We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

6,900
Open access books available

186,000
International authors and editors

200M
Downloads

Our authors are among the

154
Countries delivered to

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

CLARIVATE ANALYTICS
BOOK CITATION INDEX
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Introductory Chapter: Multi-Agent Systems

Jorge Rocha, Inês Boavida-Portugal and
Eduardo Gomes

Additional information is available at the end of the chapter

1. Introduction

Agents, or more precisely intelligent agents, are a novel paradigm for software applications development, supporting the simulation of complex individual interactions. Moreover, agent-based computing has been welcomed as the latest paradigm shift in software development as well as the new software revolution. Presently, agents are one of the main fields of interest in computer science, artificial intelligence (AI), and complex system theory.

Intelligent (or rational) agents are used in a wide multiplicity of applications, ranging from relatively small systems, e.g. email filters, to large complex systems, such as air traffic control, bird flocking, or human social behaviour. Apparently, it may look that such particularly unlike types of system cannot have much in common. Still, nothing can be more deceitful, as in both the key concept is the agent.

Before addressing the issue of agent-based systems (ABS) development, one should try to define what terms 'agent' and ABS mean. Regrettably, there is a lack of commonly accepted definitions about key concepts in agent-based computing. Actually, there is no genuine agreement on the definition of the term agent.

As Russell and Norvig [1] stated an agent is anything that can be considered able to perceive its environment through sensors and act on this environment through actuators. To Macal [2], an agent shall have the following characteristics: (i) be identifiable—a discrete individual with a set of features and rules (mathematical or logic) that govern behaviour and decision-making capacity; (ii) be located—settled in an environment with which it interacts and also in which interacts with other agents; (iii) be goal-driven; (iv) be self-contained; and (v) be flexible, and have the ability to learn and adapt its behaviour through time-based experiences.

Most authors agree that although there are multiple definitions of the term 'agent', several attributes can be pointed out such as heterogeneity, autonomy, capacity to process and

exchange information, follow if-then rules, goal-driven, and deductive code-based units, with boundary and state. There are also core behaviours: mobility, interaction, adaptation, and bounded rationality [3].

Agents are not aggregated into homogeneous pools. Rather, agents are heterogeneous with different attributes, behaviours and rules, which may differ in multiple ways (e.g. social network, individual preferences), and over time. For instance, one can model groups of residents in a specific neighbourhood but these individuals can have heterogeneous characteristics such as age, gender, income, and living preferences, although associating with the same group.

Agents are autonomous entities that are not subjected to the influence of external direction. They are developed over a bottom-up approach and have the capacity for processing information, while sharing it with other agents, through individual-based interaction that does not suffer top-down control. However, when a new agent enters the simulation, its actions can be conditioned by pre-existing norms that have been instituted through earlier agent interactions and persisted through time steps. These interactions can be expressed by the interchange of data from one agent to another. In such a way, micro and macro level will typically coevolve without pre-defined upper level controllers (i.e., bottom-up approach).

An agent may be goal-driven and takes independent actions to reach its goals. Thus, agents compare behaviour outcomes to its goals and adapt responses in the future. An agent's behaviour can be described by simple if-then rules that used to describe the theoretical assumptions of agent behaviour, in the form of computational procedures that lead to goal achievement. These procedures constitute a plan for achieving agents' objectives.

ABS is one where an agent is used as key abstraction. Theoretically, an ABS could be conceptualized in terms of agents and still be implemented without using any software consistent to agents. There is an obvious parallelism with object-oriented software, where it is fully conceivable to design a system based on objects, and implement it without getting use of object-oriented software. Nevertheless, this would be counterproductive or, at least, unusual. The same happens with ABS, where users expect agents to be designed and implemented using agent paradigm (e.g., using specific agent based software).

One should note that an ABS may have any non-zero amount of agents. A multi-agent system (MAS), designed and implemented by means of several interacting agents, is more general and pointedly more complex than the unitary (single case) agent. In real world, there are various number of situations where the single-agent case is suitable. A good example is the expert assistant, where an agent acts like an expert assistant to a user attempting to fulfil some task on a computer.

