7-2010

Secrecy-preserving Query Answering for Instance Checking in EL

Jia Tao
Iowa State University, jtao510@gmail.com

Giora Slutzki
Iowa State University, slutzki@iastate.edu

Vasant Honavar
Iowa State University

Recommended Citation

Tao, Jia; Slutzki, Giora; and Honavar, Vasant, "Secrecy-preserving Query Answering for Instance Checking in EL" (2010). *Computer Science Technical Reports*. 234.
http://lib.dr.iastate.edu/cs_techreports/234

Secrecy-preserving Query Answering for Instance Checking in EL

Abstract
We consider the problem of answering queries against a knowledge base (KB) using secrets, whenever it is possible to do so without compromising secrets. We study query answering against EL knowledge bases. We provide a polynomial time algorithm that, given an EL KB Sigma, a set S of secrets to be protected and a query q or the form C(a) or r(a,b), outputs ``Yes'' whenever Sigma entails q and the answer to q, together with the answers to any previous queries answered by the KB, does not allow the querying agent to deduce any of the secrets in S. This approach allows more flexible information sharing than is possible with traditional access control mechanisms.

Keywords
Description Logic, Secrecy, EL

Disciplines
Artificial Intelligence and Robotics

Secrecy-preserving Query Answering for Instance Checking in $\mathcal{EL}$

Jia Tao, Giora Slutzki and Vasant Honavar

TR #10-03a
June 2010, revised July 2010

This is a corrected version for TR10-03.

Department of Computer Science
226 Atanasoff Hall
Iowa State University
Ames, Iowa 50010-1040, USA

Secrecy-preserving Query Answering for Instance Checking in $\mathcal{EL}$

Jia Tao, Giora Slutzki, and Vasant Honavar

Iowa State University, Ames, IA, USA

Abstract. We consider the problem of answering queries against a knowledge base (KB) using secrets, whenever it is possible to do so without compromising secrets. We study query answering against $\mathcal{EL}$ knowledge bases. We provide a polynomial time algorithm that, given an $\mathcal{EL}$ KB Σ, a set $\mathbb{S}$ of secrets to be protected and a query q or the form $C(a)$ or $r(a, b)$, outputs "Yes" whenever $\Sigma \models q$ and the answer to q, together with the answers to any previous queries answered by the KB, does not allow the querying agent to deduce any of the secrets in $\mathbb{S}$. This approach allows more flexible information sharing than is possible with traditional access control mechanisms.

1 Introduction

The rapid expansion of the World Wide Web and the widespread use of distributed databases and networked information systems offer unprecedented opportunities for productive interaction and collaboration among individuals and organizations in virtually every area of human endeavor. However, the need to share information has to be balanced against the need to protect sensitive information. The following example illustrates one such scenario.

Example 1. **(Healthcare)** (a simplified version adapted from [15]): Suppose that Jane*'s* mother Jill had breast cancer. Dr. James, Jane's physician, who is aware of Jane's family history, concludes that Jane has a significant risk of developing breast cancer. He asks her to undergo genetic screening for BRCA1 mutation (which is linked to an increased risk of breast cancer) to determine the extent to which Jane is at risk of developing breast cancer. Suppose Jane tests positive for BRCA1 mutation. Dr. James proceeds to prescribe her a certain drug that he knows is effective at reducing the breast cancer risk for patients with BRCA1 mutation. Jane purchases the medications from her pharmacy and wants to get reimbursed for the cost of her prescription by her insurance company. If her insurance company finds out that she has tested positive for BRCA1 mutation or that she has been prescribed certain drug(s) for breast cancer, Jane risks losing her health insurance. In this setting, the knowledge base (KB) needs to be able to certify to the insurance company that Jane qualifies for reimbursement for a drug that is covered by her insurance policy without revealing the fact that she is on such drugs. ■

The preceding example illustrates the need for algorithms that can, given a knowledge base Σ and a set $\mathbb{S}$ of secrets (perhaps specified using some secrecy policy[1]), answer queries against Σ, using secrets if necessary, whenever it is possible to do so without compromising their confidentiality. Barring a few exceptions (see Section 5), most existing approaches to information protection simply *forbid* the use of secret information in answering queries. The *privacy-preserving reasoning* framework introduced in [5] was motivated by the need to alleviate, at least in part, this important limitation of current methods for information sharing

[1] Upon choosing an underlying language to express the information in the KB, a mechanism is needed to transform the secrecy policies into secrets expressed by the chosen language. Such transformation is out of the scope of this paper.

(or conversely, information protection) in the simple setting of *hierarchical* knowledge bases (KBs) under the *open world assumption* (OWA)[2]. Such KBs may contain scientific, medical, economic information, or military intelligence, etc. Our secrecy-preserving reasoning framework builds on, and substantially extends, the privacy-preserving reasoning framework introduced by Bao et al. [5] (where the focus was on protecting some class-subclass relationships in hierarchical KBs).

In general, the answer to a query q against a KB Σ can be "Yes" (i.e., q can be inferred from Σ), "No" ($\neg q$ can be inferred from Σ) or "Unknown" (e.g., because of the incompleteness of Σ). We assume cooperative as opposed to adversarial scenarios in which the KB does not *lie*. However, whenever truthfully answering a query risks compromising secrets in $\mathbb{S}$, the reasoner associated with the KB is allowed to hide the answer to the query by feigning ignorance, i.e., answering the query as "Unknown". Given a set of secrets which we call the *secrecy set* $\mathbb{S}$, it is clear that, to protect $\mathbb{S}$, answers to queries about secrets in $\mathbb{S}$ will be "Unknown". However, we will show that, in general, it is not sufficient to protect only $\mathbb{S}$ since truthful answers to certain queries (that are not in $\mathbb{S}$) may reveal some information in $\mathbb{S}$. Therefore, we must protect a superset of $\mathbb{S}$, which we call an *envelope* of $\mathbb{S}$, such that the querying agent who has no access to the envelope will not be able to deduce any information in $\mathbb{S}$.

In this paper, we investigate secrecy-preserving query answering with $\mathcal{EL}$ [4], which is one of the simplest DLs that is both computationally tractable [11, 2, 19] and practically useful [4, 22]. For example, the medical ontology Snomed ct [27] and large parts of the medical ontology Galen [24] can be expressed in $\mathcal{EL}$. We provide algorithms to answer queries against an $\mathcal{EL}$ KB that use, but not reveal, the information that is designated as secret. Because of the open world assumption and the fact that the language of $\mathcal{EL}$ does not include negation, the answer to a query can only be "Yes" or "Unknown".

To answer queries posed to the KB, we utilize a *secrecy maintenance system* that consists of: a finite set of consequences of the KB Σ, denoted by $\mathcal{A}^*$, and a secrecy envelope $\mathbb{S} \subseteq \mathbb{E}_{\mathbb{S}} \subseteq \mathcal{A}^*$. The answer to a query q is censored by the reasoner if $q \in \mathbb{E}_{\mathbb{S}}$.

It is easy to see that a secrecy envelope always exists. For instance, $\mathcal{A}^*$ constitutes an envelope for any secrecy set $\mathbb{S} \subseteq \mathcal{A}^*$. A key challenge is to *develop strategies that can be used by the KB to respond to queries as informatively as possible (i.e., using an envelope that is as small as possible) without compromising secrets that the KB is obliged to protect.* Unfortunately, computing a minimum envelope is NP-hard.

We compute $\mathcal{A}^*$ using the (usual) tableau expansion rules. To compute $\mathbb{E}_{\mathbb{S}}$, we introduce the following idea. From each original expansion rule, we construct a corresponding *inverse expansion rule*. We show that the inverted system of expansion rules generates an envelope of $\mathbb{S}$. To the best of our knowledge, the idea of constructing a secrecy envelope by inverting the tableau expansion rules is novel. Furthermore, we introduce a couple of useful optimizations that help reduce the size of an envelope.

The rest of the paper is organized as follows. Section 2 introduces the secrecy-preserving framework. Section 3 initializes the secrecy maintenance system. We provide a tableau algorithm for computing the consequences of the KB and two tableau algorithms for computing secrecy envelopes. Section 4 discusses how to retrieve answers to queries. Section 5 concludes with a summary, a discussion of related work, and an outline of some directions for further research.

[2] Under the closed world assumption a statement that cannot be inferred from the KB to be true, is presumed to be false. Under the open world assumption, the status of a statement that cannot be inferred from the KB is presumed to be unknown, *not necessarily* false.

2 Preliminaries

2.1 Syntax and Semantics

The non-logical signature of the $\mathcal{EL}$ description language includes three mutually disjoint sets: a set of *concept names* $N_{\mathcal{C}}$, a set of *role names* $N_{\mathcal{R}}$ and a set of *individual names* $N_{\mathcal{O}}$. The syntax of $\mathcal{EL}$ is defined by specifying *expressions* and *formulae*. $\mathcal{EL}$ expressions consist of the set of *role names* $N_{\mathcal{R}}$ and the set of *concepts* $\mathcal{C}$ which is recursively defined as follows:

$$C, D \longrightarrow A \mid \top \mid C \sqcap D \mid \exists r.C$$

where $A \in N_{\mathcal{C}}$, $\top$ is the *top symbol*, $C, D \in \mathcal{C}$ and $r \in N_{\mathcal{R}}$. In this paper we will consider three kinds of $\mathcal{EL}$ formulae: *assertions* of the form $C(a)$ or $r(a, b)$, *definitions* of the form $A \doteq D$ and *general concept inclusions (GCI)* of the form $C \sqsubseteq D$ where $a, b \in N_{\mathcal{O}}$, $C, D \in \mathcal{C}$, $r \in N_{\mathcal{R}}$ and $A \in N_{\mathcal{C}}$.

The semantics of $\mathcal{EL}$ is defined by using an *interpretation* $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ where Δ is a non-empty domain and $\cdot^{\mathcal{I}}$ is a function that maps each individual name to an element in Δ, each concept name to a subset of Δ and each role name to a subset of $\Delta \times \Delta$. The interpretation of concept expressions is extended recursively as follows: for all $r \in N_{\mathcal{R}}$ and $C, D \in \mathcal{C}$, $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, $(\exists r.C)^{\mathcal{I}} = \{a \in \Delta \mid \exists b \in \Delta : (a, b) \in r^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\}$. For a finite set of symbols $N \subset N_{\mathcal{C}} \cup N_{\mathcal{R}} \cup N_{\mathcal{O}}$, we define an *interpretation $\mathcal{I}$ restricted to N* to be $\mathcal{I}_N = \langle \Delta, \cdot^{\mathcal{I}}|_N \rangle$.

