Developing an Open and Adaptive Agent Architecture to Support Multidisciplinary Decision Making

Liang Xiao
Hubei University of Technology
lx@mail.hbut.edu.cn

John Fox
University of Oxford
john.fox@eng.ox.ac.uk

Hong Zhu
Oxford Brookes University
hzhu@brookes.ac.uk

Abstract—Patient care is becoming increasingly complex and multidisciplinary for many conditions, notably cancer and chronic diseases, in which a care team participates in and shares responsibility for the patient's care. Providing IT support for joint clinical decision making in a distributed environment raises some unique challenges: 1) clinical specialists, located in their own working environments, need to be able to group together wherever necessary; 2) new clinical evidence and guidelines, published by healthcare authorities and subject to continuous revision, need to be shared and enacted by the care team, as automatically as possible; 3) decision points, distributed in the environment, need to refer consistently the same set of guidelines. In this paper we propose an open and adaptive Multi-Agent System architecture to resolve these challenges.

Keywords—Multidisciplinary Decision Making; Multi-Agent Systems; Open and Adaptive Software Architecture

1 Introduction & Motivation

1.1 Multidisciplinary healthcare and collaborative decision-making

It is now recognised that *Clinical Decision Support Systems* (CDSS) could effectively retrieve up-to-date medical knowledge and help to interpret clinical data at the point of care, and this may assist clinicians in keeping their knowledge up to date and improving compliance of clinical decision-making with evidence-based guidelines. The need and demand for this capability from both managers and patients are increasing [5]. The work on CDSS has been reported in a number of literatures, some are found more successful within the limit of specific sites and for specific needs but many fail to achieve the promised improvement in clinical outcome in routine use. The design, development and delivery of CDSS remain a grand challenge, as identified by leading experts in this field [10].

To complicate matters, modern healthcare practice has seen *joint decision-making* becoming more common, as groups of specialists are increasingly involved in and share responsibilities for care. In complex diseases such as cancer, there can be many significant decision points, the responsibilities for which may be distributed among specialist doctors, nursing staff, GPs and even patients themselves. In modelling a breast cancer diagnosis and treatment pathway, for example, we have found that there may be 65 or more significant decision points, each of which may lead on to many distinct tasks and workflows to achieve clinical goals [7] [8].

As the actors participate in a patient's care from different places and at different times, decision-making often needs to be supported in a coordinated and collaborative way. An *open and adaptive system architecture* is required in order to flexibly choreograph workflows and decision-making throughout a "patient journey", from detection and diagnosis to treatment planning, management and follow-up. Major challenges arise in delivering this kind of architecture, including

1. Clinical evidence and guidelines for addressing any given clinical problem are continuously increasing and improving (see the controversy and criticism about NICE guidelines in [11] [12]). Such new knowledge needs to be rapidly disseminated and used – though often isn't at present.
2. Clinical collaboration is required as some clinical objectives may be difficult to achieve by a single clinician working alone.
3. Collaborations are "learning opportunities". Data about successful practices should be accumulated and knowledge extracted for future reuse.

A major goal of our research is to design an open and adaptive software architecture that addresses the above challenges while providing support to knowledge dissemination and decision-making. The CDSS of this architecture will be more versatile than its conventional counterparts, and capable of incorporating new clinical guidelines and adapting to changing collaboration structure wherever necessary and as automatically as possible.

We introduce in Section 2 a generic agent paradigm towards collaborative decision support. A set of agent-oriented design patterns supporting such a unified framework are discussed in Section 3. In Section 4 we demonstrate the application of the approach using a breast cancer referral example. Finally we conclude the paper in Section 5.

2 GNAPB: A generic paradigm for coordinated agent conduct and decision making in society

An agent society resembles a human society in having similar notions of inter-relating norms which govern the conducts of individuals or groups, as in collaborative decision-making. An agent is seen here as being composed of five parts, in which **Norms** apply across the entire society, **Agreements** between groups or organisations, and **Goals**, **Plans** and **Beliefs** are relevant to individuals.

Agent (Role): {Goals, Norms, Agreements, Plans, Beliefs}

Definitions

Goals are the states that an agent wants to bring about in the environment.

