Counting Strategies for the Probabilistic Description Logic $\mathcal{ALC}^{\mathsf{ME}}$ Under the Principle of Maximum Entropy*

Marco Wilhelm[1][0000−0003−0266−2334], Gabriele Kern-Isberner[1][0000−0001−8689−5391], Andreas Ecke[2][0000−0001−9971−8925], and Franz Baader[2][0000−0002−4049−221X]

[1] Dept. of Computer Science, TU Dortmund, Dortmund, Germany
[2] Dept. of Computer Science, TU Dresden, Dresden, Germany

Abstract. We present $\mathcal{ALC}^{\mathsf{ME}}$, a probabilistic variant of the Description Logic $\mathcal{ALC}$ that allows for representing and processing conditional statements of the form "if E holds, then F follows with probability p" under the principle of maximum entropy. Probabilities are understood as degrees of belief and formally interpreted by the aggregating semantics. We prove that both checking consistency and drawing inferences based on approximations of the maximum entropy distribution is possible in $\mathcal{ALC}^{\mathsf{ME}}$ in time polynomial in the domain size. A major problem for probabilistic reasoning from such conditional knowledge bases is to count models and individuals. To achieve our complexity results, we develop sophisticated counting strategies on interpretations aggregated with respect to the so-called conditional impacts of types, which refine their conditional structure.

Keywords: Probabilistic Description Logics · Aggregating Semantics · Principle of Maximum Entropy · Domain-lifted Inference

1 Introduction

Description Logics [1] are a well-investigated family of logic-based knowledge representation languages that are tailored towards representing *terminological* knowledge, i.e. knowledge about concepts, which can then be used to state *facts* about individuals and objects in a concrete situation. In many application domains, like medicine, knowledge is, however, not always certain, which motivates the development of extensions that can deal with uncertainty. In this paper, we present the *probabilistic Description Logic* $\mathcal{ALC}^{\mathsf{ME}}$, which allows to represent and process uncertain knowledge using conditional statements of the form "if E holds, then F follows with probability p". Probabilities are understood as degrees of belief based on the *aggregating semantics* [9]. This semantic generalizes the statistical interpretation of conditional probabilities by combining it with

* This work was supported by the German Research Foundation (DFG) within the Research Unit FOR 1513 "Hybrid Reasoning for Intelligent Systems".

subjective probabilities based on probability distributions over possible worlds. Basically, in a fixed world $\mathcal{I}$, the conditional $(F|E)$ can be evaluated statistically by considering the number of individuals that verify the conditional (i.e., belong to E and F) and dividing this number by the number of individuals to which the conditional applies (i.e., the elements of E). In the aggregation semantics, this is not done independently for each world. Instead, one first sums up these numbers over all possible worlds, weighted with the probability of the respective world, both in the numerator and in the denominator, and only then divides the resulting sums by each other. The semantics obtained this way therefore mimics statistical probabilities from a subjective point of view. This is in contrast to other approaches for probabilistic Description Logics, which handle either subjective [10] or statistical probabilities [13], or are essentially classical terminologies over probabilistic databases [4].

Due to this combination of statistical and subjective probabilities, the models of $\mathcal{ALC}^{\mathsf{ME}}$-knowledge bases are probability distributions over a set of interpretations that serve as possible worlds. In order to ensure that the possible worlds have the same scope and that counting elements with certain properties leads to well-defined natural numbers, we assume that all the interpretations have the same *fixed finite domain*. However, reasoning on all models of an $\mathcal{ALC}^{\mathsf{ME}}$-knowledge base is not productive due to the vast number of such models. Thus, for reasoning purposes, we select among all models of the knowledge base the distinct model with *maximum entropy* [12]. This MaxEnt distribution is known to be the only model fulfilling some evident common sense principles that can be summarized by the main idea that "essentially similar problems should have essentially similar solutions" [11]. In general, however, the MaxEnt distribution is a real-valued function without a finite, closed-form representation. In fact, from a computational point of view, it is the solution of a nonlinear optimization problem, and thus approximations with values in the rational numbers must be used.

The main result shown in this paper is that all required computations can be done in time polynomial in the chosen domain size. First, we show that checking consistency of $\mathcal{ALC}^{\mathsf{ME}}$-knowledge bases is possible in time polynomial in the domain size. A consistent $\mathcal{ALC}^{\mathsf{ME}}$-knowledge base always has a MaxEnt model. Second, we prove that, once an approximation of this distribution is determined, inferences can be drawn exactly from this approximation, and these inferences can again be computed in time polynomial in the domain size. Investigating the complexity with respect to the domain size is a fundamental problem in probabilistic reasoning as the domain size is usually *the* crucial quantity in application domains. Inferences that can be drawn in time polynomial in the domain size are known as *domain-lifted inferences* [6]. The problem of drawing inferences in a domain-lifted manner is non-trivial since probability distributions are defined over possible worlds, the number of which is exponential in the domain size. Thus, our complexity results require sophisticated strategies of aggregating and counting interpretations. More precisely, we capture the fact that interpretations with the same *conditional structure* [8] have the same impact on the aggregating semantics and the MaxEnt distribution, and we refine the notion of conditional structures of interpretations to *conditional impacts* of *types* [14, 15], which enables the use of efficient counting strategies.

The rest of the paper is organized as follows. In Section 2, we introduce syntax and semantic of the Description Logic $\mathcal{ALC}^{\mathsf{ME}}$. We prove that checking consistency and drawing inferences from approximations of the maximum entropy distribution are possible in $\mathcal{ALC}^{\mathsf{ME}}$ in time polynomial in the domain size in Section 5. For this, we first discuss how interpretations can be aggregated into equivalence classes based on conditional structures and types (Section 3), and then show how these equivalence classes and their cardinalities can be determined efficiently (Section 4).

2 The Description Logic $\mathcal{ALC}^{\mathsf{ME}}$

We present $\mathcal{ALC}^{\mathsf{ME}}$, a probabilistic conditional extension of the terminological part of the Description Logic $\mathcal{ALC}$. The semantics of $\mathcal{ALC}^{\mathsf{ME}}$ is based on the *aggregating semantics* [16] and the *principle of maximum entropy* [12].

Let $\mathcal{N}_C$ and $\mathcal{N}_R$ be disjoint finite sets of concept and role names, respectively. A *concept* is either a concept name or of the form

$$\top, \quad \bot, \quad \neg C, \quad C \sqcap D, \quad C \sqcup D, \quad \exists r.C, \quad \forall r.C,$$

where C and D are concepts and r is a role name. The set of all *subconcepts* of a concept C, i.e. the concepts C is built of, is denoted by $\mathsf{sub}(C)$.

An *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is a tuple consisting of a non-empty set $\Delta^{\mathcal{I}}$ called *domain* and an *interpretation function* $\cdot^{\mathcal{I}}$ that maps every $C \in \mathcal{N}_C$ to a subset $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and every $r \in \mathcal{N}_R$ to a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The interpretation of arbitrary concepts is recursively defined as

- $\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$ and $\bot^{\mathcal{I}} = \emptyset$,
- $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$, $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, and $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$,
- $(\exists r.C)^{\mathcal{I}} = \{a \in \Delta^{\mathcal{I}} \mid \exists b \in \Delta^{\mathcal{I}} : (a, b) \in r^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\}$, and
- $(\forall r.C)^{\mathcal{I}} = \{a \in \Delta^{\mathcal{I}} \mid \forall b \in \Delta^{\mathcal{I}} : (a, b) \in r^{\mathcal{I}} \rightarrow b \in C^{\mathcal{I}}\}$.

