Sources of Complexity in the Design of Healthcare Systems: Autonomy vs. Governance

A. Taleb-Bendiab
David England
Martin Randles
Phil Miseldine
Karen Murphy

School of Computing and Mathematical Sciences
Liverpool John Moores University
Byrom St
Liverpool L3 3AF

Contact email: d.england@livjm.ac.uk
http://www.cms.livjm.ac.uk/2nrich

Abstract: In this paper we look at decision support for post-operative breast cancer care. Our main concerns are to support autonomy of decision making whilst maintaining the governance and reliability of the decision making process. We describe the context of our work in the wider medical setting. We then present a set of decision support tools based on the situation calculus as a means of maintaining the integrity of rule bases underlying the decision making system

Keywords: decision support, breast cancer care, autonomy, self-governing systems

Introduction

Our project on decision support for post-operative breast cancer care raises a number of interdisciplinary questions in a complex and emotive area. The project is collaboration between computer scientists, statisticians and clinicians which itself is a complex arrangement. The nature of the subject involves life-threatening decisions. Clinicians and patients are faced with the most difficult decisions. From an HCI viewpoint we are also faced with supporting and not supplanting clinician's judgments. There are also wider issues of the implications of our studies of historical clinical data and what that might mean for future approaches to prognosis. In an earlier paper we discussed our approach to process understanding in breast cancer care. We described how decisions on post-operative breast cancer treatment are currently governed by a set of medical guidelines including the National Institute for Clinical Evidence (NICE) guidelines (Nice 2005) in which the decision process is as follows: the clinician uses available data with a staging method to define the patient risk category. The risk category is then used to determine the type of treatment that the patient would be offered. The guidelines approach has been investigated by Woolf (Woolf 1999) who concluded that guidelines can cause benefit or harm and that the clinical guidelines need to be "*Rigorously developed evidence based guidelines [to] minimize the potential harms*".

The aim of our project is three-fold.

1. To understand and support the clinician's decision making process within a set (or sets) of clinical guidelines;
2. To analyze historical data of breast cancer care to produce new rules for treatment choice.
3. To bring these two processes together so that treatment decisions can be guided by an evidence-based computer system.

The evidence-based approach to the delivery of medical care has gained wide recognition within the healthcare community, advocating that decision-making should use current knowledge and clinical evidence from systematic research (Rosenberg 1995). In breast cancer care, there are currently a few staging methods in widespread use by clinicians, namely the Nottingham and Manchester staging systems. However, there is no standard method to support oncologists' decision-making processes as to how and when to include new evidence, and how to validate emerging local patient data patterns or other models and practices from elsewhere. This means that there is no standard way to ensure that clinicians are following accepted guidelines or deviating from them. There may be valid clinical reasons why a standard decision path is not chosen (e.g. the age or infirmity of the patient) but these decisions are not currently recorded in a standard way.

In this paper we wish to address some of the sources of complexity in the design of healthcare systems. More specifically we are interested in safe and reliable decision support for post-operative breast cancer care and the wider lessons we can learn from our experiences. There are many contributory factors to the complexity of designing these systems, some of which we shall discuss below. These include:

- Environment and Context: The aims of national initiatives such as NPFit/PACIT in driving down costs and errors etc.
- The Autonomy of clinicians versus the governance requirements of clinical audit
- Resource limitations – staffing drugs, radiology etc
- The limitations of Medical Knowledge and how that knowledge evolves
- The co-evolution of knowledge and the needs for system validation and safety
- The Requirements for safe solutions and effective treatment
- The (moving) Requirements for knowledge elicitation
- The abilities of Computer Scientists to encode medical knowledge
- The limitations of data mining approaches as a means to supporting evidence-based, decision making.

We will concentrate on our approach to tool support of the decision complexities using the situation calculus (Reiter 1991).

