

Answering Vague Queries in Fuzzy DL-Lite⁻

Umberto Straccia
ISTI- CNR, Italy
straccia@isti.cnr.it

December 22, 2005

Abstract

Fuzzy Description Logics (fuzzy DLs) have been proposed as a mean to describe structured knowledge with vague concepts. Unlike classical DLs, where an answer to a query is a set of tuples that satisfy a query, in fuzzy DLs an answer is a set of tuples ranked according to the degree they satisfy the query.

In this paper, we consider fuzzy DL-Lite⁻. We show how to compute efficiently the top- k answers of a complex query (i.e. conjunctive queries) over a huge set of instances.

Category: F.4.1: Mathematical Logic and Formal Languages: Mathematical Logic: [Logic and constraint programming]

Category: I.2.3: Artificial Intelligence: Deduction and Theorem Proving: [Logic programming]

Terms: Theory

Keywords: Fuzzy Description Logics, top- k query answering

1 Introduction

In the last decade a substantial amount of work has been carried out in the context of *Description Logics* (DLs) [1]. DLs are a logical reconstruction of the so-called frame-based knowledge representation languages, with the aim of providing a simple well-established Tarski-style declarative semantics to capture the meaning of the most popular features of structured representation of knowledge. Nowadays, DLs have gained even more popularity due to their application in the context of the *Semantic Web* [7]. DLs play a particular role as they are essentially the theoretical counterpart of the *Web Ontology Language OWL DL*, the state of the art language to specify ontologies [7].

Fuzzy DLs [11, 13] extend classical DLs by allowing to deal with *fuzzy/imprecise concepts*. A major problem fuzzy DLs have to face with is that, unlike classical DLs, where an answer to a query is a set of tuples that satisfy a query, in fuzzy DLs an answer is a set of tuples ranked according to the degree they satisfy the

Interpretation			
\mathcal{I}	=	$(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, $\Delta^{\mathcal{I}}$ is the domain	\wedge = t-norm
$C^{\mathcal{I}}$:	$\Delta^{\mathcal{I}} \rightarrow [0, 1]_{\mathbb{Q}}$	\vee = s-norm
$R^{\mathcal{I}}$:	$\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \rightarrow [0, 1]_{\mathbb{Q}}$	\neg = negation
			\Rightarrow = implication

Concepts				
Syntax		Semantics		Examples
$C, D \rightarrow$	\top	(top concept)	$\top^{\mathcal{I}}(x) = 1$	
	\perp	(bottom concept)	$\perp^{\mathcal{I}}(x) = 0$	
	A	(atomic concept)	$A^{\mathcal{I}}(x) \in [0, 1]$	Human
	$C \sqcap D$	(concept conjunction)	$(C_1 \sqcap C_2)^{\mathcal{I}}(x) = C_1^{\mathcal{I}}(x) \wedge C_2^{\mathcal{I}}(x)$	Human \sqcap Male
	$C \sqcup D$	(concept disjunction)	$(C_1 \sqcup C_2)^{\mathcal{I}}(x) = C_1^{\mathcal{I}}(x) \vee C_2^{\mathcal{I}}(x)$	Nice \sqcup Rich
	$\neg C$	(concept negation)	$(\neg C)^{\mathcal{I}}(x) = \neg C^{\mathcal{I}}(x)$	\neg Male
	$\exists R.C$	(existential quantification)	$(\exists R.C)^{\mathcal{I}}(x) = \sup_{y \in \Delta^{\mathcal{I}}} \{R^{\mathcal{I}}(x, y) \wedge C^{\mathcal{I}}(y)\}$	\exists has_child.Blond
	$\forall R.C$	(universal quantification)	$(\forall R.C)^{\mathcal{I}}(u) = \inf_{y \in \Delta^{\mathcal{I}}} \{R^{\mathcal{I}}(x, y) \Rightarrow C^{\mathcal{I}}(y)\}$	\forall has_child.Human

Fuzzy assertions: $\langle a:C, n \rangle$, $\mathcal{I} \models \langle a:C, n \rangle$ iff $C^{\mathcal{I}}(a^{\mathcal{I}}) \geq n$. An example is $\langle \text{John:Happy_Father}, 0.7 \rangle$.
 Related to roles, we have assertions $\langle (a, b):R, n \rangle$ with semantics $\mathcal{I} \models \langle (a, b):R, n \rangle$ iff $R^{\mathcal{I}}(a^{\mathcal{I}}, b^{\mathcal{I}}) \geq n$. An example is $\langle (\text{John}, \text{Mary}):Loves}, 0.6 \rangle$.

Inclusion axioms: $C \sqsubseteq D$, $\mathcal{I} \models C \sqsubseteq D$ iff $\forall x \in \Delta^{\mathcal{I}}. C^{\mathcal{I}}(x) \leq D^{\mathcal{I}}(x)$.
 Alternative, using residuated implication $\forall x \in \Delta^{\mathcal{I}}. i(C^{\mathcal{I}}(x), D^{\mathcal{I}}(x)) = 1$. An example is,
 $\text{Happy_Father} \sqsubseteq \text{Man} \sqcap \exists \text{has_child.Female}$

Knowledge base: $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, where \mathcal{T} (\mathcal{A}) is a finite set of inclusion axioms (fuzzy assertions).
 $\mathcal{I} \models \mathcal{K}$ iff $\mathcal{I} \models \mathcal{T}$ and $\mathcal{I} \models \mathcal{A}$.

