Q-Table compression for reinforcement learning

Leonardo Amado, Felipe Meneguzzi
(PUC-RS, Brazil)

Reinforcement learning algorithms are often used to compute agents capable of acting in environments with no prior knowledge. However, these algorithms struggle to converge in environments with large branching factors and their large resulting state-spaces. In this paper, we develop an approach to compress the number of entries in a Q-value table using a deep auto-encoder that uses a binary representation of the state, compressing it to a smaller representation. Our approach focuses specifically in mitigating the large branching factor problem resulting from dealing with multi-agent scenarios.
We apply these techniques in the scenario of the MicroRTS Real-Time Strategy (RTS) game, where both state space and branching factor are a problem. To reduce the branching factor on the MicroRTS game, we use separate instances of Q-Learning for each agent, and we separate agents by roles. We combine both the auto-encoder and the separate instances of Q-Learning, creating an AI capable of playing the MicroRTS game.
We empirically evaluate an implementation of the technique to control agents in an RTS game scenario where classical reinforcement learning fails. Our results shows that our approach is able to compete against pre-coded AIs and other works in the literature.