

Answering Fuzzy Queries over Fuzzy DL-Lite Ontologies

Gabriella PASI, Rafael PEÑALOZA

University of Milano-Bicocca, Italy

(e-mail: {gabriella.pasi,rafael.penaloz}@unimib.it)

submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003

Abstract

A prominent problem in knowledge representation is how to answer queries taking into account also the implicit consequences of an ontology representing domain knowledge. While this problem has been widely studied within the realm of description logic ontologies, it has been surprisingly neglected within the context of vague or imprecise knowledge, particularly from the point of view of mathematical fuzzy logic. In this paper we study the problem of answering conjunctive queries and threshold queries w.r.t. ontologies in fuzzy DL-Lite. Specifically, we show through a rewriting approach that threshold query answering w.r.t. consistent ontologies remains in AC^0 in data complexity, but that conjunctive query answering is highly dependent on the selected triangular norm, which has an impact on the underlying semantics. For the idempotent Gödel t -norm, we provide an effective method based on a reduction to the classical case.

1 Introduction

Description logics (DLs) (Baader et al. 2007) are a well-known and widely used family of knowledge representation formalisms that, thanks to their clear syntax and formal semantics, have been used to represent and deal with the knowledge of various representation domains. Among the many members of this family, a subfamily of languages with a limited expressivity, known as the DL-Lite family (Calvanese et al. 2007) has a prominent role. In fact, simplifying a bit, DL-Lite was originally designed with the goal of including background knowledge within the task of answering queries, and avoiding the need for an explicit enumeration of all the facts that are implicitly implied by the domain knowledge.

Consider for example a touristic scenario, which includes information about museums, monuments, restaurants, and pubs. Knowing that museums and monuments are touristic attractions, and that restaurants and pubs are eateries, one can immediately deduce that the modern art museum and the peace monument are touristic attractions, and that the Irish pub is an eatery, without having to make this knowledge explicit. A user may thus ask for e.g., a *tourist attraction that contains an eatery*. Using classical query answering techniques (Ortiz and Šimkus 2012), all attractions that satisfy this requirement can be efficiently retrieved.

Being based on classical logic, DLs in general and DL-Lite in particular are unable to handle imprecise or vague knowledge effectively. In our touristic scenario, for instance, we may want to extend the knowledge with some additional properties of the objects of interest. For example, a tourist in a hurry may want to visit the *popular* attractions first;

or a backpacker on a budget may be more interested in finding *cheap* eateries. Note that *cheap* and *popular* are two vague notions that do not allow for any precise definition. In a simplistic scenario, cheapness may be defined in terms of the mean cost for a meal, but even then, it is impossible to specify a precise price-point where an eatery stops being cheap; moreover this is also a subjective notion. The case of popularity is even worse, as there is no obvious proxy for it.

To solve this issue, fuzzy extensions of DLs have been widely studied; see for example (Borgwardt 2014; Borgwardt and Peñaloza 2017; Bobillo et al. 2015; Lukasiewicz and Straccia 2008; Cerami 2012) and references therein. In essence, fuzzy logic (Hájek 1998) extends classical logic by allowing truth degrees in the interval $[0, 1]$ for the propositions that contain fuzzy (or vague) predicates. One can thus say, e.g., that the modern art museum is popular to a degree of 0.8 meaning, intuitively, that it is popular, but more popular attractions may exist.

Interestingly, although fuzzy DLs and their reasoning services have been widely studied, the task of answering queries based on fuzzy ontologies has been mostly ignored. Most of the earlier work from this point of view was carried out by Straccia and Pan. Specifically, Straccia (Straccia 2006) studied the problem of computing the answers with highest degree on a query w.r.t. some background knowledge. This was followed by Pan et al. (Pan et al. 2007), who considered more complex queries to be answered. While from some perspective these works seem to cover the whole area of query answering, they were based on the so-called Zadeh semantics, which does not have adequate properties from a mathematical logic point of view (Hájek 1998). Another limitation of all these approaches is that they allowed only the facts in the ontology to be graded, but restricted the terminological knowledge to be crisp (i.e., hold fully). Other work considering query answering in fuzzy DLs includes (Straccia 2012), where the k answers with the highest degree are retrieved. This latter work is closer to our approach but has several limitations. Perhaps the most obvious is that its semantics follows a closed-world assumption, even in the case of background knowledge. In addition, background knowledge is interpreted as a *rule*, where the degrees of the body of an axiom define the degree of the head, but knowledge about the head cannot be used to infer knowledge about the body. We, in change, use the open world assumption, as typical in knowledge representation, and use the logical interpretation of axioms.

Later on, Turhan and Mailis studied the problem of query answering w.r.t. background knowledge from the point of view of fuzzy logic (Hájek 1998), where the semantics are based on the properties of continuous triangular norms (Klement et al. 2000). They developed a technique for computing the satisfaction degrees of conjunctive queries when the semantics were based on the Gödel t-norm (Mailis and Turhan 2014). This technique, which is based on the construction of a classical query, was later implemented and shown to be effective in (Mailis et al. 2015). However, it still suffered from two main drawbacks: (i) it was only capable to handle the idempotent (Gödel) t-norm, and (ii) terminological knowledge had to still be precise, allowing no graded axioms. The latter condition is essential for the correctness of their approach: their reduction is unable to keep track of the degrees used by the terminological axioms, as this would require an unbounded memory use.

In this paper, we study the problem of query answering w.r.t. DL-Lite ontologies, filling out the gaps left by the existing work. To be more explicit, our work is the first to consider

adequate semantics from the mathematical fuzzy logic point of view, alongside graded axioms stating vague knowledge beyond just vague data. We start by considering the kind of conjunctive queries studied by Turhan and Mailis, but allowing the association of numeric degrees also in the TBox. Interestingly, although this is a generalization of the previously studied setting, we are able to develop a much simpler method, which does not rely on rewriting, but rather on a reduction to a classical query answering scenario. The method is based on the idea of *cut ontologies*, where all knowledge holding to a low degree is ignored. Hence, we obtain a more robust and easier to maintain approach than previous work. Still considering the Gödel t-norm, we considered the case of threshold queries, also left open in previous work, in which every conjunct in a query is assigned a different degree. In this case, a direct reduction to classical query answering does not work, but we were able to adapt the classical rewriting methods to handle the degrees effectively.

The final part of the paper considers other t-norms as the underlying semantics for the fuzzy constructors. In this case, we show through several examples that conjunctive queries cannot be easily handled, but we identify some special cases where queries can be effectively answered. On the other hand, we show that we can still apply the rewriting technique to answer threshold queries, even for non-idempotent t-norms. This is a surprising result because in the idempotent scenario threshold queries are a generalization of conjunctive queries.

Some of the results in this paper were previously published in (Pasi and Peñaloza 2020). In addition to full proofs, deeper explanations, and examples, here we extend that previous work by handling threshold queries, including the full rewriting technique from Section 5. We also provide better results for non-idempotent t-norms, and highlight some of the problems of combining conjunctions and non-idempotent t-norms in the appendix.

2 Preliminaries

We briefly introduce the syntax and semantics of fuzzy DL-Lite_R and other related notions that will be important for this paper. Let N_C , N_R , and N_I be three mutually disjoint sets whose elements are called *concept names*, *role names*, and *individual names*, respectively. The sets of DL-Lite_R *concepts* and *roles* are built through the grammar rules:

$$\begin{array}{ll} B ::= A \mid \exists Q & C ::= B \mid \neg B \\ Q ::= P \mid P^- & R ::= Q \mid \neg Q \end{array}$$

where $A \in N_C$ and $P \in N_R$. Concepts of the form B and roles of the form Q are called *basic*, and all others are called *general*.

Definition 1 (ontology)

A *fuzzy DL-Lite_R TBox* is a finite set of *fuzzy axioms* of the form $\langle B \sqsubseteq C, d \rangle$ and $\langle Q \sqsubseteq R, d \rangle$, where d is a number in $[0, 1]$. An axiom is *positive* if it does not have negation on its right-hand side and *negative* otherwise. A *fuzzy DL-Lite_R ABox* is a finite set of *fuzzy assertions* of the form $\langle B(a), d \rangle$ and $\langle P(a, b), d \rangle$, where $a, b \in N_I$. A *fuzzy DL-Lite_R ontology* is a pair of the form $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ where \mathcal{T} is a TBox and \mathcal{A} is an ABox.

Note that negations can never occur on the left-hand side of an axiom. In the remainder

Table 1. *The three fundamental continuous t-norms and related operations*

Name	$d \otimes e$	$d \Rightarrow e$	$\ominus d$
Gödel	$\min\{d, e\}$	$\begin{cases} 1 & d \leq e \\ e & \text{otherwise} \end{cases}$	$\begin{cases} 1 & d = 0 \\ 0 & \text{otherwise} \end{cases}$
Lukasiewicz	$\max\{d + e - 1, 0\}$	$\min\{1 - d + e, 1\}$	$1 - d$
product	$d \cdot e$	$\begin{cases} 1 & d \leq e \\ e/d & \text{otherwise} \end{cases}$	$\begin{cases} 1 & d = 0 \\ 0 & \text{otherwise} \end{cases}$

of this paper, we will mostly exclude the qualifiers “fuzzy,” and “DL-Lite” and simply refer to axioms, ontologies, etc.

The semantics of fuzzy DL-Lite_R is based on fuzzy interpretations, which provide a *membership degree* or for objects belonging to the different concept and role names. Formally, following the basics of classical description logics, concept names are interpreted as fuzzy unary relations, and role names are interpreted as fuzzy binary relations. To fully define this semantics in the presence of other constructors according to fuzzy logic, we need the notion of a triangular norm (or *t-norm* for short).

Definition 2 (t-norm)

A *t-norm* \otimes is a binary operator over the real interval $[0, 1]$ that is commutative, associative, monotonic, and has 1 as the neutral element; i.e., $1 \otimes x = x$ for all $x \in [0, 1]$ (Klement et al. 2000).

Triangular norms are used to generalize the logical conjunction to handle truth degrees that take values from the interval $[0, 1]$. Every continuous t-norm defines a unique *residuum* \Rightarrow where $f \otimes d \leq e$ iff $f \leq d \Rightarrow e$. The residuum interprets implications. With the help of this operation, it is also possible to interpret other logical operators such as negation ($\ominus d := d \Rightarrow 0$). The three basic continuous t-norms are the *Gödel*, *Lukasiewicz*, and *product* t-norms, which are defined, with their residua and negations in Table 1. These t-norms are the “fundamental” ones in the sense that every other continuous t-norm is isomorphic to the ordinal sum of copies of them (Hájek 1998; Mostert and Shields 1957). Hence, as usual, we focus our study on these three t-norms.

Note that the residuum always satisfies that $d \Rightarrow e = 1$ iff $d \leq e$, and that in the Gödel and product t-norms the negation is annihilating in the sense that it maps to 0 any positive value, while the negation of 0 is 1. In particular, this means that the negation is not *involution*; that is, $\ominus \ominus d \neq d$ in general. In contrast, the negation operator for the Lukasiewicz t-norm is involutive. In addition, the Lukasiewicz t-norm is the only t-norm (up to isomorphism) with the property that for every $x \in (0, 1)$ there exists a $y \in (0, 1)$ such that $x \otimes y = 0$. Specifically, this y is $1 - x$. In other words, the Lukasiewicz t-norm is *nilpotent*. From now on, unless specified explicitly otherwise, we assume that we have an arbitrary, but fixed, t-norm \otimes which underlies the operators used. When the t-norm becomes relevant in the following sections, we will often use G, Π , and L as prefixes to express that the underlying t-norm is Gödel, product, or Lukasiewicz, respectively, as usual in the literature.

We can now formally define the semantics of the logic. An *interpretation* is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set called the *domain*, and $\cdot^{\mathcal{I}}$ is the *interpretation function* which maps: (i) every individual name $a \in N_I$ to an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$; (ii) every

concept name $A \in N_C$ to a function $A^{\mathcal{I}} : \Delta^{\mathcal{I}} \rightarrow [0, 1]$; and (iii) every role name $P \in N_R$ to a function $P^{\mathcal{I}} : \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \rightarrow [0, 1]$. That is, concept names are interpreted as fuzzy unary relations and role names are interpreted as fuzzy binary relations over $\Delta^{\mathcal{I}}$. The interpretation function is extended to other constructors with the help of the t-norm operators as follows. For every $\delta, \eta \in \Delta^{\mathcal{I}}$,

$$\begin{aligned} (\exists Q)^{\mathcal{I}}(\delta) &:= \sup_{\delta' \in \Delta^{\mathcal{I}}} Q^{\mathcal{I}}(\delta, \delta') & (\neg B)^{\mathcal{I}}(\delta) &:= \ominus B^{\mathcal{I}}(\delta) & (\top)^{\mathcal{I}}(\delta) &:= 1 \\ (P^-)^{\mathcal{I}}(\delta, \eta) &:= P^{\mathcal{I}}(\eta, \delta) & (\neg Q)^{\mathcal{I}}(\delta, \eta) &:= \ominus Q^{\mathcal{I}}(\delta, \eta) \end{aligned}$$

The interpretation \mathcal{I} satisfies the axiom

- $\langle B \sqsubseteq C, d \rangle$ iff $B^{\mathcal{I}}(\delta) \Rightarrow C^{\mathcal{I}}(\delta) \geq d$ holds for every $\delta \in \Delta^{\mathcal{I}}$; and
- $\langle Q \sqsubseteq R, d \rangle$ iff $Q^{\mathcal{I}}(\delta, \eta) \Rightarrow R^{\mathcal{I}}(\delta, \eta) \geq d$ holds for every $\delta, \eta \in \Delta^{\mathcal{I}}$

It is a *model* of the TBox \mathcal{T} if it satisfies all axioms in \mathcal{T} . \mathcal{I} satisfies the assertion

- $\langle B(a), d \rangle$ iff $B^{\mathcal{I}}(a^{\mathcal{I}}) \geq d$;
- $\langle P(a, b), d \rangle$ iff $P^{\mathcal{I}}(a^{\mathcal{I}}, b^{\mathcal{I}}) \geq d$.

It is a *model* of the ABox \mathcal{A} if it satisfies all axioms in \mathcal{A} , and it is a *model* of the ontology $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ if it is a model of \mathcal{T} and of \mathcal{A} .

