

# Agent-Based Applications in Health Care

John Nealon<sup>1</sup> and Antonio Moreno<sup>2</sup>

<sup>1</sup> Department of Computing, Oxford Brookes University,  
Oxford OX33 1HX, UK

`johnnealon@brookes.ac.uk`

<sup>2</sup> Computer Science and Mathematics Department,  
Universitat Rovira i Virgili, ETSE Campus Sescelades,  
Av. dels Pasos Catalans 26, 43007-Tarragona, Spain

`amoreno@etse.urv.es`

**Abstract.** In this paper we introduce the main issues related to the deployment of agent-based systems in health care. First, we comment on the characteristics of health care problems and we argue that multi-agent systems are a good choice to tackle problems with these features. This belief is supported with a number of exemplar applications of agent-based systems in medical domains. We also discuss several lines of research that have to be covered before multi-agent systems can be successfully deployed in real health care settings. We conclude that multi-agent systems do have an increasingly important role to play in health care domains, because they significantly enhance our ability to model, design and build complex, distributed health care software systems.

## 1 Introduction

Health care at all levels - local, regional, national and international - is a vast open environment characterized by shared and distributed decision making and management of care, requiring the communication of complex and diverse forms of information between a variety of clinical and other settings, as well as the coordination between groups of health care professionals with very different skills and roles. It is the aim of health care software systems to operate effectively in this environment, in order to meet the information needs of patients and health care providers. Practitioners in health care environments, in particular, require that the information is both timely and error-free, such that recommendations or decisions offered by the software systems are secure and trustworthy [1].

There is a growing interest in the application of agent-based techniques [2], [3] to problems in the medical domain<sup>1</sup>. In this paper we explain which are the characteristics of the problems in this area, and we argue that multi-agent systems are indeed an interesting tool to solve them, since the usual properties of intelligent agents match quite precisely with our needs in this field (basically with the requirement of having autonomous intelligent proactive collaborative entities in a distributed environment). This volume is another empirical confirmation of this claim, since it presents an extensive list of agent-based systems that are being

developed across Europe to solve a wide range of problems in health care (from the management of organ and tissue transplants to the provision of personalised access to medical information to diabetic patients). All these systems are being developed by the members of the AgentCities.NET Working Group on Health Care Applications [7].

Despite the adequacy of multi-agent systems in the building of health care systems, it must be stressed that there are still many research topics that have to be thoroughly studied before multi-agent systems may be successfully deployed in real health care settings. Section 3 offers a summary of the main current research and development topics relevant to the application of agents to health care.

## 2 Agent-Based Systems in Health Care

### 2.1 Agents and Multi-Agent Systems

An agent is a software entity that applies Artificial Intelligence techniques to choose the best set of actions to perform in order to reach a goal specified by the user. It is normally assumed that they have the following properties:

- They can react timely and flexibly to the dynamic and unexpected changes in their environment.
- They have an autonomous and independent behaviour, which is not controlled by any external entity.
- They can take the initiative and perform proactively actions that may help them to reach their goals.
- They can communicate with users or other agents. Thus, they can exchange information, engage in complex negotiations, and coordinate their activities to cooperate in the joint resolution of a problem.
- Agents usually have reasoning, planning and learning capabilities that allow them to display an intelligent behaviour.

A multi-agent system [2], [3] might be defined as a collection of autonomous agents that communicate between themselves to coordinate their activities in order to be able to solve collectively a problem that could not be tackled by any agent individually. In recent years it has been argued that multi-agent systems may be considered as the latest software engineering paradigm [8], [9]. This kind of system may be used in domains with the following features:

- The knowledge required to solve the problem is spatially distributed in different locations.
- Several entities, while keeping their autonomous behaviour, have to join their problem-solving abilities to be able to solve a complex problem.
- The problems in the domain may be decomposed in different sub-problems, even if they have some kind of inter-dependencies.

## 2.2 Characteristics of Problems in the Health Care Domain

It is interesting to note that many problems that appear in health care share a number of similarities. Noting these similarities is a good first step towards finding a framework that may be used to approach most of these problems. Among the most important characteristics of problems in the medical field we can cite the following:

