

IEEE Transactions on Computational Intelligence and AI in Games

Special Issue on Deep/Reinforcement Learning and Games

Guest Editors: I-Chen Wu, Chang-Shing Lee, Yuandong Tian, and Martin Müller

Submission Deadline: April 15, 2017

Deep Learning (DL) and Reinforcement Learning (RL) have been applied with great success to many games, including Go and Atari 2600 games. Monte Carlo Tree Search (MCTS), developed in 2006, can be viewed as a kind of online RL. This technique has greatly improved the level of Go-playing programs. MCTS has since become the state of the art for many other games including Hex, Havannah, and General Game Playing, and has found much success in applications as diverse as scheduling, unit commitment problems, and probabilistic planning.

Deep learning has transformed fields such as image and video recognition and speech understanding. In computer games, DL started making its mark in 2014, when teams from University of Edinburgh and Google DeepMind independently applied Deep Convolutional Neural Networks (DCNNs) to the problem of expert move prediction in Go. Clark and Storkey's DCNN achieved a move prediction rate of 44%, exceeding all previously published results. DeepMind's publication followed soon after, with a DCNN that reached 55%.

The combination of DL and RL led to great advances in Atari 2600 game playing, and to the ultimate breakthrough in computer Go. What is the larger impact of these new techniques? For which games do they succeed or fail? How can they be extended to new applications? How can they be combined with other approaches? The purpose of this special issue is to publish high quality papers reporting the latest research covering the theory and practice of DL/RL/DRL methods applied to games. Topics include but are not limited to:

- MCTS and reinforcement learning
- Deep/reinforcement learning for all kinds of games, including board games, card games, video games, general game playing, etc.
- Deep/reinforcement learning for procedural content generation (PCG)
- Deep/reinforcement learning for modeling players/designers
- Deep/reinforcement learning for game analytics
- Deep learning neural net architectures
- Online and offline deep reinforcement learning methods
- Training and testing issues for deep learning, such as transfer learning, dropout, regularization to avoid overfitting, adaptive learning rates, momentum, selection of training data sets, etc.
- Approximation methods for deep learning
- Hybrid deep learning approaches
- Real world applications
- DL-based knowledge representation models for games

Authors should follow normal TCIAIG guidelines for their submissions, but clearly identify their papers for this special issue during the submission process. See <http://www.ieee-cis.org/pubs/tciaig/> for author information. Extended versions of previously published conference papers are welcome, provided that the journal paper is a significant extension, and is accompanied by a cover letter explaining the additional contribution. Short papers or correspondences describing novel experimental results are also welcome. The deadline for submissions is **April 15, 2017**.