CSE 573 Au 24 Written Assignment

Name:

This assignment is take home and is due on **Thursday December 5th at 11:59 pm**. Please submit directly to Gradescope. This assignment should not take significantly longer than 3 hours to complete if you have already carefully studied all of course material. Studying while doing the assignment may take longer. :)

This homework is open book and open notes, but you must complete all of the work yourself with no help from others. Please feel free to post clarification questions to the course message board, but please do not discuss solutions.

Partial Credit: If you show your work and *briefly* describe your approach to the longer questions, we will happily give partially credit, where possible. We reserve the right to take off points for overly long answers. Please do not just write everything you can think of for each problem.

Scores						
Q.1 (45)	**Q.2 (40)**	**Q.3 (30)**	**Q.4 (30)**	**Q.5 (20)**	**Q.6 (35)**	**Total (200)**

Question 1 – True/False – 45 points

Circle the correct answer for each True / False question.

1. True / False – A* Tree Search requires a consistent heuristic for optimality. (3 pt)

2. True / False – Minimax is optimal against perfect opponents. (3 pt)

3. True / False – There exist problems for which an admissible heuristic cannot be found. (3 pt)

4. True / False – Uniform cost search with costs of 1 for all transitions is the same as depth first search. (3 pt)

5. True / False – Policy Iteration always finds the optimal policy, when run to convergence. (3 pt)

6. True / False – Higher values for the discount (γ) will, in general, cause value iteration to converge more slowly. (3pt)

7. True / False – For MDPs, adapting the policy to depend on the previous state, in addition to the current state, can lead to higher expected reward. (3pt)

8. True / False – Graph search can sometimes expand more nodes than tree search. (3pt)

9. True / False – When using Naive Bayes with Laplace smoothing, if the training error is low but validation error is much higher, we should decrease the value of smoothing strength k. (3 pt)

10. True / False – When using features to represent the Q-function it is guaranteed that this feature-based Q-learning finds the same Q-function, Q^*, as would be found when using a tabular representation for the Q-function. (3 pt)

11. True / False – Given no independence assumptions, $P(A|B,C) = \frac{P(B|A,C)P(A|C)}{P(B|C)}$. (3 pt)

12. True / False – The Viterbi algorithm for HMMs has polynomial time complexity in the number of states. (3 pt)

13. True / False – The number of parameters in a Bayesian network grows exponentially with the highest in-degree (number of parents) of a node in the network. (3 pt)

14. True / False – When solving a HMM with states S and evidence E, it is possible for variable elimination and the forward algorithm to reach different solutions for $P(S_t|E_{1:t})$ for some timestep t. (3 pt)

15. True / False – The initial random assignment of data points in the K-means clustering algorithm does not influence the algorithm's outcomes. The resulting clusters will be the same every time the algorithm is executed. (3 pt)

Question 2 – Short Answer – 40 points

These short answer questions can be answered with a few sentences each. Please be brief, we will subtract points for very long responses (e.g. more than a sentence or two for each part of the question).

1. Short Answer – Briefly describe the relationship between admissible and consistent heuristics. When would you use each, and why? (5 pts)

2. Short Answer – Briefly describe the difference between UCS and A* search. When would you prefer to use each, and why? (5 pts)

3. Short Answer – Briefly describe the difference between value iteration and policy iteration. Describe conditions under which one algorithm might be preferred to the other, in practice. (5 pts)

4. Short Answer – Describe the naive Bayes bag-of-words model for document classification. Draw the Bayes net graph, annotate the class label node and the feature nodes, describe what the domain of the features is, describe any properties of the conditional probability tables that are specific to the bag-of-words (and not necessarily true in all naive Bayes models). For simplicity it is OK to assume that every document in consideration has exactly N words and that words come from a dictionary of D words. (5 pts)

5. Short Answer – Briefly describe how you would decide which algorithm to use for answering queries to a Bayesian network. What is the key property of the network that, if known, would best help you make the appropriate decision. (5 pts)

6. Short Answer – In machine learning, explain generalization and over-fitting. Describe an experimental setup that correctly measures generalization. Assume that your algorithm has one hyperparameter that must be set. (5 pts)

7. Short Answer – Briefly describe when you would prefer to report precision and recall for a learned classifier, instead of accuracy. (5 its)

8. Short Answer – Briefly describe the difference between model-free reinforcement learning and model-based reinforcement learning (5 its)

