A knowledge-rich distributed decision support framework: a case study for brain tumour diagnosis

Dupplaw, D; Croitoru, M; Dasmahapatra, S; Gibb, A; Gonzalez-Velez, H; Lurgi, M; Hu, B; Lewis, P; Peet, Andrew

Document Version
Publisher's PDF, also known as Version of record

A knowledge-rich distributed decision support framework: a case study for brain tumour diagnosis

DAVID DUPPLAW[1], MADALINA CROITORU[1,†],
SRINANDAN DASMAHAPATRA[1], ALEX GIBB[2],
HORACIO GONZÁLEZ-VÉLEZ[3], MIGUEL LURGI[4], BO HU[1,‡],
PAUL LEWIS[1] and ANDREW PEET[5]

[1]*Intelligence, Agents, Multimedia Group, School of Electronics and Computer Science, University of Southampton, Highfield, Southampton SO17 1BJ, UK;*
e-mail: dpd@ecs.soton.ac.uk, croitoru@lirmm.fr, sd@ecs.soton.ac.uk, bo01.hu@sap.com, phl@ecs.soton.ac.uk;
[2]*University of Birmingham, Birmingham B15 2TT, UK;*
e-mail: a.j.gibb@bham.ac.uk;
[3]*School of Computing and IDEAS Research Institute, Robert Gordon University, St Andrew Street, Aberdeen AB25 1HG, UK;*
e-mail: h.gonzalez-velez@rgu.ac.uk;
[4]*MicroArt, SL. Parc Científic de Barcelona C/Baldiri Reixac, 4-6 – 08028 Barcelona, Spain;*
e-mail: mlurgi@microart.cat;
[5]*University of Birmingham and Birmingham Children's Hospital, Birmingham B4 6NH, UK;*
e-mail: acpeet@doctors.org.uk

Abstract

The HealthAgents project aims to provide a decision support system for brain tumour diagnosis using a collaborative network of distributed agents. The goal is that through the aggregation of the small data sets available at individual hospitals, much better decision support classifiers can be created and made available to the hospitals taking part. In this paper, we describe the technicalities of the HealthAgents framework, in particular how the interoperability of the various agents is managed using semantic web technologies. On the broad scale the architecture is based around distributed *data-mart* agents that provide ontological access to hospitals' underlying data that has been anonymized and processed from proprietary formats into a canonical format. Classifier producers have agents that gather the global data from participating hospitals such that classifiers can be created and deployed as agents. The design on a microscale has each agent built upon a generic-layered framework that provides the common agent program code, allowing rapid development of agents for the system. We believe that our framework provides a well-engineered, agent-based approach to data sharing in a medical context. It can provide a better basis on which to investigate the effectiveness of new classification techniques for brain tumour diagnosis.

1 Introduction

Brain tumours present a particular challenge to healthcare systems and medical research. Although they occur in people of all ages, they are a particularly common form of cancer in children and young adults. The young age of the patients, the comparatively high mortality rate and the burden of associated neuro-disability ensures that these tumours assume an importance that exceeds their overall incidence. Improved methods for diagnosis, treatment and understanding of brain tumours are an important goal of clinical and biomedical research.

---

† Madalina Croitoru is now at LIRMM, 161 rue ADA, F34392 Montpellier Cedex 5, Montpellier, France.
‡ Bo Hu is now at SAP Research CEC Belfast, BT37 0QB, UK.

Magnetic resonance imaging (MRI) is the standard investigation performed during diagnosis and allows the presence of the lesion to be confirmed as well as its location, size and structure determined. However, there are many questions that conventional MRI cannot reliably answer. In particular, the tumour type and grade is often poorly determined, and the diagnosis is subsequently obtained by removing a sample of the tumour and performing histopathology. However, surgery has an associated risk for the patient, and improving the diagnostic and prognostic information available from non-invasive methods is an important objective.

Imaging research is currently focussing on techniques that yield information on the biological properties of tumours. One of the most promising of these is magnetic resonance spectroscopy (MRS), a method that measures levels of various chemicals in tissues. Previous studies have shown MRS profiles to be a powerful characteristic of brain tumours, aiding non-invasive diagnosis and providing novel biomarkers of prognosis. However, the method has not met with universal acceptance among clinicians. One of the barriers to the widespread clinical use of MRS lies in the difficulties faced when interpreting MRS data.

Radiologists are experts in interpreting images, which is primarily done through visual inspection. MRS data is difficult to interpret by visual inspection and the most accurate method of analysis lies in pattern recognition of the spectra. Making such techniques available in a clinical environment is challenging. Sophisticated software is required to present the results in a manner acceptable to clinicians working in a busy clinical environment. Furthermore, large numbers of cases in each diagnostic category are required to build robust pattern recognition classifiers and this necessitates the collection of data from multiple centres.

In HealthAgents (Arús *et al.*, 2006; González-Vélez *et al.*, 2009), we have created an agent-based architecture that allows the sharing of anonymized MRS patient data, providing a basis on which to create classifiers that are able to produce better results than those that are limited to a hospital's local data. By providing better diagnosis of tumours based on non-invasive techniques, the hope is that prognosis and treatment can be better managed.

