CSEP 573 Final Exam – Feb 22, 2024

Name:

This exam is take home and is due on **Wednesday March 13th at 5:00 pm**. Please submit directly to Gradescope. This exam should not take significantly longer than 3 hours to complete if you have already carefully studied all of course material. Studying while taking the exam may take longer. :)

This exam 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.

Name: Please do not forget to write your name in the space above!

Scores						
Q.1 (30)	**Q.2 (30)**	**Q.3 (15)**	**Q.4 (15)**	**Q.5 (35)**	**Q.6 (15)**	**Total (140)**

Question 1 – True/False – 30 points

Circle the correct answer for each True / False question.

1. 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)

2. True / False – For answering conditional queries in Bayesian networks, rejection sampling has generally been observed to provide worse estimates that likelihood weighting (when given the same number of samples). (3 pt)

3. 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)

4. True / False – The choice of the variable ordering in variable elimination does not change the correctness of the algorithm (you will always get the correct answer for any ordering). (3 pt)

5. True / False – Inference by enumeration can produce incorrect results if the Bayes network is dense (has many edges). (3 pt)

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

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

8. 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)

9. 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)

10. 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 – 30 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 – 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)

2. 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)

3. 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)

4. Short Answer – Briefly describe a situation in which you would use Bayes rule, and why, from the examples we saw in class. (5 pts)

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

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

Question 3 – Perceptron – 15 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? [5 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 4 – Hidden Markov Models: Tricky Coins – 15 points

Consider the following random process. A magician has two coins, each of which has an unknown type. They can either be fair coins (50/50 odds of heads vs tails), or trick coins that either (1) have heads on both sides or (2) have tails on both sides. A priori, each coin is equally likely to be any of the three possible types.

At every time step, the magician randomly picks a coin (without showing you which one was selected), flips it, and shows you the result. However, unfortunately, the magician only shows you the coin very briefly, and 10% of the time you make a mistake when you observe the true side of the coin (e.g. you see heads when it was actually tails).

1. Model this process as an HMM. Specify all of the necessary parameters. You do not have to write out all of the probability distributions explicitly, but be careful to specify what values they would have if you did the full enumeration. What conditional independences hold in this HMM? [10 pts]

2. Consider the Markov model that would result if you ran the process above and always observed heads. What is the stationary distribution of this model? [5pts]

Question 5 – Bayesian Networks – 35 points

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

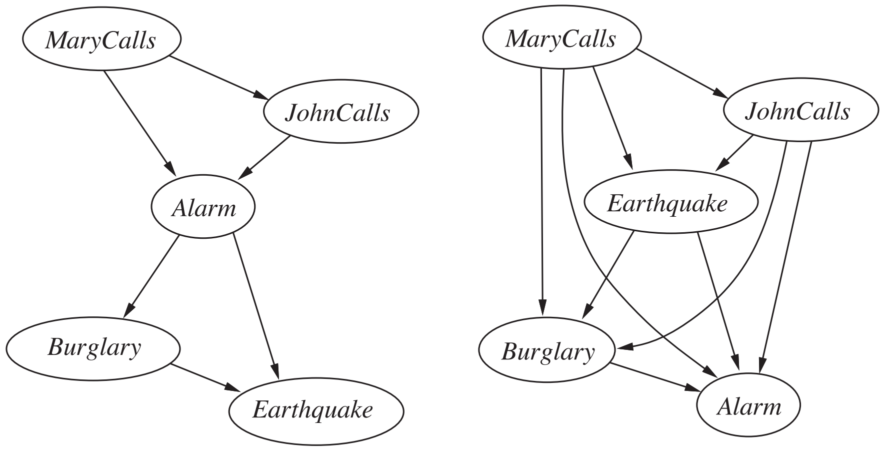

(a) (b)

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]

Question 6 – Reinforcement Learning – 15 points

Consider the grid-world given below and an agent who is trying to learn the optimal policy. States are named as (x-coordinate, y-coordinate) with horizontal axis x and vertical axis y, and the state after exiting is Done. Actions are North, South, East, West, and Exit denoted as N, S, E, W, and X for short. The Exit action can only be taken from shaded states, and Exit is the only action available in the shaded states. Rewards are only awarded for taking the Exit action from one of the shaded states. Taking this action moves the agent to the Done state, and the MDP terminates. Assume $\gamma = 1$ and $\alpha = 0.5$ for all calculations. In Q-Learning, all values are initialized to zero.

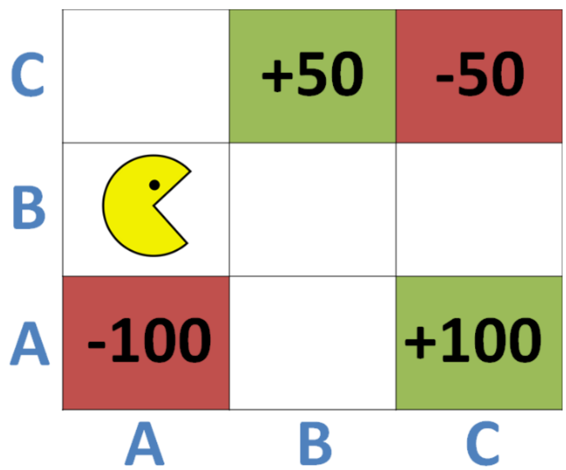

Now, assume the agent starts from (A, B) and observes the following sequence of episodes. Each step is a tuple containing (s, a, s', r).

	Episode 1	**Episode 2**	**Episode 3**	**Episode 4**	**Episode 5**
Step 1	(A,B), N, (A,C), 0	(A,B), E, (B,B), 0	(A,B), E, (B,B), 0	(A,B), E, (B,B), 0	(A,B), E, (B,B), 0
Step 2	(A,C), S, (A,B), 0	(B,B), E, (C,B), 0	(B,B), S, (B,A), 0	(B,B), E, (C,B), 0	(B,B), S, (B,A), 0
Step 3	(A,B), N, (A,C), 0	(C,B), N, (C,C), 0	(B,A), E, (C,A), 0	(C,B), W, (B,B), 0	(B,A), N, (B,B), 0
Step 4	(A,C), E, (B,C), 0	(C,C), X, Done, -50	(C,A), X, Done, +100	(B,B), N, (B,C), 0	(B,B), N, (B,C), 0
Step 5	(B,C), X, Done, +50			(B,C), X, Done, +50	(B,C), X, Done, +50

1. Fill in the following Q-values obtained from direct evaluation from the samples? (10pt)
   (a) $Q((A, B), N) =$

   (b) $Q((B, B), E) =$

2. Which Q values are non-zero after running q-learning with the episodes above? (5pt)