



Carnegie Mellon University

Computer Science Department

Speaking Skills Talk

How AI Has Learned to Play Atari Games

Steven Osman

Monday, November 12, 2018

12:00 pm

GHC 6501

Advances in deep learning have enabled many new applications over the past few years including solutions to such challenges as image classification, machine translation and speech recognition. A particular form of deep reinforcement learning, deep Q-learning has allowed an AI to learn to play Atari games through trial and error. Reinforcement learning is a form of machine learning in which, rather than being given a series of examples to learn from (an approach known as supervised learning), a model is trained by repeatedly refining the model in an attempt to maximize a reward. This reward could be positive, such as a point given each time a task is successfully completed, or negative, such as points taken away for failure or for consuming too many resources. Q-learning is an approach to teaching an agent to solve a problem by picking the best action given its current state, for example, deciding what move to make given a chess board in a particular configuration.

This talk will explain the first end to end approach to learning to play Atari games which used an algorithm called Deep Q-learning with Experience Replay (Mnih et al. 2013). First, a brief introduction to Q-learning and deep reinforcement learning will be provided, and then an explanation of how they were combined with a form of memory (the experience replay buffer) to solve the complex task of playing games. This algorithm learns to play seven Atari 2600 games with no knowledge of the internals of the games themselves. The input is just the pixels on the screen and the current score, and the output is what move to make at each moment in time.