A Security Model and its Application to a Distributed Decision Support System for Healthcare

Liang Xiao[1], Javier Vicente[3], Carlos Sáez[3], Andrew Peet[2], Alex Gibb[2], Paul Lewis[1], Srinandan Dasmahapatra[1], Madalina Croitoru[1], Horacio González-Vélez[4], Magí Lluch i Ariet[5], David Dupplaw[1]
[1]*University of Southampton, UK*
[2]*University of Birmingham, UK*
[3]*ITACA, Spain*
[4]*University of Edinburgh, UK*
[5]*MicroArt, Spain*

Abstract

A distributed decision support system involving multiple clinical centres is crucial to the diagnosis of rare diseases. Although sharing of valid diagnosed cases can facilitate later decision making, possibly from geographically different centres, the released information could reveal patient privacy if it is not properly protected. Clinical centres may have to impose their distinct regulations and rules that govern the use of their data externally. The collaboration of centres, therefore, must respect the collective policies and ideally, serve users the most appropriate and useful resources possible in the system according to the past experience. In this way, the system's value is entrusted and even elevated through continuous collaboration. We present in this paper a link-anonymised data scheme and in addition to that, a security model that together enforce privacy data security and secure resource access for distributed clinical centres. Our illustration of the approach involves a prototype medical decision support system, HealthAgents, for brain tumour diagnosis.

1. Introduction

Distributed decision making systems are becoming increasingly useful and important for the efficient sharing of data and services amongst collaborative partners. Use of these systems, based around distributed processing, requires the security design to promote trust. The internet infrastructure promotes open transferring of data which in itself is not a safe environment. Well-studied and publicly available data encryption algorithms can alleviate this problem when incorporated into the system messaging network. The data transmitted in these systems requires secure anonymisation processes. Further, the data access requires careful management to allow different levels of access rights of users distributed amongst multiple organisations. These organisations need to use resources from others and prevent their own resources from unauthorised use. If a system is over restrictive in resource access control then the system is not useful. If a system is not sufficiently restrictive then the organisations' privacy data is in danger of being exposed. This paper investigates data anonymisation and the access control required for the protection of critical resources in collaborative systems.

2. HealthAgents overview and link-anonymised data scheme for preserving privacy

Brain tumours are still an important cause of morbidity and mortality in Europe [1]. The current gold standard classification of brain tumours by biopsy and histopathological analysis involves invasive surgical procedure and incurs a risk of 2.4-3.5% morbidity and 0.2-0.8% mortality, in addition to healthcare costs and stress to patients. There is a need to improve brain tumour classification, and to provide non-invasive methods for brain tumour diagnosis and prognosis, to aid patient management and treatment.

The HealthAgents project [2], funded by the EU's Sixth Framework Programme, aims to build the world's largest distributed data warehouse of brain tumour cases data. The multi-disciplinary collaboration involves seven educational and research institutions, two SMEs, as well as some subcontractor hospitals and external expertise groups. These groups are spanned over Belgium, Italy, Spain, and the United Kingdom. HealthAgents inherits the achievements of its predecessor INTERPRET [3] and is related to the ongoing eTUMOUR [4] project. It plans to create a multi-agent distributed Decision Support System (d-DSS) based on novel medical imaging and laboratory tests to help determine the

diagnosis and prognosis of brain tumours. Novel medical imaging techniques, such as magnetic resonance spectroscopy (MRS), and laboratory techniques, such as gene expression arrays, promise to deliver these advances. These techniques suffer from a complexity of interpretation which has hindered their incorporation into routine clinical practice. However, they provide an excellent test bed for the development of a computer aided decision support system. Furthermore, the rarity of many brain tumour types requires that information must be sought from many hospitals. The use of a distributed system for data collection and management is, as a result, a necessity.

Prior to incorporation into clinical practice new methods must be fully tested within a clinical trials setting. Such trials are subject not only to data protection laws but also regulations governing clinical trials including ethical approval and informed consent of the participants. For multinational projects, ethical approval is devolved for regional bodies without any coordinated or uniform decision making and so data gathered from different centres may be subject to different restrictions. Allowing for flexibility within the data security model is therefore essential.