Norms are statements that must hold for all agents in the society at any given time.

Agreements are protocols that govern the behaviour of agents working together to achieve goals.

Plans are collections of tasks that an individual intends to carry out to achieve its **Goals**, such as **Enquiries**, **Actions**, **Decisions** and including sub-plans.

Beliefs are states which an agent holds to be true, particularly with respect to the environment and other agents.

When instantiated in healthcare this generic agent paradigm provides the following more specific schema:

Agent (Role): {

Clinical Problems to Solve *(the diagnosis or treatment of a health problem, etc.)*,

Clinical Guidelines *(published references or strategies for GP, etc.)***,**

Clinical Interaction Protocols *(the standard ways that sequences of clinical tasks unfold over time, which are carried out by cooperating clinicians)***,**

Clinical Plans *(intended clinical tasks in a logical order, consultation and intervention, etc.)*,

Clinical Opinions *(the interpretation of clinical situations, such as patient conditions)* **}**

Firstly, an agent may want to collect certain data, solve a particular diagnosis problem, or achieve other kinds of clinical target. Such are the **Goal** in focus appropriate to the agent's **Roles**.

Secondly, **Roles** are associated with addressing particular types of problems, rather than specific tasks. For this reason, a versatile agent behavioural model will be employed, appropriate to fit in any actual **Norms** which are continuously improving in qualty and incrasing in quantity.

Thirdly, agents with different capabilities may commit to solve complex clinical problems together, under mutually agreed protocols or **Agreements**. The relevant set of clinical guidelines, or the **Norms** respected in the society, can then be distributed via the enactment of **Agreements**. The agents representing a multidisciplinary team, will base their behaviour upon the guidelines distributed to them, and in a coordinated manner set out by the protocols.

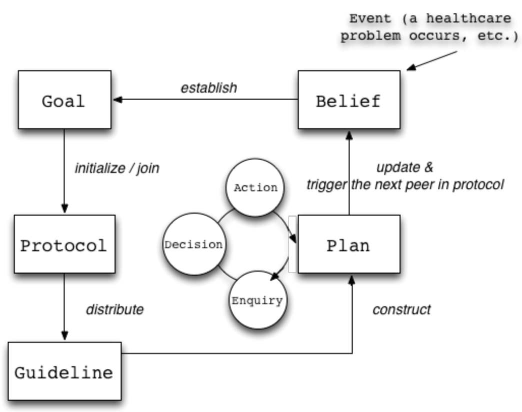

Fig. 1. An agent-oriented model for joint care delivery and decision-making

In the **GNAPB** paradigm, agents take a set of *agent society acts* and follow a process as illustrated in Figure 1 to collaboratively deliver care and make decisions, cyclically: The driving force of the cycle is always that an agent observes some environmental change which is inconsistent with its current goals and requires some actions: **Goals** provide a bridge from **Beliefs** to **Plans**. The occurrence of an **Event** (reported symptoms, etc.) triggers an agent to *establish* a new **Goal** (diagnose the problem, etc.) or *revise* an existing one. To accomplish this, the agent needs to *initialise* an interaction protocol (as it cannot achieve it alone). Later other participants will be invited to *join* this protocol, under a binding **Agreement** for the group to solve the problem together. The enactment of the protocol will include a process which *distributes* relevant guidelines for participant agents to execute. Under the Agreement, each agent is required to *respect* the relevant parts of the guidelines as **Norms** that govern their behaviour. To bring that kind of agreed duties and responsibilities into effect, agents follow a process to *construct* their individual **Plans** which themselves include a cycle of **Enquiry**, **Decision**, and **Action**.

Upon completion of a Plan, agents *update* their **Beliefs** (symptoms have been diagnosed, etc.), and may trigger the next agents to continue within the current protocol or trigger an entirely new cycle (start a further treatment process, etc.). Throughout these processes society members accumulate and share experience and evidence, and reach consensus in order to *classify* Norms with respect to quality or even *propose* new **Norms** or *obsolete* invalid ones. A guideline ranking and evaluation mechanism as such may help self-optimisation. We may call the above a distributed version of the Domino model.

