Ontological privacy support for the medical domain

Michael Hecker, Tharam Dillon
eXel Lab[1], University of Technology, Sydney
{mhecker,tharam}@it.uts.edu.au

Abstract

This work uses an ontological approach for the medical domain to derive the level of privacy for concepts specific to that domain. The authors have previously produced a generic ontology for privacy. In this paper, we describe a mapping from the general (abstract and application domain independent) privacy ontology to a domain specific level for the medical domain and how the general concepts influence the specific instances, while medical terms and concepts provide actual values for privacy principles that essentially create the level of privacy a user experiences.

Keywords: privacy, ontology, healthcare, HealthConnect, health records

1. Introduction

Privacy in computing and communication has many aspects and issues like legislation, technologies or user perceived privacy, which is the aspect of privacy a user experiences for data related to him or her. The "real" world as opposed to the digital world has to deal with privacy issues as well, but as usually more effort is required to gather and less effort to secure this information. It is not seen as such a big issue compared with the ability to collect, store and process information in the digital world. However, legal implications directly apply to both, the real world and the digital world, hence requiring parties to collect and use data carefully on a need to know basis only.

Although privacy issues are ubiquitous amongst almost every domain, actual instances and concepts as well as their influence on the overall level of privacy for "interactions" are naturally domain specific and must be addressed accordingly. In this paper, we describe the methodology to map the previously developed generic ontology [6][1] to the specific ontology for the medical domain.

2. Motivation and issues

2.1 Privacy

Privacy is considered one of the most important issues nowadays with easy collection, aggregation, linkage and storage facilities available. The Internet provides users with the ability to collect, store and share this kind of information easily, but lacks a cohesive structure making it more difficult to link data together by an automated process. However, other vast data sources (e.g. corporate databases) exist that are much more structured and allow their users to generate much clearer pictures about individuals. Therefore, it becomes more and more difficult to control others access to information about oneself.

The concept of privacy seems to be an endogenous conception, as every person has a different idea about what it means and how it should be implemented to achieve it. Therefore, it is necessary to find some common properties to build a basic foundation. Starting with a common dictionary definition, privacy would be "freedom from unauthorized intrusion" [1]. Similar technically imprecise definitions can be found in

[1] http://exel.it.uts.edu.au

other dictionaries such as Oxford English Dictionary or dictionary.com. Therefore, it is necessary to consider the definition of privacy by people experienced in that domain. Naturally, it has evolved over time and started with the expression of "the right to be let alone" [2], expressed by two lawyers in 1890. However, such a definition is also not very precise yet nor very usable nowadays, as one does not necessarily want to be left alone just to "experience" or "have" privacy. A better definition comes from Privacilla [3], a website related to privacy related policies and defines it as "the subjective condition a person experiences when two factors are in place. First, he or she must have the power to control information about him- or herself. Second, he or she must exercise that control consistent with his or her interests and values". This definition describes privacy in a much better and precise way and it sounds more logical to be in control of information related to oneself than just be left alone. A similar statement has been made by the Privacy Commissioner of Canada, defining privacy as "the right to control access to one's person and to information about oneself" [4]. Trying to make it even clearer, privacy is not about information itself, but the control of that information by a cognitive entity, which is related to it. In order to distinguish such information from other "normal" information, we call this type of information, which is about someone and could potentially identify someone, "Personal Identifiable Information" (PII). Hence, if information cannot be linked to or is about a certain entity, which could have potential interests in controlling it, privacy matters usually do not apply.

Previously, privacy protection has been tried to be accomplished by utilising mechanisms that control access to personal identifiable information. However, it is not the data subject, which is the entity the data is related to, to control access to personal identifiable information, but an "authorised entity" controlling or maintaining the system where the data is store. Needless to say that such an "authorised entity" would have a great deal of control over the information, its release and access. Even more, the data subject might not know or even have authorised that entity to regulate access to its data, but just accepted the fact explicitly or implicitly that there is some sort of protection for its data. This can also be seen as an implicit trust in such an entity to do the "right thing" with the information made available to that system by the entity.

