

Due: Wednesday 04/08/2020 at 11:59pm (submit via Gradescope).

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

Submission: Your submission should be a PDF that matches this template. Each page of the PDF should align with the corresponding page of the template (page 1 has name/collaborators, question 1 begins on page 2, etc.). **Do not reorder, split, combine, or add extra pages.** The intention is that you print out the template, write on the page in pen/pencil, and then scan or take pictures of the pages to make your submission. You may also fill out this template digitally (e.g. using a tablet.)

First name	
Last name	
SID	
Collaborators	


For staff use only:

Q1. Quadcopter: Spectator	/30
Q2. Quadcopter: Data Analyst	/36
Q3. Quadcopter: Pilot	/27
Total	/93

Q1. [30 pts] Quadcopter: Spectator

Flying a quadcopter can be modeled using a Bayes Net with the following variables:

- W (weather) $\in \{\text{clear, cloudy, rainy}\}$
- S (signal strength) $\in \{\text{strong, medium, weak}\}$
- P (true position) = (x, y, z, θ) where x, y, z **each** can take on values $\in \{0, 1, 2, 3, 4\}$ and θ can take on values $\in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$
- R (reading of the position) = (x, y, z, θ) where x, y, z **each** can take on values $\in \{0, 1, 2, 3, 4\}$ and θ can take on values $\in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$
- C (control from the pilot) $\in \{\text{forward, backward, rotate left, rotate right, ascend, descend}\}$ (6 controls in total)
- A (smart alarm to warn pilot if that control could cause a collision) $\in \{\text{bad, good}\}$



(a) Representation

(i) [2 pts] What is N_p , where N_p is the domain size of the variable P ? Please explain your answer.

Answer: $N_p =$

Explanation:

(ii) [2 pts] Please list **all** of the Conditional Probability Tables that are needed in order to represent the Bayes Net above. Note that there are 6 of them.

(iii) [1 pt] What is the size of the Conditional Probability Table for R ? You may use N_p in your answer.

Now, assume that we look at this setup from the perspective of Spencer – a spectator who can observe A and W . Spencer observes $A=\text{bad}$ and $W=\text{clear}$, and he now wants to infer the signal strength. In BN terminology, he wants to calculate $P(S|A = \text{bad}, W = \text{clear})$.

(b) [5 pts] Inference by Enumeration

If Spencer chooses to solve for this quantity using inference by enumeration, what is the biggest **factor** that must be calculated along the way, and what is its **size**? You may use N_p in your answer. Please show your work.

Biggest factor:

Size of factor:

(c) [5 pts] Inference by Variable Elimination: Order 1

Spencer chooses to solve for this quantity by performing variable elimination in the order of $R - P - C$. Answer the following prompts to work your way through this procedure.

(1a) First, we need to eliminate R . Which factors (from the 6 CPTs above) are involved?

(1b) Describe how you eliminate the variable of interest by multiplication of those factors. What conditional probability **factor** results from this step?

(2a) Second, we need to eliminate P . Which factors are involved?

(2b) Describe how you eliminate the variable of interest by multiplication of those factors. What conditional probability **factor** results from this step?

(3a) Third, we need to eliminate C . Which factor/s are involved?

(3b) Describe how you eliminate the variable of interest by multiplication of those factors. What conditional probability **factor** results from this step?

(4) List the 3 conditional probability factors that you calculated as a result of the 3 elimination steps above, along with their domain sizes. You may use N_p in your answer. Which factor is the biggest? Is this bigger or smaller than the biggest factor from the “inference by enumeration” approach?

(5) List the **1** unused conditional probability factor from the 3 that you calculated above, and also list the **2** resulting conditional probability factors from the 6 original CPTs.

(6) Finally, let's solve for the original quantity of interest: $P(S|A = \text{bad}, W = \text{clear})$. After writing the equations to show how to use the factors from (5) in order to solve for $f(S|A = \text{bad}, W = \text{clear})$, don't forget to write how to turn that into a probability $P(S|A = \text{bad}, W = \text{clear})$.
Hint: use Bayes Rule, and use the 3 resulting factors that you listed in the previous question.