MAS is a computer-based environment made of multiple interacting intelligent agents. MAS are preferably used in solving problems that are difficult (or impossible) for an individual agent. As with agents, there is no categorical definition of MAS so let us focus on one that is relatively consensual. In Stone and Veloso [4], MAS is defined as 'a loosely coupled network of problem-solving entities (agents) that work together to find answers to problems that are beyond the individual capabilities or knowledge of each entity (agent)'.

As in MAS, agent-based model (ABM) also consists of interacting agents within a specific environment. ABM is known by different names due to its wide variety of applications, which could refer to entirely diverse methodologies. It can also be called a multi-agent system (MAS) or agent-based system (ABS).

In a computer science (e.g. AI), ABM usually states a computer-based method for studying the (inter)actions of a set of autonomous entities. In non-computing–related scientific domains as in social sciences, ABM could refer to an actor in the social world and be called agent-based social simulation (ABSS). Davidsson [5], using different combinations of focus areas (e.g. agent-based computing, computer simulation, and social sciences), further subdivides ABSS into three categories: (i) social aspects of agent systems (SAAS); (ii) multi-agent based simulation (MABS); and (iii) social simulation (SocSim).

In other domains (e.g. transportation ecological science or life science), ABM mostly refer to an individual-based model or a self-sufficient computing method. Although ABM is a wide ranging paradigm applied in totally different manners in all types of scientific domains, eventually all its subtle differences meet together under the domain of agent-based computing [6].

Even though there is a significant overlay, MAS not necessarily means the same as ABM. The objective of an ABM is to search for descriptive insights into the agents' (not necessarily intelligent) collective behaviour following simple rules (typical of natural systems) rather than solving particular engineering problems. ABM is more often used in the sciences, whereas MAS is frequently applied in engineering- and technology-related issues [6]. Hereafter the designation MAS will be use as a general term covering all agent related semantics discussed in previous paragraphs.

2. MAS characteristics

An agent could refer to different components as applications have different objectives and lay down in different paradigms [7]. One can see an agent as being part of a program (e.g. model, system, or subsystem) or any type of independent entity (e.g. organization or individuals). Each agent is programmed to react to other agents and to its computational environment, with respect to behaviour rules from primitive reaction decisions to complex adaptive AI. However, one may believe that the majority of researchers should be in general agreement with Wooldridge and Jennings [8] who defined an agent as a piece of hardware or a software-based computer system that entail the following properties:

- Reactivity, in the sense agents have the perception of their environment and respond quickly to changes that may occur.

- Pro-activity, not being limited to acting in response to the environment, agents are able to take the initiative and show behaviour driven by objectives.

- Social skills. The agents are able to interact/communicate (cognitive model) with other agents (and possibly humans) through a given Agent Communication Language (ACL) and establishing connections between their autonomous objectives and the spatial context.

These properties are somewhat difficult to identify than it may seem at first sight. Autonomy, although consensual in the agents' community and essential in Wooldridge and Jennings agent definition, can never be fully achieved. It is clear that the agent has to be created and put into operation by a human (or another agent). The assumption that the agent action will not have an end is also not entirely valid. Of course, under the aegis of current science, the agent will have a limited lifetime and a final action. On the other hand, although autonomy (i.e., actions are carried out without human interference) is essential to the agent, usually human-agent interaction is desirable or even essential. It is usual to build agents that behave autonomously but are also able to take orders or instructions from humans.

To build purely reactive agents can be a simple task, but is not entirely desirable. A purely reactive agent would react to changes in environment consecutively without seeking to achieve its medium or long-term objectives, i.e. display goal-oriented behaviour. One should define agents capable of balancing reactive with proactive behaviour. Nonetheless, the difficulty in balancing these two types of behaviour is very high [9].

The pro-activeness is simple to get in functional systems. Still, this simplicity only applies if a static environment is considered, i.e. it does not change during the accomplishment of a given procedure or function. In addition, the agent should have all the information it needs to run that procedure or function, without any uncertainty in the environment. However, these assumptions are not valid for most environments. For dynamic and not fully accessible environments, the agents must be able to react to changes in the environment and reason if the original objectives are still valid, due to the changes in the environment while performing a given procedure. This means that the agents have to be reactive and therefore able to quickly adapt to phase-shifts in the environment.

