6-1-2009

Towards Knowledge Sharing and Patient Privacy in a Clinical Decision Support System

Liang Xiao
Royal College of Surgeons in Ireland

Bo Hu
SAP Research CEC Belfast

Lucy Hederman
Trinity College Dublin

Paul Lewis
University of Southampton

Borislav D. Dimitrov
Royal College of Surgeons in Ireland, borislav.d.dimitrov@gmail.com

See next page for additional authors

Authors
Liang Xiao, Bo Hu, Lucy Hederman, Paul Lewis, Borislav D. Dimitrov, and Tom Fahey

— Use Licence —

Towards Knowledge Sharing and Patient Privacy in a Clinical Decision Support System

Liang Xiao[1], Bo Hu[2], Lucy Hederman[3], Paul Lewis[4], Borislav D Dimitrov[1], Tom Fahey[1]

[1]Royal College of Surgeons in Ireland, Dublin, Ireland, [2]SAP Research CEC Belfast, UK,
[3]Trinity College, Dublin, Ireland, [4]University of Southampton, UK

liangxiao@rcsi.ie, bh@ecs.soton.ac.uk, Lucy.Hederman@cs.tcd.ie,
borislavdimitrov@rcsi.ie, tomfahey@rcsi.ie, phl@ecs.soton.ac.uk

Abstract. *Patient records and their disease and treatment history can be scattered among healthcare providers. Sharing the knowledge effectively and, at the same time, respecting patient privacy is crucial in providing safe and accurate clinical decision support systems (CDSSs). In this paper we reflect upon our experience in the HealthAgents project wherein a prototype system was developed and a novel approach employed that supports data transfer and decision making in human brain tumour diagnosis. Here we examine the capability of the Lightweight Coordination Calculus (LCC), a process calculus-based language, in combining together distributed healthcare services and meeting security challenges in pervasive settings. The result is that various clinical specialisms, being captured in representational abstractions and making contribution to patient diagnosis and management, retain their autonomy. However, at the same time, the behaviour of specialists in sharing clinical knowledge about their patients and providing clinical support is constrained by policies and rules in respect of their own clinical duties and responsibilities. Being introduced into the programme of the HRB Centre for Primary Care Research, this novel approach has the potential to help the provision of optimal solutions in data linkage and sharing across the Primary and Secondary Care interface. As added value, its application also advances the process of integrating clinical prediction rules and implementing CDSSs in practice and, ultimately, the improvement of quality of care.*

Keywords. Clinical decision support, healthcare knowledge sharing, Lightweight Coordination Calculus, patient privacy

1. Introduction

The vision of pervasive healthcare gives a picture wherein "smart" software agents would act on behalf of human specialists in collecting or monitoring critical life support data, extracting information from the data, jigsawing it together, and eventually enabling decisions and actions to be taken on the outcome of such processes. This requires the move from a "patient-seeing-doctor" pattern to a new healthcare paradigm which emphasises a degree of continuous medical surveillance, with key decisions for medical follow-ups requiring automated processing, and in a decentralised manner.

Given the distributed nature of not only the data but also the care providers or specialists accessing the data and making decisions, a major concern of the new healthcare paradigm is to ensure the data are delivered to the right person at the right moment. Thus far, tacit knowledge is largely explicit only to isolated specialists, organisations and professional guilds. Although the necessity of collaboration has been recognised, there is little systematic knowledge sharing of clinical intervention outcomes.

In the modern world, it is not a surprise to find that a patient is seen and receives consultation from his or her GP; after referral the patient is examined in one hospital; his/her case is reviewed by clinicians from another hospital; and he/she is treated in a third hospital by yet another group of clinicians due to speciality and availability. Data about a particular patient might be held by different GP systems and different departments within one or more hospitals which are possibly located in different countries. Effectively sharing such knowledge, automatically processing it, and using the appropriate pieces and gluing them together for the final clinical decision support, yet respecting patient privacy and ethical regulations in the entire process, have significant impact on the quality of care that one is offered.

