Preface

Jordi Sabater-Mir [a,*] and Paul Davidsson [b,c]

[a] *IIIA-CSIC, Campus UAB, Bellaterra, Catalonia, Spain*
E-mail: jsabater@iiia.csic.es
[b] *School of Computing, Blekinge Institute of Technology, Karlskrona, Sweden*
[c] *School of Technology, Malmö University, Malmö, Sweden*
E-mail: Paul.Davidsson@bth.se

1. Introduction

The European Workshop on Multi-Agent Systems (EUMAS) is with no doubt one of the most important meetings for the European research community in the area. Since 2003 this workshop series joins every year around 100 researchers coming mainly from European countries but also from the rest of the world. The workshop has no formal proceedings and is used as a platform to present both established research and preliminary ideas in this relatively young area.

The special issue you have in your hands contains a selection of the best papers presented at EUMAS 2009 in Ayia Napa, Cyprus. From 62 papers submitted to the workshop (coming from 18 different countries), 48 were selected by the program committee after peer-review for oral presentation during the event. For this special issue we invited the authors of the papers with the highest review scores to submit an improved and extended version of their paper. A second review process of these extended versions resulted in the selection of the 6 papers you can find in this special issue.

One of the key aspects that makes the multi-agent paradigm so popular is that it opens new challenges for traditional and established AI areas (like learning, reasoning or planning to name just a few). Old mechanisms need to be adapted and new ones developed to satisfy the demand in open environments where multiple (intelligent) entities interact and evolve. The articles in this special issue reflect, at a small scale, the diversity of the multi-agent community.

The article "Generalized learning automata for multi-agent reinforcement learning" by Yann-Michaël De Hauwere, Peter Vrancx and Ann Nowé, falls in the area of learning. The use of reinforcement learning techniques as a mechanism for learning policies in single-agent scenarios has obtained good results so far. However when multiple agents are present, the same nature of these techniques makes them unsuitable due to the exponential increase in the action-state complexity. In this article the authors propose the use of a Generalized Learning Automata (an associative reinforcement learning unit that maps inputs to actions given an internal state) as a possible solution and introduces a two-level framework that mixes traditional reinforcement learning with multi-agent learning when necessary.

The study of norms and normative systems in the multi-agent systems community is of great importance. Any society, and virtual societies are not an exception, needs mechanisms (like for instance normative systems) to regulate the behaviour of their members. In the article "Norm internalization in artificial societies", Giulia Andrighetto, Daniel Villatoro and Rosaria Conte first present a motivating theory taken from philosophy, psychology and sociology that explains the process of norm internalization. Norm internalization is presented as a cognitive mechanism that is able to represent behaviours not yet identified by the multi-agent community. The theory is then used as a base to introduce a detailed design of the normative architecture of the agents. This architecture, called EMIL-I-A, allows agents to achieve a whole normative experience, starting from the identification of the norm, its recognition, self-adaptation of the beliefs of the importance of the norm, and finally, the internalization, translated in the agents decision making as automatisms. In the experimental section, a proof-of-concept simulation based on the Prisoner's Dilemma is presented, showing how agents successfully perform the internalization process orchestrated by the *Salience*

*Corresponding author.

(a sort of measure of relevance) of the norm. Following the work on norms, in their article "Representation and monitoring of commitments and norms using OWL", Nicoletta Fornara and Marco Colombetti explore the use of the OWL Web Ontology Language together with SWRL to specify obligations and prohibitions so a standard reasoner can be used for monitoring these obligations and prohibitions at run-time.

Another topic that has received a lot of attention in the multi-agent systems community and that is inherent to the social nature of the area is coordination and especially its emergence. Norman Salazar, Juan A. Rodriguez-Aguilar and Josep L. Arcos propose a robust and completely distributed spreading mechanism for convention emergence that facilitates coordination in multi-agent systems. They validate the robustness of the mechanism through its ability of dealing with the different issues inherent to multi-agent systems: (i) the existence of multiple convention alternatives with different levels of utility; (ii) the dynamicity of the environment; and (iii) the possibility of unreliable communication between agents.

A problem in most real-world multi-agent applications is that of incomplete information, e.g., when communication between agents is delayed or fails. In the article "Speculative constraint processing for hierarchical agents" by Hiroshi Hosobe, Ken Satoh, Jiefei Ma, Alessandra Russo and Krysia Broda, this problem is addressed using speculative computation. It allows a system to compute tentative solutions by using default knowledge about agents. In the paper, an existing framework for speculative constraint process-ing for master–slave systems is extended to support hierarchically structured multi-agent systems. An operational model and a prototype implementation of the model are presented.

In "Verifying agents with memory is harder than it seemed", Nils Bulling and Wojciech Jamroga make a theoretical contribution by correcting a previously held belief concerning a variant of alternating-time temporal logic (ATL+). Much research has focused on the way such logics can be used for the verification of multi-agent systems, mainly by model checking. In the article they prove that model checking ATL+ for agents that use strategies with memory is in fact PSPACE-complete. On a more positive note, they show that fairness constraints can be added to ATL+ without further increasing the complexity of model checking.

We hope that this selection of excellent papers gives a good sample of the work presented at EUMAS-2009. There are many people without whom neither the workshop nor this special issue would have been possible. We would like to thank the members of the EUMAS advisory board for selecting us to chair the workshop, Antonis Kakas and his team for the local organization which resulted in a very successful event, the programme committee members for helping us in the selection of the best papers, and of course all authors and participants at EUMAS-2009. For this special issue in particular, we are much grateful to the editor-in-chief, Enrico Giunchiglia, for giving us the opportunity to edit this issue of the journal, the reviewers for their great work, and all the authors for contributing such outstanding papers.