Conference Papers

tPOT: People Oriented Technology

2008-01-01

Patient-Centred Laboratory Validation Using Software Agents

John McGrory
Technological University Dublin, john.mcgrory@tudublin.ie

Jane Grimson
Trinity College, jane.grimson@tcd.ie

Frank Clarke
Technological University Dublin, frank.clarke@tudublin.ie

See next page for additional authors

Recommended Citation

Authors

John McGrory, Jane Grimson, Frank Clarke, and Peter Gaffney

PATIENT-CENTRED LABORATORY VALIDATION USING SOFTWARE AGENTS

John McGrory, Frank Clarke
Dublin Institute of Technology, Dublin 8, Ireland
john.mcgrory@dit.ie, frank.clarke@dit.ie

Jane Grimson
School of Computing, Trinity College Dublin, Dublin 2, Ireland
jane.grimson@tcd.ie

Peter Gaffney
Adelaide & Meath Hospital, Incorporating the National Children's Hospital (AMNCH), Tallaght, Dublin 24, Ireland
peter.gaffney@amnch.ie

Abstract: Guidelines are self-contained documents which healthcare professionals reference to obtain knowledge about a specific condition or process. They interface with these documents and apply known facts about specific patients to gain useful supportive information to aid in developing a diagnosis or manage a condition. To automate this process a series of Standard Operating Procedures (SOP) and workflow processes are constructed using the contents of these documents in order to manage the validation flow of a patient sample. These processes decompose the guidelines into workflow plans, which are then called using condition triggers controlled by a centralised management engine. The software BDI agent offers an alternative dynamic which more closely matches the modus operandi of narrative based medical guidelines. An agent's beliefs capture information attributes, plans capture the deliberative and action attributes, and desire captures the motivational attributes of the guideline in a self-contained autonomous software module. Agents acting on behalf of guidelines which overlap and interweave in similar domains can collaborate and coordinate in a loosely coupled fashion without the need for an all encompassing centralised plan.

1 INTRODUCTION

The primary role of a clinical laboratory is to support frontline healthcare professionals who are licensed to deal with patients (McLoughlin, 2006). Their function is to accept appropriate patient samples, analyse them, and report their findings back to the ordering clinician (Marshall et al., 1995). The reported results are not considered a diagnosis, but used by clinicians to deliver patient specific care. In the majority of cases these results are use to aid in planning treatments, quantifying medication amounts and monitoring patient responses, which could all have a detrimental affect on the patient if incorrect information was used (Witte et al., 1997). Therefore, the single most important activity performed by the laboratory technologists is to ensure their generated results are valid and plausible for the specific patient from whom the sample was taken.

Clinical guidelines are condition focused documents through which domain specific aims, goals, procedures, plans and normal reference ranges are disseminated to healthcare professionals. The purpose of these documents are to guide the reader, and streamline activities around a particular medical condition or process using evidence based supportive information. When a clinical or laboratory guideline is developed by an expert group they focus on best practice for the specific condition or process. They include all relevant knowledge,

logic and motivational aspects they deem necessary to adequately describe the domain.

Clinicians and laboratory technologists care for patients not diseases or processes, therefore it is their responsibility to filter through these guidelines acting on a patient's behalf. They must try interface with these documents, to make use of the maximum decision-making support for healthcare delivery based on the known facts about their individual patient. To automate this process of searching through guidelines on a patient's behalf, the laboratory technologists in association with the clinicians construct a series of Standard Operating Procedures (SOP), and workflow processes using the contents of these documents in order to manage the validation flow of a patient sample. This is accomplished by decomposing each guideline in to a series of separate workflow activity paths. Then develop a set of centralised management rules to link these activities based on the presented patient data. However, these procedures are not truly patient-centred but process-centred. The guidelines knowledge, logic and motivation can no longer be accessed as a standalone resource, but as a series of workflow triggers managed by a centralised software package, which no longer resembles the author's guideline. This process is fundamentally different to the true operation of a medical guideline, where guidelines are used to provide supporting information based on their holistic view of the domain, rather than a series of linked activities relating to a process. So is there an alternative approach where the process can be distributed and the guidelines retain logic, knowledge and motivation as a standalone self-contained unit.

