CS 188 Spring 2021 | Introduction to Artificial Intelligence | Final Review HMMs/VPI

Q1. HMMs

Consider a process where there are transitions among a finite set of states $s_1, \cdots, s_k$ over time steps $i = 1, \cdots, N$. Let the random variables $X_1, \cdots, X_N$ represent the state of the system at each time step and be generated as follows:

- Sample the initial state s from an initial distribution $P_1(X_1)$, and set $i = 1$
- Repeat the following:
  1. Sample a duration d from a duration distribution P_D over the integers $\{1, \cdots, M\}$, where M is the maximum duration.
  2. Remain in the current state s for the next d time steps, i.e., set
  $$x_i = x_{i+1} = \cdots = x_{i+d-1} = s \qquad (1)$$
  3. Sample a successor state s' from a transition distribution $P_T(X_t|X_{t-1} = s)$ over the other states $s' \neq s$ (so there are no self transitions)
  4. Assign $i = i + d$ and $s = s'$.

This process continues indefinitely, but we only observe the first N time steps.

(a) Assuming that all three states s_1, s_2, s_3 are different, what is the probability of the sample sequence $s_1, s_1, s_2, s_2, s_2, s_3, s_3$? Write an algebraic expression. Assume $M \geq 3$.

$$p_1(s_1)p_D(2)p_T(s_2|s_1)p_D(3)p(s_3|s_2)(1 - p_D(1)) \qquad (2)$$

At each time step i we observe a noisy version of the state X_i that we denote Y_i and is produced via a conditional distribution $P_E(Y_i|X_i)$.

(b) Only in this subquestion assume that $N > M$. Let $X_1, \cdots, X_N$ and $Y_1, \cdots, Y_N$ random variables defined as above. What is the maximum index $i \leq N - 1$ so that $X_1 \perp\!\!\!\perp X_N | X_i, X_{i+1}, \cdots, X_{N-1}$ is guaranteed?
$i = N - M$

(c) Only in this subquestion, assume the max duration $M = 2$, and P_D uniform over $\{1, 2\}$ and each x_i is in an alphabet $\{a, b\}$. For $(X_1, X_2, X_3, X_4, X_5, Y_1, Y_2, Y_3, Y_4, Y_5)$ draw a Bayes Net over these 10 random variables with the property that removing any of the edges would yield a Bayes net inconsistent with the given distribution.

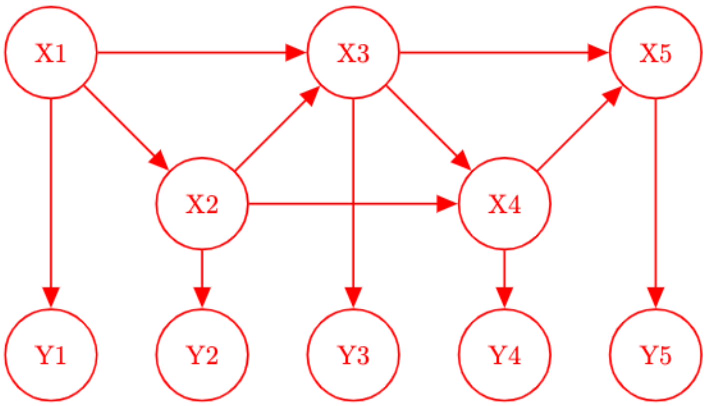

(d) In this part we will explore how to write the described process as an HMM with an extended state space. Write the states $z = (s, t)$ where s is a state of the original system and t represents the time elapsed in that state. For example, the state sequence $s_1, s_1, s_1, s_2, s_3, s_3$ would be represented as $(s_1, 1), (s_1, 2), (s_1, 3), (s_2, 1), (s_3, 1), (s_3, 2)$.

Answer all of the following in terms of the parameters $P_1(X_1), P_D(d), P_T(X_{j+1}|X_j), P_E(Y_i|X_i), k$ (total number of possible states), N and M (max duration).