- It is very usual that the knowledge required to solve a problem is spatially distributed in different locations. For instance, the problem of patient scheduling [10, 11, 12] consists in scheduling the different tasks to be performed on a hospitalised patient (e.g. a number of different analysis and tests). Normally, each unit of the hospital keeps its own information about the patients hospitalised in that unit and about the schedule of the activities with the unit's equipment. There are units that provide services to all the other units of the hospital, such as X-rays or blood tests. It is not trivial to coordinate the schedule of different tests to be performed in different units, especially if, apart from the temporal restrictions derived from the separate location of different units, there are also medical restrictions among the tests (e.g. one test has to be performed at least two hours after another test).
- The solution of a problem involves the coordination of the effort of different individuals with different skills and functions, usually without the supervision of a single centralised coordinator. The provision of health care typically involves a number of individuals - inpatients, outpatients, physicians, nurses, carers, social workers, managers, receptionists, etc. - located in many different places. Patients could be at home, at work, in hospital, or on vacation, while the health care providers are often at a number of institutions or locations within institutions, providing services such as diagnosis, treatment, surgery, laboratory testing, radiography, and so on. All these people must coordinate their activities to provide the best possible treatment to the patient.
- Health care problems are quite complex, and finding standard software engineering solutions for them is not straightforward. For instance, coordinating the process of organ transplant in a country is not an easy task [13]. When a new organ is available, it is necessary to find very quickly the most appropriate receptor of the organ, which may be located in a medical centre hundreds of miles away from the donor's location. Furthermore, as commented above, each hospital keeps the data of the patients of that hospital who are in the waiting list for a certain type of organ. It would be quite difficult to design and implement a centralised complex system (e.g. a standard decision aid expert system) to solve this coordination problem.
- In the last few years there has been a shift in health care practice towards health care promotion, shared patient-provider decision-making and managed care, creating an increased demand for information and online services [14, 15]. The shared decisions and actions of all concerned need to be coordinated to make sure that the care is efficient and effective. To facilitate

this decision-making task, and to ensure the communication and coordination processes run smoothly, software systems are needed that will reduce errors in diagnosis and treatment, deliver health care to remote locations, improve medical training and education, and make health care information more accessible to patients, their families and carers, alike. To meet these needs the software systems must be proactive in anticipating the information and knowledge needs of users and deliver it in a timely manner, support synchronous and asynchronous communication, and facilitate collaborative decision making between the various individuals involved in the process of managing and delivering health care services.

- There is a great amount of medical knowledge available on the Internet. It is necessary to provide ways of accessing the most relevant information as easily, flexibly and timely as possible. This access to medical information is necessary both for medical practitioners and for patients. In the former case, they have to be aware of all the new medicines, techniques and treatments appearing in their field of expertise; however, most practitioners lack the time to search for all this information and filter the one they need exactly. In the latter, citizens demand a more active role in the management of their care, and they want to find relevant information as easily as possible (but they usually lack the skills to search for the most adequate knowledge relevant to their specific personal needs). In both cases it is important to receive appropriate information from useful and reliable sources in a proactive way, without having to devote time and effort to look for, analyse, evaluate and filter it.

### 2.3 Adequacy of Agent-Based Systems to Health Care Problems

We would like to argue in this paper that multi-agent systems offer an appropriate tool to tackle this kind of problems. Some reasons to support this claim are the following:

- The components of a multi-agent system may be running in different machines, located in many different places. Each of the agents may keep part of the knowledge required to solve the problem, such as patient records held in different departments within a hospital or in several hospitals, clinics and surgeries, in an insurance company, or in government organisations. Therefore, multi-agent systems offer a natural way of attacking inherently distributed problems.
- One of the main properties of an intelligent agent is sociability. Agents are able to communicate between themselves, using some kind of agent communication language, in order to exchange any kind of information. In that way they can engage in complex dialogues, in which they can negotiate, coordinate their actions and collaborate in the solution of a problem (e.g. different units of a hospital may collaborate in the process of patient scheduling [10, 11, 12]).



- When a problem is too complex to be solved in a single system, it is usual to decompose it in subproblems (which will probably not be totally independent of each other). In multi-agent systems there are techniques of distributed problem solving [16], in which a group of agents may dynamically discuss how to partition a problem, how to distribute the different subtasks to be solved among them, how to exchange information to solve possible dependences between partial solutions, and how to combine the partial results into the solution of the original problem. Thus, multi-agent systems can handle the complexity of solutions through decomposition, modelling and organising the interrelationships between components.
- Agents can also be used to provide information to doctors and patients. There are information agents (also called Internet agents, [17], that are specialised in retrieving information from different sources, analysing the obtained data, selecting the information in which a user is especially interested, filtering redundant or irrelevant information, and presenting it to the user with an interface adapted to the user's preferences.
- Another important property of agents is their proactivity; their ability to perform tasks that may be beneficial for the user, even if he/she has not explicitly demanded those tasks to be executed. Using this property they may find relevant information and show it to the user before he/she has to request it. For instance, if it knows that the user has had heart problems in the past and might need this information urgently, a personal agent that also knows that the user is about to travel abroad could look for information about the medical centres in the towns to be visited that have a cardiology department.
- The basic characteristic of an intelligent agent is its autonomy. Each agent takes its own decisions, based on its internal state and the information that it receives from the environment. Therefore, agents offer an ideal paradigm to implement systems in which each component models the behaviour of a separate entity, that wants to keep its autonomy and independence from the rest of the system (e.g. each unit of the hospital may keep its private data, or each hospital may use a different policy to rank the patients that are waiting for an organ transplant).

Thus, we think that the basic properties of intelligent agents (autonomy, proactivity, social ability) and the features of multi-agent systems (management of distributed information, communication and coordination between separate autonomous entities) suggest that they offer a good option to consider when trying to solve problems in health care domains.