Let C, D, E, F be concepts and let $p \in [0, 1]$. An expression of the form $C \sqsubseteq D$ is called a *concept inclusion*, and an expression of the form $(F|E)[p]$ is called a *(probabilistic) conditional*. For computational issues, we assume p to be a rational number. Concept inclusions $C \sqsubseteq D$ represent strict knowledge ("every individual that has property C must also have property D") while conditionals $(F|E)[p]$ act as uncertain beliefs ("if E holds for an individual, then F follows with probability p").

An interpretation $\mathcal{I}$ is a *model* of a concept inclusion, written $\mathcal{I} \models C \sqsubseteq D$, iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. The semantics of conditionals is based on probability distributions over possible worlds. For this, we require a fixed finite domain $\Delta = \Delta^{\mathcal{I}}$ for all interpretations as part of the input. The interpretations serve as possible worlds, thus the fixed finite domain guarantees that all possible worlds have the same scope. We denote the set of all interpretations $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ with $\mathfrak{I}^{\Delta}$ and the set of all probability distributions $\mathcal{P} : \mathfrak{I}^{\Delta} \rightarrow [0, 1]$ with $\mathfrak{P}^{\Delta}$.

Definition 1 (Aggregating Semantics). *A probability distribution* $\mathcal{P} \in \mathfrak{P}^{\Delta}$ *is a* (probabilistic) model *of a concept inclusion* $C \sqsubseteq D$*, written* $\mathcal{P} \models C \sqsubseteq D$*, iff*

$$\mathcal{I} \not\models C \sqsubseteq D \quad \Rightarrow \quad \mathcal{P}(\mathcal{I}) = 0 \qquad \forall \mathcal{I} \in \mathfrak{I}^{\Delta},$$

and of a conditional $(F|E)[p]$*, written* $\mathcal{P} \models (F|E)[p]$*, iff*

$$\frac{\sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} |E^{\mathcal{I}} \cap F^{\mathcal{I}}| \cdot \mathcal{P}(\mathcal{I})}{\sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} |E^{\mathcal{I}}| \cdot \mathcal{P}(\mathcal{I})} = p. \tag{1}$$

Concept inclusions are interpreted as hard constraints in the obvious manner: if a concept inclusion does not hold in an interpretation $\mathcal{I}$, then $\mathcal{I}$ has probability zero. Whether a concept inclusion holds in $\mathcal{I}$ can be decided independently of the probability distribution. The interpretation of conditionals is an adaption of the *aggregating semantics* [16] and needs more explanation. The core idea is to capture the definition of conditional probabilities by a probability-weighted sum of the number of individuals b for which the conditional $(F|E)$ is *verified* (i.e., $b \in |E^{\mathcal{I}} \cap F^{\mathcal{I}}|$) divided by a probability-weighted sum of the number of individuals a for which the conditional is *applicable* (i.e., $a \in |E^{\mathcal{I}}|$). Hence, the aggregating semantics mimics statistical probabilities from a subjective point of view, and probabilities can be understood as degrees of belief in accordance with type 2 probabilities in the classification of Halpern [7].

The aggregating semantics constitutes the main difference to the approaches in [10] and [13]: while there is no probabilistic semantics for terminological knowledge in [10], conditionals are interpreted in [13] purely statistically by the relative frequencies $|E^{\mathcal{I}} \cap F^{\mathcal{I}}|/|E^{\mathcal{I}}|$ in every single interpretation $\mathcal{I}$.

A *knowledge base* $\mathcal{R} = (\mathcal{T}, \mathcal{C})$ consists of a finite set of concept inclusions $\mathcal{T}$ and a finite set of conditionals $\mathcal{C} = \{(F_1|E_1)[p_1], \ldots, (F_n|E_n)[p_n]\}$. Without loss of generality, we make the following assumptions:

1. Knowledge bases contain concepts that are built using the constructors negation ($\neg C$), conjunction ($C \sqcap D$), and existential restriction ($\exists r.C$) only. In addition, we disallow double negation. For the rest of the paper, whenever the negation of an already negated concept is mentioned, we mean the concept itself.
2. Concepts in existential restrictions $\exists r.C$ are concept names. Otherwise, replace C by a fresh concept name A and add $C \sqsubseteq A$ and $A \sqsubseteq C$ to $\mathcal{T}$.
3. Probabilities of conditionals $(F|E)[p] \in \mathcal{C}$ satisfy $0 < p < 1$. This is without loss of generality, because $(F|E)[1]$ and $E \sqsubseteq F$ as well as $(F|E)[p]$ and $(\neg F|E)[1-p]$ (and hence $(F|E)[0]$ and $E \sqsubseteq \neg F$) are semantically equivalent.

We also require the notion of the *signature* of a knowledge base $\mathcal{R}$. In particular, we denote the set of all concept names that are mentioned in $\mathcal{R}$ with $\mathsf{sig}_C(\mathcal{R})$, and the set of all role names that are mentioned in $\mathcal{R}$ with $\mathsf{sig}_R(\mathcal{R})$.

A probability distribution $\mathcal{P} \in \mathfrak{P}^{\Delta}$ is a *model* of a knowledge base $\mathcal{R} = (\mathcal{T}, \mathcal{C})$, written $\mathcal{P} \models \mathcal{R}$, iff it is a model of every concept inclusion in $\mathcal{T}$ and of every conditional in $\mathcal{C}$. A knowledge base with at least one model is called *consistent*.

Knowledge bases with $\mathcal{C} = \emptyset$ are equivalent to $\mathcal{ALC}$-TBoxes (cf. [3]) and allow for classical entailment. In particular, our probabilistic notion of consistency then coincides with the classical one.

Example 1. Consider the following knowledge of an agent. Every person that is generous certainly is wealthy. Otherwise, she would not have anything to spend. And every wealthy person most likely is successful in her career or has a generous patron. Of course, the latter is uncertain as, for example, persons could also become wealthy because of luck in a lottery, etc. Further, wealthy persons typically are not generous. We represent this knowledge by the concept inclusion

$$\mathfrak{c}_1 : \mathsf{Generous} \sqsubseteq \mathsf{Wealthy}$$

and the conditionals

$$\mathfrak{r}_1 : (\neg\mathsf{Successful} \sqcap \neg\exists\mathsf{patron}.\mathsf{Generous}|\mathsf{Wealthy})[0.1],$$
$$\mathfrak{r}_2 : (\neg\mathsf{Generous}|\mathsf{Wealthy})[0.8],$$

and consider the knowledge base $\mathcal{R}_W = (\{\mathfrak{c}_1\}, \{\mathfrak{r}_1, \mathfrak{r}_2\})$ later on. Note that $\mathfrak{r}_1$ is equivalent to the conditional $(\mathsf{Successful} \sqcup \exists\mathsf{patron}.\mathsf{Generous}|\mathsf{Wealthy})[0.9]$.

Consistent probabilistic knowledge bases typically have infinitely many models even for a fixed finite domain. Instead of reasoning w.r.t. all models, it is often more useful to reason w.r.t. a fixed model since reasoning based on the whole set of models leads to monotonic and often uninformative inferences. Any selected model $\mathcal{P}$ yields the inference relation

$$\mathcal{R} \models_{\mathcal{P}} \begin{cases} C \sqsubseteq D & \text{iff} \quad \mathcal{P} \models C \sqsubseteq D, \\ (F|E)[p] & \text{iff} \quad \mathcal{P} \models (F|E)[p]. \end{cases} \tag{2}$$

From a commonsense point of view, the *maximum entropy distribution* is the most appropriate choice of model [11]. For every consistent knowledge base, the maximum entropy distribution exists and is unique.