Clinical trials commonly use data from which personal information (e.g. name, address, date of birth) is removed but to which a unique patient identifier is added, often termed link-anonymised data. Such a scheme has the advantage of having a high chance of preserving patient anonymity whilst allowing data from the same patient to be added at a later date. This scheme also allows a specific patient's data to be located and removed from the project at any time they request, a condition usually imposed by ethics committees. Full patient records are kept for clinical purposes within the treating hospital and with the patient's permission may be used to generate and periodically update the clinical trials data.

Clinical trials are usually supported by a centralised database where the link-anonymised data is stored. This allows the patients to be reassured that their data will be afforded a high level of security and allows regulatory bodies ease of access to inspect the processes in place. For a distributed system, similarly robust arrangements must be designed to reassure ethics committees and patients that the data is secure. However, achieving this is a significant challenge and here we discuss a potential model for achieving this together with the necessary technical requirements and their proposed solutions. Each data collecting centre could have an associated link-anonymised database as approved by their appropriate ethics committee. Patient identifiers could then be kept along with the clinical patient record in the treating hospital. These databases need be the only databases kept within the system giving a truly distributed data-warehouse. The limited data required for analysis could then be subject to stringent anonymisation processes and sent to a small number of specific sites for processing, for example the production of classifiers. In this way, the distributed nature of the system could be preserved whilst allowing appropriate regulatory access to data repositories. Security systems will need to be in place which can allow each centre to potentially limit the type of data transmitted and the locations it is transmitted to.

3. The need for an enhanced security model

While complete patient records may be accessed only by hospitals and local nodes, link anonymised records may be exchanged between a limited numbers of centres producing classifiers. Furthermore, only limited amounts of data which can be considered as totally anonymised may be accessed outside the closed project network. A model shown in Figure 1 illustrates such a data protection model in a multi-layered fashion.

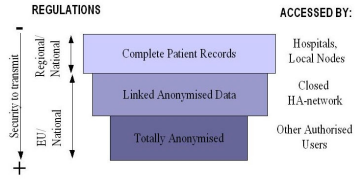

Figure 1. Prototype secure data protection model for HealthAgents

Apart from the link-anonymised data scheme, the mechanism used by the system for decision making itself offers a further level of protection to privacy data. In the system, cases are processed and tumour classifiers produced while the patient privacy is preserved. This is because cases are normally only known to the classifier producer software (agents). In the tumour diagnosis processes, the produced classifier software (agents) as opposed to specific cases are used for decision making. If no such classifier is available a new one may be produced using the available cases. In any case, no private patient data that is involved in the production of classifiers will be revealed to the clinical users.

The classifiers, used for differentiating tumour type, grade, or character, are produced by using different pattern recognition methods and data trained using the available cases. If new clinical centres, with their local case databases, join the existing collaborating centres, they can employ the classification services based on the validated data available from around the network, as well as providing new brain tumour cases for the distributed data warehouse. New classifiers can then be produced or existing ones improved using these new relevant data available. Figure 2 shows the HealthAgents network.

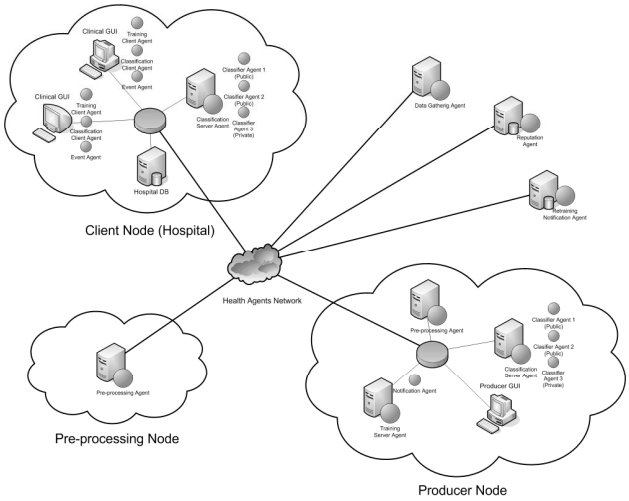

Figure 2. The HealthAgents network

The patient's private data sent from the hospital is protected by the link-anonymised data scheme and its exposure to users minimised by the classification mechanism. This, however, does not render the system safe. Maliciously or accidentally, users may create low quality classifiers, or assign unmatched ranking values to classifiers. This could happen if an inexperienced clinician, with good motivation, trains classifiers or updates their dynamic performance using low quality spectra (signal-to-noise ratio lower than 10, etc.). The use of these classifiers distracts the process of supporting diagnosis and is untrustworthy. Therefore, in addition to the private data protection scheme, a mechanism must be in place for the access control of the critical system resources. This is to avoid abuse or misuse of them by those without authorisation or sufficient privileges. Yet it should be sufficiently flexible for resource sharing among collaborative parties.

