Assignment 7 in CSE 415, Winter 2020*

by the Staff of CSE 415

This is due Wednesday, March 11, via Gradescope at 11:59 PM. Prepare a PDF file with your answers and upload it to Gradescope. The PDF file can created however you like. For example it can be from a scan of a printout of the assignment document onto which you have hand-written your answers. Or it can be from Word or Latex file with your answers. It can even be from photos of your handwritten answers on plain paper. You don't have to include the questions themselves, but it is fine to do so. In any case, it must be very clear to read, and it must be obvious and easy for each grader where to find your solutions to the exercises.

As with Assignment 4, this is an *individual work* assignment. Collaboration is not permitted.

Do the following exercises. These are intended to take 10-20 minutes each if you know how to do them. Each is worth 10 to 20 points. The total of possible points is 120. Names of responsible staff members are given for each question.

If corrections or clarifications to the problems have to be given, this will happen in the ED discussion forum under topic "Assignment 7."

Last name: ____________________, first name: ____________________

Student number: ____________________

* Version 1.3, with corrections to ex. 6c and 6e, as well as 8a and 8f.

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	0.7	0
s_1	x	s_2	0.3	0
s_1	y	s_1	0	2
s_1	y	s_2	1	7
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	0.4	25
s_2	x	s_2	0.6	20
s_2	y	s_1	1	0
s_2	y	s_2	0	0
s_2	z	s_1	0.8	10
s_2	z	s_2	0.2	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 (Bindita)

(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.

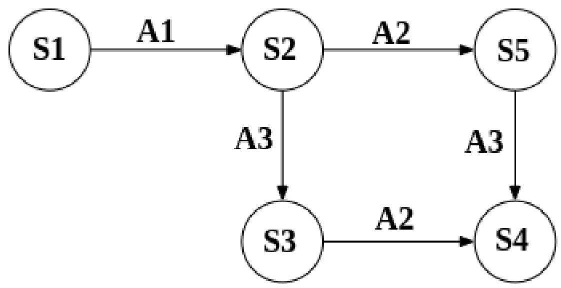

(a) (6 points) Consider the following episodes performed in this state space. The experience tuples are of the form [s, a, s', r], where the agent starts in state s, performs action a, ends up in state s', and receives immediate reward r, which is determined by the state entered. Let $\gamma = 1.0$ for this MDP. Fill in the values computed by the Q-learning algorithm with a learning rate of = 0.3. All Q values are initially 0, and you should fill out each row using values you have computed in previous rows.

[S2, A2, S5, 3]	Q(S2, A2) =
[S5, A3, S4, 5]	Q(S5, A3) =
[S2, A3, S3, -4]	Q(S2, A3) =
[S3, A2, S4, 7]	Q(S3, A2) =
[S2, A3, S5, 2]	Q(S2, A3) =
[S2, A2, S3, -2]	Q(S2, A2) =

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

$T(S2, A2, S3) =$ $\qquad$ $T(S2, A3, S3) =$

$T(S2, A2, S5) =$ $\qquad$ $T(S3, A2, S4) =$

$T(S2, A3, S5) =$ $\qquad$ $T(S5, A3, S4) =$

(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*	0.65
cloudy	*sun*	0.15
clear	*rain*	0.05
clear	*sun*	0.15

(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, H, W respectively. Is it possible to compute $P(C, R, T, H, W)$ 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 (Steve)

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

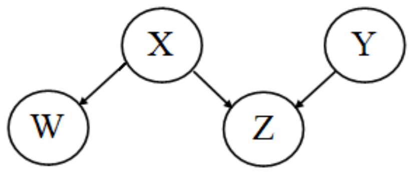

Random variable W has a domain with two values $\{w_1, w_2\}$; the domain for X has three values: $\{x_1, x_2, x_3)$; Y's domain has two values: $\{y_1, y_2\}$; and Z's domain has four values: $\{z_1, z_2, z_3, z_4\}$.

(a) (3 points) Give a formula for the joint distribution of all four random variables, in terms of the marginals (e.g., $P(X = x_i)$), 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 X and Y) 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).

