Intelligent Agents and Environment

Alfredo Garro[1], Alberto Falcone

Department of Informatics, Modeling, Electronics and System Engineering, University of Calabria - via P. Bucci, cubo 41/C, 87036 Rende (Italy)

Matteo Baldoni, Cristina Baroglio

Department of Computer Science, University of Turin - Corso Svizzera 185, 10149 Turin (Italy)

Federico Bergenti

Department of Mathematical, Physical and Computer Sciences, University of Parma - Parco Area delle Scienze 53/A, 43124 Parma (Italy)

Stefano Mariani

Department of Sciences and Methods for Engineering - University of Modena e Reggio Emilia - Via Amendola 2, 42122 Reggio Emilia (Italy)

Andrea Omicini

Department of Computer Science and Engineering, Alma Mater Studiorum - University of Bologna - Via Zamboni 33, 40126 Bologna (Italy)

Giuseppe Vizzari

Department of Computer Sciences, Systems and Communications, University of Milano-Bicocca - Piazza dell'Ateneo Nuovo 1, 20126 Milan (Italy)

Abstract

This chapter aims to provide a panorama on the fundamentals related to Agents. Specifically, the most popular definitions behind the concept of Intelligent Agents and the main properties that characterize an agent are discussed along with different agent-based models. Furthermore, the concepts of Environment as well as its properties and the role that it plays in the context of the agent paradigm is presented. Actions and Interactions among agents and environment are also

[1]Corresponding author

Preprint submitted to Encyclopedia of Bioinformatics and Computational Biology *February 9, 2024*

discussed before contextualizing the role of Intelligent Agents in the Computational Biology domain.

Keywords: Intelligent Agent, Environment, Biological Systems, Natural Systems, Artificial Systems, Autonomy, Cooperation

1. Introduction

The analysis of complex systems, i.e. systems consisting of several interdependent and interacting entities that determine the system behavior, requires the exploitation of new and more effective solutions able to face with aspects ranging from the definition of appropriate modeling formalisms to the use of advanced system analysis methods. Such kind of systems can be classified in two big categories: Artificial Systems (ASs) that are built by humans and Natural Systems (NSs), already existing in nature without the intervention of humans. In particular,

- examples of complex Artificial Systems are Cyber-Physical Systems (CPSs) (Chen and Lu, 2018) or System of Systems (SoSs) (Garro and Tundis, 2015). Here different components belonging to various application domains (such as Software, Mechanical, Electrical and Electromechanical and so on), which are natively conceived for working in isolation for a specific purpose, are integrated in a common environment in order to achieve one or more complex goals, e.g. Sampigethaya and Poovendran (2012); Zhabelova and Vyatkin (2012).

- examples of Natural Systems are represented by Biological Systems (BSs) (Anderson, 2015), in which a complex network of biological entities belonging to different biological subsystems (such as nervous system, circulatory system, respiratory systems and so on) work together in a synergistic manner.

Concerning the second example, Intelligent Agents, which can be defined as entities situated in an environment, able to act upon it and interact with each other, to achieve specific goals (Russell and Norvig, 2003), represent a very promising solution in computational biology. They allow the development of complex applications centered on theoretical methods, for supporting data-analysis based on mathematical modeling and computational simulation techniques for observing social systems and studying biological phenomena on the basis of the so-called emergent behavior (Seekhao et al., 2016; Adamatti, 2016). Indeed, typically, in these systems it is not enough to observe and analyze the state and output of the single system or individual entity, but it is necessary to observe the way they interact and cooperate, in order to capture particular dynamics resulting from their interactions, which define their emergent behaviors.

It is clear that for such class of systems, Intelligent Agents provide a suitable solving approach, thanks to their key features of autonomy and cooperation. A key role, in the adoption of Intelligent Agents, is played by the environment that represents the "problem space" in which the agents operate and in which the agents represent one possible resolution path. It is important to note that the environment can be partially or fully observable. The observability lets the agent retrieve information and to compute them, in order to take decisions and consequent actions on the basis of the perceived information. Moreover, the environment can be deterministic or stochastic. As a consequence, if the environment is deterministic, its properties are well known and none of them is random, that means that the output of the model is fully determined by the value of its parameter a by its initial conditions. If the environment is stochastic, then, some randomness and uncertainty is present in it.