Table 1: fuzzy \mathcal{ALC} .

query. This poses a new challenge in case we have to deal with a huge amount of instances. Indeed, virtually every tuple of the knowledge base satisfies a query to a degree and, thus, has to be ranked. This is not feasible in practice.

In this work, we consider fuzzy DL-Lite⁻. We show how to compute efficiently the top- k answers of a complex query (i.e. conjunctive queries) over a huge set of instances, making the approach appealing for real world scenarios.

We proceed as follows. We first introduce the main notions related to fuzzy DLs and then define DL-Lite⁻. In Section 3 we show how to answer queries. Section 4 concludes and outlines future research.

1.1 Basics of fuzzy DLs

DLs [1] are a family of logics for representing structured knowledge. Each logic is identified by a name made of labels, which identify the operators allowed in that logic. Major DLs are the so-called logic \mathcal{ALC} [10] and is used as a reference language whenever new concepts are introduced in DLs, $\mathcal{SHOIN}(\mathcal{D})$, which is the logic behind the ontology description language OWL DL and $\mathcal{SHIF}(\mathcal{D})$, which is the logic behind OWL LITE, a slightly less expressive language than OWL DL (see [7]).

Fuzzy DLs [11, 13] extend classical DLs by allowing to deal with fuzzy/imprecise concepts (for more on fuzzy DLs, see [11, 12, 13]). The main idea underlying

	Lukasiewicz Logic	Gödel Logic	Product Logic	“Zadeh semantics”
$\neg x$	$1 - x$	if $x = 0$ then 1 else 0	if $x = 0$ then 1 else 0	$1 - x$
$x \wedge y$	$\max(x + y - 1, 0)$	$\min(x, y)$	$x \cdot y$	$\min(x, y)$
$x \vee y$	$\min(x + y, 1)$	$\max(x, y)$	$x + y - x \cdot y$	$\max(x, y)$
$x \Rightarrow y$	if $x \leq y$ then 1 else $1 - x + y$	if $x \leq y$ then 1 else y	if $x \leq y$ then 1 else y/x	$\max(1 - x, y)$

Table 2: Usual connective interpretation combinations.

fuzzy DLs is that an assertion $a:C$, stating that the constant a is an instance of concept C , rather being interpreted as either true or false, will be mapped into a truth value $n \in [0, 1]$. The intended meaning is that n indicates to which extent ‘ a is a C ’. For illustrative purposes, we recall here fuzzy \mathcal{ALC} (see Table 1). *Concepts* and *roles* denote unary and binary predicates respectively. From a syntax point of view, concept forming operators allow to build more complex concepts, starting from so-called *atomic* concepts. *Fuzzy assertions* allow to state that an individual is an instance of concept at least to a given degree in $[0, 1]_{\mathbb{Q}} = [0, 1] \cap \mathbb{Q}$ (a rational number in $[0, 1]$). Similarly, for roles. *Inclusion axioms* allow to state inclusion relationships among concepts. From a semantics point of view, *interpretations* map concept (resp. roles) into functions over the domain $\Delta^{\mathcal{I}}$ (resp. $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$) into $[0, 1]_{\mathbb{Q}}$. The interpretation of complex concepts is given in terms of so-called t-norm (interpreting conjunction), s-norms (interpreting disjunction), negation function (interpreting negation), and implication function (interpreting implication) [6]. Table 2 reports most used combinations of norms.

Given a fuzzy KB \mathcal{K} , and a fuzzy assertion α (resp. an inclusion axiom $C \sqsubseteq D$), we say that \mathcal{K} *entails* α (resp. $C \sqsubseteq D$), denoted $\mathcal{K} \models \alpha$ (resp. $\mathcal{K} \models C \sqsubseteq D$), iff each model of \mathcal{K} satisfies α (resp. $C \sqsubseteq D$). Finally, given \mathcal{K} and a fuzzy assertion α , it is of interest to compute α ’s best lower and upper truth value bounds. The *greatest lower bound* of α w.r.t. \mathcal{K} (denoted $glb(\mathcal{K}, \alpha)$) is $glb(\mathcal{K}, \alpha) = \sup\{n \mid \mathcal{K} \models \langle \alpha, n \rangle\}$ where $\sup \emptyset = 0$. Determining the *glb* is called the *Best Truth Value Bound* (BTVB) problem. Finally, the basic inference problems are:

Consistency: Check if a fuzzy KB is consistent, i.e. has a model?

Subsumption: structure knowledge, compute the taxonomy, i.e. $\mathcal{K} \models C \sqsubseteq D$?

Entailment: Check if a is instance of C to degree $\geq n$, i.e. $\mathcal{K} \models \langle a:C, n \rangle$?

BTVB: Best Truth Value Bound problem, i.e. determine $glb(\mathcal{K}, a:C) = \sup\{n \mid \mathcal{K} \models \langle a:C, n \rangle\}$?

Fuzzy retrieval: Retrieve the top- k ranked constants that instantiate C w.r.t. the best truth value bound, i.e. find the top- k ranked constants of the set $ans(\mathcal{K}, C) = \{\langle a, glb(\mathcal{K}, a:C) \rangle\}$

In [11] decision procedures for the satisfiability, the entailment and the BTVB problem are given for fuzzy \mathcal{ALC} , but with the restrictions on the form of terminological axioms and terminologies. Also, “Zadeh semantics” is used for interpreting the connectives. In [12] a more efficient method is presented for the BTVB problem and covers Lukasiewicz semantics as well. Also the language is more expressive than fuzzy \mathcal{ALC} as it allows to explicitly represent concept membership functions and concept modifiers.