The nine partners of the HealthAgents project are MicroArt, Universitat de València, Universitat Autònoma de Barcelona (UAB), Institute for the Applications of Advanced Information and Communication Technologies (ITACA), Pharma Quality Europe, Katholieke Universiteit Leuven, University of Birmingham, University of Edinburgh and University of Southampton. Being a multinational, multicentre project, the HealthAgents architecture is required to be platform independent and scalable by design. It has therefore been devised so that any new functionality can be easily integrated into the general context by implementation of new agents.

There are many other systems designed for distributed storage and sharing of data, but most would not be suitable to be used in a medical context. Indeed, it is absolutely vital that patient data is kept secure at all times and the HealthAgents system has been designed such that all of the data sets involved can be anonymized; that is, data that contain no direct link to the patient's identity. The architecture also provides various security enhancements to avoid the unauthorized access of data.

In the following sections, we describe our contribution to this area of research. First, we motivate our research approach in Section 2, complemented with the user requirements in Section 3, and the layered, agent-based architecture—that is, the basis of the framework—which is delineated in Section 4. Furthermore, we properly describe the HealthAgents framework in Section 5 and finish by providing some evaluative comments and concluding remarks in Sections 6 and 7, respectively. There is a section defining our nomenclature at the end of the paper.

2 Motivation

Widely conceived as parameterizable concurrent constructs, software agents possess four intrinsic characteristics (Brugali & Sycara, 2000):

- An intrinsic knowledge-based state that can evolve dynamically;
- Changing reasoning capabilities that determine their internal behaviour through constraints or goals;

- A communication status that allows them to interact with other agents or human entities; and,
- A unique identity with roaming and service-advertising capacities.

In this vein, agent-based application frameworks follow a defined lifespan, where the architectural choices are first formulated with platform and domain-knowledge independence, and the elemental ones implementing concurrency and logic mechanisms are then devised to determine implementation particulars (Brugali *et al.*, 1997).

On the one hand, agent frameworks for medical purposes have been extensively documented in the literature (Merelli *et al.*, 2007; Annicchiarico *et al.*, 2008), but its use has been traditionally confined to patient control and information (Huang *et al.*, 1995).

Lanzola *et al.* (1999) use a prototype of a hospital information system, mainly oriented towards the management of patients with leukaemia, using an agent-based model. Although the architecture falls short of showing multicentre interaction or classification capabilities, its single framework approach is interesting *per se*.

The Artemis Project (Boniface *et al.*, 2005) has previously researched the transfer of critical patient data (mainly clinical data) in a medical context. It used a peer-to-peer-like architecture and used the web service security protocols for the transfer of patient records. Similar to HealthAgents, the project has investigated ontological means of data representation.

The MIAKT (medical images and advanced knowledge technologies) Project designed and used a knowledge-based distributed decision support system to support the diagnosis and management of breast cancer cases using ontologies and Web services. The framework developed during the project was capable of being used in different applications (Shadbolt *et al.*, 2004).

The intensive care agent platform concentrates on the intensive care side on a single-node (intra-hospital) basis (Turck *et al.*, 2007), providing little evidence of dynamic distributed multicentre interaction.

On the other hand, from a neuro-oncological point of view, despite the advances in the application of MRI and MRS techniques for diagnosis (Rees, 2003), there exist a few research endeavours that report successful automated diagnosis of brain malignancies.

Having created a single-node decision support system using MRS classification supported on histopathological diagnosis, the Interpret Project (Tate *et al.*, 2006) has centred its efforts on the clinical aspects of the diagnosis and reported successful results for a handful of common malignancies.

Georgiadis *et al.* (2008) report the successful discrimination between metastatic and primary brain tumours using MRI. They use distinct machine learning classifiers on a static patient record repository, but provide no evidence on the online classification via multiple repositories.

The eTumour project implements stringent quality control mechanisms for patient records (Wright *et al.*, 2008) but has reported few advances from a generic software architecture standpoint.

That is to say, scant research has been devoted to deal with cancer and neuro-degenerative disease diagnosis using a generic distributed agent-based framework for the secure exchange and classification of MRS patient data. HealthAgents expands previous approaches to computer-aided brain tumour diagnosis with a distributed multicentre agent architecture, an *in vivo* classification method with negotiation, an additional number of cases located in different centres across Europe and a Web-based user interface. By working with hospitals and ensuring the security of the MRS data passing through the system, HealthAgents is being used in clinical environments.

3 User requirements

The diagnosis and management of brain cancer currently depends on the histological examination of a brain biopsy, and the optimum treatment varies with the tumour type and with its grade; that is, its degree of malignancy. Both histology and grade have to be assessed using established but partly subjective criteria by a skilled neuropathologist.

Stereotactic brain biopsy has significant risks, and, moreover, histology is not accurate enough in some cases, although it is still considered the gold standard methodology. Current radiological

methods do not adequately distinguish between the large number of recognized types of brain tumours (nearly one hundred) and thus it is necessary to perform a histopathological diagnosis on a biopsy.

Normally, most patients undergo surgery for cytoreduction and decompression and the resulting biopsy is available for analysis. However, there are a few cases in which surgical resection or biopsy would not be considered, such as in very elderly or infirm patients with obvious malignant lesions, or patients with very slow-growing tumours in vital parts of the brain. In addition, there are certain pathologies (lymphomas and brain abscesses) in which pre-operative diagnostic certainty would avoid the open surgery step. An additional non-invasive method for accurately diagnosing and grading brain tumours would be a major advance in those cases.