Agent oriented architectures operate on similar principles to elements found in human decision-making by combining attributes (beliefs), methods (plans) and desires (goals). The BDI agent approach in particular is based on the principle of a belief capturing the informational attributes, the desire capturing motivational attributes and the intention capturing the deliberative attributes of an agent (Rao et al., 1995). Therefore, agents can be considered self-contained knowledge sources (KS), with a social communication interface and have the ability to act autonomously, or as part of a larger group. In research completed by the authors it was shown that a software agent can successfully capture and be encoded with the knowledge, logic and motivation of a guideline (McGrory_a et al., 2008). In additional research completed by the authors it was shown that although agent communications provide a facility to transmit data between agents, it

is also used to provide a social and collaborative aspect (McGrory_b et al., 2008). This allows the separate agents work in groups and collaborate on shared goals. This later research also demonstrated that agent communication was capable of being adapted to comply with a medical standard for communication (i.e. CEN ENV 13606-4:1999).

The thrust of this paper is to illustrate that software agents offer an alternative approach to reproduce the function of medical guidelines than the more commonly used centralised approaches. This paper also presents an operation of a framework which allows these agents coordinate and collaborate to validate a patient sample in a distributed fashion, without the need for a centralised all encompassing plan.

2 ANALYSIS OF GUIDELINE REPRESENTATION FORMATS

The traditional approach to combining separate bodies of knowledge (such as guidelines) together is to decompose the knowledge and logic into separate workflow activities and link these activities together using a centralised inference engine. Three commonly used techniques are rule base, direct coupling or blackboard systems.

The rule-base approaches are designed around a nodal tree, where expert knowledge in the form of a workflow activity is the branch, and the selecting of a particular branch at each node is based on patient information or process data. Although selecting rules based on presented facts during execution can be indicative of an illness, the rules which link them directly to a diagnosis do not reflect anything deeper than a casual understanding of human physiology. These systems are centralised and the original guideline knowledge is now absorbed within a labyrinth of rules.

Direct coupling architectures are made up of a group of separate expert knowledge modules. Each expert knowledge module contains local storage, a KS and a control switch to link to the other software modules according to their data-flow requirements using a direct call or link (Corkill, 2003). Complications arise when specific modules are subject to change and/or when the ordering of module control switching cannot be determined until run-time (Kavanagh et al., 2002). As the system expands and evolves the links change and the process becomes unwieldy and unmanageable. In addition to the aforementioned issues, the direct

coupling model does not provide a clear representation of the overall problem, and there is nothing more than relationship links used.

The blackboard model is based around three components: KS, control element and the blackboard (Turban et al., 2005). The KS is an expert at solving specific elements of the overall problem. The blackboard, acts as a central repository for data, partial solutions and control information. The blackboard also acts as a communication medium for the transfer of information, and a KS triggering mechanism. The control element directs the problem-solving process by allowing KS's to respond to blackboard changes, and it selects the most appropriate KS to be executed next, as shown in Figure 1. After completing a task the KS reports back to the blackboard and returns control to the control element. KS's are not aware, and cannot communicate with other KS's directly. They know nothing about the other experts (e.g., what parameters they use, what processes they perform, or what services they provide). The blackboard architecture tends to be a labyrinth of different configurations, levels of abstraction, and partial solutions which are orchestrated to provide a flexible problem solving mechanism. The blackboard system eliminates the communication issues raised by the directly coupled monolithic model, and gives a representation of the problem to be solved to all participants. But the blackboard does not have the capacity to indicate how group members can collaborate to solve a problem, but can only select from partial solutions it already possesses.
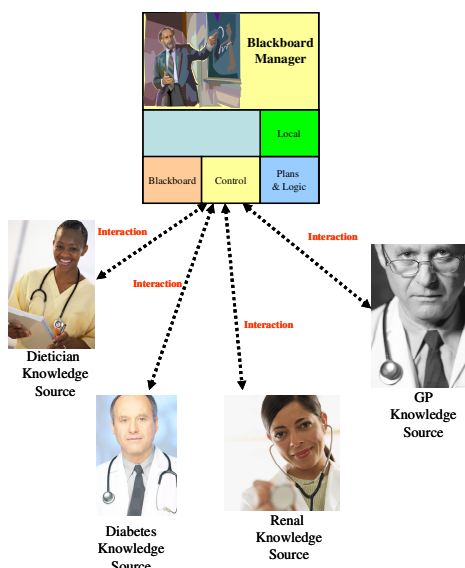