3 Agent-oriented design patterns for joint clinical decision support

Applying the **GNAPB** paradigm, as shown in Figure 1, to a multidisciplinary team suggests the sharing, distribution, execution, and coordination of clinical guidelines in an agent society, in that order: 1) the decomposition and distribution of guidelines to a group of peers which later commit to achieving their shared **Goals**; 2) the interpretation of distributed parts of guidelines by agents as **Norms**; 3) the execution of role-specific parts of the guidelines including individual decision making processes by constructing **Plans**; 4) the coordination of agents and the aggregation of results under **Agreements** and 5) the completion of guidelines adding to peers' internal **Beliefs**. Several agent-oriented design patterns support such a paradigm for later implementation.

Table 1. Agent-oriented design patterns for joint decision support in society

	Resolving Component	**Pattern**	**Society Notion**	**Problem Solved**
Agent-Protocol Subscription Pattern	Role	Subscription	Society Agreement	Distribution of Guideline
Agent-Guideline Interpretation Pattern	Behavioural Rule & Production Rule	Interpretation	Society Norm	Interpretation of Guideline Portion
Agent-Agent Coordination Pattern	Speech Act	Messaging	Society Performative Act	Coordinated Enactment of Protocol & Adaptive Interaction
Agent-Component Binding Pattern	Behavioural Rule	Runtime Invocation	Local Agreement	Adaptive Computation
Agent-Decision Making Pattern	Production Rule	Argumentation	Local Plan	Reasoning

Agent-Protocol Subscription Pattern (Society Agreements)

An interaction protocol specifies the collective behaviour that a group of peers work together, each playing a corresponding **Role**, to accomplish a shared goal, e.g. solving a comprehensive health problem. Roles are characteristics that distinguish one agent from another. At runtime, when an agent subscribes to an interaction protocol and assumes one of the roles required in the protocol, that agent gains the capabilities and commits to the responsibilities associated with that role. Agents can dynamically *subscribe* to new interaction protocols and assume the required roles. As members of different teams and in different settings, they have different problem solving capabilities. The design of agents as flexible protocol subscribers enables grouping and re-grouping of clinical expertise towards emerging multidisciplinary healthcare and provides a diverse range of services.

Agent-Guideline Interpretation Pattern (Society Norms)

When an agent participates in an interaction protocol, a portion of a clinical guideline(s) related to its role is assigned to this agent for execution – and at a different time, a different guideline may come into effect. In order to enable agents to understand guidelines and behave upon them without predefined constraints or limitation, **Behavioural Rules** are employed as a uniform container to which knowledge can be filled in and from which the required behaviour translated. Therefore, guidelines are transformed into standard rule formats, distributed to agents, and agents always follow the same pattern to *interpret* their required behaviour, dynamically. These rules have the same structural scheme but runtime instances have unique contents encoded and matched with the exact role behaviour. The design of agents as versatile problem-solver and coupled with a uniform rule scheme that advises on runtime behaviour ensures the emerging clinical guidelines will be taken into effect immediately, with minimum system re-development overhead.

Agent-Agent Coordination Pattern & Agent-Component Binding Pattern

Agents exchange data and knowledge by **Messaging**, using a common set of performative acts, data dictionaries, and message encoding tags. This allows agents to "speak the same language" in the agent society, though in their independent domains they may use private datasets, components, and applications for local computation and decision-making. Lower level data interoperability between partner sites will be of less concern, since agents exchange knowledge and *coordinate* actions via message passing, and interoperate in such higher levels as achieving shared goals. Issues about data sharing between Primary Care and Secondary Care might be alleviated in this way.

Agents can also make private arrangements with their local components, web services, or event agents wherever their computational capabilities best match the actual needs. Such are the contracts that *bind* agents with their local components. An agent may swap an old component with a new one or choose from alternatives opportunistically, based on their capabilities, performance, cost-efficiency, and other attributes. This might be achieved via human experts re-configuring the binding contracts in **Behavioural Rules** and agents interpreting these rules and binding with the desired components dynamically at runtime.