2.2 Knowledge Bases

A finite non-empty set of assertions is called an *ABox*. A finite set of definitions and GCIs is called a *TBox*. An ABox $\mathcal{A}$ and a TBox $\mathcal{T}$ whose concepts and roles belong to the language $\mathcal{EL}$ form an $\mathcal{EL}$-knowledge base $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$. A TBox $\mathcal{T}$ is *normalized* [10] if $\mathcal{T}$ contains only GCIs all of which are of one of the following forms: $A \sqsubseteq B$, $A_1 \sqcap A_2 \sqsubseteq B$, $A \sqsubseteq \exists r.B$ or $\exists r.A \sqsubseteq B$ where $A, A_1, A_2, B \in N_{\mathcal{C}} \cup \{\top\}$. It was shown that transforming a TBox into such a normal form can be accomplished in polynomial time [10]. From now on, we will assume that all the TBoxes are in normal form.

Definition 1. *Let $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ be a knowledge base, $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ an interpretation, $C, D \in \mathcal{C}$, $r \in N_{\mathcal{R}}$ and $a, b \in N_{\mathcal{O}}$. $\mathcal{I}$* satisfies *$C(a)$, $r(a, b)$, or $C \sqsubseteq D$ if, respectively, $a^{\mathcal{I}} \in C^{\mathcal{I}}$, $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$, or $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. $\mathcal{I}$ is a* model *of Σ if it satisfies all the assertions in $\mathcal{A}$ and all the GCIs in $\mathcal{T}$. Let α be an assertion or a GCI. We say that Σ* entails *α, written as $\Sigma \vDash \alpha$, if all models of Σ satisfy α.*

We denote by N_Σ all the symbols appearing in Σ and by $\mathcal{O}_\Sigma$ the set of individual names appearing in Σ. Note that $\mathcal{O}_\Sigma \subset N_{\mathcal{O}} \cap N_\Sigma$ and $N_\Sigma \setminus \mathcal{O}_\Sigma \subset N_{\mathcal{C}} \cup N_{\mathcal{R}}$.

2.3 Motivating Example

Example 2. (Example 1, continued.) Let $\Sigma_1 = \langle \mathcal{A}_1, \mathcal{T}_1 \rangle$ be a KB that contains information on the patients, their health history, the prescriptions that they get from the physicians and their insurance information. The scenario described in Example 1 can be more formally specified in the description logic $\mathcal{EL}$ as follows:

1. $\exists$is_child.A $\sqsubseteq$ CancerRisk
2. HasMutBRCA1 $\sqsubseteq$ $\exists$has_pres.CancerDrug
3. $\exists$has_pres.CancerDrug $\sqsubseteq$ CancerRisk

4. $\exists$has_pres.CoveredDrug $\sqsubseteq$ Reimburse
5. CancerDrug $\sqsubseteq$ CoveredDrug
6. A $\sqsubseteq$ Woman
7. A $\sqsubseteq$ HasCancer
8. Woman$\sqcap$HasCancer $\sqsubseteq$ A
9. Woman(Jill)
10. HasCancer(Jill)
11. is_child(Jane, Jill)
12. HasMutBRCA1(Jane)

The GCIs 1-8 form a subset of $\mathcal{T}_1$ (in normal form) and the assertions 9-12 form a subset of $\mathcal{A}_1$. In order for Jane to get reimbursed, when the query Reimburse(Jane) is posed to the KB, the answer should be "Yes". However, in order to protect Jane's privacy, the query CancerRisk(Jane) should be answered "Unknown". ■

2.4 The Secrecy-preserving Query Answering Problem

Given a knowledge base Σ and a finite secrecy set $\mathbb{S}$, the basic goal is to answer queries while preserving secrecy. It is obvious that protecting only secrets in $\mathbb{S}$ is not enough. As shown in Example 2, to protect Jane's privacy, the query CancerRisk(Jane) should be answered "Unknown". However, by only keeping CancerRisk(Jane) secret, the fact that Jane has cancer risk can still be inferred by statements 12, 2 and 3. Therefore, the secrecy-preserving query answering problem is to find a superset of $\mathbb{S}$, which we call the *secrecy envelope* of $\mathbb{S}$, denoted by $\mathbb{E}_{\mathbb{S}}$, so that by protecting $\mathbb{E}_{\mathbb{S}}$, the querying agent cannot conclude anything in $\mathbb{S}$.

2.5 The Secrecy-preserving Query Answering Framework

The secrecy-preserving query answering framework is based on the OWA. Under OWA, what cannot be inferred is considered unknown, rather than false. To protect confidential information, queries that relate to secrets may be answered as "Unknown". The idea is that when the answer to a query is "Unknown", the querying agent is not able to distinguish between (a) the answer to the query is truly unknown, or (b) the answer is being protected for reasons of secrecy.

Our goal is to provide a decision procedure for answering queries while preserving secrecy. The framework contains following components and it is illustrated in Fig. 1:

- A knowledge base $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$.
- A secrecy set $\mathbb{S}$. We assume that a secrecy set $\mathbb{S}$ is given as a finite set of assertions that contain only symbols from N_{Σ}.
- Associated with Σ, there is a reasoner $\mathfrak{R}$ that is complete. $\mathfrak{R}$ is used to answer queries by checking whether the query can be inferred from Σ and if it can, whether answering "Yes" will reveal secrets from $\mathbb{S}$. The specific tasks are:
  - To compute the set SubC of sub-expressions of all concepts and roles appearing in Σ or $\mathbb{S}$.
  - To compute the set of all assertional consequences of Σ restricted to SubC. This set is called the *assertional closure of* Σ and it is denoted by $\mathcal{A}^*$. We assume that $\mathbb{S} \subseteq \mathcal{A}^*$.
  - To compute the secrecy envelope $\mathbb{E}_{\mathbb{S}} \subseteq \mathcal{A}^*$, a set of assertions that is a superset of $\mathbb{S}$, which if truthfully answered, may reveal some secret(s) in $\mathbb{S}$.

- To answer queries. If a query cannot be inferred from the knowledge base, the answer to the query is simply "Unknown". If it can be inferred and it is not in $\mathbb{E}_{\mathbb{S}}$, the final answer is "Yes"; otherwise, the answer is "Unknown".

– A querying agent who asks queries of the form $C(a)$ or $r(a, b)$. We assume that the querying agent has computational access only to the signature of the knowledge base, i.e., its queries are over N_{Σ}. We also assume that the querying agent has the same reasoning capacity as $\mathfrak{R}$. Since we assume that $\mathfrak{R}$ is complete, this is not a restriction. The querying agent may log the history of all the answers to its queries and draw conclusions from it. Moreover, we assume that the querying agent has access to the TBox $\mathcal{T}$.

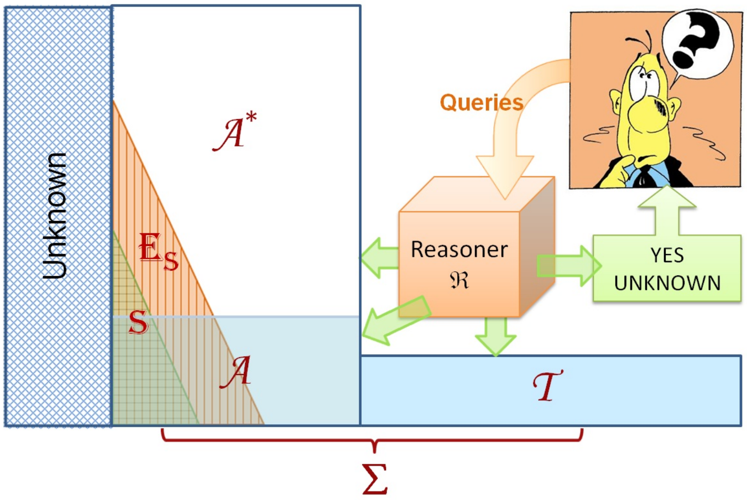

Fig. 1. The Secrecy-preserving Query-answering Framework

Example 3. (Example 2, cont.) For the given KB $\Sigma_1 = \langle \mathcal{A}_1, \mathcal{T}_1 \rangle$, consider the secrecy set $\mathbb{S}_1 = \{$CancerRisk(Jane)$\}$. Here, the querying agent is the insurance company and the queries include CancerRisk(Jane) and Reimburse(Jane). Because CancerRisk(Jane) can be inferred from statements 12, 2 and 3 in $\mathcal{A}_1$, at least one of these assertions, e.g., statement 12, should be put into the envelope. Thus, HasMutBRCA1(Jane)$\in \mathbb{E}_{\mathbb{S}1}$. ■

2.6 Secrecy Maintenance System

Before any query is posed to Σ, we compute a set of sub-expressions of all the concepts and roles appearing in Σ and $\mathbb{S}$, denoted by $Sub\mathcal{C}$. We initialize $\mathcal{A}^*$, the assertional closure of Σ, and then compute the secrecy envelope $\mathbb{E}_{\mathbb{S}}$, both restricted to $Sub\mathcal{C}$. $\mathcal{A}^*$ and $\mathbb{E}_{\mathbb{S}}$ form a *secrecy maintenance system.*

Once $\mathcal{A}^*$ and $\mathbb{E}_{\mathbb{S}}$ have been computed, if $C \in Sub\mathcal{C}$, we can answer the query $C(a)$ in linear time depending on its membership of $\mathcal{A}^*$ and $\mathbb{E}_{\mathbb{S}}$. Otherwise, we need to expand $Sub\mathcal{C}$ by adding sub-expressions of C that are not in $Sub\mathcal{C}$ and update the consequences $\mathcal{A}^*$ as well as $\mathbb{E}_{\mathbb{S}}$ accordingly.

3 Initializing Secrecy Maintenance System

In this section, we discuss the initialization of the secrecy maintenance system in detail.