Specific details about Intelligent Agents, Environment and Interactions are provided in the rest of the Chapter.

2. Agents and Intelligent Agents

Many but similar definitions of Agents have been provided in the literature. A weaker but more general definition of Agents, that could be suited to describe the highly heterogeneous approaches in the agent-based computing context, is to see an Agent as an autonomous entity, having the ability to decide the actions to be carried out in the environment and interactions to be established with other agents, according to its perceptions and internal states.

In artificial intelligence, an Intelligent Agent (IA) is an autonomous entity that observes through sensors and acts upon an environment using actuators (i.e., it is an agent) and directs its activity towards achieving goals (i.e., it is "rational", as defined in economics). Intelligent Agents may also learn or use knowledge to achieve their goals. They may be very simple or very complex: a reflex machine such as a thermostat is an Intelligent Agent.

The most popular definitions of IA are provided (i) by Smith et al. (1994), who state that "An agent is a persistent software entity dedicated to a specific purpose.", (ii) by Hayes-Roth (1995) who says that "Intelligent Agents continuously perform three functions: perception of dynamic conditions in the environment, action to affect conditions in the environment, and reasoning to interpret perceptions, solve problems, draw inferences, and determine actions"; (iii) by IBM, "Intelligent Agents are software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in so doing, employ some knowledge or representation of the user's goals or desires.".

According to the vision described in Russell and Norvig (2003), an IA can be seen as an entity in a program or environment capable of generating action.

From a more technical perspective, an IA is an entity in a program or environment capable of generating actions (Magedanz, 1996). It uses perception of the status of the environment in order to make decisions about specific actions to take. The perception is represented by the capability or sensitiveness, which is typically achieved by sensors, whereas actions are the reaction to a particular status of phenomena that can depend on the most recent perception or on the entire history (sequence of perceptions). An IA uses and provides functions. A function can be a mathematical function that maps a sequence of perceptions into one or more actions, which is implemented as an agent program. The part of the agent that is in charge of taking an action is called an actuator.

Other important characteristics of an IA are the following one:

- *rationality*: an IA is supposed to act in order to achieve its goals and does not act in such a way as to prevent its goals being achieved, at least insofar as its beliefs permit. A rational agent is one that can take the right decision in every situation on the basis of a set of criteria/testbed, used to measure the level of performance in terms of the success of the agent's behavior. Such performance measures should be based on the desired effect of the agent on the environment. In particular, the agent's rational behavior depends on (i) the performance measure that defines success; (ii) the agent's knowledge of the environment; (iii) the action that it is capable of performing; and, (iv) the current sequence of perceptions. In general, for every possible perception sequence, the agent is expected to take an action that will maximize its performance measure.

- *benevolence*: an IA does not have conflicting goals, and it will always strive to satisfy the requests made of it. This characteristic simplifies modeling an agent by assuming that its goals always align with the tasks or requests entrusted to it. Consequently, conflicts of interest or competing goals between agents are ignored, allowing for a more direct conceptualization of agent behavior. This means that an agent, according to the assumption, does not face internal struggles between different goals. Benevolence also introduces an element of predictability into the behavior of an IA that can be advantageous in specific contexts, especially when simplicity and clarity in agent behavior are prioritized.

- *veracity*: an IA processes, manages and provides accurate and reliable information, which allows for more reliable decision-making processes, explaining how results were achieved. The agent maintains the integrity of the data, used to derive information, throughout its lifecycle, recognizing and mitigating biases in the data to ensure fair and unbiased results.

Based on such characteristics, well-known IA models have been identified such as Simple Reflex, Model-Based Reflex, Goal-Based, Utility-Based, and Learning.