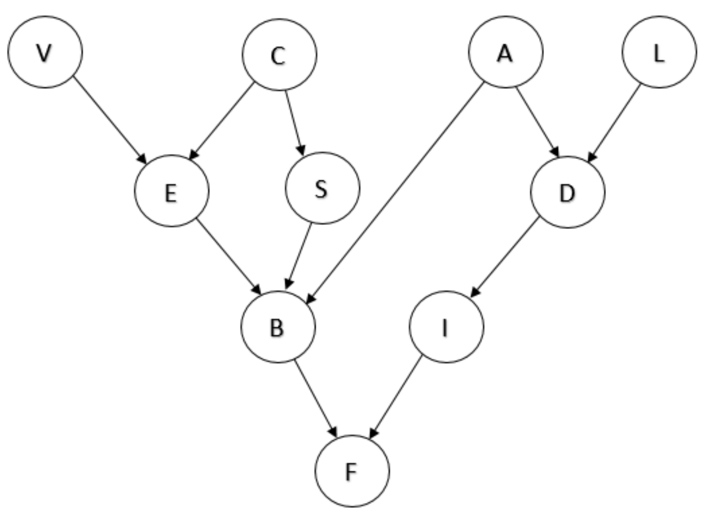

For each of the following statements, indicate whether (True) or not (False) the topology of the net guarantees that that the statement is true. If False, identify a path ("undirected") through which influence propagates between the two random variables being considered. (Be sure that the path follows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

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

(b) $D \perp\!\!\!\perp E \mid F, I$

(c) $B \perp\!\!\!\perp L \mid A, I$

(d) $A \perp\!\!\!\perp I \mid D, F$

(e) $L \perp\!\!\!\perp V \mid D, E, F$

(f) $V \perp\!\!\!\perp L \mid A, D$

(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 (Shivam)

For all parts of this question perceptrons should output 1 if $\sum w_i x_i \geq \theta$ 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 a NOR gate for the two inputs. A reminder that NOR gates only output 1 when both inputs are 0; 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 greater than or equal to 5. Draw another perceptron that outputs 1 if the input is less than or equal to -5.

(c) (3 points) Using the previous perceptrons, create a two-layer perceptron that outputs 1 if $|x| \leq 5$, and 0 otherwise. (correction)

(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 5. Assume that the initial weight vector is: $\langle w_0, w_1 \rangle = \langle 0, 1 \rangle$. 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: $(\langle 2, 8 \rangle, 1)$. The output for this example should be 1, since 8 does exceed 2 by at least 5. (correction) Starting with the weights already learned in the first step, determine what the adjusted weights should be after this new example has been processed once.

7 Pros and Cons of Advanced AI (Emilia)

(10 points) Facial recognition has been in the news recently due to law enforcement agencies in the US adopting technology marketed by Clearview AI. Clearview AI's facial recognition system relies on a database of billions of images scraped from millions of sites including Facebook, YouTube, and Venmo. In this question, imagine you've been asked to make a presentation on the pros and cons of such a system.

(See: https://bit.ly/39o1fBv for more information.)

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 Markov Models (Steve)

(20 points) According to an unnamed source, the stock market can be modeled using a Markov model, where there are two states "bull" and "bear." The dynamics of the model are given:

S_{t-1}	S_t	$P(S_t \mid S_{t-1})$
bull	bull	0.8
bull	bear	0.2
bear	bull	0.4
bear	bear	0.6

(a) (2 points) Suppose it's given that S_0 = bull. Compute the probability that S_2 = bull.

(b) (4 points) Compute the stationary probabilities for bull and bear.

(c) (2 points)Now suppose that whenever it's a bull market, a certain company's (Acme, Inc) stock stays the same or rises in value with probability 0.7 and falls in value with probably 0.3.

When it's a bear market Acme's stock value stays the same or rises with probability 0.1 and falls with probability 0.9.

Suppose an observer cannot directly tell whether the state of the stock market is bull or bear, but can only see whether Acme's stock is "rising" or "falling."

State S	Observation Q	$P(Q \mid S)$
bull	rising	0.7
bull	falling	0.3
bear	rising	0.1
bear	falling	0.9

Suppose $P(S_0 = \text{bull}) = 0.5$. If the observation at time 1 is "rising," what is the belief in S=bull right after the observation?

