Journal of Biomedical Informatics 41 (2008) 1028–1040

Contents lists available at ScienceDirect

Journal of Biomedical Informatics

journal homepage: www.elsevier.com/locate/yjbin

Situation-Based Access Control: Privacy management via modeling of patient data access scenarios

Mor Peleg [a], Dizza Beimel [b,*], Dov Dori [b], Yaron Denekamp [c]

[a] Department of Management Information systems, University of Haifa, Israel
[b] Information Systems Area, Faculty of Industrial Engineering and Management, Technion-Israel Institute of Technology, Haifa 32000, Israel
[c] Carmel Medical Center, Faculty of Medicine, Technion, Israel Institute of Technology, Haifa, Israel

ARTICLE INFO

Article history:
Received 9 November 2007
Available online 10 April 2008

Keywords:
Privacy preservation
Access control
Conceptual model
Ontology
Object-Process Methodology

ABSTRACT

Access control is a central problem in privacy management. A common practice in controlling access to sensitive data, such as electronic health records (EHRs), is Role-Based Access Control (RBAC). RBAC is limited as it does not account for the circumstances under which access to sensitive data is requested. Following a qualitative study that elicited access scenarios, we used Object-Process Methodology to structure the scenarios and conceive a Situation-Based Access Control (SitBAC) model. SitBAC is a conceptual model, which defines scenarios where patient's data access is permitted or denied. The main concept underlying this model is the Situation Schema, which is a pattern consisting of the entities Data-Requestor, Patient, EHR, Access Task, Legal-Authorization, and Response, along with their properties and relations. The various data access scenarios are expressed via Situation Instances. While we focus on the medical domain, the model is generic and can be adapted to other domains.

1. Introduction

During the last two decades, we have been witnessing a gradual transfer from paper-based to electronic-based organization of information. At the same time, many organizations have become connected via the Internet, making information in general and data concerning people's private lives in particular accessible to unauthorized vulnerable access. Personal details can be (and often are) collected, recorded, stored, processed, and transferred to unknown third parties. This reality is in conflict with the fundamental human right to control one's own personal information. The gap between the ease of access to one's personal details and the human desire to control this access is the root cause of the privacy management and assurance problem.

What is privacy? Jones [1] has defined privacy as "... not having things known about you that you don't choose to have known, or at least you know that they are known and by whom." Following Jones' perception, privacy can be interpreted as a human desire to keep certain personal details confidential. Indeed, people usually choose not to disclose their personal details unless they have to, in particular when it comes to their financial and health conditions. This view is based on findings of a survey conducted in Australia [2], where the respondents were asked to rank the types of personal data, which they prefer not to reveal. Financial and health information were at the top of the list.

Privacy is often considered a social construct [3,4], which is more than just an individual's desire to control access to his or her personal information. The privacy preservation problem has a major effect on human communities, as it touches upon social, cultural, economic, and political aspects. Privacy is influenced by legislation and legal changes, such as the right of free speech [3] and the right of governments to prosecute criminal activities [5], changes in technology, changes in journalistic practices [3], and private sector initiatives, such as WebTrust and other privacy-seal programs [4].

In many Western countries, the privacy issue has been considered intensively from the legal aspect. The most important legislation concerning healthcare privacy preservation is the US Privacy Rule [6], which resulted primarily from the requirements of the 1996 Health Insurance Portability and Accountability Act (HIPAA) [7]. The goal of the US Privacy Rule is to ensure that individuals' health information is properly protected while allowing the flow of health information needed to provide high-quality health care to the individual. However, in addition to the goal of providing the patient with the best medical care, there are other important goals, such as financial reimbursement and public-health requirements. For example, in public-health surveillance [8], when an outbreak is detected and anonymized data is used to trace its roots,

* Corresponding author. Address: Department of Business Administration, Ruppin Academic Center, Emek Hefer, 40250, Israel. Fax: +972 9 8987604.
E-mail address: dizza@technion.ac.il (D. Beimel).

1532-0464/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jbi.2008.03.014

some personal information may be released for further investigation.

In the case of privacy of medical data, patients are not the only stakeholders. Some stakeholders are interested in medical data for their own benefits. Quoting Welch, an M.D. from Massachusetts General Hospital in Boston [9], "Insurers use identifiable medical records for risk rating, employers use them for hiring and firing, health systems for quality assurance, pharmaceutical firms for marketing, banks for assessing loan risk, and the government for the detection of fraud." Other stakeholders who are interested in patient data for the benefit of the patient or the society include (1) hospital personnel and primary care physicians, (2) family members, (3) expert physicians, laboratory workers, health institutions, and pharmacies, (4) dietitians and alternative-medicine service providers, (5) insurance companies, (6) scientists, researchers, and public-health agencies, and (7) community agents, including fire services, police, ambulance services, and restaurants specializing in preparation of food for sick people according to their physician's orders, such as "Meals on Wheels."

In recognition of these views and following previous research works, our focus has been to develop an approach and a method for maintaining privacy within the healthcare domain in a way that would strike a balance between the needs for preserving patient privacy on one hand and quality of the patient's medical care on the other hand.

The literature referring to the privacy problem can be roughly divided into three categories:

(1) privacy preservation via *identity protection*, e.g., fingerprint recognition system [10],
(2) privacy preservation via *anonymity*, e.g., anonymizing private data that include explicit identifiers [11], and
(3) privacy preservation via restricting access to data, e.g., *access control and authorization models*.

The model we suggest belongs to the third category. We wish to achieve privacy preservation by means of access control. To this end, we are interested in identifying *scenarios* of access requests. A scenario is a specific process description of a data access request to patients' data. We express scenarios in a semi-formal, computer-interpretable way. This enables preservation of the patient's privacy by exposing only that personal data which is indeed required for providing the patient with the best medical care in any given particular scenario.

Significant parts of related works with a goal similar to ours focus on *Role-Based Access Control* (RBAC) [12,13] mechanisms. RBAC is an approach that separates users and their permissions regarding a collection of resources and places the role of the data-requestor at the center. Our approach is different than that of RBAC because it is situation-based rather than role-based. Our approach takes into account that the decision to uncover patients' health data is affected by various factors that comprise the situation of a patient's data access. Indeed, one of these factors is the role of the entity that requests the data, but it is definitely not the only one. Since there are additional situation factors, we call our access control approach *Situation-Based Access Control*, or *SitBAC* for short. Unlike RBAC, SitBAC includes abstractions for modeling the entities involved in a situation—Patient, Data-Requestor, Task, Legal-Authorization, EHR, and Response—along with their attributes and the relations among them. We argue that using the healthcare-related abstractions of SitBAC with focus on situations, rather than using the RBAC approach with focus on roles, will result in better coverage of the various privacy issues that healthcare organizations are expected to handle. Hence, SitBAC is likely to also improve the performance of organization privacy officers.

Our access control model is designed to enable implementation of the organization's policies. As a result of modeling data access scenarios, we have obtained detailed specifications of various data access situations. These precise models can serve as a basis for organizations to apply their internal policies and regulations regarding which action should be carried out by whom. For example, the details of a patient can be shown to a specialist who needs to treat her only if there is a referral of the patient to that specialist and if the specialist is accessing the information from his usual working place. By modeling scenarios such as this we ensure the maintenance of the patient's privacy while abiding by the organization's regulations and enabling high-quality care provision.

Throughout the research, we used two research methods: (1) qualitative methods to elicit the required data and (2) conceptual modeling to structure and specify the scenarios. Our long-term goal is to formalize the conceptual SitBAC model as a *situation knowledge base*, such that each new access request, represented as a situation instance, will be compared against the situation knowledge base to determine whether the request can be granted or has to be denied.

The rest of the paper is organized as follows: Section 2 covers background material, Section 3 presents the research questions, Section 4 describes the methods we used to carry out the research, Section 5 presents the results of the research, and Section 6 discusses them and the significance of our approach. Section 7 concludes with a discussion on the applicability of the research.

2. Background

Our research is motivated by preservation of privacy while accessing Electronic Health Records (EHRs). Accordingly, the review in this section covers the following issues: (1) reasons for privacy jeopardy or loss, (2) EHR systems, (3) the Role-Based Access Control model, and (4) access control approaches and work related to achieving privacy preservation by means of such methods.