National Programmes

In both the UK and US there are national initiatives to introduce greater use of IT in clinical settings. The broad aims of the NPFit (UK) and PACIT (USA) programmes are similar. They aim to streamline data processing to cut costs and reduce clinical errors. For example, it is proposed that electronic prescribing of medicines will cut costs in paperwork and reduce prescribing errors which account for a large number of patient deaths (44,000 to 98,000 deaths caused by medical errors in the USA). Both schemes aim to introduce electronic patient records, again to cut costs of paper records and reduce errors from paper-based systems. Both systems also look to more clinical governance and audit of medical processes so that medical staff are more accountable for their actions. The UK initiative is already displaying the signs of a large project out of control with the projected costs of £6Bn rising to between £18Bn and £31Bn. The lack of user centred design is evident by a recent (BBC) poll showing 75% of family doctors are not certain that NPFit will ever meets its goals. The first stage of the electronic appointment systems has largely failed to meets its use targets. However, a smaller scale introduction of region-wide IT in the Wirral was more widely accepted with 90% of family surgeries and the vast number of patients accepting the system. Thus IT systems can succeed. This is important for our work, for in order to succeed, it requires a working IT health infrastructure. Furthermore the twin goals of cost and error reduction may be mutually incompatible. As Reason points out (Reason 1997) organisations have processes for productivity and safety but circumstances will arise, either through unsafe acts or latent system weaknesses, which lead to organisational failure. Safety protocols may be violated in the name of efficiency or sets of latent weaknesses will line up to cause an accident. Many individual errors are the result of cognitive under-specification (Reason 1990) of the user's tasks. In our project we aim to over-specify and support clinical tasks by describing them in the situation calculus. This will provide a robust means of supporting decision making and ensuring that chances to decisions protocols remain valid.

Medical Knowledge and Computer Science

In a multidisciplinary project settings involving; clinicians, statisticians, computer scientists and public health specialists, our project has started by understanding the current decision-making practices as a prelude to systems' implementations. This will be evaluated using a set of small-scale controlled trials involving both patients and clinicians. The proposed method, unlike traditional decision-making techniques, including multi-criteria, will provide breast cancer clinicians and patients with a flexible decision framework adaptive to their decision practices. It will also allow for evolutions of decision models, decision resources (data) and users concerns. This novel approach will provide important insights into the development of an integrated decision support infrastructure for high assurance decision activities, which will directly contribute to one of the NHS R&D high priority area of "Medical Devices Directives for cancer patient care".

To model, validate and enact clinical guidelines a new approach, using ideas originating from research in distributed artificial intelligence, has been developed. In this formalisation the treatment recommendation process is conceived of as a human/artificial multi-agent system. The actions of the agents are governed by system norms. These norms constrain the behaviour of the actors in the system. They may relate to the control of the system in maintaining availability, fault tolerance, adaptability etc. for quality of service or they may be concerned with the output results such as guideline compliance for quality of process. In any event the goal is complete system governance within safe limits, encompassing clinical governance. It is a complex process, in itself, to produce safety critical systems. However in needing to allow for the variability and variety of clinicians' usage adds another level of complexity. The autonomy that needs to be retained by the clinician is tempered by the constraints of clinical governance. Thus if the system is required to take control of management functions the autonomy of the system's agents is compromised. However the agents are autonomous, rational and social so computational economy is best served by the agents following the norms. Norms arise from a social situation in that they involve providing rules for a community of more than one individual (Boman 1999). The norms can be associated with producing decisions, as a practical application, or with the governance of the system and clinical processes. The provision of adjustable autonomous agent behaviour is necessary to deliver the system services in a way where the majority of the administrative tasks are handled within the software itself whilst the autonomy is reconciled with the governance of the system processes. This is achieved by taking an aspect oriented approach and separating the concerns of governance from the application service concerns.

In order to logically model and reason, within this approach, the formalism of the Situation Calculus (McCarthy 1968) was used. The Situation Calculus is especially suited to the analysis of dynamic systems because there is no need for the prior enumeration of the state space so unexpected external events can be accommodated. What-if scenarios can be modelled with branching time-lines and counterfactual reasoning. Thus a formalism is provided to model both system behaviour, in mapping a set of symptoms to a set of treatment options and deliberative governance strictures so that formal reasoning techniques, such as deduction, induction or abduction, can be applied to analyse the meta-control of the system.

Although a full explanation of the Situation Calculus is outside the scope of this paper a full specification can be found in Levesque 1998. Briefly stated, the Situation Calculus is based on actions. A situation is an action history emanating from an initial situation, defined by a set of function values called fluents. Successive situations are defined by effect axioms, for the action, on the fluents. Successor state axioms together with action precondition axioms give a complete representation that mostly solves the frame problem (Reiter 1991).