3.1 Computing $\boldsymbol{\mathit{Sub}\mathcal{C}}$

$\mathit{Sub}\mathcal{C}$, the set of certain sub-expressions of all the concepts and roles appearing in Σ or $\mathbb{S}$, is defined as follows:

- if $C(a) \in \mathcal{A} \cup \mathbb{S}$, then $C \in \mathit{Sub}\mathcal{C}$;
- if $C \sqsubseteq D \in \mathcal{T}$, then $\{C, D\} \subseteq \mathit{Sub}\mathcal{C}$;
- if $r(a, b) \in \mathcal{A} \cup \mathbb{S}$, then $r \in \mathit{Sub}\mathcal{C}$;
- if $C_1 \sqcap \cdots \sqcap C_k \in \mathit{Sub}\mathcal{C}$ where either $C_i \in N_{\mathcal{C}}$ or $C_i = \exists r.C$, then $C_i \in \mathit{Sub}\mathcal{C}$ $(1 \leq i \leq k)$;
- if $\exists r.C \in \mathit{Sub}\mathcal{C}$, then $\{r, C\} \subseteq \mathit{Sub}\mathcal{C}$;
- if $\exists r.C \in \mathit{Sub}\mathcal{C}$ and $C \sqsubseteq D \in \mathcal{T}$ or $D \sqsubseteq C \in \mathcal{T}$, then $\exists r.D \in \mathit{Sub}\mathcal{C}$.

Note that $\mathit{Sub}\mathcal{C}$ does not contain all the sub-expressions of concepts appearing in Σ or $\mathbb{S}$. For example, if $C_1 \sqcap C_2 \sqcap C_3(a) \in \mathcal{A}$, then $\{C_1, C_2, C_3, C_1 \sqcap C_2 \sqcap C_3\} \subseteq \mathit{Sub}\mathcal{C}$. However, $C_1 \sqcap C_2 \notin \mathit{Sub}\mathcal{C}$ unless there is an assertion $C_1 \sqcap C_2(a) \in \mathcal{A} \cup \mathbb{S}$. If a query $C(a)$ comes along where $C \notin \mathit{Sub}\mathcal{C}$, it will be added into $\mathit{Sub}\mathcal{C}$. As such, the secrecy maintenance system is built up gradually depending on the history of queries. Also note that the initial size of $\mathit{Sub}\mathcal{C}$ is linear in the size of the knowledge base Σ plus the size of the secrecy set $\mathbb{S}$.

3.2 Computing $\boldsymbol{\mathcal{A}^*}$

The ABox $\mathcal{A}^*$ is initialized as $\mathcal{A}$ and expanded by recursively applying assertion expansion rules listed in Fig. 2. We say that $\mathcal{A}^*$ is *assertionally closed* or that it is an *assertional closure of* Σ if no assertion expansion rule is applicable. The set of all the individual names appearing in $\mathcal{A}^*$ is denoted by $\mathcal{O}^*$. It is initialized as $\mathcal{O}_\Sigma$ and is expanded with applications of the $\exists_2^{\mathcal{A}}$-rule. An individual a is said to be *fresh* (at a particular time during the expansion process) if $a \in N_{\mathcal{O}} \setminus \mathcal{O}^*$ (at that time). An individual $a \in \mathcal{O}^*$ is *blocked* by an individual $b \in \mathcal{O}^*$ if $a \in \mathcal{O}^* \setminus \mathcal{O}_\Sigma$, b is either in $\mathcal{O}_\Sigma$ or b was picked earlier than a (during the expansion process), and $\{C \mid C(a) \in \mathcal{A}^*\} \subseteq \{C' \mid C'(b) \in \mathcal{A}^*\}$. Recall that we have assumed that the querying agent has computational access only to the signature of the knowledge base. In particular, the querying agent cannot ask any queries that involve individual names in $\mathcal{O}^* \setminus \mathcal{O}_\Sigma$. This is referred to as *Hidden Names Assumption* (HNA).

We denote by Λ the tableau algorithm which nondeterministically applies assertion expansion rules until no further applications are possible. Since each expansion rule can be applied polynomially many times (in the size of $\mathit{Sub}\mathcal{C}$), the computation of $\mathcal{A}^*$ can be done in polynomial time. When an execution of Λ terminates, we have an assertionally closed ABox $\mathcal{A}^*$. Note that different executions of Λ may result different $\mathcal{A}^*$'s which differ only in the individual names that have been freshly chosen during the executions of Λ.

Let $\mathcal{I}^1 = \langle \Delta, \cdot^{\mathcal{I}^1} \rangle$, $\mathcal{I}^2 = \langle \Delta, \cdot^{\mathcal{I}^2} \rangle$ be two interpretations and $N_2 \subset N_1$ be finite subsets of $N_{\mathcal{C}} \cup N_{\mathcal{R}} \cup N_{\mathcal{O}}$ such that $N_1 \setminus N_2 \subset N_{\mathcal{O}}$. A model $\mathcal{I}^1_{N_1} = \langle \Delta, \cdot^{\mathcal{I}^1}|_{N_1} \rangle$ is a *semantic extension* of a model $\mathcal{I}^2_{N_2} = \langle \Delta, \cdot^{\mathcal{I}^2}|_{N_2} \rangle$ if $(\mathcal{I}^1_{N_1})_{N_2} = \mathcal{I}^2_{N_2}$. The following theorem shows the soundness the tableau algorithm Λ.

Theorem 1. *(Soundness) Let* $\mathcal{A}^*$ *is an assertionally closed ABox obtained from* Σ *by applying* Λ. $\forall C \in \mathcal{C} \cap \mathit{Sub}\mathcal{C}$, $\forall a \in \mathcal{O}^*$, *if* $C(a) \in \mathcal{A}^*$, *then for every model* $\mathcal{I}_{N_\Sigma}$ *of* Σ, *there is a semantic extension of* $\mathcal{I}_{N_\Sigma}$ *that satisfies* $C(a)$.

$\sqcap_1^{\mathcal{A}}$ -rule: if $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathcal{A}^*$ and $C_i(a) \notin \mathcal{A}^*$,
then $\mathcal{A}^* := \mathcal{A}^* \cup \{C_i(a)\}$ where $1 \leq i \leq k$;
$\sqcap_2^{\mathcal{A}}$ -rule: if $\{C_1(a), ..., C_k(a)\} \subseteq \mathcal{A}^*, C_1 \sqcap \cdots \sqcap C_k \in \mathit{Sub}\mathcal{C}$
and $C_1 \sqcap \cdots \sqcap C_k(a) \notin \mathcal{A}^*$, then $\mathcal{A}^* := \mathcal{A}^* \cup \{C_1 \sqcap \cdots \sqcap C_k(a)\}$;
$\exists_1^{\mathcal{A}}$ -rule: if $\{r(a,b), C(b)\} \subseteq \mathcal{A}^*, \exists r.C \in \mathit{Sub}\mathcal{C}$ and $\exists r.C(a) \notin \mathcal{A}^*$,
then $\mathcal{A}^* := \mathcal{A}^* \cup \{\exists r.C(a)\}$;
$\exists_2^{\mathcal{A}}$ -rule: if $\exists r.C(a) \in \mathcal{A}^*, a$ is not blocked and for all $b \in \mathcal{O}^*, \{r(a,b), C(b)\} \nsubseteq \mathcal{A}^*$,
then $\mathcal{A}^* := \mathcal{A}^* \cup \{r(a,c), C(c)\}$ where c is fresh, and $\mathcal{O}^* := \mathcal{O}^* \cup \{c\}$;
$\sqsubseteq^{\mathcal{T}}$ -rule: if $C(a) \in \mathcal{A}^*, C \sqsubseteq D \in \mathcal{T}$ and $D(a) \notin \mathcal{A}^*$, then $\mathcal{A}^* := \mathcal{A}^* \cup \{D(a)\}$;

Fig. 2. Assertion Expansion Rules

Proof. Let $\mathcal{I}_{N_\Sigma} = \langle \Delta, \cdot^{\mathcal{I}}|_{N_\Sigma} \rangle$ be an arbitrary model of Σ. We need to show that after applying each expansion rule, there is a semantic extension of $\mathcal{I}_{N_\Sigma}$ that satisfies new assertion(s) being added to $\mathcal{A}^*$. We prove it by induction on the construction of $\mathcal{A}^*$. The base case is when $C(a) \in \mathcal{A}$. Since $\mathcal{I}_{N_\Sigma}$ is a model of Σ, $a^{\mathcal{I}} \in C^{\mathcal{I}}$. For the induction step, we use $\mathcal{A}'$, $\mathcal{O}'$ and $\mathcal{I}' = \langle \Delta, \cdot^{\mathcal{I}'} \rangle$ to denote the ABox before the application of each expansion rule, the set of individual names appearing in $\mathcal{A}'$, and the model of $\langle \mathcal{A}', \mathcal{T} \rangle$, respectively. We also denote by $\mathcal{A}''$ the ABox after the application of each expansion rule and by $\mathcal{O}''$ the set of individual names appearing in $\mathcal{A}''$. Note that except for the case 4, $\mathcal{O}'' = \mathcal{O}'$.

