

Mixing Description Logics in Privacy-Preserving Ontology Publishing

Franz Baader^{*[0000-0002-4049-221X]} and Adrian
Nuradiansyah^{**[0000-0002-9047-7624]}
firstname.lastname@tu-dresden.de

Theoretical Computer Science, TU Dresden, Dresden, Germany

Abstract. In previous work, we have investigated privacy-preserving publishing of Description Logic (DL) ontologies in a setting where the knowledge about individuals to be published is an \mathcal{EL} instance store, and both the privacy policy and the possible background knowledge of an attacker are represented by concepts of the DL \mathcal{EL} . We have introduced the notions of compliance of a concept with a policy and of safety of a concept for a policy, and have shown how, in the context mentioned above, optimal compliant (safe) generalizations of a given \mathcal{EL} concept can be computed. In the present paper, we consider a modified setting where we assume that the background knowledge of the attacker is given by a DL different from the one in which the knowledge to be published and the safety policies are formulated. In particular, we investigate the situations where the attacker’s knowledge is given by an \mathcal{FL}_0 or an $\mathcal{FL}\mathcal{E}$ concept. In both cases, we show how optimal safe generalizations can be computed. Whereas the complexity of this computation is the same (ExpTime) as in our previous results for the case of \mathcal{FL}_0 , it turns out to be actually lower (polynomial) for the more expressive DL $\mathcal{FL}\mathcal{E}$.

1 Introduction

Description Logics (DLs) [3] are a well-investigated family of logic-based knowledge representation languages, which are frequently used to formalize ontologies for application domains such as biology and medicine [13]. To define the important notions of such an application domain as formal concepts, DLs state necessary and sufficient conditions for an individual to belong to a concept. For example, in the DL \mathcal{EL} [2], which is, e.g., employed to define the large medical ontology SNOMED CT,¹ the concept of all male patients that are seen by a female doctor working in the oncology department can be formalized as

$$C = Patient \sqcap Male \sqcap \exists seen_by.(Doctor \sqcap Female \sqcap \exists works_in.Oncology). \quad (1)$$

When publishing information about an individual (e.g., by stating that it belongs to a concept like C in (1)), one needs to ensure that certain privacy constraints

^{*} Partially funded by DFG grant 389792660 as part of TRR 248.

^{**} Funded by DFG within the Research Training Group 1907 RoSI.

¹ see <https://www.snomed.org>

are fulfilled. These constraints are encoded as *privacy policies*, and before publishing the information one needs to check whether the information is *compliant* with these policies [11,12,5]. For example, when publishing information about hospitals, doctors, and patients, the policy may require that one should not be able to find out who are the cancer patients. This policy can, for instance, be expressed by the \mathcal{EL} concept

$$D = Patient \sqcap \exists seen_by.(Doctor \sqcap \exists works_in.Oncology). \quad (2)$$

The concept C in (1) is not compliant with this policy D since C is subsumed by D (written $C \sqsubseteq D$), i.e., if we know that an individual belongs to C , then we can conclude that it also belongs to D . Thus, part of the information contained in C needs to be removed before it can be published. In [5], we have shown how to compute optimal compliant generalizations of C , i.e., concepts C' such that $C \sqsubseteq C'$ and $C' \not\sqsubseteq D$, where optimal means that C' should be as close to C as possible. In our example, the concept $C' = Male \sqcap \exists seen_by.(Doctor \sqcap Female \sqcap \exists works_in.Oncology)$ is such an optimal compliant generalization of C . However, it is *not safe*. In fact, if we publish that John belongs to this concept, and the attacker already knows that John is a patient, then the attacker can still deduce that John belongs to the concept D in (2) since $Patient \sqcap C' \sqsubseteq D$. In general, we say that C' is not safe for the policy D if there is a compliant concept E such that $C' \sqcap E \sqsubseteq D$. Using the algorithm of [5] we can compute the (unique) optimal safe generalization C'' of C in (1) w.r.t. the policy D in (2):

$$C'' = Male \sqcap \exists seen_by.(Doctor \sqcap Female \sqcap \exists works_in.\top) \sqcap \exists seen_by.(Female \sqcap \exists works_in.Oncology).$$

In [5] it was assumed that the information to be published about individuals, the policies, and the attacker's background knowledge are all given by concepts of the DL \mathcal{EL} . In the present paper, we investigate how the notion of safety changes if the attacker's knowledge is assumed to be given by concepts formulated in a different DL. More precisely, we consider the DL \mathcal{FL}_0 , which differs from \mathcal{EL} in that value restrictions $\forall r.C$ are used instead of existential restrictions $\exists r.C$, and the DL $\mathcal{FL}\mathcal{E}$, which combines the constructors of \mathcal{EL} and \mathcal{FL}_0 . If we assume in our example that the attacker's knowledge is given by an \mathcal{FL}_0 concept, the concept C'' from above is no longer safe. In fact, we have $C'' \sqcap Patient \sqcap \forall seen_by.\forall works_in.Oncology \sqsubseteq D$. An attacker may, for example, gain this knowledge by learning that the individual in question is a patient, but is not seen by any of the doctors working in another department. The results shown in the present paper imply that the concept

$$C''' = Male \sqcap Patient \sqcap \exists seen_by.(Doctor \sqcap Female)$$

is an optimal safe generalization of C in (1) w.r.t. the policy D in (2) if the attacker's knowledge is assumed to be given by an \mathcal{FL}_0 concept.

If, instead, we assume that the attacker's knowledge is given by an $\mathcal{FL}\mathcal{E}$ concept, then C''' is no longer safe. In fact, $C''' \sqcap \forall seen_by.\exists works_in.Oncology \sqsubseteq$

D . In this setting, the results shown in this paper imply that $C'''' = Male$ is the optimal safe generalization of C in (1) w.r.t. the policy D in (2). It may be less intuitive to see how an attacker could gain this additional knowledge. But the assumption in [11,12,5] and in the present paper is that the attacker can gain *any* compliant knowledge that can be formulated in the considered formalism (i.e., $\mathcal{FL}\mathcal{E}$ concepts in the last part of our example).

