

## Abstract

Reinforcement learning attempts to train an *agent* to interact with their *environment* so as to maximize its expected future *reward*. This framework has successfully provided solutions to a variety of difficult problems. Recent advances in deep learning, a form of supervised learning with automatic feature extraction, have been a significant factor in modern reinforcement learning successes. We use the combination of deep learning and reinforcement learning, deep reinforcement learning, to address the portfolio management problem, in which an agent attempts to maximize its cumulative wealth spread over a set of assets. We apply Deep Deterministic Policy Gradient, a continuous control reinforcement learning algorithm, and introduce modifications based on auxiliary learning tasks and  $n - step$  rollouts. Further, we demonstrate its success on the learning task as compared to several standard benchmark online portfolio management algorithms.