The age of patients and brain tumour locations, for example, can be associated with tumour types. This information is useful for diagnosis. A contract signed between two clinical centres may allow some cases to be transferred to a single trusted third party but no further. The collaboration of multiple centres, which not only provide their cases but also require classifiers for their own use, requires the system to respect the access control policies individually employed by each centre. In addition, there might be global constraints applicable to shared resources. All these policies and constraints could change continuously according to the system needs. For instance, a new junior clinician who has just joined one of the collaborative centres may have no right to create a new classifier, or give a definitive diagnosis to a case that will later trigger a classifier reputation being updated. These operations could have global impact on all diagnoses across centres. But he/she may be allowed to do such operations later on when they gain more experience. The system may have to assign to different users or even the same user at different times or under different contexts, various access rights to system resources distributed amongst the centres. Moreover, after accumulative interactions, the system could possibly tell which classifiers are good and which are bad in terms of their performance, feedback being obtained from clinicians after their use of them. The system could then, ideally, always find the proper nodes where high quality classifiers are built and high quality data is supplied, and even adjust the overall interaction pattern to serve its users. Many such scenarios being considered, a model adaptive to continuous collaboration is needed, concerning not only security (access control in particular), but also trust and reputation which all have crucial global effects on the overall system. A solution centred on a particular type of agents, the YellowPagesAgent, will be discussed next.

4. An enhanced resource controllability and performance dependability model

In the heart of the HealthAgents network shown in Figure 2 is the YellowPagesAgent. The YellowPagesAgent plays a key role in agent communication of the HealthAgents system. Agents can search for other agents in the YellowPages based on agent properties and send the messages to the result of that search. Apart from the yellow page function originally designed in the system useable to all agents for looking up information, the YellowPagesAgent is envisioned a key component and a control point for the system's resource access and secure communication, as well as the continuous improvement of the system's performance and hence the value of the system.

4.1 The secure communication mechanism

Communication amongst clinical centres must be secured. This means that the messages being transported in the HealthAgents network which might contain patient privacy information or diagnosis results should not be intercepted or modified by eavesdroppers. Symmetric encryption involving secret keys is best suited for the encryption of the message contents while asymmetric encryption involving public and private key pairs for the protection of the secret keys. In the infrastructure, we make use of YellowPagesAgent for storing and managing public keys and in establishing trust relationships. Only agents who have been formally recognised and registered in the YellowPagesAgent will be regarded trustworthy and so YellowPagesAgent plays the role of Certificate Authority (CA) in the sense of

their assurance of the trustworthiness of communicating parties. Being an integral part of the framework, the YellowPagesAgent simplifies the mechanism of the secure communication.

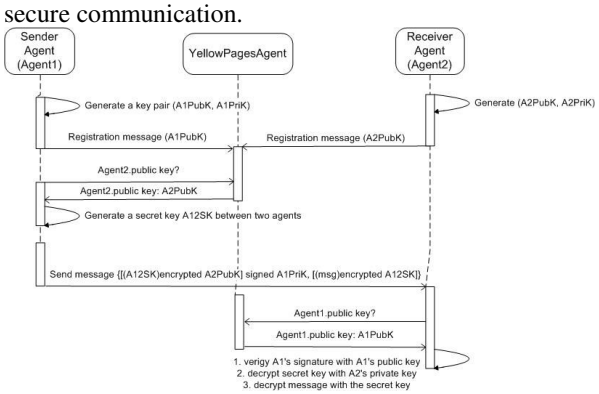

Figure 3. The secure communication scheme in HealthAgents

More specifically, Figure 3 shows a generic scenario with two agents communicate with each other. The receiver agent must at start-up stage, while it registers itself to the system via the YellowPagesAgent, generates a pair of public and private keys. The public key is obtained by the YellowPagesAgent and the private key obtained by itself. The sender agent can retrieve receiver agent's public key, at runtime, from a key store maintained by the YellowPagesAgent. Upon obtaining this public key, the sender agent generates a secret key that will be used to encrypt the plain-text message to be secured. The secret key must be shared between two agents. This can be achieved via the sender agent's encryption of the secret key using receiver agent's public key. This data with the secret key encrypted is further signed by the sender agent's private key. The secret key protected message and the private key protected secret key is encapsulated in the transmitted message. Upon receiving the message, the receiver agent reads the sender agent's signed data and verifies its identity by retrieving the public key of the sender agent from the common public key store. The data is then decrypted by the receiver agent using its own private key and thus the secret key is revealed. The encrypted message will be finally decrypted using the secret key. A common approach for implementing this scheme is the Java Cryptograph Architecture (JCA).