The social capacity of an agent is related to its ability to exchange high-level messages (and not only data-bytes without an associated meaning) and carry out processes of social interaction with other agents (and/or humans) similar to those used by humans in their daily lives, establishing collective behaviours. These processes include the coordination, cooperation, and negotiation. In order to conduct them, it is necessary to reason about the objectives of the other agents (if any) present in the environment or, at least, have notion of their existence.

It is also necessary to understand that they are also autonomous agents and do not necessarily share common goals. In this way, it may be necessary to negotiate and cooperate with other agents, eventually exchanging information and/or goods. For example, in order to convince an agent to cooperate, it might be necessary to make a payment or offer a particular good or service. In several cases, agents have opposite objectives and, therefore, are not able to carry out any cooperative process that includes them.

The equilibrium between social capacity and proactive or reactive capabilities is also of great importance. This importance is even greater in a scheme of cooperative work set by a group of agents who share a common goal. In these situations, each agent has to adapt his reaction to the events that occur in the modeling environment, both with the free will needed to perform individual tasks and with the social behaviors necessary to perform collective tasks [9].

Some researchers highlight other aspects of agency (e.g. mobility and adaptability). Indeed, agents may have supplementary features, and in specific uses, some features can be more significant than others. Yet, is the conjugation of the three main properties (reactivity, pro-activity and social skills) in a single entity that gives importance to the agent paradigm and makes the difference between agent systems and related software paradigms, e.g. object-oriented systems, distributed systems, and expert systems [10].

Franklin and Graesser [7] discuss about various definitions of agents and list some behaviours displayed by them: (i) reactive; (ii) autonomous; (iii) guided by objective(s)/pro-active; (iv) temporally continuous; (v) social/communicative; (vi) have ability to learn/adapt; (vii) mobile; (viii) flexible; and (ix) have personality. There are still other agents classification schemes based, for example, on the type of task they are running or in their architecture.

The very definition of agent satisfies only the first four identified features, as Franklin and Graesser [7] state 'an agent is a system located in an environment and is part of it. An agent understands the environment and acts on it over time. An agent has its own agenda in order to reflect its perceptions of future'. But even so, this definition is generic enough to cover from a thermostat, containing one or two sensors and extremely simple control structure, to human beings with multiple and conflicting guidelines, various sensors, various possibilities of actions, and structures of extremely complex and sophisticated control.

The concept of agent is related to rationality. According to Russel and Norvig [1], rationality is associated with four factors: (i) the performance measure that defines success criteria; (ii) the prior knowledge of the agent about the environment; (iii) the actions that the agent is capable to perform; and (iv) the sequence of agent's perceptions. These lead to the definition that 'for each possible perceptions, a rational agent must select an action that is expected to maximize its performance measure, given the evidence provided by the sequence of perceptions and any prior internal knowledge of the agent' [1].

3. MAS classification

Due to the wide range of applications, the difficulty in defining what is truly an 'intelligent agent' and to the enormous momentum that this area of research has had over the past few years, there are several synonyms of the term 'intelligent agent' created by different researchers in an attempt to better characterize their own work. Thus, it is usually found in the specialized agent literature designations such as robots, software agents (or softbots), knowbots, taskbots or userbots, personal assistants, virtual characters, and so on. Although the existence of these synonyms is perfectly understandable, sometimes obscures the concept itself making it harder to define the object in MAS research.

In order to better characterize the agents' scientific area, it is useful to divide agents into classes analysing the different typologies of agents proposed in the literature. The high number of attributes previously discussed allows to realize how difficult it is to implement an agent that incorporates all those attributes. This is also related to the fact that the characteristics of an

agent are ideally application-type dependents. The analysis of agents' attributes has been used by researchers to sort and categorize them in types. A typology is a classification by types of agents that have attributes in common.

Nwana [11] proposes a typology of agents, identifying seven distinct classification dimensions:

1. **Mobility**. Static or mobile agents. Mobile agents can be resident in the source machine or temporarily in another one.

2. **Reasoning model**. Presence or not of a type of symbolic reasoning, i.e. an agent may be reactive or purely deliberative.

