Electronic Transactions on Artificial Intelligence
Vol. 3 (1999), Section B: pp. *23-45*
http://www.ep.liu.se/ej/etai/1999/003/

Refereed Article

Kevin Irwig and Wayne Wobcke:

Multi-Agent Reinforcement Learning with Vicarious Rewards

Authors' affiliations:
KI: Bell Labs Australia, North Ryde NSW, Australia
WW: BT Laboratories, Ipswich, Suffolk, UK.

Abstract: *Reinforcement learning is the problem faced by an agent that must learn behaviour through trial-and-error interactions with a dynamicenvironment. In a multi-agent setting, the problem is often further complicated by the need to take into account the behaviour of otheragents in order to learn to perform effectively. Issues of coordination and cooperation must be addressed; in general, it is not sufficient for each agent to act selfishly in order to arrive at a globally optimal strategy. In this work, we apply the AdaptiveHeuristic Critic (AHC) and Q-learning algorithms to agents in a simple artificial multi-agent domain based on the Tileworld.We experimentally compare the performance of the AHC and Q-learning algorithms to each other as well as to a hand-coded greedy strategy. The overall result is that AHC agents perform better than the others, particularly when many other agents are present or the world is dynamic. We also examine the notion of global optimality in this system, and present a simple method of encouraging agents to learn cooperative behaviour, which we call vicarious reinforcement. The main result of this work is that agents that receive additional vicarious reinforcement perform better than selfish agents, even though the task being performed here is not inherently cooperative.*

Publication and review history

1. This work is originally presented at the Machine Intelligence 16 Workshop, 1-3 December, 1998, and the draft paper is presented to the participants.
2. The present article is *advertised* to the MI peer community on November 9, 1999.
3. The present article is *published* by Linköping University Electronic Press on December 30, 1999, and is permanently available at
   http://www.ep.liu.se/ea/cis/1999/034/
4. The article is refereed during January - May, 2000 by the Organizing Committee of the Machine Intelligence workshops, and is accepted for inclusion in the ETAI.
5. The article appears in Electronic Transactions on Artificial Intellligence, Volume 3 (1999), Section B, pages 23-45. The issue and the annual volume are made permanently available at
   http://www.ep.liu.se/ej/etai/1999/
6. Paper editions of the ETAI issue and the ETAI volume containing this artice are *republished* by the Royal Swedish Academy of Sciences.

The review policy and the quality requirements for acceptance are documented at http://www.etaij.org/info/

Article maintenance

The Electronic Transactions on Artificial Intelligence maintains an Article Interaction Page for this article at
http://www.etaij.org/ra/mi/002/
Besides the publication and review history, it contains links for contacting the author(s), as well as to amendments and other related information for the article. It is intended to keep these links up-to-date in the future.

Copyright conditions