We note that the classical notion of DL-Lite_R (Calvanese et al. 2007) is a special case of fuzzy DL-Lite_R, where all the axioms and assertions hold with degree 1. In that case, it suffices to consider interpretations which map all elements to $\{0, 1\}$ representing the classical truth values. When speaking of classical ontologies, we remove the degree and assume it implicitly to be 1.

Example 3

Consider an ontology $\mathcal{O}_{\text{exa}} = (\mathcal{T}_{\text{exa}}, \mathcal{A}_{\text{exa}})$ representing some knowledge about a touristic location. The TBox

$$\begin{aligned} \mathcal{T}_{\text{exa}} = \{ & \langle \text{Monument} \sqsubseteq \text{TouristAttraction}, 1 \rangle, \quad \langle \text{Museum} \sqsubseteq \text{TouristAttraction}, 1 \rangle, \\ & \langle \text{Pub} \sqsubseteq \text{Eatery}, 1 \rangle, \quad \langle \text{Restaurant} \sqsubseteq \text{Eatery}, 1 \rangle, \quad \langle \text{locln} \sqsubseteq \text{Near}, 1 \rangle \\ & \langle \text{Museum} \sqsubseteq \text{Popular}, 0.6 \rangle, \quad \langle \exists \text{locln} \sqsubseteq \neg \text{Cheap}, 0.5 \rangle \quad \} \end{aligned}$$

defines some notions about eateries and tourist attractions, including some vague notions in the last two axioms. For example, it expresses that museums are popular (with a degree at least 0.6), and that services located at some attraction are not cheap (with degree at least 0.5). The ABox

$$\begin{aligned} \mathcal{A}_{\text{exa}} = \{ & \langle \text{Monument}(\text{peace}), 1 \rangle, \quad \langle \text{Monument}(\text{love}), 1 \rangle, \\ & \langle \text{Museum}(\text{modernArt}), 1 \rangle, \quad \langle \text{Museum}(\text{contArt}), 1 \rangle, \quad \langle \text{Museum}(\text{comic}), 1 \rangle, \\ & \langle \text{Restaurant}(\text{sioux}), 1 \rangle, \quad \langle \text{Restaurant}(\text{gamberone}), 1 \rangle, \\ & \langle \text{Pub}(\text{irish}), 1 \rangle, \quad \langle \text{locln}(\text{sioux}, \text{modernArt}), 1 \rangle, \\ & \langle \text{Popular}(\text{comic}), 0.8 \rangle, \quad \langle \text{Cheap}(\text{irish}), 0.6 \rangle, \quad \langle \text{near}(\text{irish}, \text{comic}), 0.7 \rangle \quad \} \end{aligned}$$

provides information about the specific attractions and services provided at the location. From this information, we can deduce, for example, that the `modernArt` museum is a `TouristAttraction`, and is `Popular` to a degree at least 0.6. Under the Gödel t-norm, a possible model of \mathcal{O}_{exa} is depicted graphically in Figure 1, where any assertion not depicted




Fig. 1. A model for the ontology \mathcal{O}_{exa} from Example 3. Individual names are abbreviated to avoid cluttering, and start with a lower-case letter as customary in DLs. The shape and border of the nodes represent the crisp concepts, while vague concepts are associated to a degree.

is considered to hold to degree 0. For example, the model from Figure 1 interprets the irish pub as being **Cheap** to degree 0.7, which satisfies the constraint in the ABox requiring this degree to be at least 0.6. In addition, the **peace** monument is **Popular** to degree 0.3, even though there is no explicit requirement for this in \mathcal{O}_{exa} . Note that under this semantics, any model \mathcal{I} of \mathcal{O}_{exa} should necessarily satisfy that $\text{Cheap}^{\mathcal{I}}(\text{sioux}^{\mathcal{I}}) = 0$; this is in fact the case in the model from Figure 1. Hence, adding any assertion of the form $\langle \text{Cheap}(\text{sioux}), d \rangle$ with $d > 0$ to this ontology would make it inconsistent.

For this paper, we are interested in answering two kinds of queries. The first kind are conjunctive queries, which consider whether a combination of facts can be derived from the knowledge in an ontology. In the fuzzy setting, the degree of such derivation must also be taken into account.

Let N_V be a set of *variables*, which is disjoint from N_I , N_C , and N_R . A *term* is an element of $N_V \cup N_I$; that is, an individual name or a variable. An *atom* is an expression of the form $C(t)$ (concept atom) or $P(t_1, t_2)$ (role atom). Henceforth, \mathbf{x} and \mathbf{y} denote tuples of variables.

Definition 4 (conjunctive query)

A *conjunctive query* (CQ) is a first-order formula of the form $\exists \mathbf{y}.\phi(\mathbf{x}, \mathbf{y})$ where ϕ is a conjunction of atoms which only use the variables from \mathbf{x} and \mathbf{y} . The variables \mathbf{y} are called *existential variables*, and those in \mathbf{x} are *answer variables*. A *union of conjunctive queries* (UCQ) is a finite set of CQs that use the same answer variables. Henceforth, $\text{At}(\phi)$ denotes the set of all atoms appearing in ϕ .

As in the classical setting, an answer to a conjunctive query, or a union of conjunctive queries, is only considered when it is provided by every model of the ontology. This is usually known as a *certain answer*.

Given the CQ $q(\mathbf{x}) = \exists \mathbf{y}.\phi(\mathbf{x}, \mathbf{y})$, the interpretation \mathcal{I} , and a tuple of individuals \mathbf{a} of the same length as \mathbf{x} , a *match* is a mapping π which assigns to each $a \in N_I$ the value $a^{\mathcal{I}}$; to each variable in \mathbf{x} the corresponding element of $\mathbf{a}^{\mathcal{I}}$; and to each variable in \mathbf{y} an element $\delta \in \Delta^{\mathcal{I}}$. We extend the match π to apply to assertions as follows: $\pi(B(t)) = B(\pi(t))$ and $\pi(P(t_1, t_2)) = P(\pi(t_1), \pi(t_2))$. The *degree* of the CQ $q(\mathbf{x})$ w.r.t. the match π is

$$q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}, \pi(\mathbf{y})) := \bigotimes_{\alpha \in \text{At}(\phi)} (\pi(\alpha))^{\mathcal{I}}.$$

That is, a match maps all the variables in the query to elements of the interpretation

domain, where the tuple \mathbf{a} is used to identify the mapping of the answer variables. The satisfaction or matching degree of the query is the (fuzzy) conjunction—that is, the t-norm—of the satisfaction or matching degrees of the atoms under this mapping. From now on, $\Pi(\mathcal{I})$ denotes the set of all matches of $q(\mathbf{x})$ w.r.t. the interpretation \mathcal{I} . An important difference between classical query answering and our setting is that the fuzzy semantics provides a degree to every possible atom. Hence, in reality $\Pi(\mathcal{I})$ is always defined by the set of all tuples of individuals with length $|\mathbf{x}|$. However, the degree of these matches varies and may often be zero. For example, for the model \mathcal{I} in Figure 1 and the query $q(x) = \text{Popular}(x)$, the set of all matches $\Pi(\mathcal{I})$ assigns to the variable x any of the constants $\{\mathbf{mA}, \mathbf{cA}, \mathbf{c}, \mathbf{p}, \mathbf{l}, \mathbf{s}, \mathbf{i}, \mathbf{g}\}$ to degrees 0.7, 0.6, 0.9, 0.3, 0, 0, 0, and 0, respectively. When answering a query, one is often interested in the matches that hold to at least some degree d , as defined next.

Definition 5 (degree queries)

A tuple of individuals \mathbf{a} is an *answer* of the conjunctive query $q(\mathbf{x})$ to degree d w.r.t. the interpretation \mathcal{I} (denoted by $\mathcal{I} \models q(\mathbf{a}) \geq d$) iff $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) := \sup_{\pi \in \Pi(\mathcal{I})} q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}, \pi(\mathbf{y})) \geq d$. It is a *certain answer* (or *answer* for short) of $q(\mathbf{x})$ over the ontology \mathcal{O} to degree d (denoted by $\mathcal{O} \models q(\mathbf{a}) \geq d$) iff $\mathcal{I} \models q(\mathbf{a}) \geq d$ holds for every model \mathcal{I} of \mathcal{O} . The crisp set of certain answers of the query $q(\mathbf{x})$ w.r.t. \mathcal{O} and their degree is denoted by $\text{ans}(q(\mathbf{x}), \mathcal{O})$; that is,

$$\text{ans}(q(\mathbf{x}), \mathcal{O}) := \{(\mathbf{a}, d) \mid \mathcal{O} \models q(\mathbf{a}) \geq d \text{ and for all } d' > d, \mathcal{O} \not\models q(\mathbf{a}) \geq d'\}.$$

It is important to keep in mind that the atoms in a CQ are not graded, but simply try to match with elements in the domain as both concept and roles are interpreted as fuzzy relations (unary and binary, respectively). The use of the truth degrees in the ontology becomes relevant in the degree of the answers found. Moreover, recall that every tuple of individuals of length $|\mathbf{x}|$ belongs to $\text{ans}(q(\mathbf{x}), \mathcal{O})$, but with different associated degrees.

Returning to our example, while all individuals belong to the set $\text{ans}(q(x))$, for the query $q(x) = \text{Popular}(x)$ to some degree, the certain answers for $q(x)$ w.r.t. \mathcal{O}_{exa} to degree at least 0.6 are only **modernArt**, **contArt**, and **comic**. The latter one is the only answer to degree at least 0.8.

The second kind of query we are interested in generalises that of degree queries, when considering the Gödel semantics, by allowing a degree threshold for each of the atoms in the conjunction, rather than for the overall conjunction. We formally define this class next.

Definition 6 (threshold queries)

A *threshold atom* is an expression of the form $\alpha \geq d$, where α is an atom and $d \in [0, 1]$. A *threshold query* (TQ) is a first-order formula of the form $\exists \mathbf{y}.\phi(\mathbf{x}, \mathbf{y})$ where ϕ is a conjunction of threshold atoms using only the variables from \mathbf{x} and \mathbf{y} .

The notion of a match and an answer to a threshold query are analogous to those of degree queries, with the proviso that the degree bounds apply at the level of atoms, and not at the level of queries.

Definition 7 (TQ answer)

Given an interpretation \mathcal{I} and a tuple of individuals \mathbf{a} , the match π *satisfies* the threshold atom $\alpha \geq d$ (denoted by $\pi \models \alpha \geq d$) iff $\alpha^{\mathcal{I}} \geq d$. It *satisfies* the threshold query $q(\mathbf{x}) = \exists \mathbf{y}.\phi(\mathbf{x}, \mathbf{y})$ ($\pi \models q(\mathbf{a})$) iff $\pi \models \alpha \geq d$ holds for every threshold atom in q .

A tuple of individuals \mathbf{a} is an *answer* to the TQ $q(\mathbf{x})$ w.r.t. the interpretation \mathcal{I} ($\mathcal{I} \models q(\mathbf{a})$) iff there is a match π w.r.t. \mathbf{a} and \mathcal{I} such that $\pi \models q(\mathbf{a})$. It is a *certain answer* of $q(\mathbf{x})$ over the ontology \mathcal{O} iff for every model \mathcal{I} of \mathcal{O} it holds that $\mathcal{I} \models q(\mathbf{a})$.

Note that, differently from conjunctive queries, but in an analogous manner to degree queries, the answers to a threshold query are *not* graded. Indeed, a tuple \mathbf{a} may or may not be an answer, and we are interested in finding those tuples which satisfy the degrees at each of the threshold atoms. In a sense, threshold queries provide a more fine-grained structure to deal with the properties of interest within a query in relation to degree queries. Indeed, in a degree query, one can only provide an overall degree which should be obtained after the degrees of all the atoms are conjoined via the t-norm. In particular, for non-idempotent t-norms and large queries, this conjunction will tend to be smaller and smaller, and the degrees of the independent atoms have the same influence overall. Even when considering the idempotent Gödel t-norm, a degree query $q(\mathbf{x}) \geq d$ only expresses that all the atoms in q should hold to degree at least d (recall that the Gödel t-norm refers to the minimum operator), but it is not possible to express that some atoms should hold with a higher degree. A threshold query, on the other hand, is capable of requiring different degrees for each of the atoms.

Example 8

Suppose, in our running example, that we are interested in finding a cheap eatery that is nearby a popular tourist attraction, and that we are using the Gödel semantics. This basic query could be expressed as¹

$$q(x) = \exists y. \text{Cheap}(x), \text{Popular}(y), \text{near}(x, y).$$

Since this query considers vague concepts and roles, we want to find answers that satisfy it to at least some degree. For the degree query $q(x) \geq 0.6$, the only possible answer is the irish pub.

Suppose now that for us it is more important that the eatery is cheap than the popularity of the tourist attraction. For example, even though we are content with the tourist attraction being popular to only degree 0.6, the eatery should be cheap to a degree at least 0.8. This can be expressed through the threshold query

$$q'(x) = \exists y. \text{Cheap}(x) \geq 0.8, \text{Popular}(y) \geq 0.6, \text{near}(x, y) \geq 0.6.$$

In this case, the TQ has no answers w.r.t. the ontology \mathcal{O} . However, any answer to q' would also be an answer to $q(x) \geq 0.6$, as overall they define the same minimum over all the degrees of interest. Note that this last claim only holds for the case of the Gödel semantics. Indeed, as we will see later in this paper, for other semantics degree queries are not properly special cases of TQs.

A class of conjunctive queries of special significance is that where the tuple of answer variables \mathbf{x} is empty. This means that the answer tuple of individuals provided as an answer must also be empty. In the classical setting, these are called *Boolean queries*, because they can only return a Boolean value: true if there is a match for the existential variables in every model, and false otherwise. In the fuzzy setting, the set of answers

¹ For brevity, we conjoin the atoms in a CQ through commas (‘,’) instead of \wedge .

to such a query will only contain one element $((), d)$. Thus, in that case, we are only interested in finding the degree d , and call those queries *fuzzy queries*. This degree is the tightest value for which we can find a satisfying matching. Formally, the ontology \mathcal{O} entails the fuzzy query $q()$ to degree d iff $\mathcal{O} \models q() \geq d$ and $\mathcal{O} \not\models q() \geq d'$ for all $d' > d$. Fuzzy queries allow us to find the degree of a specific answer \mathbf{a} without having to compute $\text{ans}(q(\mathbf{x}), \mathcal{O})$: simply compute the degree of the fuzzy query $q(\mathbf{a})$.