Agent-Decision Making Pattern

Situated in a society, agents may receive notifications of emergency or routine requests and respond with suggestion for resolving these requests. The key function is a *decision-making* or planning process. In the Domino model, **Plans may** include **Enquiries** to collect information prior to reaching **Decisions**, and committing to **Actions** once those Decisions are made. In addition to such procedural knowledge as structured in Behavioural Rules, declarative statements are captured in **Production Rules**. They specify the logical relationships embedded in guidelines and may fire in a forward-chaining control structure to deduce extra knowledge. Such knowledge may supplement the enactment of Behavioural Rules, where enquired information needs to be processed, conditions checked, and actions enacted. The schemes and examples of both types of rules will be illustrated in the next section.

4 A case study: breast cancer referral

A joint decision support model in Figure 1 suggests a cycle starting when Events *occur*, Beliefs are *updated*, Goals *establishd*, Protocols *initialised*, Guidelines *distributed*, Plans *constructed*, and with multiple occurrences of the above cycle ending up with a patient being successfully diagnosed, treated, etc. Also an agent architecture and its supporting design patterns are summarised in Table 1. We illustrate below how the decision support model and the agent architecture may be integrated into a unified framework using the breast cancer referral example. The guidelines for this example are presented in Figure 2.

Urgent referral

Refer urgently patients:

- of any age with a discrete, hard lump with fixation, with or without skin tethering C
- who are female, aged 30 years and older with a discrete lump that persists after their next period, or presents after menopause C
- who are female, aged younger than 30 years:
  - with a lump that enlarges C
  - with a lump that is fixed and hard C
  - in whom there are other reasons for concern such as family history[5] D
- of any age, with previous breast cancer, who present with a further lump or suspicious symptoms C
- with unilateral eczematous skin or nipple change that does not respond to topical treatment C
- with nipple distortion of recent onset C
- with spontaneous unilateral bloody nipple discharge C
- who are male, aged 50 years and older with a unilateral, firm subareolar mass with or without nipple distortion or associated skin changes. C

Non-urgent referral

Consider non-urgent referral in:

- women aged younger than 30 years with a lump C
- patients with breast pain and no palpable abnormality, when initial treatment fails and/or with unexplained persistent symptoms. (Use of mammography in these patients is not recommended.) B(DS)

Fig. 2. Clinical guideline for breast cancer referral written in natural language [17]

1) **Preparation**. A set of guidelines from a nationally or internationally recognised repository of evidence-based recommendations (e.g. *NICE pathways* or *BMJ Clinical Evidence*) is formalising using a standardised representation. The distribution of formalised guidelines among the members of a multidisciplinary team reflects the clinical roles in the clinical centre. In the present case study, urgent and non-urgent referral criteria (from NICE CG27 shown in Figure 2) are structured and distributed to a "GP Agent".

2) **Agent-Protocol Subscription**. A breast cancer referral protocol initialises when an **Event** occurs such as a "Patient Agent" reporting an abnormal lump (shown diagrammatically in the top layer of Figure 3 and textually in Figure 4a). The protocol enacts in a way that three **Roles** (a Patient Agent, a GP Agent, and a Specialist Agent) group together, join the protocol, and commit to solve the problem together. All agents except the initialising one subscribe to the protocol when they receive invitation messages from their coordinators.

3) **Agent-Guideline Interpretation**. The GP Agent uses a **Behavioural Rule** to interpret and execute its assigned part of the guideline (shown diagrammatically in the middle layer of Figure 3 and textually in Figure 4b). On receipt of a patient report event, it updates its knowledge about this patient using a "Patient" type of component, evaluates its referral criteria using a "Referral" type of component, and finally either sends a discharge message to the Patient Agent or an urgent referral message to a "Specialist Agent", or holds it in abeyance.

4) **Agent-Agent Coordination**. An agent enacts its role behaviour by carrying out **Plans** and making decisions, and in our approach these are uniformly structured as Behavioural Rules. Their structure defines a very simple and uniform coordination pattern: an agent joins coordination when it receives an Event message(s) and invites other agents to join when it produces an Action message(s). The recurrent **Message Passing** pattern between pairs of agents determines the way that multiple roles are coordinated in interaction protocols, and decisions are made jointly, where the outcome of one decision becomes the input of another. In our example, when a Patient Agent *reports* a problem to a GP Agent (*receiving a report message*), the latter may *refer* the patient to a Specialist Agent in appropriate situations (*sending a referral request message*). In this way guidelines previously distributed across different locations are jointly enacted with regards to the coordinative multi-agent behaviour.