- *Simple Reflex*: The decision of the action to take, it is only based on the current perception. The history and the perceptions gathered in the past are neglected. This model is based on *condition-action* rules. This model works if the environment is fully observable (stateless).

- *Model-Based Reflex*: this agent model works when the world is not fully observable. As a consequence, it is important that the agent remember previous observations about the parts of the environment that cannot be observed in a particular period of time. This requires a supporting model for representing the environment (state-full).

- *Goal-Based*: this agent model aims at driving the agent to reach a specific purpose and action to be taken depends on the current state and on what it tries to accomplish (the goal). Sometimes, the goal to achieve requires a single action; in another case, the goal to be reached is complex and decomposed into multiple sub-goals, each of which requires one or a set of actions. In this case, the achievement of all subgoals subsumes achieving the main goal. Usually, in this case, strategies, planning, and sifting through a search space for possible solutions are necessary.

- *Utility-Based*: this can be seen as a supporting or a complementary model of the Goal-Based model previously described. In this case, the agent knows the utility function that is continuously monitored and exploited to estimate the distance between the goal to be achieved and the current state of the agent's goal.

- *Learning*: this agent model has the capability of enriching its "knowledge" and abilities by observing and acting consequently. This means that the agent is able to learn from past occurrences in the environment to predict the future and in some cases (pro)actively affect the environment.

The next section focuses on the concepts of environment, its role, and its main properties.

3. Environment

While the notion of Intelligent Agents is obviously central for this subject, all the most widespread definitions of agent at least mention the fact that a surrounding *environment* is present, for instance, to provide percepts and a context in which actions are attempted. Whereas the awareness of the significance of the environment in which an Intelligent Agent is situated was already present in the earliest version of the most widely adopted book on Artificial Intelligence (Russell and Norvig, 2003), within the autonomous agents and multi-agent systems research curiously, the recognition of the environment as an explicit and essential part of a multi-agent system required some time and a systematic analysis of the typical practice of researchers in the area. Weyns et al. (2007), in a foundational paper on this topic, provide the following definition:

> The environment is a first-class abstraction that provides the surrounding conditions for agents to exist and that mediates both the interaction among agents and the access to resources.

The authors also highlight the fact that the environment is a first-class abstraction for agent-oriented models, not just providing the surrounding conditions for agents to exist, but also representing an exploitable design abstraction for building multi-agent system applications.

Russell and Norvig, in the above-cited book, provided several dimensions for the characterization of an environment, and in particular, the most relevant in this context are:

- *observability*: agents can have complete or partial access to the state of the environment;
- *determinism*: in deterministic environments, agents' actions have single, guaranteed effects;

- *dynamism*: in a static environment, agents can assume that no change happens during its own deliberation;
- *discreteness*: discreteness can refer to different aspects of the environment, namely its state, the way time is represented and managed, the perceptions and actions of agents; generally, in a discrete environment, there are a fixed, finite number of actions and percepts in it.

Examples of environments and their respective characterization are shown in Table 1. Clearly, the features of the environment heavily influence the design decisions about the agent architecture; it must also be clarified that it is sometimes possible to take a simplifying but still acceptable perspective on specific aspects of an environment to actually come up with applicable and tractable solutions.

Table 1: Example of different environments and their characterization.

Example	**Observable**	**Deterministic**	**Static**	**Discrete**
Chess	Fully	Yes	Yes	Yes
Poker	Partly	No	Yes	Yes
Real-time strategy	Partly	No	No	No

Recently, Russell and Norvig also included an additional dimension of analysis that specifies if the environment includes other agents and, in this case, also the cooperative or competitive attitude of agents should be discussed (a more thorough analysis of the different types of interaction is provided by Ferber (1999)).