1. If $\sqcap_1^{\mathcal{A}}$-rule is applicable, then there is an assertion $C_1 \sqcap \cdots \sqcap C_k(a)$ in $\mathcal{A}'$ and an integer i $(1 \leq i \leq k)$ such that $C_i(a) \notin \mathcal{A}'$. After applying the rule, $C_i(a) \in \mathcal{A}''$. By IH, $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathcal{A}'$ implies that $a^{\mathcal{I}'} \in (C_1 \sqcap \cdots \sqcap C_k)^{\mathcal{I}'} = C_1^{\mathcal{I}'} \cap \cdots \cap C_k^{\mathcal{I}'}$. It follows that $a^{\mathcal{I}'} \in C_i^{\mathcal{I}'}$. Therefore, $\mathcal{I}'$ satisfies the newly added assertion.
2. If $\sqcap_2^{\mathcal{A}}$-rule is applicable, then $\{C_1(a), ..., C_k(a)\} \subseteq \mathcal{A}'$, $C_1 \sqcap \cdots \sqcap C_k(a) \notin \mathcal{A}'$ and $C_1 \sqcap \cdots \sqcap C_k \in \mathit{Sub}\mathcal{C}$. After applying the rule, $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathcal{A}''$. By IH, $\{C_1(a), ..., C_k(a)\} \subseteq \mathcal{A}'$ implies that $a^{\mathcal{I}'} \in C_i^{\mathcal{I}'} (1 \leq i \leq k)$ which is equivalent to $a^{\mathcal{I}'} \in C_1^{\mathcal{I}'} \cap \cdots \cap C_k^{\mathcal{I}'} = (C_1 \sqcap \cdots \sqcap C_k)^{\mathcal{I}'}$. It follows that $a^{\mathcal{I}'} \in (C_1 \sqcap \cdots \sqcap C_k)^{\mathcal{I}'}$ and hence $\mathcal{I}'$ satisfies the newly added assertion $C_1 \sqcap \cdots \sqcap C_k(a)$.
3. If $\exists_1^{\mathcal{A}}$-rule is applicable, then $\{r(a,b), C(b)\} \subseteq \mathcal{A}'$ and $\exists r.C(a) \in \mathit{Sub}\mathcal{C}$ and $\exists r.C(a) \notin \mathcal{A}'$. After applying the rule, $\exists r.C(a) \in \mathcal{A}''$. By IH, $\{r(a,b), C(b)\} \subseteq \mathcal{A}'$ implies that $(a^{\mathcal{I}'}, b^{\mathcal{I}'}) \in r^{\mathcal{I}'}$ and $b^{\mathcal{I}'} \in C^{\mathcal{I}'}$. It follows that $a^{\mathcal{I}'} \in (\exists r.C)^{\mathcal{I}'}$. So, $\mathcal{I}'$ satisfies the newly added assertion $\exists r.C(a)$.
4. If $\exists_2^{\mathcal{A}}$-rule is applicable, then $\exists r.C(a) \in \mathcal{A}'$, a is not blocked and for all $b \in \mathcal{O}'$, $\{r(a,b),\ C(b)\} \nsubseteq \mathcal{A}'$. In applying the rule, a fresh individual name c is picked and after the application, $\{r(a,c), C(c)\} \in \mathcal{A}''$. By IH, $\exists r.C(a) \in \mathcal{A}'$ implies that there is $d \in \Delta$ such that $(a^{\mathcal{I}'}, d) \in r^{\mathcal{I}'}$ and $d \in C^{\mathcal{I}'}$. We define an interpretation $\mathcal{J}$ such that $c^{\mathcal{J}} = d$ and $\mathcal{J}_{N_\Sigma \cup \mathcal{O}'} = \mathcal{I}'_{N_\Sigma \cup \mathcal{O}'}$. It is obvious that $\mathcal{J}_{N_\Sigma \cup \mathcal{O}''}$ is a semantic extension of $\mathcal{I}'{}_{N_\Sigma \cup \mathcal{O}'}$ and we have $(a^{\mathcal{J}}, c^{\mathcal{J}}) \in r^{\mathcal{J}}$ and $c^{\mathcal{J}} \in C^{\mathcal{J}}$.
5. If $\sqsubseteq^{\mathcal{T}}$-rule is applicable, then $C(a) \in \mathcal{A}'$, $C \sqsubseteq D \in \mathcal{T}$ and $D(a) \notin \mathcal{A}'$. After applying the rule, $D(a) \in \mathcal{A}''$. By IH, $C(a) \in \mathcal{A}'$ implies that $a^{\mathcal{I}'} \in C^{\mathcal{I}'}$. $C \sqsubseteq D \in \mathcal{T}$ implies that $C^{\mathcal{I}'} \subseteq D^{\mathcal{I}'}$. It follows that $a^{\mathcal{I}'} \in D^{\mathcal{I}'}$, meaning that the newly added assertion $D(a)$ is satisfied by $\mathcal{I}'$. ■

To prove the completeness of Λ, we define a *canonical interpretation* $\mathcal{J}_{\mathcal{A}^*} = \langle \Delta, \cdot^{\mathcal{J}} \rangle$ for the assertionally closed ABox $\mathcal{A}^*$ as follows:

- $\Delta := \mathcal{O}^*$;
- $a^{\mathcal{J}} := a$, for each $a \in \Delta$;

– $A^{\mathcal{J}} := \{a \mid A(a) \in \mathcal{A}^*\}$ where $A \in N_{\mathcal{C}} \cap Sub\mathcal{C}$;
– $r^{\mathcal{J}} := \{(a,b) \mid r(a,b) \in \mathcal{A}^*\} \cup \{(c,b) \mid a$ blocks c and $r(a,b) \in \mathcal{A}^*\}$ where $r \in N_{\mathcal{R}} \cap Sub\mathcal{C}$.

The following lemma shows that $\mathcal{J}_{\mathcal{A}^*}$ is a model of $\mathcal{A}^*$.

Lemma 1. *For each $C \in \mathcal{C} \cap Sub\mathcal{C}$ and each $a \in \mathcal{O}^*$, $C(a) \in \mathcal{A}^* \Leftrightarrow \mathcal{J}_{\mathcal{A}^*} \models C(a)$.*

Proof. ($\Longrightarrow$) Assume that $C(a) \in \mathcal{A}^*$. We argue by induction on the structure of C. The base case is when $C \in N_{\mathcal{C}} \cap Sub\mathcal{C}$. By the definition of $\mathcal{J}_{\mathcal{A}^*}$, $a^{\mathcal{J}} \in C^{\mathcal{J}}$ and hence $\mathcal{J}_{\mathcal{A}^*} \models C(a)$.

If $C = C_1 \sqcap \cdots \sqcap C_k$, then since $\mathcal{A}^*$ is assertionally closed, $\{C_1(a), ..., C_k(a)\} \subseteq \mathcal{A}^*$ due to the $\sqcap_1^{\mathcal{A}}$-rule. By IH, $C_i(a) \in \mathcal{A}^* \Rightarrow \mathcal{J}_{\mathcal{A}^*} \models C_i(a)$, $1 \leq i \leq k$. It follows that $a^{\mathcal{J}} \in C_i^{\mathcal{J}}$. Hence $a^{\mathcal{J}} \in (C_1 \sqcap \cdots \sqcap C_i)^{\mathcal{J}} = C^{\mathcal{J}}$. Therefore, $\mathcal{J}_{\mathcal{A}^*} \models C(a)$.

If $C = \exists r.C_1$, there are two cases:

– If $\exists r.C_1(a)$ has a witness $b \in \mathcal{O}^*$ such that $\{r(a,b), C_1(b)\} \subseteq \mathcal{A}^*$, then by definition of $\mathcal{J}_{\mathcal{A}^*}$, $(a^{\mathcal{J}}, b^{\mathcal{J}}) \in r^{\mathcal{J}}$. By IH, $b^{\mathcal{J}} \in C_1^{\mathcal{J}}$. It follows that $a^{\mathcal{J}} \in (\exists r.C_1)^{\mathcal{J}} = C^{\mathcal{J}}$. Therefore, $\mathcal{J}_{\mathcal{A}^*} \models C(a)$.
– If $\exists r.C_1(a)$ does not have a witness $b \in \mathcal{O}^*$ such that $\{r(a,b), C_1(b)\} \subseteq \mathcal{A}^*$, then there must exist an individual c that blocks a where $\{r(c,d), C_1(d), \exists r.C_1(c)\} \subseteq \mathcal{A}^*$. By definition of $\mathcal{J}_{\mathcal{A}^*}$, $(a^{\mathcal{J}}, d^{\mathcal{J}}) \in r^{\mathcal{J}}$. By IH, $d^{\mathcal{J}} \in C_1^{\mathcal{J}}$. It follows that $a^{\mathcal{J}} \in (\exists r.C_1)^{\mathcal{J}} = C^{\mathcal{J}}$. Therefore, $\mathcal{J}_{\mathcal{A}^*} \models C(a)$.

($\Longleftarrow$) We need to show that for each $C \in \mathcal{C} \cap Sub\mathcal{C}$ and each $a \in \mathcal{O}^*$, if $C(a) \notin \mathcal{A}^*$, then $\mathcal{J}_{\mathcal{A}^*} \not\models C(a)$. The base case is when $C \in N_{\mathcal{C}} \cap Sub\mathcal{C}$. If $C(a) \notin \mathcal{A}^*$, by the definition of $\mathcal{J}_{\mathcal{A}^*}$, $a^{\mathcal{J}} \notin C^{\mathcal{J}}$. Therefore, $\mathcal{J}_{\mathcal{A}^*} \not\models C(a)$.

If $C = C_1 \sqcap \cdots \sqcap C_k$, then since $\mathcal{A}^*$ is assertionally closed, $C(a) \notin \mathcal{A}^*$ implies $\{C_1(a), ..., C_k(a)\} \nsubseteq \mathcal{A}^*$ due to the $\sqcap_2^{\mathcal{A}}$-rule. So there is a C_i such that $C_i(a) \notin \mathcal{A}^*$, $1 \leq i \leq k$. By IH, $\mathcal{J}_{\mathcal{A}^*} \not\models C_i(a)$. It follows that $a^{\mathcal{J}} \notin C_i^{\mathcal{J}}$. Hence $a^{\mathcal{J}} \notin C_1^{\mathcal{J}} \cap \cdots \cap C_k^{\mathcal{J}} = (C_1 \sqcap \cdots \sqcap C_k)^{\mathcal{J}} = C^{\mathcal{J}}$. Therefore, $\mathcal{J}_{\mathcal{A}^*} \not\models C(a)$.

If $C = \exists r.C_1$, then since $\mathcal{A}^*$ is assertionally closed, $C(a) \notin \mathcal{A}^*$ implies that there does not exist $b \in \mathcal{O}^*$ such that $\{r(a,b), C_1(b)\} \subseteq \mathcal{A}^*$ due to the $\exists_1^{\mathcal{A}}$-rule. By the definition of $r^{\mathcal{J}}$ and IH, for each $b \in \mathcal{O}^*$, either $(a^{\mathcal{J}}, b^{\mathcal{J}}) \notin r^{\mathcal{J}}$ or $b^{\mathcal{J}} \notin C_1^{\mathcal{J}}$. It follows that $a^{\mathcal{J}} \notin (\exists r.C_1)^{\mathcal{J}} = C^{\mathcal{J}}$. Therefore, $\mathcal{J}_{\mathcal{A}^*} \not\models C(a)$. ■

Corollary 1. *In a canonical interpretation $\mathcal{J}_{\mathcal{A}^*}$, for each $C \in \mathcal{C} \cap Sub\mathcal{C}$, $C^{\mathcal{J}} = \{b \in \mathcal{O}^* \mid C(b) \in \mathcal{A}^*\}$.*

Next lemma shows that $\mathcal{J}_{\mathcal{A}^*}$ is also a model of the TBox $\mathcal{T}$.