There is a need to improve brain tumour classification, and to provide non-invasive methods for brain tumour diagnosis and prognosis, to aid patient management and treatment. Consequently, HealthAgents uses two sources of biological information that allow the non-invasive typing and grading of brain tumours to be improved:

- *In vivo* techniques: MRS is a non-invasive technique that provides biochemical information of tissue *in vivo*. MRS, coupled with conventional MRI, provides metabolite profiles of a single voxel or volume (SV) of tumour tissues. It also produces a molecular image of particular tumour metabolites in 10 min using multivoxel (MV) techniques.
- *Ex vivo* and *in vitro* techniques: high-resolution magic angle spinning (HR-MAS) is a modality of MRS that is applied to biopsies *in vitro* in order to improve characterization. DNA microarray analysis can be used to determine tumour phenotype from gene expression profiles.

The HealthAgents network is a set of sites interconnected to share resources. The main sites are located in Birmingham (United Kingdom), Barcelona (Spain) and Valencia (Spain). Birmingham is aggregating data from 50 different contributing centres, whereas Barcelona and Valencia are providing data from six and four hospitals, respectively. The Barcelona node is collecting brain tumour cases mainly from the validated database from the Interpret Project. Globally, the different databases will comprise some ~ 600 cases upon project completion. All these cases will be quality checked and verified.

Two main decision support functionalities are offered through the HealthAgents system. The first is provided by the groups carrying out pattern recognition (ITACA and The University of Leuven) who are producing tumour classifiers based on case data that is being gathered throughout the project by the medical centres. The second is the evidence-based search service (ebSS) developed by The University of Edinburgh, which is a data mining tool that complements the use of the classifiers with a more general approach to seek out and gather information, especially textual evidence from the online literature and patient information.

In spite of the advantages inherent to a centralized system, a distributed system implementation is proposed, which better supports the following system properties:

- Confidentiality: The clinical record information of each centre is private and its confidentiality must be guaranteed. The centre should maintain the control of its information diffusion. In some cases, they may decide to share this information to every HealthAgents site (all or parts of it), but in other cases they may prefer to keep it private. In a centralized implementation, the entire clinical record of the contributing cases should be made accessible.
- Robustness: Usually, in a distributed system, when a node is temporarily out of service, the global system would continue working and the other remaining nodes would be able to benefit from its use. Often, in a centralized implementation, if the server is down, all local sites must wait for the solution of the problem, thus being unable to benefit from the system.
- Speed: As each contributing centre has its own copy of a classifier, the answer in the local predictions of dDSS is sped up for each new case. This distributed architecture enables two main kinds of classifiers: those obtained from every public data case, and other bespoke specific classifiers. The HealthAgents system has been built as a distributed network of nodes, in which each contributing centre manages its own local site, with the autonomous capacity to classify new tumour cases according to both global classifiers and specific classifiers.

4 The architecture

Figure 1 shows the logical architecture of the system's data-marts (in HealthAgents parlance a data-mart is a database made available on the HealthAgents network via a software agent). All nodes are protected by firewalls and their HealthAgents connections are available in 'de-militarized zones' (DMZ; a fully firewalled network) outside of the main medical networks in which agents are free to communicate. The databases behind the firewalls are not anonymized, whereas databases in the DMZ are either link-anonymized or fully anonymized depending on local laws and norms. This is what necessitates the anonymization process when porting data to the DMZ.

Agent technology provides access to the various functionalities that a particular machine on the network can provide. For example, *data-mart agents* provide access to a database and *classifier agents* provide access to classification tools. By distributing the data and functionality of the system, the whole system is less likely to fail, if any one particular machine or node fails. To allow for agent discovery, a cooperative network of *yellow-pages agents* is used, which allows agents to submit queries based on a description of a remote agent's functionality. Classification is the main functionality provided by the network and *classifier training agents* within the network are able to

Figure 1 The logical architecture of the HealthAgents system showing the various countries and hospital data-marts involved. Data-marts in the hospitals run entirely outside of the hospitals' own networks. The British and Valencian data-marts aggregate data from multiple hospitals into a single database, unlike the Barcelona hospitals that each provide their own data-mart. The data-marts are made available on the agent network through the database agents

Figure 2   The top-level concepts in the HealthAgents Language (HAL) ontology

train classifiers from data sourced from the network matching specific requirements. These trained classifiers can then be shared on the network as *classifier agents* that provide classification for new data combinations.

The agent network leverages semantic web technologies to provide interoperability between the various agents by using ontologies. Here, we refer to an ontology as 'an explicit specification of a conceptualization' (Gruber, 1993); that is, a well-defined vocabulary that explicitly defines what certain concepts mean and their relationship with other concepts. The agent network has an ontology for its own use and allows the use of any other ontologies for the definition of data items.

The HealthAgents Language (HAL) is the definition of the language that the agents use to communicate and it is defined in an ontology using the Resource Description Framework Schema (RDFS; vocabulary of describing vocabularies). The HAL ontology defines the messages and their properties that are used in agent communication and it is dedicated to this task. For example, it contains definitions of the message to submit a query to a yellow-pages agent and a definition of the result message that will be returned. Figure 2 shows the top-level concepts in the HAL ontology.