To be more precise, just as in [5], we assume here that all knowledge about individuals (be it the one to be published or the one known by an attacker) is represented using concepts, which corresponds to what is called an instance store in [14]. The restriction compared to general DL-based ontologies is, on the hand, that relations between named individuals cannot be stated. On the other hand, there is no TBox available to state subconcept constraints. One reason for considering this restricted setting is that, when investigating a new inference problem, it is usually quite hard to start with the most general situation. Also note that the work in [11,12] concentrates on a setting where there are relations between individuals, but no instance relationships between individuals and complex concepts can be stated. Our hope is that, by combining our approaches with the ones developed in [11,12], we can make the step from instance stores to general DL ABoxes. A second reason is that, in a medical application that uses SNOMED CT as an ontology, the TBox can be reduced away by expanding concept definitions since SNOMED CT is an acyclic TBox [16]. In addition, patient data are usually annotated with SNOMED concepts, but not with SNOMED roles, which justifies considering an instance store rather than a general ABox. The main difference to our work in [5] is that we assume that the background knowledge of the attacker is formulated in a DL (\mathcal{FL}_0 or $\mathcal{FL}\mathcal{E}$) different from the one (\mathcal{EL}) in which the knowledge to be published and the safety policies are expressed. Using \mathcal{FL}_0 concepts for formulating the attacker’s knowledge makes sense since in SNOMED CT roles have implicit typing constraints [16], which are not explicitly stated using value restrictions, but which may be known to an attacker. Considering $\mathcal{FL}\mathcal{E}$ as well is interesting from an academic perspective since our results illustrate the effect that assuming a more expressive DL for the attacker’s knowledge may have on the computation of safe generalizations.

In the next section, we will introduce the DLs $\mathcal{FL}\mathcal{E}$, \mathcal{EL} , and \mathcal{FL}_0 , and then define the notions of compliance and safety used in this paper. In Section 3, we will investigate the setting where the attacker’s knowledge is assumed to be given by an \mathcal{FL}_0 concept, and in Section 4 we consider the case where $\mathcal{FL}\mathcal{E}$ is used instead.

2 Preliminaries

A wide range of DLs of different expressive power has been investigated in the literature [3]. Here, we introduce the DL $\mathcal{FL}\mathcal{E}$ and its sublogics \mathcal{EL} and \mathcal{FL}_0 .

Starting with disjoint sets N_C and N_R of *concept* and *role names*, respectively, the set of $\mathcal{FL}\mathcal{E}$ *concepts* over these names are constructed from concept names using the constructors top concept (\top), conjunction ($C \sqcap D$), existential

restriction ($\exists r.C$), and value restriction ($\forall r.D$). The sublogic \mathcal{EL} forbids the use of value restrictions, and \mathcal{FL}_0 forbids the use of existential restrictions. The *size* of an $\mathcal{FL}\mathcal{E}$ concept C is the number of occurrences of \top as well as concept and role names in C , and the *role depth* $rd(C)$ is the maximal nesting of existential restrictions and value restrictions.

The semantics of $\mathcal{FL}\mathcal{E}$ (and thus of \mathcal{EL} and \mathcal{FL}_0) is defined through *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set, called the *domain*, and $\cdot^{\mathcal{I}}$ is the *interpretation function*, which maps every $A \in N_C$ to a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and every $r \in N_R$ to a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. This function is extended to arbitrary $\mathcal{FL}\mathcal{E}$ concepts by setting

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

The $\mathcal{FL}\mathcal{E}$ concept C is *subsumed by* the $\mathcal{FL}\mathcal{E}$ concept D (written $C \sqsubseteq D$) if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ holds for all interpretations \mathcal{I} . Strict subsumption (written $C \sqsubset D$) holds if $C \sqsubseteq D$ and $D \not\sqsubseteq C$, and we say that C is *equivalent to* D (written $C \equiv D$) if $C \sqsubseteq D$ and $D \sqsubseteq C$.

Whereas subsumption between concepts (without a TBox) is polynomial in \mathcal{EL} [6] and \mathcal{FL}_0 [15], it is NP-complete for $\mathcal{FL}\mathcal{E}$ [7]. Fortunately, in the context of this paper, we are not interested in subsumption between two $\mathcal{FL}\mathcal{E}$ concepts, but only in deciding whether an $\mathcal{FL}\mathcal{E}$ concept is subsumed by an \mathcal{EL} concept. Before we can show that this problem is decidable in polynomial time, we need to introduce some notation.

We call an $\mathcal{FL}\mathcal{E}$ concept an *atom* if it is a concept name, an existential restriction, or a value restriction, and we denote the set of atoms occurring in the top-level conjunction of an $\mathcal{FL}\mathcal{E}$ concept C with $\mathbf{con}(C)$. For example, if $C = A \sqcap \exists r.(B \sqcap \exists s.A) \sqcap \forall r.\forall s.B$, then $\mathbf{con}(C) = \{A, \exists r.(B \sqcap \exists s.A), \forall r.\forall s.B\}$. Given an $\mathcal{FL}\mathcal{E}$ concept C and a role name $r \in N_R$, we define $\mathbf{filler}_r^{\forall}(C) := \{E \mid \forall r.E \in \mathbf{con}(C)\}$.

The following proposition is an easy consequence of the characterization of subsumption between $\mathcal{FL}\mathcal{E}$ concepts given in [6] (Theorem 24).

Proposition 1. *Let C be an $\mathcal{FL}\mathcal{E}$ concept and D be an \mathcal{EL} concept. Then, $C \sqsubseteq D$ holds iff*

- a.) $A \in \mathbf{con}(D)$ implies $A \in \mathbf{con}(C)$ for every concept name A , and
- b.) for every $\exists r.D' \in \mathbf{con}(D)$, there is $\exists r.C' \in \mathbf{con}(C)$ such that

$$C' \sqcap \prod \mathbf{filler}_r^{\forall}(C) \sqsubseteq D',$$

where $\prod \mathbf{filler}_r^{\forall}(C)$ denotes the conjunction of all elements of $\mathbf{filler}_r^{\forall}(C)$.

In particular, subsumption between C and D can be decided in time polynomial in the size of C and D .

By induction on the role depth, it is easy to show that the recursive characterization of subsumption between an $\mathcal{FL}\mathcal{E}$ and an \mathcal{EL} concept given in the proposition indeed yields a polynomial-time decision procedure.

We close this section by defining the notions of compliance and safety used in this paper. Following [5], we assume that policies are given not just by one \mathcal{EL} concept, but by a finite set of \mathcal{EL} concepts. We assume that these concepts are not equivalent to \top since otherwise there would not be compliant concepts.