However, a major drawback of current existing reasoning algorithm for fuzzy \mathcal{ALC} is their inefficiency to solve the fuzzy retrieval problem. To date, we have to compute for all constants in the knowledge base its best truth value bound, then order the constants according to this degree bound and then select the top- k ranked constants. This is clearly not a feasible solution if the KB deals with a huge amount of constants.

This will be addressed in the next section, but we have to rely on a simpler language than fuzzy \mathcal{ALC} . But, on the other hand we allow for more complex queries.

2 Fuzzy DL-Lite⁻

To come up with an efficient solution to the fuzzy retrieval problem, we propose to consider a sub-language of fuzzy \mathcal{ALC} . Indeed, we consider fuzzy DL-Lite⁻, a slightly weaker language than the fuzzy variant of DL-Lite [2], which has been proposed as a computationally tractable (in data complexity) DL to query large crisp databases. In fuzzy DL-Lite⁻, concepts are defined as follows:

$$\begin{aligned} B &\rightarrow A \mid \exists R \\ C &\rightarrow B \mid \neg B \\ R &\rightarrow P \mid P^- \end{aligned}$$

where A denotes an atomic concept, P denotes an atomic role. A role R can be either an atomic role P or its *inverse* P^- . B denotes a *basic concept* that can be either an atomic concept, a concept of the form $\exists R$, i.e. the standard DL construct of unqualified existential quantification (equivalent to $\exists R.\top$). C denotes a *general concept*. Note that we use negation on basic concepts only, and we do not allow for disjunction.

A fuzzy DL-Lite⁻ *knowledge base* is pair $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, where \mathcal{T} and \mathcal{A} are finite sets of fuzzy DL-Lite⁻ axioms and assertions. A *fuzzy DL-Lite⁻ axiom* is of the form

$$\begin{aligned} B &\sqsubseteq C \quad (\text{inclusion axiom}) \\ \mathbf{fun}(R) &\quad (\text{functionality axiom}) \end{aligned}$$

(functionality axiom expresses the functionality of a role), while a *fuzzy assertion* is of the form $\langle a:B, n \rangle$ (*fuzzy concept assertion*), or of the form $\langle (a, b):P, n \rangle$ (*fuzzy role assertion*). Additionally, without loss of generality, we assume that if $\langle a:B, n \rangle$ and $\langle a:B, m \rangle$ belong to a fuzzy KB, then $n = m$ (otherwise, we just discard the fuzzy concept assertion with the lower truth degree). Similarly for fuzzy role assertions.

Our language allows for querying the extensional knowledge of a KB in a much more powerful way than usual fuzzy DLs. Specifically, fuzzy DL-Lite⁻ allows for using conjunctive queries of arbitrary complexity. A *conjunctive query* q over a knowledge base \mathcal{K} is an expression of the form

$$q(\mathbf{x}) \leftarrow \exists \mathbf{y}. \text{conj}(\mathbf{x}, \mathbf{y})$$

where \mathbf{x} are the *distinguished variables*, \mathbf{y} are existentially quantified variables called the *non-distinguished variables*, and $\text{conj}(\mathbf{x}, \mathbf{y})$ is a conjunction of atoms

of the form $B(z)$, or $P(z_1, z_2)$, where B and P are respectively a basic concept and a role (but, not inverse role) in \mathcal{K} , and z, z_1, z_2 are constants in \mathcal{K} or variables in \mathbf{x} or \mathbf{y} .

The semantics of fuzzy DL-Lite⁻ is similar to fuzzy \mathcal{ALC} , and is given in terms of interpretations. The major difference is that we consider a *fixed infinite domain* Δ . We assume to have one object for each constant, denoting exactly that object. In other words, we have standard names [8], and we will not distinguish between the alphabet of constants and Δ . So, an *interpretation* is now $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ and consists of a fixed infinite domain Δ with an interpretation function $\cdot^{\mathcal{I}}$ mapping concepts and roles as for fuzzy \mathcal{ALC} , where additionally $(P^-)^{\mathcal{I}}(x, y) = P^{\mathcal{I}}(y, x)$ and for functional roles we say that a fuzzy interpretation satisfies the axiom $\text{fun}(R)$ iff for all $x \in \Delta$ there is a unique $y \in \Delta$ such that $R^{\mathcal{I}}(x, y) > 0$. Finally, we extend fuzzy interpretations to queries $q(\mathbf{x}) \leftarrow \exists \mathbf{y}. \text{conj}(\mathbf{x}, \mathbf{y})$ as follows. Of course, for a basic concept B (role R) appearing in the body of the rule, the interpretation of B (R) is $B^{\mathcal{I}}$ ($R^{\mathcal{I}}$). For $\mathbf{c}, \mathbf{c}' \in \Delta \times \dots \times \Delta$, the conjunction $\text{conj}(\mathbf{c}, \mathbf{c}')$ is interpreted as the t-norm of the conjuncts, while the existential quantifier is interpreted as sup. Therefore, for $\mathbf{c} \in \Delta \times \dots \times \Delta$ the degree of truth of $\exists \mathbf{y}. \text{conj}(\mathbf{x}, \mathbf{y})$ under \mathcal{I} is $\sup_{\mathbf{c}' \in \Delta} \{(\text{conj}(\mathbf{c}, \mathbf{c}'))^{\mathcal{I}}\}$. This is also the truth of $q^{\mathcal{I}}(\mathbf{c})$, i.e., for all $\mathbf{c} \in \Delta \times \dots \times \Delta$,

$$q^{\mathcal{I}}(\mathbf{c}) = \sup_{\mathbf{c}' \in \Delta} \{(\text{conj}(\mathbf{c}, \mathbf{c}'))^{\mathcal{I}}\}.$$