Definition 2 (Maximum Entropy Distribution). *Let $\mathcal{R}$ be a consistent knowledge base and Δ a fixed domain. The probability distribution*

$$\mathcal{P}^{\mathsf{ME}}_{\mathcal{R}} = \arg\max_{\substack{\mathcal{P} \in \mathfrak{P}^{\Delta} \\ \mathcal{P} \models \mathcal{R}}} - \sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} \mathcal{P}(\mathcal{I}) \cdot \log \mathcal{P}(\mathcal{I}) \tag{3}$$

is called the maximum entropy distribution *(also MaxEnt distribution) of $\mathcal{R}$. In (3), the convention* $0 \cdot \log 0 = 0$ *applies.*

Since it is the solution of a nonlinear optimization problem, the MaxEnt distribution can only be calculated approximately in general. This is typically done by solving the dual optimization problem (cf. [5]), which leads to

$$\mathcal{P}^{\mathsf{ME}}_{\mathcal{R}}(\mathcal{I}) = \begin{cases} \alpha_0 \cdot \prod_{i=1}^{n} \alpha_i^{f_i(\mathcal{I})} & \mathcal{I} \models \mathcal{T}, \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

where, for $i = 1, \ldots, n$, the index i refers to the i-th conditional $(F_i|E_i)[p_i]$ in $\mathcal{C}$, the feature function f_i is defined as $f_i(\mathcal{I}) = |E_i^{\mathcal{I}} \cap F_i^{\mathcal{I}}| - p_i \cdot |E_i^{\mathcal{I}}|$, α_0 is a normalizing constant, and the vector $\boldsymbol{\alpha}_{\mathcal{R}}^{\mathsf{ME}} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{R}_{>0}^n$ is a solution of the system of equations

$$\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} f_i(\mathcal{I}) \cdot \prod_{j=1}^{n} \alpha_j^{f_j(\mathcal{I})} = 0, \qquad i = 1, \ldots, n. \tag{5}$$

Given $\alpha_1, \ldots, \alpha_n$ and the feature functions, the normalization constant α_0 is defined as

$$\alpha_0 = \Big(\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} \prod_{i=1}^{n} \alpha_i^{f_i(\mathcal{I})} \Big)^{-1}. \tag{6}$$

Its rôle is to ensure that a probability distribution is obtained, i.e., that summing up the probabilities of the elements of $\mathfrak{I}^{\Delta}$ yields 1.

The system (5) can, for instance, be solved using Newton's method. Here, we do not investigate this approximation process, but assume that an approximation $\boldsymbol{\beta} \in \mathbb{Q}_{>0}^n$ of $\boldsymbol{\alpha}_{\mathcal{R}}^{\mathsf{ME}}$ is given. Then, $\boldsymbol{\beta}$ defines an approximation of $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ via

$$\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}(\mathcal{I}) = \begin{cases} \beta_0 \cdot \prod_{i=1}^{n} \beta_i^{f_i(\mathcal{I})} & \mathcal{I} \models \mathcal{T}, \\ 0 & \text{otherwise}, \end{cases} \tag{7}$$

where β_0 is a normalizing constant that is defined analogously to (6). It is easy to see that $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$ indeed is a probability distribution. In particular, $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$ is an exact model of $\mathcal{T}$ and of $\mathcal{C}$ up to a deviation depending on the precision of the approximation $\boldsymbol{\beta}$.

3 Conditional Structures and Types for $\mathcal{ALC}^{\mathsf{ME}}$

All kinds of maximum entropy calculations involve sums over interpretations. As the number of interpretations is exponential in $|\Delta|$, evaluating these sums in the naïve way is intractable. In this section, we aggregate interpretations into equivalence classes such that equivalent interpretations have the same impact on the calculations (basically, they have the same MaxEnt probability), while the number of equivalence classes is bounded polynomially in $|\Delta|$.

The *conditional structure* $\sigma_{\mathcal{R}}(\mathcal{I})$ of an interpretation $\mathcal{I}$ with respect to a knowledge base $\mathcal{R} = (\mathcal{T}, \mathcal{C})$ is a formal representation of how often the conditionals in $\mathcal{C}$ are verified and falsified in $\mathcal{I}$ [8]. Mathematically, the conditional structure

$$\sigma_{\mathcal{R}}(\mathcal{I}) = \prod_{i=1}^{n} (\mathbf{a_i^+})^{|E_i^{\mathcal{I}} \cap F_i^{\mathcal{I}}|} \cdot (\mathbf{a_i^-})^{|E_i^{\mathcal{I}} \cap (\neg F_i)^{\mathcal{I}}|} \tag{8}$$

is an element of the free Abelian group that is generated by $\mathfrak{G} = \{\mathbf{a_i^{\pm}} \mid i = 1, \ldots, n,\ \pm \in \{+, -\}\}$. The elements in $\mathfrak{G}$ indicate whether the i-th conditional is verified ($\mathbf{a_i^+}$) or falsified ($\mathbf{a_i^-}$). The frequencies of verification and falsification in $\mathcal{I}$ are respectively indicated by the exponents of $\mathbf{a_i^+}$ and $\mathbf{a_i^-}$.

Example 2. Recall $\mathcal{R}_W$ from Example 1 and consider the interpretation $\mathcal{I}$ in which each $d \in \Delta$ is wealthy ($d \in \mathsf{Wealthy}^{\mathcal{I}}$), successful ($d \in \mathsf{Successful}^{\mathcal{I}}$), but not generous ($d \notin \mathsf{Generous}^{\mathcal{I}}$). Then, $\sigma_{\mathcal{R}_W}(\mathcal{I}) = (\mathbf{a_1^-})^{|\Delta|} \cdot (\mathbf{a_2^+})^{|\Delta|}$.

Conditional structures are important for maximum entropy reasoning as the MaxEnt distribution $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ assigns the same probability to interpretations that are models of $\mathcal{T}$ and have the same conditional structure. The same holds for all approximations of $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ defined by (7) since $\sigma_{\mathcal{R}}(\mathcal{I}) = \sigma_{\mathcal{R}}(\mathcal{I}')$ implies $f_i(\mathcal{I}) = f_i(\mathcal{I}')$ for all $i = 1, \ldots, n$. We now refine conditional structures with respect to the so-called conditional impacts of types.

Definition 3 (Type [2]). *Let $\mathcal{M}$ be a set of concepts such that for every concept $C \in \mathcal{M}$ its negation is also in $\mathcal{M}$ (modulo removal of double negation). A subset τ of $\mathcal{M}$ is a* type *for $\mathcal{M}$ iff*

- *for every $C \in \mathcal{M}$, either C or $\neg C$ belongs to τ, and*
- *for every $C \sqcap D \in \mathcal{M}$ it holds that $C \sqcap D \in \tau$ iff $C, D \in \tau$.*

The set of all types for $\mathcal{M}$ is denoted by $\mathfrak{T}(\mathcal{M})$.

In particular, we are interested in types for a knowledge base $\mathcal{R} = (\mathcal{T}, \mathcal{C})$, i.e. types for $\mathfrak{T}_{\mathcal{R}} = \mathfrak{T}(\mathcal{M}_{\mathcal{R}})$ where $\mathcal{M}_{\mathcal{R}}$ is the closure under negation of the set of subconcepts of concepts occurring in $\mathcal{R}$, i.e.,

$$\mathcal{M}_{\mathcal{R}}^{+} = \bigcup_{C \sqsubseteq D \in \mathcal{T}} \Big(\mathsf{sub}(C) \cup \mathsf{sub}(D)\Big) \cup \bigcup_{(F|E)[p] \in \mathcal{C}} \Big(\mathsf{sub}(E) \cup \mathsf{sub}(F)\Big),$$

and $\mathcal{M}_{\mathcal{R}} = \mathcal{M}_{\mathcal{R}}^{+} \cup \{\neg C \mid C \in \mathcal{M}_{\mathcal{R}}^{+}\}$.

Example 3. There are 16 different types for $\mathcal{R}_W$ from Example 1 (cf. Table 1).