In the case of threshold queries, we can also consider the special case where the answer tuple \mathbf{x} is empty. In that case, as in the classical case, the only possible answer is the empty tuple (if there is a match which satisfies the query) or no answer if no such match exists. For that reason, in the case of threshold queries without answer variables we preserve the classical terminology and call them Boolean (threshold) queries.

As it is typically done for query answering in description logics, we consider two measures of complexity: *data complexity*, where only the size of the ABox (and the candidate answer, if any) is considered as part of the input, and *combined complexity* in which the size of the whole ontology (including the TBox) is taken into account.² For data complexity, it is relevant to consider sub-linear complexity classes. In particular, we consider AC^0 and LOGSPACE . For the full formal definitions, we refer the interested reader to Papadimitriou (1994) and Boppana and Sipser (1990). Here we only mention briefly that evaluation of FO-queries over a database is in AC^0 on the size of the database (Abiteboul et al. 1994) and AC^0 is strictly contained in LOGSPACE (Furst et al. 1984).

In classical DL-Lite_R, query answering w.r.t. an ontology is reduced to the standard problem of query answering over a database through a process known as query rewriting, and thus is in AC^0 w.r.t. data complexity. The main idea is to include in the query all the information that is required by the TBox, in such a way that only assertions from the ABox need to be considered. In our running example, note that there is no assertion in the ABox \mathcal{A}_{exa} which explicitly mentions a tourist attraction. We only know that the two monuments and the three museums are tourist attractions thanks to the TBox. In this case, the query rewriting approach would take the query $q(x) = \text{TouristAttraction}(x)$ and transform it into the UCQ

$$\{\text{TouristAttraction}(x), \text{Museum}(x), \text{Monument}(x)\}$$

looking “backwards” over the axioms in the TBox. The answers of this UCQ over the ABox alone are exactly those of the original query over the whole ontology.

As seen in this simple example, there are many possible choices to create the matches that comply with the TBox. Hence, this method results in a UCQ even if the original query is a simple CQ. At this point, the ABox is treated as a database, which suffices to find all the certain answers. Similarly, a special UCQ can be used to verify that the ontology is *consistent*; that is, whether it is possible to build a model for this ontology. For

² Note that our notion of *combined complexity* does *not* include the query as part of the input, but only the ontology. This view contrasts the usual database definition (and papers following it) where the combined complexity includes the query, but is in line with the terminology used in ontology-based query answering; e.g. (Artale et al. 2009). The motivation is to understand the influence of the knowledge in the complexity, abstracting from the query, which is already known to be a source of intractability already for databases. In the context of ontology-based query answering, combined complexity is typically only used in combination with simple fixed queries, which means that the query does not really have an important influence.

the full details on how these query rewritings work in classical DL-Lite_R, see (Calvanese et al. 2007). In terms of combined complexity, consistency can be decided in polynomial time; in fact, it is NLOGSPACE-complete (Artale et al. 2009).

3 The Canonical Interpretation

A very useful tool for developing techniques for answering queries in DL-Lite_R is the canonical interpretation. We first show that the same idea can be extended (with the necessary modifications) to fuzzy ontologies, independently of the t-norm underlying its semantics.

Let $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ be a DL-Lite_R ontology and assume w.l.o.g. that there are no axioms of the form $\langle \exists Q_1 \sqsubseteq \exists Q_2, d \rangle \in \mathcal{T}$; any such axiom can be substituted by the two axioms $\langle \exists Q_1 \sqsubseteq A, 1 \rangle, \langle A \sqsubseteq \exists Q_2, d \rangle$ where A is a new concept name not appearing in \mathcal{T} . The *canonical interpretation* of \mathcal{O} is the interpretation $\mathcal{I}_{\text{can}}(\mathcal{O}) = (\Delta^{\mathcal{I}_{\text{can}}}, \cdot^{\mathcal{I}_{\text{can}}})$ over the domain $\Delta^{\mathcal{I}_{\text{can}}} := N_I \cup N_N$ —where N_N is a countable set of *constants*—obtained through the following (infinite) process. Starting from the *empty* interpretation which sets $A^{\mathcal{I}_{\text{can}}}(\delta) = 0$ and $P^{\mathcal{I}_{\text{can}}}(\delta, \eta) = 0$ for every $A \in N_C, P \in N_R$ and $\delta, \eta \in \Delta^{\mathcal{I}_{\text{can}}}$, exhaustively apply the following rules:

- R1.** if $\langle A(a), d \rangle \in \mathcal{A}$ and $A^{\mathcal{I}_{\text{can}}}(a) < d$, then update the value $A^{\mathcal{I}_{\text{can}}}(a) := d$;
- R2.** if $\langle P(a, b), d \rangle \in \mathcal{A}$ and $P^{\mathcal{I}_{\text{can}}}(a, b) < d$, then update the value $P^{\mathcal{I}_{\text{can}}}(a, b) := d$;
- R3.** if $\langle A_1 \sqsubseteq A_2, d \rangle \in \mathcal{T}$ and $A_2^{\mathcal{I}_{\text{can}}}(\delta) < A_1^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$, then update $A_2^{\mathcal{I}_{\text{can}}}(\delta) := A_1^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$;
- R4.** if $\langle A \sqsubseteq \exists P, d \rangle \in \mathcal{T}$ and for every $\eta \in \Delta^{\mathcal{I}_{\text{can}}}$, $P^{\mathcal{I}_{\text{can}}}(\delta, \eta) < A^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$ holds, then select a fresh element η_0 such that $P^{\mathcal{I}_{\text{can}}}(\delta, \eta_0) = 0$ and update the value $P^{\mathcal{I}_{\text{can}}}(\delta, \eta_0) := A^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$;
- R5.** if $\langle A \sqsubseteq \exists P^-, d \rangle \in \mathcal{T}$ and for every $\eta \in \Delta^{\mathcal{I}_{\text{can}}}$ $P^{\mathcal{I}_{\text{can}}}(\eta, \delta) < A^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$ holds, then select a fresh element η_0 such that $P^{\mathcal{I}_{\text{can}}}(\eta_0, \delta) = 0$ and update the value $P^{\mathcal{I}_{\text{can}}}(\eta_0, \delta) := A^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$;
- R6.** if $\langle \exists P \sqsubseteq A, d \rangle \in \mathcal{T}$ and $\exists \eta \in \Delta^{\mathcal{I}_{\text{can}}}$ such that $A^{\mathcal{I}_{\text{can}}}(\delta) < P^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d$, then update $A^{\mathcal{I}_{\text{can}}}(\delta) := P^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d$;
- R7.** if $\langle \exists P^- \sqsubseteq A, d \rangle \in \mathcal{T}$ and $\exists \eta \in \Delta^{\mathcal{I}_{\text{can}}}$ such that $A^{\mathcal{I}_{\text{can}}}(\delta) < P^{\mathcal{I}_{\text{can}}}(\eta, \delta) \otimes d$, then update $A^{\mathcal{I}_{\text{can}}}(\delta) := P^{\mathcal{I}_{\text{can}}}(\eta, \delta) \otimes d$;
- R8.** if $\langle Q_1 \sqsubseteq Q_2, d \rangle \in \mathcal{T}$ and $Q_2^{\mathcal{I}_{\text{can}}}(\delta, \eta) < Q_1^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d$, then update $Q_2^{\mathcal{I}_{\text{can}}}(\delta, \eta)$ to the value $Q_1^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d$.

where the rules are applied in a fair manner; that is, an applicable rule is eventually triggered. The process of rule application is a monotone non-decreasing function, and as such has a least fixpoint, which is the canonical interpretation $\mathcal{I}_{\text{can}}(\mathcal{O})$.³

Intuitively, $\mathcal{I}_{\text{can}}(\mathcal{O})$ should be a minimal model of \mathcal{O} , which describes the necessary conditions of all other models of \mathcal{O} . Indeed, the first two rules ensure that the conditions imposed by the ABox are satisfied, by setting the degrees of the unary and binary relations to the smallest required value. The remaining rules guarantee that all elements of the domain satisfy the positive axioms from the TBox, and each rule is as weak as possible in satisfying these constraints. The canonical interpretation of the ontology \mathcal{O}_{exa} from Example 3 is depicted in Figure 2. Note that in general it provides a lower membership

³ By Tarski's Theorem (Tarski 1955), this fixpoint is the limit of the (fair) application of the rules starting from the smallest element; in this case, the empty interpretation as described before.




Fig. 2. The canonical interpretation for the ontology \mathcal{O}_{exa} from our running example.

degree of each individual to every concept when compared to the model from Figure 1. This intuition justifies the name of *canonical* interpretation. As in the classical case, $\mathcal{I}_{\text{can}}(\mathcal{O})$ can be homomorphically embedded in every model of \mathcal{O} , and hence be used as a representative of them all. We show a similar result with the difference that in this case, the homomorphism needs to take into account the truth degrees from the interpretation function as well.⁴ This is described in the following proposition.

Proposition 9

Let \mathcal{O} be a consistent fuzzy DL-Lite ontology, $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ be a model of \mathcal{O} , and $\mathcal{I}_{\text{can}}(\mathcal{O}) = (\Delta^{\mathcal{I}_{\text{can}}}, \cdot^{\mathcal{I}_{\text{can}}})$ its canonical interpretation. There is a function ψ from $\Delta^{\mathcal{I}_{\text{can}}}$ to $\Delta^{\mathcal{I}}$ such that:

1. for each $A \in N_C$ and $\delta \in \Delta^{\mathcal{I}_{\text{can}}}$, $A^{\mathcal{I}_{\text{can}}}(\delta) \leq A^{\mathcal{I}}(\psi(\delta))$; and
2. for each $P \in N_R$ and $\delta, \eta \in \Delta^{\mathcal{I}_{\text{can}}}$, $P^{\mathcal{I}_{\text{can}}}(\delta, \eta) \leq P^{\mathcal{I}}(\psi(\delta), \psi(\eta))$.

Proof

Let $\mathcal{O} = (\mathcal{T}, \mathcal{A})$. We construct the function ψ recursively through the rule applications that define $A^{\mathcal{I}_{\text{can}}}$, and show that the two properties from the proposition are invariant w.r.t. the rule applications. We define first $\psi(a) = a^{\mathcal{I}}$ for all $a \in N_I$. Recall that initially, $A^{\mathcal{I}_{\text{can}}}(\delta) = P^{\mathcal{I}_{\text{can}}}(\delta, \eta) = 0$ for all $A \in N_C, P \in N_R, \delta, \eta \in \Delta^{\mathcal{I}_{\text{can}}}$. Hence, the properties hold trivially in this case.

Assume now that the properties hold before a rule application; we show that they also hold afterwards by a case analysis over the rule used:

- R1.** if $\langle A(a), d \rangle \in \mathcal{A}$, since \mathcal{I} is a model of this axiom, it follows that $A^{\mathcal{I}}(a^{\mathcal{I}}) \geq d$. The rule application sets $A^{\mathcal{I}_{\text{can}}}(a) = d$ and hence $A^{\mathcal{I}_{\text{can}}}(a) \leq A^{\mathcal{I}}(\psi(a)) = A^{\mathcal{I}}(a^{\mathcal{I}})$.
- R2.** if $\langle P(a, b), d \rangle \in \mathcal{A}$, the rule application sets $P^{\mathcal{I}_{\text{can}}}(a, b) = d$. Since \mathcal{I} satisfies this axiom, it follows that $P^{\mathcal{I}}(\psi(a), \psi(b)) = P^{\mathcal{I}}(a^{\mathcal{I}}, b^{\mathcal{I}}) \geq d = P^{\mathcal{I}_{\text{can}}}(a, b)$.
- R3.** if $\langle A_1 \sqsubseteq A_2, d \rangle \in \mathcal{T}$, the rule application over a given $\delta \in \Delta^{\mathcal{I}_{\text{can}}}$ updates $A_2^{\mathcal{I}_{\text{can}}}(\delta)$ to $A_1^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$. Since \mathcal{I} satisfies this axiom, by the induction hypothesis and monotonicity of \otimes we know that $A_2^{\mathcal{I}}(\psi(\delta)) \geq A_1^{\mathcal{I}}(\psi(\delta)) \otimes d \geq A_1^{\mathcal{I}_{\text{can}}}(\delta) \otimes d = A_2^{\mathcal{I}_{\text{can}}}(\delta)$.
- R4.** if $\langle A \sqsubseteq \exists P, d \rangle \in \mathcal{T}$, let δ be the element over which the rule is applicable, and η_0 the fresh element selected by the rule application. Since \mathcal{I} is a model, we know

⁴ A careful reader will notice that the exhaustive application of the rules may produce different interpretations. We discuss this issue in further detail later in this section. For now, it suffices to know that all the possible canonical interpretations are equivalent modulo homomorphisms.

that there exists an element $\kappa \in \Delta^{\mathcal{I}}$ such that $P^{\mathcal{I}}(\psi(\delta), \kappa) \geq A^{\mathcal{I}}(\psi(\delta)) \otimes d$. We thus define $\psi(\eta_0) := \kappa$. By the induction hypothesis and monotonicity of \otimes we get $P^{\mathcal{I}}(\psi(\delta), \psi(\eta_0))P^{\mathcal{I}}(\psi(\delta), \kappa) \geq A^{\mathcal{I}}(\psi(\delta)) \otimes d \geq A^{\mathcal{I}_{\text{can}}}(\psi(\delta)) \otimes d = P^{\mathcal{I}_{\text{can}}}(\delta, \eta_0)$.