2.1. Reasons for privacy jeopardy or loss

Jones [1] has claimed that modern technology constitutes the main reason for privacy loss and reports the following technological characteristics that exacerbate the privacy problem.

(1) *Permanence and volume:* Transient data (like images) about events is increasingly captured, recorded, and kept in digital form. This immense volume of data includes not just primary data, but also secondary or associated data and metadata of all kinds, which may be very extensive.
(2) *Invisibility, accessibility, and remoteness:* Data about individuals, mostly in the public domain, is being constantly collected without getting the individuals' approval first or bringing the collection to their attention. Once collected, this data can be read by any number of people, few or many, authorized or unauthorized, conceptually near or afar.
(3) *Assembly and aggregation:* It is possible to collect data about a person from any number of sources and combine them in ways that provide additional private information.

Due to the growing use of search engines on the Web and within organizations, these characteristics are more prominent today than in the past.

2.2. Electronic Health Records

In the healthcare domain, digital data about a patient that is collected into a record is often referred to as the *Electronic Health Re-*

cord (*EHR*), Electronic Medical Record (EMR), or Computerized Patient Record (CPR). EHRs record such data items as diagnoses, hospital admissions, medications, operations, laboratory tests, imaging, and pathology data. Shabo [14] has made the distinction between *health* records and *medical* records by noting that the former are more general; they are longitudinal and cross-institutional, and may include individual health data such as lifestyle, workplace hazards, and preferences, which was not necessarily created during medical practice.

Since the EHR concept appears to be very complex, an endeavor of Standards Development Organizations (SDOs) has been to anchor the EHR to a standard. Two prominent SDOs in the healthcare domain are Health Level 7 (HL7) [15] and the European Committee for Standardization (CEN) [16]. Both have focused on creating a standard for exchanging health data, which might potentially be the basis for a future standard EHR.

The motivation for producing a protocol for health data exchange prior to a standard for EHR structure and content draws on two main insights: (1) generating one record that will hold any type of health-related data item for the entire lifetime of a person appears to be a very ambitious goal, (2) various EHR systems exist within healthcare organizations, and the immediate necessity is to enable interoperability, i.e., inter-organizational data exchange and sharing in order to increase the quality of service to patients. Simple examples include prevention of repeating orders of blood tests and reuse of historical data.

Data exchange standards, such as DICOM (Digital Imaging and Communications in Medicine) [17] and HL7 version 2 messaging standard [15], are already operational and widely implemented. DICOM is a standard for handling, storing, printing, and transmitting information in medical imaging, which has been developed and sponsored by the American College of Radiology and the National Electrical Manufacturers Association (NEMA). The HL7 version 2 messaging standard aims to support data interchange and hospital workflows. Originally created in 1987, it defines a series of electronic messages to support administrative, logistical, financial, and clinical processes.

More recently, HL7 has developed the Clinical Document Architecture (CDA), which is a specification for a medical record message that can serve as an infrastructure for exchanging medical data, such as progress notes, discharge summaries, and results of physical examinations. CEN has published the pre-standard prENV 13606 by TC.251 for electronic health record communication, which includes four parts: (1) extended architecture, (2) domain term list, (3) distribution rules, and (4) messages for the exchange of information.

2.3. The RBAC model for restricting access to personal data

Privacy via access control and authorization refers to tailored access control tools, access–control languages, or access control models that address the provision of infrastructure for developing access–control and authorization mechanisms. One of the more recognized access control models, the Role-Based Access Control (RBAC) model, was proposed in 1996 by Sandhu et al. [12,13], who adopted the "need-to-know" concept and integrated it into their model. The concept assumes that privacy is preserved as long as data access processes occur only when they are necessary for a right purpose, and minimum details are revealed along the process. Several access control languages were developed on top of this model. Since RBAC is one of the leading models in the access-control/authorization domain, the following paragraphs elaborate on this model, in particular since we evaluate our SitBAC approach by comparing it to the RBAC model.

The Role-Based Access Control (RBAC) [12,13] approach advocates distinguishing different user types and their privileges regarding a collection of resources. Instead of dealing with the privileges (permissions) of each user specifically, the users are grouped into roles, and each such group is associated with a number of privileges. Roles can be created and added as much as the system requires. Each function defined in the system (read, write, append, delete, create, etc.) can be associated with a privilege that is assigned to a role. Users are assigned roles based on their responsibilities and qualifications. Any user can be reassigned from a role and assigned to another. Roles are dynamic, so once they are created with a set of permissions, this set can be changed dynamically without additional compilation. New permissions can be added to the role and existing permissions can be deleted from it.

The core RBAC includes five basic elements: Users, Roles, Objects, Operations, and Permissions. RBAC is organized into four levels of increasing functional capabilities: (1) *Flat RBAC*, shown in Fig. 1, supports multiple users per role, multiple roles per user, multiple permissions (privileges) per role, and multiple roles per permission, (2) *Hierarchical RBAC*, which supports role hierarchies and inheritance, (3) *Constrained RBAC*, which enforces Separation of Duty (SOD), and (4) *Symmetric RBAC*, in which a requirement for permission-role review is added.

The RBAC approach is widely employed in the healthcare domain, probably because the classical healthcare roles (physician, nurse, secretary, etc.) are straightforward. The following subsection describes several works that aim to protect health-related data by means of access control with emphasis on role definitions.

2.4. Access-control-based works for protecting healthcare-related data

The majority of the works related to protecting healthcare data place the role of the data-requestor in the focus of the request for patient data access. However, as the evaluations below show, the results are not quite satisfactory. The first works recognized the need of roles [18] and emphasized that role-defining processes in an EHR system are crucial for privilege management. Next, structured roles [19] were characterized, offering solutions that follow the rationale of allowing access only to authorized entities, in accord with a table of authorizations supervised by a security committee. Later on, researchers realized that roles have a dynamic aspect in addition to their structural aspect [20]. Consequently, existing authorization concepts had to be extended in order to support particular healthcare authorization requirements. For instance, a physician is allowed to access health information about a shared patient only in case s/he is involved in the treatment of that patient. Finally, an observation was made that roles have more than one contextual variable (e.g., time, place, affiliation), which affect their behavior. Motta and Furuie [21] proposed a contextual role-based access control authorization model aiming to increase patient privacy and the confidentiality of patient data while being flexible enough to consider specific cases. They suggested defining a role hierarchy with inheritance of authorizations and modeling the types of data found in an EHR according to clinical content (e.g., demographics, prescriptions). Authorizations are defined as a 5-tuple with the following elements.

- R, the role;
- PT, the privilege type, which can be positive when an operation is allowed or negative when it is disallowed;
- Opr, the operation (or access mode);
- Obj, the object (or resource) to be protected; and
- At, the authorization type, which can be strong or weak.

The authors also proposed a technique for handling conflicts between authorizations.

Fig. 1. Elements and their relations in the Flat RBAC.

3. Research question

While we recognize privacy as a social construct that includes requirements from the patients, care providers, the payers, researchers, and public-health authorities, our focus in this research is on the patient. The goal of our research is to develop and evaluate a request–response decision mechanism for approving or denying access to requested data in a patient's electronic health record, which would meet the expectations of the patient regarding the protection of her privacy and the requirements of other stakeholders involved in patient care, including healthcare providers and administrators, such as medical secretaries. This is a non-trivial goal due to the complexity of the healthcare domain. Specifically, the following aspects make this goal difficult to attain.

(1) A typical EHR can contain large amounts of health-related data items of various types, shapes, and forms.
(2) Many health-related agents are interested in the information contained within EHRs.
(3) Various factors complicate the task of decision-making, including the place and time of the access request, the patient's consent, age, ethnic origin, and the data-requestor's workplace and authorized tasks are involved in the data access scenarios.
(4) Some access-request tasks require different degrees of anonymity in the data to be released, such as de-identifying patient information.