A type τ can be understood as the concept C_τ that is the conjunction of all concepts in τ. If $\tau \neq \tau'$ are different types, then C_τ and $C_{\tau'}$ are disjoint, i.e. $C_\tau^{\mathcal{I}} \cap C_{\tau'}^{\mathcal{I}} = \emptyset$ for all $\mathcal{I} \in \mathfrak{I}^{\Delta}$. Every concept $D \in \mathcal{M}$ can be expressed as a disjunction of such disjoint type concepts [2]:

$$D \equiv \bigsqcup_{\substack{\tau \in \mathfrak{T}(\mathcal{M}) \\ D \in \tau}} C_\tau \quad \text{and} \quad |D^{\mathcal{I}}| = \sum_{\substack{\tau \in \mathfrak{T}(\mathcal{M}) \\ D \in \tau}} |C_\tau^{\mathcal{I}}|. \tag{9}$$

Additionally, the cardinalities $|C_\tau^{\mathcal{I}}|$ of all type concepts $\tau \in \mathfrak{T}(\mathcal{M})$ sum up to $|\Delta|$:

$$\bigsqcup_{\tau \in \mathfrak{T}(\mathcal{M})} C_\tau \equiv \top \quad \text{and} \quad \sum_{\tau \in \mathfrak{T}(\mathcal{M})} |C_\tau^{\mathcal{I}}| = |\Delta|. \tag{10}$$

To prove this, let $\mathcal{I} \in \mathfrak{I}^{\Delta}$ and consider $d \in \Delta$. If we define $\tau = \{D \in \mathcal{M} \mid d \in D^{\mathcal{I}}\}$, then it is easy to see that τ is a type and that $d \in C_\tau^{\mathcal{I}}$. This shows that

τ	$\rho_{\mathcal{R}_W}(\tau)$	$\tau \models \mathfrak{c}_1$?
$\tau_1 = \{S, W, G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^-}\mathbf{a_2^-}$	yes
$\tau_2 = \{S, W, G, \neg\exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^-}\mathbf{a_2^-}$	yes
$\tau_3 = \{S, W, \neg G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^-}\mathbf{a_2^+}$	yes
$\tau_4 = \{S, W, \neg G, \neg\exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^-}\mathbf{a_2^+}$	yes
$\tau_5 = \{S, \neg W, G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{1}$	no
$\tau_6 = \{S, \neg W, G, \neg\exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{1}$	no
$\tau_7 = \{S, \neg W, \neg G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{1}$	yes
$\tau_8 = \{S, \neg W, \neg G, \neg\exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{1}$	yes
$\tau_9 = \{\neg S, W, G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^-}\mathbf{a_2^-}$	yes
$\tau_{10} = \{\neg S, W, G, \neg\exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^+}\mathbf{a_2^-}$	yes
$\tau_{11} = \{\neg S, W, \neg G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^-}\mathbf{a_2^+}$	yes
$\tau_{12} = \{\neg S, W, \neg G, \neg\exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{a_1^+}\mathbf{a_2^+}$	yes
$\tau_{13} = \{\neg S, \neg W, G, \exists p.G, \neg(\neg S \sqcap \neg\exists p.G)\}$	$\mathbf{1}$	no
$\tau_{14} = \{\neg S, \neg W, G, \neg\exists p.G, \neg S \sqcap \neg\exists p.G\}$	$\mathbf{1}$	no
$\tau_{15} = \{\neg S, \neg W, \neg G, \exists p.G, \neg S \sqcap \neg\exists p.G\}$	$\mathbf{1}$	yes
$\tau_{16} = \{\neg S, \neg W, \neg G, \neg\exists p.G, \neg S \sqcap \neg\exists p.G\}$	$\mathbf{1}$	yes

Table 1. Types, their conditional impacts w.r.t. the conditionals in $\mathcal{R}_W$ (cf. Example 1), and their satisfaction behavior w.r.t. the concept inclusion $\mathfrak{c}_1$ in $\mathcal{R}_W$. Concept and role names are abbreviated by their first letter.

$\bigcup_{\tau\in\mathfrak{T}(\mathcal{M})} C_\tau^{\mathcal{I}} \equiv \Delta$, and thus also $|\Delta| = \sum_{\tau\in\mathfrak{T}(\mathcal{M})} |C_\tau^{\mathcal{I}}|$ due to the fact that the type concepts are pairwise disjoint.

As a consequence of (10), types can be seen as characterizations of individuals through the concepts they belong to. Hence, we may say that an individual $d \in \Delta$ is of type τ in the interpretation $\mathcal{I} \in \mathfrak{I}^{\Delta}$ iff $d \in C_\tau^{\mathcal{I}}$, and two individuals are equivalent iff they are of the same type. With this, the conditional structure of interpretations (8) can be broken down to the *conditional impact* of types. We define the *conditional impact* of a type τ for a knowledge base $\mathcal{R}$ by

$$\rho_{\mathcal{R}}(\tau) = \prod_{i=1}^{n} \begin{cases} \mathbf{a_i^+} & \text{iff} \quad E_i, F_i \in \tau \\ \mathbf{a_i^-} & \text{iff} \quad E_i, \neg F_i \in \tau \\ \mathbf{1} & \text{iff} \quad \neg E_i \in \tau \end{cases} .$$

Example 4. The conditional impacts of the types for $\mathcal{R}_W$ from Example 1 are shown in Table 1.

Analogously to the conditional impact of a type, we define the *feature*

$$f_i(\tau) = \begin{cases} 1 - p_i & \text{iff} \quad E_i, F_i \in \tau \\ -p_i & \text{iff} \quad E_i, \neg F_i \in \tau \\ 0 & \text{iff} \quad \neg E_i \in \tau \end{cases} \tag{11}$$

of τ for $i = 1, \ldots, n$.

Proposition 1. *Let $\mathcal{R} = (\mathcal{T}, \mathcal{C})$ be a knowledge base. Then, for all $\mathcal{I} \in \mathfrak{I}^{\Delta}$,*

1. $\sigma_{\mathcal{R}}(\mathcal{I}) = \prod_{\tau \in \mathfrak{T}_{\mathcal{R}}} \rho_{\mathcal{R}}(\tau)^{|C_\tau^{\mathcal{I}}|}$,
2. $f_i(\mathcal{I}) = \sum_{\tau \in \mathfrak{T}_{\mathcal{R}}} |C_\tau^{\mathcal{I}}| \cdot f_i(\tau)$ *for* $i = 1, \ldots, n$.

Proof. To see that 1. holds, note that we have

$$\begin{aligned}
\sigma_{\mathcal{R}}(\mathcal{I}) &= \prod_{i=1}^{n} (\mathbf{a_i^+})^{|E_i^{\mathcal{I}} \cap F_i^{\mathcal{I}}|} \cdot (\mathbf{a_i^-})^{|E_i^{\mathcal{I}} \cap (\neg F_i)^{\mathcal{I}}|} \\
&= \prod_{i=1}^{n} (\mathbf{a_i^+})^{|(\sqcup_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}} \\ E_i, F_i \in \tau}} C_\tau)^{\mathcal{I}}|} \cdot (\mathbf{a_i^-})^{|(\sqcup_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}} \\ E_i, \neg F_i \in \tau}} C_\tau)^{\mathcal{I}}|} \\
&= \prod_{i=1}^{n} (\mathbf{a_i^+})^{\sum_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}} \\ E_i, F_i \in \tau}} |C_\tau^{\mathcal{I}}|} \cdot (\mathbf{a_i^-})^{\sum_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}} \\ E_i, \neg F_i \in \tau}} |C_\tau^{\mathcal{I}}|} \\
&= \prod_{i=1}^{n} \prod_{\tau \in \mathfrak{T}_{\mathcal{R}}} \begin{cases} (\mathbf{a_i^+})^{|C_\tau^{\mathcal{I}}|} & \text{iff} \quad E_i, F_i \in \tau \\ (\mathbf{a_i^-})^{|C_\tau^{\mathcal{I}}|} & \text{iff} \quad E_i, \neg F_i \in \tau \end{cases} \\
&= \prod_{\tau \in \mathfrak{T}_{\mathcal{R}}} \rho_{\mathcal{R}}(\tau)^{|C_\tau^{\mathcal{I}}|}.
\end{aligned}$$

The equations in 2. can be shown using the same arguments. □

Proposition 1 advises one to consolidate interpretations with the same counts $|C_\tau^{\mathcal{I}}|$ for $\tau \in \mathfrak{T}_{\mathcal{R}}$ to equivalence classes. We define $\mathcal{I} \sim_{\mathcal{R}} \mathcal{I}'$ iff $|C_\tau^{\mathcal{I}}| = |C_\tau^{\mathcal{I}'}|$ for all $\tau \in \mathfrak{T}_{\mathcal{R}}$, and obtain the following corollary.