- R5.** if $\langle A \sqsubseteq \exists P^-, d \rangle \in \mathcal{T}$, let δ be the element over which the rule is applicable, and η_0 the fresh element selected by the rule application. Since \mathcal{I} is a model, we know that there exists an element $\kappa \in \Delta^{\mathcal{I}}$ such that $P^{\mathcal{I}}(\kappa, \psi(\delta)) \geq A^{\mathcal{I}}(\psi(\delta)) \otimes d$. We thus define $\psi(\eta_0) := \kappa$. By the induction hypothesis and monotonicity of \otimes we get $P^{\mathcal{I}}(\psi(\eta_0), \psi(\delta))P^{\mathcal{I}}(\kappa, \psi(\delta)) \geq A^{\mathcal{I}}(\psi(\delta)) \otimes d \geq A^{\mathcal{I}_{\text{can}}}(\psi(\delta)) \otimes d = P^{\mathcal{I}_{\text{can}}}(\eta_0, \delta)$.
- R6.** if $\langle \exists P \sqsubseteq A, d \rangle \in \mathcal{T}$, then for the chosen $\delta, \eta \in \Delta^{\mathcal{I}_{\text{can}}}$ we have by the induction hypothesis that $A^{\mathcal{I}_{\text{can}}}(\delta) = P^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d \leq P^{\mathcal{I}}(\psi(\delta), \psi(\eta)) \otimes d \leq A^{\mathcal{I}}(\psi(\delta))$.
- R7.** if $\langle \exists P \sqsubseteq A, d \rangle \in \mathcal{T}$, then for the chosen $\delta, \eta \in \Delta^{\mathcal{I}_{\text{can}}}$ we have by the induction hypothesis that $A^{\mathcal{I}_{\text{can}}}(\delta) = P^{\mathcal{I}_{\text{can}}}(\eta, \delta) \otimes d \leq P^{\mathcal{I}}(\psi(\eta), \psi(\delta)) \otimes d \leq A^{\mathcal{I}}(\psi(\delta))$.
- R8.** if $\langle Q_1 \sqsubseteq Q_2, d \rangle \in \mathcal{T}$, the rule application over given $\delta, \eta \in \Delta^{\mathcal{I}_{\text{can}}}$ updates $Q_2^{\mathcal{I}_{\text{can}}}(\delta, \eta)$ to $Q_1^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d$. Since \mathcal{I} satisfies this axiom, by the induction hypothesis and monotonicity of \otimes we know that

$$Q_2^{\mathcal{I}}(\psi(\delta), \psi(\eta)) \geq Q_1^{\mathcal{I}}(\psi(\delta), \psi(\eta)) \otimes d \geq Q_1^{\mathcal{I}_{\text{can}}}(\delta, \eta) \otimes d = Q_2^{\mathcal{I}_{\text{can}}}(\delta, \eta).$$

Hence, the result holds after the fair application of all possible rules. \square

Importantly, note that the construction of $\mathcal{I}_{\text{can}}(\mathcal{O})$ does not take the negations into account; e.g., the axiom $\langle \exists \text{locln} \sqsubseteq \neg \text{Cheap}, 0.5 \rangle$ is never used during this construction. The effect of this is that $\mathcal{I}_{\text{can}}(\mathcal{O})$ might not be a model of \mathcal{O} at all.

Example 10

Consider the fuzzy DL-Lite_R ontology $\mathcal{O}_{\text{exa}} = (\mathcal{T}_0, \mathcal{A}_0)$ where

$$\begin{aligned} \mathcal{T}_0 &:= \{ \langle A_1 \sqsubseteq \neg A_2, 1 \rangle \}, \\ \mathcal{A}_0 &:= \{ \langle A_1(a), 0.5 \rangle, \langle A_2(a), 0.5 \rangle \}. \end{aligned}$$

Under the Gödel semantics, by application of the first rule, the canonical interpretation maps $A_1^{\mathcal{I}_{\text{can}}}(a) = A_2^{\mathcal{I}_{\text{can}}}(a) = 0.5$. However, this violates the axiom in \mathcal{T}_0 , which requires that $A_1^{\mathcal{I}_{\text{can}}}(a) \Rightarrow \neg A_2^{\mathcal{I}_{\text{can}}}(a) = 1$. That is, it requires that $A_1^{\mathcal{I}_{\text{can}}}(a) < \neg A_2^{\mathcal{I}_{\text{can}}}(a)$, which is only possible when $A_1^{\mathcal{I}_{\text{can}}}(a) = 0$ or $A_2^{\mathcal{I}_{\text{can}}}(a) = 0$. Note that a similar phenomenon could be observed also in the TBox \mathcal{T}_{exa} of our running example, which contains an axiom with a negated concept.

The issue is that the negative axioms may introduce inconsistencies, by enforcing upper bounds in the degrees used, which are not verified by the canonical interpretation; recall, in fact, that the previously described construction monotonically increases the degrees to satisfy the minimal requirements, but never verifies whether these degrees affect some upper bound. On the other hand, we can prove that, as long as there is a model, $\mathcal{I}_{\text{can}}(\mathcal{O})$ is one.

Proposition 11

$\mathcal{I}_{\text{can}}(\mathcal{O})$ is a model of \mathcal{O} iff \mathcal{O} is consistent.

Proof

The *only if* direction is trivial, hence we focus on showing that if \mathcal{O} is consistent, then $\mathcal{I}_{\text{can}}(\mathcal{O})$ is a model of \mathcal{O} . Note first that, by construction, $\mathcal{I}_{\text{can}}(\mathcal{O})$ satisfies all positive axioms. Otherwise, a rule would trigger, and since the construction applies all rules fairly until exhaustion, no rule is applicable in the resulting interpretation $\mathcal{I}_{\text{can}}(\mathcal{O})$. Hence, if $\mathcal{I}_{\text{can}}(\mathcal{O})$ is not a model of \mathcal{O} , there must exist a negative axiom of the form (i) $\langle B \sqsubseteq \neg C, d \rangle$ or (ii) $\langle Q \sqsubseteq \neg R, d \rangle$ that is not satisfied by the canonical interpretation. We consider the case (i); the other case can be treated analogously.

If $\mathcal{I}_{\text{can}}(\mathcal{O}) \not\models \langle B \sqsubseteq \neg C, d \rangle$, then there must exist an element $\delta \in \Delta^{\mathcal{I}_{\text{can}}}$ such that $B^{\mathcal{I}_{\text{can}}}(\delta) \Rightarrow (\neg C)^{\mathcal{I}_{\text{can}}}(\delta) < d$ or, equivalently, $\ominus C^{\mathcal{I}_{\text{can}}}(\delta) < B^{\mathcal{I}_{\text{can}}}(\delta) \otimes d$. Since \mathcal{O} is consistent, there must exist a model $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ of \mathcal{O} . By Proposition 9, there exists a function $\psi : \Delta^{\mathcal{I}_{\text{can}}} \rightarrow \Delta^{\mathcal{I}}$ such that, in particular, $B^{\mathcal{I}_{\text{can}}}(\delta) < B^{\mathcal{I}}(\psi(\delta))$ and $C^{\mathcal{I}_{\text{can}}}(\delta) < C^{\mathcal{I}}(\psi(\delta))$. By antitonicity of \ominus , the latter means that $\ominus C^{\mathcal{I}}(\psi(\delta)) \leq \ominus C^{\mathcal{I}_{\text{can}}}(\delta)$ and hence

$$\ominus C^{\mathcal{I}}(\psi(\delta)) \leq \ominus C^{\mathcal{I}_{\text{can}}}(\delta) < B^{\mathcal{I}_{\text{can}}}(\delta) \otimes d \leq B^{\mathcal{I}}(\psi(\delta)) \otimes d$$

But this means that $\mathcal{I} \not\models \langle B \sqsubseteq \neg C, d \rangle$, which contradicts the assumption that \mathcal{I} was a model of \mathcal{O} . \square

It can be seen that the ontology \mathcal{O}_0 from Example 10 is inconsistent under the Gödel semantics. On the other hand, under the Łukasiewicz semantics, \mathcal{O}_0 is in fact consistent which, by this proposition, means that $\mathcal{I}_{\text{can}}(\mathcal{O})$ is a model of this ontology. This is easily confirmed by recalling that the Łukasiewicz negation is involutive; that is $\ominus d = 1 - d$. In the case of the example, we have $\ominus 0.5 = 0.5$; the axiom $\langle A_1 \sqsubseteq \neg A_2, 1 \rangle$ is satisfied because $0.5 = A_1^{\mathcal{I}_{\text{can}}}(a) \leq (\neg A_2)^{\mathcal{I}_{\text{can}}}(a) = 0.5$.

The consequence of the last two propositions is that $\mathcal{I}_{\text{can}}(\mathcal{O})$ is complete for existential positive queries, and in particular for conjunctive queries and threshold queries.

Corollary 12

If \mathcal{O} is a consistent fuzzy DL-Lite_R ontology, then

1. for every CQ $q(\mathbf{x})$, answer tuple \mathbf{a} , and $d \in [0, 1]$ it holds that $\mathcal{O} \models q(\mathbf{a}) \geq d$ iff $\mathcal{I}_{\text{can}}(\mathcal{O}) \models q(\mathbf{a}) \geq d$;
2. for every TQ $q(\mathbf{x})$ and answer tuple \mathbf{a} , $\mathcal{O} \models q(\mathbf{a})$ iff $\mathcal{I}_{\text{can}}(\mathcal{O}) \models q(\mathbf{a})$.

Proof

Proposition 11 states that $\mathcal{I}_{\text{can}}(\mathcal{O})$ is a model; hence anything that does not follow from it cannot be an answer. On the other hand, Proposition 9 states that the degree of every atom in any model is at least the degree given by \mathcal{I}_{can} , and hence if a tuple is an answer in \mathcal{I}_{can} , it is also an answer in every other model. \square

A short note on the canonical interpretation

Before delving deeper into the process of answering queries (the main contribution of this paper), it is worth considering the canonical interpretation in more detail, starting with the definite article used in its naming. Indeed, although we always speak about *the* canonical interpretation, the actual structure produced is not necessarily unique, and depends on the order in which rules are chosen to apply, specially in relation to rules **R4** and **R5**, which introduce new relevant elements. This is highlighted in the following example.

Fig. 3. Two canonical interpretation constructions from the ontology in Example 13. From the empty interpretation (a), **R1** is applied to each assertion to reach (c). One can either apply **R4** to $\langle A \sqsubseteq \exists R, 0.3 \rangle$ and go through the upper branch (d) to build the interpretation (e); or to $\langle B \sqsubseteq \exists R, 0.5 \rangle$ and obtain (f) directly.

Example 13

Consider an ontology containing the axioms

$$\begin{aligned} \mathcal{A} &:= \{\langle A(a), 1 \rangle, \langle B(a), 1 \rangle\} \\ \mathcal{T} &:= \{\langle A \sqsubseteq \exists R, 0.3 \rangle, \langle B \sqsubseteq \exists R, 0.5 \rangle\} \end{aligned}$$

After applying the rule **R1** over the ABox axioms, we have an interpretation where $A^{\mathcal{I}^{\text{can}}}(a) = B^{\mathcal{I}^{\text{can}}}(a) = 1$. At this point, rule **R4** is applicable for any of the two axioms in \mathcal{T} . If we first apply it to the first axiom, we select a fresh element, e.g. η_0 and set $R(a, \eta_0) = 0.3$; at this point, the rule is still applicable to the second axiom. This application requires selecting a new fresh element (now η_1) and set $R(a, \eta_1) = 0.5$. At this point, no rules are applicable and we have a canonical interpretation.

If instead we chose first to apply the rule on the second axiom, we would choose a fresh element (say, η_2) and set $R(a, \eta_2) = 0.5$. This application immediately disallows the application of the rule to the first axiom, and hence the process stops.

Note that the two interpretations built in this example are not equivalent (see Figure 3). However, they are homomorphic in the sense specified by Proposition 9. This is not a coincidence. In fact, note that the proofs of Propositions 9 and 11 do not depend on the order of rule applications, but only on the fact that these rules were exhaustively (and fairly) applied. If the ontology is consistent, by the latter proposition the interpretation obtained is a model, regardless of the order chosen, and by the former proposition, it is homomorphic to all the interpretations which can be derived following different application orderings. In other words, the canonical interpretation is unique *up to homomorphism*. In the following, we disregard this issue and consider an arbitrary, but fixed, canonical interpretation as unique.

We now return to the issue of answering queries. Corollary 12 states that these queries can be answered through the canonical interpretation. Obviously, such an approach is impractical; in fact, impossible, because it is an infinite model constructed through an infinitary process. Additionally, we still have the burden to prove that the ontology is consistent, which is a prerequisite for the use of Corollary 12 to answer queries. Fortunately, for the Gödel and product t-norms, we can resort to existing results from the literature for this latter task.

Definition 14 (classical version)

Let $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ be a fuzzy DL-Lite ontology. The *classical version* $\widehat{\mathcal{O}}$ of \mathcal{O} is defined by $\widehat{\mathcal{O}} := (\widehat{\mathcal{T}}, \widehat{\mathcal{A}})$, where

$$\begin{aligned} \widehat{\mathcal{T}} &:= \{B \sqsubseteq C \mid \langle B \sqsubseteq C, d \rangle \in \mathcal{T}, d > 0\} \cup \{Q \sqsubseteq R \mid \langle Q \sqsubseteq R, d \rangle \in \mathcal{T}, d > 0\}, \\ \widehat{\mathcal{A}} &:= \{B(a) \mid \langle B(a), d \rangle \in \mathcal{T}, d > 0\} \cup \{P(a, b) \mid \langle P(a, b), d \rangle \in \mathcal{T}, d > 0\}. \end{aligned}$$

That is, $\widehat{\mathcal{O}}$ contains all the axioms and assertions from \mathcal{O} which hold with a positive

degree—note that any fuzzy axiom or assertion with degree 0 could be removed w.l.o.g. anyway. The following result is a direct consequence of work on more expressive fuzzy DLs (Borgwardt et al. 2015).

Proposition 15

Let \mathcal{O} be a G-DL-Lite_R or II-DL-Lite_R ontology. Then \mathcal{O} is consistent iff $\widehat{\mathcal{O}}$ is consistent.

In those cases, consistency checking can be reduced to the classical case, without the need to modify the query or the basic formulation of the ontology. For the ontology \mathcal{O}_0 in Example 10, we have $\widehat{\mathcal{O}}_0 = (\{A_1 \sqsubseteq \neg A_2\}, \{A_1(a), A_2(a)\})$, which is inconsistent in the classical case, thus showing (through Proposition 15) that it is inconsistent under the Gödel and product t-norm semantics. We note that the example also shows that Proposition 15 does not hold for the Łukasiewicz t-norm, since we have established that \mathcal{O}_0 is consistent under this semantics, although its classical version remains inconsistent under classical interpretations.

A particular consequence of Proposition 15 is that deciding consistency of G-DL-Lite_R and II-DL-Lite_R ontologies is in AC⁰ w.r.t. data complexity, and NLOGSPACE-complete w.r.t. combined complexity, where the NLOGSPACE lower bound comes from known results in classical DL-Lite_R (Artale et al. 2009). Thus adding truth degrees does not affect the complexity of this basic reasoning task. We now turn our attention to the task of query answering with the different semantics, starting with the idempotent case of the Gödel t-norm. We consider first the case of conjunctive queries, which allows for a simple solution, and then study threshold queries for which a rewriting technique is needed.