Considering the first three aspects and inspired by the work of Motta and Furuie [21], we propose a situation-based approach for developing a decision mechanism for solving the access control problem within the realm of healthcare. Our research proceeded as follows. We started by eliciting, modeling, and analyzing domain-related information by means of qualitative methods. To this end, we defined a qualitative research question based on the following three aspects: the EHR data, the interested agents, and the various factors that account for the complexity of protecting privacy in electronic healthcare records. The initial research question was formulated as follows:

How can we characterize and model requests for sensitive patient data disclosure with respect to (1) the participating entities and their relations, (2) the context of the request, and (3) the types of requested data?

The second stage of our research involved the development of a conceptual model of the healthcare access control domain. The second research question was therefore:

How can we specify the structure of the requests for disclosure of sensitive patients' data and the behavior of the human roles involved in a way that reflects the use of domain abstractions?

4. Methods

To achieve the research goal, we employed two complementary approaches: qualitative research methods were used to answer the first research question and conceptual modeling for the second one.

4.1. Qualitative research methods

Originally applied in social sciences, qualitative research methods [22] are gaining recognition as being applicable also in the healthcare informatics area [23]. The qualitative research methods we used for eliciting and analyzing data included document studies of EHRs, questionnaires, and interviews.

The three documents we examined were the actual records of three patients at a local hospital, which were obtained from a common EHR system. The records were created when the patients arrived at the emergency room, and they were used and updated while the patients were hospitalized. When the patients left the hospital, their records were archived by the EHR system. We examined the records in order to characterize the types of data items that are recorded while patients visit the hospital and the circumstances at which they are exposed. Interesting findings were discovered while examining these records. For example, in one case the patient was asked to provide medical information regarding his close family members.

The questionnaires (provided in Appendix A), administered to eight patients, were aimed mainly at identifying EHR data items deemed by the patients to be *sensitive* in the context of scenarios of data access.

Semi-structured interviews (see Appendix B) constituted the centerpiece of this study. We conducted 24 such interviews, each lasting between 30 and 45 min. The interviewees were patients and healthcare workers who needed to access EHR data in order to carry out their tasks. These workers included doctors, nurses, a secretary, a health information systems administrator, a dietitian, and a health insurance attorney. The interviewees were asked to describe various data access encounters that take place while they provide service to patients. All the interviews were transcribed.

During qualitative analysis on the transcribed interviews and the questionnaires, we highlighted *informative phrases* (e.g., types of medications, laboratory-test results, and nurse) and grouped them into *categories*. For example, medications and laboratory-test results were categorized as *EHR Sections*, whereas nurse was categorized as *Role*. Additional examples are provided in Tables 1 and 2.

4.2. Object-Process Methodology

During further analysis we identified unique scenarios of request for data access. As noted, a scenario is a process description (in unstructured text) of a patient's data access. We realized that a scenario could be structured into a pattern that we named *Situation*. Consequently, we refer to this stage as *situation-based analysis*. We used a structured analysis method for discovering and specifying the structure of the situation. To this end, we used Object-Process Methodology (OPM) [24], a holistic systems modeling and lifecycle-support approach that integrates the structural, functional, and behavioral aspects of a system in a single, unifying model. The model is expressed bi-modally in equivalent graphics and text with built-in refinement–abstraction mechanism. We have used OPM for conceptualizing and structuring the privacy-related scenarios.

OPM comprises entities and links. The three entity types are objects, processes (both referred to as "things"), and states. Objects are things that exist and can be stateful (i.e., have states). Processes

Table 1
Qualitative and situation-based analysis of Scenario 1 - The family physician

No.	Qualitative analysis		Situation-based analysis		Allowed value
	Informative Phrase	Category	Situation schema element	Super element	
1	My	Relation between the patient and the data-requestor	Data Requestor-to-Patient Relation	Relation	Relation-type: family doctor-of-patient
2	Patient	Patient	Patient	Entity	NA
3	Arrives at	Location	Location patient's attribute	Refineable	NA
		Relation between the patient's primary care clinic and his location	Refineable-Relation	Relation	Relation: Primary Care-Clinic, Location Relation-type: Equal-to
4	Primary care clinic	Workplace of the data-requestor	Workplace data-requestor's attribute	Refineable	NA
		Clinics assigned to the patient	Primary-Care-Clinic patient's attribute	Refineable	NA
5	Document	Action executed by the data-requestor	Action part-of the task	Refineable	Document
6	Encounter	EHR section	EHR Section part-of the task	Refineable	Encounter
7	His	Relation between the patient and the EHR	EHR-to-Patient Relation	Relation	Relation-type: Record-of
8	EHR	EHR	EHR	Entity	
9	Access only from the clinic	Location	Location data-requestor's attribute	Refineable	NA
		Relation between the workplace and the location of the data-requestor	Refineable-Relation	Relation	Relation: workplace, Location Relation-type: Equal-to

NA, not applicable.

Table 2
Qualitative and situation-based analyses of Scenario 2—the medical secretary

No.	Qualitative analysis		Situation-based analysis		Allowed value
	Informative Phrase	Category	Situation schema element	Super element	
1	Secretary	Role of the data-requestor	Role Attribute of the data-requestor	Refineable	Secretary
2	Department of medicine	Workplace of the data-requestor	Workplace attribute of the data-requestor	Refineable	Internal Medicine Unit
3	Hospital	Organization which owns the EHR	Organization, Ownership EHR 's attribute	Refineable	NA
4	In	Relation between medical organization units	Refineable-Relation	Relation	Relation: ownership, workplace. Relation-type: part-of
5	Transfer	Action carried out by the data-requestor	Action part-of the task	Refineable	transfer
6	Discharge-Letter	EHR section	EHR Section part-of the task	Refineable	Discharge-Letter
7	Of	Relation between the patient and the EHR	EHR-to-Patient Relation	Relation	Relation-type:Record-of
8	Patient	Patient	Patient	Entity	NA
9	During the previous three months	Retroactive access-time	Retroactive Access-Time data-requestor's attribute	Refineable	3 months
		Date of issue of a medical section in EHR	Date-Of-Issue EHR Section's attribute	Refineable	NA
		Relation between the date-of-issue and the retroactive access-time	Refineable-Relation	Relation	Relation: Date-of-Issue, Retroactive Access-Time Relation-type: Within
10	The patient was referred	Legal-Authorization	Legal- Authorization	Entity	
		Type of legal-authorization	Legal-Authorization Type legal-author's attribute	Refineable	Referral Letter
		The Medical Unit which issued the legal-authorization	Issued-by Medical Unit legal-author's attribute	Refineable	NA
		Relation between the EHR owner and the creator of the legal-authorization	Refineable-Relation	Relation	Relation: Owner-of-EHR, Issued-by Medical Unit Relation-type: Partof
11	County medical clinic	The Medical Unit to which the legal-author. was addressed	Referred-to Medical Unit legal-author's attribute	Refineable	NA
		Relation between the referral unit and the transferred unit	Refineable-Relation	Relation	Relation: Action, Legal- Author. (referred-to) Relation-type: Equal-to

NA, not applicable.

transform objects: they generate and consume objects, or affect stateful objects by changing their states. Objects and processes are of equal importance, as they complement each other in the single-model specification of the system. An online quick guide to Object-Process Methodology symbols and semantics (some of which are used in the OPM models in this paper) is available at http://www.technion.ac.il/~dizza/QuickGuide.doc. As depicted in the quick guide, links, which are the OPM elements that connect entities, are of two types: structural and procedural. OPM objects relate statically to each other via structural relations, graphically expressed as structural links. The four fundamental structural relations are aggregation–participation, generalization–specialization, exhibition–characterization, and classification–instantiation. Objects can also be structurally related to each other by unidirectional or bidirectional tagged relations, similar to association links in UML class diagrams. Structural relations specify relations between any two objects. Due to the object-process symmetry, they can also specify relations between any two processes. Conversely, proce-

dural links connect a process with an object or an object's state to specify the dynamics of the system. Procedural links include (1) transforming links: effect link, consumption link, result link, and the pair of input-output links, (2) enabling links: agent and instrument links, and (3) control links: event, condition, invocation, and time exception links.