Corollary 1. *Let $\mathcal{R}$ be a knowledge base, and let $\mathcal{I}, \mathcal{I}' \in \mathfrak{I}^{\Delta}$ with $\mathcal{I} \sim_{\mathcal{R}} \mathcal{I}'$.*

1. *Then,* $\sigma_{\mathcal{R}}(\mathcal{I}) = \sigma_{\mathcal{R}}(\mathcal{I}')$ *and* $f_i(\mathcal{I}) = f_i(\mathcal{I}')$ *for* $i = 1, \ldots, n$.
2. *If $\mathcal{R}$ is consistent and additionally $\mathcal{I}$ and $\mathcal{I}'$ are models of $\mathcal{T}$, then*
   (a) $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}(\mathcal{I}) = \mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}(\mathcal{I}')$,
   (b) $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}(\mathcal{I}) = \mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}(\mathcal{I}')$ *for any approximation* $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$ *of* $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ *defined by (7).*

We close this section with a rough estimation of the number of equivalence classes in $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$. These equivalence classes $[\mathcal{I}]_{\sim_{\mathcal{R}}}$ can differ in the numbers $|C_\tau^{\mathcal{I}}|$ for $\tau \in \mathfrak{T}_{\mathcal{R}}$, all of which can vary between zero and $|\Delta|$. Hence, $|\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}|$ is bounded by $(|\Delta| + 1)^{|\mathfrak{T}_{\mathcal{R}}|}$, which is polynomial in $|\Delta|$. Note that this bound is not sharp.

4 Counting Strategies for $\mathcal{ALC}^{\mathsf{ME}}$

We give combinatorial arguments that allow us to compute the equivalence classes in $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$ as well as their cardinalities in time polynomial in $|\Delta|$.

By definition, the equivalence classes $[\mathcal{I}]_{\sim_{\mathcal{R}}} \in \mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$ differ in the *number* of individuals from Δ that have the types $\tau \in \mathfrak{T}_{\mathcal{R}}$, i.e., that belong to $C_\tau^{\mathcal{I}}$. No other properties of these individuals are relevant. Hence, specifying all equivalence classes in $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$ is related to the combinatorial problem of classifying $|\Delta|$-many elements into $|\mathfrak{T}_{\mathcal{R}}|$-many categories. For the rest of the paper let $k = |\Delta|$, $\mathfrak{T}_{\mathcal{R}} = \{\tau_1, \ldots, \tau_m\}$, and $k_i = k(\tau_i) = |C_{\tau_i}^{\mathcal{I}}|$, if it is clear from the context which interpretation $\mathcal{I}$ is considered. Then, $[\mathcal{I}]_{\sim_{\mathcal{R}}}$ is in a one-to-one correspondence with the vector $\boldsymbol{k} = (k_1, \ldots, k_m) \in \mathbb{N}_0^m$, and we may define $[\mathcal{I}]_{\boldsymbol{k}}$ as the unique equivalence class corresponding to $\boldsymbol{k}$. Due to (10), for all $[\mathcal{I}]_{\boldsymbol{k}}$ we have that

$$\sum_{i=1}^{m} k_i = k \tag{12}$$

holds. However, not every vector $\boldsymbol{k} \in \mathbb{N}_0^m$ that satisfies (12) leads to an equivalence class in $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$. This is due to the fact that existential restrictions relate individuals to each other and may force the existence of further individuals of a certain type.

Example 5. Consider the knowledge base

$$\mathcal{R}_{smk} = (\emptyset, \{(\exists \mathsf{friend}.\mathsf{Smoker}|\mathsf{Smoker})[0.9]\}$$

stating that smokers typically have at least one friend that is a smoker, too. There are four types for $\mathcal{R}_{smk}$ (concept and role names are abbreviated by their first letter):

$$t_1 = \{\ \ S, \exists f.S\}, \qquad t_2 = \{\ \ S, \neg\exists f.S\},$$
$$t_3 = \{\neg S, \exists f.S\}, \qquad t_4 = \{\neg S, \neg\exists f.S\}.$$

If there is an individual of type t_3, i.e. a non-smoker who has a friend that smokes, then there must be a second person who is a smoker, i.e., an individual of type t_1 or t_2. Hence, $k_3 > 0$ enforces $k_1 + k_2 > 0$.

To deal with this phenomenon, we adopt the following definition from [2].

Definition 4. *Let τ be a type that contains an existential restriction $\exists r.A$, and let $\neg\exists r.B_1, \ldots, \neg\exists r.B_l$ be all the negated existential restrictions for the role r in τ. A type τ'* satisfies *$\exists r.A$ in τ iff $A, \neg B_1, \ldots, \neg B_l \in \tau'$.*

It is now easy to see that, for every type $\tau \in \mathfrak{T}_{\mathcal{R}}$ and for every existential restriction $\exists r.A \in \tau$,

$$k(\tau) = 0 \quad \text{or} \quad \sum_{\substack{\tau' \in \mathfrak{T}_{\mathcal{R}} \\ \tau' \text{ satisfies } \exists r.A \text{ in } \tau}} k(\tau') > 0 \tag{13}$$

must hold. Conversely, using ideas from [2], it is not hard to show that any vector $\boldsymbol{k}$ satisfying $\sum_{i=1}^{m} k_i = k$ and (13) is realized by an interpretation. Thus, we have

$$\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}} = \{[\mathcal{I}]_{\boldsymbol{k}} \mid \boldsymbol{k} \in \mathbb{N}_0^m,\ \sum_{i=1}^{m} k_i = k,\ \text{and (13) holds}\}. \tag{14}$$

Equation (14) allows us to enumerate the equivalence classes in $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$ in time polynomial in $|\Delta|$, as Condition (13) is independent of Δ and iterating through all $\boldsymbol{k} \in \mathbb{N}_0^m$ that satisfy $\sum_{i=1}^{m} k_i = k$ is possible in time polynomial in $|\Delta|$. Furthermore, note that we are interested in only those interpretations that satisfy all concept inclusions in $\mathcal{T}$. In these interpretations there must not exist an individual $d \in \Delta$ with $d \in C^{\mathcal{I}}$ and $d \notin D^{\mathcal{I}}$ for any $C \sqsubseteq D \in \mathcal{T}$. Due to (9) and (10), this constraint is equivalent to

$$C, \neg D \in \tau \quad \Rightarrow \quad k(\tau) = 0 \qquad \forall \tau \in \mathfrak{T}_{\mathcal{R}},\ C \sqsubseteq D \in \mathcal{T}. \tag{15}$$

We say that a type $\tau \in \mathfrak{T}_{\mathcal{R}}$ for which $C \in \tau$ implies $D \in \tau$ for all $C \sqsubseteq D \in \mathcal{T}$ *satisfies* $\mathcal{T}$, written $\tau \models \mathcal{T}$. Hence, (15) states that $k(\tau) > 0$ holds for only those types that satisfy $\mathcal{T}$. Consequently, the set of all equivalence classes of those interpretations that satisfy $\mathcal{T}$ is

$$\mathfrak{I}^{\Delta}_{\mathcal{T}}/\sim_{\mathcal{R}} = \{[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}^{\Delta}/\sim_{\mathcal{R}} \mid (15) \text{ holds}\}$$

and can be determined in time polynomial in $|\Delta|$, too.

