Principles and Practice of Multi-Agent Systems

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 & disaster management, and data mining & 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).

This issue contains selected papers from the eighteenth edition of PRIMA, which took place from 26 to 30 October 2015 in Bertinoro, FC, Italy. The conference 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 this 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.

The contributions that were eventually submitted underwent a thorough two- or three- stage reviewing procedure, resulting in the ten papers in the present collection. They constitute the most recent advances in the theory and practice of multi-agent systems, with interesting intersections with other domains.

Since agents are entities that reason and interact in a social environment, considerable research has been conducted in the last decades in multi-agent communication and dialogues. In *Multi-Party Persuasion: a Paraconsistent Approach*, Alina Powała and Barbara Dunin-Kęplicz focus on conflicts of opinions resolved by a persuasion dialogue and on semantically deeper conflicts reaching to motivations of opinions. Their investigation leads to proposing deep persuasion dialogues for the resolution of conflicting motivations of opinions, resulting in the TalkLOG formal framework.

Other long-standing domain of interest for PRIMA are resource allocation and distributed constraint satisfaction and optimization. In *Strategy-proof Cake Cutting Mechanisms for All-or-nothing Utility*, Takamasa Ihara, Shunsuke Tsuruta, Taiki Todo, Yuko Sakurai and Makoto Yokoo consider the all-or-nothing utility function as a representative example of non-additive utility. Motivated by an incompatibility between envy-freeness and Pareto efficiency when each agent has all-or-nothing utility, the authors propose two strategy-proof mechanisms that satisfy Pareto efficiency: one based on the serial dictatorship mechanism, as well as another one which satisfies envy-freeness.

In *Leximin Multiple Objective DCOPs on Factor Graphs for Preferences of Agents*, Toshihiro Matsui, Marius Silaghi, Tenda Okimoto, Katsutoshi Hirayama, Makoto Yokoo and Hiroshi Matsuo address distributed constraint optimization problems (DCOPs), enabling to formalize applications as constraint optimization problems where variables and functions are distributed among agents. Focusing on a class of Multiple Objective DCOPs called Leximin AMODCOP, where multiple objectives for individual agents are optimized based on the leximin operator, the authors address the Leximin AMODCOPs on factor graphs that directly represent n-ary functions, proposing an exact, dynamic programming-based solution on factor graphs, as well as investigate several approximate/inexact algorithms for relatively dense problems.

The need to develop powerful and expressive frameworks that leverage the potential of agent theories for practical applications inspired several agent architectures, languages and environments. In *Programming Agent Deliberation Using Procedural Reflection*, Sam Leask and Brian Logan look into agent deliberation strategies, proposing a way to embed them into agent programs, enabling application-specific BDI deliberation strategies to be programmed in a simple way.

Another prolific recent line of research in the PRIMA scope is the relation between logics and game theory. In *The Long-Term Benefits of Following Fairness Norms under Dynamics of Learning and Evolution*, Emiliano Lorini and Rohland Mhlenbernd present a game-theoretic model of guilt in relation to sensitivity to norms of fairness. The authors focus on a specific kind of fairness norm according to which a fair society should be organized so as to admit economic inequalities to the extent that they are beneficial to the less advantaged agents. The paper uses evolutionary game theory to study the impact of the sensitivity to this fairness norm on the behavior of agents, to reveal that a great sensitivity to the fairness norm is beneficial in the long term, and that cooperativeness and fairness norm sensitivity can coevolve in a population of initially solely defectors.

A different perspective is instead provided by Chanjuan Liu, Fenrong Liu and Kaile Su in their work *A Dynamic-Logical Characterization of Solutions in Sight-limited Extensive Games*. There, the authors develop the DLS logic enabling a characterization of the solutions to these games. They also propose an algorithm for model checking DLS and discuss possible applications of the logic.

Distributed and preference-based agent reasoning is where work by Tiep Le, Trao Can Son, and Enrico Pontelli is positioned. In *Multi-Context Systems with Preferences*, the authors present an extension of the Multi-Context Systems (MCS) framework to allow the encoding of preferences at the level of the contexts. The work is motivated by the observation that a naive use of preference logics can lead to undesirable outcomes. To address this issue, the authors define weakly and strongly-preferred

equilibria through the notion of ranked logics, and illustrate how such a framework enables modeling distributed configuration problems and finding explanations for distributed abductive diagnosis problems.

In a related context, in *Sequence Semantics for Modelling Reason-based Preferences* Erica Calardo, Guido Governatori, and Antonino Rotolo study how the non-classical n-ary operator $\otimes$, originally intended to capture the concept of reparative obligation, can be used in the context of social choice theory to model preferences. The authors propose a sound and complete axiomatization of a minimal modal logic for the operator, extending it with axioms suitable to model social choice consistency principles such as extension consistency and contraction consistency.

Social choice theory and judgment aggregation is also addressed by Daniele Porello in *Logics for Modelling Collective Attitudes*, where the author introduces a number of logics to reason about collective propositional attitudes that are defined by means of the majority rule. The proposed logics are based on a substructural propositional logic that allows for circumventing inconsistent outcomes. Individual and collective propositional attitudes, such as beliefs, desires, obligations, are then modeled by means of minimal modalities to ensure a number of basic principles, thus obtaining a viable consistent modeling of collective attitudes.

A different perspective on judgment aggregation is finally provided by Ewa Andrejczuk, Juan A. Rodriguez-Aguilar and Carles Sierra in *Collaborative Rankings*. There, the authors introduce the collaborative judgment ranking algorithm to account for peer opinions of agents and/or humans on objects, as well as peer judgments over those opinions. Collaborative judgment is applied to the use case of scientific paper assessment and validated over simulated data, showing that the rankings thus obtained improve current scientific paper ranking practice.

We thank the authors for putting effort into submitting and revising their papers in order to achieve their current quality, the reviewers for their care and constructive criticism, and the Editor-in-Chief of Fundamenta Informaticae, Damian Niwiński, for allowing us to publish the resulting collection of papers as a special issue of this journal. Enjoy reading!

Editors:

Qingliang Chen
Jinan University, China
tpchen@jnu.edu.cn

Paolo Torroni
Alma Mater – Università di Bologna, Italy
p.torroni@unibo.it

Serena Villata
Université Côte d'Azur, CNRS, Inria, I3S, France
villata@i3s.unice.fr