = \mathcal{B}_{\mathcal{K}_2}(C \sqsubseteq D \mid \kappa)$.*

\mathcal{ALCP} is not only representation invariant, but also language invariant. This property is of computational interest, in particular in combination with independence, that we investigate subsequently. To illustrate this, suppose that we added knowledge about bone fractures in our medical example, which is independent of the knowledge about infections. Independence guarantees that we can ignore the knowledge about infections when answering queries about bone fractures. In this way, we can decrease the size of the knowledge base. Language invariance guarantees that we can also ignore the concepts, relations and propositional variables related to the infection domain. Thus, we can decrease the size of the language. Exploiting both properties, the size of the computational problems can sometimes be decreased significantly.

Theorem 25 (Language Invariance). *Let $\mathcal{K}_1, \mathcal{K}_2$ be KBs over $\mathcal{L}^1, \mathcal{N}_C^1, \mathcal{N}_R^1$ and $\mathcal{L}^2, \mathcal{N}_C^2, \mathcal{N}_R^2$, respectively. If $\mathcal{K}_1 = \mathcal{K}_2$, $\mathcal{L}^1 \subseteq \mathcal{L}^2$, $\mathcal{N}_C^1 \subseteq \mathcal{N}_C^2$ and $\mathcal{N}_R^1 \subseteq \mathcal{N}_R^2$, then for all concepts $C, D \in \mathcal{N}_C^1$ and contexts $\kappa \in \mathcal{L}^1$, it holds that*

$$\mathcal{B}_{\mathcal{K}_1}(C \sqsubseteq D \mid \kappa) = \mathcal{B}_{\mathcal{K}_2}(C \sqsubseteq D \mid \kappa).$$

For an \mathcal{L} -TBox \mathcal{T} , we define the *signature* of \mathcal{T} to be the set $\text{sig}(\mathcal{T})$ of all concept names and role names appearing in \mathcal{T} . Likewise, $\text{sig}(\mathcal{R})$ is the set of all propositional variables appearing in \mathcal{R} . The signature of a KB $\mathcal{K} = (\mathcal{R}, \mathcal{T})$ is $\text{sig}(\mathcal{K}) := \text{sig}(\mathcal{R}) \cup \text{sig}(\mathcal{T})$.

Theorem 26 (Independence). *Let $\mathcal{K}_1, \mathcal{K}_2$ be s.t. $\text{sig}(\mathcal{K}_1) \cap \text{sig}(\mathcal{K}_2) = \emptyset$, C, D be two concepts, and κ a context where $(\text{sig}(C) \cup \text{sig}(D) \cup \text{sig}(\kappa)) \cap \text{sig}(\mathcal{K}_2) = \emptyset$. Then $\mathcal{B}(C \sqsubseteq_{\mathcal{K}_1} D \mid \kappa) = \mathcal{B}(C \sqsubseteq_{\mathcal{K}_1 \cup \mathcal{K}_2} D \mid \kappa)$.*

To conclude, we consider continuity. One important practical feature of continuous probabilistic logics is that they guarantee a numerically stable behaviour. That is, minor rounding errors due to floating-point arithmetic will not result in major errors in the computed probabilities. As demonstrated by Paris in [22], measuring the difference between probabilistic knowledge bases is subtle and is best addressed by comparing knowledge bases extensionally; i.e., with respect to their model sets. To this end, Paris considered the Blaschke metric. Formally, the *Blaschke distance* $\|S_1, S_2\|_B$ between two convex sets S_1, S_2 is defined by

$$\inf\{\delta \in \mathbb{R} \mid \forall P_1 \in S_1 \exists P_2 \in S_2 : \|P_1, P_2\|_2 \leq \delta \text{ and} \\ \forall P_2 \in S_2 \exists P_1 \in S_1 : \|P_2, P_1\|_2 \leq \delta\}$$

Intuitively, $\|S_1, S_2\|_B$ is the smallest real number d such that for each distribution in one of the sets, there is a probability distribution in the other that has distance at most d to the former. We say that a sequence of knowledge bases (\mathcal{K}_i) converges to a knowledge base \mathcal{K} iff the classical part of each \mathcal{K}_i is equivalent to the classical part of \mathcal{K} and the probabilistic part converges to the probabilistic part of \mathcal{K} . Our reasoning approach behaves indeed continuously with respect to this metric.