A custom scripting language, JBel (Miseldine, 2004) has been developed to facilitate the deployment of, the generated, Situation Calculus defined, self-governing decision agent in a grid based architecture –

Clouds (figure 1). The Clouds concept is conceived as a system with fluid and flexible boundaries that can interact and merge with similar system architectures.

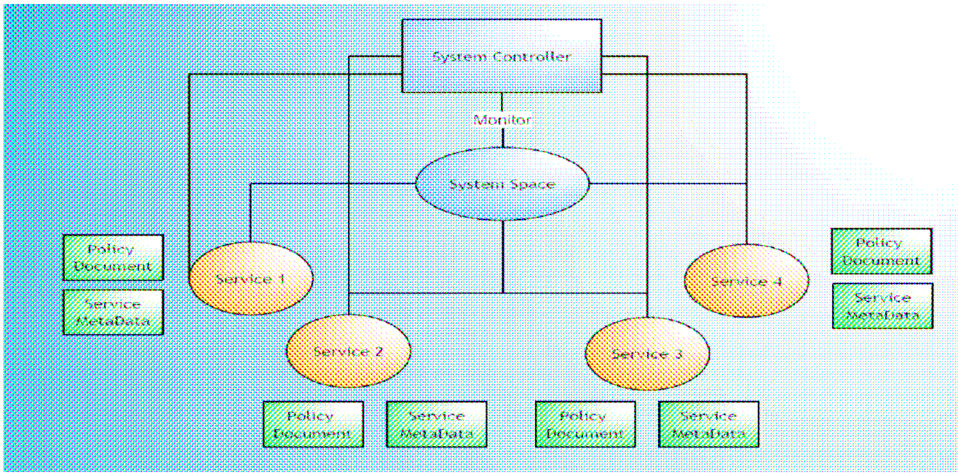

Figure 1. Clouds Architecture

The Cloud can be thought of as a federation of services (agents) and resources controlled by the system controller and discovered through the system space. The system space provides persistent data storage for service registration and state information giving the means to coordinate the application service activities. The system controller controls access to and from the individual services and resources, within a Cloud. It brokers requests to the system based on system status and governance rules, in JBel objects, derived from the logical normative deliberative process. The system controller acts as an interface to the Cloud. It can function in two roles, either as an abstraction that inspects calls between the System Space and services, or as a monitor that analyses the state information stored within the system space.