4.2 The resource access control scheme

The other layer of security in the HealthAgents system is concerned about the resource access control in the business level as opposed to the physical network level. This layer of security requires more delicate considerations where ordinary business needs shall not be compromised and the users access what they have been granted. The YellowPagesAgent constrains the collaboration pattern through the imposition of access control.

Specifically, the YellowPagesAgent can be looked up by Clinical GUI Agents which send questions to be solved and then a list of classifiers appropriate in that context will be returned. Moreover, the yellow pages can be referred to for data sources when new classifiers need to be produced. In this business infrastructure, the YellowPagesAgent maintains a list of available classifiers, along with their associated profiles including abilities (questions to be solved for clinicians), reputation, and a profile of the training data with which they were produced.

Once trained by the Training Manager Agent, new classifiers can register themselves with the YellowPagesAgent together with their profiles. Clinicians can then search for those relevant to the particular cases under consideration using the GUI Agent. Once classification results are produced, they are evaluated via comparing with the validated diagnosis results supplied by the clinicians and the reputation of classifiers is updated accordingly in the YellowPagesAgent. The next time when they are running, more accurate information about these classifiers is known to the clinicians. This process continues iteratively and the YellowPagesAgent keeps updating classifier profiles for the most accurate and efficient performance of the overall system possible.

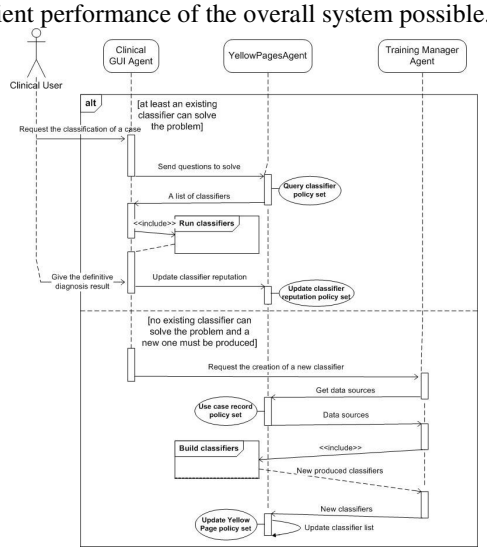

Figure 4. The sequence diagram of data and classifier access control in HealthAgents

Figure 4 shows the message passing sequence among several HealthAgents agents. The processes of running and building of classifiers are included as part of the overall diagram. The diagram illustrates the YellowPagesAgent's function in informing clinicians of classifiers and informing classifier producers of data sources for the production, as well as maintaining the reputation of classifiers. Two major alternative

interactions involving distinct YellowPagesAgent functions are differentiated and shown in the upper and lower partition of the "alt" region with their respective guards. Various security policy sets are applied in corresponding circumstances, e.g. when available classifiers are queried and, once the validated diagnosis of the case is given by the clinician, reputation values of the executed ones are updated and so YellowPagesAgent is maintained. The security constraints are usually explicitly expressed and such knowledge is subject to continuous maintenance, being in a direct human intervene process as shown in Figure 5.

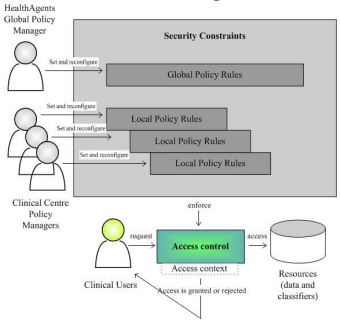

Figure 5. A scheme relating the impacting factors to data and classifier access control

The design diagram of Figure 4 indicates the global resource access control the YellowPagesAgent could impose as well as the affect it could make to the overall classification results through its own performance. We are aiming at achieving flexible management of security access and continuous performance enhancement, respectively, through the careful design of the YellowPagesAgent.