Lemma 2. *For $C, D \in \mathcal{C} \cap Sub\mathcal{C}$, $C \sqsubseteq D \in \mathcal{T} \Rightarrow \mathcal{J}_{\mathcal{A}^*} \models C \sqsubseteq D$.*

Proof. The claim is an easy consequence of $\sqsubseteq^{\mathcal{T}}$-rule and Lemma 1. For all subsumption $C \sqsubseteq D \in \mathcal{T}$, $C(a) \in \mathcal{A}^* \Rightarrow D(a) \in \mathcal{A}^*$ by $\sqsubseteq^{\mathcal{T}}$-rule. By Corollary 1, $C^{\mathcal{J}} = \{b \in \mathcal{O}^* \mid C(b) \in \mathcal{A}^*\}$ and similarly for $D^{\mathcal{J}}$. It follows that $C^{\mathcal{J}} \subseteq D^{\mathcal{J}}$. ■

Theorem 2. *(Completeness) Let $Sub\mathcal{C}$ be the set of sub-expressions obtained from a KB $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a finite set of assertions S (see Sect. 3.1.). Let $\mathcal{A}^*$ is an assertionally closed ABox obtained from Σ by applying Λ. Then for all $C \in \mathcal{C} \cap Sub\mathcal{C}$ and all $a \in \mathcal{O}^*$, $\Sigma \models C(a) \Rightarrow C(a) \in \mathcal{A}^*$.*

Proof. Suppose that $C(a)$ holds in all models of Σ. By Lemma 1, the canonical interpretation $\mathcal{J}_{\mathcal{A}^*}$ is a model of $\mathcal{A}^*$ and hence of $\mathcal{A}$. By Lemma 2, $\mathcal{J}_{\mathcal{A}^*}$ is a model of $\mathcal{T}$. It follows that $C(a)$ holds in $\mathcal{J}_{\mathcal{A}^*}$. By Lemma 1, $C(a) \in \mathcal{A}^*$. ■

Example 4. Continuing Examples 1-3 with the KB $\Sigma_1 = \langle \mathcal{A}_1, \mathcal{T}_1 \rangle$ and secrecy set $\mathbb{S}_1 = \{$CancerRisk(Jane)$\}$, by applying Λ, we obtain the assertional closure of Σ_1, denoted by $\mathcal{A}_1^*$, as follows.
$\mathcal{A}_1^* = \mathcal{A}_1 \cup \{$ A(Jill), ∃is_child.A(Jane), CancerRisk(Jane), has_pres(Jane, a), ∃has_pres.CancerDrug(Jane), CancerDrug(a), ∃has_pres.CoveredDrug(Jane), CoveredDrug(a), Reimburse(Jane) $\}$. ■

Ignoring the issue of secrecy, we point out a difference between the reasoning of the KB reasoner $\mathfrak{R}$ and that of the querying agent. Consider the assertion $\exists r.C(a) \in \mathcal{A}^*$ when a is not blocked and $\nexists b \in \mathcal{O}_\Sigma$ for which $\{r(a,b), C(b)\} \subseteq \mathcal{A}^*$. In this case $\mathfrak{R}$ picks a fresh individual name $c \notin \mathcal{O}_\Sigma$ as a witness for the inclusion $\exists r.C(a) \in \mathcal{A}^*$. The querying agent only knows the existence of the witness individual without knowing the individual name itself. Of course, for its own reasoning process, the querying agent may pick any individual name in $N_\mathcal{O} \setminus \mathcal{O}_\Sigma$, say d, and then force $r(a,d)$ and $C(d)$ to be consequences of Σ. Clearly, the reasoner $\mathfrak{R}$ and the querying agent are not aware of each other's "fresh" individual names. To differentiate the assertional closure of the KB reasoner $\mathfrak{R}$ from the reasoning of the querying agent, we will use $\cdot^+$ to denote the latter.

3.3 Secrecy in KBs

Our basic approach to designing secrecy-preserving reasoners for knowledge bases that contain sensitive knowledge is to answer "unknown" to every query whose secrecy must be protected or from which secret information could be deduced. Because of OWA, a querying agent cannot distinguish between an answer "Unknown" that results from the reasoner's incomplete information and an "Unknown" resulting from the reasoner's need to protect secret information. Since we have assumed that the querying agent can only ask queries over the vocabulary N_Σ, the information the reasoner needs to protect against the querying agent need not include assertions about individuals that are not in $\mathcal{O}_\Sigma$.

Definition 2. *Given a knowledge base $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a finite secrecy set $\mathbb{S} \subseteq \mathcal{A}^*$, a* secrecy envelope *of $\mathbb{S}$, denoted by $\mathbb{E}_\mathbb{S}$, is a superset of $\mathbb{S} \subseteq \mathbb{E}_\mathbb{S} \subseteq \mathcal{A}^*$ such that $(\mathcal{A}^* \setminus \mathbb{E}_\mathbb{S})^+ \cap \mathbb{S} = \emptyset$ where $\mathcal{A}^*$ is the assertional closure of Σ for $\mathfrak{R}$ and $(\mathcal{A}^* \setminus \mathbb{E}_\mathbb{S})^+$ is the assertional closure of the knowledge base $\langle \mathcal{A}^* \setminus \mathbb{E}_\mathbb{S}, \mathcal{T} \rangle$ for the querying agent.*

The idea behind this definition is that if the reasoner $\mathfrak{R}$ answers every query in $\mathbb{E}_\mathbb{S}$ with "Unknown" and every query in $\mathcal{A}^* \setminus \mathbb{E}_\mathbb{S}$ with "Yes", the querying agent will not be able to deduce any assertions in $\mathbb{S}$. A secrecy envelope always exists. For example, $\mathcal{A}^*$ is one such envelope. However, in order to be as informative as possible, we aim to make $\mathbb{E}_\mathbb{S}$ as small as possible. Unfortunately, to compute a minimum envelope, i.e., an envelope with the smallest cardinality is hard. Specifically, the decision version of the problem of computing minimum envelopes is NP-complete.

An instance of the Minimum Secrecy Envelope (MSE) problem contains a triple $\langle \Sigma = \langle \mathcal{A}, \mathcal{T} \rangle, \mathbb{S}, K \rangle$ where Σ is a knowledge base, $\mathcal{A}^*$ is the set of consequences of Σ restricted to $Sub\mathcal{C}$, $\mathbb{S} \subseteq \mathcal{A}^*$ is a secrecy set, and $K \leq |\mathcal{A}^*|$ is a nonnegative integer. The question is "Is there a secrecy envelope $\mathbb{E}$ such that $\mathbb{S} \subseteq \mathbb{E} \subseteq \mathcal{A}^*$ and $|\mathbb{E} \setminus \mathbb{S}| \leq K$?"

Given a set of assertions $\mathbb{E}' \supseteq \mathbb{S}$, we can verify (a) whether $\mathbb{E}'$ is an envelope by recalculating $(\mathcal{A}^* \setminus \mathbb{E}')^+$ and checking that it contains no assertions in $\mathbb{S}$, and (b) whether $|\mathbb{E}' \setminus \mathbb{S}| \leq K$. Both tasks are doable in polynomial time.

To show that the MSE problem is NP-hard, we reduce the Hitting Set (HS) problem to the MSE problem. An instance of HS consists of a collection $\mathcal{M}$ of subsets of a finite set S and a positive integer $K \leq |S|$. The question is "Is there a subset $S' \subseteq S$ with $|S'| \leq K$ such that S' contains at least one element from each set in $\mathcal{M}$?" W.l.o.g., we may assume that every set in $\mathcal{M}$ has at least two elements.

Given an instance of HS, we construct an instance of MSE, using the same constant K, as follows:

- $N_\mathcal{O} = \{a\}$, $N_\mathcal{R} = \emptyset$, $N_\mathcal{C} = S$
- $\mathbb{S} = \{A_1 \sqcap \cdots \sqcap A_m(a) \mid \{A_1, ..., A_m\} \in \mathcal{M}\}$
- $\mathcal{A} = \{A(a) \mid A \in S\} \cup \mathbb{S}$, $\mathcal{T} = \emptyset$
- $\Sigma = \langle \mathcal{A}, \emptyset \rangle$

Claim. S has a hitting set S' that hits every subset in $\mathcal{M}$ with $|S'| \leq K$ iff there is a secrecy envelope $\mathbb{E}$ such that $\mathbb{S} \subseteq \mathbb{E} \subseteq \mathcal{A}^*$ and $|\mathbb{E} \setminus \mathbb{S}| \leq K$.

Proof. Suppose that S has a hitting set S' that hits every set in $\mathcal{M}$ and $|S'| \leq K$. Then for each set $\{A_1, ..., A_m\} \in \mathcal{M}$, $\exists A_j \in S'$ $(1 \leq j \leq m)$. Let $\mathbb{E} = S' \cup \mathbb{S}$. It follows that $\forall A_1 \sqcap \cdots \sqcap A_m(a) \in \mathbb{S}$, $\exists A_j(a) \in \mathbb{E}$ $(1 \leq j \leq m)$. By construction, Σ does not involve any roles and $Sub\mathcal{C} = \{C \mid C(a) \in \mathcal{A}\}$. Therefore $\mathcal{A}^* = \mathcal{A}$, and so $\mathcal{A}^* \setminus \mathbb{E} = \mathcal{A} \setminus \mathbb{E}$. Note that $\mathcal{A} \setminus \mathbb{E}$ contains only assertions of the form $C(a)$ where $C \in N_{\mathcal{C}}$. It follows that none of assertion expansion rules is applicable to $\mathcal{A} \setminus \mathbb{E}$, implying that $(\mathcal{A}^* \setminus \mathbb{E})^+ \cap \mathbb{S} = (\mathcal{A}^* \setminus \mathbb{E}) \cap \mathbb{S} = \emptyset$. It follows that $\mathbb{E}$ is an envelope with $|\mathbb{E} \setminus \mathbb{S}| \leq |S'| \leq K$.