## 2.4 Fields of Application within Health Care

Intelligent agents have already been proposed to deal with many different kinds of problems in the health care domain (see [4], [5], [6], [7] and the rest of the chapters of this volume). Just to give a short list of examples, some of the fields in which they are already being applied are the following:

- Patient scheduling: in patient appointment scheduling, where medical procedures have become more complex and their tests and treatments more interrelated, manual and traditional software solutions have been shown to be inadequate while a multi-agent solution gave significantly improved results [10, 11, 12]. The complexities of medical appointment scheduling have been successfully formalised and implemented in an agent framework [18].
- Organ and tissue transplant management: the agent-based coordination of tissue or organ transplants across a hospital [19, 20, 21] could provide significant improvements in the time required to pull together the resources required for a transplant operation. This approach could also be feasible at a regional level where each hospital has a list of waiting patients, and when an organ is available somewhere in the region, the hospital transplant coordinators must get in touch and quickly find the most appropriate recipient [13].
- Community care: coordinating all the activities that have to be performed in order to provide an efficient health care to the citizens of a community (especially older or disabled citizens). Agents can provide remote care monitoring and information for such groups as the elderly and chronically ill. There is an elderly care management system [22] in which one agent is associated with each elderly person that is responsible for receiving medical data, giving reminders to the person, and alerting the medical centre if something is wrong.
- Information access: the deluge of medical information available on the Internet has led to the development of information agents to collect and organise this information, such as the Multi-Agent Retrieval Vagabond on Information Networks (MARVIN) [23], developed by the Health On the Net Foundation and the Swiss Institute of Bioinformatics, or a multi-agent system that helps to manage the UK National Electronic Library for Communicable Diseases [24]. An information agent, based on a user profile, has proactively performed the role of locating, assessing, retrieving, filtering and presenting information from many distributed sources on a periodic basis [25]. An intelligent user interface to adapt to a clinicians requirements, specialism and the characteristics of diabetic patients whose records are being accessed uses a multi-agent framework to coordinate these possibly conflicting requirements [26]. It has also been implemented a multi-agent system that provides mobile users with information about the medical centres or the doctors available in a particular town, and that lets the user access his/her medical record or book a visit to be examined by a doctor [27].
- Decision support systems: A distributed decision support system based on the multi-agent paradigm can monitor the status of a hospitalised patient and help to diagnose the state of the patient [28], or support co-operative medical decision- making [29], [30].
- Training: agents can help to improve medical training and education in distance- learning tutoring systems [31].
- Internal hospital tasks: Patient information retrieval and workflow management using agent communication techniques and medical ontologies is being

applied to the management of patients suffering from stroke [32, 33]. A cooperative multi-agent framework can support the heterogeneous transaction workflow process among the people involved in patient care management [34]. Multi-agent systems have also been suggested for monitoring the application of medical protocols [35], or controlling the usage of restricted use antibiotics [36].

- Senior citizen care: a group of special interest for the application of agent-based systems are the senior and the disabled citizens, to whom this technology could be useful to help to increase their ability to lead an independent life. Those agent systems have to be devised to provide aid in carrying out activities of daily living, and health care maintenance. In addition, they will provide links to the outside world, including entertainment and information, and will facilitate communication with family and the environment. Their functions may include standardised behavioural assessments useful in medical monitoring. These kinds of tools may be used to facilitate the health care and social interaction of senior citizens person, and may delay their institutionalisation by prolonging the period of relative independence [37]. A European IST project, TeleCARE, aims to design and develop a configurable agent-based framework for virtual communities focused on supporting assistance to elderly people employing tele-supervision and tele-assistance [38].

### 3 Research and Development Challenges

There are several issues that must be addressed if multi-agent systems are to be successfully deployed in real world health care applications. A discussion of the most significant of these follows.

#### 3.1 Communication standards

Health care systems are complex, diverse and dispersed. Consequently, the development, dissemination and utilization of common communication standards, vocabularies and ontologies are and will be central to the development of multi-agent systems in health care. In the main, communication standards (sometimes called specifications or protocols) already exist. For some time the EU and US standardization bodies for health care information and communications technology have been overseeing the development of communications standards. The EUs CEN/TC 251 [39] aim is to achieve compatibility and interoperability between independent systems, to support clinical and administrative procedures, technical methods to support interoperable systems as well as requirements regarding safety, security and quality. The two most well known US standardization bodies, the American Society for Testing and Materials Committee on Healthcare Informatics (ASTM E31) [40] and Health Level Seven (HL7) [41], are involved in similar work. ASTM E31 is developing standards related to the

architecture, content, storage, security, confidentiality, functionality, and communication of information. HL7 is mainly concerned with protocol specifications for application level communications among health data acquisition, processing, and handling systems. Its scope is wide in that it attempts to meet the communication requirements of entire health care organisations, while most other efforts focus on the requirements of a particular department, whereas most standards organisations produce standards for a particular healthcare domain such as pharmacy, medical devices, imaging or insurance (claims processing) transactions. HL7's domain is clinical and administrative data. "Level Seven" refers to the highest level of the International Standards Organisation's communications model for Open Systems Interconnection. HL7 defines the data to be exchanged, the timing of the interchange, and the communication of certain errors to the application. The seventh level supports such functions as security checks, participant identification, availability checks, exchange mechanism negotiations and, most importantly, data exchange structuring.