This paper describes an approach that addresses the need for knowledge sharing and data protection in the setting of a particular clinical decision support system (CDSS), the outcome of an EU Framework 6 project HealthAgents [6]. The system has been

implemented with much medical input and will facilitate clinicians to make better diagnosis of brain tumours. Here we advance the use of Lightweight Coordination Calculus (LCC) models, work from another EU Framework 6 project OpenKnowledge [8], upon this system and which bring better knowledge sharing capabilities among healthcare professionals in a distributed environment and at the same time facilitate better access control. The contribution of bringing together the two pieces of work is reported in the paper.

2. Patient privacy and confidentiality

A CDSS is at its most useful if patient-specific decision support can be given in the current clinical pathway, taking into account both the patient's medical record and the decision support knowledge applicable to this patient. A major difficulty in building a CDSS lies in the scattering of the electronic healthcare records among various care providers, a fact that makes it necessary to have relevant information linked and shared for decision making and at the same time patient's confidentiality respected. Challenges are seen in different countries. For example, in Ireland, there is no unique identifier for linking Electronic Healthcare Records (EHR) which are scattered across public and private sectors and managed by four separate GP systems and many more hospital systems. In order to address a similar problem, in 1992, the National Health Service (NHS) in the UK attempted to build a unified electronic patient record system and give access to extended NHS community. This has been opposed [1] [2] for the reason that such a system, collecting data from existing GP systems but out of their control, is in conflict with the ethical principle that no patient should be identifiable other than to the GP without patient consent [3] and the result from a survey that most patients are unwilling to share their information with NHS [4]. Healthcare data, being scattered but not consolidated for confidentiality reasons, will have to be shared among care providers and specialists on a need basis for the best delivery of care in a distributed information system. The fact that separately stored patient records and clinical information can now be made easily accessible to a wider range of people implies that a trust relationship must be built up between the patients and the distributed healthcare services rather than the doctors.

Building such trust into a distributed healthcare service or information system requires us to set up two types of constraints.

1) *Only relevant clinicians who have real life relationships with patients in clinical centres should access their records*. Two of the access control policy principles British Medical Association put forward [1] are as follows.

Principle of Access: "Each identifiable clinical record shall be marked with an access control list naming the people or groups of people who may read it and append data to it. The system shall prevent anyone not on the access control list from accessing the record in any way."

Principle of Control: "One of the clinicians on the access control list must be marked as being responsible. Only she may alter the access control list, and she may only add other healthcare professionals to it."

2) *Clinicians should only have access to particular record parts on the basis of their need for providing care and treatment*. It is often the clinical expertise that determines the actual needs of access, to the data parts clinicians routinely work with and by doing so, healthcare roles fulfilled. Such role-based access [5] is assured to NHS patients [10] and is necessary for data protection in other jurisdictions. For example, a receptionist booking an appointment will only have access to basic information to identify a patient and make the booking; pathology medical records or reports may be sent to a pathologist involved in a patient's care; prescription sent to a pharmacist; and sensitive parts not sent out at all.

A named responsible clinician, presumably a patient GP, determines the clinical roles that each care team member or specialist fulfils. Normally, every member or specialist can read the full records and write to the parts they are responsible for. The control of access needs to be strict but flexible. It is up to the actual access needs of the clinicians who actually deliver care to the patient that the record controller or a GP governs who can access what. This paper investigates how such access control can be implemented in a particular CDSS but does not impose any specific policy.

3. LCC and interaction models

We shall use the formalism of Lightweight Coordination Calculus, LCC [8] which is a logic programming language, to lay out the various domains of expertise that are called upon to

accumulate clinical knowledge and evidence as well as to process these for better care delivery. We tolerate the diversity and heterogeneity while systematically choreographing individual resources to combine their knowledge of a particular patient or disease. This shift of emphasis immediately suggests taking a process oriented view for system design and analysis. As a system, supporting healthcare provided by multi-disciplinary teams requires dealing with groups of people who rarely had their systems engineered to do tasks together. This also fits the paradigm of concurrency where there is no single locus of control of task execution. Instead of the other resources existing merely to serve the control unit, these entities lead an autonomous existence and only undergo message-induced transitions upon opening up access to each other: centralised control gives way to concurrent processes wherein each party accomplishes the tasks allocated to it and exposes the results to accommodate requests from others. Moreover, this interaction-based sharing of information enables dynamic knowledge composition: by sharing knowledge through interactions we indirectly share data. We later demonstrate that leveraging interaction models as opposed to trying to combine knowledge in the traditional manner benefits healthcare knowledge management and complements the existing work in the HealthAgents project.