Conversely, suppose that there is a secrecy envelope $\mathbb{E}$ such that $\mathbb{S} \subseteq \mathbb{E} \subseteq \mathcal{A}^*$ and $|\mathbb{E} \setminus \mathbb{S}| \leq K$. Then, by Definition 2, $\forall A_1 \sqcap \cdots \sqcap A_m(a) \in \mathbb{S}$, $\exists A_j(a) \in \mathbb{E}$ where $1 \leq j \leq m$. $\mathbb{E} \setminus \mathbb{S}$ contains only assertions of the form $C(a)$ where $C \in N_{\mathcal{C}}$. This shows that S contains a subset $S' = \{C \mid C(a) \in \mathbb{E} \setminus \mathbb{S}\}$ that hits every set in $\mathcal{M}$ and $|S'| = |\mathbb{E} \setminus \mathbb{S}| \leq K$. ■

In what follows, we provide an algorithm that computes envelopes. Based on this algorithm as well as the HNA, we further optimize the algorithm to result a smaller envelope. To compute an envelope, we introduce the idea of inverting assertion expansion rules. For $\mathcal{EL}$ with TBox, we have five assertion expansion rules as listed in Fig. 2. For each assertion expansion rule, the resulting inverse rule is named by changing the superscript in the name of the original rule to $\mathbb{S}$. These inversion rules are called $\mathfrak{R}$-*secrecy closure rules* and are listed in Fig. 3. In Fig. 3, $\mathcal{A}^*$ is the assertional closure of Σ that has been computed previously by applying Λ; $\mathbb{E}$ is initialized as the given secrecy set, and expanded by using $\mathfrak{R}$-secrecy closure rules.

$$\begin{array}{rl}
\sqcap_1^{\mathbb{S}}\text{ -rule:} & \text{if } C_1 \sqcap \cdots \sqcap C_k(a) \in \mathcal{A}^* \setminus \mathbb{E} \text{ and } \{C_1(a), ..., C_k(a)\} \cap \mathbb{E} \neq \emptyset, \\
 & \text{then } \mathbb{E} := \mathbb{E} \cup \{C_1 \sqcap \cdots \sqcap C_k(a)\}; \\
\sqcap_2^{\mathbb{S}}\text{ -rule:} & \text{if } C_1 \sqcap \cdots \sqcap C_k(a) \in \mathbb{E} \text{ and } \{C_1(a), ..., C_k(a)\} \cap \mathbb{E} = \emptyset, \\
 & \text{then } \mathbb{E} := \mathbb{E} \cup \{C_i(a)\} \text{ where } 1 \leq i \leq k; \\
\exists_1^{\mathbb{S}}\text{ -rule:} & \text{if } \exists r.C(a) \in \mathbb{E} \text{ and } \{r(a,b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E} \text{ with } b \in \mathcal{O}^*, \\
 & \text{then } \mathbb{E} := \mathbb{E} \cup \{r(a,b)\} \text{ or } \mathbb{E} := \mathbb{E} \cup \{C(b)\}; \\
\exists_2^{\mathbb{S}}\text{ -rule:} & \text{if } \exists r.C(a) \in \mathcal{A}^* \setminus \mathbb{E}, \text{ and for all } b \in \mathcal{O}^* \text{ with } \{r(a,b), C(b)\} \subseteq \mathcal{A}^*, \\
 & \text{we have } \{r(a,b), C(b)\} \cap \mathbb{E} \neq \emptyset, \text{ then } \mathbb{E} := \mathbb{E} \cup \{\exists r.C(a)\}; \\
\sqsubseteq^{\mathbb{S}}\text{ -rule:} & \text{if } D(a) \in \mathbb{E}, C \sqsubseteq D \in \mathcal{T} \text{ and } C(a) \in \mathcal{A}^* \setminus \mathbb{E}, \text{ then } \mathbb{E} := \mathbb{E} \cup \{C(a)\}.
\end{array}$$

Fig. 3. $\mathfrak{R}$-secrecy closure rules obtained by inverting rules in Fig. 2.

We denote by $\Lambda_{\mathbb{S}}^{\mathfrak{R}}$ the tableau algorithm which nondeterministically applies the $\mathfrak{R}$-secrecy closure rules until no further rules are applicable. When no $\mathfrak{R}$-secrecy closure rule is applicable, we say that $\mathbb{E}$ is $\mathfrak{R}$-*closed*. It is clear that all executions of $\Lambda_{\mathbb{S}}^{\mathfrak{R}}$ on an input consisting of the assertional closure of a KB and a finite set of assertions terminate and that when $\Lambda_{\mathbb{S}}^{\mathfrak{R}}$ terminates, $\mathbb{E}$ is $\mathfrak{R}$-closed. It is also easy to see that $\Lambda_{\mathbb{S}}^{\mathfrak{R}}$ takes polynomial time in the size of its input.

Lemma 3. *Given $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a finite secrecy set $\mathbb{S}$, let $\mathcal{A}^*$ be the assertional closure of Σ and suppose that $\mathbb{E}$ is $\mathfrak{R}$-closed. Then $\mathcal{A}^* \setminus \mathbb{E}$ is assertionally closed w.r.t. assertion expansion rules listed in Fig. 2.*

Proof. To prove that $\mathcal{A}^* \setminus \mathbb{E}$ is assertionally closed, we need to show that no assertion expansion rules are applicable to $\mathcal{A}^* \setminus \mathbb{E}$. We prove it by contradiction.

- If $\sqcap_1^{\mathcal{A}}$-rule is applicable, then there is an assertion $C_1 \sqcap \cdots \sqcap C_k(a)$ in $\mathcal{A}^* \setminus \mathbb{E}$ and there is $C_i(a) \notin \mathcal{A}^* \setminus \mathbb{E}$ $(1 \leq i \leq k)$. $\mathcal{A}^*$ being assertionally closed implies that $C_i(a) \in \mathcal{A}^*$ and so $C_i(a) \in \mathbb{E}$. However, since $\mathbb{E}$ is $\mathfrak{R}$-closed, due to the $\sqcap_1^{\mathbb{S}}$-rule, $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathbb{E}$, contradicting $C_1 \sqcap \cdots \sqcap C_k(a) \notin \mathbb{E}$.
- If $\sqcap_2^{\mathcal{A}}$-rule is applicable, then $\{C_1(a), ..., C_k(a)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$, $C_1 \sqcap \cdots \sqcap C_k(a) \notin \mathcal{A}^* \setminus \mathbb{E}$ and $C_1 \sqcap \cdots \sqcap C_k \in \mathit{Sub}\mathcal{C}$. $\mathcal{A}^*$ being assertionally closed implies that $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathcal{A}^*$ and hence $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathbb{E}$. Since $\mathbb{E}$ is $\mathfrak{R}$-closed, due to the $\sqcap_2^{\mathbb{S}}$-rule, there is $C_i(a) \in \mathbb{E}(1 \leq i \leq k)$, contradicting $\{C_1(a), ..., C_k(a)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$.
- If $\exists_1^{\mathcal{A}}$-rule is applicable, then $\{r(a, b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$, $\exists r.C \in \mathit{Sub}\mathcal{C}$ and $\exists r.C(a) \notin \mathcal{A}^* \setminus \mathbb{E}$. $\mathcal{A}^*$ being assertionally closed implies that $\exists r.C(a) \in \mathcal{A}^*$ and so $\exists r.C(a) \in \mathbb{E}$. Since $\mathbb{E}$ is $\mathfrak{R}$-closed, due to the $\exists_1^{\mathbb{S}}$-rule, $\{r(a, b), C(b)\} \cap \mathbb{E} \neq \emptyset$, contradicting $\{r(a, b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$.
- If $\exists_2^{\mathcal{A}}$-rule is applicable, then $\exists r.C(a) \in \mathcal{A}^* \setminus \mathbb{E}$, a is not blocked and there does not exist $b \in \mathcal{O}^*$ such that $\{r(a, b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$. Together with the assumption that $\mathcal{A}^*$ is assertionally closed, it follows that for all $c \in \mathcal{O}^*$ with $\{r(a, c), C(c)\} \subseteq \mathcal{A}^*$, we have $\{r(a, c), C(c)\} \cap \mathbb{E} \neq \emptyset$. Since $\mathbb{E}$ is $\mathfrak{R}$-closed, due to the $\exists_2^{\mathbb{S}}$-rule, $\exists r.C(a) \in \mathbb{E}$, a contradiction.
- If $\sqsubseteq^{\mathcal{T}}$-rule is applicable, then $C(a) \in \mathcal{A}^* \setminus \mathbb{E}$, $C \sqsubseteq D \in \mathcal{T}$ and $D(a) \notin \mathcal{A}^* \setminus \mathbb{E}$. $\mathcal{A}^*$ being assertionally closed implies that $D(a) \in \mathcal{A}^*$. It follows that $D(a) \in \mathbb{E}$. Since $\mathbb{E}$ is $\mathfrak{R}$-closed, due to $\sqsubseteq^{\mathbb{S}}$-rule, $C(a) \in \mathbb{E}$, contradicting $C(a) \in \mathcal{A}^* \setminus \mathbb{E}$. ■

Corollary 2. *Given $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a finite secrecy set $\mathbb{S}$, let $\mathcal{A}^*$ be the assertional closure of Σ and suppose that $\mathbb{E}$ is $\mathfrak{R}$-closed. Then $(\mathcal{A}^* \setminus \mathbb{E})^+ \subseteq \mathcal{A}^* \setminus \mathbb{E}$.*

Proof. It follows from Lemma 3 that no assertion expansion rules are applicable to $\mathcal{A}^* \setminus \mathbb{E}$. ■

Corollary 3. *Given $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a finite secrecy set $\mathbb{S}$, let $\mathcal{A}^*$ be the assertional closure of Σ. If $\mathbb{E}$ is $\mathfrak{R}$-closed, then $\mathbb{E}$ is a secrecy envelope of $\mathbb{S}$.*

Proof. It follows from Corollary 2 and Definition 2. ■

It turns out that $\Lambda_{\mathbb{S}}^{\mathfrak{R}}$ algorithm, although certainly producing an envelope, may actually result an envelope that contains redundant information. This compromises our goal of informativeness. Specifically, even if $\exists_2^{\mathcal{A}}$-rule is applicable to $(\mathcal{A}^* \setminus \mathbb{E}_{\mathbb{S}})^+$, due to OWA, the querying agent can only conclude that there exists an individual d that is the witness for $\exists r.C(a)$ and that $d \notin \mathcal{O}_{\Sigma}$. However, by HNA, the querying agent has no computational access to individual names in $\mathcal{O}^* \setminus \mathcal{O}_{\Sigma}$ and hence it cannot infer any information in $\mathbb{E}_{\mathbb{S}}$. This provides a cue that when computing a secrecy envelope, the $\exists_2^{\mathbb{S}}$-rule, which inverts the $\exists_2^{\mathcal{A}}$-rule, is dispensable. Figure 4 lists a set of secrecy closure rules, called *$\mathcal{Q}$-Secrecy Closure Rules*, that can be used to compute an envelope. The $\mathcal{Q}$-Secrecy Closure Rules do not include the $\exists_2^{\mathbb{S}}$-rule, and make one further "optimization" in that the $\exists_1^{\mathbb{S}}$-rule is replaced by the $\exists^{\mathbb{S}}$-rule in which the individual name b is restricted to be in $\mathcal{O}_{\Sigma}$ rather than in $\mathcal{O}^*$.

