

Artificial Intelligence - Assignment 4

Posted Friday, November 14, 2003, 2:30pm

Due Tuesday, November 25, 2003, in class

1. [15 points] **Expected value**

Suppose you have to take an exam with ten multiple choice questions, in which each question has five possible answers. We will assume that the questions are independent of each other.

- (a) Suppose that you are given 1 point for answering a question correctly, and 0 points for a wrong answer. You have no clue about the questions, so you pick answers randomly. What is your expected score?
- (b) Suppose that your professor has a nutty grading scheme, in which you get -0.5 for a wrong answer. However, if you admit your ignorance and do not fill out any answer, you get 0 points. What is your expected score for random guessing now? What would be the best choice of action?
- (c) Suppose now that you are not completely ignorant, and you can rule out, for any question, k answers that are obviously wrong. However, between the remaining choices, you still have to guess. Write down the expected scores for $k = 1, 2, 3$. What is the best thing to do in each case?

2. [15 points] **Utility**

Russell & Norvig, pg. 610, Problem 16.3

3. [20 points] **Markov Decision Processes**

Consider the problem of selling a house. An offer of amount w_j comes at the beginning of each day, drawn from a probability distribution p_j . Any offer not accepted right away can be accepted at a later date. If the house remains unsold for one day, a maintenance cost c is incurred. There is a deadline to sell the house in N days. The objective is to maximize the price for which the house is sold, minus the maintenance cost. Model this problem as a Markov decision process. Specify the states, action, rewards and transition probabilities.

Extra credit (10 points): can you describe what the optimal policy looks like?

4. [50 points] **Programming questions**

Implement Monte Carlo search for the game of Knockabout. You should use a heuristic of your choice to pick moves for yourself, and for the opponent. Write a little report, explaining your heuristic, how many moves you generate, and how many different trajectories you average. Play your Monte Carlo player against the random player 10 times, as well as against a player that uses your heuristic greedily, and report the results. If you have a working expectiminimax player, you can also test against it.

Make sure you download the new version of the code for this assignment, because there have been significant updates!

You should submit your code and little report using handin.