CMPS 4660/6660 Reinforcement Learning – Fall 20

Lab 1 (10 points)

Due 10/06/20 before class

This lab should be implemented in Python.

In this lab, you will solve the Gambler's problem in Barto and Sutton's book (Example 4.3) using dynamic programming.

"A gambler has the opportunity to make bets on the outcomes of a sequence of coin flips. If the coin comes up heads, he wins as many dollars as he has staked on that flip; if it is tails, he loses his stake. The game ends when the gambler wins by reaching his goal of \$100, or loses by running out of money. On each flip, the gambler must decide what portion of his capital to stake, in integer numbers of dollars. This problem can be formulated as an undiscounted, episodic, finite MDP. The state is the gambler's capital, $s \in \{1, 2, ..., 99\}$ and the actions are stakes, $a \in \{0, 1, ..., \min(s, 100 - s)\}$. The reward is zero on all transitions except those on which the gambler reaches his goal, when it is $+1$. The state-value function then gives the probability of winning from each state. A policy is a mapping from levels of capital to stakes. The optimal policy maximizes the probability of reaching the goal. Let p_h denote the probability of the coin coming up heads. If p_h is known, then the entire problem is known and it can be solved, for instance, by value iteration."

Solve the Gambler's problem for $p_h = 0.25$ and $p_h = 0.55$ using **your favourite dynamic programming algorithm** (value iteration, policy iteration, linear programming, etc.) and generate figures similar to Figure 4.3 in Barto and Sutton's book. In programming, you may find it convenient to introduce two dummy states corresponding to termination with capital of 0 and 100, giving them values of 0 and 1 respectively.

You need to turn in either a complete Colab notebook (including the code for generating figures) that is runnable in Google Colab or a zip file including:

1. the source code;
2. figures similar to Figure 4.3 in Barto and Sutton's book.
3. a README file that describes how to compile and run your code and how to generate the figures. **It is your responsibility to make sure that I can compile and run your code.**

References:

1. Google Colab: `colab.research.google.com`
2. Getting Started With Google Colab: `https://towardsdatascience.com/getting-started-with-google-colab-f2fff97f594c`