2.2 Privacy and policies

As just stated, a person controlling information about others could be of great danger to the privacy of the data subjects it is controlling – remembering that privacy is about the ability of the data subject to control personally identifiable information about themselves. Thus, privacy policies have been established, accepted and are now widely used and are backed up by legislation. This gives users more confidence when providing information about themselves if it is used in a certain way or to inform them at least how it is used. Privacy policies are usually set up and governed by certain rules and regulations that apply in the territory the entity collecting information is located in, leading to different privacy policies in different regions (e.g. privacy legislation in Europe compared to Australia). The problem with privacy policies is their different semantics and their dependency on the domain they are applied to. While privacy policies within the same domain (and possibly region for regulatory reasons) may have similar structures, there is no semantic way of comparing them with each other or even evaluating the level of privacy they try to offer. They may just be (and are often) written in a certain natural language (e.g. English). This obviously creates problems with precision, clarity and interoperability, making it ambiguous for the reader who has to understand it – being a person or software (agent). Different persons would understand a privacy policy differently, depending on the complexity and clarity of the policy and naturally depending on their "knowledge" about privacy and other intrinsic factors (e.g. culture), making it a fuzzy concept, making it hard or even impossible to formulate privacy in a mathematically precise way. Thus, software agents would have even

more trouble "understanding" these kinds of privacy policies due to the lack of precision.

2.3 Privacy on the web and P3P

On the web, privacy policies have been established in a structured way with the introduction of the Platform for Privacy Preferences (P3P) [5]. Basically, this formulates certain statements about how resources (that is personal or other information) are used for what purpose, by whom and with what kind of retention. As P3P is a platform designed for websites, it covers mostly web-specific terms and is specific to that domain only (omitting the fact that extensions are possible, but their actual values are not standardised). Considering the initial issue that not all data is on or accessible via the web (in fact, a majority is actually not), other privacy policies that also cover electronic as well as non-electronic records from the same or other domains are necessary and available to be evaluated in a systematic way (by a human being or preferably automatically by a software agent).

2.4 Privacy Ontology

The basic idea behind these thoughts leads to a specified conceptualisation of the terminology "privacy", omitting internal and personal factors as it is difficult to capture them precisely. The terminology used for such a conceptualisation and formalisation is commonly known as "ontology". An ontology specific to the domain privacy as described in our previous work [6], showing the different concepts and associations between them. This helps to create interoperability as well as allows one to derive the impact or level of privacy a certain "transaction" (digital or non-digital) has upon the data-subject when they agree to enter into it. Furthermore, the data subject is not the only one to benefit from such an evaluation of privacy levels, but the other participants of a transaction as well. They can essentially use the ontology to model their (privacy) policies and procedures to comply with regulations within their domain. Also, it can help system developers that need to implement any privacy functionality or mechanisms by providing a guide to the concepts and what privacy actually refers to, without being an expert in that domain.

Such an ontology needs to be very general and contains only very abstract concepts to make it applicable to all sorts of different domains. However, as every domain also has specific concepts that need to be covered when considering privacy, a general privacy ontology cannot provide that level of detail. Hence, another ontology, mapping from the generic one to a specific domain is required to capture the domain knowledge and add concepts that would influence privacy within that domain [10]. It is highly likely that a person with sufficient knowledge of that domain is required to help create such a specific ontology, as people that may not be familiar with that domain might not cover all the issues involved. A generic privacy ontology would therefore be a template for the specific one, providing the fundamental structure, notions and principles that are ubiquitous to privacy.