An OPM model consists of a set of hierarchically organized Object-Process Diagrams (OPDs) that alleviate systems' complexity. Each OPD is obtained by in-zooming or unfolding of a thing (object or process) in its ancestor OPD. One or more new things (objects and/or processes) can be specified within a thing in an OPD that was refined from a higher-level OPD. Copies of an existing thing can be placed in any diagram, where some or all the details, such as object states or links to other things, which are unimportant in the context of the diagram, can be hidden. It is sufficient for some detail to appear once in some OPD for it to be true for the system in general even though it is not shown in any other OPD. OPCAT [25] is a software environment that supports OPM-based system development and lifecycle management.

5. Results

The scenarios that we identified included the same types of categories detected in the qualitative analysis. Moreover, there were categories that appeared in each one of the scenarios and some that were optional. Some categories appeared only once per scenario while others appeared several times. Realizing that a scenario could be structured into lower-level building blocks, we defined a *Situation* as a structured representation of a scenario of data access request. Situation is the basis for our proposed access control model, hence the name *Situation-Based Access Control*, or *SitBAC* for short.

From the scenarios that we had identified and collected, we discovered 129 distinct situations in this situation-based analysis. Most of the situations, 99 in total, were simple, while the remaining 30 were complex. To demonstrate our approach and present the results, we describe the process of analyzing two scenarios, which were extracted from two interviews. Following the scenario descriptions in Section 5.1, we present in Section 5.2 the categories identified in these two scenarios. In Section 5.3 we introduce the generic SitBAC model. The OPM models of the two scenarios are provided in Section 5.4.

5.1. The two example scenarios

One of the two scenarios, described first, was relatively simple, while the other was more complex.

5.1.1. Scenario 1—the family physician (the simple scenario)

Our simple scenario is extracted from an interview with a family physician, Dr. S. The scenario concerns an appointment of a patient with his family doctor, in which the patient complains about a medical problem. Dr. S. is obliged to document this encounter in the EHR, and this is done under the EHR Encounters section. If the medical complaint is followed by a new diagnosis, Dr. S. generates a new record under the Diagnosis section in the EHR. Dr. S. may prescribe some medications to the patient, which will result in a new record in the EHR under the Medications section. Our scenario relates only to the encounter documenting task, as described by Dr. S. Following is an excerpt from the interview with Dr. S., which provided the basis for the simple scenario (the numbers are used later on for analysis purposes).

"When *my*[1] *patient*[2] *arrives*[3] for an appointment at the *primary care clinic*[4], I need to *document*[5] this *encounter*[6] in *his*[7] *EHR*[8], which I can access *only from the clinic*[9]."

5.1.2. Scenario 2—the medical secretary (the complex scenario)

The complex scenario is taken from an interesting interview with Ms. R., a secretary of a hospital department head. Among her other duties, Ms. R. is responsible for preparing the medical documents necessary for a follow-up encounter between a patient, who has recently been hospitalized in the department, and an expert physician, to whom the patient is referred, and whose clinic is located at the hospital's county. Ms. R. generates the documents and transfers them (by fax or email) to the clinic prior to the encounter. The documents include a discharge letter with the patient's current and past diagnoses, a list of currently prescribed medications, laboratory and imaging test results, and results of pathology test, if done. Following is an excerpt from the interview with Ms. R., which provided the basis for the complex scenario.

"As the *secretary*[1] *of the department of medicine*[2] *in*[4] the *hospital*[3], I can *transfer*[5] a *discharge-letter*[6] of[7] a *patient*[8] who was hospitalized in my department *during the previous three months*[9] to a *county medical clinic*[11] if *the patient was referred*[10] to one of the physicians who works in this clinic."

5.2. Categories identified by the qualitative analysis

Table 1 presents the outcome of the qualitative analysis and the situation-based analysis of the simple scenario. The second and third columns represent the outcome of the qualitative analysis. Each informative phrase that we found is listed in the second column, while the third column presents the category of the informative phrase. The other columns represent the outcome of the situation-based analysis, which is described in detail in Section 5.3. Table 2 describes the same for the second scenario.

5.3. The SitBAC model

As explained in Section 4, the qualitative analysis resulted in a list of categories, but it did not capture the structure of a scenario. Analyzing the scenarios further to define their structure, we observed the following scenario characteristics:

(1) A scenario is a self-contained unit that exhibits certain functionality and conforms to a repeated pattern.
(2) The functions of an organizationally defined role are not always carried out by an entity assigned with that role. For example, a secretary may need to use her boss's password in order to prepare a patient's file.
(3) Depending on the circumstances, the same task may require different authorizations. For example, a hospital physician may access a patient's previous test results only if the patient is hospitalized in the physician's department, otherwise he is not allowed to do so.

Based on these observations, we decided to structure access scenarios as patterns revolving around *tasks* rather than around roles. Under different circumstances, a task can be executed by different entities in the organization. We call the generic pattern *Situation Schema* and specific scenarios that conform to it—Situation Instances, or simply *Situations*. The situation schema is analogous to an XML schema and a situation is analogous to an XML document, which contains values and is validated against its schema.

5.3.1. Situation schema elements

As presented in Table 1 and Table 2, the output of the qualitative analysis was a collection of informative phrases, grouped into categories. The categories were the basis for deriving the *situation schema*, shown in Fig. 2. The situation schema is a pattern that con-

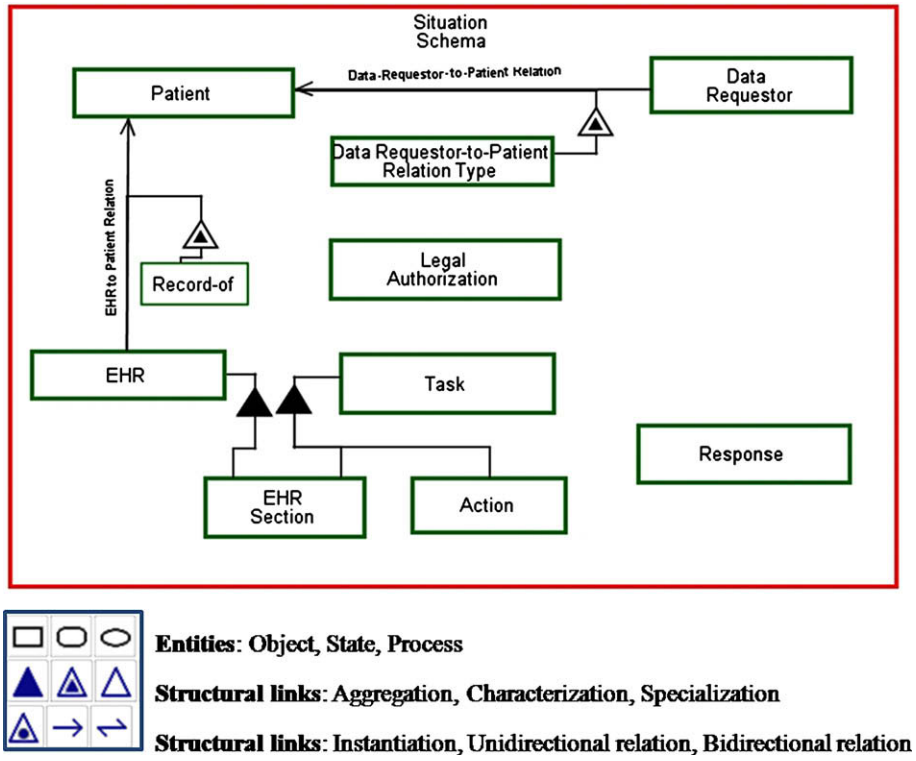

Fig. 2. The Object-Process Diagram of the situation schema (top-level view).

sists of a combination of situation schema elements, belonging to one of the following *super element* types:

(1) *Entity:* A significant object in the data access scenario.
(2) *Refineable:* An object which refines another super element, as explained below.
(3) *Relation:* A link between two entities or between two refineables.

The classification of each super element follows.

5.3.1.1. Entity. There are six types of entities in our model. These entities appear in almost all the scenarios that we have encountered:

(1) *Data-Requestor:* The (human) entity requesting access to the patient's data.
(2) *Patient:* The (human) entity who is the subject of the requested data.
(3) *EHR:* The Electronic Health Record where the patient's data is maintained.
(4) *Task:* The operation on the data that the Data-Requestor wishes to carry out.
(5) *Legal-Authorization:* A legal document authorizing the requested Task.
(6) *Response:* The data access decision made with respect to the situation in question.

