


Here you will find replicated learning rate experiments for the mountain car example from this chapter in the book. These plots compare how fast the linear gradient decent SARSA could solve this problem with various value for λ and α . As in the book for a single experiment we compute the steps per episode, after performing 20 episodes, and do this experiment 30 times to average the results. For the case with replacing traces we find



While for the case with accumulating traces we find that

[John Weatherwax](#)

Last modified: Sun May 15 08:46:34 EDT 2005