

CPSC422 Spring 2007

Assignment 3

Due: 2:00pm, Tuesday 27 February 2007.

The aim of this assignment is to learn more about reinforcement learning. You are to write code for the simple game at:

<http://www.cs.ubc.ca/spider/poole/demos/rl/sGame.html>

You are free to use any of the code for that applet.

Please read and post to the bulletin board in the course WebCT site for hints on how to modify the controllers.

Question 1

- (a) Improve on assignment 2: Find a better set of features for the game. You can base your code on the function approximation code posted to WebCT. Your set of features should be statistically significantly better (in terms of how quickly it learns) than the features in the posted solution to last assignment. There will be bonus marks for the student whose set of features performs the best on average.
- (b) Explain:
 - i) How do the state-based features that don't depend on the action help (or not) in the posted solution (as the action doesn't depend on such features)?
 - ii) Does it matter if one feature is multiplied by a constant, or if some feature is duplicated?
 - iii) Compare what happens the features have values $\{0, 1\}$ as opposed to $\{-1, 1\}$.

Question 2

In the simple game, change the Q-learning applet so that an adversary gets to choose the location of the treasure (not the probability of it appearing, but which corner it appears in). You can assume that the adversary gets to observe the state and the agents action (but not the square the agent ends up at).

Question 3

For the hawk-dove game:

		Agent 2	
		dove	hawk
Agent 1	dove	$R/2, R/2$	$0, R$
	hawk	$R, 0$	$-D, -D$

where $D > 0$ and $R > 0$, and each agent is trying to maximize its utility. Is there a Nash equilibrium with a randomized strategy? What are the probabilities? What is the expected payoff to each agent? (These should be expressed as functions of R and D). Show your calculation.

Question 4

[Note that this question is worth marks, so don't forget to do it.]

- For each question and each part in this assignment, say how long you spent on it. Was this reasonable? What did you learn?
- Tell us every person you discussed this with (including the TA and fellow students), and every external source (e.g., web site, research paper, book) you referenced to do this assignment.