The Knowledge Engineering Review, Vol. 31:1, 1–2.
doi:10.1017/S0269888915000144

Preface to the special issue: adaptive learning agents

PETER VRANCX[1], ENDA HOWLEY[2] and MATT KNUDSON[3]

[1]*Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium;*
e-mail: pvrancx@vub.ac.be;
[2]*National University of Ireland, Galway, Ireland;*
e-mail: enda.howley@nuigalway.ie;
[3]*CMU – NASA Ames Research Center, USA;*
e-mail: matt.knudson@sv.cmu.edu

1 Introduction

The field of adaptive learning agents studies systems that are capable of acting autonomously and adapting to their surroundings. It encompasses research from disciplines as diverse as artificial intelligence, software engineering, biology, as well as cognitive and social sciences. An agent is said to be an adaptive learning agent when it is capable of changing its behavior in order to react to changes in its environment and can use previous experience to improve its performance with respect to some evaluation measure.

While the development of a single learning agent may already present a serious challenge, current research frontiers also have a large focus on systems where multiple agents interact in a shared environment. Often systems are inherently decentralized, and a centralized, single-agent learning approach is not feasible. This situation may arise because data or control is physically distributed, because multiple, possibly conflicting, objectives should be met, or simply because a single centralized controller requires too many resources. Examples of such systems can be found in multi-robot set-ups, decentralized network routing, distributed load balancing, electronic auctions, traffic control and many others.

In these multiagent settings, agents not only have to deal with a dynamic environment, but also with other agents that act, learn and change over time. When agent objectives are aligned and all agents try to achieve a common goal, coordination among the agents is still required to reach optimal results. When agents have opposing goals, a clear optimal solution may no longer exist and an equilibrium between agents is generally sought. These issues have given rise to an important research track studying coordination mechanisms in multiagent learning, as is also evidenced by several of the papers in this issue.

This special issue contains selected papers from the 2012 Adaptive and Learning Agents (ALA) workshop, held as a satellite workshop at the Autonomous Agents and Multi-Agent Systems conference in Valencia, Spain. The goal of the ALA workshop is to increase awareness and interest in adaptive agent research, encourage collaboration and provide a representative overview of current research in the area of adaptive learning agents. It aims at bringing together not only different areas of computer science (e.g. agent architectures, reinforcement learning and evolutionary algorithms), but also different fields studying similar concepts (e.g. game theory, bio-inspired control and mechanism design). The workshop serves as an interdisciplinary forum for the discussion of ongoing or completed work in adaptive learning agents and multiagent systems.

2 Contents of the issue

This special issue contains six papers, carefully selected out of 40 initial workshop submissions. All papers were presented at the ALA 2012 workshop and were thoroughly reviewed and revised over two separate review rounds. The result provides an excellent overview of the current research directions and state-of-the-art within the adaptive learning agent community.

In the first paper of the special issue, entitled ‘Combining reward shaping and hierarchies for scaling to large multiagent systems’, Chris HolmesParker, Adrian K. Agogino and Kagan Tumer study coordination mechanisms in multiagent systems. They show that combining two well-known coordination mechanisms, hierarchical organization and shaped difference rewards, improves coordination and scalability in large multiagent systems. Their results not only offer real performance improvements for large multiagent systems, but also indicate that studying the circumstances in which existing coordination mechanisms can be combined is a promising future research direction.

In ‘Honesty and deception in populations of selfish, adaptive individuals’, David Catteeuw and Bernard Manderick investigate the requirements for honest signaling to arise in populations of selfish, adaptive agents. By analyzing learning dynamics they show how honest signaling can emerge through adaptive processes and demonstrate the existence of settings where honest signaling is an equilibrium, but not necessarily the final learning outcome.

The third paper of the issue is ‘Overcoming incorrect knowledge in plan-based reward shaping’ by Kyriakos Efthymiadis, Sam Devlin and Daniel Kudenko. In this paper, the authors extend previous research on plan-based reward shaping with a knowledge revision scheme that allows agents to overcome incorrect knowledge used by the higher-level planning layers. They provide a generic method for identifying, verifying and revising the incorrect knowledge and demonstrate how their method can help the learning agent to reach the optimal policy.

‘Plan-based reward shaping for multi-agent reinforcement learning’ by Sam Devlin and Daniel Kudenko studies plan-based reward shaping in multiagent settings. The paper demonstrates that shaping agent rewards using either individual or joint high-level plans can significantly improve performance. In addition, the authors investigate the effects of conflicts in agent plans and show how these problems can be overcome.

Yann-Michaël De Hauwere, Sam Devlin, Daniel Kudenko and Ann Nowé also study reward shaping for multiagent systems in ‘Context-sensitive reward shaping for sparse interaction multi-agent systems’. They introduce context-sensitive shaping functions for multiagent systems with sparse interactions. The authors combine these shaping functions with an algorithm that can detect when explicit coordination between agents is combined and validate their approach in the setting of air traffic control.

The final paper of the issue is ‘A reinforcement learning approach to coordinate exploration with limited communication in continuous action games’ by Abdel Rodríguez, Peter Vrancx and Ann Nowé. In their paper the authors introduce a coordination mechanism for repeated continuous action games. They show how agents can efficiently explore different attractors in their joint action space, without the need to continuously observe the strategies or rewards of other agents. The proposed coordination mechanism allows agents to estimate the basin of attraction for a learning outcome based only on the evolution of their individual reward signal. This allows agents to reduce their strategy space and explore different attractors.