Then we say that: (i) a fuzzy DL-Lite⁻ knowledge base $\mathcal{K} = \langle \mathcal{I}, \mathcal{A} \rangle$ entails $q(\mathbf{c})$ to degree n , written $\mathcal{K} \models \langle q(\mathbf{c}), n \rangle$ iff for any model \mathcal{I} of \mathcal{K} , $q^{\mathcal{I}}(\mathbf{c}) \geq n$; and of course, (ii) the *greatest lower bound* of $q(\mathbf{c})$ w.r.t. \mathcal{K} (denoted $\text{glb}(\mathcal{K}, q(\mathbf{c}))$) is

$$\text{glb}(\mathcal{K}, q(\mathbf{c})) = \sup\{n \mid \mathcal{K} \models \langle q(\mathbf{c}), n \rangle\}.$$

As fuzzy DL-Lite⁻ deals with conjunctive queries, the basic reasoning services that mainly concerns us is the fuzzy knowledge base satisfiability problem and the fuzzy retrieval problem, where this latter is defined as:

Fuzzy retrieval: Retrieve the top- k ranked tuples \mathbf{c} that instantiate the query q w.r.t. the best truth value bound, i.e. find the top- k ranked tuples of the answer set of q , denoted

$$\text{ans}_k(\mathcal{K}, q) = \text{Top}_k(\{\langle \mathbf{c}, \text{glb}(\mathcal{K}, q(\mathbf{c})) \rangle\}).$$

Example 1 Suppose we have information about hotels and conference locations and the distance between them, as shown in the table below. Assume we have a function, which measures the closeness degree between hotels and conference locations, depending on the distance. We may ask to find hotels close to a

conference location, i.e. rank the hotels according to their degree of closeness.

$$\begin{aligned}
\text{Hotel} &\sqsubseteq \exists \text{hasLoc} \\
\text{Conference} &\sqsubseteq \exists \text{hasLoc} \\
\text{Hotel} &\sqsubseteq \neg \text{Conference} \\
\text{Close}(hl, cl) &= \max(0, 1 - \frac{\text{distance}(hl.\text{hasLoc}, cl.\text{hasLoc})}{1000})
\end{aligned}$$

HotelID	hasLoc	ConferenceID	hasLoc
h1	h11	c1	c11
h2	h12	c2	c12
⋮	⋮	⋮	⋮

hasLoc	hasLoc	distance	hasLoc	hasLoc	close
h11	c11	300	h11	c11	0.7
h11	c12	500	h11	c12	0.5
h12	c11	750	h12	c11	0.25
h12	c12	800	h12	c12	0.2
⋮	⋮		⋮	⋮	

We may express our information need using the query (*c1 is our conference location*)

$$\begin{aligned}
q(h) &\leftarrow \text{hasLocation}(h, hl) \wedge \text{hasLocation}(c1, cl) \\
&\quad \wedge \text{close}(hl, cl) .
\end{aligned}$$

Then we want to retrieve the top-*k* answers, according to the degree of closeness. It is not feasible to compute all degrees first and then rank them (there may be a huge amount of hotels and conference locations).

Despite the simplicity of its language and the specific form of inclusion axioms allowed, fuzzy DL-Lite⁻ is able to capture the main notions (though not all, obviously) of both ontologies, and of conceptual modelling formalisms used in databases and software engineering (i.e., ER and UML class diagrams). In particular, fuzzy DL-Lite⁻ axioms allow us to specify *ISA*, e.g., stating that concept A_1 is subsumed by concept A_2 , using $A_1 \sqsubseteq A_2$; *disjointness*, e.g., between concepts A_1 and A_2 , using $A_1 \sqsubseteq \neg A_2$; *role-typing*, e.g., stating that the first (resp., second) component of the relation R is an instance of A_1 (resp., A_2), using $\exists P \sqsubseteq A_1$ (resp., $\exists P^- \sqsubseteq A_2$); *participation constraints*, e.g., stating that all instances of concept A participate to the relation P as the first (resp., second) component, using $A \sqsubseteq \exists P$ (resp., $A \sqsubseteq \exists P^-$); *non-participation constraints*, using $A \sqsubseteq \neg \exists P$ and $A \sqsubseteq \neg \exists P^-$; *functionality restrictions* on relations, using $\text{fun}(R)$. Additionally, observe that fuzzy DL-Lite⁻ does allow for cyclic axioms. A calculus for inclusion axioms is yet unknown for more expressive fuzzy DLs such as fuzzy \mathcal{ALC} . Notice that fuzzy DL-Lite⁻ is a strict subset of fuzzy OWL Lite and, thus of OWL DL [7], which presents some constructs (e.g., some kinds of role restrictions) that are non expressible in fuzzy DL-Lite⁻ (and that make reasoning in fuzzy OWL Lite non-tractable in general).

