Noname manuscript No.
(will be inserted by the editor)

Reasoning in Description Logic Ontologies for Privacy Management

Adrian Nuradiansyah

Received: date / Accepted: date

1 Introduction

Ontologies have been well-known for sharing common understanding of the structure of information in various application domains, e.g., Semantic Web [7] or medicine [3]. Some real examples of ontologies that have been published are [3, 14]. Ontologies are also used to provide more semantics to formally describe a meaning of the data. In contrast to relational databases, information stored in ontologies are mainly assumed to be incomplete, which means that we can deduce additional facts from the ontologies, which are not explicitly stated there. Another difference with the database paradigm is that ontologies normally employ *open-world assumption* stating that knowledge that is not explicitly stored in the ontologies and cannot be inferred, is neither assumed to be true nor to be false.

In addition to semantic features, ontologies are formulated using languages that are much more expressive than database schema languages. One of the most common languages formalizing ontology is the Web Ontology Language (OWL)[1] that has been standardized by W3C and frequently used in many application domains. This standardization results in a connection with a family of knowledge representation languages, called Description Logics (DLs) (see [1]), that are known as a fragment of first-order logic.

However, all these attracting features of ontologies are still prone to privacy violations. Assume that there is a DL ontology $\mathcal{O}$ containing the following information:

$$\mathcal{O} := \{\, \exists \mathsf{seenBy}.\mathsf{Oncologist} \sqsubseteq \exists \mathsf{suffer}.\mathsf{Cancer},\\ \exists \mathsf{seenBy}.\mathsf{Oncologist}(x) \,\}.$$

The first axiom of the form of *general concept inclusion* (GCI) states that if someone is seen by an oncologist, then he/she suffers from cancer, while the second axiom says that the individual x, whose name is anonymous, is seen by an oncologist. Now, suppose there is a privacy policy $\mathcal{P}$ expressing that the public is not allowed to know any disease of any individual of the ontology $\mathcal{O}$. It is also emphasized that $\mathcal{O}$ should satisfy $\mathcal{P}$ before $\mathcal{O}$ is published. However, the derivable fact "x suffers from cancer" from $\mathcal{O}$ means that $\mathcal{O}$ does not obey $\mathcal{P}$.

In the doctoral dissertation [12], we covered three anticipation steps that need to be considered before publishing an ontology. First, we ask if the confidential information of individuals, such as their identity, is kept hidden or not w.r.t. an ontology. If it is the case that *a privacy breach is detected*, then we deal with the second step where one needs to *repair the ontology* such that the sensitive properties of individuals are not possible to be disclosed unauthorizedly, but at the same time the utility value of the ontology remains preserved. The last step is to guarantee if the solution for ontology repairs can *avoid a linkage attack* from possible attackers that have extra knowledge from different sources. In particular, this sort of attacks may occur when the combination of the repaired ontology with the background knowledge of the attackers still violates a privacy policy.

Indeed, considering such malicious attacks for ontologies in the last step above has also been recently investigated in different contexts, such as in the area of linked data [6, 9] or in the area of ontology-based data integration (OBDI) [4]. In the context of privacy in DL ontologies, to the best of our knowledge, the studies of preserving identity or reckoning such linkage attacks were still unexplored, whereas the studies of ontology repairs have been carried out by e.g., [10, 16] with different settings and motivations. A summary of these new studies as well as their results in [12] is presented in this article.

Adrian Nuradiansyah
Institute of Theoretical Computer Science,
Technische Universität Dresden, Dresden, Germany
E-mail: adrian.nuradiansyah@tu-dresden.de

[1] See https://www.w3.org/TR/owl-features/

2 Detecting Privacy Breaches

The work on detecting if there is a privacy leak occurring in an ontology has been investigated in many literatures, such as [5, 15]. Most of the previous work designed their privacy-preserving ontology system by conceiving the privacy policy as a property of individuals. A privacy policy is represented as a query and an ontology is said to be *compliant* with a privacy policy if none of the sensitive answers to the query representing the policy hold in the ontology. In other words, we cannot deduce that an individual (or a tuple of individuals) is a member of the sensitive answers to certain queries representing a privacy policy. In the medical example we had in the previous section, we may say that "suffering from cancer" is a property of the anonymous x that needs to be protected.