Example 6. Recall $\mathcal{R}_W$ from Example 1. All types $\tau \in \mathfrak{T}_{\mathcal{R}_W}$ satisfy $\mathcal{T}_W$ except for τ_5, τ_6, τ_{13}, and τ_{14} (cf. Table 1).

It still remains to determine the cardinalities $|[\mathcal{I}]_{\boldsymbol{k}}|$. These cardinalities depend on two factors. First, the k individuals in Δ have to be allocated to the types for $\mathcal{R}$. This is the combinatorial problem of classifying elements into categories mentioned at the beginning of this section, and for which there are

$$\binom{k}{k_1, \dots, k_m} = \frac{k!}{k_1! \cdots k_m!}$$

many possibilities if $k_i = |\tau_i|$ for every $\tau_i \in \mathfrak{T}_{\mathcal{R}}$. Put differently, one can also say that this is the task of specifying $C_\tau^{\mathcal{I}}$ for every $\tau \in \mathfrak{T}_{\mathcal{R}}$ when previously only the cardinalities $|C_\tau^{\mathcal{I}}|$ were known.

Second, once this allocation is given, the sets $\mathcal{C}_\tau^{\mathcal{I}}$ for every $\tau \in \mathfrak{T}_{\mathcal{R}}$ still do not determine a unique interpretation. There remains some degree of freedom when picking a single interpretation from $[\mathcal{I}]_{\boldsymbol{k}}$. To see this, recall that an interpretation $\mathcal{I} \in \mathfrak{I}^{\Delta}$ is fully specified iff for all concept names $C \in \mathcal{N}_C$ and for all role names $r \in \mathcal{N}_R$ the sets $C^{\mathcal{I}}$ and $r^{\mathcal{I}}$ are fixed. As every concept name A that is mentioned in $\mathcal{R}$ also occurs in every single type for $\mathcal{R}$ as either A or $\neg A$, the sets $A^{\mathcal{I}}$ for these concept names are uniquely determined by the types. However, this does not hold for concept names that are not mentioned in $\mathcal{R}$. Actually, given a concept name in $\mathcal{N}_C \setminus \mathsf{sig}_C(\mathcal{R})$, one can choose freely for every individual in Δ whether it belongs to this concept name or not. There are $2^{k \cdot |\mathcal{N}_C \setminus \mathsf{sig}_C(\mathcal{R})|}$ possibilities of allocating the k individuals in Δ to the concepts in $\mathcal{N}_C \setminus \mathsf{sig}_C(\mathcal{R})$. Determining the degree of freedom that arises from role memberships is more difficult. Again, for the roles that are not mentioned in $\mathcal{R}$, there is free choice such that there are $2^{k^2 \cdot |\mathcal{N}_R \setminus \mathsf{sig}_R(\mathcal{R})|}$ possible combinations of allocating k^2 many tuples of individuals to them. For the membership to roles that are mentioned in $\mathcal{R}$, we first define the degree of freedom of a role and discuss it afterwards.

Definition 5. *Let $\mathcal{R} = (\mathcal{T}, \mathcal{C})$ be a knowledge base, and let $\tau \in \mathfrak{T}_{\mathcal{R}}$ be a type. Further let $\exists r.A_1, \ldots, \exists r.A_l$ be all the existential restrictions and let $\neg\exists r.B_1, \ldots, \neg\exists r.B_h$ be all the negated existential restrictions for the role r in τ. We define the* degree of freedom *of r in τ with respect to $[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$ as*

$$\phi_{\boldsymbol{k}}(\tau, r) = \Big(\sum_{\mathfrak{I} \subseteq \{1,\ldots,l\}} (-1)^{|\mathfrak{I}|} \cdot \prod_{\substack{j=1,\ldots,m, \\ \neg B_1,\ldots,\neg B_h \in \tau_j, \\ \neg A_i \in \tau_j \ \forall i \in \mathfrak{I}}} 2^{k_j} \Big)^{k(\tau)}. \tag{16}$$

Definition 5 is a generalization of Definition 4 that takes counting aspects into account by making use of the well-known inclusion-exclusion principle. In this way, it keeps track of which individual guarantees that a certain existential restriction holds. To understand the good behavior of Definition 5, assume that there is *no* positive existential restriction $\exists r.A$ for r in τ. Then, for every $d \in A_{\tau}^{\mathcal{I}}$ and for every individual d' in any $A_{\tau'}^{\mathcal{I}}$ with $\neg B_1, \ldots, \neg B_h \in \tau'$, whether $(d, d') \in r^{\mathcal{I}}$ or not can be chosen freely, which results in the factor $(2^{k(\tau')})^{k(\tau)}$ in $\phi_{\boldsymbol{k}}(\tau, r)$. Now, assume there is one (positive) existential restriction $\exists r.A$ in τ. For individuals $d' \in \tau'$ with τ' such that $\neg A, \neg B_1, \ldots, \neg B_h \in \tau'$, again the belonging of (d, d') to $r^{\mathcal{I}}$ is optional. However, there must be at least one individual d'' among the individuals of a type τ'' with $A, \neg B_1, \ldots, \neg B_h \in \tau''$ such that $(d, d'') \in r^{\mathcal{I}}$. This results in the degree of freedom

$$\begin{aligned} \phi_{\boldsymbol{k}}(\tau, r) = \Big(& \prod_{\substack{\tau' \in \mathfrak{T}_{\mathcal{R}} \\ \neg B_1,\ldots,\neg B_h \in \tau'}} 2^{k(\tau')} - \prod_{\substack{\tau' \in \mathfrak{T}_{\mathcal{R}} \\ \neg B_1,\ldots,\neg B_h, \neg A \in \tau'}} 2^{k(\tau')} \Big)^{k(\tau)} \\ = \Big(& (\prod_{\substack{\tau' \in \mathfrak{T}_{\mathcal{R}} \\ \neg B_1,\ldots,\neg B_h, \neg A \in \tau'}} 2^{k(\tau')}) \cdot (\prod_{\substack{\tau' \in \mathfrak{T}_{\mathcal{R}} \\ \neg B_1,\ldots,\neg B_h, A \in \tau'}} 2^{k(\tau')} - 1) \Big)^{k(\tau)}. \end{aligned}$$

If there are more than one (positive) existential restrictions, then all of them could be satisfied by the same tuple of individuals. Alternatively, there may exist several tuples of individuals each satisfying only some of the restrictions. Then, a combination of tuples is needed to satisfy all of the existential restrictions. This makes the application of the inclusion-exclusion principle necessary.

Altogether, for every $[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$, one has

$$\begin{aligned} |[\mathcal{I}]_{\boldsymbol{k}}| = & \binom{k}{k_1, \ldots, k_m} \cdot \Big(\prod_{j=1}^{m} \prod_{r \in \mathsf{sig}_R(\mathcal{R})} \phi_{\boldsymbol{k}}(\tau_j, r) \Big) \\ & \cdot 2^{(|\mathcal{N}_C \setminus \mathsf{sig}_C(\mathcal{R})|) \cdot k} \cdot 2^{(|\mathcal{N}_R \setminus \mathsf{sig}_R(\mathcal{R})|) \cdot k^2}, \end{aligned} \tag{17}$$

which can be calculated in time polynomial in $|\Delta|$.