Definition 1. *An \mathcal{EL} policy is a finite set $\mathcal{P} = \{D_1, \dots, D_p\}$ of \mathcal{EL} concepts that are not equivalent to \top . Let C be an \mathcal{EL} concept, $Q \in \{\exists, \forall, \forall\exists\}$, and $\mathcal{L}_\exists = \mathcal{EL}$, $\mathcal{L}_\forall = \mathcal{FL}_0$, $\mathcal{L}_{\forall\exists} = \mathcal{FL}\mathcal{E}$. We say that*

- the \mathcal{L}_Q concept C' is compliant with \mathcal{P} if $C' \not\sqsubseteq D_i$ for all $i = 1, \dots, p$,
- the \mathcal{EL} concept C' is
 - Q -safe for \mathcal{P} if for all \mathcal{L}_Q concepts C'' that are compliant with \mathcal{P} , $C' \sqcap C''$ is also compliant with \mathcal{P} , i.e., $C' \sqcap C'' \not\sqsubseteq D_i$ for all $i = 1, \dots, p$,
 - a Q -safe generalization of C for \mathcal{P} if $C \sqsubseteq C'$ and C' is Q -safe for \mathcal{P} ,
 - an optimal Q -safe generalization of C for \mathcal{P} if it is a Q -safe generalization of C for \mathcal{P} , and there is no Q -safe generalization C'' of C for \mathcal{P} such that $C'' \sqsubset C'$.

Let C, C' be \mathcal{EL} concepts, \mathcal{P} be an \mathcal{EL} policy, and $Q \in \{\exists, \forall, \forall\exists\}$. The Q -safety problem asks whether C is Q -safe for \mathcal{P} and the Q -optimality problem asks whether C' is an optimal Q -safe generalization of C for \mathcal{P} .

We call an \mathcal{EL} policy \mathcal{P} *redundancy-free* if \mathcal{P} does not contain distinct concepts D, D' such that $D \sqsubseteq D'$. We can without loss of generality restrict our attention to redundancy-free policies since removing redundant concepts (i.e., concepts $D' \in \mathcal{P}$ such that there is $D \in \mathcal{P} \setminus \{D'\}$ with $D \sqsubseteq D'$) does not change the sets of compliant and safe concepts (see Lemma 2 in [5]).

Compliance for \mathcal{EL} concepts and \exists -safety have been investigated in detail in [5]. In the following two sections, we investigate \forall -safety and $\forall\exists$ -safety.

3 Investigating \forall -Safety

In this section, we first characterize \forall -safety and prove that this characterization can be decided in polynomial time. Then, we show how to compute optimal \forall -safe generalizations of a given \mathcal{EL} concept. Finally, we address the \forall -optimality problem.

Characterizing \forall -safety First, note that a value restriction can never imply an existential restriction. Thus, if C'' is an \mathcal{FL}_0 concept and D an \mathcal{EL} concept of role depth > 0 , then $C'' \not\sqsubseteq D$. This shows that an \mathcal{FL}_0 concept C'' is compliant with any \mathcal{EL} policy that does not contain a concept of role depth 0.

Using this observation and Proposition 1, we can characterize \forall -safety for redundancy-free policies as follows.

Proposition 2. *Let C be an \mathcal{EL} concept and \mathcal{P} a redundancy-free \mathcal{EL} policy. Then C is \forall -safe for \mathcal{P} iff the following two conditions hold for all $D \in \mathcal{P}$:*

- 1.) *if $\text{rd}(D) = 0$, then $\text{con}(C) \cap \text{con}(D) = \emptyset$,*
- 2.) *if $\text{rd}(D) > 0$, then there is $\exists r.D' \in \text{con}(D)$ such that*
 - a.) *if $\text{rd}(D') = 0$, then there is no concept of the form $\exists r.C' \in \text{con}(C)$,*
 - b.) *if $\text{rd}(D') > 0$, then for all $\exists r.C' \in \text{con}(C)$, C' is \forall -safe for $\{D'\}$.*

Proof. To show the *if-direction*, assume that C is not \forall -safe for \mathcal{P} . Then, there is an \mathcal{EL} concept $D \in \mathcal{P}$ and an \mathcal{FL}_0 concept C'' that complies with \mathcal{P} such that $C \sqcap C'' \sqsubseteq D$. Since $C \sqcap C''$ is an $\mathcal{FL}\mathcal{E}$ concept, Proposition 1 applies to this subsumption. First, we consider the case where $\text{rd}(D) = 0$. Proposition 1 implies that every concept name $A \in \text{con}(D)$ is contained in $\text{con}(C) \cup \text{con}(C'')$. However, since C'' complies with \mathcal{P} , we have $C'' \not\sqsubseteq D$, and hence there must be an $A \in \text{con}(D)$ that is not contained in $\text{con}(C'')$. Consequently this A must belong to $\text{con}(C)$, and thus property 1.) above is violated.

Now, consider the case where $\text{rd}(D) > 0$, i.e., there is an existential restriction $\exists r.D' \in \text{con}(D)$. By Proposition 1 and since C is an \mathcal{EL} and C'' an \mathcal{FL}_0 concept, $C \sqcap C'' \sqsubseteq D$ implies that there is an existential restriction $\exists r.C' \in \text{con}(C)$ such that $C' \sqcap \prod \text{filler}_r^\forall(C'') \sqsubseteq D'$. If $\text{rd}(D') = 0$, then this clearly violates 2a.). If $\text{rd}(D') > 0$, then 2b.) is violated since $\prod \text{filler}_r^\forall(C'')$ then cannot be subsumed by D' , and thus $C' \sqcap \prod \text{filler}_r^\forall(C'') \sqsubseteq D'$ shows that C' is not \forall -safe for $\{D'\}$.

To show the *only-if-direction*, we assume that one of the conditions 1.) or 2.) is violated, and prove that this implies that C is not \forall -safe for \mathcal{P} .

First, assume that 1.) is violated, i.e., there is $D \in \mathcal{P}$ such that $\text{rd}(D) = 0$ and there is $A \in \text{con}(C) \cap \text{con}(D)$. Then, $C'' := \prod(\text{con}(D) \setminus \{A\})$ is an \mathcal{FL}_0 concept that complies with D , and satisfies $C \sqcap C'' \sqsubseteq D$. To conclude that C is not \forall -safe for \mathcal{P} , it remains to show that C'' also complies with all $\hat{D} \in \mathcal{P} \setminus \{D\}$. However, if we assume that $C'' \sqsubseteq \hat{D}$ for some $\hat{D} \in \mathcal{P} \setminus \{D\}$, then the fact that $D \sqsubseteq C''$ implies $D \sqsubseteq \hat{D}$, which contradicts our assumption that \mathcal{P} is redundancy-free.