As argued by [8], one of the property of individuals that should be importantly protected is their identity, which has not been formally considered, at least, in the context of DL ontologies. Suppose that the ontology $\mathcal{O}$ is now extended to an ontology $\mathcal{O}_1$ by adding the following axioms:

$$\exists \mathsf{suffer}.\mathsf{Cancer} \equiv \{\mathsf{JOHN}, \mathsf{LINDA}\},$$
$$\mathsf{Female} \sqsubseteq \neg \mathsf{Male},$$
$$\mathsf{Male}(x), \mathsf{Male}(\mathsf{JOHN}), \mathsf{Female}(\mathsf{LINDA}).$$

Note that the first axiom, including a DL concept called *(one-of) nominals*, tells that people suffering from cancer are only John and Linda. Now, if we do reasoning over $\mathcal{O}_1$, then a consequence, such as $\{x\} \equiv \{\mathsf{JOHN}, \mathsf{LINDA}\}$, can be inferred, stating that the anonymous x is either John or Linda. In fact, the only male in that set of individuals is John, and thus we can infer that x is actually John w.r.t. $\mathcal{O}_1$, which means that the identity of x is now revealed.

To this end, we introduced the *identity problem* asking whether two individuals a, b are equal w.r.t. an ontology in general. We show that this problem is trivial for all DLs that are fragments of first-order logic without equality since we can always deduce that no equality between two individuals w.r.t. ontologies formulated in those DLs. Then, we introduced *DLs with equality power*, which DLs with nominals, number restrictions, functional roles, or functional dependencies belong to, and in which the identity problem is non-trivial. We showed that for these DLs the identity problem has the same complexity as the instance problem, which is the problem asking if an individual a is an instance of a DL concept C w.r.t. an ontology $\mathcal{O}$.

Theorem 1 *For all DLs with equality power, the identity problem can be reduced to the instance problem.*

We extended the identity problem to a role-based access control setting and to a setting where an attacker does not want to know the exact identity of an anonymous individual x and it is sufficient for him to deduce that the identity of x belongs to a set of known individuals of cardinality smaller than k. We showed that problems in both settings can be reduced to the instance problem for DLs with equality power.

Learning the problems above, we see that the identity problem and its extensions can eventually be reduced to classical reasoning problems in DLs. This means that in the following sections, we do not need to specifically consider the privacy policy to be written as queries asking for identity, but also can be standard queries, such as subsumption queries, instance queries, or even conjunctive queries.

3 Repairing Ontologies

If one can derive sensitive information about individuals from an ontology $\mathcal{O}$, then it makes sense to repair $\mathcal{O}$ such that the secret consequences α are no longer entailed by the ontology repair $\mathcal{O}'$. We additionally require that $\mathcal{O}'$ should be implied by $\mathcal{O}$. Such a repair is optimal if there is no repair $\mathcal{O}''$ that strictly implies $\mathcal{O}'$. However, we show that optimal repairs need not always exists.

In DL communities, initial motivations for repairing ontologies came from a question on why a consequence computed by a DL reasoner actually follows from an ontology. This initiated the work on computing so-called *justifications* in [2, 13], i.e., minimal subsets of the ontology that have the consequence in question. Considering all justifications, one may construct a hitting set of the justifications, i.e., a set of axioms containing at least one axiom from each justification. Removing minimal hitting sets yields maximum subsets of the ontology that do not entail the consequence.

However, the main problem with this approach is that removing complete axioms may also eliminate consequences that are actually wanted. Instead, we proposed to replace axioms directly by weaker ones, i.e., axioms that have less consequences . This motivates us to introduce the notion of *gentle repair* [12]. In this approach, we generally weaken one axiom from each justification such that the modified justifications no longer have the consequence.

