CSE 473 Final Exam – June 3, 2016

Name:

This exam is take home and is due on **Friday June 10th at 5:00 pm**. You can submit directly to GradeScope or to the course staff. 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.

If you show your work and *briefly* describe your approach to the longer questions, we will happily give partially credit, where possible.

There are 10 pages in this exam.

Question 1 – True/False – 30 points

Circle the correct answer for each True / False question. If you think a question is ambiguous, please add a very short explanation of the interpretation you are making, and we will do our best to grade accordingly.

1. True / False – Adding edges to a Bayes Net allows it to encode a bigger space of possible distributions. (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 – Naive Bayes models always encode incorrect independence assumptions. (3 pt)

4. True / False – Approximate Q-Learning with feature-vectors will always converge to the optimal policy. (3 pt)

5. True / False – $P(A|B,C) * P(B|C) * P(C) = P(A,B,C)$, given no independence assumptions. (3 pt)

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

7. True / False – $P(B|A,C) * P(C|A,B)/P(B,C) = P(A|B,C)$, given no independence assumptions(3 pt)

8. 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)

9. True / False – The forward algorithm for HMMs has exponential time complexity in the number of states. (3 pt)

10. True / False – The number of parameters in a Bayesian network grows exponentially with the highest out degree (number of children) of a node in the network. (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 – 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)

2. Short Answer – Describe an experimental setup that correctly measures generalization in machine learning. Assume that you are given a fixed amount of data and your algorithm has one hyperparameter that must be set. (5 pts)

3. Short Answer – Briefly describe the conditions that are required for Q-learning to converge to the optimal policy. (5 pts)

4. Short Answer – Briefly describe a sign of overfitting in Naive Bayes learning, and how it can be avoided. (5 pts)

5. Short Answer – Briefly describe the pros and cons of using the forward algorithm vs. a particle filter for HMMs. When would you use each and why? (5 its)

6. Short Answer – Briefly describe the difference between outcomes and events in joint probability models. (5 its)

Question 3 – Markov Models – 15 points

One of the traffic lights on University Ave has gone faulty, and while it can still display one of three colors (red, yellow, or green) at a time, it no longer does so in the order you would expect. One UW student has figured out that the light transitions according to a Markov model given by the following CPT:

C_t	C_{t+1}	$P(C_{t+1}\|C_t)$
Red	Red	0.8
Red	Yellow	0.1
Red	Green	0.1
Yellow	Red	0.2
Yellow	Yellow	0.3
Yellow	Green	0.5
Green	Red	0.0
Green	Yellow	0.7
Green	Green	0.3

When this fault occurred the light reset, and it followed the initial distribution of $P(Red) = 0.5$, $P(Green) = 0.5$, and $P(Yellow) = 0.0$ (consider this to be $t = 0$).

1. What is the probability that the light produces the sequence Green, Yellow, Green, Yellow, Red? (5 pts)

2. Fill in the values of the probabilities below, using the forward algorithm for Markov models. (10pts)

	$t = 0$	$t = 1$	$t = 2$
$P(C_t = Red)$			
$P(C_t = Yellow)$			
$P(C_t = Green)$			

Question 4 – Hidden Markov Models – 15 points

Continuing from Q3, now imagine that a driver with very poor eyesight is driving up the street and cannot discern the actual color of the light. However, they can make some assumptions based on whether the brightness of the light is bright or dark. Consider the following emission distribution:

C_t	B_t	$P(B_t \mid C_t)$
Red	Bright	0.25
Red	Dark	0.75
Yellow	Bright	0.8
Yellow	Dark	0.2
Green	Bright	0.5
Green	Dark	0.5

For this problem, assume that the driver initially has the belief at time t = 0 that $B(Red) = 0.3$, $B(Yellow) = 0.3$, and $B(Green) = 0.4$ and that the world behaves as an HMM with emissions and transition defined as above.

1. If at time t = 1 we observe B_1 = Bright, and at time t = 2 we observe B_2 = Dark. What is the drivers belief at time t = 2 after considering the most recent observation? (15 pts)

Question 4 – Bayesian Networks – 25 points

Consider the following Bayesian network:

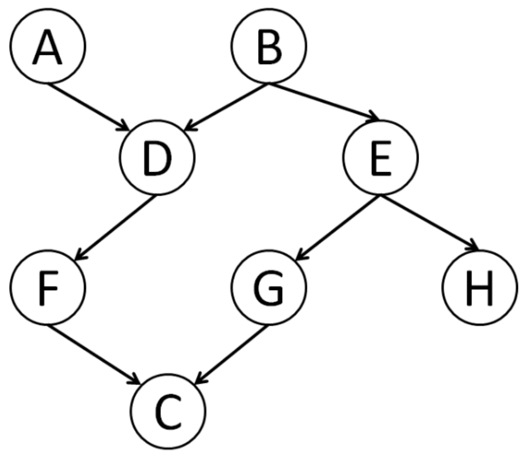

1. Which of the following are guaranteed to be true without making any additional conditional independence assumptions, other than those implied by the graph? (Circle all true statements) [10 pts]:

   (a) $P(E|G, C) = P(E|G)$

   (b) $P(A, B|C) = P(A|C) * P(B|C)$

   (c) $P(B, G|C) = P(B|C) * P(G|C)$

   (d) $P(A|E = e) = P(A)$

   (e) $P(E, C|H) = P(E|H) * P(C|H)$

2. Continuing with the same Bayes Net, lets say we are now interested in finding P(B, G — +h, -c) using variable elimination.

   (a) What factors do we start with after incorporating evidence? (5pt)

(b) If we start by eliminating E, we get a new factor. Write this factor as a function of the necessary variables along with the equation for the factor (hint: the equation should include a sum over the variable you are eliminating): (5pt)

(c) Assume that, near the end of the variable elimination we have only the follow factors remaining: $P(B, G, +h)$, $P(-c|G)$. Write an equation for computing the final query distribution $P(B, G| + h, -c)$ as a function of these factors. (5pt).

Question 5 – 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.

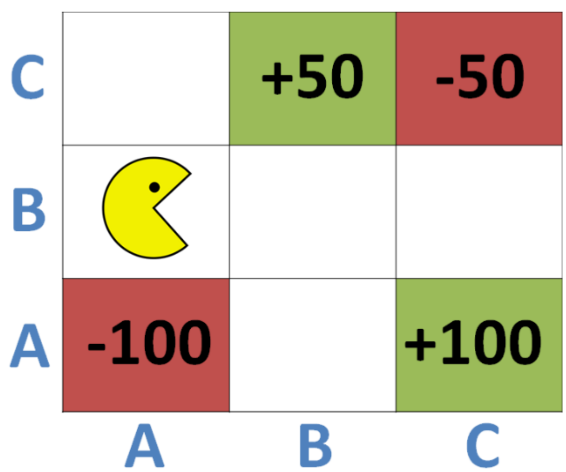

1. What are the following optimal values: (5pt):

   (a) $V^*((A, C)) =$

   (b) $V^*((B, B)) =$

Now, assume the agent starts from (A, B) and observes the following sequence of episodes from runs of the agent through this grid-world. Each line in an Episode 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

2. Fill in the following Q-values obtained from direct evaluation from the samples? (5pt)
   (a) $Q((A, B), N) =$

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

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