### 3.2 Ontologies

Ontologies are being developed specifically for health care applications. The best known examples are GALEN [42], the Unified Medical Language System (UMLS) [43], Systematized Nomenclature of Human and Veterinary Medicine (SNOMED) [44]. The aim of GALEN is to produce a computer-based multilingual coding system for medicine. The GALEN Programme is developing a clinical terminology - the GALEN Common Reference Model for medical concepts. The medical concepts represented can be represented using a scheme that can both be manipulated by computers and accessible to health care professionals. The representation scheme that is being used to build the GALEN Common Reference Model is known as GRAIL - the GALEN Representation And Integration Language. It is intended for use by clinical application builders, both when developing clinical applications, and as a run-time resource when those applications are in service. The U.S. National Library of Medicine (NLM)'s UMLS project develops and distributes multi-purpose, electronic "Knowledge Sources" and associated lexical programs, for use in the development of systems concerned with patient records, digital libraries, Web and bibliographic retrieval, natural language processing, and decision support. The UMLS project is a long-term NLM research and development effort designed to facilitate the retrieval and integration of information from multiple machine-readable biomedical information sources. Major barriers to effective retrieval and integration of information from these sources include the variety of vocabularies and classifications used in different sources and by different users and the sheer number and wide distribution of potentially relevant information sources. These barriers deter health care professionals and researchers from using available machine-readable information and also hamper the development of effective search interfaces that might assist these users. The UMLS is a relational database connecting by concept over 60 vocabularies, thesauri, medical problem lists, etc. Its purpose is to make it easy



for health professionals, medical librarians, and researchers to retrieve and integrate information from different machine-readable sources as computer-based patient records, databanks, bibliographic and full text databases, and expert systems. SNOMED's design is based on the premise that a detailed and specific nomenclature is essential to accurately reflect, in computer readable format, the complexity and diversity of information found in a patient record. SNOMED is designed for applications such as telemedicine, population-based outcomes analysis, cost-effectiveness studies, practice guidelines and the integration of electronic medical record information into a single data structure. Nevertheless, the story on ontologies is not as clear as that on communication standards. Although bioinformatics and health care informatics are fields that have active communities developing ontologies, as we have seen with GALEN, UMLS and SNOMED, their use has lagged behind their potential, despite the huge drive by health care professionals to bring health care information into clinical workstations and onto the Internet. Why is this? GALEN provides a common terminology that is currently of limited scope, while UMLS lacks a strong organisational structure, and SNOMED provides only diagnosis nomenclature and codification. Those who are building health care applications are using ontologies that have been or are being developed to meet specific needs, each with its own representation of the world, suitable to the purpose it has been developed for. There is as yet no common ontology. The experience of system developers [45] strongly suggests that the development of a single or a small number of ontologies for the health care domains is non-productive since no single domain requires such an ontology, and the overlap between ontologies is often minimal. Perhaps, this challenge will eventually be taken up in the mould of CYC [46], but that is not likely in the immediate future. Alternatively, machine learning techniques for the automatic construction of ontologies are being developed [47, 48, 49, 50]. In the meantime, researchers working in related areas are sharing and extending existing ontologies.

### 3.3 Security and privacy issues

Security is concerned with the protection of information from unauthorised access while stored and communicated, and privacy protection with the avoidance of unnecessary identification. The increasing dependence on information and communication technologies in health care organisations in order to collect, transmit, store, and assess data has brought the dual issues of security and privacy to the fore. It is accepted that medical data, such as genome information, medical records and other personal information must be treated with the greatest respect with regard to privacy and privacy [51]. The European Union has been enforcing the protection of medical data of individuals since 1995 [52], and in many countries, such as in Spain, these recommendations are among the most modern state laws that address the specific protection of this kind of data. Recently a US law that regulates the treatment of medical data was passed [53]. This law establishes which rules must be followed by the software business in this field, and the fines that will be applied to those that do not comply with the



regulations. This text, however, does not force any specific standard for storing or transmitting data. The usual properties of confidentiality, integrity and non-repudiation should be guaranteed in any agent-based health-care system. The use of cryptographic methods is also important to protect the access to data while it is being transmitted between agents (some issues related to secure communication are discussed in [54]). Frameworks that help to construct multi-agent systems are beginning to address security issues. For instance, an authentication mechanism is being added to Jade [55, 56, 57]. A user will give a username and a password to enter the system; then the system will provide an identity certificate that may be used by the agents associated with that user. This certificate will include aspects such as the identities of the user and of the emitting entity, identification of the algorithms used to protect the certificate, and its validity period. It is also suggested in [58] that agents should also have authorization certificates that allow them controlled access to particular resources; these permissions could be obtained from the user at creation time or by delegation of other agents of the system. In summary, it is obvious that the research on secure access to data, such as authentication of users, delivery and use of certificates, cryptographic methods, or security in wireless communications, will be fundamental to ensure that agents may safely deal with medical data and only authorised users may access or update this kind of information.

