

Principles and practice of multi-agent systems

QINGLIANG CHEN¹, PAOLO TORRONI² and SERENA VILLATA³

¹*Department of Computer Science, Jinan University, Guangzhou 510632, China;*
e-mail: tpchen@jnu.edu.cn;

²*Department of Computer Science and Engineering, University of Bologna, Bologna 40136, Italy;*
e-mail: paolo.torroni@unibo.it;

³*Centre National de la Recherche Scientifique (CNRS), I3S Laboratory, Sophia Antipolis 06560, France;*
e-mail: villata@i3s.unice.fr

Preface

Agent-based computing addresses the challenges in managing distributed computing systems and networks through monitoring, communication, consensus-based decision-making and coordinated actuation. As a result, intelligent agents and multi-agent systems have demonstrated the capability to use intelligence, knowledge representation and reasoning, and other social metaphors like trust, game and institution, not only to address real-world problems in a human-like way but also to transcend human performance. This has had a transformative impact in many application domains, particularly in e-commerce, and also in planning, logistics, manufacturing, robotics, decision support, transportation, entertainment, emergency relief and disaster management, data mining and analytics. As one of the largest and still growing research fields of computer science, agent-based computing today remains a unique enabler of inter-, multi- and trans-disciplinary research.

The International Conference on Conference on Principles and Practice of Multi-agent Systems (PRIMA) originally started in 1998 as a regional (Asia-Pacific) workshop and in the last decade it grew to become one of the leading and influential scientific conferences for research on multi-agent systems. Each year, PRIMA brings together active researchers, developers and practitioners from both academia and industry to showcase, share and promote research in several domains, ranging from foundations of agent theory and engineering aspects of agent systems, to emerging interdisciplinary areas of agent-based research. Previous successful editions were held in Nagoya, Japan (2009), Kolkata, India (2010), Wollongong, Australia (2011), Kuching, Malaysia (2012), Dunedin, New Zealand (2014) and Gold Coast, Australia (2014). The last two editions were held in Phuket, Thailand (2016) and in Nice, France (2017).

The eighteenth edition of PRIMA took place from 26 to 30 October 2015 in Bertinoro, FC, Italy and received 94 submissions from 30 countries. From the 29 papers presented at the conference, the best theoretically-oriented ones were invited to submit an extended version for the special issue with *Fundamenta Informaticae*, while another special issue with the *Journal of Agent-Oriented Software Engineering* has been organized around the more practically-oriented papers.

This special issue collects a selection of survey contributions intended to provide a snapshot of trends and hot topics in PRIMA research.

In particular, the paper *Answer Set Programming and Agents* by Abeer Dyoub, Stefania Costantini and Giovanni De Gasperis discusses the potential role of Answer Set Programming (ASP) in the context of approaches to the development of agents and multi-agent systems especially in the realm of computational logic. After shortly recalling the main agent-oriented frameworks, the authors introduce ASP and discuss the usefulness of a potential integration of the two paradigms in a modular heterogeneous framework, and the feasibility of such integration.

Mehdi Dastani, Paolo Torroni and Neil Yorke-Smith, in their paper *Monitoring Norms: A Multi-Disciplinary Survey* discuss the notion of monitoring in normative systems. If a norm is defined as an expected behavior in a social setting, an issue is how to determine whether an individual or a group are

adhering to this expected behavior. This process is called monitoring, and it is fundamental for processes of accountability, enforcement, regulation and sanctioning. Starting with a broad focus and narrowing to the multi-agent systems literature, the authors address four key questions: what is monitoring, what is monitored, who does the monitoring and how the monitoring is accomplished.

In *Predictive Models and Abstract Argumentation: The Case of High-Complexity Semantics*, Mauro Vallati, Federico Cerutti and Massimiliano Giacomin describe how predictive models can be positively exploited in abstract argumentation. On one side, they show that predictive models are effective for performing algorithm selection to determine which approach is better to enumerate the preferred extensions of a given argumentation framework. On the other side, they show that predictive models predict significant aspects of the solution to the preferred extensions enumeration problem. By exploiting an extensive set of argumentation framework features, the proposed approach is able to provide an accurate prediction about which algorithm would be faster on a given problem instance, as well as the structure of the solution, where the complete knowledge of such structure would require a computationally hard problem to be solved.