(d) (2 points) Suppose at time 2, the observation is "falling". what is the belief in S = bull right after that observation? (This belief will take into consideration the previous belief you computed above.)

(e) (2 points) Suppose that the actual state sequence for the first four time steps is bear, bear, bear, bull

What is the probably of observing the sequence (starting at $t = 1$) rising, falling, rising?

(f) (3 points) Now, what if the state sequence was bull, bear, bear, bull. What is the probability of observing the same sequence of stock changes as above?

(g) (2 points) Which of these two state sequences is more likely, given that sequence of observations?

(h) (3 points) Is there another state sequence that is even more likely? Explain.

(i) (optional, not for credit) Use the Viterbi algorithm to compute the most likely state sequence for this observation sequence. You should draw a trellis diagram. Assume that the initial probability distribution is $P(S_0 = \text{bull}) = 0.5$; $P(S_0 = \text{bear}) = 0.5$. A very concise and relatively clear video presentation of a simple Viterbi algorithm example is one by Luis Serrano at:

`https://www.youtube.com/watch?v=mHEKZ8jv2SY`

9 Probabilistic Context-Free Grammars (Steve)

(15 points) Consider the sentence "Let's stop controlling people." This can be interpreted to be about us controlling other people, or it can be interpreted to be about trying to stop people who have a controlling nature. With the probabilistic context-free grammar given below, compare two parses that correspond to each of these interpretations, 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	0.5	__.__	VP	::=	VB NP	0.4	__.__
S	::=	VP	0.5	__.__	VB	::=	let	0.2	__.__
NP	::=	ADJ NP	0.5	__.__	PRP	::=	's	0.3	__.__
NP	::=	PRP	0.3	__.__	VB	::=	stop	0.1	__.__
NP	::=	NNS	0.1	__.__	VBG	::=	controlling	0.3	__.__
VP	::=	VB S	0.2	__.__	NNS	::=	people	0.2	__.__
VP	::=	VBG NP	0.3	__.__	ADJ	::=	controlling	0.1	__.__

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

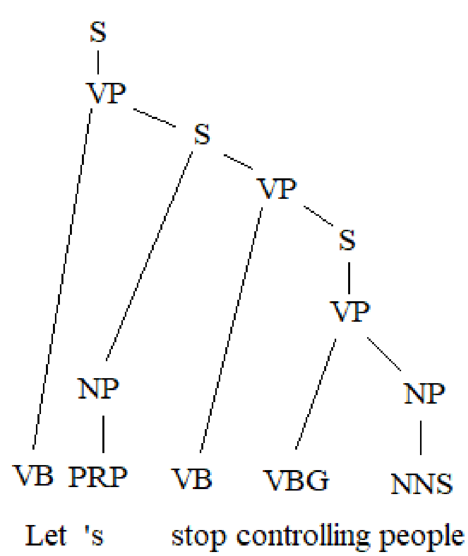

(c) (5 points) Here is the second parse. Compute its score.

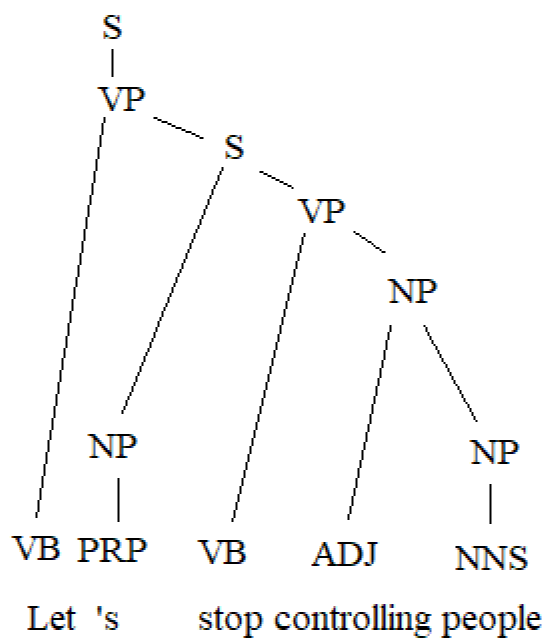

(d) (2 points) Convert each score back to a probability and write them here as P1 and P2. Then tell which parse is more probable.