A Logical Approach for Preserving Confidentiality in Shared Knowledge Bases

Erika Guetti Suca and Flávio Soares Corrêa da Silva

[1] Institute of Mathematics and Statistics – University of São Paulo – Brazil

{fcs, eguetti}@ime.usp.br

Abstract. *The control of interconnection mechanisms in shared knowledge bases is important to ensure that sensitive information is not extracted in an inappropriate way from connected bases. Our goal is to propose a logical model to specify mechanisms for query control, reasoning and evolution of knowledge bases and their corresponding ontologies, ensuring the confidentiality of information whenever appropriate. We describe the techniques currently in development to address this problem and show their capabilities and limitations. Finally, we introduce the requirements for a software tool to allow a designer of knowledge bases to define sensitive data elements and implement mechanisms to ensure their confidentiality.*

1. Introduction

Around 17% of Brazilian companies have suffered from cybercriminal actions. More than 33% of Brazilian companies still do not have an appropriate security system and they claim not to have adequate knowledge about this subject [Maia and Whitehead 2014]. Many cybercrimes do not require advanced technology and are caused by human errors in preservation of sensitive data. Technology itself cannot solve all issues related to confidentiality, yet it is important to information systems designers to take responsibility for the privacy of managed data. For that, access to data must be designed and implemented to prevent confidentiality breaches.

We aim to protect the possibility to infer confidential information and improperly extract it from the connected knowledge bases. We have formally defined the problem of confidentiality, which can be summarized as how to generate a published views knowledge base that preserves the confidentiality given an insecure knowledge base.

We want to develop a software tool, based on our model, that can automatically identify and suggest changes in rules or axioms of the knowledge bases that lead directly to the disclosure of a secret. Assisting the knowledge engineer to improve the design of the knowledge base overcoming possible types of attacks.

Our purpose is to apply confidentiality preservation techniques that minimize the risk of breaches following heuristics based on properties of knowledge bases and their planned use. We consider than every data publishing scenario has its own assumptions and requirements on the data holder, the data recipients and the data publishing purpose. Therefore, we will implement two of the main approaches in preserving confidentiality: logic-based and anonymization.

The present article is organized as follows: in section 2 we review related work, in section 3 we present motivating examples, in section 4 we describe our proposal, in 5 we

explain the expected outputs of our work and finally in section 6 we present a discussion and our conclusions.

2. Confidentiality Problem

We approach the confidentiality problem from two perspectives: inference control strategies and the inference modeling problem. Inference control strategies are based on inference time: (1) During knowledge base design time – the main advantages is that it is usually fast since it only considers the database schema and the corresponding constraints without the actual instances, but the evolution of data is generally not covered in the model. (2) During query-processing time – it provides maximal data availability because all disclosed data can be evaluated to verify the existence of an inference channel. However, it is usually more expensive and time consuming than the design time approach and affects system usage.

In the present work we explore in greater detail the second approach.

In the inference modeling problem we aim at the implementation of two main approaches: (1) Logic-based – policy requirements are enforced when the user requests access to information by means of a query. In the field of ontologies this technique is called Controlled Query Evaluation(CQE) [Bonatti et al. 2015, Eldora et al. 2011]. The main advantages are the clear formalization and decidability results, as well as the independence of the application domain. However, logic-based systems have high complexity, making it expensive for large Web applications. (2) Anonymization based – it refers to privacy preserving data publishing (PPDP) assuming that sensitive data must be retained for data analysis. The main attack and privacy models come from database techniques based on Statistical Disclosure Control(SDC) [Fung et al. 2010]. In knowledge bases the anonymization has been applied mainly in the medical area [Grau and Kostylev 2016], [Domingo-Ferrer et al. 2013]. This model has a smaller complexity than logic-based models, it works on specific domains, with little formalization of the techniques.

3. Motivating Examples

We propose two examples to motivate the need to preserve confidentiality, but we can easily imagine similar needs in many other scenarios. The two examples considered here relate to healthcare.