### 3.4 Safety critical issues

Agent-based systems are increasingly being used to support decision making in health care, a domain intrinsically uncertain and hazardous. Clinicians are well aware of these risks, and a recurring theme in physicians' criticism of health care computer systems is one of doubts concerning reliability, and the associated acceptance criteria for reliability which are related to risk and safety considerations. Consequently, everyone working in this area should be aware that very small system errors or misunderstanding in the specification of systems can lead to catastrophic consequences [1]. There exist powerful theoretical models of agency [59, 60, 61], and techniques for the specification and formalisation of agent-based systems, and techniques to analyse and manipulate such specifications, utilising symbolic information to determine appropriate behaviour, are beginning to emerge, based on well-established research from AI [62, 63, 64]. For example, tableaux for multi-modal belief or knowledge logics [65], tableaux for multi-modal belief or knowledge logics with linear-time temporal logic [66], tableaux for belief-desire-intention logics with either linear or branching-time temporal logics [59], and resolution for knowledge logics with linear-time temporal logics [67].

### 3.5 Legal issues

The distributed aspect of multi-agent systems and the lack of electronic borders facilitate the construction of multi-agent systems in which agents are representing organisations and individuals from different countries. If a multi-agent system

is used to enhance the provision of health care to citizens of a wide area, such as coordinating the management of organ transplants in different countries, one of the most difficult issues to consider is the fact that agents should conform to the local, national and international regulations in their area. The use of deontic logic would be appropriate to formalise and reason about the permissions and obligations of each agent in a system. A very promising research direction in this area is the definition of electronic institutions (see e.g. [68], [69]). An electronic institution includes a performative structure (a graph of scenes, where agent-agent interactions take place), a dialogical framework (which comprises an ontology, a set of illocutions and protocols to use them through conversation graphs) and a set of norms, which determine obligations that an agent may acquire through its actions. An e-institution prescribes the actions an agent may take and where, and imposes limits on the questions and answers which form the conversation between agents. Thus, an institution simplifies the task of preparing an agent for a negotiation task, since the range of the discourse is predefined in the dialogical framework and the rules which must be adhered to for the negotiation to complete satisfactorily are laid down in the performative structure.

### 3.6 Social acceptance

Citizens continuously demand more control over their medical information. They would like to have permanent access to it, such as via the Internet and digital TV, but they want to make sure that only those authorised to do so can have access to it, and only when appropriate. Individuals need to be confident that this information will not be disclosed, either on purpose or inadvertently, to third parties, such as government, employers, insurance companies, marketing companies, pharmaceutical companies, who may then use it for their own purposes. This feeling implies, aside from the security issues mentioned above, the need to build a relation of trust between citizens and agents that provide the access to sensitive information [54]. We consider that the AgentCities initiative may be a very important step in this direction; as it will promote the construction of platforms all around the world that will provide agent-based services, the human users of these systems will start to use them and gain confidence in the relegation of tasks to autonomous entities. Work on trust [70] provides an interesting analysis of relevant issues in this area.

### 3.7 Professional acceptance

In general, it has been observed that health care professionals are quite reluctant to accept and use new technologies. In the first place, they usually have a very busy schedule, so they lack the time to be aware of the latest advances in technologies and how they could be used to reduce their workload. They refuse to use new tools if they are not integrated smoothly into their daily workflow. They also often mention the lack of time and personnel to convert all the required medical data into an electronic format, so that it can be easily accessed and

managed - medical records are usually hand written and distributed in different departments of a medical centre. Some doctors also mention the hype built around Artificial Intelligence and, especially, expert systems, twenty years ago, which did not live up to their expectations, and they may reasonably argue that the intelligent autonomous agent paradigm, so fashionable today, may also fail to deliver real world results. Security in the access to data is also a common concern for health care professionals. Agent systems have the potential to introduce many innovations in the way in which computer systems respond to clinicians. For example, they could identify the user by behaviour or even voice, adapt the user interface to their way of working, and document diagnoses and therapies.

### 3.8 General issues

There are also many technical problems associated with the development of multi- agent systems in any domain (not especially related to health care). These include user expectations and acceptance, the lack of universally accepted standard agent communication languages, protocols and architectures (although FIPA seems to be clearly leading these efforts and their suggestions are starting to become a de facto standard), how to describe the services offered by agents, how to discover the presence of agents that provide a given service, how to guarantee the identity of an agent that is making a certain request, security, safety and trust issues, how to implement agents that offer wireless access to services (e.g. from mobile phones or PDAs, which will probably be the key towards a massive use of agent technology in the near future), how to handle properly the interactions between software agents and humans, and integration with pre-existing health-care systems.