In the secure access perspective, clinicians with certain access rights should only access the proper resources and do the proper operations. The YellowPagesAgent may reject access to private classifiers (e. g. a classifier trained exclusively with data from one and only hospital, as opposite to a public one, trained with data from all the hospitals in the network) from external centres. Also, the YellowPagesAgent may reject classification production requests or classifier reputation updating requests from certain clinicians. But such response should by no means be fixed. Instead, it should use the up-to-date policies to reflect the current security needs. A security policy model elaborated in Section 5 will discuss in more details a solution to the outlined security infrastructure.

4.3 The performance and system usability enhancement

In the performance and system usability enhancement perspective, the appropriate classifiers for use depends on many factors, including not only the performance, but also the similarity of the new case under classification with the ones used for training the classifier. Again, a flexible model must be available to the YellowPagesAgent to serve classification requests, transparently to end users. Briefly, the suitability of classifiers being used for classification depends on the followings and these are illustrated in Figure 6.

- Similarity of the new case with the training set from which the classifier is derived.
- Static performance of the classifier. The classifier the evaluation is based on the accuracy, the balanced error rate, or the geometric mean of success as calculated after the training. This is generally obtained with an independent test set or, if not available, using techniques such as cross validation.
- Dynamic performance of the classifier. This is the performance of the classifier with the 'unseen' cases the clinicians launch for orientation purposes. The answer given by the classifier is compared with the diagnosis the clinicians give once they are sure of it.
- Use level of the classifier.
- Evolution of all the previous factors during the classifier's life.

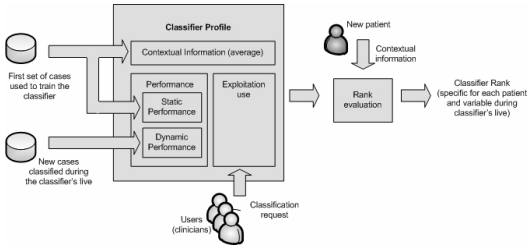

Figure 6. A scheme relating the impacting factors to classifier reputation/ranking

The knowledge accumulated in the running processes of the HealthAgents system is therefore of significant value to the automatic improvement of system performance. Comparatively, such knowledge is implicitly incremented as opposed to explicitly specified as Figure 5 shows.

To enhance the management of both reputation and security perspectives of the system, data, classifiers, security policies, and even people distributed among multiple clinical centres must collaborate in a manner that respects disparate impacting factors and take into account their dynamics. Note that both the "Security Constraints" in Figure 5 and the "Classifier Profiles" in Figure 6 are managed by the YellowPagesAgent. These can be seen as two sets of metadata or knowledgebase that the agent uses in working between users and resources. The "Security Constraints" are used in the first place to justify if users can access resources and if so, then the "Classifier Profiles" are used to choose the suitable resources for authorised users. The principle is

simple, in the distributed environment, users should have limited access to resources but if they do have the access rights they should be offered the best possible services. The system must be aware of such facts as who can access what, as well as which are the best services to provide in that context, if such an access is permitted. The building and maintenance of such metadata is of primary importance to the proper running of the distributed collaborative system and achieving its full value. Considering the public key store in Section 4.1, the repository maintained by yellow pages includes a total of three types of metadata that enable the YellowPagesAgent to play three types of roles, contributing to a secure and dependable healthcare system.

An approach that enables the YellowPagesAgent to behave securely and adaptively via the Adaptive Agent Model [5] [6], being part of an integrated methodology and providing a secure resource access control mechanism that supports the infrastructure outlined in Section 4.2, is now discussed in the following section.

5. A solution based on an adaptive security policy model

Table 1. Two example interactions in HealthAgents with their security implication

Interaction between YellowPages Agent and clinicians	Interaction description	Security implication
Clinicians want to solve questions	The clinician sends a MRS case along with diagnose (or differential) he/she thinks suitable to the case (tumour or non-tumour, aggressive or non-aggressive, glioblastoma or metastasis, etc.). Classifiers trained with MRS data that can discriminate among tumour classes or grades (so answer various questions) have been previously registered in the yellow pages, including descriptions of their capabilities, reputation, and data they have been trained upon.	Queries about classifiers are answered only if the clinician has access rights to the required classifiers in the classifier directory.
		Secure access of yellow page directory service: query classifiers
Clinicians want to give feedback after using classifiers	When the correct diagnosis is known for the case, the clinician will update the case record in the database with this information. The classification results produced by the selected classifiers will be evaluated and their reputation (dynamic performance) updated in the yellow pages so that other clinicians will have better knowledge about how good each of these classifiers is in later use.	The clinician must be authorised to update the case record as well as classifier reputation.
		Secure access of yellow page directory service: update classifier reputation values