5.3.1.2. Refineable. In OPM, objects, relationships, and processes (the latter are not used in our Situation model) may be complex or simple. Complex objects and relationships are refined by other objects, which we refer to as *refineables*. A refineable is of one of the following three types:

(1) *Part:* The refineable relates to the refinee with an aggregation–participation relation. For example, the entity EHR is a refinee that consists of several EHR Sections, so each EHR Section is a refineable of type part.
(2) *Attribute:* The refineable relates to the refinee with an exhibition–characterization relation. For example, the entity Patient is a refinee that exhibits the Age attribute, so Age is a refineable of type attribute.
(3) *Specialization:* The refineable relates to the refinee with a generalization–specialization relation. For example, the entity Organization is a refinee that specializes to Medical Site, so Medical Site is a refineable of type specialization.

The refineables are structured hierarchically. A refineable can be assigned with an *allowed value*. The allowed values are grouped into sets, which are assigned to the corresponding refineable.

5.3.1.3. Relation. A relation can exist only either between two entities or between two refineables. The relation is refined by a *relation type*, which indicates the nature of the relation, e.g., `Family-Doctor-of-Patient`. The two relation kinds are discussed below.

(1) *Entity-to-entity relation:* Relations between pairs of the six entities. Most of the possible pair combinations are meaningless (e.g., the Response entity and the EHR entity cannot establish a relation). There are two possible entity pairs between which an entity-to-entity relation exists:
    (i) The *EHR–Patient* entity pair, in which case the relation type is assigned with `record-of`.
    (ii) The *Data-Requestor–Patient* entity pair, in which case the relation type is assigned with `Family-Doctor-of-Patient`, `Gynecologist-of-Patient`, etc.

(2) *Refineable-to-refineable relation:* Relationships between refineables exist only when the two refineables share a common ancestor in the refineables' hierarchical structure. For example, the patient's *Location* refineable can be related via an Equal-to relation type to a data-requestor's *Workplace* refineable, since both refineables share a common ancestor:

Medical Site. The relation type can have one of the following values: `equal-to`, `different-from`, `greater-than`, `less-than`, and `within`.

Further details are provided in Section 5.3.4.

5.3.2. Situation schema

A *Situation Schema* is a generic pattern corresponding to data access scenarios. It consists of situation schema elements—entities, refineables, and relations, and it abides by the following rules:

(1) The task entity and the response entity are mandatory elements that must be present at each situation schema.
(2) The situation schema explicitly defines the set of refineables assigned to each entity and to each relation, so the entities and the relations that participate in a situation can only be refined by a subset of its corresponding refineables set.

Fig. 2 is a top-level view of the situation schema. Unfolding the entities and the relations exposes the refineables associated with the entity or the relation. Fig. 3 is an example of unfolding of the data-requestor entity of Fig. 2.

5.3.3. Situation

Each patient's data access scenario can be modeled as a situation. Each situation is an instance of Situation Schema. As such, it includes:

(1) at least the mandatory entities,
(2) for each entity, a subset of its refineables set, and
(3) optional relations between entities and between refineables.

For each refineable participating in the situation, at least one of the two following requirements must be fulfilled:

(1) An allowed value is assigned to the refineable. This value is taken from the set of allowed values for that refineable in the situation schema. Fig. 4 illustrates allowed values that can be assigned to the EHR Section refineable, showing that it can assume one of the following three possible values: `Identification Section`, `Demographic Section`, and `Medical Section`.
(2) The refineable participates in a refineable-to-refineable relation instance with another refineable. The relation complies with the rule that the two participating refineables must share a common ancestor in the refineables' hierarchical structure.

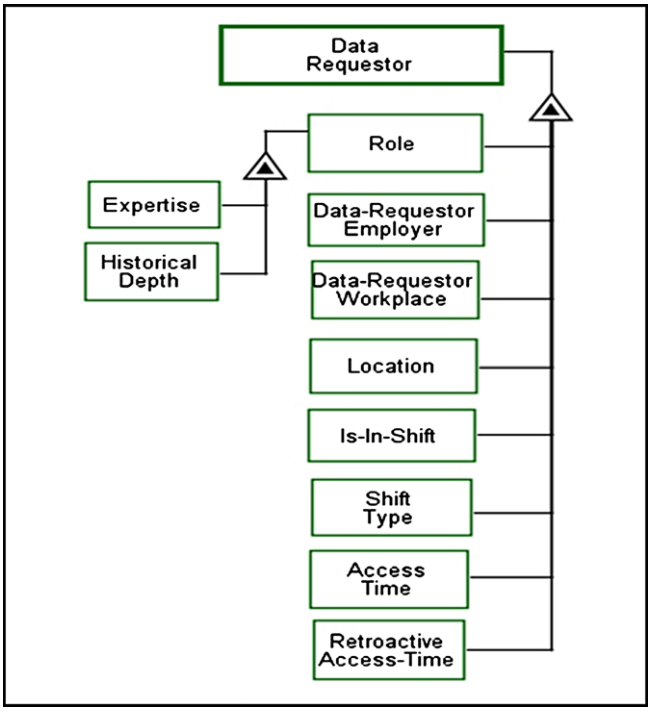

Fig. 3. An unfolded view of the data-requestor.

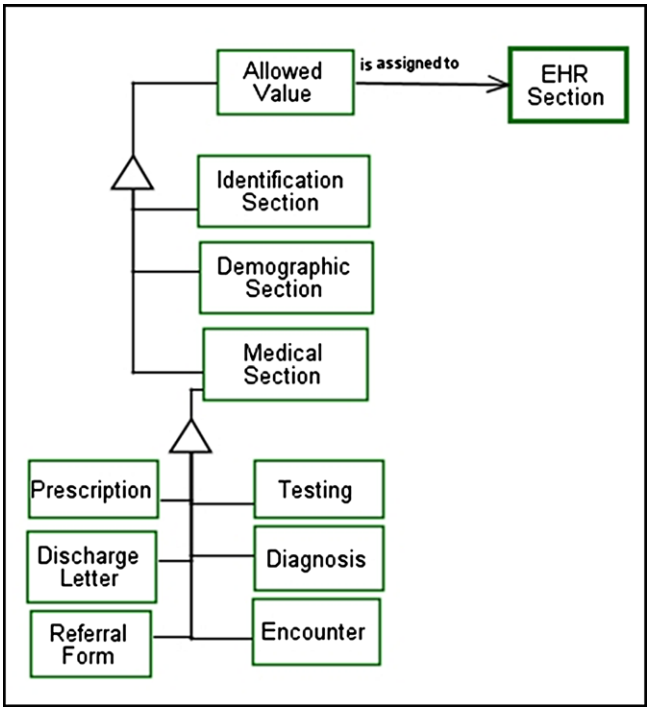

Fig. 4. Allowed values, assigned to EHR Section.

Based on the above definitions, we can now refer more precisely to columns 4–6 in Tables 1 and 2, which represent the outcome of the situation-based analysis. The fourth column lists the situation schema element. The fifth column lists the mapping of each situation schema element to a super element, and the sixth column provides the related allowed value (if it exists). The OPM model associated with the first scenario is provided in Fig. 5, while the OPM model for the second one is shown in Fig. 6 and in Fig. 7.

5.3.4. Additional details on situation schema elements

In this section, we provide more details about the entities, their refineables, possible relations, and other findings that we concluded from the data that we had collected and analyzed.

Patient has attribute refineables such as *Age*, *Gender*, *Location* (e.g., home, medical clinic), *Workplace*, *Ethnic Origin*, and an *Is-a-Celebrity* Boolean attribute, indicating whether the patient is a publicly known figure.

Data-Requestor has attribute refineables such as *Role*, *Workplace*, *Employer*, *Location*, *Access-Time*, *Shift-Type*, and the Boolean attribute *Is-in-Shift*. Location might be different from the data-requestor's usual workplace refineable in some cases. For example, the data-requestor who works in the pediatric department at a hospital might be trying to access a patient's data while he is located in the emergency room.

