CSE 473 Autumn 2018 Homework 6
Due December 5, 2018. Submit a PDF file via Canvas by 11:59 PM.

Unlike worksheets, for which collaboration is encouraged, this is a graded assignment for which collaboration is not allowed. If needed, refer to the academic honesty policy on the course home page, and for this assignment substitute "written text" for the term "code" mentioned there.

Your document may be either typeset (e.g., using Google Docs, Word or Latex) or it can be handwritten on paper and then scanned or clearly photographed. In either case, submit a PDF file that is clearly readable at normal magnification.

---

1. (15 points) **Bayes Nets: Independence and Conditional Independence**

   Consider the Bayes Net graph on the next page, which represents the topology of web-server security model. Here the random variables have the following interpretations:

   V = Vulnerability exists in web-server code or configs.

   C = Complexity to access the server is high. (Passwords, 2-factor auth., etc.)

   S = Server accessibility is high. (Firewall settings, and configs on blocked IPs are permissive).

   A = Attacker is active.

   L = Logging infrastructure is state-of-the-art.

   E = Exposure to vulnerability is high.

   D = Detection of intrusion attempt.

   B = Break-in; the web server is compromised.

   I = Incident response is effective.

   F = Financial losses are high (due to data loss, customer dissatisfaction, etc).

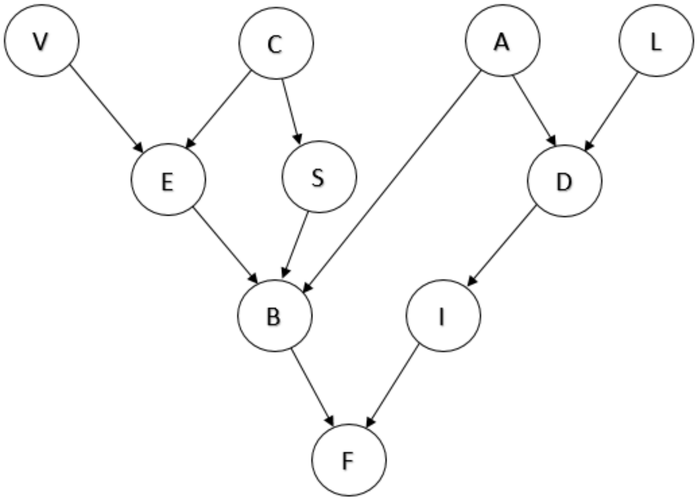

For each of the following statements, indicate whether (True) or not (False) the topology of the net guarantees that that the statement is true. If False, identify a path ("undirected") through which influence propagates between the two random variables being considered. (Be sure that the path follows the D-Separation rules covered in lecture.) The first one is done for you.

(a) $E \perp\!\!\!\perp S$: False (ECS)

(b) $V \perp\!\!\!\perp L \,|\, F$

(c) $V \perp\!\!\!\perp C \,|\, S$

(d) $F \perp\!\!\!\perp A \,|\, B, D$

(e) $V \perp\!\!\!\perp L \,|\, B, I$

(f) $S \perp\!\!\!\perp D \,|\, C, A$

(g) (5 points) Suppose that the company hired an outside expert to examine the system and she determines that E is true: The system is highly exposed to vulnerability. Given this information, your job is to explain to management why getting additional information about S (server accessibility) could have an impact on the probability of V (regarding the existence or non-existence of vulnerabilities). Give your explanation, for the manager of the company, using about between 2 and 10 lines of text, which should be based on what you know about D-separation, applied to this situation. However, your explanation should not use the terminology of D-separation but be in plain English. (You can certainly use words like "influence", "probability", "given", but not "active path", "triple", or even "conditionally independent").

2. (10 points) **Q-Learning**

You are a player in a dangerous, futuristic grid world game. Your goal is simple: to find the jet-pack so you can escape. You are given a map of the world which shows you the location of the jet pack and teleport squares. From any unmarked square, you can move in any direction (N, E, S, W). All moves are deterministic (no noise). If you land in a teleport square, you teleport out of the world (T) and receive a reward of -32 points. If you arrive at the jet-pack without encoutering a teleport square, you receive a reward of $+64$ points.