Theorem 27 (Continuity). *Let (\mathcal{K}_i) be a convergent sequence of KBs with limit \mathcal{K} and $\mathcal{B}_{\mathcal{K}_i}(C \sqsubseteq D \mid \kappa) = [\ell_i, u_i]$. If $\mathcal{B}_{\mathcal{K}}(C \sqsubseteq D \mid \kappa) = [\ell, u]$, then (ℓ_i) converges to ℓ and (u_i) converges to u (with respect to the usual topology on \mathbb{R}).*

6 Related Work

Relational probabilistic logical approaches can be roughly divided into those that consider probability distributions over the domain, those that consider probability distributions over possible worlds and those that combine both ideas [10]. Our framework belongs to the second group. Maximum entropy reasoning in propositional probabilistic logics has been discussed extensively, e.g., in [13, 22], and various extensions to first-order languages have been considered in recent years [3, 4, 14, 15]. In these works, the domain is restricted to a finite number of constants or bounded in the limit. We circumvent the need to do so by combining a classical first-order logic with unbounded domain with a probabilistic logic with fixed domain.

Many probabilistic DLs have also been considered in the last decades [16, 18, 19]. Our approach is closest to Bayesian DLs [5, 6] and DISPONTE [26]. The greatest difference with the former lies in the fact that \mathcal{ALCP} KBs do not require a complete specification of the probability distribution, but only a set of probabilistic constraints. Moreover, the previous formalisms consider only the sceptical degree of belief, while we are interested in the full belief interval. In contrast to DISPONTE, \mathcal{ALCP} is capable of expressing both, logical and probabilistic

dependencies between the axioms in a KB; in addition, DISPONTE requires all uncertainty degrees to be assigned as mutually independent point probabilities, while \mathcal{ALCP} allows for a more flexible specification.

7 Conclusions

We have introduced the probabilistic DL \mathcal{ALCP} , which extends the classical DL \mathcal{ALC} with the capability of expressing and reasoning about uncertain contextual knowledge defined through the principle of maximum entropy. Effective reasoning methods were developed using the decoupling between the logical and the probabilistic components of \mathcal{ALCP} KBs. We also studied the properties of this logic in relation to other probabilistic logics.

We plan to extend this work in several directions. First, instead of considering the ME-model, we could reason over all probability distributions that satisfy our probabilistic constraints similar to [12, 17, 20]. This will result in larger belief intervals in general. A smaller interval is preferable since it corresponds to a more precise degree of belief. However, when using all probability distributions the size of the interval can be a good indicator for the variation of the possible beliefs in our query with respect to the knowledge base.

In some applications it is also useful to allow more expressive propositional or relational context languages like those proposed in [4, 7, 15, 24]. Similarly, we can consider other DLs for our concept language. Indeed, \mathcal{ALC} was chosen as a prototypical DL for studying the basic properties of our framework. Including additional constructors into the formalism should be relatively simple. In contrast, considering other reasoning problems beyond subsumption is less straightforward. Recall, for instance, that if an \mathcal{ALCP} KB \mathcal{K} contains an inconsistent context with positive probability, then \mathcal{K} has no models. It is thus unclear how to handle the probability of consistency of a KB.

Practical reasoning with \mathcal{ALCP} can be currently performed by combining existing ME-reasoners⁴ with any \mathcal{ALC} -reasoner⁵ according to Algorithm 1. Clearly, such an approach can still be further optimized. We are working on combining the classical and probabilistic reasoning parts in more sophisticated ways.

