

Due: Wednesday 03/30/2022 at 10:59pm (submit via Gradescope).

Policy: Can be solved in groups (acknowledge collaborators) but must be written up individually

Submission: It is recommended that your submission be a PDF that matches this template. You may also fill out this template digitally (e.g. using a tablet). **However, if you do not use this template, you will still need to write down the below four fields on the first page of your submission.**

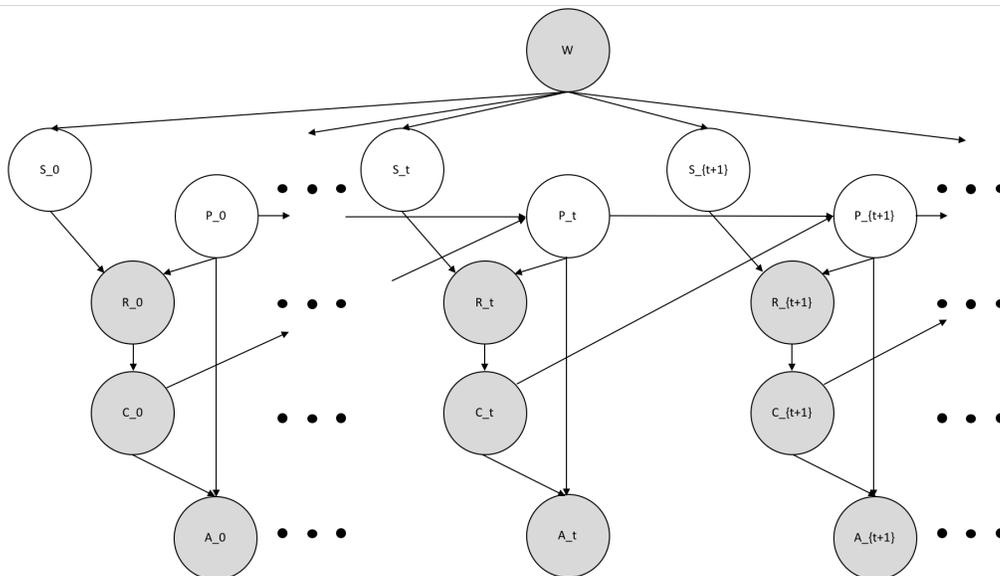
First name	
Last name	
SID	
Collaborators	

For staff use only:

Q1.	Quadcopter: Pilot	/27
Q2.	MangoBot Human Detector	/8
	Total	/35

Q1. [27 pts] Quadcopter: Pilot

In this question, we look at the same setup as the previous homework, but we now look at it from the perspective of Paul, a quadcopter pilot who can **observe W (weather), R (reading of position), C (control from the pilot), and A (smart alarm warning)**. As before, suppose weather (W) does not change throughout the quadcopter's flight.



(a) Forward Algorithm: The real deal

(i) [2 pts] Now that the only hidden states are S_t and P_t , is this graph a well-behaving HMM (where $E_{t+1} \perp\!\!\!\perp E_t \mid X_{t+1}$ and $X_{t+1} \perp\!\!\!\perp E_t \mid X_t$, recall that X is the hidden variable and E is the evidence variable)? Please explain your reasoning.

(ii) [4 pts] What is the time-elapsd **prediction** formula from time-step t to time-step $t + 1$? Be sure to include all hidden states and observed states, and show how to assemble the formula from the conditional probability tables corresponding to the graph. Denote $B(S_t, P_t) = P(S_t, P_t \mid W, R_{0:t}, C_{0:t}, A_{0:t})$. Find $B'(S_{t+1}, P_{t+1}) = P(S_{t+1}, P_{t+1} \mid W, R_{0:t}, C_{0:t}, A_{0:t})$. Hint: refer to the “Filtering algorithm” lecture slides.

- $B'(S_{t+1}, P_{t+1}) = \max_{s_t} \sum_{p_t} P(S_{t+1} \mid W) * P(P_{t+1} \mid p_t, c_t) * B(S_t, P_t)$
- $B'(S_{t+1}, P_{t+1}) = \sum_{s_t} \sum_{p_t} P(S_{t+1} \mid S_t) * P(P_{t+1} \mid p_t) * B(S_t, P_t)$
- $B'(S_{t+1}, P_{t+1}) = \sum_{s_t} \sum_{p_t} P(R_{t+1} \mid S_{t+1}, P_{t+1}) P(S_{t+1} \mid S_t) * P(P_{t+1} \mid p_t) * B(S_t, P_t)$
- $B'(S_{t+1}, P_{t+1}) = \sum_{s_t} \sum_{p_t} P(S_{t+1} \mid S_t) * P(P_{t+1} \mid p_t) P(P_{t+1} \mid c_t) * B(S_t, P_t)$
- $B'(S_{t+1}, P_{t+1}) = \sum_{s_t} \sum_{p_t} P(S_{t+1} \mid W) * P(P_{t+1} \mid p_t, c_t) * B(S_t, P_t)$
- $B'(S_{t+1}, P_{t+1}) = \sum_{s_t} \max_{p_t} P(S_{t+1} \mid W) * P(P_{t+1} \mid p_t, c_t) * B(S_t, P_t)$