3 Query Answering

We discuss reasoning in fuzzy DL-Lite⁻. We concentrate on the basic reasoning task in the context of using ontologies to access large data repositories. Namely

that of answering (conjunctive) queries over a fuzzy DL-Lite⁻ knowledge base \mathcal{K} . To this end, we first have to check if \mathcal{K} is satisfiable, as querying an inconsistent KB does not make sense in our case. We provide a simple method for checking satisfiability of \mathcal{K} . Second, we present an efficient top- k query answering procedure. *For the sake of this paper, we fix the semantics of fuzzy DL-Lite⁻ compliant to Lukasiewicz logic.*

The main feature of fuzzy DL-Lite⁻ is that we can almost follow the reasoning part of its crisp counterpart DL-Lite⁻. So, we start by preparing our knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ for effective management. That means, we first normalize it into a suitable form and then store the data in \mathcal{A} into a relational database.

The *normalization* of $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is obtained by transforming \mathcal{K} as follows. \mathcal{A} is expanded by adding to \mathcal{A} the assertions $\langle a:\exists P, n \rangle$ and $\langle b:\exists P^-, n \rangle$ for each $\langle (a, b):P, n \rangle \in \mathcal{A}$. Concerning \mathcal{T} , it contains axioms of the form (i) $B_1 \sqsubseteq B_2$, where B_1 and B_2 are basic concepts (i.e., each of them is either an atomic or an existential concept), which we call *positive inclusions* (PIs); (ii) $B_1 \sqsubseteq \neg B_2$, where B_1 and B_2 are basic concepts, which we call *negative inclusions* (NIs); (iii) functionality axioms on roles of the form $\text{fun}(P)$ or $\text{fun}(P^-)$. Then, \mathcal{T} is expanded by computing all (nontrivial) NIs between basic concepts implied by \mathcal{T} . More precisely, \mathcal{T} is closed with respect to the following inference rule: if $B_1 \sqsubseteq B_2$ occurs in \mathcal{T} and either $B_2 \sqsubseteq \neg B_3$ or $B_3 \sqsubseteq \neg B_2$ occurs in \mathcal{T} (where B_1, B_2, B_3 are arbitrary basic concepts), then add $B_1 \sqsubseteq \neg B_3$ to \mathcal{T} . It can be shown that, after the above closure of \mathcal{T} , for every pair of basic concepts B_1, B_2 , we have that $\mathcal{T} \models B_1 \sqsubseteq \neg B_2$ iff either $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$ or $B_2 \sqsubseteq \neg B_1 \in \mathcal{T}$. This is not surprising as in Lukasiewicz logic the law of contraposition holds: $x \Rightarrow y \equiv \neg x \Rightarrow \neg y$ [6]. Finally, it is easy to show that the normalization process transforms \mathcal{K} into a model preserving form. In the following, without loss of generality we assume that every concept name or role name occurring in \mathcal{A} also occurs in \mathcal{T} . It is worth mentioning that with respect to the crisp variant DL-Lite, we do not consider axioms of the form $B \sqsubseteq C_1 \sqcap C_2$, as we cannot normalize them into $B \sqsubseteq C_1$ and $B \sqsubseteq C_2$, i.e. they are not equivalent. For instance, $B \sqsubseteq A \sqcap \neg A$ means that B is equivalent to $B \sqsubseteq \perp$. This is not captured by $B \sqsubseteq A$ and $B \sqsubseteq \neg A$.

Once \mathcal{K} is normalized, we store it under the control of a Data Base Management System (DBMS), in order to effectively manage and retrieve constants. To this aim, we construct a relational database which faithfully represents a normalized ABox \mathcal{A} . More precisely, (i) for each basic concept B occurring in \mathcal{A} , we define a relational table tab_B of arity 2, such that $\langle a, n \rangle \in tab_B$ iff $\langle a:B, n \rangle \in \mathcal{A}$; and (ii) for each role P occurring in \mathcal{A} , we define a relational table tab_P of arity 3, such that $\langle a, b, n \rangle \in tab_P$ iff $\langle (a, b):P, n \rangle \in \mathcal{A}$. We denote with $\text{DB}(\mathcal{A})$ the relational database thus constructed. We assume that the tuples are stored in $\text{DB}(\mathcal{A})$ in decreasing order according to the truth degree weight.

KB satisfiability. To check the satisfiability of a normalized fuzzy KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, we verify the following conditions: (i) there exists a NI $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$

and a constant a such that the fuzzy assertions $\langle a:B_1, n \rangle$ and $\langle a:B_2, m \rangle$ belong to \mathcal{A} with $n + m > 1$; (ii) there exists an axiom $\text{fun}(P)$ (respectively, $\text{fun}(P^-)$) in \mathcal{T} and three constants a, b, c such that both $\langle (a, b):P, n \rangle$ and $\langle (a, c):P, m \rangle$ (resp., $\langle (b, a):P, n \rangle$ and $\langle (c, a):P, m \rangle$) belong to \mathcal{A} with $n, m > 0$. Informally, the first condition corresponds to checking whether \mathcal{A} explicitly contradicts some NI in \mathcal{T} , and the second condition corresponds to check whether \mathcal{A} violates some functionality axiom in \mathcal{T} . If one of the conditions below holds, then \mathcal{K} is not satisfiable. Otherwise, \mathcal{K} is satisfiable.