Figure 1, Blackboard Management Communications

3 DESIGN OF A GUIDELINE AGENT

The agent approach is based on the principle that each agent can represent a single guideline. It captures all the guidelines knowledge, logic and motivation. In addition to this the agent has a rich communication facility where data and social interaction between separate agents can take place. But how does each agent know what the other is doing since there is no centralised all encompassing plan? If the blackboard is simply a repository of information, although layered to some degree, there is no absolute necessity for it to be in a single location. Therefore, it is possible to replicate a copy of the blackboard within each Autonomous Socialising Knowledge agent (ASK-agent) as shown in Figure 2.


Figure 2, Autonomous Socialising Knowledge agent model

Each ASK-agent now contains a localised blackboard, knowledge source, localised beliefs, localised control and its own inference engine. The retention of the motivational component of the guideline within the ASK-agent is fundamentally different to the centralised approach of other systems. This allows the ASK-agent to act autonomously on behalf of the guideline in a self-contained capacity. When patient specific information is presented to the individual agents, they have the ability to apply their encoded knowledge and logic, and provide a supportive response based solely on that information. Using this approach an ASK-agent module can make use of the maximum supportive response from the other separate ASK-agent's based on the known facts

about the individual patient. By providing a framework which allows separate ASK-agent broadcast supportive communications to each other, the agent approach offers the opportunity for the data to be validated in a patient-centred fashion. But how can these separate, autonomous, self-contained ASK-agent modules share data, work in groups or collaborate to solve a problem.

3.1 Agent to agent based activity

A theory which can aid issues relating to collaborating guidelines is Activity Theory (AT) which emanated through the social sciences. AT focuses on the collaborative nature of separate autonomous systems such as individuals (Engestrom et al., 1999), on which agents are based, and have the capability to perform certain tasks as part of a group. Agents synthesise human decision-making through their goal, plan and belief elements, but do not explicitly detail how they can socialise or collaborate. AT in itself does not provide an output which can be exactly transposed into computer software, but does provide a useful framework based around interfacing interaction and collaboration of software modules. These interfaces can be used to develop an increased sense of interaction and collaboration ability in autonomous modules using a software program independent approach.

The structure of human activity according to Engestrom can be compartmentalised using *rules*, *community*, *subject*, *object*, *division of labour* and *instruments*. AT is an iterative process where an activity is developed from a simple low level activity to a higher level activity.
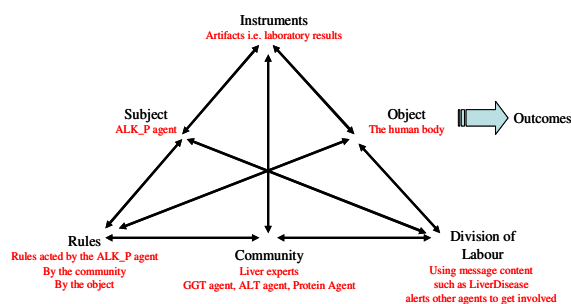
Figure 3, Low level activity

In terms of medical guidelines the iterative process dynamic exists by virtue of the design of guidelines and their focus on a condition, disease or organ. Consider for example, a low level implementation activity being the guideline behind the validation of a single analyte result, say Alkaline

Phosphates as shown in Figure 3. A higher level implementation is where the result is combined with some other single analyte results, such as Bilirubin and GGT, to perform and aid in the reporting of a Liver Function Test as shown in Figure 4. The Liver Function Test is then part of a higher level suite of tests for other medical disorder classifications.
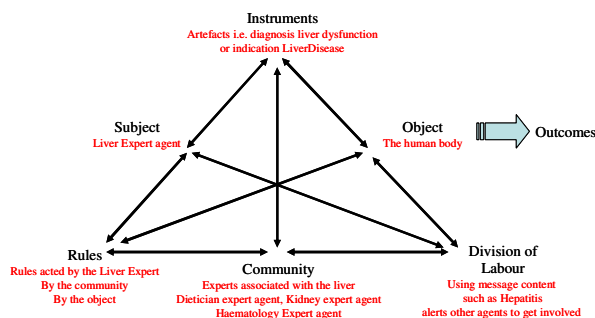
Figure 4, High level activity