Example 1: Anonymizing Healthcare Data

Consider the raw patient data in Table 1 where each record represents a surgery case with the patient specific information. *Job*, *Sex*, and *Age* are quasi-identifying(QID) attributes, this attributes uniquely identify an individual. The hospital wants to release the Table 1 for the purpose of classification analysis on the class attribute, *Transfuse*, which has two values, *YES* and *NO*, indicating whether or not the patient has received blood transfusion. Without a loss of generality, we assume that the only sensitive value in *Surgery* is *Transgender*. Table 2 shows the data after the anonymization using the LKC-privacy model [Fung et al. 2010] and after processing in order to generalize the records into equivalence groups so that each group contains at least k records with respect to some QID attributes. The general intuition of LKC-privacy is to ensure that every combination of values in QID with maximum length L, they are shared by at least K records, and the

Table 1. Raw patient data

	Quasi-identifier (QID)			*Class*	*Sensitive*
ID	**Job**	**Sex**	**Age**	**Transfuse**	**Surgery**
1	Janitor	M	34	Yes	Transgender
2	Doctor	M	58	No	Plastic
3	Mover	M	34	Yes	Transgender
4	Lawyer	M	24	No	Vascular
5	Mover	M	58	No	Urology
6	Doctor	M	24	No	Urology
7	Lawyer	F	58	No	Plastic
8	Carpenter	F	63	Yes	Vascular
9	Technician	F	63	Yes	Plastic

Table 2. Anonymous data (L=2, K=2, C = 0.5)

	Quasi-identifier (QID)			*Class*	*Sensitive*
ID	**Job**	**Sex**	**Age**	**Transfuse**	**Surgery**
1	Non-Technical	M	[30,60)	Yes	Transgender
2	Professional	M	[30,60)	No	Plastic
3	Non-Technical	M	[30,60)	Yes	Transgender
4	Professional	M	[1,30)	No	Vascular
5	Non-Technical	M	[30,60)	No	Urology
6	Professional	M	[1,30)	No	Urology
7	Professional	F	[30,60)	No	Plastic
8	Technical	F	[60,99)	Yes	Vascular
9	Technical	F	[60,99)	Yes	Plastic

confidence of inferring any sensitive values in S is not greater than C , where L , K , C are thresholds and S is a set of sensitive values specified by the data holder(the hospital). In this way, the sensitive values in each qid group are diversified enough to disorient confident inferences [Mohammed et al. 2009]. There are several anonymization models, depending on the requirements in the publication of the data.

Example 2: Protecting Confidentiality Across Several Institutions

The citizen Jane needs to take a certain preventive medicine for breast cancer. Suppose Jane does not want her physician or the pharmacy to supply the details of the prescription to her health insurance company because she does not want to risk an increase in her health insurance premium on the basis of the fact that medicine she has been prescribed is intended for use by women who are believed to have a high risk of developing breast cancer. In such a setting, in order for Jane to be reimbursed by her insurance company, the pharmacy needs to be able to certify to the insurance company, through a trusted third party, that Jane has indeed incurred a medical expense that is covered by her insurance policy [Bao et al. 2007]. In a simple way, consider K_P, the Pharmacy Knowledge Base and K_I as Insurance Company Knowledge Base, S_J are Jane's secrets, such that $K_P \cap K_I \not\subseteq S_J$.

4. Proposal

We have developed a simple formal confidentiality model M adapted from the main attack and privacy models found in the literature [Cuenca Grau and Horrocks 2008, Bonatti et al. 2015]. We consider a single knowledge base as the union of several knowledge bases, and the notion of *logical consequence*, for all knowledge base K will be denoted by $C_n(K)$.

We contemplate the problem of confidentiality involves two sub problems, namely (1) the secure publishing of data based on query-processing and (2) the secure evolution of data.