- (iii) [4 pts] How do we include the observation **update** at time-step $t + 1$? Be sure to include all hidden states and observed states, and show how to assemble the update from the conditional probability tables corresponding to the graph. Denote $B'(S_{t+1}, P_{t+1}) = P(S_{t+1}, P_{t+1} | W, R_{0:t}, C_{0:t}, A_{0:t})$, find $B(S_{t+1}, P_{t+1})$
- $B(S_{t+1}, P_{t+1}) = P(R_{t+1} | S_{t+1}, P_{t+1}) * P(C_{t+1} | R_{t+1}) * P(A_{t+1} | C_{t+1}) * B'(S_{t+1}, P_{t+1})$
 - $B(S_{t+1}, P_{t+1}) = P(R_{t+1} | S_{t+1}, P_{t+1}) * P(C_{t+1} | R_{t+1}) * B'(S_{t+1}, P_{t+1})$
 - $B(S_{t+1}, P_{t+1}) = P(R_{t+1} | S_{t+1}, P_{t+1}) * P(C_{t+1} | R_{t+1}) * P(A_{t+1} | C_{t+1}, P_{t+1}) * B'(S_{t+1}, P_{t+1})$
 - $B(S_{t+1}, P_{t+1}) \propto P(R_{t+1} | S_{t+1}, P_{t+1}) * P(C_{t+1} | R_{t+1}) * P(A_{t+1} | C_{t+1}, P_{t+1}) * B'(S_{t+1}, P_{t+1})$
 - $B(S_{t+1}, P_{t+1}) \propto P(R_{t+1} | S_{t+1}, P_{t+1}) * P(C_{t+1} | R_{t+1}) * B'(S_{t+1}, P_{t+1})$

(b) Consider a simpler scenario where we only track the 2D position (x, y) of the quadcopter. Paul, the pilot, wants to infer the quadcopter's true position P as accurately as possible.

- x, y **each** can take on values $\in \{0, 1, 2\}$.
- We have four controls: forward, backward, left, and right.
- Let variable E_R be Paul's estimate of the current position, and this variable depends on the reading R . The utility is based on the difference between the estimate of current position E_R and the actual position P : $U(P, E_R) = -\|P - E_R\|_x - \|P - E_R\|_y$, in dollars.
- We consider only one time step. In that time step the **reading R is** $(1, 0)$ and that the weather is cloudy.
- Under cloudy weather, the signal strength can take on 2 values with equal probability: weak and strong. The signal strengths correspond to the following errors in readings:
 - Weak: The reading R returns a random number (for each position element) sampled uniformly from the domain of possible positions.
 - Strong: The reading R is identical to the true position.

Answer the following questions:

- (i) [2 pts] Among the hidden variables S and P , Which variable should intuitively have the greatest VPI? Explain your answer. You should not do any calculations for this part.

Paul's coworker offers to tell him the signal strength (S) in exchange for some cash.

- (ii) [3 pts] Suppose the signal strength is strong. Given the current reading R , what is the Maximum Expected Utility after knowing this information of S ?

(iii) [3 pts] Suppose the signal strength is weak. Given the current reading R , what is the Maximum Expected Utility after knowing this information of S ?

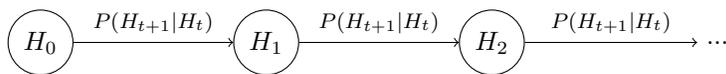
(iv) [3 pts] Considering the possibility of both signal strength, how much should Paul should pay to know this information of S ?

(c) (i) [3 pts] Suppose your coworker only tells you the signal strength with probability q , and with probability $1 - q$, they don't tell you the signal strength even after payment. How much would you be willing to pay in this scenario? Your result should contain q .

(ii) [3 pts] How much would you pay to know the true position (P)?

