

Assignment 7 in CSE 415, Winter 2024

by the Staff of CSE 415

This is due Wednesday, March 6 via Gradescope at 11:59 PM. Prepare a PDF file with your answers and upload it to Gradescope. The PDF file can be 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.

Prepare your answers in a neat, easy-to-read PDF. If you choose to typeset your answers in Latex using the template file for this document, please put your answers in [blue](#) while leaving the original text black.

Do the following exercises. These are intended to take 15-30 minutes each if you know how to do them. Each is worth 20 points. The total of possible points is 100.

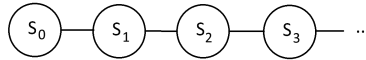
If corrections or clarifications to the problems have to be given, this will happen in the ED discussion forum under topic "Assignment 7."

Last name: _____, first name: _____

Student number: _____

1 Markov Models

(20 points) Consider the Markov Model for the stock market shown below.



The transition probabilities for this model are given as:

S_t	S_{t+1}	$P(S_{t+1} S_t)$
rising	rising	0.65
rising	falling	0.35
falling	rising	0.5
falling	falling	0.5

- (a) (5 points) Assume $P(S_0 = \text{rising}) = 1.0$ and $P(S_0 = \text{falling}) = 0.0$. Using the miniforward algorithm, compute the $P(S_2)$ distribution.
- (b) (2 points) Continuing with the miniforward algorithm, compute the $P(S_3)$ distribution.
- (c) (3 points) Now assume $P(S_0 = \text{rising}) = 0.0$ and $P(S_0 = \text{falling}) = 1.0$. Compute the $P(S_2)$ distribution.
- (d) (10 points) Compute the stationary distribution $P(S_\infty)$. Show your work.

2 D-Separation

(20 points) Consider the Bayes Net graph β below, 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.

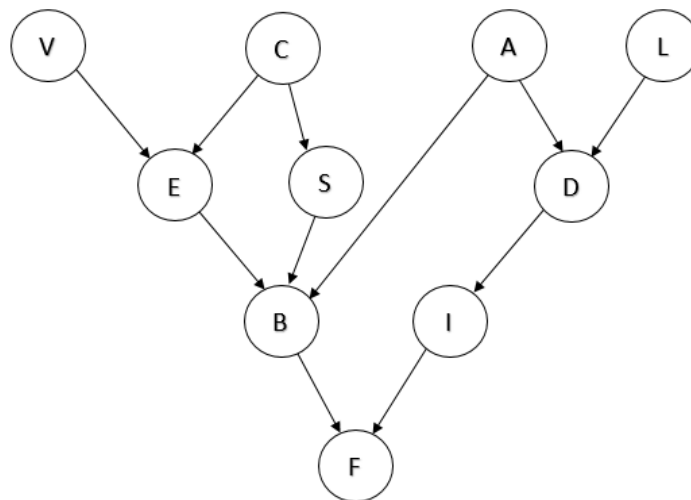
E = Exposure to vulnerability is high.

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).



Let β' be the undirected graph obtained from β by removing the arrowheads from the edges of β . By an “undirected path” in β we mean any path in β' . A “loop-free” path is any path in which no vertex is repeated.

- (a) (5 points) List all loop-free undirected paths from A to C in the graph β .
- (b) (5 points) Suppose random variables F and S are observed, and no others are observed. Then which (if any) of those paths would be active paths? Justify your answer.

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.

(c) $C \perp\!\!\!\perp A$: True (both are roots and no common descendants have been observed)

(d) (1 point) $B \perp\!\!\!\perp I \mid A$

(e) (1 point) $V \perp\!\!\!\perp D \mid E, B$

(f) (1 point) $C \perp\!\!\!\perp D \mid A, I$

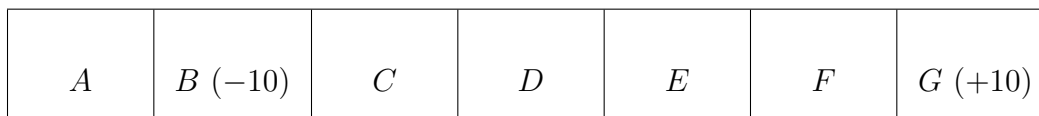
(g) (1 point) $C \perp\!\!\!\perp L \mid V, E, F, I$

(h) (1 point) $V \perp\!\!\!\perp F \mid E, B$

(i) (5 points) Suppose that the company hired an outside expert to examine the system and she determines that B, E, and I are true: a Break-in has occurred (the web server is compromised), Exposure to vulnerability is high, and Incident response is effective. Given this information, your job is to explain to management why getting additional information about C (Complexity to access the server is high) could have an impact on the probability of L (logging infrastructure is state-of-the-art). Give your explanation, for the manager of the company, using about 3 to 12 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”).

3 Reinforcement Learning

Use the Markov Decision Process (MDP) diagram below to help you as you answer the following questions about Value Iteration and Q Learning.



- (a) (10 points) Use Q Learning to estimate the true q values of this MDP given the observed state transitions below. For each state transition, indicate which q value is changing by giving the state and action associated with that value and giving its new value. You may wish to keep your own diagram of the current values while you work through the given transitions.

You may assume that at states $\{A, C, D, E, F\}$ there are two allowed actions, {Left, Right}, at states $\{B, G\}$, the only allowable action is the Exit action, and at the Terminal state there are no valid actions. Make sure to process the transitions in the order provided. All Q values are initialized to 0. Use a discount factor of $\gamma = 0.8$ and a learning rate of $\alpha = 0.5$.

Note that although a description of the MDP parameters are given in the next part of this question, Q Learning does not rely on knowing these parameters and this problem may not use the same rules and parameters.