EHR is composed of many *EHR Section* part-of refineables. In order to demonstrate the SitBAC model principles we adopted a simple model of an EHR, presented by Motta and Furuie [21]. EHR

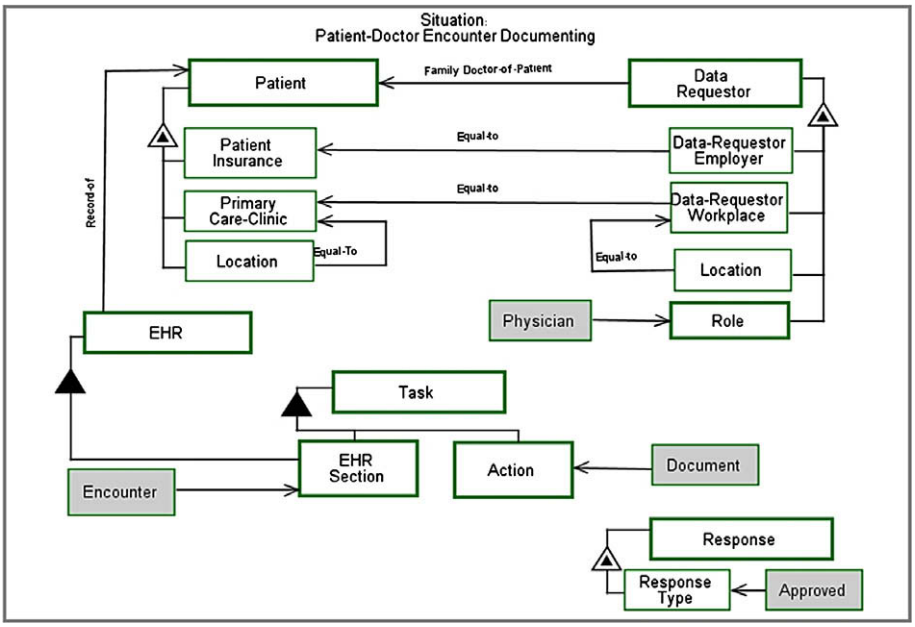

Fig. 5. Object-Process Diagram of the family physician situation.

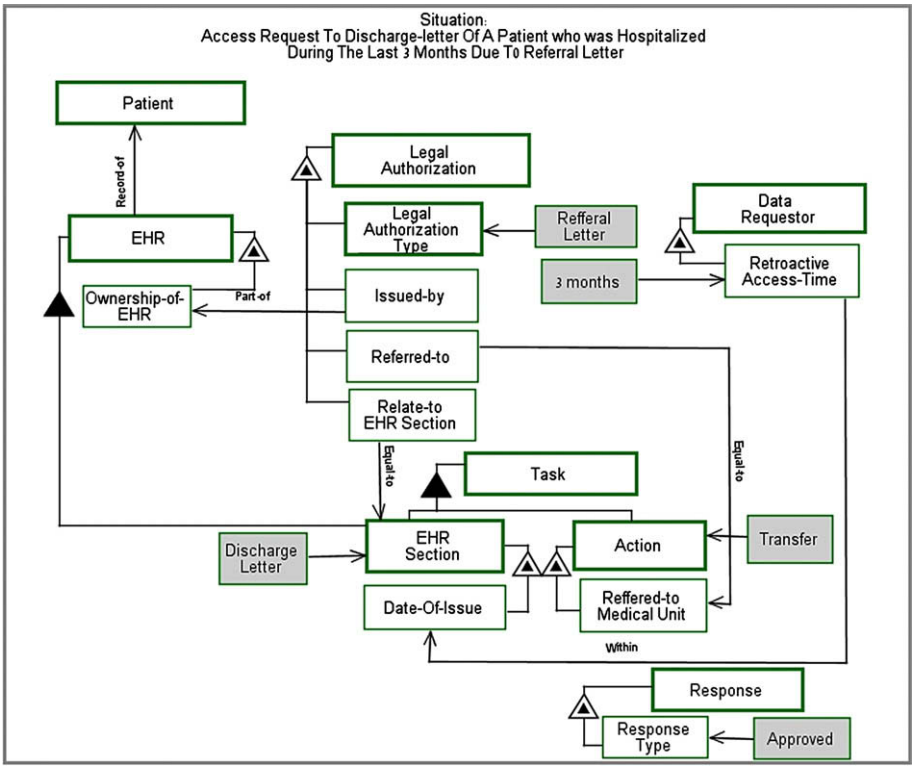

Fig. 6. Object-Process Diagram of the medical secretary situation—the single version.

models such that of HL7 [15] are much more complex than the one we present in this paper. As noted, according to our model, the EHR Section refineable can assume one of the following allowed values: `Identification Section`, `Demographic Section`, and `Medical Section`. The `Medical Section` allowed value preserves most of the patient's sensitive data. Specialized medical sections that were frequently mentioned in the questionnaires and in the interviews are `Prescription`, `Diagnosis`, `Test`, and `Encounter`.

Task is composed of two part-of refineables: an *Action* and an *EHR Section*. Action can be assigned with one of the following allowed values: `View` (which is the most frequent action), `Document`, `Update`, and `Delete`. The Action refineable can be further refined, for example by the *Referred-to-Role* or *Referred-to Medical Unit* attribute refineables.

Legal-Authorization represents the need for formal legal-authorization, especially when the scenario describes an access request from an entity not affiliated with the organization maintaining

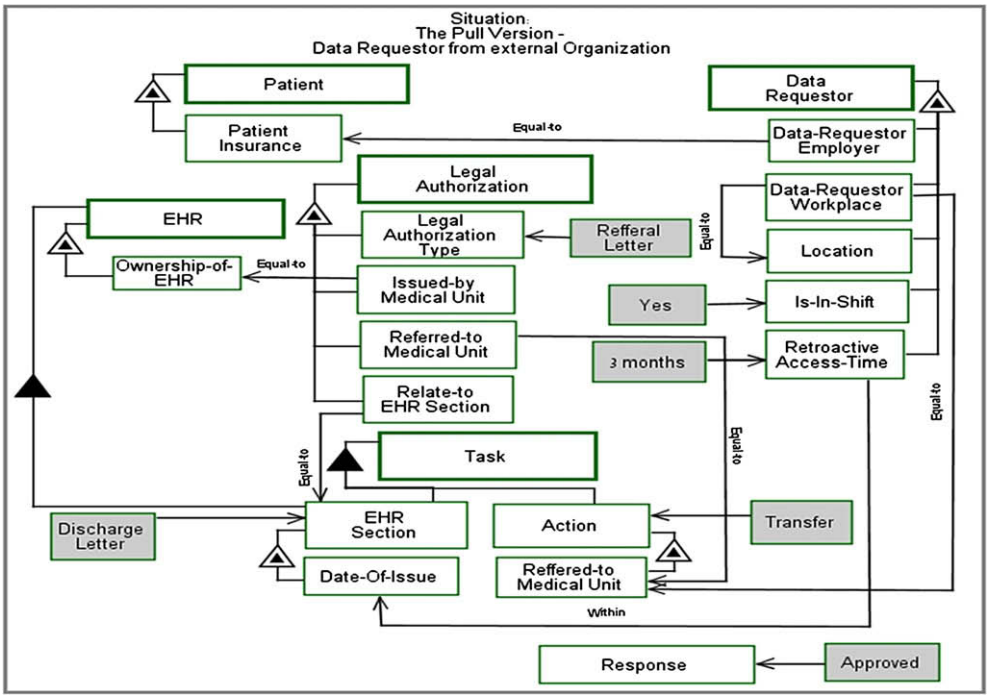

Fig. 7. Object-Process Diagram of the medical secretary situation—the pull version.

the patient's data. Legal-Authorization is refined by attribute refineables, including *Legal-Authorization type, Issued-by Medical Unit, Referred-to Medical unit*, and *related-to EHR Section*.

Response is refined by one attribute refineable named *Response Type*. The response type can be assigned with one of the following three allowed values: `Approved`, `Denied` or `Unknown`. The `Denied` and the `Approved` allowed values are used by the situation's creator to indicate whether a request for data access is approved or denied.

