

Bill Ray
Learning to Play Hammurabi

Two types of machine learning algorithms, along with one heuristic-based approach were used to play the game, Hammurabi. Two reinforcement learning-based algorithms and two genetic algorithms attempted to play Hammurabi. The reinforcement learning algorithms used were Q-Learning and Policy Gradients. The genetic algorithms used were Neuroevolution of augmenting topologies (NEAT) and Bill Ray's Artificially Intelligent Networks (BRAIN). Hammurabi's large state and actions spaces presented a challenge for the reinforcement learning algorithms and it would have been interesting if the reinforcement learning solutions could have been modified to overcome this challenge and perform well. The three different types of algorithms were compared, and the genetic algorithms outperformed the reinforcement learning algorithms and the heuristic-based approach.