Concerning only the upper part of the Figure 4, the scenario of running classifiers where classifiers are queried and their reputation updated can be described in Table 1, giving the security implication. Two distinct policy sets may have to be applied in the query and updating conditions. These security policies must be integrated into the business functions of the Multi-Agent System with its actual functions being intact during the configuration of policies. Adaptive Agent Model (AAM) provides such a framework for seamless integration. AAM is a methodology that guides the building of an interaction and computation model [7] to drive adaptive agent system behaviour. The model originates from business requirements, is interpreted and executed dynamically by agents at runtime, and under continuous maintenance by business people. Existing components and services can be reused to support agents to execute business requirements captured in the model. Tools have been developed to support the documentation and the maintenance of the model. Reaction Rules (RRs) and Policy Rules (PRs) are central model elements that compose the model and reside in separate knowledge repositories shown in the scenario of Figure 7.

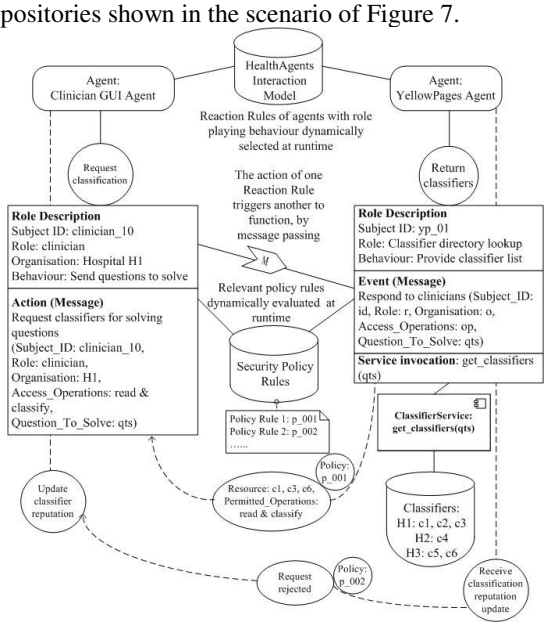

Figure 7. Agent interaction model with security policy rule application, using the AAM [5] [6]

Each RR guides an agent its individual responsibility of service provision function in reaction to other agents and in this way that agent contributes its own capability to the overall interaction process. The RR has the following scheme.

Reaction Rule Scheme:
{(Name, Interaction, Owner), Event, Processing, (Condition, Action)$_n$}

In Figure 7, YellowPagesAgent uses its RR "Return classifiers" to react to Clinician GUI Agent's "Request classification", the RR being part of the first interaction shown in Table 1 and formalised in Figure 8. A "ClassifierService" facilitates the agent to retrieve appropriate classifiers in various conditions. If the clinician has sufficient privilege, he/she will have access

to all classifiers that matches the current case or those that can answer the questions that have to be solved.

```
<reaction_rule>
    <name>Return_Classifier</name>
    <interaction-process>Use_Classifiers</interaction-process>
    <owner-agent>YellowPages_Agent</owner-agent>
    <event>
        <message>
            <from>Clinician_GUI_Agent.Request_Classification</from>
            <content>
                <request>
                    <subject>
                        <ID>clinician_10</ID>
                        <role>junior_clinician</role>
                        <organisation>Hospital_H1</organisation>
                    </subject>
                    <access_operation>read and classify</access_operation>
                    <question_to_solve>qts</question_to_solve>
                </request>
            </content>
        </message>
    </event>
    <processing>
        classifier_list = ClassifierService.
            get_classifiers(requestMsg.getContent(qts))
    </processing>
    <condition>
        ReadClassify_Policy_Set.permission() == true
    </condition>
    <action>
        <message>
            <to>Clinician_GUI_Agent.Request_Classification</to>
            <content>
                <return>
                    <classifiers>classifier_list</classifiers>
                </return>
            </content>
        </message>
    </action>
</reaction_rule>
```

Figure 8. An example Reaction Rule

One or several PRs may be collectively selected and applied in such processes to reflect global business policies that must be complied in service fulfilment, shared by all agents in the system. The original PR scheme has been adapted and specialised as security PR, enforcing security constraints during RR function [7] [8]. The security PR has the following scheme.

