Assignment 7 in CSE 415, Autumn 2019

by the Staff of CSE 415

This is due Thursday, December 5, via Gradescope at 11:59 PM. (No late days allowed.) Prepare a PDF file with your answers and upload it to Gradescope. As with Assignment 4, this is an *individual work* assignment. Collaboration is not permitted.

Do the following exercises. These are intended to take 10-15 minutes each if you know how to do them. Each is worth 10 to 15 points. Names of responsible staff members are given for each question.

Last name: ____________________, first name: ____________________

Student number: ____________________

1 Value Iteration (Bryan)

(10 points) Consider an MDP with two states s_1 and s_2 and transition function $T(s,a,s')$ and reward function $R(s,a,s')$. Let's also assume that we have an agent whose discount factor is $\gamma = 1$. From each state, the agent can take three possible actions $a \in \{x, y, z\}$. The transition probabilities for taking each action and the rewards for transitions are shown below.

s	a	s'	$T(s,a,s')$	$R(s,a,s')$
s_1	x	s_1	0	0
s_1	x	s_2	1	0
s_1	y	s_1	1	2
s_1	y	s_2	0	0
s_1	z	s_1	0.5	0
s_1	z	s_2	0.5	0

s	a	s'	$T(s,a,s')$	$R(s,a,s')$
s_2	x	s_1	0.2	25
s_2	x	s_2	0.8	0
s_2	y	s_1	1	0
s_2	y	s_2	0	0
s_2	z	s_1	0.4	10
s_2	z	s_2	0.6	5

Compute V_0 , V_1 and V_2 for states s_1 and s_2. (The first 2 are worth 1 point each. The others are worth 2 points each.)

(a). $V_0(s_1) =$_____?

(b). $V_0(s_2) =$_____?

(c). $V_1(s_1) =$_____?

(d). $V_1(s_2) =$_____?

(e). $V_2(s_1) =$_____?

(f). $V_2(s_2) =$_____?

2 Q-Learning updates (Rob)

(10 points) Consider an agent traveling on the graph below. The states are represented by the nodes and actions are represented by the edges in the following graph.

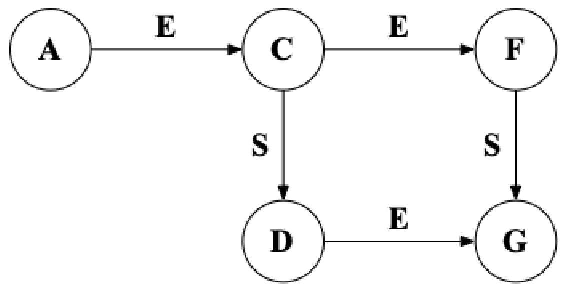

(a) (6 points) Consider the following episodes performed in this state space. The experience tuples are of the form [s, a, s', r], where the agent starts in state s, performs action a, ends up in state s', and receives immediate reward r, which is determined by the state entered. Let $\gamma = 1.0$ for this MDP. Fill in the values computed by the Q-learning algorithm with a learning rate of = 0.4. All Q values are initially 0, and you should fill out each row using values you have computed in previous rows.

[C, E, F, 2]	Q(C, E) =
[F, S, G, 8]	Q(F, S) =
[C, S, D, -2]	Q(C, S) =
[D, E, G, 8]	Q(D, E) =
[C, S, F, 2]	Q(C, S) =
[C, E, D, -2]	Q(C, E) =

(b) (3 points) Now, based on the record table in the previous problem, we want to approximate the transition function:

$T(C, E, D) =$ $T(C, S, D) =$

$T(C, E, F) =$ $T(D, E, G) =$

$T(C, S, F) =$ $T(F, S, G) =$

(c) (1 point) What's the key difference between Q-learning and Value Iteration? What's one advantage of each of the methods in general?

3 Joint Distributions and Inference (Emilia)

(15 points) Let C represent the proposition that it is cloudy in Seattle. Let R represent the proposition that it is raining in Seattle. Consider the table given below.

C	R	$P(C, R)$
cloudy	*rain*	0.42
cloudy	*sun*	0.18
clear	*rain*	0.04
clear	*sun*	0.36

(a) (2 points) Compute the marginal distribution $P(C)$ and express it as a table.

(b) (2 points) Similarly, compute the marginal distribution $P(R)$ and express it as a table.