Second, assume that 2.) is violated. Then there is $D \in \mathcal{P}$ such that $\text{rd}(D) > 0$ and for all $\exists r.D' \in \text{con}(D)$ we have

- if $\text{rd}(D') = 0$, then there is a concept of the form $\exists r.C' \in \text{con}(C)$, and
- if $\text{rd}(D') > 0$, then there is $\exists r.C' \in \text{con}(C)$ such that the concept C' is not \forall -safe for $\{D'\}$.

We define the concept C'' as follows: $\prod_{(1)} A \sqcap \prod_{(2)} \forall r.D' \sqcap \prod_{(3)} \forall r.F$, where

- (1) $A \in \text{con}(D)$;
- (2) $r \in N_R, \exists r.D' \in \text{con}(D)$, and $\text{rd}(D') = 0$;
- (3) $r \in N_R, \exists r.D' \in \text{con}(D), \text{rd}(D') > 0, \exists r.C' \in \text{con}(C)$, and F is an \mathcal{FL}_0 concept complying with D' , but satisfying $C' \sqcap F \sqsubseteq D'$.

Note that C'' is an \mathcal{FL}_0 concept that is compliant with \mathcal{P} . To see the latter, assume that $\hat{D} \in \mathcal{P}$. If $\text{rd}(\hat{D}) > 0$, then $C'' \not\sqsubseteq \hat{D}$ since an \mathcal{FL}_0 concept cannot

imply an existential restriction. If $\text{rd}(\hat{D}) = 0$, then $C'' \sqsubseteq \hat{D}$ would imply $D \sqsubseteq \hat{D}$, which contradicts our assumption that \mathcal{P} is redundancy-free.

It remains to prove that $C \sqcap C'' \sqsubseteq D$, which we show using Proposition 1. First note that, by the construction of C'' , each concept name $A \in \text{con}(D)$ satisfies $A \in \text{con}(C'')$, and thus $A \in \text{con}(C \sqcap C'')$. Second, consider an existential restriction $\exists r.D' \in \text{con}(D)$. If $\text{rd}(D') = 0$, then there is $\exists r.C' \in \text{con}(C)$, but also $\forall r.D' \in \text{con}(C'')$. Thus we have $C' \sqcap \prod \text{filler}_r^\forall(C \sqcap C'') \sqsubseteq C' \sqcap D' \sqsubseteq D'$, as required by Proposition 1. If $\text{rd}(D') > 0$, then we have $\exists r.C' \in \text{con}(C)$ for an \mathcal{EL} concept C' that is not \forall -safe for $\{D'\}$. In addition, $\forall r.F \in \text{con}(C'')$, where F is an \mathcal{FL}_0 concept such that $C' \sqcap F \sqsubseteq D'$. Consequently, we have $C' \sqcap \prod \text{filler}_r^\forall(C \sqcap C'') \sqsubseteq C' \sqcap F \sqsubseteq D'$. \square

Using induction on the role depth, it is easy to show that the characterization of \forall -safety stated in the above proposition can be decided in polynomial time.

Theorem 1. *The \forall -safety problem is in P.*

Since 1.) and 2.) in Proposition 2 are formulated for each $D \in \mathcal{P}$ separately, the following lemma is an immediate consequence of this proposition.

Lemma 1. *Let C be an \mathcal{EL} concept and \mathcal{P} an \mathcal{EL} policy. Then C is \forall -safe for \mathcal{P} iff C is \forall -safe for $\{D\}$ for all $D \in \mathcal{P}$.*

In this lemma, we have dispensed with the restriction that \mathcal{P} is redundancy-free. This is admissible since we can first remove redundant elements from \mathcal{P} and then apply the proposition to the redundancy-free policy obtained this way. Note that this fact will become important in the proof of Lemma 3, since there we cannot assume that the policy \mathcal{P}_i defined there is redundancy-free.

Computing optimal \forall -safe generalizations Before we can describe our approach for computing optimal \forall -safe generalizations, we need to introduce some notation. Given an \mathcal{EL} concept D such that $\text{rd}(D) > 0$, the set $\text{con}^\exists(D)$ consists of the elements of $\text{con}(D)$ that are existential restrictions.

Definition 2. *Let D_1, \dots, D_p be \mathcal{EL} concepts of role depth greater than zero. We say that $H \subseteq \text{con}^\exists(D_1) \cup \dots \cup \text{con}^\exists(D_p)$ is a hitting set of $\text{con}^\exists(D_1), \dots, \text{con}^\exists(D_p)$ if $H \cap \text{con}(D_i) \neq \emptyset$ for every $i = 1, \dots, p$. This hitting set is minimal if, for all $H' \subset H$, H' is not a hitting set.*

We will show that the set defined below contains all optimal \forall -safe generalizations of C for \mathcal{P} .

Definition 3. *Let C be an \mathcal{EL} concept and $\mathcal{P} = \{D_1, \dots, D_n\}$ a redundancy-free \mathcal{EL} policy. The set $\text{SSG}(C, \mathcal{P})$ of all specific \forall -safe generalizations of C for \mathcal{P} consists of the concepts C' that are obtained from C as follows:*

- If C is \forall -safe for \mathcal{P} , then $\text{SSG}(C, \mathcal{P}) = \{C\}$.
- Otherwise, perform the following steps:

- For all concept names $A \in \mathbf{con}(C)$ such that $A \in \mathbf{con}(D)$, where $D \in \mathcal{P}$ and $\mathbf{rd}(D) = 0$, remove A from $\mathbf{con}(C)$.
- If D_{j_1}, \dots, D_{j_p} are all concepts in \mathcal{P} such that $\mathbf{rd}(D_{j_\nu}) > 0$, then construct a minimal hitting set H of $\mathbf{con}^\exists(D_{j_1}), \dots, \mathbf{con}^\exists(D_{j_p})$ and do the following:
 - * For all $\exists r.E \in \mathbf{con}(C)$ such that there is a concept of the form $\exists r.D'$ in H with $\mathbf{rd}(D') = 0$, remove $\exists r.E$ from $\mathbf{con}(C)$.
 - * For each $\exists r_i.C_i \in \mathbf{con}(C)$ that was not removed in the previous step, consider the set

$$\mathcal{P}_i := \{D' \mid \exists r_i.D' \in H \text{ and } \mathbf{rd}(D') > 0\}.$$

If $\mathcal{P}_i \neq \emptyset$, then replace $\exists r_i.C_i$ in $\mathbf{con}(C)$ with $\prod \exists r_i.F$, where $F \in \text{SSG}(C_i, \mathcal{P}_i)$. If \mathcal{P}_i is empty, then leave $\exists r_i.C_i$ as it is.

