
BORA YONGACOGLU, Queen's University

Using Reinforcement Learning to Solve a Decentralized Control Problem

Reinforcement learning algorithms have been used to solve single agent Markov decision processes even when the agent does not have access to the transition probabilities or the cost function. The availability of feedback, in the form of short term cost realizations and one-step transition observations, allows the agent to learn the environment. We consider a generalization of this problem, in which there are multiple decision makers with the same interests but different local information. Learning in games is considerably more difficult than learning in a stationary environment, as the presence of other agents who are learning and changing their behaviour leads to unreliable feedback. We present an algorithm that converges to optimal behaviour.