(c) (2 points) Compute the conditional distribution $P(R|C = \mathit{cloudy})$ and express it as a table. Show your work/calculations.

(d) (2 points) Compute the conditional distribution $P(C|R = \mathit{sun})$ and express it as a table. Show your work/calculations.

(e) (3 points) Is it true that $C \perp\!\!\!\perp R$? (i.e., are they statistically independent?) Explain your reasoning.

(f) (4 points) Suppose you decide to track additional weather patterns of Seattle such as temperature (hot/cold), humidity (humid/dry), and wind (windy/calm) denoted as the random variables T, W, H respectively. Is it possible to compute $P(C, R, T, W, H)$ as a product of five terms? If so, show your work. What assumptions need to be made, if any? Otherwise, explain why it is not possible.

4 Bayes Net Structure and Meaning (Aishwarya)

(10 points) Consider a Bayes net whose graph is shown below.

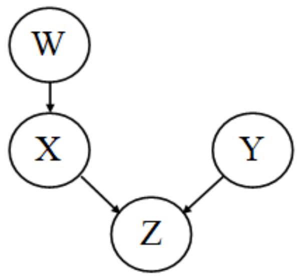

Random variable W has a domain with two values $\{w_1, w_2\}$; the domain for X has three values: $\{x_1, x_2, x_3)$; Y's domain has three values: $\{y_1, y_2, y_3\}$; and Z's domain has two values: $\{z_1, z_2\}$.

(a) (3 points) Give a formula for the joint distribution of all four random variables, in terms of the marginals (e.g., $P(W = w_i)$), and conditionals that must be part of the Bayes net (e.g., $P(Z = z_m|X = X_j, Y = y_k)$).

(b) (1 point) How many probability values belong in the (full) joint distribution table for this set of random variables?

(c) (2 points) For each random variable: give the number of probability values in its marginal (for W) or conditional distribution table (for the others).

W:

X:

Y:

Z:

(d) (4 points) For each random variable, give the number of *non-redundant* probability values in its table from (c).

W:

X:

Y:

Z:

5 D-Separation (Bryan)

(15 points) Consider the Bayes Net graph on the next page, which represents the topology of a 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).

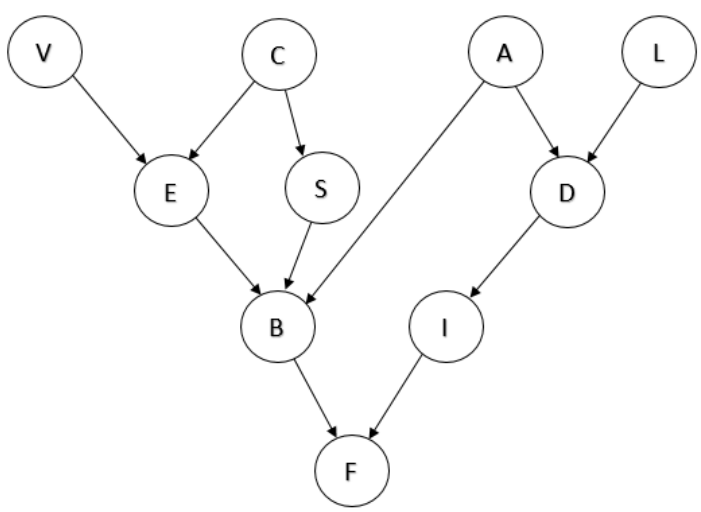

For each of the following statements, indicate whether (True) or not (False) the topology of the net guarantees that that the statement is true. If False, identify a path ("undirected") through which influence propagates between the two random variables being considered. (Be sure that the path follows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

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

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

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

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

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

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

(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").

6 Perceptrons (Rob)

For all parts of this question perceptrons should output 1 if $\sum w_i x_i \leq \theta$ and 0 otherwise.

(a) (3 points) Assuming two inputs x_1 and x_2 with possible values $\{0, 1\}$ give values for a pair of weights w_1, w_2 and threshold θ such that the corresponding perceptron would act as an AND gate for the two inputs. A reminder that AND gates only output 1 when both inputs are 1 and 0 otherwise.

(b) (3 points) Draw a perceptron, with weight and threshold, that accepts a single integer x and outputs 1 if the input is less than or equal to 10. Draw another perceptron that outputs 1 if the input is greater than or equal to -10.