Data-level constructs can be defined by importing ontologies into the HAL messages. To allow the communication of data relating specifically to brain tumour diagnosis, we have developed a domain ontology describing brain-tumour-related concepts. This nomenclature is called the HA-dom (HealthAgents domain) ontology, and is an amalgamation of various sources, including the World Health Organization's (WHO) brain tumour classification (Weiss, 1994).

As data-marts on the system are operated by hospitals and sites with relational database expertise, we have designed the data-marts to operate via an ontological mapping between HA-dom and a relational database using D2RQ (Bizer, 2006). This allows the agent system to fully utilize the flexibility of semantic Web technologies while also allowing the data sites to utilize their existing human expertise for relational database management. By designing the agent system to utilize semantic Web querying mechanisms (such as Resource Description Query Language (RDQL) or Simple Protocol and RDF Query Language (SPARQL)) we build in maximum flexibility for the integration of different functionality as the network grows, as well as providing the ability to run more advanced reasoning over the data.

The organization of the network is such that each node has its own yellow-pages agent. This allows the agents on the local platform to find other agents on the same platform with a search to the local yellow pages. The local yellow-pages agent is connected to a network of yellow-pages agents and it shares its data with the network. The network also shares its data periodically.

Figure 3 The basic layer structure of an agent

After a while every yellow-pages agent knows about all other agents in the network, minimizing the search delay. More about this organization is described in Section 5.1.

It is important to emphasize that a fully functional prototype of the HealthAgents decision support system has been fully used, and has produced clinically relevant classification results (González-Vélez *et al.*, 2009).

5 Framework

The HealthAgents framework provides the implementation of the agent framework on which the HealthAgents architecture is deployed. The whole design rationale behind the agent framework was to provide a means for the integration of new functionalities into the network with minimum programming effort. This allows less technical centres to provide agents for their specialized functionality, as well as minimizing the development time of new functionalities. The agent framework has also been built to be agnostic to the specific application that runs on it, so although it has been designed and built with the application of brain tumour decision support at the forefront, it is applicable in any domain, medical or other. In fact, the ease with which agents are created indicates that it would be a good candidate for the basis of other projects requiring distributed networks.

In this section, we describe the microlevel design of an agent in the network. All agents in the network use the same design that is based on a layered design pattern.

The general pattern is shown in Figure 3. The pattern enforces incoming messages to filter up through the various layers to get to the top where the agent's functionality is provided. From the bottom to the top, the layers are as follows:

- Messaging interface: This is the network-level interface that is used to gather incoming messages and pass them on to the security layer. This layer is multithreaded so that the agent can receive many incoming messages. This layer is also used to deliver messages onto the network. The messaging layer is also responsible for the parsing of incoming messages into a canonical format that the other layers of the agent can read.
- Security layer: This layer provides a filter for incoming messages, which ensures that they are suitable for processing by the agent. This layer can include certification checks, message typing checks and other simple filters.
- Business methods: This layer is provided as part of the agent implementation and controls the flow of authenticated messages from the lower levels into the agent functionality. It deals with the routing of messages that are identified as being intermediate messages in an ongoing conversation, and handles framework-level messaging such as pings and pongs.
- Application programming interface (API): This layer is provided as a simplified API to which a developer can write their own agent functionality. This layer provides the functionality of the specific agents in the HealthAgents network—the data-marts, the classifiers, the yellow pages, etc.

The messaging layer is distinctive from the general business methods of the framework and is abstracted through a defined API, such that any message passing system can be used. All application code is programmed on top of the HealthAgents framework allowing this to be migrated to any agent platform by reimplementing its agent (message passing) layer. For the HealthAgents project it was

decided that the implementation of the agent layer would be provided by the Java Agent Development Environment (JADE), a ready-made, open-source, agent platform. The messaging layer API implementation is provided through a class that, in this case, uses the JADE agent platform for sending basic messages that conform to the ACL (Agent Control Language) defined by the Foundation for Intelligent Physical Agents (FIPA; Bellifemine *et al.*, 2001). The messaging layer is also responsible for ensuring the secure passage of messages from one agent to another. The implementation we have provided uses a standard public-key methodology for the encryption of messages, which provides confidentiality (cannot be snooped upon), integrity (cannot be changed en route), identification (through digital signatures) and non-repudiation (through digital signature verification).

The use of HAL for conversation follows the same philosophy, which is to separate the communication language of the HealthAgents agents from language used by the agent platform. Applying abstractions on top of the agent platform minimizes any extra reasoning we may automatically gain from the underlying agent platform, but economically, it makes good business sense to minimize dependency on external libraries. That said, for maximum flexibility the HealthAgents platform defines an API that allows any language to be used for communication. An implementation of this API provides a parser for HAL utilizing Sesame (Broekstra *et al.*, 2002), allowing messages to be passed around the system using RDF (Beckett, 2004; as synthesized in TURTLE (Beckett, 2007)). Using semantic Web technologies in this way makes it easier to surmount the challenges of interoperability among the agents and interoperability among the wide variety of data sources using different schemas. As the vocabulary that the agents can understand gets larger, the semantic Web technologies provide the extensibility through reasoning over ontologies.

Security is a crucial part of the system and each message arriving at an agent has to be authenticated by the 'security guard' that resides in the security layer. Again, this is realized as an API that can be implemented in different ways depending on the necessary level of security, the type of the agent or the location. This interface provides independence of the agent platform's secure passage functionality that may or may not be used. We are implementing a security layer that uses digital certificates to securely identify the route by which messages reached the agent and, using policy rules, to allow or deny access to the incoming message.

