

Extra Practice for Final

Reinforcement Learning Problem

You are designing an AI system to play a game involving several slot machines. Each turn, the AI must play the slot machine it is at and then it must take an action that moves it to an adjacent location. (Note that, in this problem, the agent is not allowed to repeat a slot machine two times in a row.) The probability distribution for each slot machine is known (see below), and the starting location of the agent is at the left-most slot machine. Each turn, the agent loses \$1 as the cost of playing the slot machine. (Note that you are allowed to have a negative score, here.)

Slot 1	Slot 2	Slot 3	Slot 4
\$0 – 50%	\$0-80%	\$0-60%	\$0-99%
\$1 – 30%	\$5 – 20%	\$1-30%	\$150-1%
\$2 – 20%		\$4-10%	

Create and fill out a q-learning table (with three rounds) for this game, using the same approach as we used in class. (Note that, to find the immediate reward for a state, you can sum up the rewards multiplied by their corresponding probabilities.)

Markov Chain/Hidden Markov Model Problem

Suppose that the menu of the Rat each day differs probabilistically, based on the Markov Model properties. If scones are served on a given day, the probability that they'll be served the next day is 85%. If scones are not served, the probability that they'll be served on the next day is 60%.

- A. Given this information, and given that on Monday scones were served at the Rat, what is the probability that scones will be served on Wednesday? What is the probability that scones are served every weekday (Monday through Friday) this week?

- B. Now, let's expand the problem a little: If scones are served, the probability of cookies also being served is 5%. If scones are not served, the probability of cookies being served is 80%.

- C. If you see people on campus eating cookies on Wednesday and Thursday (and not on Tuesday or Friday), what is the likelihood that scones were served in the cafeteria on those days, using the forward algorithm?

- D. Using the previous problem, how does applying the forward-backward algorithm modify the observed probability of scones for Tuesday?

Neural Networks problem

In this problem, we are required to create a perceptron that takes in three inputs (with int values ranging from 0 to 10) and returns 1 if the inputs add up to 15 (or more).

Using the initial weights of -1, 2, 3, and 4 and the training samples provided below, what are the weights for the neural network after two epochs? (For this problem, assume we're using a threshold activation function, where the value returned is 1 if the inputs * the weights add up to 0 or more.)

Training Samples

X1	X2	X3	Desired output
8	6	0	0
4	4	9	1
0	8	8	1
5	0	7	0

You should be prepared for problems relating to the first half of the semester, including:

- Bayes nets and probabilistic reasoning
- Dijkstra's algorithm, A* algorithm, greedy best-first search
- Minimax (with both alpha/beta pruning and with heuristics)
- Which approach that we've talked about so far is best for a specific context and why
- Heuristic design (including consistent/admissible), problem setup, and discussion about the state space.

You should also be prepared for problems we've worked on since the first exam:

- Statistical inference and naïve Bayes classifiers
- Markov Chains and Hidden Markov Models
- Reinforcement Learning strategies, including q-learning and v-learning algorithms
- Neural networks

This exam will be comprehensive-I recommend working through the problems from this guide and from the previous review guide on the course website. The homeworks (and homework solutions) should also be very helpful in preparing for this exam.

For this exam, you are allowed two pages of handwritten notes (front and back). Calculators are not allowed for this exam.