(d) [5 pts] Inference by Variable Elimination: Order 2

Spencer chooses to solve for this quantity by performing variable elimination, but this time, he wants to do elimination in the order of $P - C - R$. Answer the following prompts to work your way through this procedure.

(1a) First, we need to eliminate P . Which factors (from the 6 CPTs above) are involved?

(1b) Show the multiplication of those factors and eliminate the variable of interest. What conditional probability **factor** results from this step?

(2a) Second, we need to eliminate C . Which factors are involved? Recall that you might want to use the factor that resulted from the previous step, but you should not reuse factors that you already used.

(2b) Show the multiplication of those factors and eliminate the variable of interest. What conditional probability **factor** results from this step?


(3a) Third, we need to eliminate R . Which factor/s are involved? Recall that you might want to use the factor that resulted from the previous step, but you should not reuse factors that you already used.

(3b) Show the multiplication of those factors and eliminate the variable of interest. What conditional probability **factor** results from this step?

(4) List the 3 conditional probability factors that you calculated as a result of the 3 elimination steps above, along with their domain sizes. You may use N_p in your answer. Which factor is the biggest? Is this bigger or smaller than the biggest factor from the “inference by enumeration” approach? Is this bigger or smaller than the biggest factor from R-P-C elimination order?

(e) D-Separation


- (i) [5 pts] Which variable (in addition to W and A), if observed, would make $S \perp\!\!\!\perp P$? Please shade in the variable that fits the criteria and run the D-separation algorithm on the resulting graph to support your answer. If no such variable exists, please provide an explanation for why each candidate fails.



(ii) [5 pts] Ivan claims that there exist two unshaded variables (which are NOT directly connected) such that if you know the value of all **other** variables, then those two variables are guaranteed to be independent. Is Ivan right?

Yes, and I will shade in all but two variables in the graph below, and run D-separation algorithm to prove that those two variables are guaranteed to be independent conditioned on all other variables in the graph.

No, there is no such two variables, and I will provide a proof below.