Definition 1 *Let M be a Simple Confidentiality Model (SCM) as follows:*

- *KB, is a knowledge base.*
- *U, is a set of users of KB. Different users access to different views of the KB.*
- *For all $u \in U$:*
  - *S_u is a finite set of secrecies that should not be disclosed to the user u.*
  - *The queries from one user u are answered using a view of the knowledge base $KB_u \subseteq KB$. KB_u is a secure view if $C_n(KB_u) \cap S_u = \emptyset$.*
  - *A view KB_u is maximal secure if it is secure and there exists no K' such as $KB_u \subset K' \subseteq KB$ and $C_n(K') \cap S_u = \emptyset$.*
  - *BK_u, is the set of statements that features the background knowledge of user u.*
- *f is a filtering function that maps for each $u \in U$, a view $V \subseteq C_n(KB_u)$.*

Responsibility for the publication of data is given by the function f of Definition 1, f implements a confidentiality strategy.

Definition 2 *Secure Publishing: f is secure if for all $u \in U$ and $s \in S_u$, there exists $K \in KB$, such that K is the answer to a query Q of user u:*

- *$f(K, u)$ generates secure views KB_u using a specific confidentiality preservation algorithm, and*
- *$s \notin C_n(K \cup BK_u)$.*

Definition 3 *Secure Evolution: Given $KB = (S, D)$, S is the knowledge base scheme and D represents the data sets. The evolution $KB = (S, D)$ to $KB' = (S', D')$ is secure w.r.t. Q and a secure view V if the confidentiality of $KB = (S, D)$ entails $KB' = (S', D')$ with a secure view V', assuming that V' was generated using the same definitions of view V. We can distinguish two types of evolution: when evolution happens in S or when the change occurs in D.*

- *Suppose S does not contain the schema β and $S' = S \cup \{\beta\}$. Then $KB' = (S', D)$ does not take break confidentiality since S' not introduce any correlation with any secrecies of KB.*
- *Confidentiality is independent of evolution in D, since it is not related to some secret query.*

5. Expected Outcomes and Results

We are working on the properties and general requirements for a software tool for the performance evaluation of the confidentiality preservation techniques and heuristic strategies for assurance of confidentiality. Assisting the knowledge engineer to correct the design of the knowledge base identifying axioms that involve to the disclosure of a secret. Beyond the improvement of our model, some specific expected outcomes of our work include:

- Systematization of heuristic evaluation of the types and confidentiality preservation techniques.
- Developing methods that allow the users to judge the correctness of the data need to support flexible conflict resolution.

- Systematization of techniques to evaluate the completeness and correctness of our results.
- Implementation of a software tool to ontologies to ensure the confidentiality of selected pieces of information. The tool can be included as a plugin of Protégé[1], allowing to assess the strengths and weaknesses of the implementations of techniques and features of ontologies provided by the user input.
- Comparative analysis of the indicators provided by the metrics of the studied heuristic techniques.
- As a initial case study, we will developed the examples presented in this paper.

Figure 1. Confidentiality Preservation Tool.