Interestingly, the algorithm verifies such conditions by posing to $\text{DB}(\mathcal{A})$ suitable conjunctive queries expressed in SQL. For instance, condition (i) holds for a given NI $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$ iff the query $q(x) \leftarrow \text{tab}_{B_1}(x) \wedge \text{tab}_{B_2}(x)$ has a non-zero degree answer in $\text{DB}(\mathcal{A})$ ($\langle \mathbf{c}, n \rangle \in \text{ans}(\mathcal{K}, q)$ and $n > 0$), while condition (ii) holds for $\text{fun}(R)$ iff the query $q(x) \leftarrow \text{tab}_R(x, y) \wedge \text{tab}_R(x, z) \wedge y \neq z$ has a non-zero degree answer in $\text{DB}(\mathcal{A})$, where \neq is the "not equal" predicate. Notice that the algorithm does not consider the PIs occurring in \mathcal{T} during its execution. Indeed, it can be shown as in [2] that PIs do not affect the consistency of a fuzzy DL-Lite⁻ KB, if \mathcal{T} is normalized.

Query answering. We now describe our query answering procedure. It closely follows [2]. The process is divided into two steps: (i) First, by considering \mathcal{T} only, the user query q is *reformulated* into a set of conjunctive queries $r(q, \mathcal{T})$. (ii) Then, the reformulated queries in $r(q, \mathcal{T})$ are evaluated over \mathcal{A} only (considered as a database), producing the requested answer $\text{ans}_k(\mathcal{K}, q)$.

In the following, we illustrate our approach from a technical point of view.

We start with the query reformulation step. We say that an argument of an atom in a query is *bound* if it corresponds to either a distinguished variable or a shared variable, i.e., a variable occurring at least twice in the query body, or a constant, while we say that it is *unbound* if it corresponds to a non-distinguished non-shared variable (as usual, we use the symbol "-" to represent non-distinguished non-shared variables). Notice that, an atom of the form $(\exists P)(x)$ (resp. $(\exists P^-)(x)$) has the same meaning as $P(x, -)$ (resp. $P(-, x)$). For ease of exposition, in the following we will use the latter form only. A PI axiom τ is *applicable* to an atom $B(x)$, if τ has B in its right-hand side, and I is applicable to an atom $P(x_1, x_2)$, if either (i) $x_2 = -$ and the right-hand side of τ is $\exists P$, or (ii) $x_1 = -$ and the right-hand side of τ is $\exists P^-$. Roughly speaking, an inclusion τ is applicable to an atom g if all bound arguments of g are propagated by τ . Obviously, since all PIs in \mathcal{T} are unary, they are never applicable to atoms with two bound arguments. We indicate with $gr(g; \tau)$ the atom obtained from the atom g by applying the inclusion axiom τ , i.e., if $g = B_1(x)$ (resp., $g = P_1(x, -)$ or $g = P_1(-, x)$) and $\tau = B_2 \sqsubseteq B_1$ (resp., $\tau = B_2 \sqsubseteq \exists P_1$) or $\tau = B_2 \sqsubseteq \exists P_1^-$, we have:

- $gr(g, \tau) = P_2(x, -)$, if $B_2 = \exists P_2$;
- $gr(g, \tau) = P_2(-, x)$, if $B_2 = \exists P_2^-$;
- $gr(g, \tau) = A(x)$, if $B_2 = A$, where A is a basic concept.

We are now ready to recall the query reformulation algorithm `QueryRef` [2]. Given a conjunctive query q and a set of axioms \mathcal{T} , the algorithm reformulates q in terms of a set of queries $r(q, \mathcal{T})$, which then can be evaluated over $\text{DB}(\mathcal{A})$.

Algorithm 1 `QueryRef`(q, \mathcal{T})

Input: Conjunctive query q , fuzzy DL-Lite⁻ axioms \mathcal{T} .
Output: Set of reformulated conjunctive queries $r(q, \mathcal{T})$.

```

1:  $r(q, \mathcal{T}) := \{q\}$ 
2: repeat
3:    $S = r(q, \mathcal{T})$ 
4:   for all  $q \in S$  do
5:     for all  $g \in q$  do
6:       if  $\tau \in \mathcal{T}$  is applicable to  $g$  then
7:          $r(q, \mathcal{T}) := r(q, \mathcal{T}) \cup \{q[g/gr(g, \tau)]\}$ 
8:       for all  $g_1, g_2 \in q$  do
9:         if  $g_1$  and  $g_2$  unify then
10:           $r(q, \mathcal{T}) := r(q, \mathcal{T}) \cup \{\kappa(\text{reduce}(q, g_1, g_2))\}$ 
11: until  $S = r(q, \mathcal{T})$ 
12: return  $r(q, \mathcal{T})$ 

```

In the algorithm, $q[g/g']$ denotes the query obtained from q by replacing the atom g with a new atom g' . Informally, the algorithm first reformulates the atoms of each query $q \in S$, and produces a new query for each atom reformulation (step 5.). Roughly speaking, PIs are used as rewriting rules, applied from right to left, that allow to compile away in the reformulation the knowledge of \mathcal{T} that is relevant for answering q .

At step 8., for each pair of atoms g_1, g_2 that unify, the algorithm computes the query $q' = \text{reduce}(q, g_1, g_2)$, by applying to q the most general unifier between g_1 and g_2 . Due to the unification, variables that were bound in q may become unbound in q' . Hence, PIs that were not applicable to atoms of q , may become applicable to atoms of q' (in the next executions of step (5)). Function κ applied to q' replaces with $_$ each unbound variable in q' . This concludes the query reformulation step. In the following, we show with a small example the behaviour of the query reformulation algorithm.

