

Special issue on adaptive and learning agents 2017

PATRICK MANNION¹, ANNA HARUTYUNYAN² and KAUSHIK SUBRAMANIAN³

¹*Department of Computer Science & Applied Physics, Galway-Mayo Institute of Technology, Dublin Road, Galway H91 T8NW, Ireland;*

e-mail: patrick.mannion@gmit.ie;

²*DeepMind, 6 Pancras Square, London NIC 4AG, UK;*

e-mail: harutyunyan@google.com;

³*College of Computing, Georgia Institute of Technology, 801 Atlantic Drive, Atlanta, GA 30332, USA;*

e-mail: ksubrama@cc.gatech.edu

1 Introduction

The Adaptive and Learning Agents (ALA) community aims to develop agent-based systems which are autonomous and which employ learning and adaptation to achieve their specified goals. Inspiration for the design of these systems is drawn from diverse fields such as multi-agent systems, game theory, evolutionary computation, multi-objective optimisation, machine learning and cognitive science, to name a few.

While developing single-agent solutions is challenging for many application domains, much of the work presented at the ALA workshop focuses on systems where multiple agents interact in a common environment. Distributed decision making and control is a feature of real-world applications such as air traffic control, autonomous vehicles, electricity generation, multi-robot systems and electronic auctions. For many of these applications, cooperation between agents is essential to achieve good performance; therefore, methods to encourage agent coordination are well-studied within the ALA community.

Improving the learning speed of agents is another common theme which is explored at the workshop. Each additional state feature which is tracked leads to an exponential increase in the number of state-action values which must be learned; this limits the applicability of ALA to large real-world domains. Techniques such as reward shaping, function approximation and integrating (expert) human knowledge have all been successfully applied to mitigate sample complexity in applications of ALA.

This special issue features selected papers from the ALA 2017 workshop, which was held during May 2017 at the International Conference on Autonomous Agents and Multiagent Systems in São Paulo, Brazil. The goals of the ALA workshop 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 five papers, which were selected out of 22 initial submissions to the ALA 2017 workshop. All papers were initially presented at the workshop, before being extended and reviewed again for this special issue. These articles provide a thorough overview of current research directions which are being explored within the ALA community.

The first paper *Leveraging human knowledge in tabular reinforcement learning: a study of human subjects* by Rosenfeld et al. (2018) investigates how different methods for injecting human knowledge into reinforcement learning agents are applied by human designers of varying levels of knowledge and skill. Experiments are conducted in the Simple Robotic Soccer, Pursuit and Mario domains which compare the effectiveness of function approximation, reward shaping and the authors' own method State Action Similarity Solutions (SASS). The empirical results demonstrate that a combination of reward shaping and SASS can achieve good performance in each of the three test domains.

The second paper *Towards life-long adaptive agents: a hybrid planning paradigm for combining domain knowledge with reinforcement learning* by Parashar et al. (2018) considers the problem of task planning for long-living agents situated in dynamic environments. The authors propose a multi-layered agent architecture which uses meta-reasoning to control hierarchical task planning and situated learning, while monitoring expectations generated by a plan against world observations and forming goals and rewards for the reinforcement learner. The authors demonstrate the efficacy of their approach using the Minecraft and Gazebo domains.

The third paper *Language Independent Recommender Agent (LIRA)* by Yucel and Sen (2018) introduces an agent-based approach to creating recommender systems. LIRA constructs agents for each user, which run regression algorithms and build trust relations on texts from different sources. The proposed method is tested in conjunction with several different regression algorithms on the ‘Amazon Reviews’ and ‘GoodReads’ data sets, and is demonstrated to perform well when compared to the RegSVD and HFT algorithms.

The fourth paper *Q-table compression for reinforcement learning* by Rosa Amado and Meneguzzi (2018) proposes a method to reduce the number of entries in a Q-value table by using a deep autoencoder. Multi-agent reinforcement learning where agents share experience updates is also applied to mitigate the large branching factors which are present when controlling teams of units in real-time strategy (RTS) games. The authors conduct a series of experiments in the MicroRTS domain, which demonstrate that their proposed approach performs well against various hand-coded strategies and against Monte Carlo tree search.

Finally, the paper *Reward Shaping for knowledge-based multi-objective multi-agent reinforcement learning* by Mannion et al. (2018) deals with the issue of credit assignment for reinforcement learning in multi-objective multi-agent systems. Two different reward shaping approaches are considered: difference rewards (D) and potential-based reward shaping ($PBRS$). The authors compare the performance of D and $PBRS$ in a benchmark domain, as well as an electricity generator scheduling problem. The paper concludes with a discussion on the merits and limitations of each reward shaping approach in the context of multi-objective multi-agent systems.

Acknowledgement

The ALA 2017 organisers would like to extend their thanks to all who served as reviewers for the workshop, and to the KER editors Prof. Peter McBurney and Prof. Simon Parsons for facilitating this special issue.

References

- Mannion, P., Devlin, S., Duggan, J. & Howley, E. 2018. Reward shaping for knowledge-based multi-objective multi-agent reinforcement learning. *The Knowledge Engineering Review*, X, XX-XX.
- Parashar, P., Goel, A. K., Sheneman, B. & Christensen, H. 2018. Towards life-long adaptive agents: a hybrid planning paradigm for combining domain knowledge with reinforcement learning. *The Knowledge Engineering Review*, X, XX-XX.
- Rosa Amado, L. & Meneguzzi, F. 2018. Q-table compression for reinforcement learning. *The Knowledge Engineering Review*, X, XX-XX.
- Rosenfeld, A., Cohen, M., Taylor, M. & Kraus, S. 2018. Leveraging human knowledge in tabular reinforcement learning: a study of human subjects. *The Knowledge Engineering Review*, X, XX-XX.
- Yucel, O. & Sen, S. 2018. Language independent recommender agent. *The Knowledge Engineering Review*, X, XX-XX.