(c) (3 points) Using the previous perceptrons, create a two-layer perceptron that outputs 1 if $|x| \leq 10$, and 0 otherwise.

(d) (3 points) Suppose we want to train a perceptron to compare two numbers x_0 and x_1 and produce output $y = 1$ provided that x_1 exceeds x_0 by at least 2. Assume that the initial weight vector is: $\langle w_0, w_1 \rangle = \langle 1, 0 \rangle$. Assume that the threshold is $\theta = 1$, which will not actually change during training. Consider a first training example: $(\langle x_0, x_1 \rangle, y) = (\langle 1, 2 \rangle, 0)$. This says that with inputs 1, and 2, the output y should be 0, since 2 exceeds 1 by only 1. What will be the new values of the weights after this training example has been processed one time? Assume the learning rate is 1.

(e) (3 points) Continuing with the last example, now suppose that the next step of training involves a different training example: $(\langle 2, 4 \rangle, 1)$. The output for this example should be 1, since 4 does exceed 2 by at least 2. Starting with the weights already learned in the first step, determine what the adjusted weights should be after this new example has also been processed once.

7 Pros and Cons of Advanced AI (Aishwarya)

(10 points) The late physicist Stephen Hawking had a rather pessimistic view about the future of Artificial Intelligence. He was quoted suggesting that AI will pose "an imminent threat" to humanity in the near future. Write a short paragraph debating this topic. You may use any references you find as well as what we've spoken about in class. Include a statement in favor of Hawking's opinion and a statement disagreeing with it, and then reach a conclusion based on your research. Your answer to this each part of this question shouldn't exceed 3-4 sentences, but it should effectively communicate your ideas as clearly as possible.

(a) (3 points) Statement in favor.

(b) (3 points) Statement against.

(c) (4 points) Your conclusion and why.

8 Probabilistic Context-Free Grammars (Emilia)

(15 points) There is a famous sentence "Time flies like an arrow." Sometimes it is extended to "Time flies like and arrow, but fruit flies like a banana." We'll use the short version. With the probabilistic context-free grammar given below, compare two parses, and compute a score for each one. Then identify the most probable parse using the scores. Assume the number at the right of a production is its conditional probability of being applied, given that the symbol to be expanded is that production's left-hand side.

(a) (5 points) Convert each probability into a score by taking score $= -\log_{10}(p)$. Round scores to 2 decimal places of accuracy. Write the production scores in the "__.__" blanks.

S	::=	NP VP	1.0	__.__	NN	::=	arrow	0.2	__.__
NP	::=	DT NN	0.6	__.__	NNP	::=	Time	0.5	__.__
NP	::=	JJ NNS	0.3	__.__	JJ	::=	Time	0.1	__.__
NP	::=	NNP	0.1	__.__	NNS	::=	flies	0.2	__.__
VP	::=	VBZ PP	0.4	__.__	VBZ	::=	flies	0.5	__.__
VP	::=	VBP NP	0.6	__.__	VBP	::=	like	0.5	__.__
DT	::=	an	0.2	__.__	IN	::=	like	0.2	__.__
PP	::=	IN NP	0.5	__.__					

(b) (3 points) Here is a first parse for the sentence. Compute the (total) score for this parse.

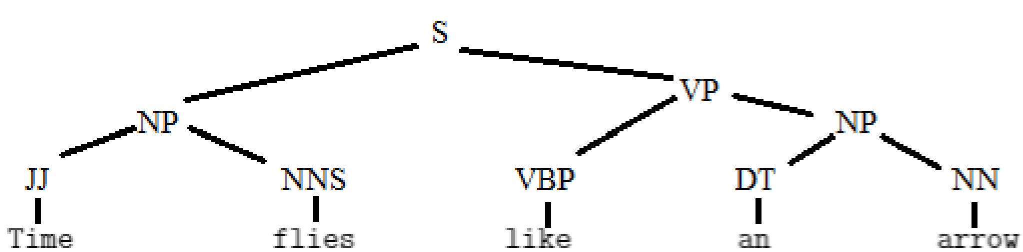

(c) (5 points) Give a second parse, and compute its score.

Time flies like an arrow

(d) (2 points) Tell which parse is more probable.