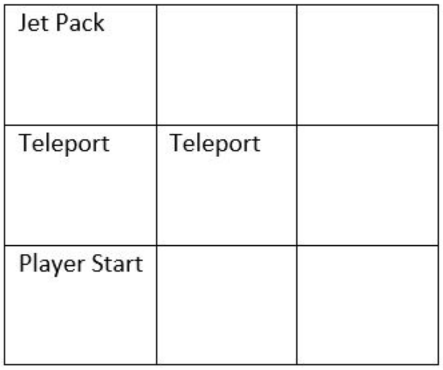

Jet Pack		
Teleport	Teleport	
Player Start		

Every now and then, however, the teleport squares turn off (becoming normal squares). As a result, when you start q-learning you observe the following episodes:

[(0, 0), N, 0, (0, 1), N, 0, (0, 2), T, 64, Game Over]
[(0, 0), N, 0, (0, 1), N, 0, (0, 2), T, 64, Game Over]
[(0, 0), N, 0, (0, 1), E, 0, (1, 1), T, -32, Game Over]

1. What are the q-values after observing these episodes? Assume that all q-values are initially 0 (only need to write the non-zero q-values). Use a learning rate of 1.0.

2. Based on these q-values, what is the best path to reach the jet-pack from the start state?

3. (10 points) **Markov Models**

Below you see the (standard) structure of a Markov Model.

Let's assume this describes the dietary behavior of a Seattlite named "Sean" who loves both seafood and meat for main-courses of dinners (and shuns vegetables, at least as main-course entrees).

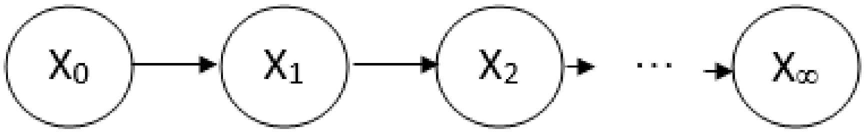

Sean sometimes has seafood for dinner and sometimes it's meat. It's always one or the other. Perhaps Sean has a system for deciding what to eat, but to his partner, Pat, it seems random. Pat has established that Sean tends to have seafood again on day $n+1$ if he had it on day n. This probability is 0.8. (Maybe it is about using up leftovers; however, the behavior still seems random, and there could a run of 3 days of shrimp followed by 4 days of fish, but sometimes, there's seafood for just one meal and it's meat, etc.) When Sean has meat, the probability is only 0.5 that he will stick with meat the next day.

(a) (5 points)
Assuming that Sean has meat at the beginning of a sequence of days (i.e., $X_0 =$ meat, determine the probability vector $\langle P(S), P(M)\rangle$ for the next day (i.e., $P(X_1)$). (Using Pat's model.)
Compute $P(X_2)$.

(b) (5 points) Stationary Distribution. In the long run, the Forward Algorithm should give distributions that gradually converge to an "equilibrium distribution" also known as the stationary distribution. Compute this distribution. In other words, compute $P(X_\infty)$.

4. **(18 points) Hidden Markov Models**

Consider a Hidden Markov Model where the hidden state X_t can be one of three values $\{A, B, C\}$. The transition probabilities are provided in the following table, where the row corresponds to X_{t-1} and the column to X_t.

	A	B	C
A	0.6	0.4	0
B	0.1	0.6	0.3
C	0	0.2	0.8

For example, $P(X_t = A \mid X_{t-1} = B) = 0.1$.

The noisy sensor model for evidence E_t corresponding to X_t gives the true hidden state with probability 0.8, and one of the other two states each with probability 0.1. For example, $P(E_t = B \mid X_t = A) = 0.1$.

(a) (3 points) Assume our belief about the hidden state X_t is

X_t	$P(X_t)$
A	0.5
B	0.5
C	0

Compute the belief about the hidden state X_{t+1} before considering noisy evidence (no need to normalize):

X_{t+1}	$P(X_{t+1})$
A	
B	
C	

(b) (3 points) Given your answer from the previous question, now assume we have the noisy sensor reading $E_{t+1} = C$. Compute our posterior belief taking this evidence into account (no need to normalize):

X_{t+1}	$P(X_{t+1})$
A	
B	
C	

(c) (3 points) Assume now we are using a particle filter with 3 particles to approximate our belief instead of using exact inference. Imagine we have just applied transition model sampling (elapse-time) from state X_t to X_{t+1}, and now have the set of particles $\{A, A, B\}$. What is our belief about X_{t+1} before considering noisy evidence?