5 Consistency Check and Drawing Inferences in $\mathcal{ALC}^{\mathsf{ME}}$

We build upon the results from Section 3 and Section 4 and prove that both checking consistency and drawing inferences from approximations of the MaxEnt distribution is possible in $\mathcal{ALC}^{\mathsf{ME}}$ in time polynomial in $|\Delta|$.

Proposition 2. *Let $\mathcal{R}$ be a knowledge base and Δ a finite domain. Then, consistency of $\mathcal{R}$ in a probabilistic model with domain Δ can be checked in time polynomial in $|\Delta|$.*

Proof. The knowledge base $\mathcal{R} = (\mathcal{T}, \mathcal{C})$ is consistent iff there is a probability distribution $\mathcal{P} \in \mathfrak{P}^{\Delta}$ such that $\mathcal{I} \not\models \mathcal{T}$ implies $\mathcal{P}(\mathcal{I}) = 0$ for all $\mathcal{I} \in \mathfrak{I}^{\Delta}$, and

$$\frac{\sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} |E_i^{\mathcal{I}} \cap F_i^{\mathcal{I}}| \cdot \mathcal{P}(\mathcal{I})}{\sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} |E_i^{\mathcal{I}}| \cdot \mathcal{P}(\mathcal{I})} = p_i, \quad i = 1, \ldots, n.$$

Alternatively, $\mathcal{R}$ is consistent iff the MaxEnt distribution $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ exists. Hence, it is legitimate to limit the search space to any subset of $\mathfrak{P}^{\Delta}$ that contains $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ when searching for a model of $\mathcal{R}$. Thus, it is sufficient to search for a model of $\mathcal{R}$ that satisfies $\mathcal{P}(\mathcal{I}) = \mathcal{P}(\mathcal{I}')$ if $\mathcal{I} \sim_{\mathcal{R}} \mathcal{I}'$ and $\mathcal{I}, \mathcal{I}' \models \mathcal{T}$, like $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ does. In other words, it is sufficient to find a probability distribution $\mathcal{P} : \mathfrak{I}_{\mathcal{T}}^{\Delta}/\sim_{\mathcal{R}} \to [0, 1]$ that satisfies

$$\frac{\sum_{[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}_{\mathcal{T}}^{\Delta}/\sim_{\mathcal{R}}} \left(\sum_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}} \\ E_i, F_i \in \tau}} k(\tau) \right) \cdot \mathcal{P}([\mathcal{I}]_{\boldsymbol{k}})}{\sum_{[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}_{\mathcal{T}}^{\Delta}/\sim_{\mathcal{R}}} \left(\sum_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}} \\ E_i \in \tau}} k(\tau) \right) \cdot \mathcal{P}([\mathcal{I}]_{\boldsymbol{k}})} = p_i, \quad i = 1, \ldots, n. \tag{18}$$

Then, $\mathcal{P}$ can be extended to a probability distribution on $\mathfrak{I}^{\Delta}$ and thereby to a model of $\mathcal{R}$ by defining for all $\mathcal{I} \in \mathfrak{I}^{\Delta}$

$$\mathcal{P}(\mathcal{I}) = \begin{cases} \mathcal{P}([\mathcal{I}]_{\boldsymbol{k}}) \cdot (|[\mathcal{I}]_{\boldsymbol{k}}|)^{-1} & [\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}_{\mathcal{T}}^{\Delta}/\sim_{\mathcal{R}} \\ 0 & \text{otherwise.} \end{cases},$$

The equations in (18) and the conditions $0 \leq \mathcal{P}([\mathcal{I}]_{\boldsymbol{k}}) \leq 1$ for all $[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}_{\mathcal{T}}^{\Delta}/\sim_{\mathcal{R}}$ can easily be transformed into a system of linear inequalities with integer coefficients. Both the number of inequalities and the number of variables of this system is in $\mathcal{O}(|\mathfrak{I}_{\mathcal{T}}^{\Delta}/\sim_{\mathcal{R}}|)$ and, hence, polynomially bounded in $|\Delta|$. It follows, that satisfiability of this system can be decided in time polynomial in $|\Delta|$. □

Proposition 3. *Let $\mathcal{R}$ be a consistent knowledge base, $\boldsymbol{\beta} \in \mathbb{Q}_{>0}^{n}$, and let C, D, E, F be concepts.*

1. *Calculating the probability p for which $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}} \models (F|E)[p]$ holds, and*
2. *deciding whether $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}} \models C \sqsubseteq D$ holds*

is possible in time polynomial in $|\Delta|$.

Proof. As $\mathcal{P}^{\boldsymbol{\beta}}_{\mathcal{R}} \models C \sqsubseteq D$ iff $\mathcal{P}^{\boldsymbol{\beta}}_{\mathcal{R}} \models (D|C)[1]$, the second statement of the proposition follows from the first. To prove the first statement, we write $p_i = \frac{s_i}{t_i}$ with $s_i, t_i \in \mathbb{N}_{>0}$ for $i = 1, \dots, n$, and $q = (F|E)[p]$. Then,

$$p = \frac{\sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} |E^{\mathcal{I}} \cap F^{\mathcal{I}}| \cdot \mathcal{P}^{\boldsymbol{\beta}}_{\mathcal{R}}(\mathcal{I})}{\sum_{\mathcal{I} \in \mathfrak{I}^{\Delta}} |E^{\mathcal{I}}| \cdot \mathcal{P}^{\boldsymbol{\beta}}_{\mathcal{R}}(\mathcal{I})} = \frac{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} |E^{\mathcal{I}} \cap F^{\mathcal{I}}| \cdot \beta_0 \cdot \prod_{i=1}^{n} \beta_i^{f_i(\mathcal{I})}}{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} |E^{\mathcal{I}}| \cdot \beta_0 \cdot \prod_{i=1}^{n} \beta_i^{f_i(\mathcal{I})}}$$

$$= \frac{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} |E^{\mathcal{I}} \cap F^{\mathcal{I}}| \cdot \prod_{i=1}^{n} \beta_i^{|E^{\mathcal{I}} \cap F^{\mathcal{I}}| - \frac{s_i}{t_i} \cdot |E^{\mathcal{I}}|}}{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} |E^{\mathcal{I}}| \cdot \prod_{i=1}^{n} \beta_i^{|E^{\mathcal{I}} \cap F^{\mathcal{I}}| - \frac{s_i}{t_i} \cdot |E^{\mathcal{I}}|}}$$

$$= \frac{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} |E^{\mathcal{I}} \cap F^{\mathcal{I}}| \cdot \prod_{i=1}^{n} \beta_i^{t_i \cdot |E^{\mathcal{I}} \cap F^{\mathcal{I}}| - s_i \cdot |E^{\mathcal{I}}| + s_i \cdot |\Delta|}}{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} |E^{\mathcal{I}}| \cdot \prod_{i=1}^{n} \beta_i^{t_i \cdot |E^{\mathcal{I}} \cap F^{\mathcal{I}}| - s_i \cdot |E^{\mathcal{I}}| + s_i \cdot |\Delta|}}. \tag{19}$$

Note that

$$t_i \cdot |E_i^{\mathcal{I}} \cap F_i^{\mathcal{I}}| - s_i \cdot |E_i^{\mathcal{I}}| + s_i \cdot |\Delta| \geq 0 \qquad \forall \mathcal{I} \in \mathfrak{I}^{\Delta},\ i = 1, \dots, n.$$

Hence, the last fraction in (19) mentions sums over products of integers ($|E^{\mathcal{I}} \cap F^{\mathcal{I}}|$ and $|E^{\mathcal{I}}|$, respectively) and rational numbers (β_i) with integer exponents only and can be computed exactly.