Interoperability is one of the challenges that needed to be addressed when constructing our distributed system. More specifically, communication can only be established if every agent can understand the languages used by others. Such a mutual understanding can be achieved by either forcing global consensus and regulating every participant to comply with it or recommending a reference vocabulary and requiring individuals to map their local vocabulary against the reference vocabulary. In HealthAgents, we have developed HealthAgents Domain Ontology (HADOM) for interoperation at a data level, but hospitals are free to use their own schemas internally. A wide variety of databases exists within the HealthAgents consortium and having a shared conceptualization of the domain in an ontology allows these databases' schemas to be mapped onto the domain vocabulary to provide a consistent view of them from the agents' point of view.

5.1 Directory services

For agents to be able to communicate with each other they need to find each other's location. Agents are identified using agent identifiers that are specific to the underlying message-passing platform. The framework abstracts these into an API such that agent identifiers can be communicated without having to resort to understanding underlying agent identifiers.

The yellow-pages agent is a directory of agent identifiers. These agent identifiers are indexed by the abilities that the agents have advertised themselves as having. The abilities are defined in the HAL ontology and also the data-level ontology. When agents start up, they advertise their services to their local yellow pages, which stores their record and flags it as authoritative. An authoritative record is the record created by the yellow pages to which an agent originally registered.

The yellow-pages agents in the system then form a network that shares data in a similar manner to the internet DNS system. The local yellow-pages agent is configured such that it has the agent
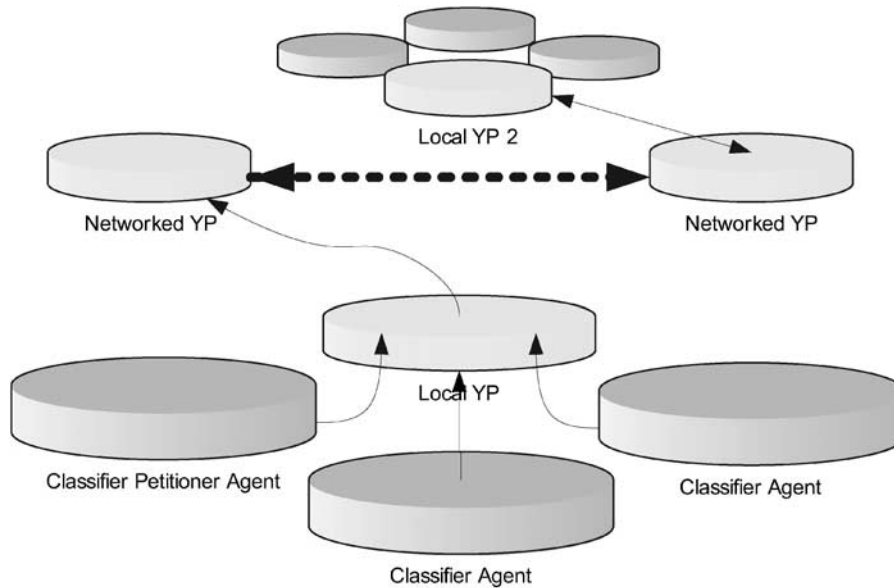
Figure 4 Network of yellow-pages agents

identifier of one of the yellow-pages agents that are outside of the hospital nodes, creating a persistent network of directory services, as shown in Figure 4. The local yellow pages regularly sends its directory to the other yellow pages it knows of. These receive the data and store the entries as non-authoritative. That way it is always known that non-authoritative data may be out of date, as the currency of the directory depends on the period of synchronization. The networked yellow-pages agents then pass on their aggregated data to the other yellow-pages agents it knows of until the complete network of yellow-pages agents are all synchronized.

The currency of any yellow-pages directory is tested regularly by pinging all of the agents in the directory. Agents that are still available respond (with 'pong' messages) and those that do not respond within a given timeout period are removed from the yellow-pages directory (see Figure 5). The response to pings is handled by the business methods layer of the agent design, and is thus handled automatically, without any intervention needed from an agent developer.

Before a yellow-pages agent receives a data synchronization message from its network neighbourhood, it may be unable to answer queries from its local agents. In this case, the local yellow-pages agent forwards the search query on to one or more networked yellow-pages agents and then aggregates and caches the results.

The yellow pages is an integral part of the agent framework, that is, the system does not run without it. However, this dependency meant that it was possible to make the system easier to program. The abstract classes of the framework have built-in methods for searching the yellow pages. The code snippet below shows the Java for searching for a classifier-petitioner agent. A set is created containing the abilities of the agent for which we are searching (the ability is that it has the type ClassifierPetitionerAgent, as defined by HAL) and the *searchYellowPages* method does all the work of creating the appropriate messages and dealing with the waiting for the asynchronous result from the network.

```
HashSet<AgentAbility> s2 = new HashSet<AgentAbility>();
s2.add( new SpecificAgentAbility(
      new Concept(Ontology.AA_HAS_TYPE),
      new Concept(Ontology.CLASSIFIER_PETITIONER_AGENT)));
Set<AgentIdentifier> petitioners=
   getMessagingService().searchYellowPages(s2);
```