Data-Requestor-to-Patient Relation is a type of entity-to-entity relation. It indicates a long-term relation between the two entities participating in the relation, where the nature of the relation is expressed via the relation type. In principle, there might be a short-term relation between the patient and the data-requestor. However, based on our qualitative findings, long-term relations are usually required to be reflected within the situation instances. Moreover, a long-term relation (e.g., between a patient and her family doctor) is usually expressed within an organization's database. Thus, long-term relations are items that can be validated against the organization's database. We identified few long-term relations, such as a patient's family doctor, a patient's gynecologist, and a patient's psychologist.

Relation between refineables indicates a sustainable relationship between two refineables that characterizes the situation. Examples for such relations include *data-requestor's location equal-to data-requestor's workplace*, and *patient's insurance equal-to data-requestor's employer*.

5.4. Modeling scenarios as situations via OPM

In this section, we use OPM to model the situations associated with the family physician scenario and the more complex medical secretary scenario. Each of the situations is specified via an OPD. In order to keep the OPDs simple and clear, the relations in the OPDs are expressed via links, where the allowed values, which are assigned to the relation type, appear above the link. Also, when a refineable is assigned with an allowed value, the value appears as a gray-colored object, which is connected to the refineable through a link.

5.4.1. The family physician scenario model

Fig. 5 presents the situation of the simple scenario, documenting the encounter between a doctor and her patient. We can observe that the EHR record participates in a refineable relation with the patient, where the relation type is assigned with a `Record-of` allowed value. The patient has a relation with the data-requestor typed as `Family Doctor-of-Patient`. As the OPD specifies, the data-requestor performs an encounter-documenting task, in which the EHR Section of the task is assigned with an allowed value `Encounter` and the task's action is assigned with an allowed value `Document`. The fact that the data-requestor needs to be located in his workplace while executing the task is expressed via the *Location* attribute refineable of the data-requestor, which participates in a refineable relation with the data-requestor's *Workplace* attribute refineable. The relation type is assigned with an `Equal-to` allowed value. The patient, in this case, also needs to be located in the data-requestor's primary care clinic. This fact is expressed via the relation between the patient's *Location* attribute refineable and the patient's *Primary-Care-Clinic* attribute refineable, where the relation type is assigned with `Equal-to`.

The OPD includes few refineables that are not explicitly mentioned in the scenario's text. For example, the fact that the patient is insured by an insurance company, which happens to be the data-requestor's employer, is expressed via a refineable relation. This fact is a known implied assumption. Thus, the OPD can be validated against the situation schema, shown in Fig. 2. As noted, the situation schema consists of situation schema elements only, while the situation, represented in the OPD of Fig. 5, includes also allowed values. Note that the situation complies with the requirement that each of its refineables is assigned with an allowed value or participates in a refineable-to-refineable relation with another refineable, where the two refineables share a common ancestor within the refineables' hierarchic structure.

5.4.2. The Medical Secretary scenario model

The Medical Secretary scenario describes the disclosure of a discharge letter of a patient who was hospitalized, for the sake of

treatment continuance. This scenario can be interpreted in two ways: the *push* version and the *pull* version. In the push version, the hospital department's secretary executes the discharge-letter transfer task, while in the pull version, the data-requestor can be a medical staff member who works in the county medical clinic (e.g., a secretary or a nurse) to which the patient is referred, who is authorized to perform the transfer task. These two possible interpretations demonstrate that the same task can be performed by different roles in different organizations. Table 2 assumes the push version.

In essence, the pull version and the push version are derived from a single scenario, which does not involve a role for the data-requestor. SitBAC can express such a role-less scenario within its model, while RBAC cannot, as in RBAC two rules (one for pulling and one for pushing) need to be been defined, scoped with the two corresponding roles.

The single conceptual situation is illustrated in Fig. 6. Two important observations are noted with respect to this situation: (1) the data-requestor entity is missing its *Role* refineable, and (2) the focus of this situation is on the task (transferring the discharge letter) and on the legal-authorization, indicating that a formal authorization is mandatory in order to perform the task. In this case, the `Referral Letter` is an allowed value that is assigned to the *legal-authorization type*.

Examining Fig. 6, we observe that the EHR relates to the patient, and the relation type is assigned with the `Record-of` allowed value. The EHR has the attribute refineable *Ownership-of-EHR*, which indicates its ownership. The *Ownership-of-EHR* attribute refineable has a relation with the *Issued-by* attribute refineable of the legal-authorization entity, where the relation type is assigned with a `Part-of` allowed value. The last relation indicates that the legal-authorization has to be issued by the organization who owns the EHR containing the discharge letter. The data-requestor is refined with a *Retroactive Access-Time* attribute refineable, which indicates the period of time during which the data-requestor is allowed to retroactively access the patient's data. This attribute has to be assigned with the allowed value `3 months`, and has to relate with relation type `Within` to the *Date-of-Issue* attribute refineable of the EHR Section. Finally, the *EHR Section* is assigned with a `Discharge Letter` allowed value.

If we wish to refine this situation by explicitly specifying the data-requestors that are allowed to execute this transfer task and their refineables, we can create two more situations that inherit from the single-version situation, presented in Fig. 6. The inherited situations are copies of the single version, which, in addition, include the required data-requestor. Fig. 7 illustrates the pull version, which includes all the elements that appear in the single version along with additional refineables of the data-requestor. The patient is refined by a *Patient Insurance* attribute refineable, which relates to the data-requestor's *Employer* attribute refineable. The relation type is assigned with the `Equal-to` allowed value.

6. Discussion

Our work has been motivated by the need to preserve privacy in the healthcare domain. Accordingly, our leading design principle was that access to specific sensitive data be allowed based on circumstances that match predefined patterns. This guideline has led us to the design of the SitBAC model, consisting of a *Situation Schema* and a collection of *Situation Instances*. Our model enables expressing access control restrictions by structurally representing the scenarios of request for data access as situations with allowed values assigned to them that can be validated against the situation schema. A collection of situations can be used to formulate the organizational policies and rules of access control permissions, and may be dynamically changed as the organization reviews its policies.

The SitBAC model was established after we elicited and analyzed data through qualitative research methods, which include document (EHR) studies, questionnaires, and interviews. The interviews, especially those conducted with interviewees who provide paramedical or administration services, provided the most significant part of the data. This is so because interviewees who provide medical services directly (mainly doctors) need frequent access to all the data in the EHR in order to treat their patients, usually independently of the circumstances. Thus, creating the situation for such cases is relatively simple. However, paramedics or administration health workers do not usually need access to the entire patient's data set on a regular basis while they provide service to the patient. When access to data is restricted, the circumstances play an important role in the process of decision-making regarding access permission. This, in turn, results in complex situations.

While asking our interviewees and respondents to point out what health data they considered to be sensitive, we got many answers reflecting different views concerning data sensitivity. Nevertheless, most of the respondents indicated gynecological data, mental health data, and data about sexually transmitted diseases as highly sensitive categories. This rating is indeed reflected in EHR systems, where gynecological units and mental health units usually cannot be accessed by users outside of these units, unlike the norm in other departments. However, there is no direct reflection of the data items that we identified as sensitive in the SitBAC model. This is because the model can protect any data item that the patient prefers not to reveal by adding a legal-authorization in a form of patient consent. Such a situation may be defined only if proper legislation exists for such patient requirement.

We noticed that the interviewees stipulated specific actions as they were describing their daily tasks (e.g., ordering an MRI, transmitting a test lab by email, etc.). As a result, we got detailed descriptions of the tasks carried out by the different roles in the organization within each scenario. Admittedly, some of the scenarios described people in the organization who were using other people's passwords in order to perform data access tasks they were asked to perform in spite of the fact that the information system in their organization was supposed to prevent them from doing so. This adds value to our research, as our model can also help the organization to represent, identify, monitor, and enforce its internal regulations regarding tasks involving data access on a task-related rather than a role-related basis.

