Nicolas
Tobler

Student	Nicolas Tobler
Examiner	Prof. Dr. Guido Schuster
Subject Area	Software and Systems

Reinforcement Learning

Theory and Experiments

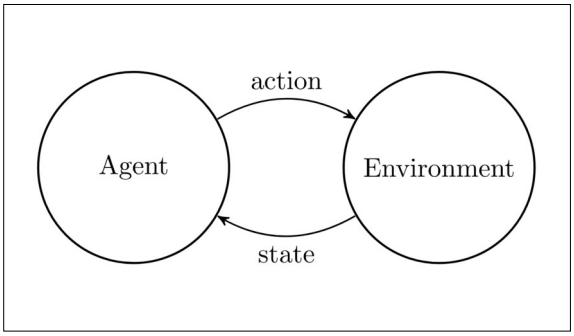

Basic idea of reinforcement learning
Own presentment

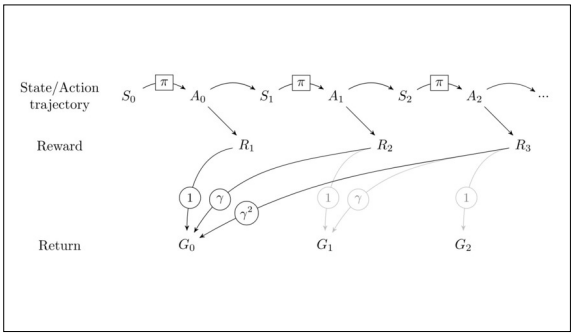

Illustration of a state action trajectory emitting rewards which accumulate to a return
Own presentment

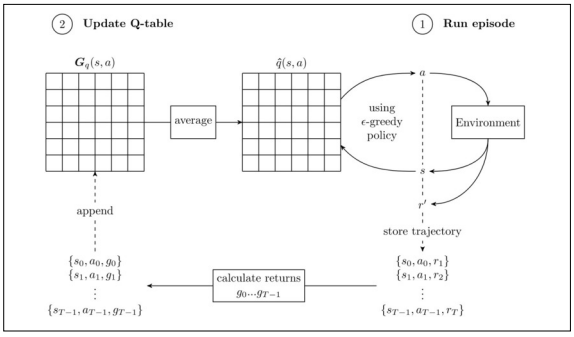

Illustration of the tabular Monte Carlo Control algorithm
Own presentment

Introduction: Reinforcement learning is a type of machine learning that learns from experience. This experience is gathered by interaction of a reinforcement learning agent with an environment. If the agent performs good actions and reaches a goal, it is rewarded by the environment and can improve its strategy.

In the last few years, a lot of progress in the field of reinforcement learning has been achieved. Reinforcement learning methods have reached superhuman performance in board games, such as chess and Go, as well as in Atari games. However until now, only few real world applications have been found.

The objective of this work is to examine the theory of reinforcement learning to gain a profound understanding of it. In addition, different state-of-the-art algorithms should be implemented and tested. Using the insights gained throughout this work, it should be possible to estimate the usability of reinforcement learning.

Approach: Several types of reinforcement learning algorithms have been implemented and tested on various environments. These environments include own creations and online published environments which commonly serve as benchmarks. Our own environments include an intra-logistics optimization problem, as well as the Swiss card game Jass.

Conclusion: It turned out to be surprisingly difficult to apply the reinforcement learning theory. Even the effort needed to reproduce the results on the games is immense.

Reinforcement learning needs a lot of computational power and has a lot of different hyperparameters which need fine-tuning. It has not been possible to achieve a reasonable performance on real world problems. The examined algorithms have learned reasonable behaviors in most situations but failed in others. Therefore, these algorithms may be better suited for control problems, where wrong decisions can be corrected.

This thesis is written as an introduction to reinforcement learning, containing the theory and derivations of the most common algorithms. Working through the theory helps to understand the ideas of reinforcement learning.