Before it is actually possible to think about the specific privacy concepts and issues of a certain domain, e.g. healthcare, it is necessary to begin with a formulation of generic concepts first. These are likely to be domain independent and are abstract enough to support this. Generally, legislative documents provide a solid foundation for those concepts and are usually covered by the individual Privacy Acts of different nations. Privacy notions and concepts are specified by the "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data" [7] and used by us as one of the actual sources, as privacy legislation in the European Union is more advanced (more protective) than in many other countries. However, a more comprehensive and concise guide of those rules has been compiled by the PRIME [8] and PISA [9] projects describing the essential principles involved in the process of privacy. Issues and principles that have been found are: a) Intention and notification, b)

Transparency, c) Finality principle, d) Legitimate grounds of processing, e) Quality, f) Data subject's rights, g) Processing by a processor, h) Security, i) Accountability, j) Consent, k) Limited Linkability (e.g. data mining), l) Openness, m) Anonymity and n) Transfer of personal data outside the EU (in general to countries with different privacy protection laws).

2.5 Privacy issues in the medical domain

As every domain has specific issues and additional concepts with privacy, so does the medical domain. Generally speaking, privacy is about health records, both electronic and non-electronic ones. However, the medical area deals with very sensitive information and has very specific issues no other domain comprehends. These issues also greatly influence the importance of certain privacy principles, for example, security could be considered less of an issue in the health domain than consent. The reason for that is the *protection of vital interests of a data subject*. This means that health records cannot be made totally secure (e.g. encrypted) in a way that only the data subject can access it, as emergency situations may require access to those records when the data subject cannot authorise access (e.g. the data subject is unconscious and his or her life is at stake). Consent on the other hand is much more important here, as it is vital for the data subject to specify what entities may access the data under what circumstances (emergencies excepted). Needless to say, consent needs to be backed up by a strong legislative framework and an intact implementation, but the influential character of those two principles is different from other domains.

3. Ontological privacy support

This section gives a brief overview of the generic privacy ontology as outlined in [6] to show the different concepts used and how they are linked to each other to support privacy, as well as the different privacy principles and their influence on the overall level of privacy.

Going back to the definition of privacy, the general idea is to control the access and use of personal identifiable information by the data subject or an entity authorised by them. Therefore, the ontology is based around the concepts of "Data Subject", "Resource" and "ResourceUser" that accesses this "Resource". Every concept of the generic privacy ontology has certain attributes, which may or may not contain actual values that describe how the concept influences the overall level of privacy. Actual values may not be possible (as just stated), as it is domain unspecific and therefore not necessarily known how big the actual impact may be. However, a relative value is assigned, which describes how a particular principle within a concept relates to principles in related concepts. For example, the concept of "Resource" can have different types of "ResourceElement", which can be of type "IdentifyingResourceElement", "Non-IdentifyingResourceElement", "AnonymousResourceElement", or "PseudoAnonymousResourceElement". In this case, an AnonymousResourceElement would have a higher relative value for the principle of anonymity (m) than the concept of PseudoAnonymousResourceElement. The actual values are not determined or used here, as they are to be defined by a domain expert, when creating an actual instance of the generic ontology for a specific domain (e.g. healthcare). Hence, a domain expert would create actual subtypes or subconcepts of a certain concept and assign actual values as well as their influential "character" on the overall level of privacy. Due to the nature of privacy, these values are all fuzzy in the sense that there is no mathematically precise definition on how a concept is assessed and values are assigned to the principles, but the domain expert would have sufficient knowledge to specify at least fuzzy ones, which is adequate. Once certain subconcepts have been created, the domain expert has to check whether there are links to other concepts that may require consideration as well by going through the ontology like a map. For example, if a certain ResourceUser has been defined, the ontology would suggest or even implicitly force one to check the "Territory" this

ResourceUser is located in to determine the "Judicature" of it and what "PrivacyLaws" apply. This will then reflect back to the overall level of privacy, as a Judicature with weak privacy protection laws would most likely result in an overall weaker level of privacy, if that ResourceUser is involved in a transaction utilising the resources of the data subject.

This general ontology has been entered into Protégé, which we use to generate the relevant OWL code for implementation, testing and evaluation.

4. Ontology instance for the medical domain

Firstly, this section shows the principles for using the generic ontology in the medical domain in general. We emphasize the classification of concepts that are used in this domain. Afterwards, we describe a specific example by classifying an "interpreter" in the medical domain.