3. **Agent function**. The main function assumed by the agent, such as an information search agent (looking for information for a given user on the Internet) or interface (which facilitates the interaction man-machine of a given application).

4. **Autonomy**. Agents operate without any direct human or other agents intervention, they have some kind of control over their actions and internal state, and they are able to exchange information with other agents. Agents are not subjected to the influence of external direction.

5. **Cooperation**. Realization of cooperative actions with other agents.

6. **Learning**. Inclusion or not of learning capabilities in the agent (individual, evolutionary, and social).

7. **Hybrid features**. These combine two or more different behaviour philosophies in the same agent.

Combining the characteristics of autonomy, cooperation, and learning, Nwana [11] derives four types of agents: (i) collaborative agents; (ii) collaborative agents with learning (memory) ability; (iii) interface agents; and (iv) truly intelligent agents (**Figure 1**). It is important to note that the bounds of this classification should not be interpreted as a strict and well-defined fact.

After establishing a typology, Nwana [11] defined seven categories of agents according to their architecture and function: (i) collaborative agents; (ii) interface agents; (iii) mobile agents, (iv) information agents, (v) reactive agents, (vi) hybrid agents, and (vii) intelligent agents. Franklin and Graesser [7] believe that an agent, by definition, must be a continuous process execution and presented the taxonomy represented diagrammatically in **Figure 2**, which divides the autonomous agents into three main groups: biological, robotic, and computational agents.

Agent-based applications can be classified through many orthogonal dimensions. They can be classified by the type of the agent, by the technology used to implement the agent, or by the application domain itself. We will focus on the latest one since it is the one that fits best the objectives and structure of this book. The aim of this classification scheme is simply to give a visual understanding of the scale and variety of agent applications.

These include, among others, information research, personal assistants, e-mail management, control of electrical energy, telecommunications network management, traffic management,

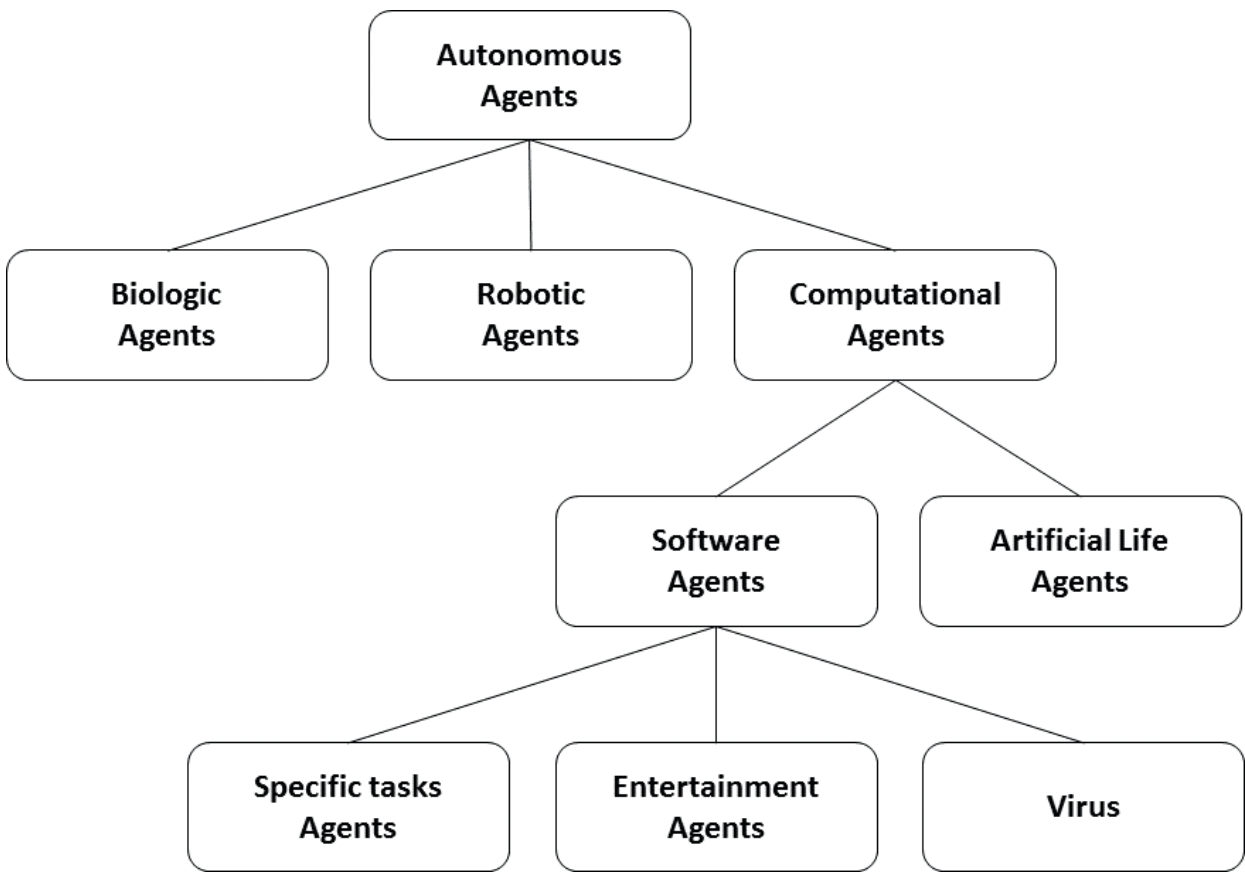**Figure 1.** Agents' categories defined by Nwana [11].