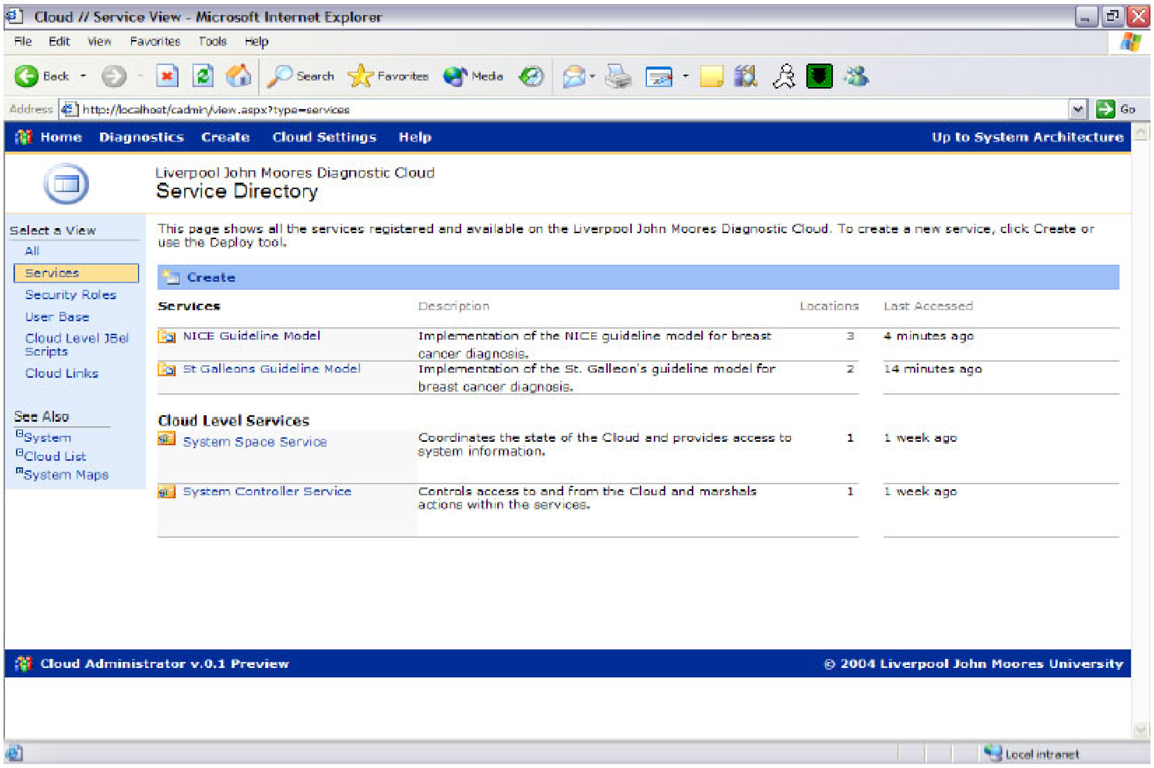

Figure 2. Clouds Decision Rule Governance Tool

In the current research prototype the Cloud environment (figure 2) produces an interactive dialogue in which the clinician enters the clinical data and is given one or more recommended treatments. At this stage of development we are looking at two approaches to critiquing the recommended outcomes. Firstly, by using a critiquing user interface (Gray 1996) where the clinician enters the data and the proposed treatment and the system returns with a recommended treatment. Secondly, where "borderline" or difficult treatment cases are explored using more than one set of clinical guidelines to see if they converge or diverge in their recommendations. This is a similar approach to voting systems used in safety critical systems (Mackie 2000). Critiquing at the user interface involves guidance by one knowledge base whereas the second form of critiquing requires the use of multiple *and* different knowledge bases.

The development of adaptive software, in healthcare systems, relies on the specification and abstraction of norms that collectively constitute a decision process as expressed logically through Situation Calculus. In addition, to allow safe re-engineering in the distributed context, the representation of such norms must be designed so that they can be adapted at runtime without alteration to their published interface signature to maintain modularity.

The methodology, used in JBel, to express behavioural rules within an introspective framework at runtime is structured in terms of conditional statements, and variable assignments which are parsed and evaluated from script form into object form, making it suitable for representing Situation Calculus rules. The resulting JBel object encapsulates the logic and assignments expressed within the script in object notation, allowing it to be inspected, modified, recompiled, and re-evaluated at runtime.

Object serialisation techniques, allowing an object to be stored, transferred and retrieved from information repositories, can be applied to the JBel objects, allowing the safe and efficient distribution of decision models.

The design of decision processes involves a process workflow and behavioural rules that determine the movement through the process within the model. Decisions within the workflow are linked to specific behaviour defined within the JBel script, with the behaviour determining the future flow of the model. Thus, with decision logic encapsulated within a JBel Script, changes to the structure of the model are separated from changes to its logical construction, yielding separate and abstracted decision model architectures. This high level of runtime accessibility permits the complete separation of the presentation layer with the underlying process model.

Behavioural rules collectively assemble a profile of accepted behaviour for the system they describe by associating logical conditions with a set of actions to take place upon their success or failure. In the context of medical decision support, such profiling allows formal modelling of clinician concerns, introducing high level assurance for their fulfillment.

In conjunction with the rule-based approach we also have a data mining interface which looks at historical patient data. Our statistician colleagues have used a neural networks approach to find patterns in breast cancer survival data (Jarman 2004). The outcomes from this analysis have been fed into a rule induction package to produce a set of human-understandable rules. There are two intended usages of the data-mining interface in Clouds. Firstly, as a support for clinical governance, in that, we can compare the explicit rules from published guidelines with the induced rules and see if the published rules are being followed in practice. Secondly, we can perform "What-if?" analyses so that, within the patient/clinician dialogue about treatment, a doctor can illustrate to the patient that there are other patients in the historical data with a profile similar to theirs. Thus they can discuss alternative treatments in cases where either, the decision system outcomes are ambiguous, or the patient's situation requires a different treatment from that recommended. We are currently modelling the clinical-patient dialogue within Clouds so that we can record alternative outcomes and their reasoning.

