

Special issue on adaptive and learning agents 2019

PATRICK MANNION¹ , PATRICK MACALPINE², BEI PENG³ and ROXANA RĂDULESCU⁴

¹*School of Computer Science, National University of Ireland Galway, University Road, Galway H91 TK33, Ireland*
e-mail: patrick.mannion@nuigalway.ie

²*Microsoft Research, Microsoft Building 99, 14820 NE 36th Street, Redmond, WA 98052, USA*
e-mail: patmac@microsoft.com

³*Department of Computer Science, University of Oxford, Parks Road, Oxford OX1 3QD, UK*
e-mail: bei.peng@cs.ox.ac.uk

⁴*Artificial Intelligence Lab, Vrije Universiteit Brussel, Pleinlaan 2, Ixelles 1050, Belgium*
e-mail: roxana.radulescu@vub.be

1 Introduction

The goal of the Adaptive and Learning Agents (ALA) community is to develop autonomous agent-based systems that rely on learning and adaptation techniques to accomplish their tasks. The development of such solutions is an interdisciplinary effort that lies at the intersection of fields such as game theory, multi-agent systems, reinforcement learning, cognitive science, evolutionary computation, multi-objective optimisation, and many more.

Numerous real-world applications, such as traffic optimisation, (smart) grid management, air traffic control, autonomous vehicles, and resource management, require distributed decision-making and control. Furthermore, such applications bring forward many challenges such as the need for cooperation and collaboration among agents, learning how to compromise among different preferred outcomes, and exchanging knowledge in order to reach a common goal. The ALA community aims to study methods suitable for all these challenges and focuses on both fundamental and application-oriented research.

Autonomous decision-making approaches for single and multi-agent systems have shown a steep increase in competence for performing complex tasks through the use of techniques such as neural-based function approximation, knowledge exchange among tasks, and knowledge exchange among agents present in the same environment. Other approaches studied and presented by the ALA community include methods for (expert) human knowledge integration and learning within a multi-objective multi-agent framework.

This special issue features selected papers from the 11th Adaptive and Learning Agents Workshop (ALA 2019), which was held on 13 and 14 May 2019 at the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2019) in Montreal, Canada. The goals of the ALA workshop series are to increase awareness of and interest in adaptive agent research, to encourage collaboration, and to provide a representative overview of current research in the area of ALA. The workshop serves as an interdisciplinary forum for the discussion of ongoing or completed work in ALA and multi-agent systems.

2 Contents of the special issue

This special issue contains 4 papers, which were selected out of 36 initial submissions to the ALA 2019 workshop. All papers were initially presented at the workshop, before being extended and reviewed again for this special issue. These articles provide an overview of current research directions which are being explored within the ALA community.

The first paper *Adaptable and stable decentralized task allocation for hierarchical domains* by Kazakova and Sukthankar (2020) extends a previous model of observed insect behavior (StimHab), to

allow agents to self-allocate to hierarchical sets of tasks. The proposed method is highly scalable as individual agents do not communicate with each other and are not aware of the capabilities, preferences, or circumstances of the other agents in the system. Experiments with a team of 1000 agents in a hierarchical patrolling task demonstrate the efficacy of the authors' approach. Agents using StimHab display a minimal amount of task switching after the initial adaption period, indicating that highly specialized behaviors are reached upon convergence.

The second paper *A utility-based analysis of equilibria in multi-objective normal form games* by Rădulescu *et al.* (2020) explores the effect of nonlinear utility functions on the set of Nash and correlated equilibria in multi-objective normal form games, under different optimization criteria: expected scalarized returns and scalarized expected returns (SER). The authors conclude that there are fundamental differences between these optimization criteria when considering nonlinear utility functions and that equilibria do not always exist under SER in this setting.

The third paper *Domain adaptation-based transfer learning using adversarial networks* by Shoeleh *et al.* (2020) proposes $DATL_{AN}$, a technique that leverages adversarial domain adaptation principles to discover and transfer related skills between source and target reinforcement learning tasks. Experimental results from several maze environments and from Starcraft: Brood War demonstrate that $DATL_{AN}$ can effectively transfer learned skills from a source task to a target task, leading to increased learning speed and improved final performance compared to a baseline SARSA agent without experience transfer.

The final paper *Learning self-play agents for combinatorial optimization problems* by Xu and Lieberherr (2020) introduces Zermelo Gamification (ZG), a technique that may be used to transform combinatorial optimization problems into two-player finite games. The transformed problem may then be solved by a neural Monte Carlo Tree Search (MCTS) agent learning through the process of self-play. A trained self-play agent may be used to reveal the solution (or show the non-existence of any solution) of the original problem through competitions against itself based on the learned strategy. The authors apply their approach to the highest safe rung problem, and empirical results show that self-play neural MCTS can find optimal strategies for several different ZG-transformed problem instances.

Acknowledgements

The ALA 2019 organisers would like to extend their thanks to all who served as reviewers for the workshop and special issue and to the Cambridge University Press staff and the KER editors Prof. Peter McBurney and Prof. Simon Parsons for facilitating this special issue.

References

- Kazakova, V. A. & Sukthankar, G. R. 2020. Adaptable and stable decentralized task allocation for hierarchical domains. *The Knowledge Engineering Review* **35**.
- Rădulescu, R., Mannion, P., Zhang, Y., Roijers, D. M. & Nowé, A. 2020. A utility-based analysis of equilibria in multi-objective normal form games. *The Knowledge Engineering Review* **35**.
- Shoeleh, F., Yadollahi, M. M. & Asadpour, M. 2020. Domain adaptation-based transfer learning using adversarial networks. *The Knowledge Engineering Review* **35**, e7.
- Xu, R. & Lieberherr, K. 2020. Learning self-play agents for combinatorial optimization problems. *The Knowledge Engineering Review* **35**, e11.