Figure 2. Agents' categories defined by Franklin and Graesser [7].

underwater exploration, control of vehicles and spacecraft, computer-integrated manufacturing, air traffic management, transport management, trading, e-commerce, financial transactions and exchanges, training of teams, games, entertainment, and virtual characters. Wooldridge and Jennings [8] draw one of the first attempts to classify agents in such a way (**Table 1**).

Actual MAS are dynamically being applied in various fields of knowledge. Practical cases of uses include [12] the modelling of organizational behaviour, team working, supply chain, consumer behaviour, social networks, distributed computing, transportation, and environmental studies. MAS have also been applied [13] to several social and society fields, comprising population dynamics, epidemics outbreaks, biological applications, civilization development, and military applications.

Macal and North [3] categorized all of these MAS applications into two types:

1. **Minimalist models**—Based on a set of idealized assumptions, it is designed to capture only the most salient features of a system. These are exploratory electronic laboratories, involving resources at computer modelling level, in which a wide range of assumptions can be varied over a large number of experimental simulations.

2. **Decision support systems** (**DSS**)—Tend to be large-scale applications, it is designed to answer a broad range of real-world policy questions, making efforts to support stakeholders in their decision-making activities. These models are distinguished by including real data and having passed some degree of validation to establish credibility in their results.

Table 2 summarizes a list of MAS applications drawn by Macal and North [3].

Industrial applications	Commercial applications
• Process control	• Information management
• Manufacturing	• Electronic commerce
• Traffic and transportation systems	• Business process management
Medical applications	**Entertainment**
• Patient monitoring	• Games
• Health care	• Interactive theatre and cinema

Table 1. MAS applications according to Wooldridge and Jennings [8].

Business and organizations	Society and culture
• Manufacturing operations	• Ancient civilizations
• Supply chains	• Civil disobedience
• Consumer markets	• Social determinants of terrorism
• Insurance industry	• Organizational networks
Economics	**Military**
• Artificial financial markets	• Command and control
• Trade networks	• Force-on-force
Infrastructure	**Biology**
• Electric power markets	• Population dynamics
• Transportation	• Ecological networks
• Hydrogen infrastructure	• Animal group behaviour
	• Cell behaviour and subcellular processes
Crowds	
• Pedestrian movement	
• Evacuation modelling	

Table 2. MAS applications according to Macal and North [3].

4. Conclusions

It is our strong belief that the concept of an agent as an independent rational decision maker has great value, not only in AI but also for conventional computer science. Likewise, the recent developments in MAS research, enabling agents to cooperate and negotiate, will surely be of vital importance in the future. The cognition and knowledge of MAS and the recognition of its applications continue to expand in line with its quick advances.