The more recent analysis, provided by Weyns et al., considers it from the perspective of a design abstraction supporting the activities of the modeler or engineer instead of describing the inherent features of an environment. Their analysis considers that the environment can provide three different levels of support:

- *basic level*, essentially just enabling the agents to directly access their deployment context;

- *abstraction level*, filling the conceptual gap between the agent abstraction and low-level details of the deployment context (e.g., wrapping physical or software resources and providing access at agents' level of abstraction);

- *interaction/mediation level*, supporting both forms of regulation to the above mentioned resources, as well as mediating the interaction among agents to support forms of coordination.

Whereas most agent-oriented platforms provide an abstraction level support, the support to the interaction/mediation level is generally not as comprehensive and systematic, as testified by works discussing meta-models for multi-agent systems explicitly including abstractions enabling this kind of high-level support (see, e.g., Omicini et al. (2008a)).

Weyns et al. also stress the fact that the environment has a fundamental role in *structuring* the system: this is particularly relevant for the sake of applications like the simulation of biological entities in which the defined model actually needs to represent a physical spatial structure (e.g., portions of tissues of a human body), but it can also be relevant in situations in which the model must consider other conceptual structures such as organizational or societal ones (e.g., roles, groups, permissions, policies). Whenever the computational model needs to represent a physical environment, and to explicitly consider its spatial aspects and even dynamical processing taking place in even without involving the modeled agents, the need for a precise and systematic model of perception and action is even more apparent than in other situations. From this perspective, models like the one described by Ferber and Müller (1996) represent relevant examples of a specific form of high-level support supplied by the environment model.

Finally, once again, especially but not exclusively in the biological context, it is often the case that the modeler needs to consider distinct but related dynamics that are more reasonably or effectively represented by employing different spatial or temporal scales. The overall multi-agent model could, therefore, include different environmental representations at different scales, potentially characterized by different features according to the above-introduced schema, or the overall model could even employ completely different styles in a hybrid approach (as discussed by Dada and Mendes (2011)): the different dynamics must then be properly coupled using some form of interaction among the different sub-models and scales.

4. Actions and Interactions

The notion of *action* directly contributes to the definition of agent: literally, *the one who acts*. Understanding the reciprocal *dependencies* between individual agents and their surrounding environment – either physical or computational, there including the space-time fabric – then amounts to understanding which kind of actions produce which kind of *effects* – either intended or not – and on *whom*—here, either another agent or (a portion of) the environment.

The distinction regarding the kind of actions taken as a reference throughout the chapter is made by focusing on the *purpose* of an action (Kirsh and Maglio, 1994):

- *epistemic actions* are meant to acquire/release information, and may have or not a direct practical effect, either intended or not;
- *practical actions*, on the contrary, are meant to directly affect a subject, and may have or not have a direct epistemic effect, either intended or not.

From this stems the distinction w.r.t. the kind of effects caused by actions:

epistemic effects directly cause to acquire/release information

practical effects directly cause a change in the subject

As the reader may notice, epistemic actions may have practical effects (possibly, not intended, and indirectly), and practical actions may have epistemic effects in turn (again, possibly not intended and indirectly). Despite how odd this could seem, the popular distinction categorizing actions in either *communicative* or *practical* ones have been proven to be misleading by Conte and

Castelfranchi (1995), argumenting that practical actions are in all respect communicative actions too when they have an intended (although possibly implicit) communicative effect.

Anyway, epistemic actions are *mostly* based on *communication*, be it explicit or not, thus likely require (FIPA, 1996)—at the very least:

- a *content language*, that is, a language for "talking about things";
- a set of *communicative acts*, that is, the acts through which "communication happens".

Besides, fruitful communication among *computational* agents likely requires the content language to be shared among participants in a "conversation", the communicative acts to have a well-defined shared semantics, and the conversations to adhere to prescribed interaction protocols guaranteeing some desired properties.