As an example from the Situation Calculus a treatment decision may be handled via the successor state axiom:

$$\begin{aligned} &NICEtreatment(patient,\ tamoxifen, do(a,s)) \Leftrightarrow [\ NICEtreatment(patient,\ tamoxifen,s) \wedge \\ &\quad \neg\exists treatment(\ a=nice\_treament\_decision(patient,\ treatment) \wedge \\ &\quad (\ treatment \neq tamoxifen))] \vee \\ &\quad [a=nice\_treatment\_decision(patient,tamoxifen)] \end{aligned}$$

with the action precondition axiom:

$$\begin{aligned} &poss(nice\_treatment\_decision(patient,tamoxifen),s) \Rightarrow (oesreceptor(patient,s)=pos) \wedge \\ &\quad (menostatus(patient,s)=post) \end{aligned}$$

It is then possible to log compliance to provide data for system update and evolution:

$$\begin{aligned} &compliance(patient,\ treatment,service\_decision,\ do(a,\ s)) \Leftrightarrow \\ &\quad [\ compliance(patient,\ \ treatment,\ service\_decision,\ s) \wedge \\ &\quad a \neq treament\_decision(patient,\ treatment1)] \vee \\ &\quad [a=treatment\_decision(patient,treatment) \wedge service\text{-}decision(patient,\ s)=treatment] \end{aligned}$$

This is the method by which the system representation and reasoning model can be realised for both self-governance and application level services. The service level concerns, the NICE guidelines in the above formalism, are handled separately from the governance concerns of process quality monitoring.

However in both cases the deliberation is modelled in the Situation Calculus then directly implemented into the JBel scripts (figure 3). The deliberation required for the governance allows the services (agents) to act autonomously within safe limits whilst the deliberation to produce a guideline-based decision is completely specified by the rule base derived from the NICE guidelines.