Security Policy Rule Scheme:
{Subject (Id, Role, Organisation), Access Operation (Op), Access Context (Co), Resource (Id, Type)}

```
<ReadClassify_Policy_Set>
    <Policy id="p_001">
        <Affection>
            <Role>clinician</Role>
        </Affection>
        <Permission description="This clinician can use any classifier">
            <Subject id="clinician_10">
                <Role>clinician</Role>
                <Organisation>H1</Organisation>
            </Subject>
            <Access_Operations>
                <Access_Operation>read</Access_Operation>
                <Access_Operation>classify</Access_Operation>
            </Access_Operations>
            <Resource>
                <Type>classifier</Type>
            </Resource>
        </Permission>
    </Policy>
</ReadClassify_Policy_Set>
```

Figure 9. An example security Policy Rule

Agents, assigned with specific roles, have their access constrained by instances of policy rules of such a scheme, in the dimensions of (1) resources they can access; (2) operations they can perform upon resources; and (3) access contexts. The YellowPagesAgent is in charge of access to all the available classifiers and responds to requests for obtaining classifiers as well as updating their reputation. Not all clinicians can do both of these tasks. One security PR is formalised in Figure 9.

The constraints identified on the PR (as shown in Figure 9) are enforced upon agents with their normal functions identified on the RR (as shown in Figure 8). This is achieved through an integrated role concept [8]. Here agent's functional duty role as described by the Agent-Oriented Software Engineering (AOSE) [8] and agent's social rights role as described by the Role-Based Access Control (RBAC) [10] are seamlessly integrated. The principle is:

A role plays its **functional duties** *if and only if* its **social rights** allow it to do so.
Functional duty role *is dynamically bound with* **social right role** before any agent plays the integrated role.

Clinicians may play the role of "Request classification" and "Update classifier reputation" functionally. But it depends on their roles (as well as associated ID and Organisation) assigned socially they may or may not indeed perform the claimed functions. Figure 7 demonstrates two Security Policy Rules being defined. Shown in Figure 9, p_001 says this clinician can perform classification by using any existing classifier across centres. Assume p_002 says only senior clinicians can update classifier reputation. Then, when a junior clinician comes to use the system, he/she will be allowed to do the "classify" operation with the classifier c1, c3, and c6 returned which match the current case. But he/she will be rejected of the reputation updating request since its precondition of policy rule satisfaction fails. Both RR and security PR have been formalised in the XML model so that agents can uniformly perform their duties if they are found having the appropriate rights [5] [6] [8].

Our approach complies with Object Management Group's (OMG) Model Driven Architecture (MDA) paradigm [11]. The model that drives agent behaviour, functionally and securely, is associated with but external to the agents. In the special case of the YellowPagesAgent, this means the agent that manages the system critical resource of classifiers always applies the security policy rules from the up-to-date rule set that are relevant to the current resource access requests. This process is carried out on the fly. Consequently, once new security policy rules become available which reflect the current security needs, they can immediately become effective in the running system. Such an infrastructure allows the agent system to behave adaptively and easy to maintain. More importantly, (changing) security needs are respected and put into effect instantly.

6. Conclusions

We have presented in this paper a link-anonymised data scheme for private data protection. In addition, we have offered an approach for secure communication as well as a more advanced secure data access control mechanism. In this way, no private information may be revealed to unauthorised people, and authorised people

can only access critical system resources with the permitted power.

Our approach contributes a security model driven software architecture. The integration of the security model into the functional interaction model allows agents to dynamically evaluate and apply the appropriate security policies before they perform their actual capabilities, a behaviour being driven by the combinational model. Changes to security needs have no effect to the rest of the system concerning core system functionalities. This allows the definition of any number of policies after the system has been developed, security requirements not entangled with others.

The agent technology is promising in both the building of a distributed decision support system for healthcare and ensuring its security. On one hand, agents have the capabilities for representing different services required by the system and providing the backbone to ensure the distribution of data. On the other hand, their abstraction of different processes where resources are accessed can be under the security control if appropriate measures are imposed upon them.