Question 3 – Game Trees – 30 points

Consider the following game tree, which has min (down triangle), max (up triangle), and expectation (circle) nodes:

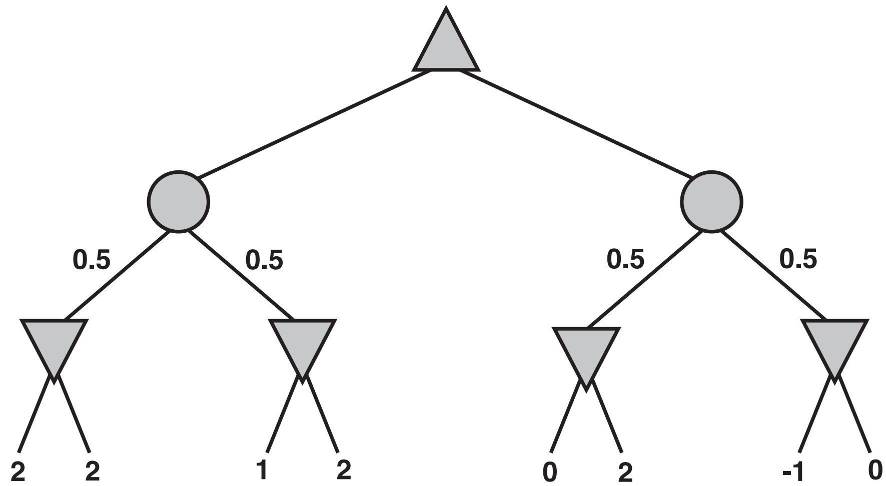

1. In the figure above, label each tree node with its value (a real number). [7 pts]

2. In the figure above, circle the edge associated with the optimal action at each choice point. [7 pts]

3. If we knew the values of the first six leaves (from left), would we need to evaluate the seventh and eighth leaves? Why or why not? [5 pts]

4. Suppose the values of leaf nodes are known to be in the range $[-2, 2]$, inclusive. Assume that we evaluate the nodes from left to right in a depth first manner. Can we now avoid expanding the whole tree? If so, why? Circle all of the nodes that would need to be evaluated (include them all if necessary). [11 pts]

Question 4 – MDP Exploration – 30 points

Pacman is now a CS student at UW. He finds himself in a (very) simplified, *deterministic* grid world MDP representation of UW depicted below, starting at state S. Pacman can take actions up, down, left or right. If an action moves him into a wall, he will stay in the same state. States C and P represent the CS building and the Party building respectively (their labels both appear above the relevant grid square). Pacman will study at the CS building or he will party at the Party building. At states C and P Pacman can take the exit action to receive the indicated reward and enter the terminal state, E. $R(s, a, s') = 0$ otherwise. Once in the terminal state the game is over and no actions can be taken. Let the discount factor $\gamma = \frac{1}{2}$ for this problem, unless otherwise specified.

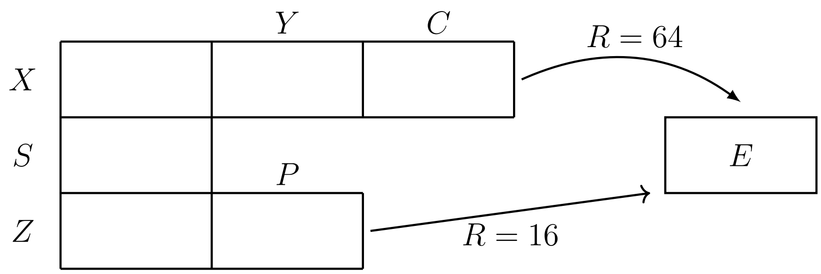

1. What is the optimal policy for Pacman on the grid world above? [5 pts]

2. Given an arbirtrary initial policy, what is the maximum number of iterations k it will take before $V^{\pi_k}(S) = V^{\pi^*}(S)$? [5 pts]

3. What are all the values that $V^{\pi_k}(S)$ will take on during the entire process of Policy Iteration given every possible initial policy. [5 pts]

4. Let us mess with Pacman's ability to study. Your task is to change some of the MDP parameters so that Pacman no longer desires to visit the CS building. S is where Pacman starts (the square to the right of the label S). All subquestions are independent of each other so consider each change on its own.