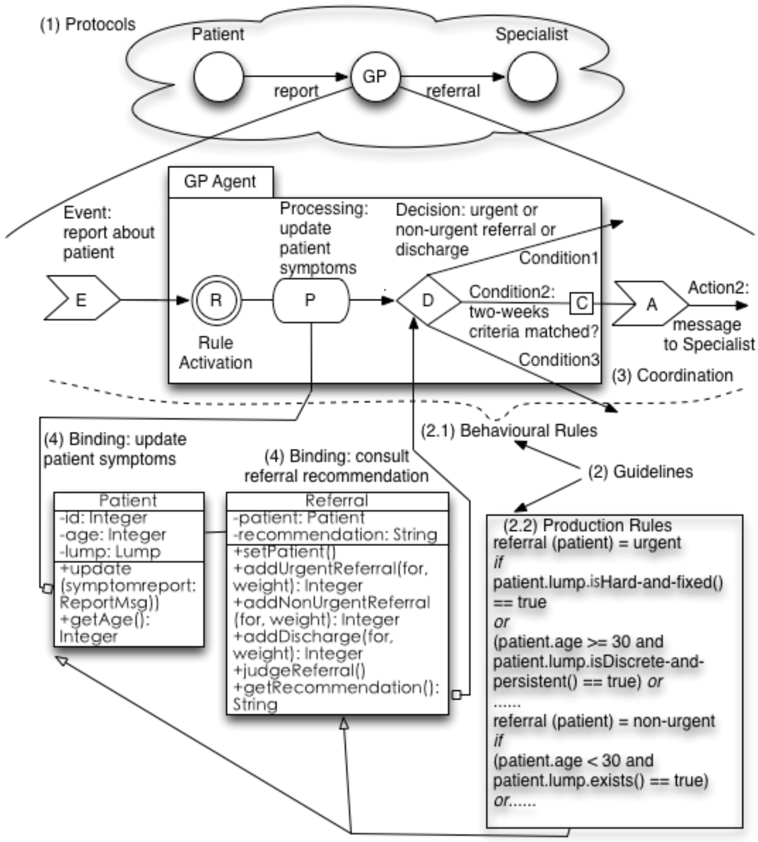

Fig. 3. Agents use Behavioural Rules and Production Rules to execute guidelines, under the agreements of Interaction Protocols and over the support of local Components

5) **Agent-Component Binding**. A Patient and a Referral component are bound to the GP Agent as contracts configured in the **Behavioural Rule** (shown diagrammatically in the bottom layer of Figure 3 and textually in Figure 4b). These have computational capabilities and invoked by the GP Agent for event processing, decision-making, and outcome production. A Patient instance is updated with the reported symptoms and conditions (result of event processing). A Referral instance, in an association relationship with the Patient instance, will have its attribute values deduced with the assistance of Production Rules. This will be useful to judge the patient referral criteria (condition checking). The interplay of the two types of rules and the components is essential to decision-making.

```
- <CancerReferral>
  - <protocol-initialiser>
      <name>Patient</name>
      <event>report a lump</event>
    </protocol-initialiser>
  - <roles>
    - <role>
        <name>GP</name>
        <port>RequestPatientInfo</port>
        <port>SendReferralDetails</port>
      </role>
    - <role>
        <name>Specialist</name>
        <port>ReceiveReferralDetails</port>
        <port>NA</port>
      - <property>
          <specialisation>breast cancer</specialisation>
        </property>
      </role>
    </roles>
  - <connector>
      <name>BreastCancerReferral</name>
      <role>referral</role>
      <role>referee</role>
    </connector>
  - <contracts>
    - <binding>
        GP.SendReferralDetails to BreastCancerReferral.referral
      </binding>
    - <binding>
        Specialist.ReceiveReferralDetails to BreastCancerReferral.re
      </binding>
    </contracts>
  </CancerReferral>
```

```
- <BehaviouralRule>
    <role>GP</role>
    <protocol>CancerReferral</protocol>
  - <component>
      <instance>patient</instance>
      <type>Patient</type>
    </component>
  - <component>
      <instance>referral</instance>
      <type>Referral</type>
    </component>
  - <event>
      <sender>Patient</sender>
    - <message>
        <type>patient.symptomreport();</type>
        <content>patient.symptomreportinXML();</content>
      </message>
    </event>
  - <processings>
      <processing>patient.update(thisMessage);</processing>
      <processing>referral.setPatient(thisPatient);</processing>
      <processing>referral.judgeReferral();</processing>
    </processings>
  - <decision-tree>
    - <branch>
      - <condition>
          referral.getRecommendation().equals("urgent");
        </condition>
      - <action>
          <receiver>Specialist</receiver>
          <message-content>referral.getReferralDetails();</message-conte
        </action>
      </branch>
    - <branch>
      - <condition>
          referral.getRecommendation().equals("non-urgent");
        </condition>
        <action> referral.hold(); </action>
      </branch>
    - <branch>
      - <condition>
          referral.getRecommendation().equals("discharge");
        </condition>
      - <action>
          <receiver>Patient</receiver>
          <message-content>referral.discharge();</message-content>
        </action>
      </branch>
    </decision-tree>
  </BehaviouralRule>
```