Figure 5 Message diagram displayed in real time during the execution of the system showing the yellow-pages directory currency checking messages. The example shows five pings being sent and only four pongs being returned. The status screen of the yellow-pages agent is shown during and after the ping-pong mechanism to show that the database agent that did not respond has been removed from the directory (the message in the bottom right of the status screen also indicated the removal of one agent)

5.2 Database communication

The HADOM is used as the common referencing point among different hospitals maintaining their own vocabularies and database schemas. Both the integrity and the independence of legacy databases is fully respected. The discrepancy between such schemas is resolved by dedicated mappings between each individual schema and the HADOM ontology. By referring to (e.g. with rdf:seeAlso) or making HADOM concepts equivalent to (e.g. with owl:sameAs) those from the external resources, any new advances and updates in the external resources can be reflected in HADOM. Thus far, HADOM is based on publicly available documents and information gathered from various sources, such as the WHO classification of tumours affecting the central nervous system (Weiss, 1994; Tatter, 1996) and other relevant sources (DeAngelis, 2001).

These documents are based on universal knowledge of the domain and may not be strictly applicable in HealthAgents or not sufficient enough to reflect all the necessary project-specific domain restrictions. Its use as a mediator in the interoperability of agents allows easier integration of functionalities into the system than it might otherwise have been in a purely relational database system. There is a remote agent in the prototype that acts as a proxy to the database.

All data-related information in the messages in the HealthAgents network is referenced to the HADOM ontology. A database agent transforms these messages into SQL (Structured Query Language) queries that are executed against a relational database.

6 Evaluation

The deployment of the agent framework to create a European-wide decision support system for brain tumour diagnosis is still underway, but the core functionality is now being tested with the data assembled by the data-collecting sites during the development of the agent framework.

USOU and MicroArt are functioning as a small 2-node network of yellow-pages agents such that all joining nodes may eventually determine the location of all agents on the network.

The graphical user interface that has been developed at MicroArt for providing access to the agent framework, the data and all the classifiers, has been deployed at hospitals in the United Kingdom and Spain.

So far, all but one of the hospitals are running data-marts, and the classifier producers at ITACA have deployed classifiers and an agent specifically for gathering data for producing new classifiers.

Tests are in the early stages, but it has been shown that data from one hospital can be classified by a classifier in a distant node, when that hospital has no previous knowledge of the classifier node, because the network of yellow pages. We have also shown that, with appropriate permissions, the classifier producers can gather data from data-marts as and when they require it for creating new classifiers.

6.1 Implementation

At the time of writing, the HealthAgents system is being deployed in various sites around Europe. Hospitals in Birmingham, United Kingdom, Barcelona, Spain and Valencia, Spain have installed the data-mart node and are gathering MRS data to be used for creating classifiers. Classifier nodes in ITACA and KULeuven providing pre-built classifiers that were trained on the data from the Interpret Project are online. Such large-scale aggregation of data is not only novel, but also somewhat unusual in the field of medicine.

However, the distributed nature of the aggregation means that the entrycosts for joining a system are very low and the amount of 'buy-in' required is very small; that is, hospitals, classification producers and other tool providers can leave and join as they see fit (within some specified boundaries). Indeed, the consortium has had much interest in the system from parties external to the project, and the Province Hospital of Jiangsu has already agreed to join the network and provide MRS data to the distributed data-mart. As previously described, other decision support systems based on MRS data do not have the flexibility of the HealthAgents system.

The agent-based distribution of the network provides low-entry requirements while also providing the flexibility to reconfigure networks as necessary, improving network robustness. The Interpret system was not based on semantic Web technologies, but the use of a domain ontology in HealthAgents provides a well-defined description of the data that eases interoperability issues, particularly when hospitals that already have their own data schema join the network. The use of the ontology to define the data also provides future proofing for the data schema. During the project, the World Health Organization released a new version of their vocabulary and it was because the technology of ontologies provides great extensibility that we were able to extend the vocabulary to cover the new terms.

We have not provided any quantitative evaluation of the agent framework's speed; this is a considered decision. Although we could measure the performance of the agent network for answering queries across the system, speed variances caused by the reliability of the network when using a European-wide distributed network would not show anything indicative of the performance of the framework itself. The main speed bottleneck of any particular agent is the parsing of an RDF message into objects. As we are using Sesame for this task, the performance of an agent's parsing routine depends on the performance of Sesame.

Nevertheless, to give an idea about the performance of the network:

Classifications are usually returned within 10 seconds. Most of that time is network related. On a single desktop-specification machine, a classification request and reply (via the Classifier

Petitioner Agent) happens in under 2 seconds. We have also found that the message encryption routines impart a large overhead (around 100 %) on communication performance.

For a qualitative clinical evaluation, the HealthAgents consortium are conducting a TAM (Technological Acceptance Model) study, which strives to gauge the acceptance of the system from the point of view of the clinicians that use the system.

6.2 Discussion

The design of the agent architecture was primarily based around the need to provide aggregation of data from many different disparate locations and institutions. The technological know-how of these institutions may be limited, so the architecture was designed in such a way as to minimize the necessary familiarity with agent-based and distributed systems. The partners in our project that were building agents based on the framework have commented that the framework is easy to use and this is borne out by the large number of distinct agents that have been created during the project by non-agent experts. Debugging of truly distributed systems such as HealthAgents system is very difficult and we encountered some challenges when deploying the system Europe-wide. We built some debugging tools (such as the message diagram and the yellow-pages agent status screen shown in Figure 5) that really helped to diagnose some problems and generally we found that agents performed as expected.