As an illustration, we define an ontology $\mathcal{O}_2$ consisting of the following axioms:

$$\exists \mathsf{seenBy}.(\mathsf{Doctor} \sqcap \exists \mathsf{worksIn}.\mathsf{Oncology}) \sqsubseteq \exists \mathsf{suffer}.\mathsf{Cancer},$$
$$\exists \mathsf{worksIn}.\mathsf{Nuclear} \sqsubseteq \exists \mathsf{seenBy}.(\mathsf{Doctor} \sqcap \exists \mathsf{worksIn}.\mathsf{Oncology}),$$
$$\exists \mathsf{worksIn}.\mathsf{Nuclear}(\mathsf{LINDA}).$$

Using the same policy $\mathcal{P}$, we can see that $\mathcal{O}_2$ does not comply with $\mathcal{P}$. If we are only allowed to modify the second axiom and this modification is based on the axiom removal, then the modified ontology is compliant with $\mathcal{P}$. This implies that consequences, such as "every worker of a nuclear

company is seen by a doctor" or "every worker of a nuclear company is seen by someone working in an oncology department" are gone. Suppose that such consequences are useful. Thus, to retain them, while achieving compliance property at the same time, we weaken the second axiom to

$$\exists \mathsf{worksIn}.\mathsf{Nuclear} \sqsubseteq \exists \mathsf{seenBy}.\mathsf{Doctor} \sqcap \exists \mathsf{seenBy}.\exists \mathsf{worksIn}.\mathsf{Oncology}$$

so that the modified ontology is now being compliant with $\mathcal{P}$ without losing the wanted consequences.

The next theorem states two important results we have in our gentle repair framework.

Theorem 2 *The following results hold in our gentle repair framework:*

1. *The gentle repair approach needs to be iterated,*
2. *At most exponentially many iterations are needed to reach a gentle repair.*

What it means by the first result is that applying this approach once does not necessarily remove the unwanted consequence.

Instead of allowing arbitrary ways to weaken axioms, we introduce the notion of *weakening relation* $\succ$, which is formally a binary relation such that for each $(\beta,\gamma) \in \succ$, the axiom γ is weaker than β. Intuitively, the larger the weakening relation is, the smaller are weakening steps needed to reach a gentle repair, which means that more iterations are performed. Moreover, several conditions of weakening relations are introduced to equip the gentle repair approach with better algorithmic properties that can, for instance, guarantee linear or polynomial number of iterations.

In a situation where we have a justification J, an (unwanted) consequence α, and an axiom $\beta \in J$, if all conditions of such weakening relations $\succ$ are satisfied, then by a search along the one-step relation

$$\succ_1 := \{(\beta_1,\beta_2) \mid \beta_1 \succ \beta_2 \text{ and } \nexists \beta_3 \text{ s.t. } \beta_1 \succ \beta_3 \succ \beta_2\}$$

one can find a maximally strong weakening of β, which is an axiom γ such that $\beta \succ \gamma$ and $(J \setminus \{\beta\}) \cup \{\gamma\} \not\models \alpha$, and there is no stronger axiom δ, where $\beta \succ \delta \succ \gamma$, with this property. Intuitively, using a maximally strong weakening, the ontology is changed in a minimal way.

For applying the above repair framework, we focused to the DL $\mathcal{EL}$ and introduced specific weakening relations for it, defined based on generalizing the right-hand side of GCIs semantically ($\succ^{sub}$) and syntactically ($\succ^{syn}$). Considering $\mathcal{O}_2$ as an $\mathcal{EL}$ ontology, the way we weaken the axiom in $\mathcal{O}_2$ illustrated above is based on the use of the weakening relation $\succ^{sub}$. If we apply $\succ^{syn}$ to the second axiom in $\mathcal{O}_2$, then we cannot split existential restrictions as what $\succ^{sub}$ can do, but we may have weaker axioms such as

$$\exists \mathsf{worksIn}.\mathsf{Nuclear} \sqsubseteq \exists \mathsf{seenBy}.\mathsf{Doctor} \text{ or}$$
$$\exists \mathsf{worksIn}.\mathsf{Nuclear} \sqsubseteq \exists \mathsf{seenBy}.\exists \mathsf{worksIn}.\mathsf{Oncology}$$