## 4 Conclusion

In this position paper we want to argue that multi-agent systems have a set of characteristics that make them appropriate to be used to improve the provision of health care to citizens:

- They may be integrated with existing applications, for example agents may access a database to obtain the information about the patients of a certain hospital [13].
- The agents in a multi-agent system may be running in different locations, for example there may be an agent associated to each department of a medical centre [27], or an agent associated to each person that is included in a health care program in a certain community [22].
- The standards provided by FIPA [71] and the FIPA-compliant frameworks of multi-agent system development are reaching a level of maturity that make it feasible to think in world wide applications that coordinate the activities of health care in different countries, for example coordinating the management of organ transplant with a whole European perspective, following ideas similar to the ones suggested in [13].

- The autonomy of each agent in a multi-agent system permits to maintain the independent views of each modelled actor, for example each agency involved in the provision of health care to a community, such as social workers, health care professionals or emergency services may have different private policies that determine their relationship with other agents and their individual decisions [22].
- Information agents may help both citizens and health care professionals to obtain up-to-date and relevant health care information from Internet, for example see [23].
- Agents may help to address the growing demand of patient-centred management of medical data, for example it is feasible to think about the possibility of having personal agents which are able to get in touch with agents at a medical centre to receive information about their medical record or to make an appointment to be visited by a doctor [27].

## References

1. Fox, J., Das, S. Safe and sound: Artificial Intelligence in hazardous applications. AAAI Press/MIT Press, 1st edition, (2000)
2. Weiss, G. Multi-agent systems. A modern approach to Distributed Artificial Intelligence. M.I.T. Press, (1999)
3. Wooldridge, M. An introduction to Multiagent systems. John Wiley, Chichester, (2002)
4. Shankararaman, V. (ed.) Proceedings of the Workshop on Agents in Health Care, at the 4th International Conference on Autonomous Agents - AA 2000.. Barcelona, (2000)
5. Corts, U., Fox, F., Moreno, A. (eds.). Proceedings of the Workshop on Agent Applications in Health Care, at the 15th European Conference on Artificial Intelligence, ECAI 2002, Lyon, (2002). A revised and extended version of the best papers will appear in 2003 in a special issue of AI-Communications
6. Moreno, A. and Garbay, C. (eds.), Special issue of Artificial Intelligence in Medicine devoted to Software Agents in Health Care. (2003) (in press)
7. AgentCities Working Group on Health Care Applications. More information available at the web page <http://www.cms.brookes.ac.uk/hcwg>
8. Petrie, C. Agent-based software engineering. In: Agent-Oriented Software Engineering. Lecture Notes in Artificial Intelligence, Vol. 1957. Springer-Verlag, Berlin, (2001), 58-76
9. Jennings, N. On agent-based software engineering. Artificial Intelligence 117, (2000), 277-296
10. Decker, K., Li, J. Coordinated hospital patient scheduling. Proceedings of the 3rd International Conference on Multi-Agent Systems, ICMAS-98. Paris, France, (1998)
11. Kumar, A.D., Kumar, A.R., Kekre, S., Prietula, M.J., Ow, P.S. Multi-agent systems and organizational structure: the support of hospital patient scheduling. In: Proceedings of the Leading Edge in Production and Operations Management, South Carolina, USA, (1989)
12. Marinagi, C. et. al. Continual Planning and scheduling for managing patient tests in hospital laboratories, AIM 2000, (2000), 139-154.