It remains to show that (19) can be calculated in time polynomial in $|\Delta|$. To prove this, we aggregate interpretations into equivalence classes as discussed in Section 3. However, we have to modify the set of types the equivalence classes are based on since the query conditional q may mention additional subconcepts that are not considered by the types in $\mathfrak{T}_{\mathcal{R}}$. Let $\mathcal{M}_q^+ = \{C \mid C \in \mathsf{sub}(E) \cup \mathsf{sub}(F)\}$, $\mathcal{M}_q = \mathcal{M}_q^+ \cup \{\neg C \mid C \in \mathcal{M}_q^+\}$, and $\mathfrak{T}_{\mathcal{R}}^q = \mathfrak{T}(\mathcal{M}_{\mathcal{R}} \cup \mathcal{M}_q)$. For interpretations $\mathcal{I}, \mathcal{I}' \in \mathfrak{I}^{\Delta}$, we define the equivalence relation $\mathcal{I} \sim_{\mathcal{R}}^q \mathcal{I}'$ iff $C_\tau^{\mathcal{I}} = C_\tau^{\mathcal{I}'}$ for all $\tau \in \mathfrak{T}_{\mathcal{R}}^q$ in analogy to $\sim_{\mathcal{R}}$. Every type $\tau \in \mathfrak{T}_{\mathcal{R}}^q$ is a refinement of a unique type $\tau' \in \mathfrak{T}_{\mathcal{R}}$, i.e. $\tau' \subseteq \tau$, and we may define $\rho_{\mathcal{R}}(\tau') = \rho_{\mathcal{R}}(\tau)$. In plain words, τ' inherits its conditional impact from τ. Accordingly, we define $f_i(\tau') = f_i(\tau)$ for $i = 1, \dots, n$. Then Proposition 1 as well as Corollary 1 still hold when replacing the underlying set of types $\mathfrak{T}_{\mathcal{R}}$ by $\mathfrak{T}_{\mathcal{R}}^q$. Also, the counting strategies and the complexity results for $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}^q$ are the same as for $\mathfrak{I}^{\Delta}/\sim_{\mathcal{R}}$. Hence, (19) can be simplified to

$$p = \frac{\sum_{[\mathcal{I}]_{\boldsymbol{k}} \in \mathfrak{I}^{\Delta}_{\mathcal{T}}/\sim_{\mathcal{R}}^q} k_i^+ \cdot \prod_{i=1}^{n} \beta_i^{t_i \cdot k_i^+ - s_i \cdot k_i^o + s_i \cdot |\Delta|}}{\sum_{\substack{\mathcal{I} \in \mathfrak{I}^{\Delta} \\ \mathcal{I} \models \mathcal{T}}} k_i^o \cdot \prod_{i=1}^{n} \beta_i^{t_i \cdot k_i^+ - s_i \cdot k_i^o + s_i \cdot |\Delta|}}$$

where $k_i^+ = \sum_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}}^q \\ E_i, F_i \in \tau}} k(\tau)$ and $k_i^o = \sum_{\substack{\tau \in \mathfrak{T}_{\mathcal{R}}^q \\ E_i \in \tau}} k(\tau)$. This fraction can clearly be calculated in time polynomial in $|\Delta|$. $\square$

Proposition 3 states that inferences in $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$ are domain-lifted (cf. [6]). Hence, the message of Proposition 3 is that the crucial part of drawing inferences at maximum entropy from $\mathcal{R}$ according to (2) is the approximation of $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ by $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$. Once this approximation is given, all further calculations can be performed *exactly* without additional inaccuracies and in time polynomial in $|\Delta|$.

6 Conclusion and Future Work

We have presented $\mathcal{ALC}^{\mathsf{ME}}$, a probabilistic variant of the Description Logic $\mathcal{ALC}$, which allows one to express uncertain knowledge by probabilistic conditional statements of the form "if E holds, then F is true with probability p". Probabilities are understood as degrees of beliefs and a reasoner's belief state is established by the principle of maximum entropy based on the aggregating semantics. We have proved that both checking consistency and drawing inferences from approximations of the maximum entropy distribution is possible in $\mathcal{ALC}^{\mathsf{ME}}$ in time polynomial in the domain size $|\Delta|$.

In future work, we want to investigate, on the one hand, complexity results for approximate inference at maximum entropy in $\mathcal{ALC}^{\mathsf{ME}}$. For this, we need error estimations and complexity results for calculating approximations $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$ of the maximum entropy distribution $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ in addition to the results presented here. Note that the size of the equation system that is generated as input for the methods used to approximate $\mathcal{P}_{\mathcal{R}}^{\mathsf{ME}}$ by $\mathcal{P}_{\mathcal{R}}^{\boldsymbol{\beta}}$ (cf. (5)) can be bounded polynomially in $|\Delta|$, using the same counting strategies as presented in Section 4.

On the other hand, we want to extend our complexity results to more general $\mathcal{ALC}^{\mathsf{ME}}$-knowledge bases containing also assertional knowledge, and to Description Logics that are more expressive than $\mathcal{ALC}$.

Finally, we intend to make a more fine-grained complexity analysis that investigates the complexity of reasoning not only w.r.t. the domain size, but also in terms of the size of the knowledge base.

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 (2003)
2. Baader, F., Ecke, A.: Extending the description logic ALC with more expressive cardinality constraints on concepts. In: Proceedings of the 3rd Global Conference on Artificial Intelligence (GCAI). pp. 6–19. EasyChair (2017)
3. Baader, F., Horrocks, I., Sattler, U.: Description Logics. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) Handbook of Knowledge Representation. Elsevier (2007)
4. Baader, F., Koopmann, P., Turhan, A.: Using ontologies to query probabilistic numerical data. In: Proceedings of the 11th International Symposium on Frontiers of Combining Systems (FroCoS). Lecture Notes in Computer Science, vol. 10483, pp. 77–94. Springer (2017)

5. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press (2004)
6. Van den Broeck, G., Taghipour, N., Meert, W., Davis, J., De Raedt, L.: Lifted probabilistic inference by first-order knowledge compilation. In: Proceedings of the 22th International Joint Conference on Artificial Intelligence (IJCAI). pp. 2178–2185. AAAI Press (2011)
7. Halpern, J.Y.: An analysis of first-order logics of probability. Artificial Intelligence **46**(3), 311–350 (1990)
8. Kern-Isberner, G.: Conditionals in Nonmonotonic Reasoning and Belief Revision. Springer (2001)
9. Kern-Isberner, G., Thimm, M.: Novel semantical approaches to relational probabilistic conditionals. In: Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning (KR). pp. 382–392. AAAI Press (2010)
10. Lutz, C., Schröder, L.: Probabilistic description logics for subjective uncertainty. In: Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR). pp. 393–403. AAAI Press (2010)
11. Paris, J.B.: Common sense and maximum entropy. Synthese **117**(1), 75–93 (1999)
12. Paris, J.B.: The Uncertain Reasoner's Companion: A Mathematical Perspective. Cambridge University Press (2006)
13. Peñaloza, R., Potyka, N.: Towards statistical reasoning in description logics over finite domains. In: Proceedings of the 11th International Conference on Scalable Uncertainty Management (SUM). Lecture Notes in Computer Science, vol. 10564, pp. 280–294. Springer (2017)
14. Pratt, V.R.: Models of program logics. In: Proceedings of the 20th Annual Symposium on Foundations of Computer Science (FOCS). pp. 115–122. IEEE Computer Society (1979)
15. Rudolph, S., Krötzsch, M., Hitzler, P.: Type-elimination-based reasoning for the description logic SHIQbs using decision diagrams and disjunctive datalog. Logical Methods in Computer Science **8**(1) (2012)
16. Thimm, M., Kern-Isberner, G.: On probabilistic inference in relational conditional logics. Logic Journal of the IGPL **20**(5), 872–908 (2012)