(i) What is $P(Z_1)$?

$$P(x_1, t) = \begin{cases} P_1(x_1) & \text{if } t = 1 \\ 0 & \text{o.w.} \end{cases} \tag{3}$$

(ii) What is $P(Z_{i+1}|Z_i)$? Hint: You will need to break this into cases where the transition function will behave differently.

$$P(X_{i+1}, t_{i+1}|X_i, t_i) = \begin{cases} P_D(d \geq t_i + 1|d \geq t_i) & \text{when } X_{i+1} = X_i \text{ and } t_{i+1} = t_i + 1 \text{ and } t_{i+1} \leq M \\ P_T(X_{i+1}|X_i)P_D(d = t_i|d \geq t_i) & \text{when } X_{i+1} \neq X_i \text{ and } t_{i+1} = 1 \\ 0 & \text{o.w.} \end{cases}$$

Where $P_D(d \geq t_i + 1|d \geq t_i) = P_D(d \geq t_i + 1)/P_D(d \geq t_i)$.
Being in X_i, t_i, we know that d was drawn $d \geq t_i$. Conditioning on this fact, we have two choices, if $d > t_i$ then the next state is $X_{i+1} = X_i$, and if $d = t_i$ then $X_{i+1} \neq X_i$ drawn from the transition distribution and $t_{i+1} = 1$.
(4)

(iii) What is $P(Y_i|Z_i)$?
$p(Y_i|X_i, t_i) = P_E(Y_i|X_i)$

(e) In this question we explore how to write an algorithm to compute $P(X_N|y_1, \cdots, y_N)$ using the particular structure of this process.

Write $P(X_t|y_1, \cdots, y_{t-1})$ in terms of other factors. Construct an answer by checking the correct boxes below:

$$P(X_t|y_1, \cdots, y_{t-1}) = \underline{\quad\textbf{(i)}\quad} \;\underline{\quad\textbf{(ii)}\quad}\; \underline{\quad\textbf{(iii)}\quad}$$

(i)
- ● $\sum_{i=1}^{k} \sum_{d=1}^{M} \sum_{d'=1}^{M}$
- ○ $\sum_{i=1}^{k} \sum_{d=1}^{M}$
- ○ $\sum_{i=1}^{k}$
- ○ $\sum_{d=1}^{M}$

(ii)
- ○ $P(Z_t = (X_t, d)|Z_{t-1} = (s_i, d))$
- ○ $P(X_t|X_{t-1} = s_i)$
- ○ $P(X_t|X_{t-1} = s_d)$
- ● $P(Z_t = (X_t, d')|Z_{t-1} = (s_i, d))$

(iii)
- ○ $P(Z_{t-1} = (s_d, i)|y_1, \cdots, y_{t-1})$
- ○ $P(X_{t-1} = s_d|y_1, \cdots, y_{t-1})$
- ● $P(Z_{t-1} = (s_i, d)|y_1, \cdots, y_{t-1})$
- ○ $P(X_{t-1} = s_i|y_1, \cdots, y_{t-1})$

Q2. Value of Perfect Information

Consider the setup shown in the figure below, involving a robotic plant-watering system with some mysterious random forces involved. Here, there are 4 main items at play.
(1) The robot (R) can choose to move either left (l) or right (r). Its chosen action pushes a water pellet into the corresponding opening.
(2) The random switch (S) is arbitrarily in one of two possible positions $\{s_0, s_1\}$. When in position (s_0), it accepts a water pellet only from the (l) tube. When in position (s_1), it accepts a water pellet only from the (r) tube.
(3) A controllable three-way switch (T) can be chosen to be placed in one of three possible positions $\{t_0, t_1, t_2\}$.
(4) A plant (P) is arbitrarily located in one of three possible locations $\{p_0, p_1, p_2\}$. When in position p_i, it can only be successfully watered if the corresponding tube t_i has been selected **and** if the water pellet was sent in a direction that was indeed accepted by the first switch (S).

Finally, in this problem, utility (U) is 1 when the plant successfully receives the water pellet, and 0 otherwise.