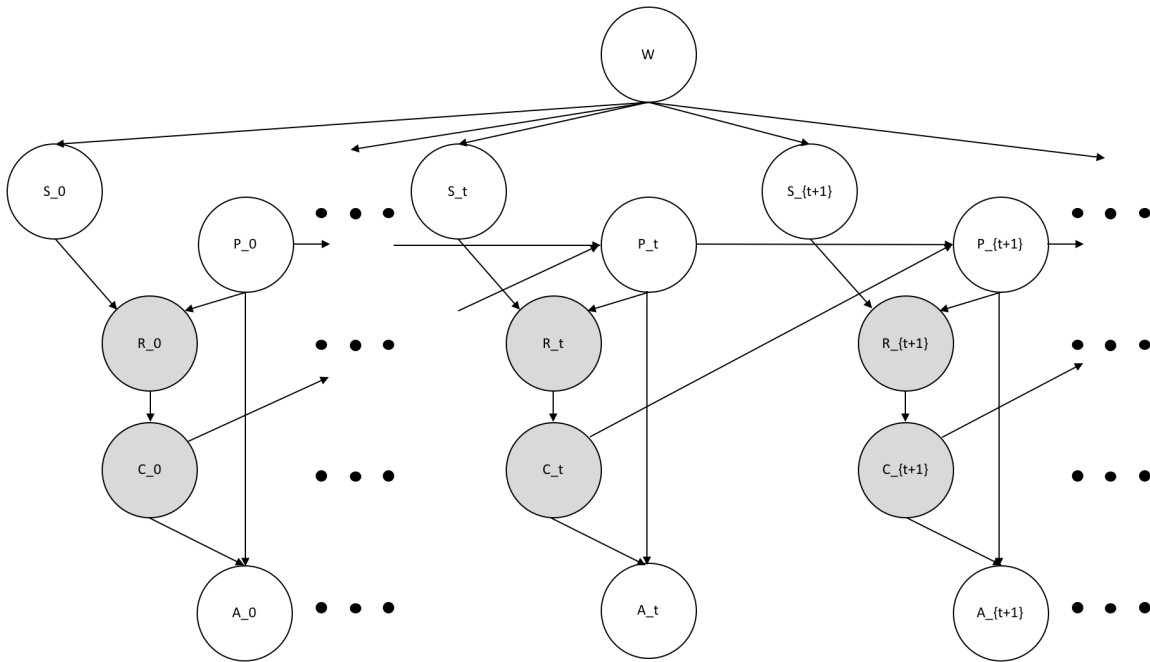


Q2. [36 pts] Quadcopter: Data Analyst

In this question, we look at the setup from the previous problem, but we consider the quadcopter flight over time. Here, flight can be considered in discrete time-steps: $t \in 0, 1, 2, \dots, N - 1$ with, for example, P_t representing the true position P at discrete time-step t . Suppose the weather (W) does not change throughout the quadcopter flight.

One key thing to note here is that there are edges going between time t and time $t + 1$: The true position at time $t + 1$ depends on the true position at time t as well as the control input from time t .

Let's look at this setup from the perspective of Diana, a data analyst who can only **observe** the output from a data-logger, which stores **R (reading of position)** and **C (control from the pilot)**.



(a) Hidden Markov Model

- (i) [4 pts] List all the hidden variables and observed variables in this setup. In a few sentences, how is this setup different from the vanilla Hidden Markov Model you saw in lecture? You should identify at least 2 major differences.

Hidden variables:

Observed variables:

Differences:

- (ii) [3 pts] As a data analyst, Diana's responsibility is to infer the true positions of the quadcopter throughout its flight. In other words, she wants to find a list of true positions $p_0, p_1, p_2, \dots, p_{N-1}$ that are the most likely to have happened, given the recorded readings $r_0, r_1, r_2, \dots, r_{N-1}$ and controls $c_0, c_1, c_2, \dots, c_{N-1}$.

Write down the probability that Diana tries to maximize in terms of a **joint probability**, and interpret the meaning of that probability. Note that the objective that you write below is such that Diana is solving the following problem: $\max_{p_0, p_1, \dots, p_N}$ (maximization objective).

Maximization objective:

Explanation:

- (iii) [3 pts] Morris, a colleague of Diana's, points out that maximizing the joint probability is the same as maximizing a **conditional probability** where all evidence (r_0, r_1, r_2, \dots and c_0, c_1, c_2, \dots) are moved to the right of the conditional bar. Is Morris right?
- Yes, and I will provide a proof/explanation below.
 - No, and I will provide a counter example below.

(b) The Markov Property

- (i) [5 pts] In this setup, conditioned on all observed evidence, does the sequence P_0, P_2, \dots, P_{N-2} follow the Markov property? Please justify your answer.

- (ii) [5 pts] In this setup, conditioned on all observed evidence, does the sequence S_0, S_2, \dots, S_{N-2} follow the Markov property? Please justify your answer.

(c) Forward Algorithm Proxy

Conner, a colleague of Diana's, would like to use this model (with the R_t and C_t observations) to perform something analogous to the forward algorithm for HMMs to infer the true positions. Let's analyze below the effects that certain decisions can have on the outcome of running the forward algorithm.

Note that when we say to **not include** some variable in the algorithm, we mean that we marginalize/sum out that variable. For example, if we do not want to include W in the algorithm, then we replace $P(S_t|W)$ everywhere with $P(S_t)$, where $P(S_t) = \sum_W P(S_t|W)P(W)$.

- (i) [4 pts] He argues that since W (weather) does not depend on time, and is not something he is directly interested in, he does not need to include it in the forward algorithm. What effect does not including W in the forward algorithm have on (a) the accuracy of the resulting belief state calculations, and on (b) the efficiency of calculations? Please justify your answer.