First, we show that the elements of $\text{SSG}(C, \mathcal{P})$ are indeed \forall -safe generalization.

Lemma 2. *If $C' \in \text{SSG}(C, \mathcal{P})$, then C' is a \forall -safe generalization of C for \mathcal{P} .*

Proof. First, we show that $C \sqsubseteq C'$. This is an easy consequence of the fact that, when constructing C' from C , atoms from the top-level conjunction of C are kept unchanged, removed, or generalized. The only non-trivial case is when $\exists r_i.C_i$ in $\mathbf{con}(C)$ is replaced with $\prod \exists r_i.F$, where F ranges over the elements of $\text{SSG}(C_i, \mathcal{P}_i)$. By induction on the role depth, we know that $C_i \sqsubseteq F$ for all $F \in \text{SSG}(C_i, \mathcal{P}_i)$, and thus $\exists r.C_i \sqsubseteq \exists r.F$.

To prove that C' is safe for \mathcal{P} , we use the characterization given in Proposition 2. Thus, let $D \in \mathcal{P}$. If $\mathbf{rd}(D) = 0$, then $\mathbf{con}(D)$ is a set of concept names, and each of them has been removed in the construction of C' . Thus, $\mathbf{con}(C') \cap \mathbf{con}(D) = \emptyset$, as required by 1.) in Proposition 2.

If $\mathbf{rd}(D) > 0$, then the minimal hitting set H used in the construction of C' contains an existential restriction $\exists r.\hat{D} \in \mathbf{con}(D)$. If $\mathbf{rd}(\hat{D}) = 0$, then all existential restrictions for the role r are removed from the top-level conjunction of C , and thus 2a.) of Proposition 2 is satisfied. Finally, consider the case where $\mathbf{rd}(\hat{D}) > 0$. If $\exists r.E \in \mathbf{con}(C')$, then there is $\exists r_i.C_i \in \mathbf{con}(C)$ such that

$$\mathcal{P}_i = \{D' \mid \exists r_i.D' \in H \text{ and } \mathbf{rd}(D') > 0\} \neq \emptyset,$$

and $r = r_i$ and $E \in \text{SSG}(C_i, \mathcal{P}_i)$. Note that $\hat{D} \in \mathcal{P}_i$, and thus $\mathcal{P}_i = \emptyset$ is not possible for an existential restriction $\exists r_i.C_i \in \mathbf{con}(C)$ with $r_i = r$. Induction (over the role depth) yields that E is \forall -safe for \mathcal{P}_i , and thus for its subset $\{\hat{D}\}$. Hence, 2b.) of Proposition 2 is satisfied. \square

However, $\text{SSG}(C, \mathcal{P})$ may also contain \forall -safe generalizations C' of C for \mathcal{P} that are not optimal, as demonstrated by the following example.

Example 1. Let $C = \exists r_1.(A \sqcap B) \sqcap \exists r_2.B \sqcap \exists r_3.A$ and $\mathcal{P} = \{D_1, D_2\}$, where

$$D_1 = \exists r_1.A \sqcap \exists r_2.\top \quad \text{and} \quad D_2 = \exists r_1.B \sqcap \exists r_3.\top.$$

We have $C \sqsubseteq D_1$ and $C \sqsubseteq D_2$, and thus C is not even compliant, let alone \forall -safe, for \mathcal{P} . Applying the construction of Definition 3 to C and \mathcal{P} , we first construct the minimal hitting set $H_1 = \{\exists r_1.A, \exists r_1.B\}$ of $\text{con}^\exists(D_1)$ and $\text{con}^\exists(D_2)$. Since $\text{rd}(A) = 0 = \text{rd}(B)$ we remove the atom $\exists r_1.(A \sqcap B)$ from $\text{con}(C)$, which yields the concept $C'_1 = \exists r_2.B \sqcap \exists r_3.A \in \text{SSG}(C, \mathcal{P})$.

If we take the minimal hitting set $H_2 = \{\exists r_1.A, \exists r_3.\top\}$ instead, then we need to remove the atoms $\exists r_1.(A \sqcap B)$ and $\exists r_3.A$ from $\text{con}(C)$, which yields $C'_2 = \exists r_2.B \in \text{SSG}(C, \mathcal{P})$. Since $C'_1 \sqsubset C'_2$, the concept C'_2 cannot be optimal.

The next lemma states that every \forall -safe generalization of C subsumes some element of $\text{SSG}(C, \mathcal{P})$.

Lemma 3. *For all \forall -safe generalization C'' of C for \mathcal{P} , there is $C' \in \text{SSG}(C, \mathcal{P})$ such that $C' \sqsubseteq C''$.*

Proof. If C is \forall -safe for \mathcal{P} , then obviously $C \in \text{SSG}(C, \mathcal{P})$ and we have $C \sqsubseteq C''$. Thus, let us assume that C is not \forall -safe for \mathcal{P} . Since C'' is a \forall -safe generalization of C for \mathcal{P} , we have $C \sqsubseteq C''$ and C'' satisfies the properties 1.) and 2.) in Proposition 2. Due to 1.), $\text{con}(C'')$ contains no concept name A such that $A \in \text{con}(D)$ for some $D \in \mathcal{P}$ with $\text{rd}(D) = 0$. In addition, for all $D_{j_\nu} \in \mathcal{P}$ such that $\text{rd}(D_{j_\nu}) > 0$, there is $\exists r.G_{j_\nu} \in \text{con}(D_{j_\nu})$ such that 2a.) or 2b.) of Proposition 2 holds. The set $H' := \{G_{j_1}, \dots, G_{j_p}\}$ is a hitting set of the sets $\text{con}^\exists(D_{j_1}), \dots, \text{con}^\exists(D_{j_p})$ considered in Definition 3. Thus, there is a minimal hitting set H of $\text{con}^\exists(D_{i_1}), \dots, \text{con}^\exists(D_{i_q})$ such that $H \subseteq H'$. Let C' be the element of $\text{SSG}(C, \mathcal{P})$ that is constructed by using H . We show that $C' \sqsubseteq C''$ using Proposition 1.