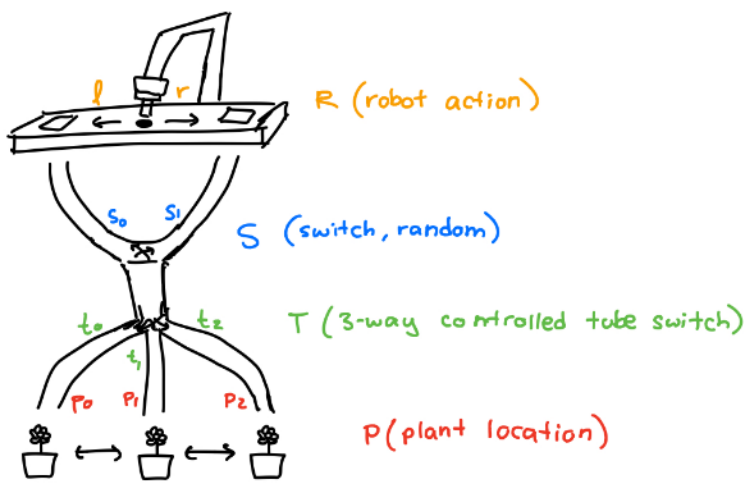

(a) Let's first set this problem up as a decision network.

(i) Which of the following decision networks correctly describe the problem described above? Select all that apply. Recall the conventions from the lecture notes:

action nodes as rectangles 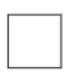, chance nodes as ovals , and utility nodes as diamonds 

.

■ (A)

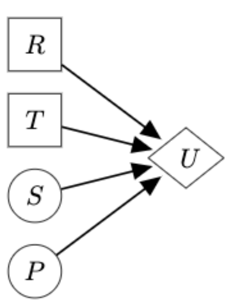

☐ (B)

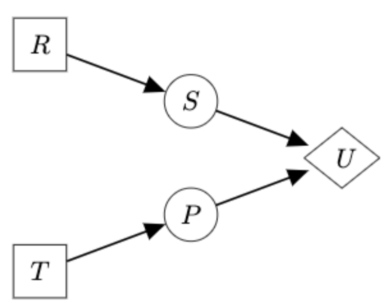

☐ (C)

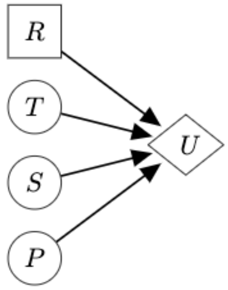

☐ (D)

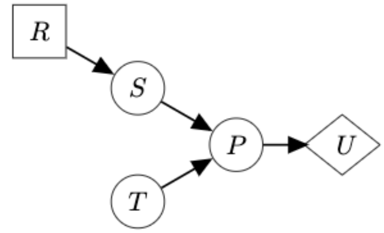

R and T are actions you can choose so they should be rectangles, S and P are random outcomes so they should be ovals. All the variables are involved in the calculation of U but they do not influence each other directly, so A has to be the answer.

(ii) Fill in the following probability tables, given that there is an equal chance of being at each of their possible locations.

S	$P(S)$
s_0	$\frac{1}{2}$
s_1	$\frac{1}{2}$

P	$P(P)$
p_0	$\frac{1}{3}$
p_1	$\frac{1}{3}$
p_2	$\frac{1}{3}$

(b) Before selecting your actions, suppose that someone could tell you the value of either S or P. Follow the steps below to calculate the **maximum expected utility (MEU)** when knowing S, or when knowing P. Then, decide which one you would prefer to be told.

(i) What is $MEU(S)$?

