
Introduction to Artificial Intelligence Review HMMs Solutions

Q1. HMMs

Consider a process where there are transitions among a finite set of states s_1, \dots, s_k over time steps $i = 1, \dots, N$. Let the random variables X_1, \dots, 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 $\{1, \dots, M\}$, where M is the maximum duration.
 2. Remain in the current state s for the next d time steps, i.e., set
$$x_i = x_{i+1} = \dots = x_{i+d-1} = s \quad (1)$$
 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 $M \geq 3$.

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

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, \dots, X_N and Y_1, \dots, Y_N random variables defined as above. What is the maximum index $i \leq N - 1$ so that $X_1 \perp\!\!\!\perp X_N | X_i, X_{i+1}, \dots, X_{N-1}$ 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 $P(Z_1)$?

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

- (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 $P(X_N|y_1, \dots, y_N)$ using the particular structure of this process.

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

$P(X_t y_1, \dots, y_{t-1}) =$		<div style="display: inline-block; width: 150px; border-bottom: 1px solid black; text-align: center;">(i)</div> <div style="display: inline-block; width: 150px; border-bottom: 1px solid black; text-align: center;">(ii)</div> <div style="display: inline-block; width: 150px; border-bottom: 1px solid black; text-align: center;">(iii)</div>
(i)	<input checked="" type="radio"/> $\sum_{i=1}^k \sum_{d=1}^M \sum_{d'=1}^M$ <input type="radio"/> $\sum_{i=1}^k \sum_{d=1}^M$	<input type="radio"/> $\sum_{i=1}^k$ <input type="radio"/> $\sum_{d=1}^M$
(ii)	<input type="radio"/> $P(Z_t = (X_t, d) Z_{t-1} = (s_i, d))$ <input type="radio"/> $P(X_t X_{t-1} = s_i)$	<input type="radio"/> $P(X_t X_{t-1} = s_d)$ <input checked="" type="radio"/> $P(Z_t = (X_t, d') Z_{t-1} = (s_i, d))$
(iii)	<input type="radio"/> $P(Z_{t-1} = (s_d, i) y_1, \dots, y_{t-1})$ <input type="radio"/> $P(X_{t-1} = s_d y_1, \dots, y_{t-1})$	<input checked="" type="radio"/> $P(Z_{t-1} = (s_i, d) y_1, \dots, y_{t-1})$ <input type="radio"/> $P(X_{t-1} = s_i y_1, \dots, y_{t-1})$

Q2. Planning ahead with HMMs

Pacman is tired of using HMMs to estimate the location of ghosts. He wants to use HMMs to plan what actions to take in order to maximize his utility. Pacman uses the HMM (drawn to the right) of length T to model the planning problem. In the HMM, $X_{1:T}$ is the sequence of hidden states of Pacman's world, $A_{1:T}$ are actions Pacman can take, and U_t is the utility Pacman receives at the particular hidden state X_t . Notice that there are no evidence variables, and utilities are not discounted.



- (a) The belief at time t is defined as $B_t(X_t) = p(X_t|a_{1:t})$. The forward algorithm update has the following form:

$$B_t(X_t) = \underline{\hspace{1cm}} \text{ (i) } \underline{\hspace{1cm}} \text{ (ii) } \underline{\hspace{1cm}} B_{t-1}(x_{t-1}).$$

Complete the expression by choosing the option that fills in each blank.

- (i) ☐ $\max_{x_{t-1}}$ ☒ $\sum_{x_{t-1}}$ ☐ \max_{x_t} ☐ \sum_{x_t} ☐ 1
(ii) ☐ $p(X_t|x_{t-1})$ ☐ $p(X_t|x_{t-1})p(X_t|a_t)$ ☐ $p(X_t)$ ☒ $p(X_t|x_{t-1}, a_t)$ ☐ 1
☐ None of the above combinations is correct

$$\begin{aligned} B_t(X_t) &= p(X_t|a_{1:t}) \\ &= \sum_{x_{t-1}} p(X_t|x_{t-1}, a_t) p(x_{t-1}|a_{1:t-1}) \\ &= \sum_{x_{t-1}} p(X_t|x_{t-1}, a_t) B_{t-1}(x_{t-1}) \end{aligned}$$

- (b) Pacman would like to take actions $A_{1:T}$ that maximizes the expected sum of utilities, which has the following form:

$$\text{MEU}_{1:T} = \underline{\hspace{1cm}} \text{ (i) } \underline{\hspace{1cm}} \text{ (ii) } \underline{\hspace{1cm}} \text{ (iii) } \underline{\hspace{1cm}} \text{ (iv) } \underline{\hspace{1cm}} \text{ (v) }$$

Complete the expression by choosing the option that fills in each blank.