Q2. [8 pts] MangoBot Human Detector

Your startup company MangoBot wants to build robots that delivers packages on the road. One core module of the robot's software is to detect whether a human is standing in front of it. We model the presence of humans with a Markov model:



where $H_t \in \{0, 1\}$ corresponds to a human being absent or present respectively. The initial distribution and the transition probabilities are given as follows:

H_0	$P(H_0)$
0	p
1	$1 - p$

H_t	H_{t+1}	$P(H_{t+1} H_t)$
0	0	0.9
0	1	0.1
1	0	0.8
1	1	0.2

(a) Express the following quantities in terms of p :

(i) [1 pt] $P(H_1 = 1) =$

(ii) [1 pt] $\lim_{t \rightarrow \infty} P(H_t = 0) =$

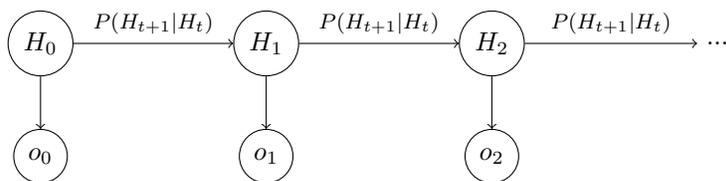
(b) The first-order Markov assumption in the model above can be inaccurate in real-world situations. Some potential ways to improve the model are listed below. For each option, determine whether it is possible to rewrite the process as a first-order Markov process, potentially with a different state representation.

(i) [1 pt] H_t depends not only on H_{t-1} but also on H_{t-2} . Yes No

(ii) [1 pt] H_t depends not only on H_{t-1} but also on $H_{t-2}, H_{t-3}, \dots, H_{t-k}$ for some fixed $k \geq 3$. Yes No

(iii) [1 pt] H_t depends not only on H_{t-1} but also on $H_{t-2}, H_{t-3}, \dots, H_1, H_0$. Yes No

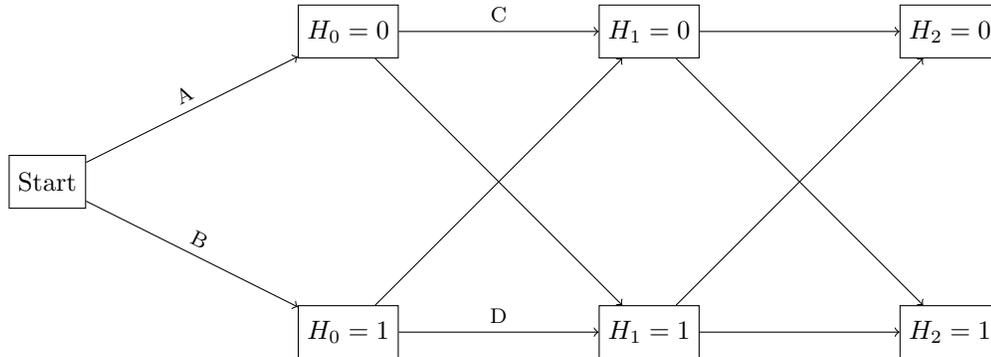
To make things simple, we stick to the original first-order Markov chain formulation. To make the detection more accurate, the company built a sensor that returns an observation O_t each time step as a noisy measurement of the unknown H_t . The new model is illustrated in the figure, and the relationship between H_t and O_t is provided in the table below.



H_t	O_t	$P(O_t H_t)$
0	0	0.8
0	1	0.2
1	0	0.3
1	1	0.7

(c) Based on the observed sensor values o_0, o_1, \dots, o_t , we now want the robot to find the most likely sequence H_0, H_1, \dots, H_t indicating the presence/absence of a human up to the current time.

- (i) [1 pt] Suppose that $[o_0, o_1, o_2] = [0, 1, 1]$ are observed. The following "trellis diagram" shows the possible state transitions. Fill in the values for the arcs labeled A, B, C, and D with the product of the transition probability and the observation likelihood for the destination state. The values may depend on p .



- (ii) [1 pt] There are two possible most likely state sequences, depending on the value of p . Complete the following (Write the sequence as "x,y,z" (without quotes), where x, y, z are either 0 or 1):
Hint: it might be helpful to complete the labelling of the trellis diagram above.

• When $p < \boxed{}$, the most likely sequence H_0, H_1, H_2 is $\boxed{}$.

• Otherwise, the most likely sequence H_0, H_1, H_2 is $\boxed{}$.

- (d) [1 pt] True or False: For a fixed p value and observations $\{o_0, o_1, o_2\}$ in general, H_1^* , the most likely value for H_1 , is always the same as the value of H_1 in the most likely sequence H_0, H_1, H_2 . True False