Before studying how to answer queries over fuzzy DL-Lite_R ontologies and its complexity, we note that in the case that an ontology is classical—i.e., it uses only degree 1 in all its axioms—its canonical interpretation constructed as described in this section is equivalent to the classical canonical interpretation from (Calvanese et al. 2007). This fact will be used in the following sections.

4 Answering Conjunctive Queries over Gödel Ontologies

For this and the following section, we are always considering the Gödel t-norm as the underlying operator for interpreting all fuzzy statements, and in particular the conjunctive queries.

The Gödel semantics are very limited in their expressivity. On the one hand, we have seen that $\ominus d \in \{0, 1\}$ for all $d \in [0, 1]$. This means that whenever we have an axiom of the form $\langle B \sqsubseteq \neg B', d \rangle$ or $\langle Q \sqsubseteq \neg Q', d \rangle$ with $d > 0$, we are in fact saying that for every element $\delta \in \Delta^{\mathcal{I}}$, if $B^{\mathcal{I}}(\delta) > 0$, then $B'^{\mathcal{I}}(\delta) = 0$ —because in this case $\ominus B'^{\mathcal{I}}(\delta) = 1$, which is the only possible way of satisfying the axiom. A similar argument holds for role axioms. Thus, for this section we can assume w.l.o.g. that all negative axioms hold with degree 1; i.e., they are of the form $\langle B \sqsubseteq \neg B', 1 \rangle$ or $\langle Q \sqsubseteq \neg Q', 1 \rangle$. On the other hand, a positive axiom of the form $\langle B \sqsubseteq B', d \rangle$ requires that for every $\delta \in \Delta^{\mathcal{I}}$, $B^{\mathcal{I}}(\delta) \geq \min\{B'^{\mathcal{I}}(\delta), d\}$. That is, the only way to guarantee that an atom gets a high degree is to use axioms with a high degree. We use these facts to reduce reasoning tasks in this setting to the classical DL-Lite_R scenario.

Consider a consistent G-DL-Lite_R ontology \mathcal{O} . We can decide a lower bound for the degree of a CQ simply by querying a *cut* of \mathcal{O} .

Definition 16 (cut ontology)

Given a value $\theta \in (0, 1]$, the θ -cut of the ontology \mathcal{O} is defined as the sub-ontology $\mathcal{O}_{\geq\theta} := (\mathcal{T}_{\geq\theta}, \mathcal{A}_{\geq\theta})$ where

$$\begin{aligned}\mathcal{T}_{\geq\theta} &:= \{\langle\gamma, e\rangle \in \mathcal{T} \mid e \geq \theta\}, \\ \mathcal{A}_{\geq\theta} &:= \{\langle\alpha, e\rangle \in \mathcal{A} \mid e \geq \theta\}.\end{aligned}$$

That is, $\mathcal{O}_{\geq\theta}$ is the subontology containing only the axioms and assertions that hold to degree at least θ . To show that θ -cuts suffice for answering queries, we use the canonical interpretation.

Note that including new axioms or assertions to an ontology would result in an update of the canonical interpretation which only increases the degree of some of the elements of the domain. More precisely, if $\mathcal{I}_{\text{can}}(\mathcal{O})$ is the canonical interpretation of $\mathcal{O} = (\mathcal{T}, \mathcal{A})$, then the canonical interpretation of $\mathcal{O}' = (\mathcal{T} \cup \{\langle B \sqsubseteq C, d \rangle\}, \mathcal{A})$ is the result of applying the construction rules starting from $\mathcal{I}_{\text{can}}(\mathcal{O})$. This holds because the resulting canonical interpretation is not dependent on the order in which rules are applied (and hence axioms taken into account) as long as this is done fairly.⁵ Since $\mathcal{I}_{\text{can}}(\mathcal{O})$ has already applied all the rules on axioms of \mathcal{O} exhaustively, the only remaining rule applications will be based on the new axiom $\langle B \sqsubseteq C, d \rangle$ and new applications over \mathcal{T} arising from it. Under the Gödel semantics, all the updates increase the interpretation function up to the value d ; that is, if $\mathcal{I}'_{\text{can}}$ is the interpretation function of $\mathcal{I}_{\text{can}}(\mathcal{O}')$, the difference between $\mathcal{I}_{\text{can}}(\mathcal{O})$ and $\mathcal{I}_{\text{can}}(\mathcal{O}')$ is that there exist some elements such that $A^{\mathcal{I}_{\text{can}}}(\delta) < A^{\mathcal{I}'_{\text{can}}}(\delta) = d$, and similarly for roles there exist some pairs δ, η such that $P^{\mathcal{I}_{\text{can}}}(\delta, \eta) < P^{\mathcal{I}'_{\text{can}}}(\delta, \eta) = d$. For all others, the degrees remain unchanged. Moreover, if d_0 is the smallest degree appearing in the ontology \mathcal{O} , then its canonical interpretation uses only truth degrees in $\{0\} \cup [d_0, 1]$; that is, no truth degree in $(0, d_0)$ appears in $\mathcal{I}_{\text{can}}(\mathcal{O})$. With these insights we are ready to produce our first results. Recall, once again, that for the rest of this section, we always consider that the semantics is based on the Gödel t-norm; i.e., we have a G-DL-Lite_R ontology.

Lemma 17

Let \mathcal{O} be a consistent G-DL-Lite_R ontology, $q(\mathbf{x})$ a query, \mathbf{a} a tuple of individuals, and $\theta \in (0, 1]$. Then $\mathcal{O} \models q(\mathbf{a}) \geq \theta$ iff $\mathcal{O}_{\geq\theta} \models q(\mathbf{a}) \geq \theta$.

Proof

Since $\mathcal{O}_{\geq\theta} \subseteq \mathcal{O}$, every model of \mathcal{O} is also a model of $\mathcal{O}_{\geq\theta}$. Hence, if $\mathcal{O}_{\geq\theta} \models q(\mathbf{a}) \geq \theta$, then $\mathcal{O} \models q(\mathbf{a}) \geq \theta$.

For the converse, assume that $\mathcal{O}_{\geq\theta} \not\models q(\mathbf{a}) \geq \theta$. By Corollary 12, this means that $\mathcal{I}_{\text{can}}(\mathcal{O}_{\geq\theta}) \not\models q(\mathbf{a}) \geq \theta$. That is, $q^{\mathcal{I}_{\text{can}}}(\mathbf{a}^{\mathcal{I}_{\text{can}}}) < \theta$. Let $\mathcal{I}_{\text{can}}(\mathcal{O}) = (\Delta^{\mathcal{I}_{\text{can}}}, \mathcal{I}'_{\text{can}})$ be the canonical interpretation of \mathcal{O} . Recall that the difference between \mathcal{O} and $\mathcal{O}_{\geq\theta}$ is that the former has some additional axioms with degrees smaller than θ . As argued before, this means that the difference between $\mathcal{I}_{\text{can}}(\mathcal{O})$ and $\mathcal{I}_{\text{can}}(\mathcal{O}_{\geq\theta})$ are just some degrees, which are all smaller than θ ; that is, for every $A \in N_C$, $P \in N_R$, and $\delta, \eta \in \Delta^{\mathcal{I}'_{\text{can}}}$, if $A^{\mathcal{I}'_{\text{can}}}(\delta) \geq \theta$, then $A^{\mathcal{I}_{\text{can}}}(\delta) \geq \theta$ and if $P^{\mathcal{I}'_{\text{can}}}(\delta, \eta) \geq \theta$, then $P^{\mathcal{I}_{\text{can}}}(\delta, \eta) \geq \theta$. By assumption, this means that $q^{\mathcal{I}'_{\text{can}}}(\mathbf{a}^{\mathcal{I}'_{\text{can}}}) < \theta$, and hence $\mathcal{I}_{\text{can}}(\mathcal{O}) \not\models q(\mathbf{a}) \geq \theta$. Thus, $\mathcal{O} \not\models q(\mathbf{a}) \geq \theta$. \square

⁵ Formally, different rule application orderings yield homomorphic interpretations. See the discussion in the previous section.

What this lemma states is that in order to find a lower bound for the degree of a query, one can ignore all the axioms and assertions that provide a smaller degree than the bound we are interested in. However, one still needs to answer a query for a fuzzy ontology ($\mathcal{O}_{\geq\theta}$ is still fuzzy), for which we still do not have any effective method. The following lemma solves this issue, considering the classical version of this ontology.

Lemma 18

Let \mathcal{O} be a consistent G-DL-Lite_R ontology such that $\mathcal{O}_{\geq\theta} = \mathcal{O}$ for some $\theta > 0$. Then, $\mathcal{O} \models q(\mathbf{a}) \geq \theta$ iff $\widehat{\mathcal{O}} \models q(\mathbf{a})$.

Proof

Every model of $\widehat{\mathcal{O}}$ is also a model of \mathcal{O} , with the additional property that the interpretation function maps all elements to $\{0, 1\}$. If $\mathcal{O} \models q(\mathbf{a}) \geq \theta > 0$, then for every model \mathcal{I} of $\widehat{\mathcal{O}}$ it holds that $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) \geq \theta > 0$, and thus $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) = 1$, which means that $\widehat{\mathcal{O}} \models q(\mathbf{a})$.

Conversely, if $\widehat{\mathcal{O}} \models q(\mathbf{a})$, the canonical interpretation $\mathcal{I}_{\text{can}}(\mathcal{O})$ must be such that $q^{\mathcal{I}_{\text{can}}}(\mathbf{a}^{\mathcal{I}_{\text{can}}}) > 0$; but as argued before, since \mathcal{O} only has axioms and assertions with degrees $\geq \theta$, it must be the case that all degrees of $\mathcal{I}_{\text{can}}(\mathcal{O})$ are in $\{0\} \cup [\theta, 1]$, and hence $q^{\mathcal{I}_{\text{can}}}(\mathbf{a}^{\mathcal{I}_{\text{can}}}) \geq \theta$. This implies, by Corollary 12 that $\mathcal{O} \models q(\mathbf{a}) \geq \theta$. \square

Note that the condition of this lemma, which requires that $\mathcal{O}_{\geq\theta} = \mathcal{O}$, is only stating that all the degrees in the ontology \mathcal{O} are at least θ . That condition is immediately satisfied by a cut ontology, and hence the lemma can be applied directly to it.

Lemmas 17 and 18 together provide a method for reducing answering degree queries over G-DL-Lite_R ontologies to query answering in classical DL-Lite_R.

Theorem 19

If \mathcal{O} is a consistent G-DL-Lite_R ontology and $\theta > 0$, then it holds that $\mathcal{O} \models q(\mathbf{a}) \geq \theta$ iff $\widehat{\mathcal{O}}_{\geq\theta} \models q(\mathbf{a})$.

This means that we can use a standard ontology-based query answering system to answer fuzzy queries in DL-Lite_R as well. Note that the approach proposed by Theorem 19 can only decide whether the degree of an answer to a query is at least θ , but it needs the value $\theta \in (0, 1]$ as a parameter. If, instead, we are interested in computing the degree of an answer, or $\text{ans}(q(\mathbf{x}), \mathcal{O})$, we can still use a classical query answering method as an underlying black-box aid as described next.

Since the TBox \mathcal{T} and the ABox \mathcal{A} which compose the ontology \mathcal{O} are both finite, the set $\mathcal{D} := \{d \mid \langle \alpha, d \rangle \in \mathcal{T} \cup \mathcal{A}\}$ of degrees appearing in the ontology is also finite; in fact, its size is bounded by the size of \mathcal{O} . Hence, we can assume that \mathcal{D} is of the form $\mathcal{D} = \{d_0, d_1, \dots, d_n, d_{n+1}\}$ where $d_0 \geq 0, d_{n+1} = 1$ and for all $i, 0 \leq i \leq n, d_i < d_{i+1}$. In order to find the degree of an answer \mathbf{a} to a query q , we proceed as follows: starting from $i := n + 1$, we iteratively ask the query $\mathcal{O}_{\geq d_i} \models q(\mathbf{a})$ and decrease i until the query is answered affirmatively, or i becomes 0 (see Algorithm 1). In the former case, d_i is the degree for $q(\mathbf{a})$; in the latter, the degree is 0—i.e., \mathbf{a} is not an answer of q .⁶

During the execution of this algorithm, each classical query needed at line 3 can be executed in AC⁰ (and in particular in LOGSPACE) in the size of the data; i.e., the

⁶ Note that the algorithm can be made more efficient using a binary search, instead of a linear decrease of available degrees. We chose this presentation to provide a clear association with Corollary 21.

Algorithm 1: Compute the degree of an answer to a query

Data: Ontology \mathcal{O} , query q , answer \mathbf{a} , $\mathcal{D} = \{d_0, d_1, \dots, d_{n+1}\}$
Result: The degree of $q(\mathbf{a})$ w.r.t. \mathcal{O}

- 1 $i \leftarrow n + 1$
- 2 $\mathcal{N} \leftarrow \widehat{\mathcal{O}}_{\geq 1}$
- 3 **while** $\mathcal{N} \not\models q(\mathbf{a})$ **and** $i > 0$ **do**
- 4 $i \leftarrow i - 1$
- 5 $\mathcal{N} \leftarrow \widehat{\mathcal{O}}_{\geq d_i}$
- 6 **return** d_i

ABox as shown in (Artale et al. 2009). The iterations in the loop do not affect the overall space used, as one can simply produce a new query every time and clean up the previous information. Overall, this means that the degree of an answer can be computed in LOGSPACE in data complexity, using a classical query answering engine.

Corollary 20

The degree of an answer \mathbf{a} to a query q w.r.t. the G-DL-Lite_R ontology \mathcal{O} is computable in logarithmic space w.r.t. the size of the ABox (i.e., in data complexity).

We will later see that this upper bound can indeed be reduced to AC⁰ by seeing a degree query as a special case of a threshold query. However, the method that provides a tight complexity bound requires a new implementation of the rewriting approach, with all its associated optimizations, in contrast to the method from Algorithm 1, which can simply call any existing classical tool; e.g. (De Giacomo et al. 2012; Calvanese et al. 2015).