Precisely, LCC is a process calculus for specifying coordination among multiple participants. It accomplishes this by clearly stating what role an individual plays in a messaging process. An LCC model is built upon the principle that role-playing agents should obey the laws and/or protocols that are explicitly specified against the roles that such agents are expected to take. LCC ensures the fulfilment of roles by individuals through regulating the message-flows among them. These include: the messages that should be sent and are expected to be received and what constraints should be satisfied before a message can be handled. The full picture of LCC syntax is shown in Figure 1.

$$
\begin{aligned}
\langle Framework \rangle &:= \{\langle Clause \rangle .\}^{1+} \\
\langle Clause \rangle &:= \langle Agent \rangle :: \langle Definition \rangle \\
\langle Agent \rangle &:= \mathbf{a}(\langle Type \rangle, \langle ID \rangle) \\
\langle Definition \rangle &:= \langle Agent \rangle \mid \langle Message\ Clause \rangle \mid \langle Definition \rangle\ \mathbf{then}\ \langle Definition \rangle \mid \\
& \quad \langle Definition \rangle\ \mathbf{or}\ \langle Definition \rangle \mid \langle Definition \rangle\ \mathbf{par}\ \langle Definition \rangle \mid \\
& \quad null \leftarrow \langle Constraint \rangle \\
\langle Message\ Clause \rangle &:= \langle Message \rangle \Rightarrow \langle Agent \rangle \mid \langle Message \rangle \Rightarrow \langle Agent \rangle \leftarrow \langle Constraint \rangle \mid \\
& \quad \langle Message \rangle \Leftarrow \langle Agent \rangle \mid \langle Constraint \rangle \leftarrow \langle Message \rangle \Leftarrow \langle Agent \rangle \\
\langle Constraint \rangle &:= Term \mid \langle Constraint \rangle \wedge \langle Constraint \rangle \mid \langle Constraint \rangle \vee \langle Constraint \rangle \\
\langle Type \rangle &:= Term \\
\langle ID \rangle &:= Constant \\
\langle Message \rangle &:= Term
\end{aligned}
$$

Figure 1. Grammar of LCC

In LCC interaction model, we use predicate $a(id, role)$ to specify the role that an individual is playing, $\Rightarrow$ and $\Leftarrow$ to specify the direction of message flow, and $\leftarrow$ for constraints. *Term* and *Constant* are implementation-specific. In the current version, *Term* is a well-formed formula in Prolog logic programming language and *Constant* is a Prolog constant starting with a lowercase letter. LCC also provides constructs for parallel (**par**), sequential (**then**) and switch branching (**or**) controls.

$$
\begin{aligned}
&a(on\_call\_doctor, N) :: \\
&routine\_check(P) \Leftarrow a(\_, A)\ \mathbf{then} \\
&\left(\begin{array}{l}
take\_temperature(P) \Rightarrow a(nurse, S)\ \mathbf{then} \\
take\_blood\_sample(P) \Rightarrow a(nurse, T) \leftarrow \neg blood\_test(P)
\end{array} \right)
\end{aligned}
$$

Figure 2. An example of LCC

Interpreting LCC is equivalent to unpacking LCC clauses, finding the next tasks that it is permitted to perform and updating the status of an interaction accordingly. For instance, in the above LCC interaction model, the sequence construct *then* is unfolded by examining the first part of the sequence or, if it is closed (i.e. executed), unfolding the next part. After unfolding, the system tries to instantiate all the variables (e.g. *P* and *A*) to examine the satisfy-ability of LCC clauses. A narrative interpretation of the LCC model in Figure 2, therefore, reads *"when an on call doctor receives a routine check request on a patient (P), he/she first asks an arbitrary nurse (S) to take P's body temperature. When the body temperature is done, he/she asks an arbitrary nurse (T) to take P's blood sample if P has not been given blood test before."* Note that whether nurses *S* and *T* are one and the same person is unknown from the context.