Radu-Casian Mihailescu, Paul Davidsson, Ulrik Eklund and Jan A. Persson, in their paper *A Survey and Taxonomy on Intelligent Surveillance from a System Perspective*, discuss the current state-of-the-art and provides researchers with an overview of existing surveillance solutions, analyzing their properties as a system and drawing attention to relevant challenges when developing, deploying and managing them. In an effort to understand the development of the advanced solutions, based on their most distinctive characteristics, the authors propose a taxonomy for surveillance systems to help classify them and reveal gaps in existing research.

The paper *Towards Adaptive Multi-Robot Systems: Attempts for Self-Organization and Self-Adaptation* by Christopher-Eyk Hrabia, Marco Lutzenberger and Sahin Albayrak, presents a comprehensive overview over state-of-the-art solutions in emergent systems, self-organization, self-adaptation and robotics. The development of complex systems ensembles that operate in uncertain environments is a major challenge, as system designers are not able to fully specify the system during specification and development and before it is being deployed. Natural swarm systems enjoy similar characteristics, yet, being self-adaptive and being able to self-organize, these systems show beneficial emergent behavior. Similar concepts can be extremely helpful for artificial systems, especially when it comes to multi-robot scenarios, which require such solution in order to be applicable to highly uncertain real-world application.

In *The Composition and Formation of Effective Teams. Computer Science Meets Organisational Psychology*, Ewa Andrejczuk, Rita Berger, Juan Rodriguez-Aguilar, Carles Sierra and Víctor Marín-Puchades review the contributions in both the computer science literature and the organisational psychology literature dealing with the composition and formation of effective teams. On the one side, they aim at identifying the strengths and weaknesses of the contributions made by these two diverse bodies of research. On the other hand, they aim at identifying cross-fertilization opportunities that help both disciplines benefit from one another.

The paper *Dimensions in Programming Multi-Agent Systems* by Olivier Boissier, Rafael H. Bordini, Jomi F. Hubner and Alessandro Ricci, discusses the fact that research on multi-agent systems moved from agent-oriented programming towards the idea of multi-agent oriented programming, in which a multi-agent system program is designed and developed using together concepts and first-class design and programming abstractions that come from various 'dimensions', for example, agent, interaction, environment and organization. In this paper, the authors provide a practical overview of multi-agent oriented programming using JaCaMo, a platform for multi-agent oriented programming built on top of Jason for programming beliefs, desires, intentions agents, going through a sequence of programming steps of increasing complexity.

Conor Muldoon, Michael J. O'Grady and Gregory M.P. O'Hare discuss in their paper *A Survey of Incentive Engineering for Crowdsourcing* game theoretic approaches to the problem in the crowdsourcing domain and places it in the context of the wider research landscape. With the growth of the Internet, crowdsourcing has become a popular way to perform intelligence tasks that hitherto would be either performed internally within an organization or not undertaken due to prohibitive costs and the lack of an appropriate communications infrastructure. In crowdsourcing systems, whereby multiple agents are not

under the direct control of a system designer, it cannot be assumed that agents will act in a manner that is consistent with the objectives of the system designer or principal agents. In situations whereby agents goals are to maximize their return in crowdsourcing systems that offer financial or other rewards, strategies will be adopted by agents to game the system if appropriate mitigating measures are not put in place.

In *A Review of Learning Planning Action Models*, Ankuj Arora, Humbert Fiorino, Damien Pellier, Marc Merivier and Sylvie Pesty present a survey of the machine learning techniques applied for learning planning action models. Automated planning has been a continuous field of study since the 1960s, since the notion of accomplishing a task using an ordered set of actions resonates with almost every known activity domain. However, as we move from toy domains closer to the complex real world, these actions become increasingly difficult to codify. The reasons range from intense laborious effort, to intricacies so barely identifiable, that programming them is a challenge that presents itself much later in the process. In such domains, planners now leverage recent advancements in machine learning to learn action models. This learning provides an opportunity for the evolution of the model towards a version more consistent and adapted to its environment, augmenting the probability of success of the plans. It is also a conscious effort to decrease laborious manual coding and increase quality.

The 10 papers in the present collection are the result of a careful selection process followed by a thorough two-stage reviewing procedure. They represent a significant contribution to the state of the art and the understanding of the theory and practice of multi-agent systems, as well as offer interesting connections with other domains.

We thank the authors for the considerable effort put into producing papers of such a high quality, the reviewers for their care and constructive criticism, and the Co-Editors-in-Chief of the Knowledge Engineering Review, Peter McBurney and Simon Parsons, for allowing us to publish the resulting collection of papers as a special issue of this journal. Enjoy reading!