Computing the whole set of pairs $\text{ans}(q(\mathbf{x}), \mathcal{O})$ is a more complex task. Although we can follow an approach similar to Algorithm 1, where the answers to $q(\mathbf{x})$ are computed for each ontology $\widehat{\mathcal{O}}_{\geq d_i}$, in order to assign the appropriate degree to each answer, we need to either keep track of all the answers found so far, or add a negated query which excludes the answers with a higher degree. In both cases, we require a different approach and a potential larger use of memory. On the other hand, the whole set of answers $\text{ans}(q(\mathbf{x}), \mathcal{O})$ will usually contain many answers that hold with a very low degree, which may not be of much interest to the user making the query. When dealing with degrees, a more meaningful task is to find the k answers with the highest degree, for some natural number k ; i.e., the *top- k answers* of q .

Algorithm 1 once again suggests a way to compute the top- k answers. As in the algorithm, one starts with the highest possible degree, and expands the classical ontology by including the axioms and assertions with a lower degree. The difference is that one stops now when the query returns at least k tuples as answers. At that point, the tuples found are those with the highest degree for the query. As before, each of these queries can be answered in AC⁰ in data complexity, which yields a LOGSPACE upper bound for answering top- k queries in data complexity.

Corollary 21

Top- k queries over consistent G-DL-Lite_R ontologies can be answered in logarithmic space w.r.t. the size of the ABox.

5 Threshold Queries over Gödel Semantics

We now turn our attention to the case of threshold queries, but keeping the assumption of the Gödel semantics in place. The first thing to notice when considering threshold queries is that the simple approach developed in the previous section, where one calls a classical query answering engine over a cut subontology, cannot work. Indeed, as each atom needs to be satisfied potentially to a different degree, there is no one cut that can suffice to answer them all. Indeed, we have already seen a TQ in Example 8 which has no answers even though the natural cut ontology provides one answer.

When considering Boolean threshold queries, it may be tempting to simply try to verify each threshold atom separately through a cut ontology. However, such an approach is not sound due to the existentially quantified variables which need to be associated to a (potentially anonymous) individual. This problem is not new, as it arises already for conjunctive queries over classical databases.

To answer these queries, we will adapt the query rewriting technique from the classical setting. The underlying idea is essentially the same, as described previously in this paper, where an atom $B(x)$ may be substituted by an atom $C(x)$ if the TBox contains the axiom $C \sqsubseteq B$. However, one has to be careful with the degrees used. In fact, some axioms may not be applied during the rewriting, if their degree is not large enough.

Example 22

Consider once again the TBox \mathcal{T}_{exa} from Example 3, and suppose that we are interested in finding all popular attractions, up to a given degree $d \in [0, 1]$; that is, we have the query $q(x) = \text{Popular}(x) \geq d$. The TBox contains the axiom $\langle \text{Museum} \sqsubseteq \text{Popular}, 0.6 \rangle$. This means that answers to $q(x)$ w.r.t. this TBox should also include the answers to $\text{Museum}(x)$, but this depends on the value of d , as we explain next.

Suppose that $d > 0.6$; e.g., if we have $q(x) = \text{Popular}(x) \geq 0.7$. For an individual a and any model \mathcal{I} of the ontology, we have no guarantee that $\text{Popular}^{\mathcal{I}}(a^{\mathcal{I}}) \geq 0.7$ regardless of the degree of $\text{Museum}^{\mathcal{I}}(a^{\mathcal{I}})$. Indeed, even if $\text{Museum}^{\mathcal{I}}(a^{\mathcal{I}}) = 1$, the only thing that can be guaranteed is that the degree of a belonging to Popular is at least 0.6, which does not suffice to become a positive answer to the query. Hence, there is no need to include Museum in the rewriting.

Suppose now that $d \leq 0.6$; e.g., with the query $q(x) = \text{Popular}(x) \geq 0.5$. In this case, we note that every individual a such that $\text{Museum}^{\mathcal{I}}(a^{\mathcal{I}}) \geq 0.5$ must satisfy also that $\text{Popular}^{\mathcal{I}}(a^{\mathcal{I}}) \geq 0.5$. Indeed, recall that under the Gödel semantics, $f \Rightarrow e$ is either e if $e \leq f$ or 1 otherwise. Since \mathcal{I} satisfies the axiom $\langle \text{Museum} \sqsubseteq \text{Popular}, 0.6 \rangle$, whenever $f = \text{Museum}^{\mathcal{I}}(a^{\mathcal{I}}) \geq 0.5$ holds, we know that the degree e of $\text{Popular}^{\mathcal{I}}(a^{\mathcal{I}})$ must be such that $f \Rightarrow e \geq 0.6$. If $e \leq f$, this can only be true if $e \geq 0.6 > 0.5$. Otherwise, we know that $e > f \geq 0.5$, and hence any individual belonging to the concept Museum to degree at least 0.5 is an answer to the query $q(x)$.

This example shows that during the rewriting process, we only need to consider the axioms that hold to a degree greater than the threshold of the current atom of interest. During the rewriting step, the original threshold is preserved regardless of the bound from the axioms. We now proceed to describe the rewriting process in detail, following the ideas developed originally for classical DL-Lite_R and other members of the DL-Lite family through the `PerfectRef` algorithm. To aid understanding from readers knowledgeable with

- If $\gamma = A(x) \geq d$ and $\alpha = \langle A_1 \sqsubseteq A, e \rangle$, then $gr(\gamma, \alpha) = A_1(x) \geq d$
- If $\gamma = A(x) \geq d$ and $\alpha = \langle \exists P \sqsubseteq A, e \rangle$, then $gr(\gamma, \alpha) = P(x, _) \geq d$
- If $\gamma = A(x) \geq d$ and $\alpha = \langle \exists P^- \sqsubseteq A, e \rangle$, then $gr(\gamma, \alpha) = P(_, x) \geq d$
- If $\gamma = P(x, _) \geq d$ and $\alpha = \langle A \sqsubseteq \exists P, e \rangle$, then $gr(\gamma, \alpha) = A(x) \geq d$
- If $\gamma = P(x, _) \geq d$ and $\alpha = \langle \exists P_1 \sqsubseteq \exists P, e \rangle$, then $gr(\gamma, \alpha) = P_1(x, _) \geq d$
- If $\gamma = P(x, _) \geq d$ and $\alpha = \langle \exists P_1^- \sqsubseteq \exists P, e \rangle$, then $gr(\gamma, \alpha) = P_1(_, x) \geq d$
- If $\gamma = P(_, x) \geq d$ and $\alpha = \langle A \sqsubseteq \exists P^-, e \rangle$, then $gr(\gamma, \alpha) = A(x) \geq d$
- If $\gamma = P(_, x) \geq d$ and $\alpha = \langle \exists P_1 \sqsubseteq \exists P^-, e \rangle$, then $gr(\gamma, \alpha) = P_1(x, _) \geq d$
- If $\gamma = P(_, x) \geq d$ and $\alpha = \langle \exists P_1^- \sqsubseteq \exists P^-, e \rangle$, then $gr(\gamma, \alpha) = P_1(_, x) \geq d$
- If $\gamma = P(x_1, x_2) \geq d$ and $\alpha \in \{ \langle P_1 \sqsubseteq P, e \rangle, \langle P_1^- \sqsubseteq P^-, e \rangle \}$ then $gr(\gamma, \alpha) = P_1(x_1, x_2) \geq d$
- If $\gamma = P(x_1, x_2) \geq d$ and $\alpha \in \{ \langle P_1 \sqsubseteq P^-, e \rangle, \langle P_1^- \sqsubseteq P, e \rangle \}$ then $gr(\gamma, \alpha) = P_1(x_2, x_1) \geq d$

Fig. 4. The result $gr(\gamma, \alpha)$ of applying the axiom α to the threshold atom γ .

the original method, we preserve as much of the terminology from (Calvanese et al. 2007) as possible. From now on, in a query $q(\mathbf{x}) = \varphi(\mathbf{x}, \mathbf{y})$, we call all the variables in \mathbf{x} *distinguished*, and any variable that appears at least twice within a query *shared*. Note that there is no need to keep track of the variables that are not distinguished nor shared (from now on, called undistinguished, unshared variables); it is only relevant that they can be adequately assigned a value. Hence, those variables will be denoted by an underscore ('_'), and use $y = _$ to express that y is one such variable.

Definition 23 (applicability)

An axiom α is *applicable* to the threshold atom $A(x) \geq d$ iff α is of the form $\langle C \sqsubseteq A, e \rangle$ and $d \leq e$. It is *applicable* to the threshold atom $P(x_1, x_2)$ iff either (i) $x_2 = _$ and α is of the form $\langle C \sqsubseteq \exists P, e \rangle$ with $d \leq e$; (ii) $x_1 = _$ and α is of the form $\langle C \sqsubseteq \exists P^-, e \rangle$; or (iii) α is of the form $\langle Q \sqsubseteq P, e \rangle$ or $\langle Q \sqsubseteq P^-, e \rangle$ with $d \leq e$.

If α is applicable to the threshold atom γ , the *result* of the application is the atom $gr(\gamma, \alpha)$ defined through the rules in Figure 4.

The PerfectRef algorithm constructs a union of threshold queries by iteratively substituting atoms γ for which an axiom α is applicable, with the result $gr(\gamma, \alpha)$ of the application. This follows the idea of tracing backwards the axioms in order to absorb the TBox into the query which was previously outlined. The pseudocode for PerfectRef is more formally described in Algorithm 2. In the algorithm, $q[\gamma, \eta]$ is the query resulting from substituting in q the atom γ with the atom η . The function $reduce(p, \gamma_1, \gamma_2)$ called in line 10 simply returns the query obtained by applying the most general unifier between γ_1 and γ_2 to p . For unification, all nondistinguished, unshared variables are considered different. For simplicity, we always assume that all nondistinguished, unshared variables are known, and hence call them $_$ when testing applicability.

Note that, just as in the classical case, the application of the *reduce* function is necessary to guarantee correctness of the rewriting. Specifically, a variable that is bound in a query p may become unbound after the unification process, which may allow more axioms to be applied for the rewriting.

Once again, the algorithm takes as input a threshold query q , and returns a union of threshold queries T which is constructed by taking into account the information from the TBox \mathcal{T} . The importance of this rewriting is that at this point, the answers to the original query q w.r.t. an ontology $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ can be obtained by applying the query T to the ABox \mathcal{A} , seen as a standard database.

Algorithm 2: PerfectRef

Data: Threshold query q , G-DL-Lite_R TBox \mathcal{T}
Result: Union of threshold queries T

```

1  $T \leftarrow \{q\}$ 
2 repeat
3    $T' \leftarrow T$ 
4   for each  $p \in T'$  do
5     for each  $\gamma \in p$ , and each  $\alpha \in \mathcal{T}$  do
6       if  $\alpha$  is applicable to  $\gamma$  then
7          $T \leftarrow T \cup \{p[\gamma/gr(\gamma, \alpha)]\}$ 
8       for each  $\gamma_1, \gamma_2 \in p$  do
9         if  $\gamma_1$  and  $\gamma_2$  unify then
10         $T \leftarrow T \cup \{reduce(p, \gamma_1, \gamma_2)\}$ 
11 until  $T' = T$ 
12 return  $T$ 
    
```

Let $db(\mathcal{A})$ be the ABox \mathcal{A} seen as a database. Note that since we have fuzzy assertions, the database will contain binary relations (representing concept assertions) and ternary relations (representing the role assertions), where the last element of the relation is the degree; a number in the interval $[0, 1]$. Under this view, a threshold query can also be seen as a conjunctive query, taking into account the inequalities in the selection. Given a union of threshold queries T , $UCQ(T)$ denotes the fact that T is being read as a UCQ in this sense. Given an ABox \mathcal{A} and a union of TQs T , we denote by $\text{ans}(db(\mathcal{A}), UCQ(T))$ the set of answers to T w.r.t. $db(\mathcal{A})$ from a database perspective. We also denote by $\text{ans}(q, \mathcal{O})$ the set of answers to the TQ q w.r.t. the ontology \mathcal{O} . We then obtain the following result.

Theorem 24

Let $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ be a consistent G-DL-Lite_R ontology, q a TQ, and T the union of TQs obtained through the rewriting. Then $\text{ans}(q, \mathcal{O}) = \text{ans}(db(\mathcal{A}), UCQ(T))$.

A consequence of Theorem 24 is that, in terms of data complexity, answering a TQ w.r.t. a DL-Lite_R ontology is at most as costly as answering a CQ over a database. Indeed, note that although the query q is transformed into a larger UCQ, the data itself remains unchanged. This yields the following result.

Theorem 25

Answering threshold queries w.r.t. consistent G-DL-Lite_R ontologies is in AC^0 w.r.t. data complexity.

Before finishing this section, we return to a question on complexity left open in the previous section; namely, the precise complexity of finding the degree of an answer to a conjunctive query. To answer this question, we first note that under the Gödel semantics, we can always see a degree query as a special case of a threshold query.

Given a CQ q , let $\text{At}(q)$ be the set of all the atoms in q . For a degree $d \in [0, 1]$, we can

define the TQ $TQ(q, d) = \bigwedge_{\gamma \in \text{At}(q)} \gamma \geq d$. That is, $TQ(q, d)$ uses the same atoms as q , but assigns a minimum degree of d to each of them. Since the Gödel semantics interprets the conjunction through the minimum operator, any answer of $TQ(q)$ yields a degree of at least d to the original query q .

Lemma 26

Let \mathcal{O} be a consistent G-DL-Lite_R ontology, q a CQ, \mathbf{a} an answer tuple, and $d \in [0, 1]$. It holds that $\mathcal{O} \models q(\mathbf{a}) \geq d$ iff $\mathcal{O} \models TQ(q(\mathbf{a}), d)$.

In order to find the degree of an answer, we can simply add as an answer variable after the rewriting one that looks at the degrees from the database $db(\mathcal{A})$. This does not affect the overall data complexity, and hence remains in AC⁰.

Corollary 27

Answering conjunctive queries w.r.t. consistent G-DL-Lite_R ontologies is in AC⁰ in data complexity.

This finishes our analysis of the Gödel t-norm, which also provides our main results. In the following section we briefly visit the case where the underlying t-norm is not idempotent, and showcase that in general dealing with such semantics becomes harder.

6 Non-idempotent t-norms

We now move our attention to the t-norms that are not idempotent; in particular the product and Łukasiewicz t-norms. Unfortunately, as we will see, the correctness of the reductions and algorithms presented in the previous sections rely strongly on the idempotency of the Gödel t-norm, and does not transfer directly to the other cases. However, at least for the product t-norm, it is still possible to answer some kinds of queries efficiently.