4.1 General classification principles

As described in the last section, the generic ontology requires the help of a domain expert to create a domain-specific extension, which elaborates the actual concepts of that domain. Furthermore, it is necessary to assess these concepts with regards to privacy and assign proper albeit fuzzy values for the influence and the actual values of the privacy level for that concept. Obviously, the medical domain is very sensitive with regards to information and their use and disclosure, usually making them very strictly regulated. To create an instance of the privacy ontology for the medical domain, we consulted the official Privacy Manual of the Department of Health, NSW [11] and extracted necessary and required concepts. Nevertheless, we only focus on major concepts and won't elaborate small sub-concepts that may have only little additional influence on the privacy level as our purpose here is to provide an exposition of the methodology and not describe the whole ontology by itself.

Some of the concepts found with regards to personal information are shown in Figure 1 with their relative level of sensitivity.

Concept/Resource	Sensitivity Level
Mental Health	High
Genetic Information	High
Information collected in conjunction with health services	Low
Images (e.g. X-Ray)	Medium

Table 1: Resources & Level of Sensitivity

The level of sensitivity for data in the medical domain is also relative to the level of sensitivity (meaning importance of privacy) of the domain itself, meaning that the high sensitivity of that domain makes a resource/concept with medium level of sensitivity in that domain still highly sensitive itself, but less sensitive than a highly sensitive resource in that domain. Figure 1 demonstrates this by comparing a highly sensitive resource in the demographic domain (e.g. date of birth) with a low sensitive resource (e.g. Information collected in conjunction with health services) in the medical domain. As one can see, the highly sensitive information of the demographic domain is still lower than the low one of the medical domain – relatively at least.

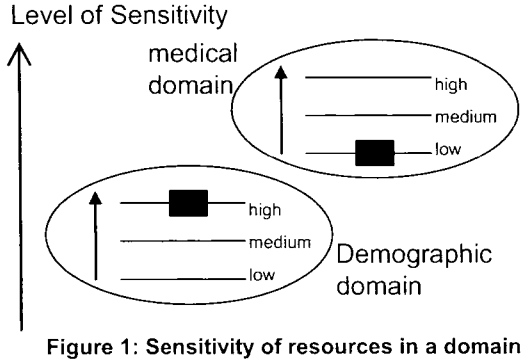

Figure 1: Sensitivity of resources in a domain

When combining different resources with different sensitivity levels as a single resource, the level of sensitivity is defined as the function $s_{r1..n} = \rho(s_{r1},...,s_{rn})$ where r depicts the resource and s_r stands for the level of sensitivity of r. The function ρ maps the sensitivity levels of

the resources $r_1 \ldots r_n$ in such a way that $\forall_k k = r_1 \ldots r_n \rightarrow s_{r1..n} > s_k$. This simply states that the level of sensitivity of a combined resource is always higher than the sensitivity of any of its components. Assuming that we have sensitivity levels on a scale from 0 to 10, where 10 is the highest level, single concepts would be assigned a level between 1 and 9. The role of the function ρ is to combine those single values to combined levels, to generate values greater than 9, but never actually reach 10, which is an upper boundary that is only reached when the sensitivity of an unlimited number of resources is combined. The lower boundary (zero) is assumed never to be reached, as perfect privacy does not exist in general.

When we classify resources with different sensitivity levels, this is done in a fuzzy way and may differ from data subject to data subject. It is generally assumed however that the majority of users will have a similar approach to classify information - at least in terms of relativity within the domain (e.g. mental health is more sensitive than the question if one is an organ donor). If a data subject is not satisfied with the sensitivity levels a domain expert has assigned to the different concepts, alteration is possible in order to accommodate the data subject.

Classifying other concepts, e.g. an interpreter or third party health provider is generally more difficult, as one has to classify first what exactly that concept does, how it interacts and what other concepts it is linked to.

4.2 Example of classifying a specific concept in the medical domain

We will now show an example for such a classification for the concept of "Interpreter", which would be done by a domain expert to create this domain specific concept (regardless of the domain it actually is). However, we will only show the changes that happen to the level of privacy when the interpreter is added and not the whole scenario, as this would be fairly lengthy and cannot be described in the limited space available in this paper.