We showed that $\succ^{syn} \subseteq \succ^{sub}$, which means that it takes larger weakening steps to reach a gentle repair using $\succ^{syn}$. Complexity wise, $\succ^{syn}$ behaves better than $\succ^{sub}$. For instance, the length of one-step weakening chain for each β, i.e., $\beta \succ_1^{syn} \beta_1 \succ_1^{syn} \beta_2 \succ_1^{syn} \ldots$, is linearly bounded in the size of β. In contrast, one can only provide a non-elementary bound for $\succ_1^{sub}$. Furthermore, we showed that maximally strongest weakenings can be effectively computed using both weakening relations in $\mathcal{EL}$. In particular, one (all) maximally strongest weakening(s) can be computed in polynomial (exponential) time w.r.t. $\succ^{syn}$.

4 Avoiding Linkage Attacks

The framework for ontology repair we described above results in a modified ontology that is *policy-compliant*. However, this property is still not enough if there is an attacker's knowledge that is different with our policy compliant ontology and it turns out that the combination of these two knowledge is again not compliant with the policy.

To address the issue above, we considered the *policy-safety* property adopted from [9], which requires that the combination of the published ontology with any other compliant ontology is again compliant w.r.t. the policy. Now, consider that this setting is applied to a specific type of $\mathcal{EL}$ ontologies, called $\mathcal{EL}$ instance stores, that has no individual relationships [11], no GCIs, and only contains instance axioms $C(a)$. It means that all the information about an individual a is given by an $\mathcal{EL}$ concept C. Then, a policy is given by a set of instance queries, i.e., by $\mathcal{EL}$ concepts $D_1,\ldots,D_n$.

Using another medical example, suppose that the policy only consists of a concept

$$D = \mathsf{Patient} \sqcap \exists \mathsf{seenBy}.(\mathsf{Doctor} \sqcap \exists \mathsf{worksIn}.\mathsf{Oncology}),$$

which says that one should not able to find out who are the patients that are seen by a doctor working in an oncology department. Moreover, it is published that John is an instance of the concept

$$C = \mathsf{Male} \sqcap \exists \mathsf{seenBy}.(\mathsf{Doctor} \sqcap \exists \mathsf{worksIn}.\mathsf{Oncology}),$$

which is compliant with the policy since $C \not\sqsubseteq D$. However, it is not safe if there is an attacker knowing that John is a patient since if this knowledge is combined with C, then $C \sqcap \mathsf{Patient} \sqsubseteq D$. In contrast, the concept

$$C' = \mathsf{Male} \sqcap \exists \mathsf{seenBy}.(\mathsf{Doctor} \sqcap \exists \mathsf{worksIn}.\top),$$

where $C \sqsubseteq C'$, is a safe generalization of C, as shown in [12], in particular when the attacker's knowledge is encoded as an $\mathcal{EL}$ concept.

The generalization process illustrated above leads us to the *optimality* property asking if the modified ontology or

concept is compliant (safe) w.r.t. a policy and changes the original ontology in a minimal way. To this end, we developed algorithms for computing optimal compliant (safe) generalizations of $\mathcal{EL}$ concepts w.r.t. $\mathcal{EL}$ policies. When dealing with different expressiveness of attackers' knowledge, such algorithms may have different complexity results as stated in the following theorem.

Theorem 3 *Given an $\mathcal{EL}$ concept C and an $\mathcal{EL}$ policy $\mathcal{P}$, we can compute the optimal compliant (safe) generalizations of $\mathcal{EL}$ concept w.r.t. $\mathcal{P}$*

1. *in ExpTime if the attacker's knowledge is written as an $\mathcal{EL}$ concept, and*
2. *in PTime if the attacker's knowledge is written in the more expressive $\mathcal{FLE}$ concept.*

Likewise, if we view optimality as a decision problem, then a coNP upper bound is given for both compliance and safety when the attacker's knowledge is written as an $\mathcal{EL}$ concept, but when it is modeled as an $\mathcal{FLE}$ concept, this optimality problem becomes PTime.