During the process of conceptualizing and modeling the SitBAC using OPM, we have noticed that the process starts by asking *what* information needs to be disclosed to whom rather than *who* has to do it. This is a fundamental difference between our approach and that of RBAC. In RBAC, when a user requests data access, his *role* is used in order to retrieve the access privileges defined for that role. In our approach, the role is only one of the factors that are considered for obtaining a decision for a data access request, and even not a mandatory one, as exemplified in Fig. 6.

Being a generalization or a superset of RBAC, the SitBAC model can express RBAC rules, because each one of the RBAC elements (Users, Roles, Objects, Operations, and Permissions) is expressed by one of the SitBAC elements. SitBAC also enables expressing scenario factors such as the patient's age or the data-requestor's workplace, which need to be considered while health data is about to be exposed. Such factors were discussed earlier by Motta and Furuie [21], who proposed to extend the RBAC model by including contex-

tual variables, such as patient's status (e.g., inpatient or outpatient) and user's location-of-access (e.g., emergency room or clinics area). With respect to our research, the work of Motta and Furuie has three main limitations:

(1) Since it extends the RBAC model, the role element remains mandatory and it serves as the mediator between the user and the task.
(2) The Motta and Furuie extension was restricted to the description of available cases within one hospital, while the vision of our model is to enable expressing scenarios of request for data access executed by one organization while the required data is maintained by another organization. We have modeled such a case via analyzing the complex scenario in Section 5.4.2. Many of the complex scenarios that we collected involve interoperability. This occurs in cases where a consortium of health organizations shares patients' data. Our model can express scenarios which involve indirect access requests, which might occur if the data-requestor has no direct relationship with the patient. For example, a researcher who uses patient data for his research needs permission from a Legal-Authorization entity. OPDs of this example and of another example can be found at http://www.technion.ac.il/~dizza/IndirectAccessRequestsSituations.doc.
(3) The Motta and Furuie extension was implemented via logical, cryptic-looking rules (e.g., "aPatCod in patCtx.in_patients"), where the contexts are semi-structured, while we intend to implement our model via a formal ontology. The situation schema elements and the allowed values in our planned implementation will be part of the ontology, and therefore fully structured.

Our model is reinforced by the work of Patel et al. [26], who reported about how experts' decision-making processes are based on understanding a given situation and acting upon their experience, relying on strategies of situation recognition.

One limitation of our research is that since the qualitative part of the research was patient-centric, it did not include all the possible stakeholders, such as employers or community service providers, with their different goals. The SitBAC model is an outcome of analyzing the 129 scenarios that we had elicited via the qualitative research. Since we did not interview all the possible stakeholders, some elements may be missing from our model. However, as explained above, since SitBAC is a superset of RBAC, it is capable of representing other scenarios, such as the indirect access scenarios discussed above.

We are currently working on the next stage of our research, which includes a formal representation of the SitBAC model and its situations as a knowledge base via OWL-based ontology. Another aspect of our ongoing research is creating a situation-similarity algorithm, which gets as input an access request in a form of a situation instance, searches the knowledge base for a similar situation, and based on a similarity measure, approves or declines the access request. We plan to use a large medical center as a case study for creating an organizational situation knowledge base. This case study would also help us estimate the number of situations needed for establishing organizational access control policies. In order to represent the current access control authorizations in organizations that use RBAC, we will need to define at least one default situation for each role. We will do so by specifying the role of the Data Requestor entity and the sections of the EMR. Then, we will elicit from the organization more specific situations for which non-default access control policies would be defined.

7. Conclusions

We have developed and exemplified a Situation-Based Access Control (SitBAC) model, which is designed for expressing scenarios of request for patient's data access as a basis to preservation of the patient's privacy. The model is generic and can be adapted to domains other than healthcare, e.g., the banking domain. The strengths of our model are in its ability to (1) structurally specify scenarios of patient's data access via situation models, (2) represent a situation where the data-requestor definition is partial (e.g., the role is missing), and (3) represent scenarios where the data-requestor and the required data do not belong to the same organization. Apart from protecting patients' privacy, SitBAC could also potentially be used by organizations to assess the adherence of its employees to regulations concerning data access tasks.

Acknowledgment

We thank our anonymous reviewers for their helpful and contributing comments.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jbi.2008.03.014.

References

[1] Jones KS. Privacy: What's different now? Interdiscip Sci Rev 2003;28(4):287–92.
[2] Morgan R. Community Attitudes to Privacy: Office of the Australian Federal Privacy Commissioner; 2001.
[3] Schauer F. Free speech and the social construction of privacy. Social Res 2001;68(1):221–32.
[4] Shapiro B, Baker CR. Information technology and the social construction of information privacy. J Account Public Policy 2001;20(4–5):295–322.
[5] Federrath H. Privacy enhanced technologies: methods—markets—misuse. Trust, privacy and security in digital business, LNCS; 2005. p. 1–9.
[6] US Department of Health and Human Services. Summary of the HIPAA privacy rule, <http://www.hhs.gov/ocr/privacysummary.pdf/>; 2003 [last accessed on 1/24/2008].
[7] Workgroup for Electronic Data Interchange. HIPAA Glossary, <http://www.wedi.org/snip/public/articles/HIPAA_GLOSSARY.PDF/>; 2001. [last accessed on 1/24/2008].
[8] McMurry AJ, Gilbert CA, Reis BY, Chueh HC, Kohane IS, Mandl KD. A self-scaling, distributed information architecture for public health, research, and clinical care. J Am Med Inform Assoc 2007;14(14):527–33.
[9] Welch CA. Sacred secrets—the privacy of medical records. N Engl J Med 2001;345:371–2.
[10] Ravera L, Colombo I, Tedeschi M, Ravera A. Security and privacy at the private multispecialty hospital istituto clinico humanitas: strategy and reality. Intl J Med Inform 2004;73(3):321–4.
[11] Sweeney L. K-anonymity: a model for protecting privacy. Intl J Uncertainty Fuzziness, Knowledge Syst 2002;10(5):557–70.
[12] Sandhu RS, Coyne EJ, Youman CE. Role-based access control models. IEEE Comput 1996;29(2):38–47.
[13] Sandhu R. The NIST Model for role-based access control: toward a unified standard. In: Proceedings of the fifth ACM workshop on role-based access control; 2000. p. 47–63.
[14] Shabo A. The implications of electronic health records for personalized medicine. Pers Med 2005;2(3):251–8.
[15] Health Level Seven. HL7 Standards, <http://www.hl7.org/Library/standards.cfm/>; 2007 [last accessed on 1/24/2008].
[16] Committee European Normalisation. Health Informatics—Electronic healthcare record communication—Part 1: Extended architecture, ENV13606-1, CEN/TC251 Health Informatics Technical Committee; 2000. Available from <http://www.centc251.org/>.
[17] DICOM. Digital imaging and communications in medicine website, <http://medical.nema.org/>; 2000 [lastaccessed1/24/2008].
[18] BBBIJoMI. Authorization and access control for electronic health record systems. Intl J Med Inform 2004;73(3):251–257.
[19] France R. Security of health care records in Belgium application in a university hospital. Intl J Med Inform 2004;73(3):235–8.
[20] Haaka Mvd, Wolffa AC, Brandnera R, Dringsb P, Wannenmacherc M, Wetter T. Data security and protection in cross-institutional electronic patient records. Intl J Med Inform 2003;70(2–3):117–30.

[21] Motta G, Furuie S. A contextual role-based access control authorization model for electronic patient record. IEEE Trans Inform Technol Biomed 2003;7(3):202–7.
[22] Myers MD. Qualitative research in information systems. MIS Q 1997;21(2): 241–2.
[23] Pope C, van Royen P, Baker R. Qualitative methods in research on healthcare quality. Qual Safety Health Care 2002;11(2):148–52.
[24] Dori D. Object-Process Methodology—a holistic systems paradigm. Berlin, Heidelberg, New York: Springer-Verlag; 2002.
[25] Dori D, Reinhartz-Berger I, Strum A. Developing complex systems with object-process methodology using OPCAT. ER, 2003. p. 570–572.
[26] Patel VL, Kaufman DR, Arocha JF. Emerging paradigms of cognition in medical decision-making. J Biomed Inform 2002;35(1):52–75.