Accuracy:

Efficiency:

- (ii) [3 pts] He also argues that he does not need to include hidden state A (smart alarm warning) in the forward algorithm. What effect does not including A in the forward algorithm have on (a) the accuracy of the resulting belief state calculations, and on (b) the efficiency of calculations? Please justify your answer.

Accuracy:

Efficiency:

- (iii) [3 pts] Last but not least, Conner recalls that for the forward algorithm, one should calculate the belief at time-step t by conditioning on evidence up to $t - 1$, instead of conditioning on evidence from the entire trajectory (up to $N - 1$). Let's assume that some other algorithm allows us to use evidence from the full trajectory ($t = 0$ to $t = N - 1$) in order to infer each belief state. What is an example of a situation (in this setup, with the quadcopter variables) that illustrates that incorporating evidence from the full trajectory can result in better belief states than incorporating evidence only from the prior steps?
- If the signal strength is bad before $t - 1$, but gets better later.
 - If the signal strength is good up to $t - 1$, and the signal is lost later.
 - There isn't such example because using evidence up to $t - 1$ gives us the optimal belief.

(d) Policy Reconstruction

Emily, another colleague of Diana's, would like to use this model to reconstruct the pilot's policy from data. Let's analyze below the effects that certain decisions can have on the outcome of doing policy reconstruction.

- (i) [2 pts] Emily states that the probabilistic model for the pilot's **policy** is entirely captured in one Conditional Probability Table from the Bayes Net Representation. Which table do you think this is, and explain why this table captures the pilot's policy.

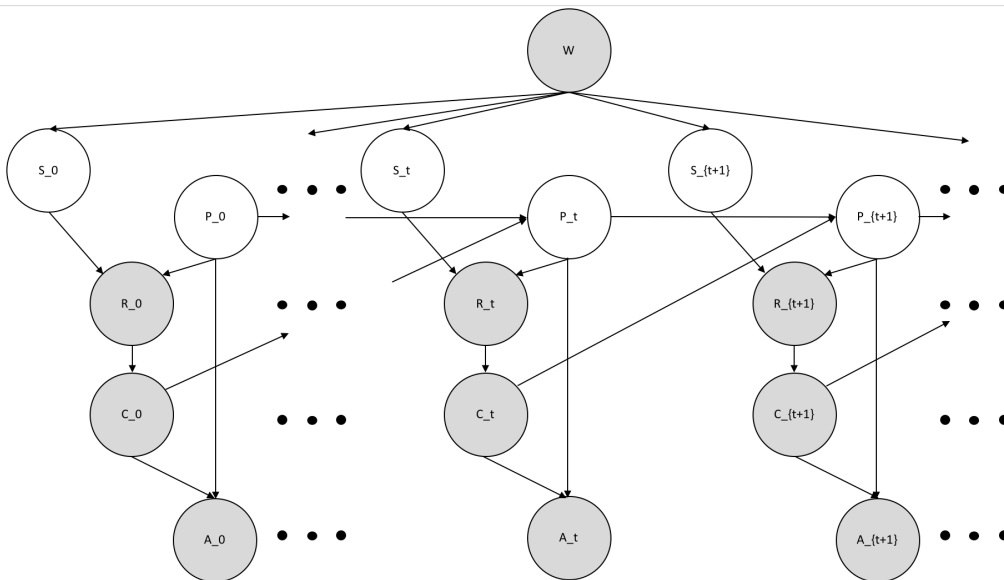
Table:

Explanation:

- (ii) [4 pts] Emily argues that if we were given a lot of data from the data logger, we could reconstruct the probabilistic model for the pilot's policy. Is she right?
- Yes, and I will provide an overview of how to reconstruct the pilot's policy from the data.
 - No, and I will provide a list of reasons for why we cannot reconstruct the policy.

Q3. [27 pts] Quadcopter: Pilot

In this question, we look at same setup from Question 2, 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. A subset of your responses from Q2 might apply here.

(ii) [4 pts] What is the **time-elapsed update** 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 update 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: you can follow the notation from lecture notes.

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

- (iii) [4 pts] What is 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 weak. 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 strong. 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 p , and with probability $1 - p$, 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 p .

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