As far as practical actions are concerned, they are *mostly* based on *practical behaviours*, thus likely require:

- *perception* (Russell and Norvig, 2003) of the acting agent surroundings for detecting the subject of the action;
- *context awareness* (Abowd et al., 1999) to perceive and understand the context in which the agent operates (including factors such as location, time, privacy and security) to make more informed decisions and provide tailored responses or actions;
- *situation recognition* (So and Sonenberg, 2004) to interpret and categorize the current state of affairs in the environment by recognizing specific situations or patterns of events. This recognition allows the agent to respond appropriately to dynamic changes in its surroundings and make decisions based on a higher-level understanding of the situation;
- *practical reasoning* grounded upon *bounded rationality* (Bratman, 1987) to plan the course of actions to undertake toward achievement of a goal

while considering feasibility, expected utility, likelihood of success, and cost of actions themselves.

Regardless of the purpose and the effects, it is apparent that actions are influenced by their subject – who it is, an agent or the environment? – as well as by their surroundings—where and when the action is taking place? Which properties may influence actions outcome, feasibility, etc.? In other words, actions are *situated* w.r.t. their *context*, which brings us to the next subsection.

4.1. Situated (inter)action

Situatedness is the property of being immersed within an *environment* (Suchman, 1987), that is, the property of being potentially influenced and, in turn, potentially capable of affecting someone or something.

Actions are then situated by definition: regardless of whether they are pragmatical or epistemic, they have a target, a source, happen at a given time (and possibly have a duration, and/or a delay), affect a given space (either virtual or physical), and cause some change (at least, if successful) either intended or not ("side effects"). Thus, agents too are situated in turn: through actions, they can be regarded as being "active" at a given time, in a precise space, for a given observer (e.g., the target of the action), either because through actions they affected someone or something, or because their course of actions has been influenced by someone or something.

It is worth noting that situatedness directly relates to the notions of *perception* and *practical reasoning*: perception is the tool by which agents and the environment become *aware* of their surroundings (Russell and Norvig, 2003), thus potentially influenced by therein activities; practical reasoning is the tool by which agents (and an intelligent environment? Mariani (2016)) can *deliberate* about how to affect something or someone (Bratman, 1987).

Inter-actions then are situated too, iterating the same reasoning: each of the participants in the interaction is situated, then the (inter-)actions they carry out *reciprocally affecting* each other are situated too—both because they are actions anyway, although (possibly) *communicative*, and because the acting participants are situated. Besides, interactions may be situated also because they are *mediated* by some means external to the participants, e.g., the environment. In this case, being the environment situated due to its very nature, any interaction it enables and constrains may be regarded as situated as well, through the properties of the environment—the flow of time, the topology of space, the existence of resources and properties with their own dynamics.

Whether (inter)-actions are mediated by the environment or not, they always represent a *social* relationship between the parties involved, which brings us to the next subsection.

4.2. Social (inter)action

Castelfranchi discusses how the complex and distributed dependencies within the agents in a MAS – mainly regarding goals (Dennett, 1971), delegation and trust (Castelfranchi and Falcone, 1998) – are fundamental to the definition of intelligence as a social construct (Castelfranchi et al., 1993). In particular, the notion of *social action* – meant to reconcile individual cognitive processes and social coordination – provides a conceptual foundation which all MAS social issues (cooperation, collaboration, competition) can be grounded on.

Not by chance, in fact, one of the first relevant acts of the Foundation for Intelligent Physical Agents (FIPA, 1996) – a world-wide organization devoted to agent-based technology – was defining a reference semantics for the FIPA Agent Communication Language (FIPA-ACL), also defining the semantics of social actions – there intended as messages exchange – which is now the standard for agent-oriented middleware (Bellifemine et al., 2001).

Along the same line, the many different kinds of social relationships expressed by social interactions have been recognized to deserve a first-class abstraction in the process of MAS engineering, especially when it comes to governing the *space of interactions*, a task which originated a whole research thread, branded *coordination models and languages* (Ciancarini, 1996). Accordingly, the SODA methodology explicitly accounts for *societies* of agents (Omicini, 2001), in line with the A&A meta-model (Omicini et al., 2008b) which adds the notion of *artefact* – in particular, coordination artefacts – as a means for agents to augment their capabilities, both practical and cognitive, and to structure the societies they live in—as well as the MAS environment.