One guideline does not cover the whole body, but more specifically focuses on an abstract conceptualisation of body components (e.g., liver function in a group of male diabetes patients). Another guideline relates to the same body component but from a different abstract conceptualisation viewpoint (e.g., kidney function in a group of male diabetes patients). Although the two guidelines are separate autonomous documents they are linked by virtue of their domain of discourse. Therefore, a link between different guidelines already exists within the guideline document itself. The overlapping knowledge is provided in two main forms. The first is in the form of similar domain knowledge that uses alternative inference mechanisms in order to derive a result (i.e. both statistical and rule-based inference engines being able to validate the same result). The second is in the form of overlapping knowledge which observes different viewpoints of the same domain. For example, the kidney filters toxins from the blood passing it to the urinary tract. As the kidney is such an integrated organ in the body there are many guidelines describing its operation from different viewpoints such as blood filtering, urinary tract, autoimmune disorders etc. Using this approach the organ disease or condition is described from different viewpoints through various guidelines. Each guideline describes different knowledge, logic and motivational aspects associated with the organ. Therefore supportive information can be exchanged between these guidelines in order to aid in describing the operation of the organ, or in the validation of a sample result.

3.2 ASK-agent model

To utilise this overlapping knowledge link a social structure was developed to manage the interfacing between agents. This social interface took the form of a mandatory set of searchable service descriptions, beliefs and actions. The service descriptions (i.e. *Name, Type, Ownership, GuidelineReference, InformationNeeded, ValidationType, EndResultType, Ontology* and *Language*) permitted each agent to be located within the agency platform through the Directory Facilitator (DF) (a feature of the Foundation for Intelligent Physical Agents (FIPA) standard offering searchable *goldenpages* facility to locate agents) (McGrory_a et al., 2008). The beliefs (i.e. *CurrentlyValidating, PlausibilityScore* and localised blackboard) permit the ASK-agent to interact with other group members. The actions relate to automated responses the ASK-agent must return to other agents when queried (e.g. *CurrentlyValidating*), and the sending of information to other agents it believes should be reported (e.g. it determined the presence of liver disease during its deliberation). Therefore, each ASK-agent only needs to know its overlapping neighbours, which it can find and interact with using the agent platforms DF and message passing. With access to supportive and overlapping knowledge it is not necessary to have a single all-encompassing rule set to manage the ASK-agents interaction.

The fundamental concept of the ASK-agent system proposed in this paper is to allow components to collaborate and share supportive information without having to explicitly disclose their position as part of the large encompassing community. An ASK-agent does not need to identify exactly what every other agent is doing; only what its neighbours (i.e. neighbours it interfaces with) are doing. To illustrate this point further and demonstrate some boundaries, consider the example of a jigsaw with 500 pieces. A jigsaw piece has two discrete dimensions: the irregular shaped edge containing four sides, and the image printed on the face. To solve the puzzle, a person directly matches individual jigsaw pieces onto the jigsaw image, say the image shown on the box. Each piece is identified using the image on its face and placed in the appropriate position. This method requires a view of the whole system to be presented before starting, but involves no greater skill than straightforward pattern matching. An alternative approach is to use a combination of the localised image on the face of the

piece and its four corners to match it to a suitable neighbour (i.e. matching the shape of the pieces together). Jigsaw assembly using these interfaces do not require the full picture to be known. Using the jigsaw example as a solution metaphor, the heart and lungs image depicted in Figure 4 is a symbolic representation of the heart and lungs as a whole, not just the image it represents. The agent was not intended to be a large all encompassing structure, but a group of loosely coupled autonomous expert knowledge sources (represented by each jigsaw piece) which could be readily and easily interfaced with as shown in Figure 4. The ASK-agent only needs to know its neighbours (i.e. the expert it interfaces with), in a similar way the jigsaw piece only needs to know another piece with similar edge profile and compatible image, not the whole picture. The ASK-agent does not need to know anything about any other piece of the jigsaw only its interfacing neighbours. The interface can be considered the ontology, overlapping facts, common laboratory results and various viewpoints of the universe of discourse. This is analogous to the jigsaw edge shape profile.