13. Aldea, A., Lpez, B., Moreno, A., Riao, D., Valls, A. A Multi-Agent System for Organ Transplant Co-ordination. In: Quaglini, S., Barahona, P., Andreassen, S. (eds.): *Artificial Intelligence in Medicine. Lecture Notes in Computer Science*, Vol. 2101, Springer-Verlag, Berlin, (2001), 413-416.
14. Silverman, B.G., Andonyadis, C., Morales, A. Web-based health care agents; the case of reminders and todos, too (R2Do2), *Artificial Intelligence in Medicine*, 14(3), (1998), 295- 316
15. Shankararaman, V., Ambrosiadou, V., Robinson, B. Agents in Medical Informatics, *IASTED International Conference on Applied Informatics*, Austria, (2000)
16. Durfee, D.H. Distributed problem solving and planning. In [2], 121-164.
17. Klusch, M. Information agent technology for the Internet: a survey. *Data and Knowledge Engineering*, Vol. 36 (3), (2001), 337-372
18. Hannebauer, M., Mller, S. Distributed Constraint Optimization for Medical Appointment Scheduling, *5th International Conference on Autonomous Agents (AGENTS 2001)*, Montreal, Canada, (2001)
19. Vzquez-Salceda, J., Corts, U. Using Agent-Mediated Institutions for the distribution of Human Tissues among Hospitals. *Advanced Course on Artificial Intelligence (ACAI-01)*, Prague, Czech Republic, (2001), 205-209
20. Vzquez-Salceda, J., Padget, J.A., Corts, U., Lpez-Navidad, A., Caballero, F. Formalizing an electronic institution for the distribution of human tissues. To appear in [6]. March 2003 (in press)
21. Moreno, A., Valls, A., Bocio, J. Management of Hospital Teams for Organ Transplants Using Multi-Agent Systems. In: Quaglini, S., Barahona, P., Andreassen, S. (eds.): *Artificial Intelligence in Medicine. Lecture Notes in Computer Science*, Vol. 2101, Springer-Verlag, Berlin, (2001), 413-416
22. Beer, M.D., Huang, W., Sixsmith, A. Using agents to build a practical implementation of the INCA-Intelligent Community Alarm- system. In: Jain, L.C., Chen, Z. and Ichalkaranje, N. (eds.): *Intelligent Agents and their applications. Studies in Fuzziness and Soft Computing*, Physica-Verlag, Berlin, (2002), 320-345
23. Baujard, O., Baujard, V., Aurel, S., Boyer, C., Appel, R.D. MARVIN, a multi-agent softbot to retrieve multilingual medical information on the Web. *Medical Informatics* 23 (3), Taylor and Francis, London, (1998), 187-191
24. Kostkova, P., Mani-Saada, J., Weinberg, J. Agent-based up-to-date data management in the National Electronic Library for Communicable Disease. In [5], 59-63.
25. Lobato, E., Shankararaman, V. PIRA: A Personalised Information Retrieval Agent, *IASTED International Conference on Artificial, Intelligence and Soft Computing*, Honolulu, (1999)
26. Marshall, P., Greenwood, S. The Use of Emergent Behaviour in a Multi-Agent System to Drive Self-Adaptation at the Interface, *Joint Web Intelligence/Intelligent Agent Technology (2001 WI/IAT)*, Japan, (2001)
27. Moreno, A., Isern, D. Accessing distributed health-care services through smart agents. *Proceedings of the 4th IEEE International Workshop on Enterprise Networking and Computing in the Health Care Industry (HealthCom 2002)*, Nancy, France, (2002), 34-41
28. Barro, S., Presedo, J., Castro, D., Fernndez Delgado, M., Fraga, S., Lama, M., Vila, J. Intelligent telemonitoring of critical-care patients. *IEEE Engineering in Medicine and Biology Magazine*, 18 (4), (1999), 80-88
29. Lanzola, G., Gatti, L., Falasconi, S., Stefanelli, M. A Framework for Building Co-operative Software Agents in Medical Applications. *Artificial Intelligence in Medicine*, 16(3), (1999), 223-249



30. Larsson, J.E., Hayes-Roth, B. Guardian: An Intelligent Autonomous Agent for Medical Monitoring and Diagnosis. *IEEE Intelligent Systems*, (1998), 58-64
31. Farias, A., Arvanitis, T.N. Building Software Agents for Training Systems: A Case Study on Radiotherapy Treatment Planning. *Knowledge-Based Systems*, (10), (1997), 161-168
32. Quaglini, S., Stefanelli, M., Cavallini, A., Micieli, G., Fassino, C., Mossa, C. Guidance-based Careflow Systems, *Artificial Intelligence in Medicine*, 20(1), (2000), 5-22
33. Marchetti, D., Lanzola, G., Stefanelli, M. An AI-based Approach to Support Communication in Health care Organizations. In: *Proc. 8th Conference on Artificial Intelligence in Medicine in Europe - AIME 2001*, Cascais, Portugal, (2001)
34. Huhns, M.N., Singh, M.P. Managing Heterogeneous Transaction Workflows with Co-operating Agents. In: Jennings, N., Wooldridge, M. (eds): *Agent Technology: Foundations, Applications and Markets*. Springer-Verlag, Berlin, (1998)
35. Alamillo, T., Alsinet, T., Bjar, R., Anstegui, C, Fernandez, C., Many, F. Automated monitoring of medical protocols: a secure and distributed architecture. To appear in [6], March 2003 (in press)
36. Godo, L., Puyol-Gruart, J., Sabater, J., Torra, V. A multi-agent system approach for monitoring the prescription of restricted use antibiotics. To appear in [6], March 2003 (in press)
37. Corts, U. et al. Assistive technologies for the disabled and for the new generation of senior citizens: the e-tools architecture. In [5], 18-22
38. Camarinha-Matos, L.M., Afsarmanesh, H. Virtual Communities and Elderly Support, *Advances in Automation, Multimedia and Video Systems, and Modern Computer Science*, V.V. Kluev, C.E. DAttellis, N. E. Mastorakis (eds.), WSES, (2001), 279-284
39. CEN/TC251 (European Committee for Standardization/TC251), <http://www.centc251.org>
40. ASTM E31: <http://www.astm.org/COMMIT/COMMITTEE/E31.htm>
41. HL7 (Health Level 7), <http://hl7.org>
42. OpenGALEN: <http://www.opengalen.org>
43. UMLS, <http://www.nlm.nih.gov/research/umls>
44. SNOMED, <http://www.snomed.org>
45. Agentcities Working Group on HealthCare meeting, Agentcities ID2, Lisbon, September (2002). More information available at the web page <http://www.cms.brookes.ac.uk/hcwg>
46. Cyc Knowledge Base, <http://www.cyc.org>
47. A. Madche, S. Staab. Mining Ontologies from Text. In: Dieng, R., Corby, O. *EKAU-2000 - 12th International Conference on Knowledge Engineering and Knowledge Management*, Juan-les-Pins, France. Lecture Notes in AI. Springer-Verlag, Berlin (2000)
48. A. Madche, S. Staab. Ontology Learning for the Semantic Web. *IEEE Intelligent Systems*. 16(2), (2001), 72-79
49. Kietz, J-U., Maedche, A., Volz, R. A Method for Semi-Automatic Ontology Acquisition from a Corporate Intranet, In: *Proc. of the EKAU'2000 Workshop Ontologies and Texts*, (2000)
50. Faure D., Nedellec, C. Knowledge Acquisition of Predicate-Argument Structures from technical Texts using Machine Learning. In: *Proceedings of Current Developments in Knowledge Acquisition*. In: Fensel D., Studer, R. (eds.): *EKAU-99*. Springer-Verlag, Karlsruhe, (1999), 329-334