However, there are still a few small areas in the framework where a little tittivating could improve the overall simplicity of its use. In particular, the creation of the HAL ontology was effectively a distributed task, as programmers in the project found that they required new predicates by which they could refer to values. In general, HAL was built in a relatively *ad hoc* manner. That is not to say it is unworthy, but to say that should further programmers require new ways to refer to data structures, they will be required to understand the basics of semantic Web technologies at the very least. RDF is also the main cause for concern in the performance of the communications. The parsing of the messages into the Sesame RDF store (which also calculates all the inference) is one of the main bottlenecks of the framework. Of course, its performance is proportional to the speed of the machine on which it executes, thus throwing computing power at the issue can ease the problem. We have also designed the language centre of the framework to be a plugin; therefore, should new ways of dealing with graphical data become available or industrialization of the framework require that the communication language be made non-graphical, either can be easily achieved through implementation of one API. Of course, using RDF plays two ways—its extensibility means that it is possible for new programmers to easily extend the language beyond its original scope; something that is much more difficult to achieve with the relational and functional means of creating language schemas.

Another way we hope to improve the agent framework is through agent workflows. Currently, the workflow of an agent must be handled explicitly by the programmer of the agent by arranging their code such that it performs the appropriate workflow. This is not a trivial task for some agents, for example, where certain messages must be handled prior to other messages, and it makes the maintenance of agent code much more difficult. We researched agent workflow systems and built a module for controlling an agent's messaging agenda using the agent workflow language called Lightweight Coordination Calculus (LCC; Robertson, 2004). As LCC is a very simple-to-understand language, the building of agents would become much easier for the programmers. Indeed, it would be possible for simple agents to be created by non-programmers. Conversion programs exist for converting other agent workflow languages to LCC so in that way we encompass a wider gamut of the workflow community. LCC agent workflow language Agents following LCC workflows were tested in controlled situations and correctly followed the appropriate workflow. In the test scenario, the framework was extended to make these workflows as simple to use as possible. Unfortunately, because of, time constraints, the workflow language was not used in the final HealthAgents system; however, we still hope to distribute this as part of the framework codebase and perhaps in the future the HealthAgents communication system will move to a workflow-based execution environment.

7 Conclusions

The fully distributed agent-based design for the data-marts is novel in the field of computer-based decision support. It is important to realize that previous brain tumour decision support systems have been centralized, making them limited in both the range of data available to them and their extensibility. Therefore, because the HealthAgents system is distributed, new data that becomes available on the network will help to increase the accuracy of the decision support, thereby making the system more valuable as the community builds. In addition, the layered approach of the agent framework makes the development of new agents very easy and therefore the possibilities for the extension of the system become greater than other decision support systems. We have shown that the agent design is highly extensible, which means that agents beyond the remit of the project can be deployed within the system easily and quickly, which could lead to a collaborative network of agents for many different medical decision support scenarios.

The use of semantic Web technologies has been an important factor in the design of the system such that the messages and the data payload of the messages are formally represented. With the increasing uptake of RDF, the system stands in good stead for interoperability with new and future systems. HADOM provides an important role in linking the nomenclature of hospitals across Europe into a formalized vocabulary that can be used to map the schemas of existing patient record systems in hospitals into the system. HAL provides an extensible, hierarchical language with which agents can communicate with each other. Together, HADOM and HAL provide a novel way in which decision support data can be viewed, which shows that the extensibility of semantic Web technologies makes them ideal for distributed systems and distributed development environments.

The aim of the HealthAgents project is to build a distributed decision support system for brain tumour diagnosi, and in this paper we have presented an extensible agent architecture that uses research-led methods to provide this novel application.

Acknowledgements

References

Annicchiarico, R., Cortés, U. & Urdiales, C. (eds) 2008. Agent technology and e-health. In *Whitestein Series in Software Agent Technologies and Autonomic Computing*. Birkhäuser Verlag.

Arús, C., Celda, B., Dasmahapatra, S., Dupplaw, D., González-Vélez, H., van Huffel, S., Lewis, P., Lluch i Ariet, M., Mier, M., Peet, A. & Robles, M. 2006. On the design of a web-based decision support system for brain tumour diagnosis using distributed agents. In *2006 IEEE/WIC/ACM International Conference on Web Intelligence & Intelligent Agent Technology (WI-IAT 2006 Workshop)*, IEEE, 208–211.

Beckett, D. 2004. Modernising semantic web markup. In *XML Europe 2004*, Dumbill E. (ed.). IDEA Alliance, `http://www.idealliance.org/papers/dx_xmle04/papers/03-08-03/03-08-03.html`.

Beckett, D. 2007. Turtle—Terse RDF triple language. Manual (version 20 November 2007), ILRT University of Bristol, `http://www.dajobe.org/2004/01/turtle/`.

Bellifemine, F., Poggi, A. & Rimassa, G. 2001. JADE: a FIPA2000 compliant agent development environment. In *AGENTS'01: Fifth International Conference on Autonomous Agents*. ACM Press, Montreal, Canada, 216–217.