○ 0 ○ $\frac{1}{9}$ ○ $\frac{1}{6}$ ○ $\frac{1}{4}$ ● $\frac{1}{3}$ ○ $\frac{1}{2}$
○ $\frac{2}{3}$ ○ $\frac{3}{4}$ ○ $\frac{5}{6}$ ○ 1 ○ None of the above

Note: can definitely answer this with intuition (and no math).
$= \frac{1}{2}MEU(S = s_0) + \frac{1}{2}MEU(S = s_1)$
$= \frac{1}{2}(max_t(EU(S = s_0, T = t_0), EU(S = s_0, T = t_1), EU(S = s_0, T = t_2)))$
$+ \frac{1}{2}(max_t(EU(S = s_1, T = t_0), EU(S = s_1, T = t_1), EU(S = s_1, T = t_2)))$
$= \frac{1}{2}(max(1/3, 1/3, 1/3)) + \frac{1}{2}(max(1/3, 1/3, 1/3))$
$= \frac{1}{3}$

(ii) What is $MEU(P)$?

○ 0 ○ $\frac{1}{9}$ ○ $\frac{1}{6}$ ○ $\frac{1}{4}$ ○ $\frac{1}{3}$ ● $\frac{1}{2}$
○ $\frac{2}{3}$ ○ $\frac{3}{4}$ ○ $\frac{5}{6}$ ○ 1 ○ None of the above

Note: can definitely answer this with intuition (and no math).
$= \frac{1}{3}MEU(P = p_0) + \frac{1}{3}MEU(P = p_1) + \frac{1}{3}MEU(P = p_2)$
$= \frac{1}{3}(max_R(EU(P = p_0, R = l), EU(P = p_0, R = r))) +$
$\frac{1}{3}(max_R(EU(P = p_1, R = l), EU(P = p_1, R = r))) +$
$\frac{1}{3}(max_R(EU(P = p_2, R = l), EU(P = p_2, R = r)))$
$= \frac{1}{3}(max(1/2, 1/2)) +$
$\frac{1}{3}(max(1/2, 1/2)) +$
$\frac{1}{3}(max(1/2, 1/2))$
$= \frac{1}{2}$

(iii) Would you prefer to be told S or P? ○ S ● P

P since it has a higher MEU (and therefore a higher VPI).

(c) **(i)** What is $MEU(S, P)$?

○ 0 ○ $\frac{1}{9}$ ○ $\frac{1}{6}$ ○ $\frac{1}{4}$ ○ $\frac{1}{3}$ ○ $\frac{1}{2}$
○ $\frac{2}{3}$ ○ $\frac{3}{4}$ ○ $\frac{5}{6}$ ● 1 ○ None of the above

1, because you have enough information to definitely get the water pellet to the plant.

(ii) In this problem, does $VPI(S, P) = VPI(S) + VPI(P)$? ○ Yes ● No

$VPI(S, P) = MEU(S, P) - MEU(none)$

$VPI(S) = MEU(S) - MEU(none)$
$VPI(P) = MEU(P) - MEU(none)$
Answer is NO, because $MEU(S, P) = 1$, $MEU(S) = \frac{1}{3}$, and $MEU(P) = \frac{1}{2}$.

(iii) In general, does $VPI(a, b) = VPI(a) + VPI(b)$? Select all of the statements below which are true.

○ Yes, because of the additive property.
○ Yes, because the order in which we observe the variables does not matter.
○ Yes, but the reason is not listed.
● No, because the value of knowing each variable can be dependent on whether or not we know the other one.
○ No, because the order in which we observe the variables matters.
○ No, but the reason is not listed.

(d) For each of the following new variables introduced to this problem, what would the corresponding VPI of that variable be?

(i) A new variable X indicates the weather outside, which affects the overall health of the plant.

○ VPI(X)< 0 ● VPI(X)= 0 ○ VPI(X)> 0

The health of the plant does not affect the utility so the VPI is 0.

(ii) A new variable X indicates the weather outside, which affects the metal of switch S such that when it's hot outside, the switch is most likely to remain in position s_0 with probability 0.9 (and goes to s_1 with probability 0.1).

○ VPI(X)< 0 ○ VPI(X)= 0 ● VPI(X)> 0

This will allow us to predict which direction to move the robot with more accuracy so the VPI is greater than 0.