First, consider a concept name $A \in \text{con}(C'')$. Since $C \sqsubseteq C''$, we know that $A \in \text{con}(C)$. In addition, as mentioned above, $\text{con}(C'')$ contains no concept name A such that $A \in \text{con}(D)$ for some $D \in \mathcal{P}$ with $\text{rd}(D) = 0$. Consequently, when constructing C' from C , the concept name A is not removed, which yields $A \in \text{con}(C')$.

Second, consider an existential restriction $\exists r.E \in \text{con}(C'')$. Since $C \sqsubseteq C''$, there is $\exists r.C_i \in \text{con}(C)$ such that $C_i \sqsubseteq E$. If $\exists r.C_i$ is not removed or generalized when constructing C' , then $\exists r.C_i \in \text{con}(C')$, and we are done. If $\exists r.C_i$ is removed from $\text{con}(C)$ to construct C' , then there is $\exists r.D' \in H \subseteq H'$ such that $\text{rd}(D') = 0$. By the definition of H' , we thus know that $\exists r.D'$ must satisfy 2a.) of Proposition 2. But then $\exists r.E \in \text{con}(C'')$ would not be possible.

Finally, if $\exists r.C_i$ is generalized in the construction of C' from C by replacing it with $\prod_{F \in \text{SSG}(C_i, \mathcal{P}_i)} \exists r.F$, then we know that \mathcal{P}_i is non-empty. Now, consider an element D' of \mathcal{P}_i . Then, $\exists r.D' \in H \subseteq H'$ and $\text{rd}(D') > 0$ imply that $\exists r.D'$ satisfies 2b.) of Proposition 2. Since $\exists r.E \in \text{con}(C'')$, we thus know that E is \forall -safe for $\{D'\}$. Since this is true for all elements D' of \mathcal{P}_i , Lemma 1 yields that E is \forall -safe for \mathcal{P}_i . Together with $C_i \sqsubseteq E$, this shows that E is a \forall -safe generalization of C_i for \mathcal{P}_i , and thus induction yields that there is $F \in \text{SSG}(C_i, \mathcal{P}_i)$ such that $F \sqsubseteq E$. Since $\exists r.F \in \text{con}(C')$, this concludes our proof that $C' \sqsubseteq C''$. \square

The following proposition states that all optimal \forall -safe generalizations of C for \mathcal{P} are contained in $\text{SSG}(C, \mathcal{P})$.

Proposition 3. *Let C be an \mathcal{EL} concept and $\mathcal{P} = \{D_1, \dots, D_p\}$ a redundancy-free \mathcal{EL} policy. If C'' is an optimal \forall -safe generalization of C for \mathcal{P} , then $C'' \in \text{SSG}(C, \mathcal{P})$ (up to equivalence).*

Proof. Given an optimal \forall -safe generalization C'' of C for \mathcal{P} , Lemma 3 yields an element $C' \in \text{SSG}(C, \mathcal{P})$ such that $C' \sqsubseteq C''$. By Lemma 2, $C \sqsubseteq C'$ and C' is \forall -safe for \mathcal{P} . Thus, optimality of C'' implies that $C' \equiv C''$. \square

The following theorem is an easy consequence of this proposition and the definition of $\text{SSG}(C, \mathcal{P})$.

Theorem 2. *Let C be an \mathcal{EL} concept and $\mathcal{P} = \{D_1, \dots, D_p\}$ a redundancy-free \mathcal{EL} policy. The cardinality of the set of all optimal \forall -safe generalization of C for \mathcal{P} is at most exponential, and each of its elements has at most exponential size. Additionally, the set of all optimal \forall -safe generalizations of C for \mathcal{P} can be computed in exponential time.*

Proof. It is sufficient to show that the set $\text{SSG}(C, \mathcal{P})$ satisfies the properties stated above. The cardinality of $\text{SSG}(C, \mathcal{P})$ is at most exponential since in Definition 3 at most exponentially many minimal hitting sets are considered, and each such set yields exactly one element of $\text{SSG}(C, \mathcal{P})$. Moreover, the size of each element C' in $\text{SSG}(C, \mathcal{P})$ may become exponential (but not more) since, during constructing C' , we may need to compute a conjunction of at most exponentially many existential restrictions, where each of them has at most exponential size (by induction). To compute the set of all optimal \forall -safe generalizations of C for \mathcal{P} , we need to remove all concepts in $\text{SSG}(C, \mathcal{P})$ that are not minimal w.r.t. subsumption. This requires exponentially many (polynomial) subsumption tests on exponentially large concepts. \square

The following example shows that an algorithm for computing all optimal \forall -safe generalizations cannot be better than exponential in the worst case.

Example 2. Let $C = \exists r_1.(\exists s_1.\top \sqcap \exists s_2.\top) \sqcap \dots \sqcap \exists r_n.(\exists s_1.\top \sqcap \exists s_2.\top)$ and $\mathcal{P} = \{\exists r_i.\exists s_1.\top \sqcap \exists r_i.\exists s_2.\top \mid 1 \leq i \leq n\}$. Using Definition 3, we obtain that $\text{SSG}(C, \mathcal{P})$ consists of the exponentially many concepts $\exists r_1.\exists s_{j_1}.\top \sqcap \dots \sqcap \exists r_n.\exists s_{j_n}.\top$, where $j_i \in \{1, 2\}$ for all $i = 1, \dots, n$. Since these concepts are incomparable w.r.t. subsumption, they are exactly the optimal \forall -safe generalization of C for \mathcal{P} .

The \forall -optimality problem Given \mathcal{EL} concepts C, C' such that $C \sqsubseteq C'$ and an \mathcal{EL} policy \mathcal{P} , the \forall -optimality problem asks whether C' is an optimal \forall -safe generalization of C for \mathcal{P} . Since \forall -safety is a polynomial, upward-closed property (see Definition 5 in [5]), Theorem 3 in [5] yields the following complexity upper-bound for this problem.

Proposition 4. *The \forall -optimality problem is in coNP.*

Similarly to the case of the \exists -optimality problem in [5], we do not know whether the \forall -optimality problem is also coNP-hard. But we can show that it is at least as hard as the Hypergraph Duality Problem [8], called DUAL. Note that this problem is in coNP, but conjectured to be neither in P nor coNP-hard [9,10].