First recall that Proposition 15 holds for the product t-norm as well. Hence, deciding consistency of a Π -DL-Lite_R ontology remains reducible to the classical case and thus, efficient. We now show with simple examples that the other results do not transfer so easily.

Example 28

Consider the ontology $\mathcal{O}_{\text{exb}} := (\mathcal{T}_{\text{exb}}, \mathcal{A}_{\text{exb}})$ where $\mathcal{T}_{\text{exb}} := \{\langle A_i \sqsubseteq A_{i+1}, 0.9 \rangle \mid 0 \leq i < n\}$ and $\mathcal{A}_{\text{exb}} := \{\langle A_0(a), 1 \rangle\}$. Note that $\mathcal{O}_{\text{exb}} = (\mathcal{O}_{\text{exb}})_{\geq 0.9}$, but the degree for the query $q() = A_n(a)$ is 0.9^n which can be made arbitrarily small by making n large.

Similarly, it is not possible to find the top- k answers simply by layering the θ -cuts for decreasing values of θ until enough answers can be found.

Example 29

Let $\mathcal{O}'_{\text{exb}} := (\mathcal{T}_{\text{exb}}, \mathcal{A}'_{\text{exb}})$, where $\mathcal{A}'_{\text{exb}} := \mathcal{A}_{\text{exb}} \cup \{\langle A_n(b), 0.85 \rangle\}$ and $\mathcal{T}_{\text{exb}}, \mathcal{A}_{\text{exb}}$ are as in Example 28. The top answer for $q(x) = A_n(x)$ is b with degree 0.85, but from $(\mathcal{O}'_{\text{exb}})_{\geq 0.9}$ we already find the answer a , which is not the top one.

The main point with these examples is that, from the lack of idempotency of the t-norm \otimes , we can obtain low degrees in a match which arises from combining several axioms and assertions having a high degree. On the other hand, the product behaves well for positive values in the sense that applying the t-norm to two positive values always results in a

positive value; formally, if $d, e > 0$, then $d \otimes e > 0$. Thus, if we are only interested in knowing whether the result of a query is positive or not, there is no difference between the Gödel t-norm and the product t-norm.

Definition 30

A tuple \mathbf{a} is a *positive answer* to the query $q(\mathbf{x})$ w.r.t. the ontology \mathcal{O} (denoted by $\mathcal{O} \models q(\mathbf{a}) > 0$) iff for every model \mathcal{I} of \mathcal{O} it holds that $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) > 0$.

Theorem 31

If \mathcal{O} is a consistent II-DL-Lite_R ontology, then $\mathcal{O} \models q(\mathbf{a}) > 0$ iff $\widehat{\mathcal{O}} \models q(\mathbf{a})$.

Proof

Every model of $\widehat{\mathcal{O}}$ is also a model of \mathcal{O} , with the additional property that the interpretation function maps all elements to $\{0, 1\}$. If $\mathcal{O} \models q(\mathbf{a}) > 0$, then for every model \mathcal{I} of $\widehat{\mathcal{O}}$ it holds that $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) > 0$ and thus $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) = 1$, which means that $\widehat{\mathcal{O}} \models q(\mathbf{a})$.

Conversely, if $\widehat{\mathcal{O}} \models q(\mathbf{a})$, then the canonical interpretation is such that $q^{\mathcal{I}_{\text{can}}}(\mathbf{a}^{\mathcal{I}_{\text{can}}}) > 0$, and hence for every model \mathcal{I} it also holds that $q^{\mathcal{I}}(\mathbf{a}^{\mathcal{I}}) > 0$. \square

This means that, for the sake of answering positive queries over the product t-norm, one can simply ignore all the truth degrees and answer a classical query using any state-of-the-art engine. In particular, this means that positive answers can be found in AC^0 in data complexity just as in the classical case.

We now briefly consider the Łukasiewicz t-norm, which is known to be the hardest to handle due to its involutive negation and nilpotence, despite being in many cases the most natural choice for fuzzy semantics (Borgwardt et al. 2017). As mentioned already, Proposition 15 does not apply to the Łukasiewicz t-norm. That is, there are consistent L-DL-Lite_R ontologies whose classical version is inconsistent (see Example 10). As a result, there is currently no known method for deciding consistency of these ontologies, let alone answering queries. The culprits for this are the involutive negation, which is weaker than the negation used in the other two t-norms, but also the nilpotence, which may combine positive degrees to produce a degree of 0. The latter also means that, even if one could check consistency, it is still not clear how to answer even positive queries.

Example 32

Consider the ontology $\mathcal{O}_2 := (\mathcal{T}_2, \mathcal{A}_2)$ where

$$\begin{aligned} \mathcal{T}_2 &:= \{ \langle A_0 \sqsubseteq A_1, 0.5 \rangle, \langle A_1 \sqsubseteq A_2, 0.5 \rangle \} \\ \mathcal{A}_2 &:= \{ \langle A_0(a), 1 \rangle \}. \end{aligned}$$

Note that \mathcal{O}_2 is consistent, but there is a model \mathcal{I} (e.g., the canonical interpretation) of this ontology which sets $A_2^{\mathcal{I}}(a^{\mathcal{I}}) = 0$. Hence, a is not a positive answer to the query $q(x) = A_2(x)$ even though it is an answer of $q(x)$ over $\widehat{\mathcal{O}}_2$.

Importantly, if we extend DL-Lite_R with the possibility of using conjunctions as constructors for complex concepts, one can show following the ideas from (Borgwardt et al. 2017; Borgwardt et al. 2014) that deciding consistency of a L-DL-Lite_R ontology is NP-hard in combined complexity even if negations are disallowed; see Appendix Appendix A for full details. In the classical case, this logic—which is called $\text{DL-Lite}_{\text{Horn}}$ —has a polynomial

time consistency problem (Artale et al. 2009). This gives an indication that dealing with \mathbb{L} -DL-Lite $_R$ may also lead to an increase in complexity.

Interestingly, the rewriting technique from Section 5 also works for other t-norms—modulo some basic modifications—when answering threshold queries. Recall, for example, that given an axiom $\langle A \sqsubseteq B, e \rangle$, and a threshold atom $B(x) \geq d$, if $e \geq d$ then the rewriting technique would substitute this atom with $A(x) \geq d$. Although this substitution is sound for the idempotent Gödel t-norm, it does not work directly for the other ones. For example, under the product t-norm, if we set $d = e = 0.9$ we note that guaranteeing $A^{\mathcal{I}}(x) \geq 0.9$ does not necessarily implies, in a model \mathcal{I} of $\langle A \sqsubseteq B, e \rangle$ that $B^{\mathcal{I}}(x) \geq 0.9$. Indeed, as long as $B^{\mathcal{I}}(x) \geq 0.81$, the axiom is satisfied in this case. A similar argument can be made for the Łukasiewicz t-norm. Hence, we need to increase the required degree for the rewritten atom.

Recall from the properties of the residuum that for every t-norm \otimes it holds that $A^{\mathcal{I}}(x) \Rightarrow B^{\mathcal{I}}(x) \geq e$ iff $B^{\mathcal{I}}(x) \geq A^{\mathcal{I}}(x) \otimes e$. Thus, to ensure that $B^{\mathcal{I}}(x) \geq d$ it suffices to guarantee that $A^{\mathcal{I}}(x) \otimes e \geq d$. In the case of the product t-norm, this is akin to the condition $A^{\mathcal{I}}(x) \geq d/e$. For the Łukasiewicz t-norm the condition translates to the inequality $A^{\mathcal{I}}(x) \geq \min\{1, d + 1 - e\}$. We can then apply the same PerfectRef algorithm, with a new definition of the function gr that changes the last degree (which is always $\geq d$ in Figure 4) with the new degree developed here. Overall, this yields the following complexity result.

Theorem 33

Answering threshold queries w.r.t. consistent DL-Lite ontologies is in AC^0 in data complexity.

Note that this theorem does not solve the problems sketched before for non-idempotent t-norms. Indeed, it is still not clear how to check for consistency of a \mathbb{L} -DL-Lite ontology. Moreover, this result cannot be used to answer CQs because the analogous to Lemma 26 does not hold. Indeed, suppose that we have a simple CQ with only two atoms:

$$q(x) = A(x) \wedge B(x).$$

To turn $q(x) \geq d$ into a TQ, we need to assign a threshold to each of the atoms. Note however that, under a non-idempotent t-norm, we cannot assign the same degree d to each atom, as their conjunction would become in fact lower than d . To be more precise consider the product t-norm and $d = 0.9$. Note that an answer to the TQ $A(x) \geq 0.9 \wedge B(x) \geq 0.9$ is *not* necessarily an answer to $q(x) \geq 0.9$ because there could be a model that assigns both atoms to degree 0.9; hence the product of those degrees is $0.81 < 0.9$. To use a TQ, we need to choose two degrees d_1, d_2 such that $d_1 \cdot d_2 = 0.9$, and construct the TQ $A(x) \geq d_1 \wedge B(x) \geq d_2$. But there are infinitely many choices to make in this regard, hence we cannot even construct a finite UTQ. Thus, unfortunately, although we are able to answer TQs efficiently (if the ontology is known to be consistent), degree queries remain an open problem for non-idempotent t-norms.

7 Conclusions

In this paper we have studied the problem of answering queries over fuzzy ontologies written in DL-Lite. Our goal was to cover the gap in this area left by previous research.

Indeed, although query answering w.r.t. ontologies is still an active topic, most work referring to fuzzy terminologies or ABoxes focused on the so-called Zadeh semantics, which does not preserve desired properties from the mathematical fuzzy logic point of view. To our knowledge, only Mailis and Turhan (Mailis and Turhan 2014; Mailis et al. 2015) have studied this problem based on t-norms, and found solutions based on the Gödel t-norm. However, they limited their approach to *classical* TBoxes. They left open the problems of dealing with graded TBoxes, handling threshold queries, and dealing with non-idempotent t-norms.

A second goal of our work was to reuse as much as possible the classical techniques, in order to avoid an implementation overhead when our algorithms will be, in future work, implemented and tested. As a result, we developed a method for answering degree queries which relies heavily on a classical query answering tool as a black box. Through this method, we can take advantage of all the existing optimisations and improvements from that area, and simply use a better tool whenever it becomes available without having to worry about the internal intricacies that make it work. In few words, our algorithm for answering CQs w.r.t. the Gödel semantics simply considers the classical version of the cut of the ontology. That is, the method ignores all axioms that hold to a degree lower than the threshold imposed, and then sees the remaining query answering question as a classical one. We emphasise that this approach works perfectly even if the TBox is graded. This means that our results improve those from (Mailis and Turhan 2014) by allowing for fuzzy TBox axioms and not requiring a new rewriting of the query.

Dealing with threshold queries, where each atom can be independently assigned a different degree, turns out to be more complex technically. In fact, we were not able to produce a direct reduction to a classical query answering problem—and it is unlikely that such a reduction could exist, given the nature of the graded axioms. However, we could still exploit the main ideas from the classical scenario, adapting the well-known **PerfectRef** method to the fuzzy scenario. In some sense **PerfectRef** absorbs the TBox into the query, forming a larger UCQ which can be answered classically, seeing the ABox as a database. In our case, we also need to take care of the degree at which the rewritten atoms should hold, when creating the new query. Under the Gödel semantics, it suffices to preserve the same degree from the original query, but for non-idempotent t-norms the degree has to be increased accordingly to avoid including spurious answers. Importantly, this shows that answering threshold queries w.r.t. consistent fuzzy ontologies is in the same complexity class (AC^0) in data complexity as for the classical case, regardless of the t-norm underlying the semantics. The only caveat is that it is not known how to verify consistency of a fuzzy DL-Lite ontology in general. The idempotency of the Gödel t-norm allowed us then to show that CQ answering w.r.t. consistent G-DL-Lite ontologies is also in AC^0 in data complexity. This latter bound does not hold for non-idempotent t-norms.

It is worth noting that the methods for answering degree and threshold queries both ultimately rely on a rewriting of the query to represent the information expressed in the TBox. While the rewriting does not affect the *data* complexity, it is well known that the UCQ obtained through **PerfectRef** may grow exponentially (Calvanese et al. 2007; Pérez-Urbina et al. 2010). This means that, depending on the instance, answering these queries may still be impractical. For that reason, different rewriting and answering techniques have been developed and tested; for example, rewriting into a Datalog program instead

of an UCQ (Gottlob and Schwentick 2012; Gottlob et al. 2011). Our approach for solving threshold queries, given its reliance on PerfectRef, suffers from the same drawbacks. In order to use more optimised approaches, it is necessary to study whether other existing rewritings can also be adapted to the fuzzy setting. On the other hand, the approach to degree queries is, as mentioned already, fully black box: we only need to call an unmodified classical query answering tool repeatedly. This allows us to directly *plug* whichever system performs best, without worrying about the implementation overhead of understanding and adapting the existing tools.

Through illustrative examples, we showed that dealing with CQs is in general harder when the underlying t-norm is not idempotent. The main issue is that there is no unique way to decide the bounds for the different degrees to which atoms should hold to satisfy the lower bound for a conjunction of atoms. The problem is exacerbated by the fact that it is not even clear how to decide consistency of ontologies under the nilpotent t-norm. Indeed, even in the absence of an ABox, the Lukasiewicz t-norm may impose upper bounds in the degrees of some assertions which are not obvious to detect, and could contradict other conditions.

As future work, we are also interested in implementing and testing our ideas, with the help of some fuzzy ontologies which will be developed for specific application domains.