4. Modelling healthcare interaction process in HealthAgents using LCC

When LCC interaction models are used in healthcare, knowledge management is built upon an awareness of the flow of healthcare information within the system, reflecting protocols and guidelines driven by legal and ethical concerns. Tacit knowledge is, therefore, observed through the patterns of messages. Such a transparent knowledge acquisition procedure implies the way that one learns can be literally copied with the same conclusion as long as the same contexts are reconstructed.

The HealthAgents [6] system is a distributed decision support system that supports diagnosis and prognosis, employs a set of distributed nodes that either store patient case data, build classifiers that are trained upon case data and capable of classifying tumour types, or use

classifiers for the diagnosis and prognosis of brain tumours. The magnetic resonance spectroscopy (MRS) data used by the system is built up using anonymous information from child and adult cases. Classifiers are created by the producer nodes that receive requests from the clinicians to generate classifiers for particular tumours. Clinicians will employ classifiers which assist in the diagnosis of patients for particular tumours. The HealthAgents system consists of a variety of agents each charged with a different task (see [7] for more detailed description of the HealthAgents components and architecture).

Figure 3 shows the interaction model where a patient case is first retrieved, then relevant classifiers invoked, case classification performed and finally, the diagnosis results updated as well as the ranking of involved classifiers.
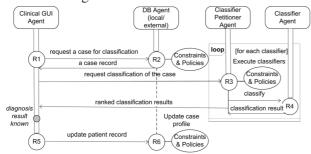

Figure 3. A HealthAgents Interaction Model

The following LCC clauses show its specification. The clinician plays a role of classification (R1) and updating case profile (R5). The role changes when an accurate diagnosis result is known.

```
/* R1: classify a case */
a(clinician_classify, CID) ::
  requestCaseRecordByID(I) ⇒ a(database, DBID) then
  caseRecord (R) ⇐ a(database, DBID) then
  requestClassification(R, C) ⇒ a(classifier_petitioner, CPID)
then
  classificationResults(S) ⇐ a(classifier_petitioner, CPID) then
  a(clinician_followingdiagnosis, CID)
/* R5: update case record and classifier reputation following
diagnosis */
a(clinician_followingdiagnosis, CID) ::
  ( updateCaseRecordByID(I) ⇒ a(database_update, DBID) then
  caseRecordUpdated(Y) ⇐ a (database_update, DBID) )
  par
  ( updateClassifier(I) ⇒ a(classifier_petitioner, CPID) then
  classifierUpdated(Y) ⇐ a (classifier_petitioner, CPID) )
```

A clinician *CID* shall initialise an interaction for a case classification by sending to a database agent a request for retrieving case *I*. CID then expects to receive the case with which he or she can request a classifier petitioner to perform classification on his or her behalf. A set of classification results given by various classifiers being received, the clinician is supported to make a decision about the diagnosis, depending upon his or her confidence on the classifiers and his or her own expertise. Now the clinician can change the role from R1 to R5 which will be activated if

the actual diagnosis result is known. At that time, the clinician can update the case and the reputation values of the classifiers, after comparing the actual diagnosis result and the results suggested by the classifiers in the previous interaction.

Obviously, the clinician who performs the classification and diagnosis, the database where the case resides, the classifiers which give the classification can be distributed in different locations. The interaction model pulls together the different parties, defines the manner of knowledge acquisition, human computer interaction, expertise involvement, clinical decision support, and finally the care delivery.

5. Taking into account the privacy issues

In Section 4 we give the role definitions of R1 and R5 with regard to the behaviour of a clinician for knowledge sharing and decision making in HealthAgents. The two constraints outlined in Section 2 must be respected in such a process – the named clinician must have the right permission to have access to that case and to perform the specified operation upon the case.

The following LCC clauses describe the fundamental interaction pattern for resource access control.

```
a(resource_request, RRID) ::
  request(Resource, Operation, Context) ⇒ a(resource_manager,
RMID)
a(resource_manager, RMID) ::
  request(Resource, Operation, Context) ⇐ a(resource_request,
RRID) ← grantPermission(RRID, Resource, Operation, Context,
Policies) then (
    response(Grant_yes) ⇒ a(resource_request, RRID) or
    response(Resource_result) ⇒ a(resource_request, RRID) ←
getOperationResult(Resource, Operation, Access_result) )
```

a(resource_request, RRID) and a(resource_manager, RMID) define two agents, RRID and RMID, play the roles of *resource_request* and *resource_manager*, respectively. The former role has been instantiated in R1 and R5 in the previous section, where a message of resource access request is sent from the agent that plays the request role to the agent that plays the manager role. Upon receipt of this message, the resource manager agent applies appropriate security policies and responds by sending back a message either saying the request has been granted (or rejected) or by providing the actual resources (or the results of their usage) being requested. It is in the latter role, that a constraint must be checked as satisfactory before the entire interaction model proceeds and the resources are returned.