References

1. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, 2nd edn. (2007)
2. Baader, F., Knechtel, M., Peñaloza, R.: Context-dependent views to axioms and consequences of semantic web ontologies. *J. of Web Semantics* 12–13, 22–40 (2012)
3. Barnett, O., Paris, J.B.: Maximum entropy inference with quantified knowledge. *Logic Journal of the IGPL* 16(1), 85–98 (2008)

⁴ <https://www.fernuni-hagen.de/wbs/research/log4kr/>

⁵ <http://owl.cs.manchester.ac.uk/tools/list-of-reasoners/>

4. Beierle, C., Kern-Isberner, G., Finthammer, M., Potyka, N.: Extending and completing probabilistic knowledge and beliefs without bias. *KI* 29(3), 255–262 (2015)
5. Ceylan, I.I., Peñaloza, R.: The Bayesian description logic BEL. In: Proc. of IJCAR 2014. LNCS, vol. 8562, pp. 480–494. Springer (2014)
6. d’Amato, C., Fanizzi, N., Lukasiewicz, T.: Tractable reasoning with Bayesian description logics. In: Proc. SUM 2008. LNCS, vol. 5291, pp. 146–159. Springer (2008)
7. De Bona, G., Cozman, F.G., Finger, M.: Towards classifying propositional probabilistic logics. *J. of Applied Logic* 12(3), 349–368 (2014)
8. Domingos, P.M., Lowd, D.: Markov Logic: An Interface Layer for Artificial Intelligence. Synthesis Lectures on Artif. Intel. and Machine Learning, Morgan & Claypool Publishers (2009)
9. Donini, F.M., Massacci, F.: ExpTime tableaux for \mathcal{ALC} . *Artif. Intel.* 124(1), 87–138 (2000)
10. Halpern, J.Y.: An analysis of first-order logics of probability. *Artif. Intel.* 46, 311–350 (1990)
11. Halpern, J.Y., Koller, D.: Representation dependence in probabilistic inference. *JAIR* pp. 319–356 (2004)
12. Hansen, P., Perron, S.: Merging the local and global approaches to probabilistic satisfiability. *Intern. J. of Approx. Reasoning* 47(2), 125 – 140 (2008)
13. Kern-Isberner, G.: Conditionals in nonmonotonic reasoning and belief revision. Springer, LNAI 2087 (2001)
14. Kern-Isberner, G., Thimm, M.: Novel semantical approaches to relational probabilistic conditionals. In: Proc. KR 2010. pp. 382–391. AAAI Press (2010)
15. Kern-Isberner, G., Lukasiewicz, T.: Combining probabilistic logic programming with the power of maximum entropy. *Artif. Intell.* 157(1-2), 139–202 (2004)
16. Klinov, P., Parsia, B.: A hybrid method for probabilistic satisfiability. In: Proc. CADE 2011. LNCS, vol. 6803, pp. 354–368. Springer (2011)
17. Lukasiewicz, T.: Probabilistic deduction with conditional constraints over basic events. *JAIR* 10, 380–391 (1999)
18. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the semantic web. *J. of Web Semantics* 6(4), 291–308 (2008)
19. Lutz, C., Schröder, L.: Probabilistic description logics for subjective uncertainty. In: Proc. KR 2010. AAAI Press (2010)
20. Nilsson, N.J.: Probabilistic logic. *Artificial Intelligence* 28, 71–88 (February 1986)
21. Nocedal, J., Wright, S.J.: Numerical Optimization. Springer, 2nd edn. (2006)
22. Paris, J.: The uncertain reasoner’s companion – A mathematical perspective. Cambridge University Press (1994)
23. Peñaloza, R., Potyka, N.: Probabilistic reasoning in the description logic alcp with the principle of maximum entropy (full version). CoRR abs/1606.09521 (2016), <http://arxiv.org/abs/1606.09521>
24. Potyka, N.: Reasoning over linear probabilistic knowledge bases with priorities. In: Proc. SUM 2015. vol. 9310, pp. 121–136. Springer (2015)
25. Potyka, N.: Relationships between semantics for relational probabilistic conditional logics. In: Computational Models of Rationality, Essays dedicated to Gabriele Kern-Isberner. pp. 332–347. College Publications (2016)
26. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R.: Epistemic and statistical probabilistic ontologies. In: Proc. URSW-12. vol. 900, pp. 3–14. CEUR-WS (2012)
27. Schmidt-Schauß, M., Smolka, G.: Attributive concept descriptions with complements. *Artif. Intell.* 48(1), 1–26 (1991)
28. Yeung, R.W.: Information theory and network coding. Springer Science & Business Media (2008)