   (a) What discount factor forces Pacman to be indifferent between studying and partying given that he starts at state S? [5 pts]

   (b) Tweak the reward function such that Pacman will always choose partying over studying. Write a bound on $R(C, exit, E)$, that guarantees Pacman exits from P instead of C. [5 pts]

   (c) Let us make the reward of studying so distant that Pacman no longer exits from C. We'll accomplish this by adding a certain number of grid positions, x of them, in between Y and C as depicted below. Give a lower bound for x that **guarantees** Pacman does not exit from C. [5 pts]

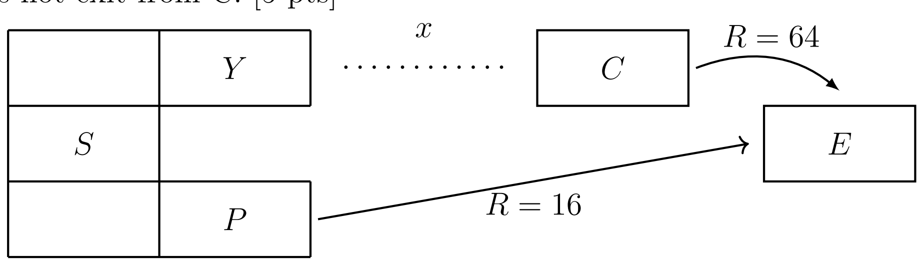

Question 5 – Perceptron – 20 points

We would like to use a perceptron to train a classifier for datasets with 2 features per point and labels +1 or -1. Consider the following labeled training data:

Features (x_1, x_2)	Label y^*
(-1,2)	1
(3,-1)	-1
(1,2)	-1
(3,1)	1

1. Our two perceptron weights have been initialized to $w_1 = 2$ and $w_2 = -2$. After processing the first point with the perceptron algorithm, what will be the updated values for these weights? [10 pts]

2. After how many steps will the perceptron algorithm converge? Write "never" if it will never converge. Note: one steps means processing one point. Points are processed in order and then repeated, until convergence. [10 pts]

Question 6 – Bayesian Networks – 35 points

Consider the following two Bayesian networks, which are variations on the alarm network we discussed in class:

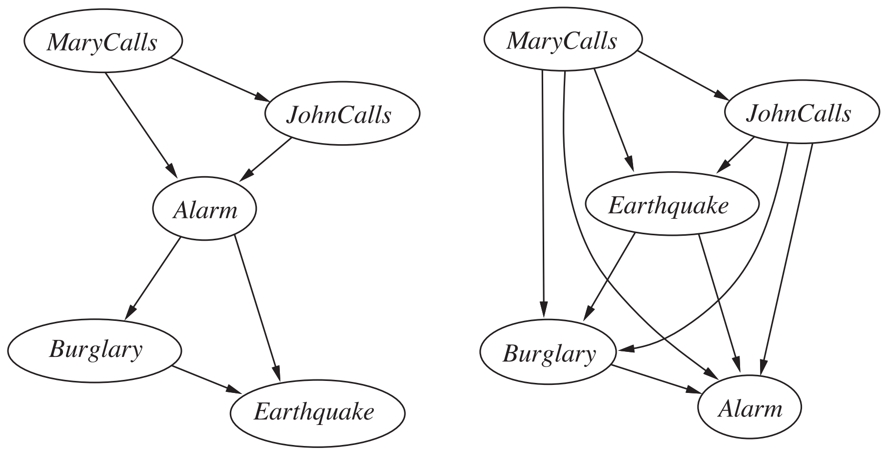

(a) (b)

1. Based on the network structure alone, which network above makes the most independence assumptions? [5 pts]

2. Draw a new Bayesian network with the same set of random variables that makes as many independence assumptions as possible. [5 pts]

3. Write down two conditional independence assumptions encoded by the structure of network (a). If there are not two, write as many as possible. [5 pts]

4. Write down two conditional independence assumptions encoded by the structure of network (b). If there are not two, write as many as possible. [5 pts]

5. If the edge between MaryCalls and Earthquake is removed from network (b), will the class of joint probability distributions that can be represented by the resulting Bayesian network be smaller or larger than that associated with the original network? Briefly explain your answer. [5 pts]

6. Simulate the execution of the variable elimination algorithm on network (a) to compute $P(Marycalls|Burglary = true)$. Since we have not given you the CPTs, you do not need to compute the entries. Instead, just list the tables that would be created and eliminated at each step of the computation. Use the most computationally efficient variable ordering. [10 pts]