Example 2 Suppose the set of inclusion axioms is $\mathcal{T} = \{\exists P_1^- \sqsubseteq A, A \sqsubseteq \exists P_1, B \sqsubseteq \exists P_2\}$. We also assume that the set of assertions \mathcal{A} is stored in the tables below (P_2 is a role, while B is a basic concept):

P_2		
a	s	1.0
b	t	0.8
c	q	0.6
d	q	0.4

B	
e	0.9
f	0.7
g	0.5
h	0.3

Assume our query is q_0 :

$$q(x) \leftarrow P_2(x, y) \wedge P_1(y, -) .$$

Then at the first execution of step 7., the algorithm inserts query q_1 , $q(x) \leftarrow P_2(x, y) \wedge A(y)$ into $r(q, \mathcal{T})$ using the axiom $A \sqsubseteq \exists P_1$. At the second execution of step 7., the algorithm inserts query q_2 , $q(x) \leftarrow P_2(x, y) \wedge P_2(-, y)$ using the axiom $\exists P_1^- \sqsubseteq A$. Since the two atoms of the second query unify, step 10. inserts the query q_3 , $q(x) \leftarrow P_2(x, -)$. At the third execution of step 7., the

algorithm inserts query q_4 , $q(x) \leftarrow B(x)$ using the axiom $B \sqsubseteq \exists P_2$. At the fourth execution of step 7., the algorithm uses q_0 and axiom $B \sqsubseteq \exists P_2$ to insert query q_5 , $q(x) \leftarrow B(x) \wedge P_1(y, -)$. At the fifth execution of step 7., the algorithm uses q_5 and axiom $A \sqsubseteq \exists P_1$ to insert query q_6 , $q(x) \leftarrow B(x) \wedge A(y)$. At the sixth execution of step 7., the algorithm uses q_6 and axiom $\exists P_1^- \sqsubseteq A$ to insert query q_7 , $q(x) \leftarrow B(x) \wedge P_2(-, y)$. Finally, at the seventh execution of step 7., the algorithm uses q_7 and axiom $B \sqsubseteq \exists P_2$ to insert query q_8 , $q(x) \leftarrow B(x) \wedge B(y)$.

We point out that we need not to evaluate all queries. Indeed, it can be verified that for each query q_i , all constants c and interpretations \mathcal{I} either $q_3^{\mathcal{I}}(c) \geq q_i^{\mathcal{I}}(c)$ or $q_4^{\mathcal{I}}(c) \geq q_i^{\mathcal{I}}(c)$. That is, we can restrict the evaluation of the set of reformulated queries to $\{q_3, q_4\}$ only. As a consequence, the top-2 answers to the original query are the tuples $\langle \mathbf{a}, 1 \rangle$, $\langle \mathbf{e}, 0.9 \rangle$, which are the top-2 ranked tuples of the union of the answer sets of q_3 and q_4 .

The main property of the query reformulation algorithm is as follows. It can be shown that

$$ans_k(\mathcal{K}, q) = \text{Top}_k(\{\langle c, \text{glb}(\mathcal{A}, q_i(c)) \mid q_i \in r(q, \mathcal{T}) \rangle\}) .$$

The above equation dictates that the set of reformulated queries $q_i \in r(q, \mathcal{T})$ can be used to find the top- k answers, by evaluating them over the set of instances \mathcal{A} only, without referring to the ontology \mathcal{T} anymore.

In the following, we show how to find the top- k answers of the union of the answer sets of conjunctive queries $q_i \in r(q, \mathcal{T})$.

We first note that each conjunctive query $q_i \in r(q, \mathcal{T})$ can easily be transformed into an SQL query expressed over $\text{DB}(\mathcal{A})$. The transformation is conceptually very simple. The only non-trivial case concerns binary atoms with unbound terms: for any atom in a query $q_i \in r(q, \mathcal{T})$ of the form $P(-, x)$, we introduce a view predicate that represents the union of $tab_P[2, 3]$ with $tab_{\exists P^-}$, where $tab_P[2, 3]$ indicates projection of tab_P on its second and third column (similarly for $P(x, -)$). All SQL queries obtained from P , together with the views introduced in the transformation can be easily dispatched to an SQL query engine and evaluated over $\text{DB}(\mathcal{A})$.

Then, a possible, naive solution to the fuzzy retrieval problem may be as follows: we compute for all $q_i \in r(q, \mathcal{T})$ the whole answer set $ans(q_i, \mathcal{A}) = \{\langle c, \text{glb}(\mathcal{A}, q_i(c)) \rangle\}$, take the union, $\bigcup_{q_i \in r(q, \mathcal{T})} ans(q_i, \mathcal{A})$, of these answer sets, order it in descending order of truth degree and then we take the top- k tuples.

A major drawback of this solution is the fact that each tuple satisfies a query to a degree and hence for any query $q_i \in r(q, \mathcal{T})$, all tuples are retrieved always. This is in practice *not feasible*, as there may be millions of tuples in the knowledge base. Restricting the answer set of each query to the tuples with non-zero degree is not satisfactory as well, as still there might be too many tuples in the answer set.