State Transitions				New Q Values		
State (s)	Action (a)	New State (s')	Reward (r)	State	Action	Q(S,A)
A	Right	B	−1			
B	Exit	Terminal	−10			
E	Right	D	−1			
D	Right	E	−1			
E	Right	F	−1			
F	Right	G	−1			
G	Exit	Terminal	10			

- (b) (10 points) Value Iteration:

Consider an MDP with the states shown in the diagram above. In this MDP, there are two allowable actions for non-goal states, {Left, Right}, whose intended effect is to move the agent to the state directly to the left or right of the current state respectively.

However, at each action with probability η (noise), the agent will instead move to the state in the opposite direction of its intended move. If at any point an action would move the agent to an invalid state, the agent instead will remain in place.

At states B and G , the only allowable action is the Exit action, which gives the reward indicated in the diagram and moves the agent to the Terminal state, which is a state with no allowable actions, indicating the end of the episode.

A few examples with $\eta = 0.3$: from state D , taking action Right, the agent will move to state E with probability 0.7 and will move to state C with probability 0.3. From state A , taking action Left, the agent will remain in state A with probability 0.7 and will move to state B with probability 0.3. From state B , the only allowed action is Exit, which gives reward -10 .

Perform three iterations of value iteration on the MDP described above, using noise $\eta = 0.2$, discount $\gamma = 0.9$, and a living reward of 0.0. The initial values, V_0 , for all the states are provided for you. Write the values V_i corresponding to each state in the boxes for that state for iterations $i = 1, 2, 3$.

V_0 :

$A: 0$	$B: 0$	$C: 0$	$D: 0$	$E: 0$	$F: 0$	$G: 0$
--------	--------	--------	--------	--------	--------	--------

V_1 :

A ____	B ____	C ____	D ____	E ____	F ____	G ____
----------	----------	----------	----------	----------	----------	----------

V_2 :

A ____	B ____	C ____	D ____	E ____	F ____	G ____
----------	----------	----------	----------	----------	----------	----------

V_3 :

A ____	B ____	C ____	D ____	E ____	F ____	G ____
----------	----------	----------	----------	----------	----------	----------

4 The Laws of Robotics

(20 points) In the 1940's, Isaac Asimov introduced a set of three laws to govern robot behavior:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

(NOTE: You might also want to take a look at this cartoon <https://xkcd.com/1613/>)

For this question, you will read Asimov's short story *Liar!* linked here:

<https://courses.cs.washington.edu/courses/cse415/24wi/uwnetid/misc/Asimov-Liar.pdf>.

It is one of Asimov's robot short stories that were collected together in the book *I, Robot*. After you have read the story, please answer the questions below.

(a) (1 points) Who was Isaac Asimov and what is his relevance to the field of Artificial Intelligence (answer in 3-4 sentences)?

(b) (2 points) Explain why the robot was lying to the humans in the story.

-
- (c) (2) Do you think the robot behaved in accordance to the Laws of Robotics by lying or should it have realized that lying could also lead to harm being done (remember to explain your reasoning)?
- (d) (4 points) Brain-computer interfaces are an area of active research today. What are some beneficial applications of these technologies (if you look online to help you answer this question, please include citations)?
- (e) (4 points) What are some concerns you would have with technology that could in some way "read" the user's mind (cite any references you consult)?

-
- (f) (1 points) Keeping in mind the requirement that the robot has to obey the Laws of Robotics, did you find the end of the story effective and/or convincing? Explain why or why not.
- (g) (3 points) If you were creating a mind-reading robot today, how might you try to program the robot to prevent such a catastrophic result? In your response, assume you are free to modify or eliminate the laws of robotics.
- (h) (3 points) Did the approach you described above require modifications to or the elimination of the laws of robotics? Do you think constraints at the hardware (e.g. robot) level are the correct approach to take to ensuring ethical uses of technology? Explain why or why not.

5 Perceptrons

(20 points) In this question we'll assume that perceptrons should output 1 if $\sum_{i=0}^n w_i x_i \geq 0$ and 0 otherwise. The weight w_0 is called the bias weight, and x_0 is called the bias input and it is always equal to 1.

- (a) (4 points) Assuming there will be two inputs x_1 and x_2 , each with possible values in $\{0, 1\}$, give values for a triple of weights $\langle w_0, w_1, w_2 \rangle$ such that the corresponding perceptron would act as a NOR gate for the two inputs. (Again, weight w_0 is the bias weight.) Note that a NOR gate outputs 1 when both its inputs are 0; it outputs 0 otherwise.
- (b) (4 points) Draw a perceptron, with weights, that accepts a single integer x_1 and outputs 1 if and only if the input is greater than or equal to -5 . Be sure to include x_0 (the bias input of value 1) and its weight w_0 in your diagram. Draw another perceptron that outputs 1 if and only if the input is less than or equal to 5.
- (c) (4 points) Using the previous perceptrons, create a two-layer perceptron that outputs 1 if $|x| \leq 5$, and 0 otherwise.
- (d) (4 points) Suppose we want to train a perceptron to compare two numbers x_1 and x_2 and produce output $y = 1$ provided that x_2 exceeds x_1 by at least 7. Assume that the initial weight vector is: $\langle w_0, w_1, w_2 \rangle = \langle 1, -1, -1 \rangle$. Consider a first training example: $(\langle x_1, x_2 \rangle, y) = (\langle 3, 11 \rangle, 1)$. This says that with inputs 3, and 11, the output y should be 1, since 11 exceeds 3 by 8. What will be the new values of the weights after this training example has been processed one time? Assume the learning rate is 2.
- (e) (4 points) Continuing with the last example, now suppose that the next step of training involves a different training example: $(\langle 2, 8 \rangle, 0)$. The output for this example should be 0, since 8 does not exceed 2 by at least 7. 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.