Proposition 5. *DUAL can be reduced in polynomial time to the \forall -optimality problem.*

The proof of this proposition, which is similar to the one of Proposition 7 in [5], is omitted due to the space constraints.

4 Investigating $\forall\exists$ -Safety

We now consider the case where the attacker's knowledge is assumed to be given by an $\mathcal{FL}\mathcal{E}$ concept. Since \mathcal{FL}_0 is a sublogic of $\mathcal{FL}\mathcal{E}$, concepts that are $\forall\exists$ -safe are also \forall -safe, but the opposite need not hold. The following proposition characterizes $\forall\exists$ -safety.

Proposition 6. *Let C be an \mathcal{EL} concept and $\mathcal{P} = \{D_1, \dots, D_p\}$ a redundancy-free \mathcal{EL} policy. Then C is $\forall\exists$ -safe for \mathcal{P} iff*

- 1.) $A \notin \text{con}(C)$ for all concept names $A \in \text{con}(D_1) \cup \dots \cup \text{con}(D_p)$, and
- 2.) for all existential restrictions $\exists r.D' \in \text{con}(D_1) \cup \dots \cup \text{con}(D_p)$, there is no concept of the form $\exists r.E$ in $\text{con}(C)$.

Proof. First, assume that C is not $\forall\exists$ -safe for \mathcal{P} . Hence, there is $D_i \in \mathcal{P}$ and an $\mathcal{FL}\mathcal{E}$ concept C'' such that C'' complies with \mathcal{P} , but $C \sqcap C'' \sqsubseteq D_i$. This subsumption implies that $A \in \text{con}(C) \cup \text{con}(C'')$ holds for all $A \in \text{con}(D_i)$. If there is an $A \in \text{con}(D_i)$ such that $A \in \text{con}(C)$, then property 1.) is violated. Otherwise, all $A \in \text{con}(D_i)$ belong to $\text{con}(C'')$. But then $C'' \not\sqsubseteq D_i$ can only be due to the fact that there is $\exists r.D' \in \text{con}(D_i)$ such that, for all $\exists r.C' \in \text{con}(C'')$, we have $C' \sqcap \text{filler}_r^\forall(C'') \not\sqsubseteq D'$. Applying Proposition 1 again to the subsumption $C \sqcap C'' \sqsubseteq D_i$ thus yields that there is $\exists r.E \in \text{con}(C)$ such that $E \sqcap \text{filler}_r^\forall(C'') \sqsubseteq D'$. Consequently, property 2.) is violated.

To show the other direction, assume that condition 1.) or 2.) is violated. If 1.) is violated, then there are $D_i \in \mathcal{P}$ and a concept name A such that $A \in \text{con}(C) \cap \text{con}(D_i)$. We modify D_i to C'' by removing A from the top-level conjunction of D_i . Then C'' is an \mathcal{EL} concept, and thus also an $\mathcal{FL}\mathcal{E}$ concept, such that $C'' \not\sqsubseteq D_i$ and $C \sqcap C'' \equiv C \sqcap D_i \sqsubseteq D_i$. Given $D \in \mathcal{P} \setminus \{D_i\}$ we have $C'' \not\sqsubseteq D$ since otherwise $D_i \sqsubseteq C'' \sqsubseteq D$ would contradict our assumption that \mathcal{P} is redundancy-free. Thus C is not $\forall\exists$ -safe for \mathcal{P} .

If condition 2.) is violated, then there are $D_i \in \mathcal{P}$ and existential restrictions $\exists r.D' \in \text{con}(D_i)$ and $\exists r.E \in \text{con}(C)$. Let C'' be obtained from D_i by replacing every existential restriction $\exists r.F$ from the top-level conjunction of D_i with the corresponding value restriction $\forall r.F$. To show that $C \sqcap C'' \sqsubseteq D_i$, it is sufficient

to show that $C \sqcap C'' \sqsubseteq \exists r.F$ for all $\exists r.F \in \text{con}(D_i)$. This is the case since $C \sqcap C'' \sqsubseteq \exists r.E \sqcap \forall r.F \sqsubseteq \exists r.(E \sqcap F) \sqsubseteq \exists r.F$.

It remains to show that C'' is compliant with \mathcal{P} , i.e., for all $D \in \mathcal{P}$ we have $C'' \not\sqsubseteq D$. If D contains an existential restriction for r , then this holds since C'' does not contain an existential restriction for r . In particular, this covers the case where $D = D_i$. If D does not contain an existential restriction for r , then the changes we made when going from D_i to C'' are not relevant for D , i.e., we have $C'' \sqsubseteq D$ iff $D_i \sqsubseteq D$. Since \mathcal{P} is redundancy-free, this yields $C'' \not\sqsubseteq D$. \square

Due to the simplicity of the conditions 1.) and 2.) in this proposition, it is now easy to show that all relevant computation or decision problems for $\forall\exists$ -safety are tractable.

Theorem 3. *Given \mathcal{EL} concepts C, C'' and a redundancy-free \mathcal{EL} policy \mathcal{P} , we*

- *can decide whether C is $\forall\exists$ -safe for \mathcal{P} ,*
- *can compute the unique optimal $\forall\exists$ -safe generalization of C for \mathcal{P} , and*
- *can decide whether C'' is an optimal $\forall\exists$ -safe generalization of C for \mathcal{P}*

in polynomial time.

Proof. First, note that the characterization of $\forall\exists$ -safety given in Proposition 6 can obviously be checked in polynomial time. Secondly, to obtain the optimal $\forall\exists$ -safe generalization of C for \mathcal{P} , we simply remove from $\text{con}(C)$ all concept names A with $A \in \text{con}(D_1) \cup \dots \cup \text{con}(D_p)$, and all existential restrictions $\exists r.E$ such that $\text{con}(D_1) \cup \dots \cup \text{con}(D_p)$ contains an existential restriction for the role r . This can clearly be done in polynomial time. Finally, to decide whether C'' is an optimal $\forall\exists$ -safe generalization of C for \mathcal{P} , apply the procedure just described to C , and check whether the resulting concept C' is equivalent to C'' . Since the subsumption problem is polynomial in \mathcal{EL} , this yields a polynomial-time decision procedure for the optimality problem. \square

5 Conclusion

We have investigated the notion of safety for a policy in the setting where the knowledge about individuals and the policy are given by \mathcal{EL} concepts, but the attacker's knowledge is assumed to be expressed in \mathcal{FL}_0 or \mathcal{FLE} . For both cases, we have characterized safety and have shown how to compute optimal safe generalizations of a given \mathcal{EL} concept w.r.t. a given \mathcal{EL} policy. For \mathcal{FL}_0 , the complexity results proved here are the same as the ones shown in [5] for the case where the attacker's knowledge is assumed to be expressed in \mathcal{EL} . Nevertheless, the characterizations of safety and the developed algorithms are, of course, different depending on whether \mathcal{FL}_0 or \mathcal{EL} is considered. For the case of \mathcal{FLE} , the characterization of safety developed in this paper is considerably simpler, and as a consequence the relevant decision and computation problems become tractable. While this may be seen as an advantage, it is actually due to the fact that, with an assumed stronger capability of the attacker, concepts need to be changed

more radically to make them safe. Thus, less knowledge can be preserved when publishing information in a privacy-preserving way.