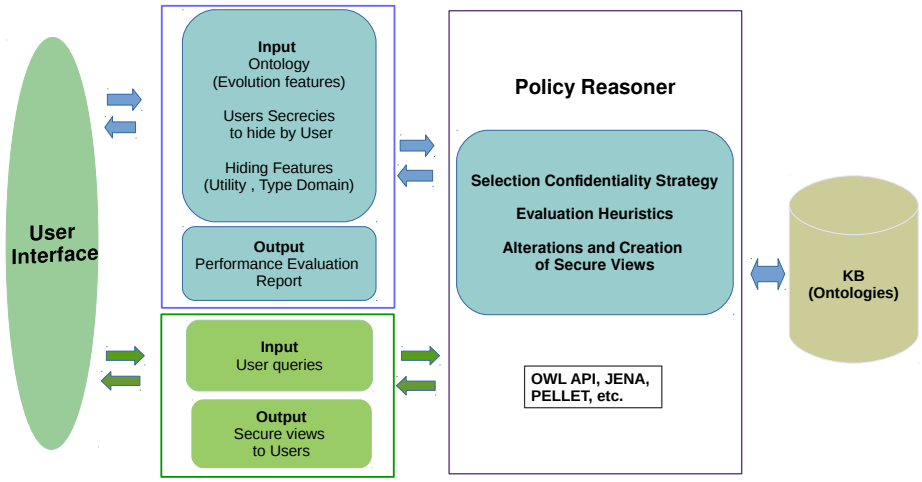

Each strategy is directly related to the purpose of the secure views of a KB. One of our goals in this project is to identify heuristics based on properties of knowledge bases and their planned use to identify and apply confidentiality preservation techniques that minimize the risk of breaches. At present, some heuristics can be sketched as follows:

- Data anonymization should be applied when:
  - The goal is to preserve the semantics of the data, omitting personal data, e.g. to establish a balance between retaining context and protecting participants.
  - The purpose is to study the properties of a data set without allowing the identification of a particular individual.
  - It is important not to interfere with the usefulness of the original knowledge base.
  - It is important to preserve the original data and reversibility of the securing process is a requirement.
- Creating secure views following logical approaches should be applied when:
  - The availability of secure views is not restricted to preserve a set of data.
  - It is totally independent of the application domain.
  - It can be adjusted to a specific technique.

In Figure 1 we have a generic description of the software tool we plan to develop in this project. The tool will be released as open source in GitHub.

[1] http://protege.stanford.edu/

6. Conclusions

In this paper we present the following items:

- We formally define the problem of generating a knowledge base that preserves confidentiality from an insecure knowledge base.
- We present the properties and general requirements for a software tool proposed for the evaluation of performance of confidentiality preservation techniques and heuristic building strategies to guarantee confidentiality in specific cases.
- A future direction of our work is to consider other forms of distortion of the knowledge base to ensure confidentiality. For example, we can explore not only remove elements of the knowledge base, but we may add new elements.

References

[Bao et al. 2007] Bao, J., Slutzki, G., and Honavar, V. (2007). Privacy-preserving reasoning on the semanticweb. pages 791–797.

[Bonatti et al. 2015] Bonatti, P. A., Petrova, I. M., and Sauro, L. (2015). Optimized construction of secure knowledge-base views. In Calvanese, D. and Konev, B., editors, *Description Logics*, volume 1350 of *CEUR Workshop Proceedings*. CEUR-WS.org.

[Cuenca Grau and Horrocks 2008] Cuenca Grau, B. and Horrocks, I. (2008). Privacy-preserving query answering in logic-based information systems. In *Proceedings of the 2008 Conference on ECAI 2008: 18th European Conference on Artificial Intelligence*, pages 40–44, Amsterdam, The Netherlands, The Netherlands. IOS Press.

[Domingo-Ferrer et al. 2013] Domingo-Ferrer, J., Sánchez, D., and Rufian-Torrell, G. (2013). Anonymization of nominal data based on semantic marginality. *Information Sciences*, 242:35 – 48.

[Eldora et al. 2011] Eldora, Knechtel, M., and Peñaloza, R. (2011). Correcting access restrictions to a consequence more flexibly. In Rosati, R., Rudolph, S., and Zakharyaschev, M., editors, *Description Logics*, volume 745. CEUR-WS.org.

[Fung et al. 2010] Fung, B. C., Wang, K., Fu, A. W.-C., and Yu, P. S. (2010). *Introduction to Privacy-Preserving Data Publishing: Concepts and Techniques*. Chapman & Hall/CRC, 1st edition.

[Grau and Kostylev 2016] Grau, B. C. and Kostylev, E. V. (2016). Logical foundations of privacy-preserving publishing of linked data. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 943–949.

[Maia and Whitehead 2014] Maia, O. and Whitehead, M. (2014). Pesquisa global sobre crimes econômicos. Technical report, PWC Global.

[Mohammed et al. 2009] Mohammed, N., Fung, B. C., Hung, P. C., and Lee, C.-k. (2009). Anonymizing healthcare data: A case study on the blood transfusion service. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1285–1294. ACM.