References

- ABITEBOUL, S., HULL, R., AND VIANU, V. 1994. *Foundations of Databases*. Addison Wesley.
- ARTALE, A., CALVANESE, D., KONTCHAKOV, R., AND ZAKHARYASCHEV, M. 2009. The DL-Lite family and relations. *Journal of Artificial Intelligence Research* 36, 1–69.
- BAADER, F., CALVANESE, D., MCGUINNESS, D., NARDI, D., AND PATEL-SCHNEIDER, P., Eds. 2007. *The Description Logic Handbook: Theory, Implementation, and Applications*, Second ed. Cambridge University Press.
- BOBILLO, F., CERAMI, M., ESTEVA, F., GARCÍA-CERDAÑA, À., PEÑALOZA, R., AND STRACCIA, U. 2015. Fuzzy description logic. In *Handbook of Mathematical Fuzzy Logic Volume 3*, P. Cintula, C. G. Fermüller, and C. Noguera, Eds. Studies in Logic, vol. 58. College Publications.
- BOPPANA, R. B. AND SIPSER, M. 1990. The complexity of finite functions. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, J. van Leeuwen, Ed. Elsevier and MIT Press, 757–804.
- BORGWARDT, S. 2014. Fuzzy description logics with general concept inclusions. Ph.D. thesis, Technische Universität Dresden, Germany.
- BORGWARDT, S., CERAMI, M., AND PEÑALOZA, R. 2014. Many-valued horn logic is hard. In *Proceedings of the First Workshop on Logics for Reasoning about Preferences, Uncertainty, and Vagueness, PRUV 2014*, T. Lukasiewicz, R. Peñaloza, and A. Turhan, Eds. CEUR Workshop Proceedings, vol. 1205. CEUR-WS.org, 52–58.
- BORGWARDT, S., CERAMI, M., AND PEÑALOZA, R. 2015. The complexity of subsumption in fuzzy EL. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, Q. Yang and M. J. Wooldridge, Eds. AAAI Press, 2812–2818.
- BORGWARDT, S., CERAMI, M., AND PEÑALOZA, R. 2017. The complexity of fuzzy EL under the Lukasiewicz t-norm. *International Journal of Approximate Reasoning* 91, 179–201.
- BORGWARDT, S., DISTEL, F., AND PEÑALOZA, R. 2015. The limits of decidability in fuzzy description logics with general concept inclusions. *Artificial Intelligence* 218, 23–55.
- BORGWARDT, S. AND PEÑALOZA, R. 2017. Fuzzy description logics - A survey. In *Proceedings of the 11th International Conference on Scalable Uncertainty Management (SUM 2017)*, S. Moral, O. Pivert, D. Sánchez, and N. Marín, Eds. LNCS, vol. 10564. Springer, 31–45.

- CALVANESE, D., COGREL, B., KALAYCI, E. G., KOMLA-EBRI, S., KONTCHAKOV, R., LANTI, D., REZK, M., RODRIGUEZ-MURO, M., AND XIAO, G. 2015. OBDA with the ontop framework. In *23rd Italian Symposium on Advanced Database Systems, SEBD 2015, Gaeta, Italy, June 14-17, 2015*, D. Lembo, R. Torlone, and A. Marrella, Eds. Curran Associates, Inc., 296–303.
- CALVANESE, D., DE GIACOMO, G., LEMBO, D., LENZERINI, M., AND ROSATI, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automated Reasoning* 39, 3, 385–429.
- CERAMI, M. 2012. Fuzzy description logics from a mathematical fuzzy logic point of view. Ph.D. thesis, University of Barcelona.
- COOK, S. A. 1971. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*. STOC '71. Association for Computing Machinery, New York, NY, USA, 151–158.
- DE GIACOMO, G., LEMBO, D., LENZERINI, M., POGGI, A., ROSATI, R., RUZZI, M., AND SAVO, D. F. 2012. MASTRO: A reasoner for effective ontology-based data access. In *Proceedings of the 1st International Workshop on OWL Reasoner Evaluation (ORE-2012), Manchester, UK, July 1st, 2012*, I. Horrocks, M. Yatskevich, and E. Jiménez-Ruiz, Eds. CEUR Workshop Proceedings, vol. 858. CEUR-WS.org.
- FURST, M. L., SAXE, J. B., AND SIPSER, M. 1984. Parity, circuits, and the polynomial-time hierarchy. *Math. Syst. Theory* 17, 1, 13–27.
- GOTTLÖB, G., ORSI, G., AND PIERIS, A. 2011. Ontological queries: Rewriting and optimization. In *Proceedings of the 27th International Conference on Data Engineering, ICDE 2011*, S. Abiteboul, K. Böhm, C. Koch, and K. Tan, Eds. IEEE Computer Society, 2–13.
- GOTTLÖB, G. AND SCHWENTICK, T. 2012. Rewriting ontological queries into small nonrecursive datalog programs. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning KR 2012*, G. Brewka, T. Eiter, and S. A. McIlraith, Eds. AAAI Press.
- HÁJEK, P. 1998. *Metamathematics of Fuzzy Logic*. Trends in Logic, vol. 4. Kluwer.
- KLEMENT, E.-P., MESIAR, R., AND PAP, E. 2000. *Triangular Norms*. Trends in Logic, vol. 8. Springer.
- LUKASIEWICZ, T. AND STRACCIA, U. 2008. Managing uncertainty and vagueness in description logics for the semantic web. *Journal of Web Semantics* 6, 4, 291–308.
- MAILIS, T. P. AND TURHAN, A. 2014. Employing DL-Lite_R reasoners for fuzzy query answering. In *Proceedings of the 4th Joint International Conference on Semantic Technology (JIST 2014)*, T. Supnithi, T. Yamaguchi, J. Z. Pan, V. Wuwongse, and M. Buranarach, Eds. LNCS, vol. 8943. Springer, 63–78.
- MAILIS, T. P., TURHAN, A., AND ZENKER, E. 2015. A pragmatic approach to answering cqs over fuzzy dl-lite-ontologies - introducing flite. In *Proceedings of the 2015 International Workshop on Description Logics (DL'15)*, D. Calvanese and B. Konev, Eds. CEUR Workshop Proceedings, vol. 1350. CEUR-WS.org.
- MOSTERT, P. S. AND SHIELDS, A. L. 1957. On the structure of semigroups on a compact manifold with boundary. *Annals of Mathematics* 65, 1, 117–143.
- ORTIZ, M. AND ŠIMKUS, M. 2012. *Reasoning and Query Answering in Description Logics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1–53.
- PAN, J. Z., STAMOU, G. B., STOILLOS, G., AND THOMAS, E. 2007. Expressive querying over fuzzy dl-lite ontologies. In *Proceedings of the 2007 International Workshop on Description Logics (DL'07)*. CEUR Workshop Proceedings, vol. 250. CEUR-WS.org.
- PAPADIMITRIOU, C. H. 1994. *Computational Complexity*. Addison Wesley.
- PASI, G. AND PEÑALOZA, R. 2020. Query answering in fuzzy dl-lite with graded axioms. In *Rules and Reasoning - 4th International Joint Conference, RuleML+RR 2020, Oslo, Norway, June 29 - July 1, 2020, Proceedings*, V. Gutiérrez-Basulto, T. Kliegr, A. Soylu, M. Giese, and D. Roman, Eds. Lecture Notes in Computer Science, vol. 12173. Springer, 39–53.

- PÉREZ-URBINA, H., MOTIK, B., AND HORROCKS, I. 2010. Tractable query answering and rewriting under description logic constraints. *J. of Applied Logic* 8, 2, 186–209.
- STRACCIA, U. 2006. Towards top-k query answering in description logics: The case of DL-Lite. In *Proceedings of the 10th European Conference in Logics in Artificial Intelligence (JELIA 2006)*, M. Fisher, W. van der Hoek, B. Konev, and A. Lisitsa, Eds. LNCS, vol. 4160. Springer, 439–451.
- STRACCIA, U. 2012. Top-k retrieval for ontology mediated access to relational databases. *Inf. Sci.* 198, 1–23.
- TARSKI, A. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* 5, 2, 285 – 309.

Appendix A NP-Hardness for the Horn Case

In this appendix we show that ontology consistency in $\text{L-DL-Lite}_{\text{Horn}}$, a DL closely related to L-DLLR , is NP-hard. The proof builds on the idea originally developed for a finitely-valued logic (Borgwardt et al. 2014),⁷ extended with the methods from (Borgwardt et al. 2017; Borgwardt et al. 2015) to deal with infinitely-valued t-norms. For brevity, we consider only a restricted version of $\text{L-DL-Lite}_{\text{Horn}}$ which already suffices to show hardness.

A $\text{L-DL-Lite}_{\text{Horn}}$ GCI is an expression of the form $\langle B \sqsubseteq C, d \rangle$ where B, C are built through the grammar $B ::= A \mid B \sqcap B \mid \perp$, where $A \in N_C$; that is, they are conjunctions of concept names, and $d \in [0, 1]$. The notion of an ABox, a TBox, and an ontology are analogous to the DL-Lite_R setting. The semantics of this logic is defined as for DL-Lite_R , with the addition that the interpretation function is defined for conjunctions as $(B \sqcap C)^{\mathcal{I}}(\delta) := B^{\mathcal{I}}(\delta) \otimes C^{\mathcal{I}}(\delta)$ and for bottom as $\perp^{\mathcal{I}}(\delta) := 0$.

We show NP-hardness through a reduction from the well-known problem of satisfiability of 3-CNF formulas (Cook 1971). Very briefly, a 3-clause is a disjunction of exactly three literals (variables or negated variables) and a 3-CNF formula is a conjunction of 3-clauses.

The main idea behind the reduction is to use an intermediate degree greater than 0 to simulate the logical truth value *false* from classical propositional logic. This allows us to simulate the disjunction from the 3-clauses through a conjunction of concepts. This idea is formalised in full detail next.

Consider the set \mathcal{V} of propositional variables. For each $v \in \mathcal{V}$ we define two concept names A_v and A'_v with the intuitive meaning that A_v stands for the valuation making v true, and A'_v for the valuation making v false. We define the function ρ which maps propositional literals to concept names as:

$$\rho(\ell) := \begin{cases} A_v & \text{if } \ell = v \in \mathcal{V} \\ A'_v & \text{if } \ell = \neg v, v \in \mathcal{V}. \end{cases}$$

This function is extended to 3-clauses by defining $\rho(\ell_1 \vee \ell_2 \vee \ell_3) := \rho(\ell_1) \sqcap \rho(\ell_2) \sqcap \rho(\ell_3)$. Abusing the notation, we identify the 3-CNF formula φ with the set of 3-clauses it contains.

⁷ In reality, we use a simplified form based on 3-SAT, rather than the more general m-SAT from that work.

Let φ be a propositional formula in 3-CNF, and $\text{var}(\varphi)$ represent the set of all propositional variables appearing in φ . We construct the ontology \mathcal{O}_φ consisting of the very simple ABox $\mathcal{A}_\varphi := \{\langle A_0(a), 1 \rangle\}$ and the TBox

$$\mathcal{T}_\varphi := \{ \langle A_v \sqcap A_v \sqcap A_v \sqsubseteq A_v \sqcap A_v \sqcap A_v \sqcap A_v, 1 \rangle, \quad (\text{A1})$$

$$\langle A'_v \sqcap A'_v \sqcap A'_v \sqsubseteq A'_v \sqcap A'_v \sqcap A'_v \sqcap A'_v, 1 \rangle, \quad (\text{A2})$$

$$\langle A_v \sqcap A'_v \sqsubseteq \perp, 1/3 \rangle, \langle A_0 \sqsubseteq A_v \sqcap A'_v, 2/3 \rangle \mid v \in \text{var}(\varphi) \} \cup \quad (\text{A3})$$

$$\{ \langle A_0 \rightarrow \rho(c), 1/3 \rangle \mid c \in \varphi \} \quad (\text{A4})$$

Theorem 34

The formula φ is satisfiable iff the ontology $\mathcal{O}_\varphi = (\mathcal{A}_\varphi, \mathcal{T}_\varphi)$ is consistent.

Proof

We start with some observations about the TBox \mathcal{T}_φ . The axioms in lines (A1) and (A2) require that every model $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$ of \mathcal{T}_φ is such that $A_v^\mathcal{I}(\delta) \in [0, 2/3] \cup \{1\}$ and $A'_v{}^\mathcal{I}(\delta) \in [0, 2/3] \cup \{1\}$ for all $\delta \in \Delta^\mathcal{I}$.⁸ Given the ABox axiom, the GCIs from line (A3) guarantee that *exactly one* between $A_v^\mathcal{I}(a^\mathcal{I})$ and $A'_v{}^\mathcal{I}(a^\mathcal{I})$ is interpreted as 1, and the other one as 2/3. Intuitively, the value 2/3 will be read as “false” and 1 as “true.” Finally, the axioms in line (A4) guarantee that for every clause $c \in \varphi$, at least one of the conjuncts in $\rho(c)$ is interpreted as 1 (i.e., “true”) at the element $a^\mathcal{I}$. Now we prove the property.

If φ is satisfiable, let \mathcal{V} be a valuation making φ true. As customary, a valuation is a subset of $\text{var}(\varphi)$ that expresses which variables are mapped to true. We construct the interpretation $\mathcal{I}_\mathcal{V} = (\{a\}, \cdot^\mathcal{I})$ having a singleton domain, where $A_0^\mathcal{I}(a) = 1$ and for each $v \in \text{var}(\varphi)$

$$A_v^\mathcal{I}(a) = \begin{cases} 1 & \text{if } v \in \mathcal{V} \\ 2/3 & \text{otherwise} \end{cases}$$

$$A'_v{}^\mathcal{I}(a) = \begin{cases} 1 & \text{if } v \notin \mathcal{V} \\ 2/3 & \text{otherwise} \end{cases}$$

It is easy to see that this interpretation satisfies \mathcal{A}_φ and all the axioms in lines (A1)–(A3) of \mathcal{T}_φ . Since \mathcal{V} is a model of φ , for every 3-clause $c \in \varphi$, there exists a literal ℓ in c that \mathcal{V} maps to true. By construction, $(\rho(\ell))^\mathcal{I}(a) = 1$, and hence, $\mathcal{I}_\mathcal{V}$ is also a model of these axioms. Hence \mathcal{O}_φ is consistent.

Conversely, let $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$ be a model of \mathcal{O}_φ and let $\delta = a^\mathcal{I}$. We construct the propositional valuation $\mathcal{V}_\mathcal{I}$ as follows: for every $v \in \text{var}(\varphi)$, $v \in \mathcal{V}$ iff $A_v^\mathcal{I}(\delta) = 1$. By construction, for every $v \notin \mathcal{V}$ it holds that $A'_v{}^\mathcal{I}(\delta) = 1$. Since \mathcal{I} is a model of \mathcal{T}_φ , it satisfies the axioms in line (A4), and hence for each clause c , there is a literal ℓ such that $(\rho(\ell))^\mathcal{I}(\delta) = 1$. This is true iff $\mathcal{V}_\mathcal{I}$ makes ℓ true, and in particular the whole clause c true. Hence $\mathcal{V}_\mathcal{I}$ satisfies φ . \square

⁸ Intuitively, if the interpretation is between 2/3 and 1, then conjoining three times will necessarily yield a greater value than conjoining four times, due to monotonicity. The only way to avoid this is to make the conjunction reach 0, or stay in 1.