A more effective solution consists in relying on existing top- k queries answering algorithms (see, e.g. [3, 4, 9]), which support efficient evaluations of ranking top- k queries in relational database systems. This gives us immediately a much more efficient method to compute $ans_k(\mathcal{K}, q)$: we compute for all $q_i \in r(q, \mathcal{T})$, the top- k answers $ans_k(\mathcal{A}, q_i)$, using e.g. the system RankSQL [9]. If both k and the number, $n_q = |r(q, \mathcal{T})|$, of reformulated queries is small, say $k \cdot n_q \leq 100$,

then we may take the union, $U(q, \mathcal{K}) = \bigcup_{q_i \in r(q, \mathcal{T})} ans_k(\mathcal{A}, q_i)$, of these top- k answer sets, order it in descending order of truth degree and then we take the top- k tuples.

As an alternative, we can avoid to compute the whole union $U(q, \mathcal{K})$, by relying on a variant of the so-called *Threshold Algorithm* (TA) [5]: let us assume that the tuples in the top- k answer set $ans_k(\mathcal{A}, q_i)$ are sorted in decreasing order with respect to the truth degree. Then we process each top- k answer set $ans_k(\mathcal{A}, q_i)$ ($q_i \in r(q, \mathcal{T})$) in parallel and top-down. (i) For each tuple \mathbf{c} seen, if its truth degree is one of the k highest we have seen, then remember object \mathbf{c} and its truth degree $t(\mathbf{c})$ (ties are broken arbitrarily, so that only k tuples and their truth degrees need to be remembered at any time). (ii) For each answer set $ans_k(\mathcal{A}, q_i)$, let t_i be the truth degree of the last tuple seen in this set. Define the threshold value θ to be $\max(t_1, \dots, t_{n_q})$. As soon as at least k tuples have been seen whose grade is at least equal to θ , then halt (indeed, any successive retrieved tuple will have truth-degree $\leq \theta$). (iii) Let Y be a set containing the k tuples that have been seen with the highest truth degrees. The output is then the graded set $\{\langle \mathbf{c}, t(\mathbf{c}) \rangle \mid \mathbf{c} \in Y\}$. This set is $ans_k(\mathcal{K}, q)$.

The above method, based on existing technology for answering top- k queries over relation databases, improves significantly the naive solution to the fuzzy retrieval problem. Though, it leaves room for improvement. Indeed, we still require to find for each query $q_i \in r(q, \mathcal{T})$ the top- k answer sets $ans_k(\mathcal{A}, q_i)$, merge and sort them and retrieve the final top- k ones. An improvement can be obtained by extending current top- k query answering procedures over relational databases, by generalizing them to the retrieval of top- k tuples of the *union* of queries and not just *one* query. We will deserve to this point more attention in future works.

4 Conclusions

Fuzzy DLs have been proposed as a mean to describe structured knowledge with vague concepts and find their natural application in the context of the Semantic Web. A major distinction of fuzzy DLs is that, an answer to a query is a set of tuples *ranked* according to the degree they satisfy the query. As a consequence, whenever we deal with a huge amount of tuples, the ranking of the answer set becomes the major problem that has to be addressed to make fuzzy DLs viable for real-world application.

In this paper we have considered fuzzy DL-Lite⁻ and have shown how to answer complex queries (in particular, conjunctive queries) efficiently over a huge set of instances. The main ingredients of our solution is a simple and effective query reformulation procedure and the use of existing top- k query answering technology over relational databases. Indeed, a user query is reformulated into a set of conjunctive queries using the inclusion axioms only and, then, the reformulated queries can be submitted to the top- k query answering engine over a relational database where the tuples have been stored. The proposed solution allows to deal with virtually millions of tuples, depending on the effectiveness

of the DBMS.

Related to future research, we envisage two directions: (i) to verify the applicability of our method to richer fuzzy DLs than DL-Lite⁻; and (ii) as, stated at the end of the previous section, to improve the core top-*k* query answering technology towards the management of the union of queries.

References

- [1] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- [2] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. DL-Lite: Tractable description logics for ontologies. In *Proc. of the 20th Nat. Conf. on Artificial Intelligence (AAAI 2005)*, 2005.
- [3] K. Chang and S. Hwang. Minimal probing: Supporting expensive predicates for top-*k* queries. In *SIGMOD Conference*, 2002.
- [4] R. Fagin. Combining fuzzy information: an overview. *SIGMOD Rec.*, 31(2):109–118, 2002.
- [5] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. In *Symposium on Principles of Database Systems*, 2001.
- [6] P. Hájek. *Metamathematics of Fuzzy Logic*. Kluwer, 1998.
- [7] I. Horrocks, P. F. Patel-Schneider, and F. van Harmelen. From SHIQ and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics*, 1(1):7–26, 2003.
- [8] H. J. Levesque and G. Lakemeyer. *The Logic of Knowledge Bases*. MIT Press, 2001.
- [9] C. Li, K. Chang, I. F. Ilyas, and S. Song. RankSQL: query algebra and optimization for relational top-*k* queries. In *SIGMOD Conference*, 2005.
- [10] M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.
- [11] U. Straccia. Reasoning within fuzzy description logics. *Journal of Artificial Intelligence Research*, 14:137–166, 2001.
- [12] U. Straccia. Description logics with fuzzy concrete domains. In *21st Conf. on Uncertainty in Artificial Intelligence (UAI-05)*, pages 559–567, 2005. AUAI Press.
- [13] U. Straccia. A fuzzy description logic for the semantic web. In Elie Sanchez, editor, *Capturing Intelligence: Fuzzy Logic and the Semantic Web*. Elsevier, 2006. To appear.