When no $\mathcal{Q}$-secrecy closure rule is applicable, we say that $\mathbb{E}$ is *$\mathcal{Q}$-closed.* We denote by $\Lambda_{\mathbb{S}}^{\mathcal{Q}}$ the tableau algorithm which nondeterministically applies the $\mathcal{Q}$-secrecy closure rules until no further rules are applicable. It is clear that all executions of $\Lambda_{\mathbb{S}}^{\mathcal{Q}}$ on an input consisting of the assertional closure of a knowledge base and a finite set of assertions terminate and that when $\Lambda_{\mathbb{S}}^{\mathcal{Q}}$ terminates, $\mathbb{E}$ is $\mathcal{Q}$-closed.

Lemma 4. *Given $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a secrecy set $\mathbb{S}$, let $\mathcal{A}^*$ be the assertional closure of Σ and suppose that $\mathbb{E}$ is $\mathcal{Q}$-closed. Then $\mathcal{A}^* \setminus \mathbb{E}$ is closed under $\sqcap_1^{\mathcal{A}}$, $\sqcap_2^{\mathcal{A}}$, $\exists_1^{\mathcal{A}}$ and $\sqsubseteq^{\mathcal{T}}$ rules for the querying agent.*

Proof. The proof that $\mathcal{A}^* \setminus \mathbb{E}$ is closed under $\sqcap_1^{\mathcal{A}}$-, $\sqcap_2^{\mathcal{A}}$- and $\sqsubseteq^{\mathcal{T}}$-rules is similar to the proof of the corresponding cases in Lemma 1. Here we prove that $\exists_1^{\mathcal{A}}$-rule is not applicable to $\mathcal{A}^* \setminus \mathbb{E}$.

If $\exists_1^{\mathcal{A}}$-rule is applicable, then $\{r(a, b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$, $\exists r.C \in \mathit{Sub}\mathcal{C}$ and $\exists r.C(a) \notin \mathcal{A}^* \setminus \mathbb{E}$. $\mathcal{A}^*$ being assertionally closed implies that $\exists r.C(a) \in \mathcal{A}^*$ and so we must have $\exists r.C(a) \in \mathbb{E}$. There are two cases:

$\sqcap_1^{\mathbb{S}}$ -rule: if $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathcal{A}^* \setminus \mathbb{E}$ and $\{C_1(a), ..., C_k(a)\} \cap \mathbb{E} \neq \emptyset$,
then $\mathbb{E} := \mathbb{E} \cup \{C_1 \sqcap \cdots \sqcap C_k(a)\}$;
$\sqcap_2^{\mathbb{S}}$ -rule: if $C_1 \sqcap \cdots \sqcap C_k(a) \in \mathbb{E}$ and $\{C_1(a), ..., C_k(a)\} \cap \mathbb{E} = \emptyset$,
then $\mathbb{E} := \mathbb{E} \cup \{C_i(a)\}$ where $1 \leq i \leq k$;
$\exists^{\mathbb{S}}$ -rule: if $\exists r.C(a) \in \mathbb{E}$ and $\{r(a,b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$ with $b \in \mathcal{O}_\Sigma$,
then $\mathbb{E} := \mathbb{E} \cup \{r(a,b)\}$ or $\mathbb{E} := \mathbb{E} \cup \{C(b)\}$;
$\sqsubseteq^{\mathbb{S}}$ -rule: if $D(a) \in \mathbb{E}, C \sqsubseteq D \in \mathcal{T}$ and $C(a) \in \mathcal{A}^* \setminus \mathbb{E}$, then $\mathbb{E} := \mathbb{E} \cup \{C(a)\}$.

Fig. 4. $\mathcal{Q}$-Secrecy Closure Rules

- If $\{a, b\} \subseteq \mathcal{O}_\Sigma$, since $\mathbb{E}$ is $\mathcal{Q}$-closed, due to the $\exists^{\mathbb{S}}$-rule, $\{r(a,b), C(b)\} \cap \mathbb{E} \neq \emptyset$, contradicting $\{r(a,b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$.
- If $\{a, b\} \nsubseteq \mathcal{O}_\Sigma$, suppose $b \notin \mathcal{O}_\Sigma$. This together with $\mathcal{A}^*$ being assertionally closed implies that there does not exist $c \in \mathcal{O}_\Sigma$ such that $\{r(a,c), C(c)\} \subseteq \mathcal{A}^*$. By HNA, the querying agent has no computational access to b. Therefore, $\exists_1^{\mathcal{A}}$-rule is not applicable for the querying agent. If $a \notin \mathcal{O}_\Sigma$, the querying agent has no computational access to any of the assertions $r(a,b)$ and $\exists r.C(a)$. ■

Theorem 3. *Given $\Sigma = \langle \mathcal{A}, \mathcal{T} \rangle$ and a finite secrecy set $\mathbb{S}$, suppose that $\mathcal{A}^*$ is the assertional closure of Σ. If $\mathbb{E}$ is $\mathcal{Q}$-closed, then $\mathbb{E}$ is a secrecy envelope of $\mathbb{S}$.*

Proof. We need to show that $(\mathcal{A}^* \setminus \mathbb{E})^+ \cap \mathbb{S} = \emptyset$. Since $\mathbb{E}$ is $\mathcal{Q}$-closed, by Lemma 2, for the querying agent, $\sqcap_1^{\mathcal{A}}$, $\sqcap_2^{\mathcal{A}}$, $\exists_1^{\mathcal{A}}$ and $\sqsubseteq^{\mathcal{T}}$ rules are not applicable to $\mathcal{A}^* \setminus \mathbb{E}$.

If $\exists_2^{\mathcal{A}}$-rule is applicable to $\mathcal{A}^* \setminus \mathbb{E}$, then $\exists r.C(a) \in \mathcal{A}^* \setminus \mathbb{E}$, a is not blocked and there does not exist $b \in \mathcal{O}^*$ such that $\{r(a,b), C(b)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$. Since $\mathcal{A}^*$ is assertionally closed and there is no witness $c \in \mathcal{O}_\Sigma$ such that $\{r(a,c), C(c)\} \subseteq \mathcal{A}^* \setminus \mathbb{E}$, there is $d \in N_{\mathcal{O}} \setminus \mathcal{O}_\Sigma$ such that $\{r(a,d), C(d)\} \subseteq \mathcal{A}^*$. By OWA, the querying agent can infer the existence of some satisfying individual. However, due to HNA, it can not infer which individual it is. It follows that $(\mathcal{A}^* \setminus \mathbb{E})^+ \cap \mathbb{S} = \emptyset$. ■

Note that the whole initialization of the secrecy maintenance system (including the computation of $Sub\mathcal{C}$, $\mathcal{A}^*$ and $\mathbb{E}_{\mathbb{S}}$) is easily seen to be doable in polynomial time in the size of the KB Σ plus the size of the given secrecy set $\mathbb{S}$.

4 Queries

In this section we assume that the three sets $Sub\mathcal{C}$, $\mathcal{A}^*$ and $\mathbb{E}_{\mathbb{S}}$ (the latter two, restricted to $Sub\mathcal{C}$) have been precomputed in the pre-query stage as described in Sect. 3.

The computation of the answer to a query of the form $C(a)$ is given in Fig. 5. The input of the secrecy-preserving query-answering procedure SPQA contains the TBox $\mathcal{T}$ in normal form, precomputed assertional closure $\mathcal{A}^*$, the query $C(a)$ and the precomputed secrecy envelope $\mathbb{E}_{\mathbb{S}}$.

Since sub-expressions in the concept C of the query $C(a)$ need not appear in $Sub\mathcal{C}$, Line 3 in the SPQA procedure computes sub-expressions of the concept C as defined in Sect. 3.1 and expand $Sub\mathcal{C}$ by adding expressions in $sub(C) \setminus Sub\mathcal{C}$. The expanded $Sub\mathcal{C}$ will be used to update $\mathcal{A}^*$ by applying assertion expansion rules (Fig. 2) until none of them is applicable, as indicated in Line 4. As a consequence, there may be applicable secrecy closure rules (Fig. 4), implying that $\mathbb{E}_{\mathbb{S}}$ may no longer be a secrecy envelope for $\mathbb{S}$. Therefore, we apply necessary secrecy closure rules exhaustively (Line 5).

```
SPQA($\mathcal{T}, \mathcal{A}^*, C(a), \mathbb{E}_{\mathbb{S}}$):
1.  if ($C \notin Sub\mathcal{C}$)
2.  {
3.    compute $sub(C)$; $Sub\mathcal{C} = Sub\mathcal{C} \cup sub(C)$;
4.    expand $\mathcal{A}^*$ to $Sub\mathcal{C}$;
5.    expand the secrecy envelope $\mathbb{E}_{\mathbb{S}}$ to $Sub\mathcal{C}$;
6.  }
7.  if ($C(a) \in \mathcal{A}^*$ and $C(a) \notin \mathbb{E}_{\mathbb{S}}$)
8.    return "Yes"
9.  else
10.   return "Unknown"
```

Fig. 5. Secrecy-preserving Query-answering Procedure

We denote by $\mathcal{A}_1^*$ the assertional closure of Σ after the application of Line 4. By Theorem 1, the newly added assertions are entailed by Σ. We also denote by $\mathbb{E}'_{\mathbb{S}}$ the expansion of $\mathbb{E}_{\mathbb{S}}$ after the application of Line 5. By Theorem 3, since $\mathbb{E}'_{\mathbb{S}}$ is $\mathcal{Q}$-closed, it is a secrecy envelope. Clearly, a single invocation of the procedure SPQA takes polynomial time (in the sum of the sizes of its arguments).

For queries of the form $r(a, b)$, the procedure is much simpler: if $r(a, b) \in \mathcal{A} \setminus \mathbb{E}_{\mathbb{S}}$, then the answer is "Yes"; otherwise, the answer is "Unknown". Here $\mathbb{E}_{\mathbb{S}}$ is the current secrecy envelope.

