Abstract

This dissertation applies human intervention and guidance to modern reinforcement learning environments. Using Open AIs' Retro application, playing the classic Super Mario Bros, the concept of Guided Learning is presented where a reinforcement learning agent can effectively 'ask for help' as it encounters stagnation. At which point a human supervisor can effectively 'guide' the agent as to how to progress beyond the point of stagnation. This guidance is then encoded as a new, separately trained neural network known as a 'Taught Response Memory' that can be recalled when another 'similar' situation arrises in the future. This dissertation applies Guided Learning on top of an evolutionary algorithm but also shows how Guided Learning is algorithm independent and can be applied in any reinforcement learning context where human intervention is applicable. The results show that the initial Guided Learning implementation gives superior rate-of-progression in both the best case and worst case scenarios and yields, on average, an increase of 136% in the rate of progression of the most fit genome with minimal human intervention. This is due to the significant reduction at which the agent must explore the solution space by effectively giving the agent more information to exploit. The results obtained show good promise for Guided Learnings potential as such results were obtained with only a partial implementation and much future work still remains.

Keywords