Firstly, it is necessary to determine where the concept of Interpreter is most likely going to be "attached" to, meaning to figure out what the next logical level of generalisation in the hierarchy might be. Using a top-down approach when going through the ontology, the concept of interpreter is clearly an entity, which represents a cognitive agent (the actual interpreter as a person or a software agent). Therefore it is governed by a certain jurisdiction, depending on the territory the person belongs to. This in return *may have* implications for the level of privacy, as the privacy regulations the interpreter is bound by, may differ from the one the hospital is regulated by. However, a domain expert would be required to determine if it differs or not. If the jurisdiction is the same as the client's one (the hospital, or a GP), the level of privacy would not be influenced. Stepping down one level in the hierarchy, it is also clear that an interpreter will have something to do with access to patient information, making them a "ResourceAccessor", which is classified as a concept that will deal with certain personal information at some stage, although this may not even see the data but just "pass" it on to some other entity. However, as the interpreter will be a recipient of personal information at some point, they are also a "ResourceUser". The latter concept is important, as a recipient of personal information is governed by a "PrivacyProcess", which itself is governed by a certain "Policy" that outlines the conditions of that process and also links to safeguards that may protect the transfer of information from a "ResourceAuthoriser" to the recipient. Furthermore, interpreters will also directly access personal information, making them a "ResourceReader" that can read information, but may not alter or delete any part of it. However, an interpreter is a very specific ResourceReader, the so called "ResourceHandler", inferring that they have "reading" access/consent and additionally may "translate" information in a way that content is not altered semantically. Another example of such a "ResourceHandler" would be a person entering data from a form (on paper) that has been filled out by a data subject into a digital system.

However, translating information into another language for example may not be completely accurate, but the interpreter should make every reasonable attempt to translate as closely as possible to the original version without adding any own opinions or values.

Now that the position of the concept "Interpreter" as an instance of a ResourceHandler has been determined within the ontology, we can now look at the levels of privacy and how the concepts influence them. We assume that the actual interpreter is a real person, working in the hospital and communicating verbally with the patient.

Being a real person (and not a software agent that translates), the interpreter is actually an entity, but as the person works for the same hospital, it is within the same jurisdiction as the other ResourceUsers (e.g. doctors, nurses), the patient may have given consent to. Therefore, the level of privacy, particularly, the privacy principle "Transfer of data outside certain legal boundaries" is not influenced and can be used from the original evaluation of the hospital (which is not shown here). If the translator were a software component, similar rules would apply. This is because it would run under the authority of the hospital and a separate check for jurisdiction would not be necessary. The patient/client will have an agreement with the interpreter, whether it is explicit or implicit. We assume an implicit agreement, which is not written down in any form. Thus, the patient would assume that it is essentially the very same one that has been established between them and the hospital (or doctor or nurse). Due to the fact that an interpreter is a ResourceReader, they are not allowed to disclose information to anyone else. This seems to be a contradiction, as an interpreter needs to "tell" the translated content to a healthcare professional. However, this can be regarded as translating the content (without alteration) and handing it back to the client, who then forwards it to the relevant doctor or nurse, which would receive the information directly normally if no interpreter were necessary. As previously stated, the ResourceUser (which the Interpreter is implicitly) is the recipient of a resource, the resource must be transferred from the client or any other ResourceAuthoriser (which does not exist in our example) to the interpreter. The communication path in our case is defined as "VerbalCommunication". As any communication path is defined as having a safeguard to protect it, we need to define one here as well. For simplicity reasons and to stay realistic, we simply assume that being in a separate room with the patient is quite a strong safeguard. We ignore issues like other patients in the same room, audio monitoring or people that could read lips. The concept of Interpreter also assigns a value to the privacy principle of "Purpose", which is essentially "Translating between languages" and nothing else. Obviously, this requires the interpreter just to translate the information provided by the client and not to do anything else with it – referring back to our previous argument about disclosure.

