

2/14/17

580-03, Spring 2017, Programming Project 2

In this homework, you will implement a Monte Carlo method. Showing that this method works will give you 7 points.

I recommend you use Open AI Gym, as discussed in class because Gabe is very familiar with this platform. However, if you would like to use a different framework (e.g., BURLAP, RL-Glue, RL-PI, etc.) you can. We just won't be able to help as much.

Recall that homework is graded on a 10-point scale. Here is a guide to how we will grade your homework. To get a 7/10, you only need to successfully implement a MC method and show that it works.

Ideas for making your report more interesting, with rough point values.

+1 point: Compare exploring starts with soft policies.

+0.5 point: How does the amount of randomness change the speed of learning?

+1 point: Implement an off-policy method. How does the speed of learning interact with the quality of the policy being followed?

+1 point: Use a different task or modify the task provided to show that MC works in stochastic domains.

+0.5 points: How does changing the rewards and/or the amount of randomness in the domain change the speed of learning?