An alternative approach is the one where social action is based on the notion of *social commitment*, a directed relationship from a debtor agent to a creditor agent to bring about a goal of interest when some given context becomes true (Singh, 2011; Marengo et al., 2011).

5. Agents in Bioinformatics and Computational Biology

It is not the aim of this chapter (nor it is considered actually possible) to provide a compact but comprehensive review of the most notable applications of agent models and technologies to bioinformatics and systems biology. Nonetheless, it is reasonable and useful to give an idea of the *categories* or areas of these applications.

First of all, as also noted in a general resource describing the state of the art and perspectives on agent-based computing (Luck et al., 2005), simulation represents an application context in which the notion of autonomous agent has become almost ubiquitous: An et al. (2009), for instance, present a review of agent-based modeling approaches to translational systems biology. However, a plethora of applications to other relevant areas and approaches could also be reported. What is worth noting is the fact that very often, the notion of agent employed in these works takes a very different perspective than the one provided by the most widely accepted definitions of Intelligent Agents, being more focused on studying the resulting or emerging behavior of the local actions and interactions of relatively simple agents than to define and employ knowledge-level agents involved in complicated patterns of coordination.

The latter, however, become instead quite relevant to design and implement solutions for managing specific parts or the overall workflow of scientists working in the field, in a more general e-science perspective, also as a consequence of results from the sub-field of Agent Oriented Software Engineering (see, e.g., Jennings (2001) and Bergenti et al. (2004)). For instance, Miles (2006) presents an example of an application in which agent-based approaches have been employed for performing data curation in the bioinformatics area, whereas Bartocci et al. (2007) describe a web-based Workflow Management System for bioinformatics that employs an agent-based middleware. Another application of Intelligent Agents in bioinformatics by Graumann et al. (2012) describes the design and functioning of an intelligent agent supporting data acquisition and real-time database searching for Shotgun Proteomics, also employing the gene ontology (GO)[2]. Other promising applications in Computational Biology involve the use of Intelligent Agents equipped with Reinforcement Learning techniques for solving specific problems ranging from Drug Discovery and Design to Personalized Medicine and Metabolic Engineering (Mahmud et al., 2018).

In Computational Biology and Bioinformatics, the convergence of Intelligent Agents and the Internet of Things (IoT) allows researchers to explore the complexity of biological systems with efficiency, speed, and depth of analysis (Merelli et al., 2007). Intelligent Agents, equipped with autonomous learning and reasoning capabilities, interact with IoT devices to collect, process and identify complex biological models using data coming from IoT devices. This includes identifying genetic sequences, protein interactions, and correlations between genetic and environmental factors. This is possible because Intelligent Agents adapt their behavioral model over time, improving the accuracy and efficiency in recognizing correlations (Gheysari and Tehrani, 2022). The integration between Intelligent Agents and IoT also leads to a plethora of applications for health care, such as Remote Patient Monitoring, Predictive Analytics for Disease Management, Hospital Asset Tracking and Management, Healthcare Supply Chain Optimization, Infection Control and Monitoring (see Yuehong et al. (2016); Amin and Hossain (2020) for comprehensive reviews on the topic).

[2]https://www.geneontology.org/docs/ontology-documentation/

6. Closing Remarks

The domain of Intelligent Agents is highly escalating at an exponential rate (Singh et al., 2017). The basic concepts offer powerful solutions for designing a variety of agent-based solutions. Since the last three decades, the advances in the domain have been significant, specifically pertaining to agent interaction and communication in complex distributed systems. The technology is now finding space in commercial applications and complex and mission-critical applications as well. Efforts to improve the computing are still going on, and many issues, such as defining the limit on trust, intelligence, and their applicability, are yet open. The full potential of agents is yet to be realized.

Abowd, G.D., Dey, A.K., Brown, P.J., Davies, N., Smith, M., Steggles, P. (1999). Towards a better understanding of context and context-awareness, in: Handheld and Ubiquitous Computing: First International Symposium, HUC'99 Karlsruhe, Germany, September 27–29, 1999 Proceedings 1, Springer. pp. 304–307.