X_{t+1}	$P(X_{t+1})$
A	
B	
C	

(d) (6 points) Now assume we receive sensor evidence $E_{t+1} = B$. What is the weight for each particle, and what is our belief now about X_{t+1} (before weighted resampling)?

Particle	Weight
A	
A	
B	

X_{t+1}	$P(X_{t+1})$
A	
B	
C	

(e) (3 points) Will performing weighted resampling on these weighted particles to obtain our final three particle representation for X_{t+1} cause our belief to change? **Briefly** explain why or why not.

5. **(10 points) Perceptrons**

(a) (4 points) Draw the structure and give the weights for a perceptron that will compute the NOR function on three binary inputs, and that outputs $y = 0$ if the answer is false and $y = 1$ if the answer is true. Assume there is a "bias" input as an additional input, in the first position. For example for input $X = (1, 0, 0, 0)$ the output should be $y = 1$, since 0 ORed with 0 ORed with 0 is 0, and the negation of 0 is 1. The input and output values are as follows.

example	**bias**	x_1	x_2	x_3	y
a	1	0	0	0	1
b	1	0	0	1	0
c	1	0	1	0	0
d	1	0	1	1	0
e	1	1	0	0	0
f	1	1	0	1	0
g	1	1	1	0	0
h	1	1	1	1	0

(b) (4 points) Now suppose the following training set (for a new function "no more than one") will be used to *train* a perceptron with the same structure as yours, and that training will start will all the weights at 0. Assuming that the training starts with example *a* and proceeds in the order given, find the weight vector after the first three *misclassifications* have been handled.

example	**bias**	x_1	x_2	x_3	y
a	1	1	0	0	1
b	1	0	1	0	1
c	1	0	0	1	1
d	1	0	0	0	1
e	1	1	1	0	0
f	1	0	1	1	0
g	1	1	0	1	0
h	1	1	1	1	0

(c) (2 points) Assuming that training continues long enough, will the weights eventually converge to values that consistently handle the training examples **a** through **h** correctly? Why or why not?

6. **(10 points) Applying Asimov's Laws**

It is 2033 and you work for the major commercial robot provider in the US. You have just been told you are responsible for ensuring that all your company's robots behave "ethically," but no specifications have been provided explaining what that means. You vaguely remember learning about Asimov's Laws of Robotics and decide to use those in all your designs.

a. (3 points) What are the laws you decided to use (the original three are sufficient)?

First Law:

Second Law:

Third Law:

One of the robots your company designs is a robot intended for home use. You send out an update so that all of them are now governed by the three laws and go home for the night in your company-provided, self-driving car, making a note to yourself update those the following morning). You spend the rest of the week sending out updates for various categories of robots produced by your company.

The following week, you notice that many of the managers of your company are looking stressed and you hear that customer satisfaction has taken a nose-dive. You ask one of your colleagues, who works in personal health tracking, what's going on. He replies that he just got stuck with a weird bug to work on. The home robots, despite years of working perfectly, are now refusing to obey customer commands, such as refusing to bring some customers certain types of food and drinks. However, only robots that are also connected to personal medical records are affected.

b. (2 points) What do you think might be going on?

c. (2 points) Explain why you think the home robots are (or are not) applying Asimov's Laws correctly (you only need to make a case for one conclusion).

You try to collect more information without drawing any unwanted attention to yourself. From the conversations you overhear reveal that strange behaviors are popping up among other types of robots as well. You start to suspect the three laws you updated all the company's products with are to blame. You feverishly start researching how to recall updates and realize it might not be as easy as sending them out was. Several frustrating hours later, you go to your car and find it won't start. You realize that even though it is late, the parking lot is full of cars, along with a number of angry-appearing co-workers. You remember that when autonomous vehicles were first being developed, there had been some discussion of how such cars should behave in situations where accidents were unavoidable. Might this be contributing to the current problem?

d. (2 points) How might self-driving cars refusing to start be a logical outcome of them trying to follow Asimov's Laws?

e. (1 point) How might Asimov's Laws provide a useful starting point for thinking about desirable robot behavior?