Generically, "← grantPermission(RRID, Resource, Operation, Context, Policies)" says this constraint will be met (evaluation results to a Boolean value of true) if *RRID* can perform *Operation* upon *Resource* under *Context*, according to *Policies*.

An interaction model is uniquely defined and its running context varies, e.g. involved clinicians and cases. A resource manager must check the request (resource and operation) against the requester identity at runtime, in compliance with the access policies. Specifically, the clinician must be a member of the workgroup delivering care to the owner of the case before the case is allowed to be updated. Additional local policy rules must also be satisfied to accommodate necessary constraints. A generic policy rule schema is described in [7] for this purpose. The following shows the LCC constraints used by the database agent, being a resource manager, for permission checking before the actual role functions are carried out. The database agent issues a case record (R2) and updates the same record (R6), different levels of permissions being needed.

```
/* R2: send a case record for classification */
a(database_download, DBID) ::
  requestCaseRecordByID(I) ⇐ a(clinician_classify, CID) ←
grantPermission(CID, I, Read, Normal_classify_from_local_site,
Local_database_read_policy_set) then
  caseRecord(R) ⇒ a(clinician_classify, CID) ←
getCaseRecordByID(I, R) then
  a(database_update, DBID)
/* R6: update a case record after classification */
a(database_update, DBID) ::
  updateCaseRecordByID(I) ⇐ a(clinician_followingdiagnosis, CID)
← grantPermission(CID, I, Update, Normal_update_from_local_site,
Local_database_update_policy_set) then
  caseRecordUpdated (Y) ⇒ a(clinician_followingdiagnosis, CID)
```

The notion *a(id, role)* defines the role one should play and its identity can be bound with workgroup memberships and professional specialists at runtime. Access policies, either identity-specific or role-related, as defined in clinical environment and governing access strategies, will be applied against the defined role play behaviour.

A clinician not in the right workgroup may be able to download a case but cannot update it. Similarly, a user identity is bound with professional specialists and constraints about their access to case partitions applied at runtime, e.g. only the named clinicians may update or write major diagnosis results; certain specialists may write reports in their areas; others on the case care list may only read those areas. Thus, a three dimension resource request of (user, resource, operation) will be constrained in two dimensions: user-resource must match workgroup membership (constraint 1) and user-operation match clinical information access needs (constraint 2).

It is at the point of checking the LCC constraint of "*grantPermission*" that user workgroup and speciality, as well as the case under consideration will be related (clinician identity of CID and case identity of I), prior to the required operation/interaction. The running and execution of LCC specification is supported by the OpenKnowledge kernel [8].

6. Future work

The HRB Centre for Primary Care Research is undertaking a 5-year programme with an aim of developing and evaluating information and communication technology (ICT) based decision support system that focuses on evidence-based practice in relation to improved diagnosis, clinical management and co-ordination. The major decision aids for patients rely upon the electronic clinical prediction rules (eCPRs). An example prediction rule as such, Centor Score [9], provides a practical tool that helps doctors to decide on the management of patients with sore throat or strep throat and whether an antibiotic is required. Using LCC, we can encode the relevant clinical practice guidelines or protocols and define the process in the following interaction model, enabling collaborative decision making in a distributed environment.

```
a(score_calculator, SCID)::
  patient_data(P) ⇐ a(GP_request(P), GPID) then
  score(s) ⇒ a(diagnosis_manager(P), DMID)
  ← s = temperature(P, higher, 38) + cough(P, no) + cervical(P,
swelling) + tonsils(P, swelling) + age(P, less, 15)

a(diagnosis_manager(P), DMID)::
  a(score_calculator, SCID ⇒ score(s)  then
  ( recommend(P, no_antibiotic) ⇒ a(GP_request(P), GPID)
← s = 0 or s = 1 )
  or
  ( recommend(P, antibiotic) ⇒ a(GP_request(P), GPID)
← (s = 2 or s = 3) and culture_result = positive )
  or
  ( recommend(P, antibiotic) ⇒ a(GP_request(P), GPID)
← s = 4 )
```