We further investigated the case where both published ontology and attackers' knowledge may contain individual relationships. If the policy is an instance query, then the complexities of the corresponding compliance, safety, and optimality problems remain the same as in our previous setting for the case of instance stores. If we upgrade the policy form to be a conjunctive query, then most of the complexity results we have lie on the second or the third level of the polynomial hierarchy.

5 Conclusions

We have seen that for each anticipation step discussed in Section 1, we contributed in introducing relevant reasoning problems in DLs and presenting frameworks, inference procedures and algorithms that provide automated support for dealing with those problems. In particular, this work is coupled with investigations on the complexity of the procedures and algorithms as well as the worst-case complexities of the problems solved by them.

Obviously, this work has many directions to be explored. For instance, probabilistic assumptions are taken into account to annotate ontology axioms and equalities between individuals that only hold with a certain probability may be derived. Additionally, avoiding ontologies from attackers' knowledge that are complete (*closed world*) or featured with integrity constraints is also a realistic crucial challenge in privacy issues nowadays.

Acknowledgements The author is deeply grateful to Franz Baader for his supports and insightful comments. This work was funded by German Research Foundation (DFG) within the Research Training Group 1907 "RoSI".

References

1. F. Baader, I. Horrocks, C. Lutz, and U. Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017.
2. F. Baader and B. Suntisrivaraporn. Debugging SNOMED CT using axiom pinpointing in the description logic EL+. In R. Cornet and K. A. Spackman, editors, *Proceedings of the 3rd International Conference on Knowledge Representation in Medicine*, 2008.
3. M. Bada, R. Stevens, C. A. Goble, Y. Gil, M. Ashburner, J. A. Blake, J. M. Cherry, M. A. Harris, and S. Lewis. A short study on the success of the gene ontology. *J. Web Semant.*, 1(2):235–240, 2004.
4. M. Benedikt, B. C. Grau, and E. V. Kostylev. Source information disclosure in ontology-based data integration. In S. P. Singh and S. Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 1056–1062. AAAI Press, 2017.
5. D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. View-based query answering over description logic ontologies. In G. Brewka and J. Lang, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 11th International Conference, KR*, pages 242–251, 2008.
6. R. Delanaux, A. Bonifati, M. Rousset, and R. Thion. RDF graph anonymization robust to data linkage. In R. Cheng, N. Mamoulis, Y. Sun, and X. Huang, editors, *Web Information Systems Engineering - WISE 2019 - 20th International Conference*, LNCS, pages 491–506. Springer, 2019.
7. B. Glimm and H. Stuckenschmidt. 15 years of semantic web: An incomplete survey. *KI*, 30(2):117–130, 2016.
8. B. C. Grau. Privacy in ontology-based information systems: A pending matter. *Semantic Web*, 1(1-2):137–141, 2010.
9. B. C. Grau and E. V. Kostylev. Logical foundations of linked data anonymisation. *J. Artif. Intell. Res.*, 64:253–314, 2019.
10. M. Horridge. *Justification based Explanation in Ontologies*. PhD thesis, University of Manchester, 2011.
11. 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*, 2004.
12. A. Nuradiansyah. *Reasoning in Description Logic Ontologies for Privacy Management*. PhD thesis, Technische Universität Dresden, 2019.
13. S. Schlobach and R. Cornet. Non-standard reasoning services for the debugging of description logic terminologies. In *IJCAI-03, Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 355–362, 2003.
14. M. Q. Stearns, C. Price, K. A. Spackman, and A. Y. Wang. SNOMED clinical terms: overview of the development process and project status. In *AMIA 2001, American Medical Informatics Association Annual Symposium*, 2001.
15. P. Stouppa and T. Studer. Data privacy for knowledge bases. In S. N. Artëmov and A. Nerode, editors, *Logical Foundations of Computer Science, International Symposium, LFCS*, pages 409–421, 2009.
16. N. Troquard, R. Confalonieri, P. Galliani, R. Peñaloza, D. Porello, and O. Kutz. Repairing ontologies via axiom weakening. In S. A. McIlraith and K. Q. Weinberger, editors, *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 1981–1988, 2018.