Several approaches have been investigated that employ agents in healthcare domains for providing security. One proposed scheme introduces the concept of heuristic security agents [14] which intercept all calls to resources and check them against behavioural rules before an "allow" or "deny" decision is made. Another method is proposed for the secure access of electronic healthcare records (EHR) which may be scattered across healthcare units [12]. A security agent will be employed per hospital site which authenticates users and controls the access to the local resources by looking at the user roles. This approach has its limitation in the types of resources it can protect as well as the use of the shared common services amongst multiple sites. A third approach introduces a central access control (CAC) system and multiple local access control (LAC) systems [13] to the similar distributed record exchange problem. CAC and LAC are Multi-Agent Systems which use authentication agents, encryption agents, and access control agents. In this architecture, the security level is determined by the weakest LAC.

All the above methods introduce agents or Multi-Agent Systems explicitly for the purpose of access control, security not being considered as part of an integrated software design by software engineers in the first place. A software system may have its functionality and usability negatively compromised if security is to be added or fixed after its implementation. As opposed to these, in our approach security policies are pluggable and maintainable in the system from the beginning. Participant agents serve core clinical business functions with associated security measures or policies applicable by the agents as behavioural constraints before their performance of normal functioning behaviour in the clinical setting. The HealthAgents system will be made more effectively and flexibly secure but work so far indicates that our approach is promising and useful to the development of a distributed decision support system for secure healthcare applications.

Acknowledgements

This work is supported under the HealthAgents and OpenKnowledge projects funded by EU Framework 6 under Grants: IST-FP6-027214 and IST-FP6-027253.

References

[1] Bray, F., Sankila, R., Ferlay, J. & Parkin., D.M., "Estimates of cancer incidence and mortality in Europe in 1995", *European Journal of Cancer*, 38(1):99-166, 2002.

[2] HealthAgents: http://www.healthagents.net/.

[3] INTERPRET: http://azizu.uab.es/INTERPRET/.

[4] eTUMOUR: http://www.etumour.net/.

[5] Xiao, L. & Greer, D., "The Agent-Rule-Class Framework for Multi-Agent Systems", *International Journal of Multiagent and Grid Systems* 2(4):325-351, IOS Press, 2006.

[6] Xiao, L. & Greer, D., "Towards Agent-oriented Model-Driven Architecture", *European Journal of Information Systems* 16(4):390-406, Palgrave, 2007.

[7] Xiao, L., Robertson, D., Croitoru, M., Lewis, P., Dashmapatra, S., Dupplaw, D. and Hu, B., "Adaptive Agent Model: an Agent Interaction and Computation Model", *Proceedings of the 31st IEEE Annual International Computer Software and Applications Conference (COMPSAC'07)*, pp. 153-158, IEEE Computer Society, 2007.

[8] Xiao, L., Peet, A., Lewis, P., Dashmapatra, S., Sáez, C., Croitoru, M., Vicente, J., Gonzalez-Velez, H. and Lluch i Ariet, M., "An Adaptive Security Model for Multi-agent Systems and Application to a Clinical Trials Environment", *Proceedings of the 31st IEEE Annual International Computer Software and Applications Conference (COMPSAC'07)*, pp. 261-266, IEEE Computer Society, 2007.

[9] Zambonelli, F., Jennings, N.R. & Wooldridge, M.J., "Developing Multiagent Systems: the Gaia Methodology", *ACM Transactions on Software Engineering and Methodology* 12(3):317-370, 2003.

[10] Sandhu, R.S., Coyne, E.J., Feinstein, H.L. & Youman, C.E., "Role-Based Access Control Models", *IEEE Computer* 29(2):38-47, 1996.

[11] Object Management Group, 250 First Ave. Suite 100, Needham, MA 02494, USA.

[12] Gritzalis, D. & Lambrinoudakis, C., "A security architecture for interconnecting health information systems", *International Journal of Medical Informatics* 73(3):305-309. Elsevier, 2004.

[13] Choe, J. & Yoo, S, "Web-based secure access from multiple patient repositories", *International Journal of Medical Informatics*, Elsevier, In press.

[14] Keese, J. & Motzo, L, "Pro-active approach to malware for healthcare information and imaging systems", *International Congress Series* 1281:943-947. Elsevier, 2005.