Fig. 4a A partial specification of a protocol “CancerReferral” & 4b Specification of a Behavioural Rule for the GP Agent

6) **Agent-Decision Making**. The original source guidelines, shown in Figure 2, guide both the establishment of procedural interrelationship between agents such as patients consult GP or referral to specialists, but also the declarative relationships between key concepts and their logical connection. The former type of knowledge is modelled in Behavioural Rules and the latter in **Production Rules**, which support the deduction of extra facts based on existing ones. In our case study, the GP Agent uses the Behavioural Rules, shown in Figure 4b, to carry out the main body of its plans, and the Production Rules to support the evaluation of Referral upon Patient. That leads to the final judgment and of most importance to the decision making process as part of the plans.

Interaction Protocol Specification

Specification of Interaction Protocols lays down the runtime *agent coordination architecture* required to achieve agents’ goals and the distribution of guidelines to specific roles. Some Architecture Description Languages (ADLs) such as ACME [18], C2, and UniCon are well recognised for describing Component-Connector based architectures and may be adapted for specifying role-based agent interactions. We show such an example in Figure 4a, an XML-based interaction protocol for the breast cancer referral case.

Behavioural Rule Formalism

Behavioural Rules define a common behavioural pattern that needs to be recognised by all agents when they carry out their plans. The rules are formulated as below.

Agent (Behavioural Rule):

{Event, Processing, Decision (Condition, Action)$_n$, Belief}

The agents' role-playing behaviour and their interaction is via runtime interpretation of Behavioural Rules. These rules guide agents in a manner similar to the enactment of a Plan in PROforma: *process* **Events** (loosely matching PROforma Enquiry but with explicit incoming messages), *make* **Decisions** (branches loosely matching PROforma candidates and argumentation), and *carry on* further **Actions** (loosely matching PROforma Actions but with explicit outgoing messages) when given **Conditions** are satisfied, and finally *update* its own **Beliefs** towards the environment. Plans are distributed across different sites and then coordinated or "choreographed". An XML-based specification of a Behavioural Rule is shown in Figure 4b. It says when a patient reports symptoms (Event), the GP Agent updates its knowledge about this patient (Processing), considers the criteria for urgent or non-urgent referral (Decision), sends the patient to a specialist or keeps the patient on a wait list or simply discharges her (three Actions matching their corresponding Conditions), and finally updates himself for this occasion (Belief).