Bizer, C. 2006. *D2RQ*. Freie Universität Berlin. `http://sites.wiwiss.fu-berlin.de/suhl/bizer/D2RQ/` (accessed 13 February 2007).

Boniface, M., Leonard, T., Surridge, M., Taylor, S., Finlay, L. & Mccorry, D. 2005. Accessing patient records within virtual healthcare organisations. In *eChallenges 2005: Innovation and the Knowledge Economy: Issues, Applications, Case Studies*, Cunningham, P. & Cunningham, M. (eds). ISBN 1-58603-563-0. IOS Press, Ljubljana, Slovenia.

Broekstra, J., Kampman, A. & van Harmelen, F. 2002. Sesame: a generic architecture for storing and querying RDF and RDF schema. In *First International Semantic Web Conference*, Lecture Notes in Computer Science **2342**, 54–68. Springer Verlag.

Brugali, D. & Sycara, K. 2000. Towards agent oriented application frameworks. *ACM Computing Survey* **32**(1), 21–27.

Brugali, D., Menga, G. & Aarsten, A. 1997. The framework life span. *Communications of the ACM* **40**(10), 65–68.

DeAngelis, L. M. 2001. Brain tumors. *New England Journal of Medicine* **344**(2), 114–123.

Georgiadis, P., Cavouras, D., Kalatzis, I., Daskalakis, A., Kagadis, G. C., Sifaki, K., Malamas, M., Nikiforidis, G. & Solomou, E. 2008. Improving brain tumor characterization on MRI by probabilistic neural networks and non-linear transformation of textural features. *Computer Methods and Programs in Biomedicine* **89**(1), 24–32.

González-Vélez, H., Mier, M., Julià-Sapé, M., Arvanitis, T. N., García-Gómez, J. M., Robles, M., Lewis, P. H., Dasmahapatra, S., Dupplaw, D., Peet, A., Arús, C., Celda, B., Van Huffel, S. & Lluch-Ariet, M. 2009. HEALTHAGENTS: distributed multi-agent brain tumor diagnosis and prognosis. *Applied Intelligence* **30**(3), 191–202.

Gruber, T. R. 1993. A translation approach to portable ontology specifications. *Knowledge Acquisition* **5**(2), 199–220.

Huang, J., Jennings, N. & Fox, J. 1995. Agent-based approach to health care management. *Applied Artificial Intelligence* **9**(4), 401–420.

Lanzola, G., Gatti, L., Falasconi, S. & Stefanelli, M. 1999. A framework for building cooperative software agents in medical applications. *Artificial Intelligence in Medicine* **16**(3), 223–249.

Merelli, E., Armano, G., Cannata, N., Corradini, F., d′Inverno, M., Doms, A., Lord, P., Martin, A., Milanesi, L., Möller, S., Schroeder, M. & Luck, M. 2007. Agents in bioinformatics, computational and systems biology. *Briefings in Bioinformatics* **8**(1), 45–59.

Rees, J. 2003. Advances in magnetic resonance imaging of brain tumours. *Current Opinion in Neurology* **16**(6), 643–650.

Robertson, D. 2004. Multi-agent coordination as distributed logic programming. In *International Conference on Logic Programming*, Lecture Notes in Computer Science **3132**, 416–430. Springer.

Shadbolt, N., Lewis, P., Dasmahapatra, S., Dupplaw, D., Hu, B. & Lewis, H. 2004. MIAKT: combining grid and web services for collaborative medical decision making. In *UK e-Science All Hands Meeting*, Cox S. (ed.). EPSRC, Nottingham, UK, 784–791. http://www.allhands.org.uk/2004/proceedings/.

Tate, A. R., Underwood, J., Acosta, D. M., Julià-Sapé, M., Majós, C., Moreno-Torres, A., Howe, F. A., van der Graaf, M., Lefournier, V., Murphy, M. M., Loosemore, A., Ladroue, C., Wesseling, P., Bosson, J. L., Cabañas, M. E., Simonetti, A. W., Gajewicz, W., Calvar, J., Capdevila, A., Wilkins, P. R., Bell, B. A., Rémy, C., Heerschap, A., Watson, D., Griffiths, J. R. & Arús, C. 2006. Development of a decision support system for diagnosis and grading of brain tumours using in vivo magnetic resonance single voxel spectra. *NMR in Biomedicine* **19**(4), 411–434.

Tatter, S. B. 1996. *The New WHO Classification of Tumors Affecting the Central Nervous System.* PhD thesis, Cornell University Medical College/The Rockefeller University. http://neurosurgery.mgh.harvard.edu/newwhobt.htm.

Turck, F. D., Decruyenaere, J., Thysebaert, P., Hoecke, S. V., Volckaert, B., Danneels, C., Colpaert, K. & Moor, G. D. 2007. Design of a flexible platform for execution of medical decision support agents in the intensive care unit. *Computers in Biology and Medicine* **37**(1), 97–112.

Weiss, S. 1994. Histological typing of soft tissue tumours. In *World Health Organization International Histological Classification of Tumours*, 2nd edn. ISBN 978-3-540-56794-3. Springer.

Wright, A. J., Arús, C., Wijnen, J. P., Moreno-Torres, A., Griffiths, J. R., Celda, B. & Howe, F. A. 2008. Automated quality control protocol for MR spectra of brain tumors. *Magnetic Resonance in Medical Sciences* **59**(6), 1274–1281.