Example 5. (Example 1-4, continued) Recall that we have a knowledge base $\Sigma_1 = \langle \mathcal{A}_1, \mathcal{T}_1 \rangle$ and the secrecy set $\mathbb{S}_1 = \{$CancerRisk(Jane)$\}$ in Example 2.

The assertional closure of Σ_1, denoted by $\mathcal{A}_1^*$, and one possible envelope $\mathbb{E}_{\mathbb{S}1}$ are listed below:

$\mathcal{A}_1^* = \mathcal{A}_1 \cup \{$ A(Jill), ∃is_child.A(Jane), CancerRisk(Jane), has_pres(Jane, a),
∃has_pres.CancerDrug(Jane), CancerDrug(a), CoveredDrug(a),
∃has_pres.CoveredDrug(Jane), Reimburse(Jane)$\}$.

$\mathbb{E}_{\mathbb{S}1} = \{$CancerRisk(Jane), is_child(Jane, Jill), HasMutBRCA1(Jane),
∃is_child.A(Jane), ∃has_pres.CancerDrug(Jane)$\}$.

If the querying agent asks the query Reimburse(Jane), Reimburse(Jane)$\in \mathcal{A}_1^* \setminus \mathbb{E}_{\mathbb{S}1}$, the answer to the query is "Yes". If the querying agent asks the query CancerRisk(Jane), since CancerRisk(Jane)$\in \mathcal{A}_1^* \cap \mathbb{E}_{\mathbb{S}1}$, the answer to the query is "Unknown". ■

5 Summary and Discussion

Summary: In this paper, we have introduced a logic-based framework for secrecy preserving query answering in $\mathcal{EL}$ knowledge bases. We have provided a polynomial time algorithm that, given an $\mathcal{EL}$ KB Σ, a set $\mathbb{S}$ of secrets to be protected and a query q, truthfully answers the query whenever: (i) $\Sigma \models q$ and (ii) the answer to q, together with the answers to any previous queries answered by the KB does not allow the querying agent to deduce any of the secrets in $\mathbb{S}$. Our approach exploits the open world semantics under which it is impossible for the querying agent to distinguish between a scenario in which the answer to the query is "Unknown" because of the incomplete knowledge of the KB or because of selective censoring of answers by the KB. Our secrecy-preserving reasoning framework builds on, and substantially extends, the privacy-preserving reasoning framework introduced by Bao et al. [5] which considered protecting class-subclass relationships in hierarchical ontologies.

Related Work: Problems of trust, privacy and security in information systems, including the web-based information systems, are topics of significant current interest. Most of the work in this area falls into four broad categories: (i) Access control mechanisms (see [18] for an overview) and logic-based policies that control access to sensitive information based on various criteria (e.g., role of the individual) (see [6] for a survey) or data encryption using cryptographic protocols to prevent of unauthorized access to sensitive information (see for example, [9]); (ii) Information confinement [20], information flow control mechanisms [14] and programming language constructs for enforcing information flow policies (see Sabelfeld and Myers [25] for a survey); (iii) Preventing disclosure of information about specific individuals from statistical or aggregate information about a population [12], or as a result of data mining [1, 21], or record linkage [23, 17] or database queries [13]; and (iv) Controlled query evaluation [26] which offers a mechanism for answering database queries without revealing secrets (see [7] for a survey). Recently, Baader et al. [3] provide an approach to generate views for users based on their access rights. Grau and Motik [16] have studied the problem of hiding a part of the signature of an ontology that is reused (via importing) by another ontology.

In contrast to access-control mechanisms, our approach permits the use of secrets in answering queries when it is possible to do so without compromising secrets. Unlike in the case of work on information flow where the main focus is on mechanisms for enforcing information flow policies or certifying compliance with such policies by procedural programs, our focus is on mechanisms for secrecy-preserving query answering against KBs. In contrast to work on preventing information about specific individuals from being disclosed by statistical queries or data mining, our focus is on preventing secrets being compromised by answers to queries that are answered using purely deductive (logical) inferences from a knowledge base. Unlike controlled query evaluation, which, with the exception of [8], has largely focused on protecting secrets in (typically relational) databases under the closed world assumption, our focus is on secrecy-preserving query answering against KBs under the open world assumption. Furthermore, we focus on cooperative as opposed to adversarial scenarios where the KB is not allowed to lie (although it is allowed to censor some of the answers). Such scenarios naturally arise in many information sharing applications of practical interest (e.g., selective information sharing among: different branches of government; patients, physicians, and health insurance providers; intelligence agencies of friendly nations, etc.).

Future Work: Some natural directions for future work include: (i) design of an efficient algorithm for computing a “tight” envelope for $\mathcal{EL}$ KBs, i.e., an envelope from which no statement can be dropped without risking the possibility of secrets being compromised (such an algorithm is of interest in light of the fact that our current algorithm is not guaranteed to produce a tight envelope and the fact that computing the minimum envelope is NP-hard); (ii) exploration of secrecy-preserving query answering algorithms in the case of more expressive e.g., $\mathcal{ALC}$, DL-Lite, and RDF KBs; (iii) investigation of secrecy-preserving query answering in settings with multiple querying agents, under various restrictions on communication among agents.

References

1. Rakesh Agrawal and Ramakrishnan Srikant. Privacy-preserving data mining. In *Proc. of the ACM SIGMOD Conference on Management of Data*, pages 439–450. ACM Press, May 2000.
2. F. Baader. Computing the least common subsumer in the description logic EL wrt terminological cycles with descriptive semantics. In *Proceedings of the 11th International Conference on Conceptual Structures, ICCS*, volume 2746, pages 117–130. Springer, 2003.
3. F. Baader, M. Knechtel, and R. Peñaloza. A generic approach for large-scale ontological reasoning in the presence of access restrictions to the ontologys axioms. *The Semantic Web-ISWC 2009*, pages 49–64, 2009.
4. Franz Baader. Terminological cycles in a description logic with existential restrictions. In *IJCAI'03: Proceedings of the 18th international joint conference on Artificial intelligence*, pages 325–330, San Francisco, CA, USA, 2003. Morgan Kaufmann Publishers Inc.

5. Jie Bao, Giora Slutzki, and Vasant Honavar. Privacy-preserving reasoning on the semantic web. In *Web Intelligence*, pages 791–797. IEEE Computer Society, 2007.
6. Elisa Bertino, L. R. Khan, Ravi S. Sandhu, and Bhavani M. Thuraisingham. Secure knowledge management: confidentiality, trust, and privacy. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 36(3):429–438, 2006.
7. Joachim Biskup and Piero A. Bonatti. Controlled query evaluation for enforcing confidentiality in complete information systems. *Int. J. Inf. Sec.*, 3(1):14–27, 2004.
8. Joachim Biskup and Torben Weibert. Keeping secrets in incomplete databases. *Int. J. Inf. Sec.*, 7(3):199–217, 2008.
9. D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In *Advances in CryptologyCRYPTO 2001*, pages 213–229. Springer, 2001.
10. S. Brandt. Reasoning in ELH wrt General Concept Inclusion Axioms. Technical report, Dresden, TU, 2004.
11. Sebastian Brandt. Polynomial time reasoning in a description logic with existential restrictions, gci axioms, and - what else? In Ramon López de Mántaras and Lorenza Saitta, editors, *ECAI*, pages 298–302. IOS Press, 2004.
12. F.Y. Chin. Security in statistical databases for queries with small counts. *ACM Transactions on Database Systems*, 3:92–104, 1978.
13. Bernardo Cuenca Grau and Ian Horrocks. Privacy-preserving query answering in logic-based information systems. In *Proc. of the 18th Eur. Conf. on Artificial Intelligence (ECAI 2008)*, 2008.
14. Dorothy E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, 1976.
15. Csilla Farkas, Alexander Brodsky, and Sushil Jajodia. Unauthorized inferences in semistructured databases. *Information Sciences*, 176:3269–3299, 2006.
16. Bernardo Cuenca Grau and Boris Motik. Pushing the Limits of Reasoning over Ontologies with Hidden Content. In *Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning*, pages 214–224, 2010.
17. Lifang Gu, Rohan Baxter, Deanne Vickers, and Chris Rainsford. Record linkage: Current practice and future directions. Technical report, CSIRO Mathematical and Information Sciences CMIS Technical Report No. 03/83, 2003.
18. Sushil Jajodia. Database security and privacy. *ACM Comput. Surv.*, 28(1):129–131, 1996.
19. Adila Krisnadhi and Carsten Lutz. Data complexity in the el family of dls. In Diego Calvanese, Enrico Franconi, Volker Haarslev, Domenico Lembo, Boris Motik, Anni-Yasmin Turhan, and Sergio Tessaris, editors, *Description Logics*, volume 250 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2007.
20. Butler W. Lampson. A note on the confinement problem. *Communications of the ACM*, 16(10):613–615, 1973.
21. Yehuda Lindell and Benny Pinkas. Privacy Preserving Data Mining. *Journal of Cryptology*, 15(3):177–206, 2002. An extended abstract appeared at the CRYPTO 2000 conference.
22. N. Novakovic. Proof-theoretic Approach to Deciding Subsumption and Computing Least Common Subsumer in EL wrt Hybrid TBoxes. In *Logics in Artificial Intelligence: 11th European Conference, JELIA 2008, Dresden, Germany, September 28-October 1, 2008. Proceedings*, page 311. Springer-Verlag New York Inc, 2008.
23. Christine M. O'Keefe, Ming Yung, Lifang Gu, and Rohan Baxter. Privacy-preserving data linkage protocols. In *WPES '04: Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 94–102, New York, NY, USA, 2004. ACM.
24. A. Rector and I. Horrocks. Experience building a large, re-usable medical ontology using a description logic with transitivity and concept inclusions. In *Proceedings of the Workshop on Ontological Engineering, AAAI Spring Symposium (AAAI97), Stanford, CA*, pages 321–325, 1997.
25. Andrei Sabelfeld and Andrew C. Myers. Language-based information-flow security. *Selected Areas in Communications, IEEE Journal on*, 21(1):5–19, Jan 2003.
26. George L. Sicherman, Wiebren de Jonge, and Reind P. van de Riet. Answering queries without revealing secrets. *ACM Trans. Database Syst.*, 8(1):41–59, 1983.
27. K. Spackman. Managing clinical terminology hierarchies using algorithmic calculation of subsumption: Experience with SNOMED-RT. *J. of the Amer. Med. Informatics Assoc.*, 2000.