The privacy principle of Quality may be influenced by an interpreter, as the accuracy of information might degrade during translation, but we assume that is negligible.

Concluding this section, one can see that an interpreter can have an impact on the overall level of privacy. However, the example tried to keep the level of privacy as high as possible, for example by using an internal interpreter, assuming similar agreement policies (albeit implicit, which can often has an impact) and talking in a separate room for security reason.

5. Example from HealthConnect

This section uses an example from HealthConnect [12] when registering a new client to show the privacy issues in the digital and non-digital world and where the ontology can help the client/patient to make decisions upon what he can expect from a level of privacy perspective. We will not show specific technical insights, but briefly describe where the ontology could be "hooked in" to help evaluating the expected outcome.

"Mr. Kevin Stephenson, a 55 year-old male with diabetes mellitus, presents for the first time to a general practitioner's surgery and informs the

receptionist that he is feeling extremely unwell. The receptionist asks if he has an existing HealthConnect record which he denies and the receptionist asks him if we would like to register. He asks Kevin to fill out the usual surgery form for a new client. This includes Kevin's demographic details (name, sex, DOB, address, contact details) and basic health information (his allergy to penicillin, diabetes diagnosis, treatment for depression and recent knee operation) and registers him into the local system. The receptionist gives Kevin a brochure that outlines the scope and benefits of HealthConnect and includes and internet address for further information and contact information for the HealthConnect helpdesk. Kevin advises that he is already familiar with HealthConnect and would like to register. The receptionist helps him complete and sign the registration and consent forms. Kevin selects the standing consent option (ie consent to authorised providers unless otherwise revoked) where access to his HealthConnect record is made available to a nominated list of providers, although he chooses to limit access to information from consultations involving mental health issues to his psychiatrist and general practitioner. He understands that in emergencies this limitation may be overridden but a review of such an action would be undertaken. The collected registration and consent information is sent to HealthConnect from the local system. Following an immediate check for an existing registration, Kevin is provisionally registered pending receipt of his signed form and issue of his HealthConnect access control token." [12]

Although this example is fairly small, it contains quite a few privacy sensitive issues. A few of those issues are related to the real world, while others refer to the digital world. For simplicity reasons, we will omit the fact that Kevin needs to be registered to the local system first but concentrate on the direct registration for HealthConnect.

In order to show that the ontology *can* be used for the non-digital world as well, we will now briefly outline how this might work. The relevant work (e.g. conceptualising terms) would be done by a domain expert upfront and not by a client. By entering the surgery, Kevin also enters the domain of health care, which requires building or utilising a mapping from the core ontology to the medical one in order to discover and use the relevant concepts. As Kevin decides to register with HealthConnect, he has to fill out the relevant form and provide some details about his person. Naturally, Kevin needs to read and understand the privacy policy of HealthConnect that is written in natural language. It is necessary to conceptualise and clarify the statements in the policy and create an appropriate mapping from the core ontology to the specific statements of the policy. The next step would be a mapping between the type of data he enters into the form, which would be demographic information and therefore information that identifies him. Furthermore, Kevin is required to use his real identity and not an arbitrary or anonymous one, which is then checked by the receptionist to make sure his information is accurate. Entering both, demographic information and his basic health information in the same form makes it a very sensitive resource, especially considering he also enters mental health data which is highly sensitive in the already sensitive domain of health (see section 4). By signing the document with his written signature, he legally certifies the document as being accurate, making the document even more sensitive, as the provider (HealthConnect) is guaranteed that the information is precise and really about him. The physical nature of the form (paper) also adds privacy implications, as paper is durable and needs to be retained during his membership with HealthConnect as legal proof of the information he provides. Therefore, safeguard issues arise that directly influence the security principle, e.g. how is the paper stored and who has access to it and how this is enforced. Furthermore, Kevin has to fill out the form in the practise and pass it on to the receptionist, which could require us to address those circumstances as well. For example, someone could read the information while he writes it or it could just lie on the desk at the counter unprotected. However, in order to stay realistic, we will not take those special issues into

consideration, but assume that filling out the form and passing it on to the receptionist is reasonably secure. The receptionist is also a "ResourceHandler" as described in section 4, as it's his/her duty to "convert" the hardcopy form into a digital one without alteration or disclosing the information to anyone else. Obviously the pathway that is used to convert the document from the hardcopy form also plays a role, e.g. if the communication with the online system is secure.