Adamatti, D. (2016). Multi-agent-based simulations applied to biological and environmental systems.

Amin, S.U., Hossain, M.S. (2020). Edge intelligence and internet of things in healthcare: A survey. Ieee Access 9, 45–59.

An, G., Mi, Q., Dutta-Moscato, J., Vodovotz, Y. (2009). Agent-based models in translational systems biology. Wiley Interdisciplinary Reviews: Systems Biology and Medicine 1, 159–171.

Anderson, P. (2015). On the origins of that most transformative of biological systems - the nervous system. Journal of Experimental Biology 218, 504–505.

Bartocci, E., Corradini, F., Merelli, E., Scortichini, L. (2007). Biowms: a web-based workflow management system for bioinformatics. BMC Bioinformatics 8.

Bellifemine, F., Poggi, A., Rimassa, G. (2001). Jade: A fipa2000 compliant agent development environment, in: Proceedings of the Fifth International Conference on Autonomous Agents, ACM, New York, NY, USA. pp. 216–217.

Bergenti, F., Gleizes, M.P., Zambonelli, F. (Eds.) (2004). Methodologies and Software Engineering for Agent Systems: The Agent-Oriented Software Engineering Handbook. Springer.

Bratman, M. (1987). Intention, Plans, and Practical Reason. Center for the Study of Language and Information.

Castelfranchi, C., Cesta, A., Conte, R., Miceli, M. (1993). Foundations for interaction: The dependence theory. Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 59–64.

Castelfranchi, C., Falcone, R. (1998). Principles of trust for mas: cognitive anatomy, social importance, and quantification, in: Proceedings International Conference on Multi Agent Systems (Cat. No.98EX160), pp. 72–79.

Chen, D., Lu, Z. (2018). A methodological framework for model-based self-management of services and components in dependable cyber-physical systems. Advances in Intelligent Systems and Computing 582, 97–105.

Ciancarini, P. (1996). Coordination models and languages as software integrators. ACM Comput. Surv. 28, 300–302.

Conte, R., Castelfranchi, C. (1995). Cognitive and social action. Psychology Press.

Dada, J.O., Mendes, P. (2011). Multi-scale modelling and simulation in systems biology. Integr. Biol. 3, 86–96.

Dennett, D.C. (1971). Intentional systems. The Journal of Philosophy 68, 87–106.

Ferber, J. (1999). Multi-Agent Systems. Addison-Wesley.

Ferber, J., Müller, J.P. (1996). Influences and reaction: a model of situated multiagent systems, in: Proceedings of Second International Conference on Multi-Agent Systems (ICMAS-96), AAAI. pp. 72–79.

FIPA (1996). Foundation for intelligent physical agents.

Garro, A., Tundis, A. (2015). On the reliability analysis of systems and sos: The ramsas method and related extensions. IEEE Systems Journal 9, 232–241.

Gheysari, M., Tehrani, M.S.S. (2022). The role of multi-agent systems in iot, in: Multi Agent Systems: Technologies and Applications towards Human-Centered. Springer, pp. 87–114.

Graumann, J., Scheltema, R.A., Zhang, Y., Cox, J., Mann, M. (2012). A framework for intelligent data acquisition and real-time database searching for shotgun proteomics. Molecular & Cellular Proteomics 11.

Hayes-Roth, B. (1995). An architecture for adaptive intelligent systems. Artificial intelligence 72, 329–365.

Jennings, N.R. (2001). An agent-based approach for building complex software systems. Communications of the ACM 44, 35–41.

Kirsh, D., Maglio, P. (1994). On distinguishing epistemic from pragmatic action. Cognitive science 18, 513–549.

Luck, M., McBurney, P., Sheory, O., Willmott, S. (Eds.) (2005). Agent Technology: Computing as Interaction. University of Southampton.