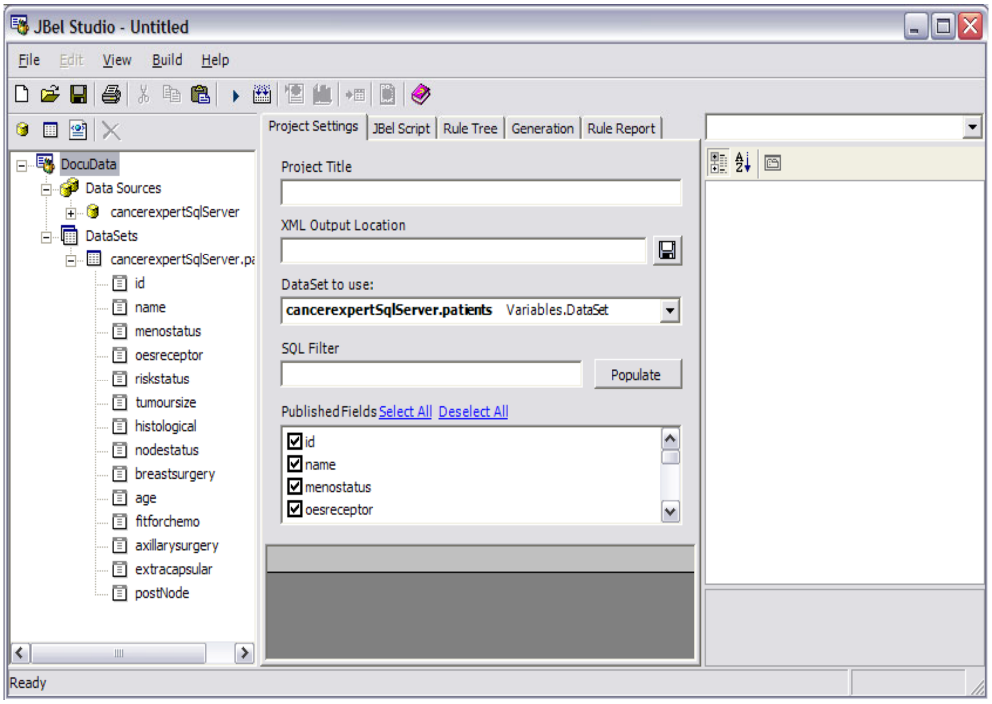

Figure 3 The JBel IDE

Conclusions

There are a number of remaining challenges to the successful deployment of the kind of system we are proposing. Practically there are implementation problems and issues with current access to IT by clinicians and their team members. Some of these issues may be resolved by the NPFit programme (NPFit 2004). However, we have already noted above problems with cost overruns, delays and issues with clinical acceptance of NPFit and its various constituent systems. From a safety viewpoint there are several issues with how we maintain the validity of our rules, whether induced or derived from guidelines, as NICE guidelines are update every five-year. Within the lifetime of the project there have been new medical developments which have been used in clinical trials. We need to ensure that our system continues to give safe recommendations after any rule updated, and the situation calculus can be used to specify and validate the dynamics of such open systems (Randles 2004).

In the wider medical viewpoint we can see applications for our approach in other areas of medicine, such as the diagnosis and treatment of lymphoma which has a similar staging model to breast cancer. We are also exploring the use of self-governing rule sets in Dentistry where the system would be used to module dentistry protocols. More broadly still we are looking at the use of our approach in the general area of self-governing and autonomous systems in a wide range of settings requiring ambient intelligence.

References

Boman M, "Norms in artificial decision making" AI and Law, 1999.

England D A. Taleb-Bendiab, A, Lisboa, P, Murphy, K, Jarman, I, "Decision Support for Post-Operative Breast Cancer Care", Coping with Complexity: Sharing New Approaches for the Design of Human-Computer Systems in Complex Settings, University of Bath, 2004

Gray PD, Draper SW, "A Unified Concept of Style and its Place in User Interface Design", Proceedings of BCS HCI 96 Conference, Cambridge University Press, 1996

Jarman, I, Lisboa, P J, "A comparative study of NPI and PLANN-ARD by prognostic accuracy and treatment allocation for breast cancer patients", ENRICH Deliverables, http://www.cms.livjm.ac.uk/2nrich/deliverables2.asp, 2004.

Levesque H. J., Pirri F. and Reiter R. "Foundations for the situation calculus". Linköping Electronic Articles in Computer and Information Science, Vol. 3(1998): nr 18. http://www.ep.liu.se/ea/cis/1998/018/

McCarthy J. and Hayes P. "Some philosophical problems from the standpoint of artificial intelligence". Machine Intelligence 4, ed: B. Meltzer and D. Michie, pp 463-502, Edinburgh University Press, Edinburgh, Scotland 1968

Mackie J, Sommerville I, "Failures of Healthcare Systems", Proceedings of the First Dependability IRC Workshop, Edinburgh, United Kingdom, March 22-23, 2000

Miseldine, P. JBel Application Development Guide, 2nrich Project. http://www.cms.livjm.ac.uk/2nrich/deliverables.asp. 2004

NICE, National Institute for Clinical Excellence, www.nice.org.uk, 2005

NPFit, National Programme for IT in the NHS, www.npfit.nhs.gov.uk, 2004

PITAC - President's Information Technology Advisory Committee, Revolutionizing Health Care Through Information Technology, June 2004

Randles M, Taleb-Bendiab A, Miseldine P, "A Stochastic Situation Calculus Modelling Approach for Autonomic Middleware", Proceedings of PGNet 2004, Liverpool John Moores University, 2004

Reason, J, "Human Error", Cambridge University Press, 1990

Reason J, "Managing the Risks of Organisational Accidents", Ashgate, 1997

Reiter R. "The frame problem in the situation calculus: a simple solution (sometimes) and a complete result for goal regression". Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy, ed: V. Lifschitz, pp359-380, Academic Press, San Diego, California 1991

Rosenberg W, McDonald A, "Evidence based medicine: an approach to clinical problem-solving", BMJ 1995;310:1122-1126

Steve H Woolf, Richard Grol, Allen Hutchinson, Martin Eccles, Jeremy Grimshaw, "Potential benefits, limitations, and harms of clinical guidelines", BMJ 1999;318:527-530

ACKNOWLEDGEMENTS

This project is collaboration between Liverpool John Moores University, the Christie Hospital, Manchester and the Linda McCartney Centre of the Royal Liverpool Hospital. It is funded by the EPSRC.