Production Rule: Fact Deduction Facilities

Behavioural Rules specify the procedural plans for agents including message processing, decision-making and collaboration, as well as the binding of components that can assist in such tasks. Declarative logic statements concerned with reasoning about candidate decision options and other matters are captured in **Production Rules**. They deduce additional component attributes based on the known set, establish connections between component instances, and support reasoning and argumentation. For example, a fragment of descriptions in Figure 2 is represented as a Production

Rule in the bottom right of Figure 3. It establishes the relationship between two component instances: if the attribute "age" of a Patient has a value over 30, and the attribute "lump" has the value "discrete-and-persistent", then a method of Referral will fire: *addUrgentReferral(for, weight)*. This will affect the aggregation of arguments, the assessment of the overall preferences over decision options (urgent, non-urgent, or discharge), and the final ordering of decision options for recommendation ("*judgeReferral*" and "*getRecommendation*", as shown in Figure 4b).

Open and Adaptive Agent Architecture

When we need to support a new kind of multidisciplinary care pathway and decision making, we may follow these design patterns: define a new protocol, make it open to clinical experts and services, let them join it as a team, pass around the relevant evidence and guidelines, and finally they coordinate and make decisions together. Our knowledge-driven architecture is open to new participants to join and adaptive to accommodate and disseminate new guidelines. This is largely because of the use of reconfigurable rules that bind agents with their cooperating partners, behavioural pattern, and assistant components later at runtime rather than design time. At the time of playing roles and making decisions these agents form a team and work to comply with the current guidelines.

5 Discussion, Conclusions, and Future Work

The growing specialisation and complex interrelationships in medicine today imply more collaborative rather than independent decisions, a process in which one decision depends upon the result of another and cannot be reached in isolation. In such collaborations, any individual specialist cannot see all the data but may share the responsibilities, which represents a risk to patient safety. We believe it is important to make the knowledge base of clinical guidelines and the interaction protocols among care specialists or services explicit, in dedicated repositories. These can be used and reused for later discovery, customisation and adaptation. The underlying software systems also need to be capable of disseminating new knowledge and using that knowledge for coordinated decision-making in a distributed environment, as flexibly and adaptively as possible. Multi-Agent Systems are a good candidate, especially when coupled with reconfigurable rules for knowledge encapsulation and runtime interpretation.

Some previous work on applying MAS to healthcare has been reviewed extensively in [14] and more recently in [21], which selected 15 most recent and important applications. Among them three are concerned with clinical decision making, where different types of agents have been proposed, in managing datasets at local clinical sites, collecting relevant data or evidence in a distributed network, interacting with end users, and so on. Major decision making solutions include case-based reasoning (Singh), pattern recognition and data training (HealthAgents), and guideline application (HeCaSe2). In these approaches and others, a focus has been put upon modelling the organisational structures and specific workflows among various types of agents, so that an agent architecture may support the real environment. In HeCaSe2 where guidelines are central and an approach most close to our own, although a Guideline Agent is proposed, authors seem to assume it can automatically accomplish its job. The approach is rather aimed at addressing agent and service interaction processes, including the mentioned Guideline Agent, Medical Record Agent, Service Agent, and User Agent, in a networked infrastructure, especially, the Catalonia medical environment. It seems thus far, no work has been dedicated to the mechanism of representation and distribution of guidelines among a multidisciplinary team, the understanding and interpretation of guidelines by an agent, and the maintenance and dissemination of new guidelines in a MAS architecture.

The CDSS development community has been called for standard service architectures and interfaces so that any EHR system can subscribe to for the needed capabilities with minimum implementation effort [10]. We believe the proposed open and adaptive agent-oriented decision support architecture may offer a reference model for researchers and developers for adaptation and fitting their own data and knowledge. It will be especially useful whereas care collaboration patterns are emerging and guidelines are improving rapidly, in a multidisciplinary care environement. The approach contributes to cost-effective IT development and maintenance and will eventually provide a shift of effort from system re-design and re-coding to knowledge reconfiguration and re-dissemination.

We used the NICE CG27 clinical guideline about early referral of suspected breast cancer in this study. Later investigation in this field will include familial breast cancer (CG14 & CG41), early and locally advanced breast cancer (CG80), and advanced breast cancer (CG81). Eventually, we will develop a unified multidisciplinary framework that covers the whole “cancer journey” for coordinated decision-making.

6 References

[1] Sutton, D.R., Fox, J., The Syntax and Semantics of the PRO*forma* guideline modelling language, Journal of the American Medical Informatics Association 10(5): 433-443, 2003.