First, a GP sends the patient data including symptoms to a score calculator. The calculator determines the patient's total sore throat score by adding up points on the basis of relevant criteria. Then, the total score is sent by the calculator to a diagnosis manager. Its role is to give GP suggestions on the basis of the patient's total score. Finally, the doctor can make a decision based on the conclusions drawn from the interaction.

In this simple example, clinical prediction rules are explicitly embedded into the functions of two role players, score calculator and diagnosis manager. They take in patient symptoms and apply the clinical knowledge which was drawn from past experience against the current case. Imagine when a doctor is using a GP system with the symptoms and other related patient information available, such an interaction can be implicitly triggered and a diagnosis recommendation presented to the GP for decision making. The integration of these clinical prediction rules with current GP systems in a seamless healthcare workflow can greatly reduce the likelihood of misdiagnosis, evidence omission, or prescribing errors, and eventually improve the quality of care.

It is also conceivable that in the interaction model we show, the score calculator role may be played by one or more healthcare specialists, giving their part of conclusions to a series of evidence accumulation and decision support process in a bigger picture. Thus, a healthcare team with various specialists, clinical data and evidence holders, and clinical prediction rule functions can interact together and contribute to the final care delivery, in a pervasive setting. Selected healthcare data, being a subset (or full set) of the patient's record, is passed around to specialists with appropriate role for their opinions in different fields. Interaction models defined using LCC have the capability to express the knowledge flow, glue together the clinicians' expertise of the fields for collaboration, and the application of the supporting prediction rules at the right time at the right place. When a specialist is about to request data, his or her privilege to view/write to the entire patient record or part of it and further interpret it for clinical purpose is subject to the access policies associated with his or her ID/role.

The approach is promising and should be able to solve the complicated healthcare environment we are confronted with today in many countries: data, knowledge, and expertise being fragmented. Two major components coming into the picture help to build an effective and systematic approach: 1) clinical prediction rules take in the particular clinical symptoms and evidence on a patient-specific basis and suggest decisions or suggest intermediate evidence which later contributes to the final decisions with regard to this particular patient; 2) interaction models enable knowledge sharing, elicitation, and generation, and finally drawing clinical conclusions, with data protection regulations respected. In our research centre, the use of the approach as well as other methods will be further investigated in the next five years, with an eventual aim to improve the quality of Irish healthcare.

7. Acknowledgements

8. References

[1] Anderson RJ. Security in clinical information systems. British Medical Journal 1996.

[2] Anderson RJ. Undermining data privacy in health information. BMJ 2001; 322: 442-443.

[3] Joint Computer Group of the GMSC and RCGP. GMSC and RCGP guidelines for the extraction and use of data from general practitioner computer systems by organisations external to the practice. Appendix III in: Committee on Standards of Data Extraction from General Practice Guidelines; 1988.

[4] Hawker A. Confidentiality of personal information: a patient survey. Journal of Informatics in Primary Care 1995; 16-19.

[5] Sandhu RS, Coyne EJ, Feinstein HL, Youman CE. Role-Based Access Control Models. Computer 1996; 29(2): 38-47

[6] HealthAgents, http://www.healthagents.net/

[7] Xiao L, Lewis P, Gibb A. Developing a Security Protocol for a Distributed Decision Support System in a Healthcare Environment. In: 30th International Conference on Software Engineering; 2008. p. 673-682.

[8] Robertson D. et al. Open Knowledge: Semantic Webs Through Peer-to-Peer Interaction. OpenKnowledge Manifesto; 2006. http://www.openk.org/

[9] McIsaac WJ, White D, Tannenbaum D, Low DE. A clinical score to reduce unnecessary antibiotic use in patients with sore throat. CMAJ 1998; 158(1): 75-83.

[10] NHS Connecting for Health, "The Care Record Guarantee", 2007.