In the future, we intend to continue this work in several directions. First, we want to extend our framework to a setting where the knowledge about individuals is given by an ABox that contains not only concept assertions, but also role assertions. As mentioned in the introduction, our idea for achieving this is to combine our approaches with the ones developed in [11,12]. In a second step, we intend to add a TBox that constrains the interpretation of concepts. This means that subsumption w.r.t. all interpretations is replaced by subsumption w.r.t. all models of the TBox. Since this subsumption relation is less easy to characterize than the characterization given in Proposition 1 for subsumption without a TBox, this will probably be a quite challenging task. As a starting point, we will also look at the case where the TBox satisfies certain cycle-restrictions (see, e.g., [1]). In addition, we intend to look also at other (combinations of) DLs in these settings.

Another interesting question is how to impose additional constraints on which knowledge is potentially available to an attacker. At the moment, we play it very safe by assuming that any knowledge expressible in the respective DL could be available to the attacker. This may, however, result in very general and uninformative safe generalizations, as in our example in the introduction when we assume the attacker’s knowledge is expressed in $\mathcal{FL}\mathcal{E}$. It might make sense to restrict the potential knowledge of the attacker about an individual to knowledge that really holds for this individuals (since we are not interested in what follows from lies about the individual), but the question is how to express this restriction formally. One possibility could be to assume that the “real world” is given by a finite interpretation, or by a knowledge base that approximates it.

A different approach for restricting the access to information contained in an ontology has been investigated in [4]. In this work, the axioms of the ontology are labeled with access restrictions, and users can only see (the consequences of) the axioms for which they have the right of access. In contrast, in our work the access restrictions are formulated on the the side of the consequences (in the form of policies) rather than on the side of the axioms. In addition, in [4] axioms are removed completely if a user does not have the right to access them, whereas in [5] and in the present work they are weakened appropriately. Nevertheless, it might be interesting to combine the two approaches, e.g., by producing variants of an axiom of different strength, and allowing a user that does not have the right to access the original axiom at least access to an appropriate weakening, which depends on the user’s access right.

References

1. F. Baader, S. Borgwardt, and B. Morawska. Extending unification in \mathcal{EL} towards general TBoxes. In *Proc. of the 13th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2012)*, pages 568–572. AAAI Press, 2012.
2. F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In L. P. Kaelbling and A. Saffiotti, editors, *Proceedings of the Nineteenth International Joint Conference*

- on *Artificial Intelligence IJCAI-05*, pages 364–369, Edinburgh, UK, 2005. Morgan-Kaufmann Publishers.
3. F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, New York, NY, USA, 2003.
 4. F. Baader, M. Knechtel, and R. Peñaloza. Context-dependent views to axioms and consequences of semantic web ontologies. *J. Web Semant.*, 12:22–40, 2012.
 5. F. Baader, F. Kriegel, and A. Nuradiansyah. Privacy-Preserving Ontology Publishing for \mathcal{EL} Instance Stores. In F. Calimeri and M. Manna, editors, *16th European Conference on Logics in Artificial Intelligence, JELIA 2019, Rende, Italy, May 8-10, 2019, Proceedings*, volume 11468 of *Lecture Notes in Computer Science*, pages 323–338. Springer, 2019.
 6. F. Baader, R. Küsters, and R. Molitor. Computing least common subsumers in description logics with existential restrictions. In T. Dean, editor, *Proc. of the 16th Int. Joint Conf. on Artificial Intelligence (IJCAI'99)*, pages 96–103, 1999.
 7. F. M. Donini, M. Lenzerini, D. Nardi, B. Hollunder, W. Nutt, and A. Marchetti-Spaccamela. The complexity of existential quantification in concept languages. *Artif. Intell.*, 53(2-3):309–327, 1992.
 8. T. Eiter and G. Gottlob. Hypergraph transversal computation and related problems in logic and AI. In S. Flesca, S. Greco, N. Leone, and G. Ianni, editors, *Logics in Artificial Intelligence, European Conference, JELIA 2002, Cosenza, Italy, September, 23-26, Proceedings*, volume 2424 of *Lecture Notes in Computer Science*, pages 549–564. Springer, 2002.
 9. M. L. Fredman and L. Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. *J. Algorithms*, 21(3):618–628, 1996.
 10. G. Gottlob and E. Malizia. Achieving new upper bounds for the hypergraph duality problem through logic. *SIAM J. Comput.*, 47(2):456–492, 2018.
 11. B. C. Grau and E. V. Kostylev. Logical foundations of privacy-preserving publishing of linked data. In D. Schuurmans and M. P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 943–949, 2016.
 12. B. C. Grau and E. V. Kostylev. Logical foundations of linked data anonymisation. *J. Artif. Intell. Res.*, 64:253–314, 2019.
 13. R. Hoehndorf, P. N. Schofield, and G. V. Gkoutos. The role of ontologies in biological and biomedical research: A functional perspective. *Brief. Bioinform.*, 16(6):1069–1080, 2015.
 14. I. Horrocks, L. Li, D. Turi, and S. Bechhofer. The instance store: DL reasoning with large numbers of individuals. In V. Haarslev and R. Möller, editors, *Proceedings of the 2004 International Workshop on Description Logics (DL2004)*, Whistler, British Columbia, Canada, June 6-8, 2004, 2004.
 15. H. J. Levesque and R. J. Brachman. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3:78–93, 1987.
 16. B. Suntisrivaraporn. *Polynomial-Time Reasoning Support for Design and Maintenance of Large-Scale Biomedical Ontologies*. Doctoral thesis, Technische Universität Dresden, Dresden, Germany, 2009.