

1 Project Description

Look at Example 6.6 on page 140 of the second edition of the Sutton and Barto book, *Reinforcement Learning: An Introduction*. The example describes a gridworld scenario where an agent must learn a policy to reach a goal configuration by travelling on the edge of a “cliff”. Please implement an RL agent that can solve this problem using both q-learning and SARSA. You should have two implementations: one version using q-learning and the other using SARSA. Please create a graph similar to that shown in the bottom of Figure 6.5 on page 141. You must work alone. You should do this assignment using python.

The PDF of the book is located here:

[HTTP://INCOMPLETEIDEAS.NET/SUTTON/BOOK/BOOKDRAFT2016SEP.PDF](http://incompleteideas.net/sutton/book/bookdraft2016sep.pdf)

This project is worth an additional 3% on your final grade.

2 Timeline and Deliverables

4/28/2017 EC project is due. The main deliverable is a PDF file that describes in a single paragraph what you did and shows your learning curve similar to that shown in the bottom of Figure 6.5. You must also submit your python code used to create the agent and plot the graph. Please submit a single zip file that contains both the code and the PDF file.