Yet, one can say that MAS are an overvalued technology of the last decades. Indeed, the current interest in agents carries with it the drawback of trying to label everything as an agent. Despite this fact, the technology has much to offer and it is imperative not to oversell it. MAS

are likely to be most useful for a specific class of applications, which exhibit the kind of properties listed in Section 2. MAS has been broadly applied in a range of disciplines that include, but not limited to biology, ecology, computer simulation, business, economic science, policy, social sciences, political science, and military studies.

This book reviews several MAS applications and depicts the concept of MAS as scoped in the literature. Those applications fall mainly into two operational fields: individual-based models that study personal transportation-related activities and behaviour and system and computational methods to study collaborative and reactive autonomous decision making.

The chapters in this book cover a wide spectrum of issues related to the applications of intelligent agents and MAS. The introductory chapter explains the basic concepts underlying MAS, followed by experience chapters that deal with specific MAS and empirical applications. Application MAS domains include: collision avoidance, automotive applications, evacuation simulation, emergence analyses, cooperative control, context awareness, data (image) mining, resilience enhancement, and the management of a single-user multi-robot. The range of author affiliations covers a significant proportion of the organizations that are currently working in MAS and powerfully proves the universal nature of this relatively new field of research.

Author details

Jorge Rocha[1]*, Inês Boavida-Portugal[2] and Eduardo Gomes[3]

*Address all correspondence to: jorge.rocha@campus.ul.pt

1 Institute of Geography and Spatial Planning, Universidade de Lisboa, Portugal

2 Department of Spatial Planning and Environment, University of Groningen, The Netherlands

3 Géographie-cités, UMR 8504, Université Paris 1 Panthéon-Sorbonne, France

References

[1] Russel S, Norvig P. Artificial Intelligence—A Modern Approach. 3rd ed. England: Pearson Education Limited; 2014. p. 1079

[2] Macal CM. Everything you need to know about agent-based modelling and simulation. Journal of Simulation. 2016;**10**(2):144-156

[3] Macal CM, North MJ. Tutorial on agent-based modelling and simulation. Journal of Simulation. 2010;**4**(3):151-162. DOI: 10.1057/jos.2010.3

[4] Stone P, Veloso M. Multiagent systems: A survey from a machine learning perspective. Autonomous Robots. 2000;**8**(3):345-383

[5] Davidsson P. Agent-based social simulation: A computer science view. Journal of Artificial Societies and Social Simulation. 2002;**5**(1):1-7

[6] Niazi M, Hussain A. Agent-based computing from multi-agent systems to agent-based models: A visual survey. Scientometrics. 2011;**89**(2):479-499

[7] Franklin S, Graesser A. Is it an agent, or just a program?: A taxonomy for autonomous agents In: Intelligent Agents III. Agent Theories, Architectures, and Languages, Proceedings of ECAI'96 Workshop (ATAL), Budapest, Hungary, August 12-13, 1996, Müller, J., Wooldridge, M.J., Jennings, N.R. (Eds.), Lecture Notes in Computer Science, Springer, 1997;**1193**:21-35

[8] Wooldridge M, Jennings NR. Intelligent agents: theory and practice. The Knowledge Engineering Review. 1995;**10**(2):115-152

[9] Reis L, Paulo L, Nuno O, Eugénio C. Situation based strategic positioning for coordinating a team of homogeneous agents. In: Hannenbauer M, Wendler J, Pagello E, editors. Balancing Reactivity and Social Deliberation in Multi-Agent System—From RoboCup to Real-World Applications. Lecture Notes in Artificial Intelligence, Berlin: Springer; 2001;**2013**:175-197

[10] Wooldridge M. An Introduction to Multi-Agent Systems. 2nd ed. England: John Wiley & Sons; 2009. p. 453

[11] Nwana H. Software agents: An overview. Knowledge Engineering Review. 1996;**11**(3):1-40

[12] Hughes HPN, Clegg CW, Robinson MA, Crowder RM. Agent-based modelling and simulation: The potential contribution to organizational psychology. Journal of Occupational and Organizational Psychology. 2012;**85**(3):487-502

[13] Aschwanden G, Wullschleger T, Muller H, Schmitt G. Agent-based evaluation of dynamic city models: A combination of human decision processes and an emission model for transportation based on acceleration and instantaneous speed. Automation in Construction. 2012;**22**:81-89