Magedanz, T., R.K.K.S. (1996). Intelligent agents: An emerging technology for next generation telecommunications. Proceedings of the IEEE Fifteenth Annual Joint Conference of the IEEE Computer Societies on Networking the Next Generation INFOCOM'96 2, 464–472.

Mahmud, M., Kaiser, M.S., Hussain, A., Vassanelli, S. (2018). Applications of deep learning and reinforcement learning to biological data. IEEE transactions on neural networks and learning systems 29, 2063–2079.

Marengo, E., Baldoni, M., Baroglio, C., Chopra, A.K., Patti, V., Singh, M.P. (2011). Commitments with Regulations: Reasoning about Safety and Control in REGULA, in: Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2011, IFAAMAS, Taipei, Taiwan. pp. 467–474.

Mariani, S. (2016). Coordination of complex sociotechnical systems: Self-organisation of knowledge in mok .

Merelli, E., Armano, G., Cannata, N., Corradini, F., d'Inverno, M., Doms, A., Lord, P., Martin, A., Milanesi, L., Möller, S., et al. (2007). Agents in bioinformatics, computational and systems biology. Briefings in bioinformatics 8, 45–59.

Miles, S. (2006). Agent-oriented data curation in bioinformatics. ITSSA 1, 43–50.

Omicini, A. (2001). SODA: Societies and Infrastructures in the Analysis and Design of Agent-Based Systems. Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 185–193.

Omicini, A., Ricci, A., Viroli, M. (2008a). Artifacts in the a&a meta-model for multi-agent systems. Autonomous Agents and Multi-Agent Systems 17, 432–456.

Omicini, A., Ricci, A., Viroli, M. (2008b). Artifacts in the a&a meta-model for multi-agent systems. Autonomous Agents and Multi-Agent Systems 17, 432–456.

Russell, S.J., Norvig, P. (2003). Artificial Intelligence: A Modern Approach. Pearson Education. 2 edition.

Sampigethaya, K., Poovendran, R. (2012). Cyber-physical integration in future aviation information systems, in: Digital Avionics Systems Conference (DASC), 2012 IEEE/AIAA 31st, IEEE. pp. 7C2–1.

Seekhao, N., Shung, C., Jaja, J., Mongeau, L., Li-Jessen, N. (2016). Real-time agent-based modeling simulation with in-situ visualization of complex biological systems: A case study on vocal fold inflammation and healing, pp. 463–472.

Singh, A., Juneja, D., Singh, R., Mukherjee, S. (2017). A thorough insight into theoretical and practical developments in multiagent systems. Int. J. Ambient Comput. Intell. 8, 23–49.

Singh, M.P. (2011). Commitments in multiagent systems some controversies, some prospects, in: Paglieri, F., Tummolini, L., Falcone, R., Miceli, M. (Eds.), The Goals of Cognition. Essays in Honor of Cristiano Castelfranchi. College Publications, London. chapter 31, pp. 601–626.

Smith, D.C., Cypher, A., Spohrer, J. (1994). Kidsim: Programming agents without a programming language. Commun. ACM 37, 54–67.

So, R., Sonenberg, L. (2004). Situation awareness in intelligent agents: Foundations for a theory of proactive agent behavior, in: Proceedings. IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2004.(IAT 2004)., IEEE. pp. 86–92.

Suchman, L.A. (1987). Plans and situated actions: The problem of human-machine communication, Cambridge university press.

Weyns, D., Omicini, A., Odell, J. (2007). Environment as a first class abstraction in multiagent systems. Autonomous Agents and Multi-Agent Systems 14, 5–30.

Yuehong, Y., Zeng, Y., Chen, X., Fan, Y. (2016). The internet of things in healthcare: An overview. Journal of Industrial Information Integration 1, 3–13.

Zhabelova, G., Vyatkin, V. (2012). Multiagent smart grid automation architecture based on iec 61850/61499 intelligent logical nodes. IEEE Transactions on Industrial Electronics 59, 2351–2362.