EC 400: Introduction to Reinforcement Learning
Fall 2021

Basic Information

Instructor: Alex Olshevsky
Contact: alexols@bu.edu
Time: MW 4:30-6:15

Course Description

Reinforcement learning is a subfield of artificial intelligence which deals with learning from repeated interactions with an environment. Reinforcement learning is the basis for state-of-the-art algorithms for playing strategy games such as Chess, Go, Backgammon, and Starcraft, as well as a number of problems throughout robotics and various fields of engineering and computer science. In this course, we will study the fundamental algorithms of reinforcement learning. Our goal will be to understand the mathematical principles underlying these algorithms and implement them on several popular benchmark problems and games.

(a) A robot hand solves a Rubik's cube.

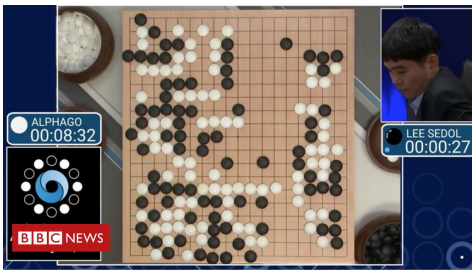

(b) AlphaGo plays the (human) world champion.

(c) Spectators watch AlphaStar play Starcraft.

(d) A quadruped navigates a rocky terrain.

Figure 1: Four systems built with reinforcement learning (subfigure captions are clickable links).

Prerequisites: Probability (EK 381) and Linear Algebra (EK 103) or equivalents.

Topics:

1. A brief review of probability and linear algebra.
2. Introduction to reinforcement learning: value iteration and policy iteration.
3. Reinforcement learning from samples: temporal difference learning and Q-learning.
4. Approximation in reinforcement learning. Linear approximations and off-policy methods.
5. Using neural networks for approximation. Expressivity, training, convergence results.
6. Policy gradient methods. Playing games with Monte Carlo tree search.

Tentative plan: The current plan is to have five homework assignments, three quizzes, and a final project. There will be no final exam.

The homeworks will contain a sequence of coding exercises culminating in the final project. These will start from elementary RL problems and will progress to writing software to play simple games at human-comparable level. The current plan is to use SuperTuxCart, which is an open source implementation of a variant of Mario Kart.

Figure 2: SuperTuxKart can be controlled by pointing in the desired direction of motion. One homework will ask you to train an algorithm which determines this direction to complete a sequence of races.

The final project will likely involve a competition among teams to write agents which play "Pacman capture the flag," a two-player variant where each player must capture the food on the opponent's side of the map.

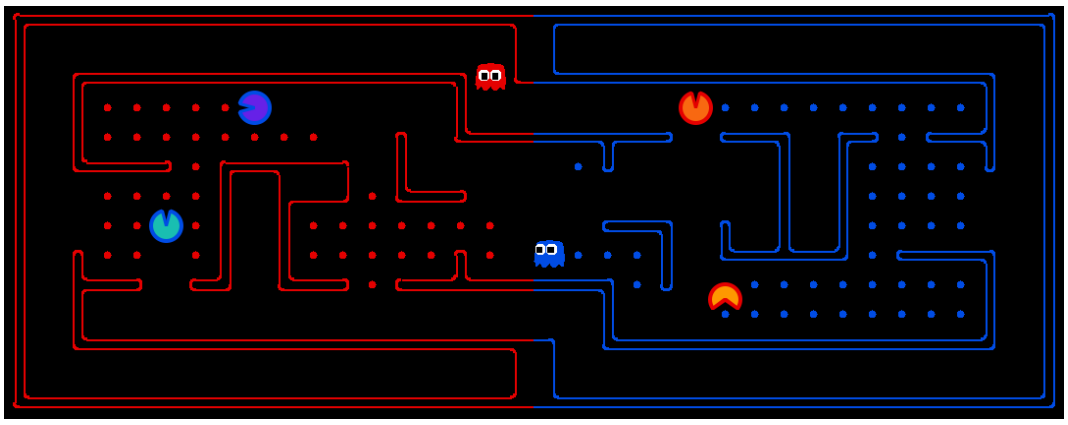

Figure 3: Screenshot from "Pacman capture the flag." The rules can be found here.