Figure 4, Jigsaw metaphor representing agent components of the heart and lungs

4 CONCLUSION

This research demonstrates the agency approach offers a facility to manage and interface with medical guidelines electronically, in a similar modus operandi to original guideline documents. This is because of the synergy between the knowledge base, plans, decisions, action, goals and the self-contained nature components between guidelines and agents.

Element	Centralised Approach	ASK-Agent
Multi-ontologies	No	Yes
Processing	Centralised	Distributed
Requirement for overlapping knowledge	None	Required to provide the links.
Addition, altering or removal of guidelines from the system	Any changes require centralised inference to be recompiled.	Each ASK-agent is independent and loaded separately.
System resilience	None	Yes, all ASK-agents have a copy of the blackboard.
Independent accessible knowledge	None. All access to information through the centralised engine.	Yes, all ASK-agents are independent. Information accessed via message passing.
Clinician having access to specific guideline knowledge	A clinician cannot access knowledge directly.	A clinician can access each ASK-agent via a message and directly access the specific guideline knowledge.
Method of collaboration	Direct links using the centralised engine.	Using the mandatory beliefs, action and descriptions in the agent platforms Directory Facilitator.

Table 1, Summary of centralised and ASK-agent approach.

The agents can be encoded to reproduce the beliefs, desires and intentions of the narrative guideline and act accurately, faithfully and autonomously on behalf of that document. This body of knowledge and logic can then be interfaced with, whenever that information needs to be accessed. The addition of activity theory and in particular the iteration model concept showed that the guideline documents already contain aspects that link them together. Using these links and the developed social communication the ASK-agents can locate, access, communicate, collaborate and coordinate activities between each other. This allows supportive information exchanges to be completed between separate expert agents about an individual patient, without the need for an all encompassing centralised plan. In cases where there is an inconsistency in held patient specific information, this agent approach offers an advanced, robust and efficient patient centred validation alternative to existing approaches. However, if overlapping knowledge between guidelines is not available the links created using this approach are not present and the separate guidelines are standalone islands of information. The guidelines knowledge, logic and motivations are still accessible as a standalone entity, but other agents would need to be created to provide the links. Developing a system using the latter approach still permits distributed processing to be accomplished, but not without a source of knowledge to provide the links. A summary of the differences between the centralised and ASK-agent approach are given in Table 1.

REFERENCES

Corkill D., 2003. *Collaborating Software Blackboard and Multi-Agent Systems & the Future*. Accessed via www.cs.umass.edu, Proceedings of the International Lisp Conference published in October 2003.

Engestrom Y., Miettinen R., Punamaki R. *Perspectives on Activity Theory*. Published by Cambridge University Press January 1999. ISBN13: 9780521437301.

Kavanagh M., Price S., "*The quest for a computerized guideline standard: The process, its history, and an evaluation of the most common and promising methods used today*", Capstone Project documentation 2002.

Marshall W.J., Bangert S.K., 1995. *Clinical Biochemistry, Metabolic and Clinical Aspects*, ISBN 0 443 043 418, published in 1995 by Churchill Livingstone.

McGrory_a J., Grimson J., Clarke F., Gaffney P., "*Agents representing medical guidelines*", HEALTHINF 2008 International Conference on Health informatics, Funchal, Madeira, Portugal, January 2008.

McGrory_b J., Grimson J., Clarke F., Gaffney P., "*Communication of medical information using agents*", HEALTHINF 2008 International Conference on Health informatics, Funchal, Madeira, Portugal, January 2008.

McLoughlin V., Millar J., Mattke S., Franca M., Jonsson P., Somekh D. and Bates D., 2006. *Selecting indicators for patient safety at the health system level in OECD countries*. International Journal for Quality in Health Care 2006 18(Supplement 1):14-20.

Rao A., Georgeff M., 1995. *BDI Agents: From Theory to Practice,* proceedings of the first international conference on multi agent systems (ICMMAS-95) San Francisco USA 1995.

Turban E., Aronson J.E., Liang TP., *Decision Support Systems and Intelligent Systems*. published in 2005 by Pearson Prentice Hall, Seven Edition, ISBN 0-13-123013-1, Chapter 10, pgs 538-570.

Witte D.L., VanNess S., Angstadt D., and Pennell B.J., 1997. *Errors, mistakes, blunders, outliers, or unacceptable results: how many?* Published by Clinical Chemistry 1997; Vol:43 pgs 1352-1356.