1135-VN-2479 **Sarah Teichman** (steichman18@amherst.edu), **Jing Li*** (gracelj@hku.hk), **Quentin Dubroff** and **Xiaoyu Xie**. *Catching a Falling Ball with Reinforcement Learning.* Preliminary report.

Reinforcement learning is an area of machine learning concerned with teaching an agent how to act so as to maximize a reward within a given environment. In this project, we compared two deep reinforcement learning methods, deep Q-learning and policy based learning, with shallow reinforcement learning, a proposed alternative to these methods that avoids the complicated network architecture of a deep neural network. We applied these methods to design a control strategy for a simple game, in which a ball falls diagonally and a paddle attempts to catch it. We compared the three methods to each other over several combinations of game parameters. We found that with a standard optimizer, Q-learning and policy based learning performed similarly for small game parameters, while shallow learning took more computation to reach a lower performance. For larger game parameters, Q-learning significantly outperformed policy based learning when they both used a standard optimizer, but when we implemented policy based learning with the recently developed Adam optimizer, their performance was comparable. We also visualized the the convolutional filters for our deep learning methods and found that the first layer of these networks responded heavily to the direction of the falling ball. (Received September 26, 2017)