Once the data is stored in the digital system, it is logically important for the client to make sure the data is only used in a way he wants it to be used. In general, one would assume that the data is protected by safeguards (e.g. encrypted), but due to the nature of electronic health records, the influence of the principle of security may have exceptions when compared to other domains. This is related to emergency situations where data must be available even if the patient is unconscious and therefore unable to give consent to access it. If health records were encrypted in such a way that they cannot be accessed in this situation, the vital interests of the data subject (e.g. life or death situation) would be "violated". For the data subject, a greater importance and influence on the overall level of privacy would therefore lie in the consent to make sure only authorised people can access it and the notification whenever personal information are used, regardless of how this is implemented by the provider, as long as it is guaranteed (and legally backed up).

6. Conclusion and Future Work

In this paper, we have presented how a medical privacy ontology derived from a generic one can support information about privacy levels during certain interactions. The medical mapping, which is an extension of the core privacy ontology, is hereby used to support the specific concepts of that domain. The actual instance of that domain provides actual values for the influence of the specific concepts to the level of privacy and the relevant privacy principles. As the medical extension only supplies additional concepts and values for that domain, the structure of the core privacy ontology provides the different interactions and core concepts to support a derivation of privacy levels experienced.

As described, the example cannot be described exhaustively due to the nature and complexity of the medical domain, requiring incorporation of a large amount of concepts and their influences on the level of privacy. Utilising expert domain knowledge, the domain specific instance of the ontology is currently built and implemented to use automatic tools for deriving the actual level of privacy a user experiences. We will apply this to HealthConnect to see how good the support for privacy for electronic health records is, where problems exist and how they might be possibly addressed.

The mapping itself is not limited to the medical domain, but is used in an exemplary manner and other domains will follow later on, especially to evaluate privacy when covering a number of different domains, while referring back to the basic concepts of privacy in the core ontology.

Subsequently, the core ontology will be used to "tweak" the privacy experience of data subjects on an individual level, as they are able to set their own influential values for the different principles. This can be applied to different and new domains, even though the user may not have actually used any privacy preferences for this new domain before, as long as the domain specific one can be mapped to the core ontology.

References

[1] Merriam-Webster Online Dictionary: "Privacy", accessed 08/2006.

[2] S. D. Warren and L. D. Brandeis, "The Right To Privacy," Harvard law review, vol. 4, pp. 193-220, 1890.

[3] Privacilla, "Privacy Fundamentals: Privacilla's Two-Part Definition of Privacy," 2003.

[4] G. Radwanski, Privacy Commissioner of Canada, "Patient Privacy in the Information Age", E-Health 2001: The Future of Health Care in Canada, May 29, 2001

[5] World Wide Web-Consortium, "Platform for Privacy Preferences (P3P) Project," 2004.

[6] M. Hecker, T. S. Dillon, "Towards a privacy ontology", Submitted to: Communications of the ACM (under review)

[7] "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data" http://www.cdt.org/privacy/eudirective/EU_Directive_.html

[8] PRIME, "Privacy and Identity Management for Europe – PRIME White Paper," 2005

[9] J. Huizenga, Handbook of Privacy and Privacy-Enhancing Technologies: The case of Intelligent Software Agents. The Hague: College bescherming persoonsgegevens, 2003.

[10] E. Chang, T. S. Dillon, F. Hussain "Trust and Reputation in Service Oriented Environments", John Wiley and Sons Ltd, Chichester, 2006.

[11] Department of Health, NSW "Privacy Manual (Version 2) – NSW Health", version 2, June 3, 2005

[12] HealthConnect Business architecture version 0.7