[2] Fox, J., Patkar, V., Thomson, R., Decision Support for Healthcare: the PROforma evidence base, Inform Prim Care 14(1): 49-54, 2006.

[3] Bury, J., Humber, M., Fox, J., Integrating Decision Support with Electronic Referrals, Proceedings of the Tenth World Congress on Health and Medical Informatics (Medinfo2001), London, UK, September 2-5, 2001.

[4] Walton, R.T. et al., Evaluation of computer support for prescribing (CAPSULE) using simulated cases, BMJ 315(7111): 791-795, 1997.

[5] Steele, R., Fox, J., Enhancing conventional web content with intelligent knowledge processing, Proceedings of the 9th Conference on Artificial Intelligence in Medicine in Europe (AIME'03), Lecture Notes in Artificial Intelligence (LNAI 2780), pp. 142-152, Springer Verlag, 2003.

[6] Fox, J., Glasspool, D., Modgil, S., A Canonical Agent Model for Healthcare Applications, IEEE Intelligent Systems 21(6): 21-28, 2006.

[7] Patkar, V., Hurt, C., Steele, R., Love, S., Purushotham, A., Williams, M., Thomson, R., Fox, J., Evidence-based guidelines and decision support services: a discussion and evaluation in triple assessment of suspected breast cancer, British Journal of Cancer 95(11): 1490–1496, 2006.

[8] Fox, J., Glasspool, D., Patkar, V., Austin, M., Black, L., South, M., Robertson, D., Vincent, C., Delivering clinical decision support services: there is nothing as practical as a good theory, Journal of Biomedical Informatics 43(5): 831-843, 2010.

[9] Fox, J., Beveridge, M. and Glasspool, D., Understanding intelligent agents: analysis and synthesis, AI Communications 16(3): 139-152, 2003.

[10] Sittig, D.F., Wright, A., Osheroff, J.A., Middletone, B., Teich, J.M., Ash, J.S., Campbell, E., Bates, D.W., Grand challenges in clinical decision support, Journal of Biomedical Informatics 41(2): 387–392, 2008.

[11] Littlejohns, P., Garner, S., Doyle, N., Macbeth, F., Barnett, D., Longson, C., 10 years of NICE: still growing and still controversial, The Lancet Oncology 10(4): 417-424, 2009.

[12] Drummond, M. and Sorenson, C., Nasty or Nice? A Perspective on the Use of Health Technology Assessment in the United Kingdom, Value in Health 12: S8–S13, 2009.

[13] Field, M.J., Lohr, K.N. (Eds), Clinical Practice Guidelines: Directions for a New Program, Institute of Medicine, Washington, DC: National Academy Press, 1990.

[14] Moreno, A., Nealon, J.L. (Eds.), Application of Software Agent Technology in the Health Care Domain, Birkhäuser, 2003.

[15] Xiao, L., Greer, D., Adaptive Agent Model: Software Adaptivity using an Agent-oriented Model Driven Architecture, Information & Software Technology 51(1): 109-137, Elsevier, 2009.

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

[17] National Institute for Health and Clinical Excellence (NICE), Referral guidelines for suspected cancer (CG27), p. 12, issued in 2005 and updated in 2011.

[18] Garlan, D., Monroe, R., Wile, D., Acme: an architecture description interchange language, Proceeding of the CASCON First Decade High Impact Papers (CASCON '10), pp. 159-173, 2010.

[19] Fox J. and Das S., Safe and Sound: Artificial Intelligence in Hazardous Applications, Jointly published by the AAAI, Menlo Park, CA, and MIT Press, Cambridge Mass., 2000.

[20] Xiao, L., Cousins, G., Fahey, T., Dimitrov, B., Hederman, L., Developing a rule-driven clinical decision support system with an extensive and adaptive architecture, Proceedings of the 14th International Conference on E-health Networking, Application & Services (HealthCom'2012), pp.250-254, 2012.

[21] Isern, D, Sánchez, D., Moreno. A., Agents applied in health care: A review, International Journal of Medical Informatics 79(3): 145-166, 2010.