51. de Moor, G. Privacy Issues and Grid Environments, HealthGRID 2003, Lyon (2003)
52. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.
53. US Department of Health and Human Services. Standards for Privacy and Individually Identifiable Health Information. Federal Register, v.65 (2000), 82461-82510.
54. Wong, H., Sycara, K. Adding security and trust to multi-agent systems. In: Proceedings of Workshop on deception, fraud and trust in agent societies, Autonomous Agents, Seattle, (1999), 149-161
55. Bellifemine, F., Poggi, A., Rimassa, G. Developing multi-agent systems with a FIPA compliant agent framework, Software Practice and Experience, 31, (2001), 103-128
56. Bellifemine, F., Poggi, A., Rimassa, G. JADE- A FIPA compliant agent framework, Proceedings of Practical Applications of Intelligent Agents and Multi-Agents, London, (1999), 97-108
57. Quaglioni, S., Stefanelli, M., Cavallini, A., Micieli, G., Fassino, C., Mossa, C. Guidance- based Careflow Systems, Artificial Intelligence in Medicine, 20(1), (2000), 5-22
58. A. Poggi, G. Rimassa and M. Tomaiuolo, Multi-user and security support for multi-agent systems. In: Proceedings of WOA 2001 (Dagli oggetti agli agenti), Omicini, A., Viroli, M. (eds.), (2001), 20-26
59. Rao, A. S. Decision Procedures for Propositional Linear-Time Belief-Desire-Intention Logic. In Proc. ATAL'95. Lecture Notes in AI, Vol. 1037, Springer-Verlag, Berlin, (1996), 102-118
60. Rao, A. S., Georgeff, M. P., Modeling Rational Agents within a BDI-Architecture. In: Proc. KR'91. Morgan Kaufmann (1991) 473-484
61. Linder, B. van, Hoek, W. van der and Meyer, J-J.Ch. Formalising Abilities and Opportunities of Agents. In: Meyer, J-J.Ch. and Treur, J. (ed.): Agent-Based Defeasible Control in Dynamic Environments. Handbook of Defeasible Reasoning and Uncertainty Management Systems. Kluwer, Dordrecht, Boston, London, (2002) 253-307
62. d'Inverno, M., Fisher, M., Lomuscio, A., Luck, M., de Rijke, M., Ryan, M. and Wooldridge, M. Formalisms for Multi-Agent Systems. In: The Knowledge Engineering Review (1997) 12, 3
63. Di Marzo Serugendo, G. A Formal Development and Validation Methodology Applied to Agent-Based Systems. Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems. (eds.) Wagner, T., Rana, O., Lecture Notes in AI, Vol. 1887, Springer-Verlag, Berlin, (2001), 214-226
64. Ndumu, D. T. and Nwana, H. S., Research and Development Challenges for Agent-Based Systems. IEE Proc. Softw. Eng. 144 (1), (1997)
65. Halpern, J. Y. and Moses, Y., A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief. Artif. Intell. Vol. 54, (1992), 319-379
66. Wooldridge, M., Dixon, C., Fisher, M.A. Tableau-Based Proof Method for Temporal Logics of Knowledge and Belief. J. Appl. Non-Classical Logics 8 (3) (1998) 225-258
67. Dixon, C., Fisher, M. and Wooldridge, M. Resolution for Temporal Logics of Knowledge. J. Logic Computat. 8 (3), (1998), 345-372

68. Esteva, M., Rodriguez, J.A., Sierra, C., Garcia P., Arcos, J.L. On the formal specifications of electronic institutions, Agent-mediated electronic commerce. The European AgentLink Perspective, Lecture Notes in AI, Vol. 1991, Dignum, F., Sierra, C. (eds.), (2001), 126- 147
69. Vzquez-Salceda. J., Corts, U. Using Agent-Mediated Institutions for the distribution of Human Tissues among Hospitals, Advanced Course on Artificial Intelligence, Praga (2001) 205-20
70. Castelfranchi, C., Falcone, R. Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In: Proceedings of the International Conference on Multi Agent Systems, Los Alamitos, (1998), 72-79
71. FIPA: Foundation for Intelligent Physical Agents, <http://www.fipa.org>