- (i) ☒ $\max_{a_{1:T}}$ ☐ \max_{a_T} ☐ $\sum_{a_{1:T}}$ ☐ \sum_{a_T} ☐ 1
(ii) ☐ \max_t ☐ $\prod_{t=1}^T$ ☒ $\sum_{t=1}^T$ ☐ \min_t ☐ 1
(iii) ☐ \sum_{x_t, a_t} ☒ \sum_{x_t} ☐ \sum_{a_t} ☐ \sum_{x_T} ☐ 1
(iv) ☐ $p(x_t|x_{t-1}, a_t)$ ☐ $p(x_t)$ ☒ $B_t(x_t)$ ☐ $B_T(x_T)$ ☐ 1
(v) ☐ U_T ☐ $\frac{1}{U_t}$ ☐ $\frac{1}{U_T}$ ☒ U_t ☐ 1
☐ None of the above combinations is correct

$$\text{MEU}_{1:T} = \max_{a_{1:T}} \sum_{t=1}^T \sum_{x_t} B_t(x_t) U_t(x_t)$$

- (c) A greedy ghost now offers to tell Pacman the values of some of the hidden states. Pacman needs your help to figure out if the ghost's information is useful. Assume that the transition function $p(x_t|x_{t-1}, a_t)$ is not deterministic. **With respect to the utility** U_t , mark all that can be True:

☒ $VPI(X_{t-1}|X_{t-2}) > 0$ ☐ $VPI(X_{t-2}|X_{t-1}) > 0$ ☒ $VPI(X_{t-1}|X_{t-2}) = 0$ ☒ $VPI(X_{t-2}|X_{t-1}) = 0$
☐ None of the above

It is always possible that $VPI = 0$. Can guarantee $VPI(E|e)$ is not greater than 0 if E is independent of $parents(U)$ given e .

- (d) Pacman notices that calculating the beliefs under this model is very slow using exact inference. He therefore decides to try out various particle filter methods to speed up inference. Order the following methods by how accurate their estimate of $B_T(X_T)$ is? If different methods give an equivalently accurate estimate, mark them as the same number.


	Most accurate		Least accurate	
Exact inference	<input checked="" type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4
Particle filtering with no resampling	<input type="radio"/> 1	<input checked="" type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4
Particle filtering with resampling before every time elapse	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input checked="" type="radio"/> 4
Particle filtering with resampling before every other time elapse	<input type="radio"/> 1	<input type="radio"/> 2	<input checked="" type="radio"/> 3	<input type="radio"/> 4

Exact inference will always be more accurate than using a particle filter. When comparing the particle filter resampling approaches, notice that because there are no observations, each particle will have weight 1. Therefore resampling when particle weights are 1 could lead to particles being lost and hence prove bad.

Q3. Particle Filtering

You've chased your arch-nemesis Leland to the Stanford quad. You enlist two robo-watchmen to help find him! The grid below shows the campus, with ID numbers to label each region. Leland will be moving around the campus. His location at time step t will be represented by random variable X_t . Your robo-watchmen will also be on campus, but their locations will be fixed. Robot 1 is always in region 1 and robot 2 is always in region 9. (See the * locations on the map.) At each time step, each robot gives you a sensor reading to help you determine where Leland is. The sensor reading of robot 1 at time step t is represented by the random variable $E_{t,1}$. Similarly, robot 2's sensor reading at time step t is $E_{t,2}$. The Bayes' Net to the right shows your model of Leland's location and your robots' sensor readings.

1*	2	3	4	5
6	7	8	9*	10
11	12	13	14	15



In each time step, Leland will either stay in the same region or move to an adjacent region. For example, the available actions from region 4 are (WEST, EAST, SOUTH, STAY). He chooses between all available actions with equal probability, regardless of where your robots are. Note: moving off the grid is not considered an available action.

Each robot will detect if Leland is in an adjacent region. For example, the regions adjacent to region 1 are 1, 2, and 6. If Leland is in an adjacent region, then the robot will report *NEAR* with probability 0.8. If Leland is not in an adjacent region, then the robot will still report *NEAR*, but with probability 0.3.

For example, if Leland is in region 1 at time step t the probability tables are:

E	$P(E_{t,1} X_t = 1)$	$P(E_{t,2} X_t = 1)$
<i>NEAR</i>	0.8	0.3
<i>FAR</i>	0.2	0.7

- (a) Suppose we are running particle filtering to track Leland's location, and we start at $t = 0$ with particles $[X = 6, X = 14, X = 9, X = 6]$. Apply a forward simulation update to each of the particles using the random numbers in the table below.

Assign region IDs to sample spaces in numerical order. For example, if, for a particular particle, there were three possible successor regions 10, 14 and 15, with associated probabilities, $P(X = 10)$, $P(X = 14)$ and $P(X = 15)$, and the random number was 0.6, then 10 should be selected if $0.6 \leq P(X = 10)$, 14 should be selected if $P(X = 10) < 0.6 < P(X = 10) + P(X = 14)$, and 15 should be selected otherwise.

Particle at $t = 0$	Random number for update	Particle after forward simulation update
$X = 6$	0.864	11
$X = 14$	0.178	9
$X = 9$	0.956	14
$X = 6$	0.790	11

- (b) Some time passes and you now have particles $[X = 6, X = 1, X = 7, X = 8]$ at the particular time step, but you have not yet incorporated your sensor readings at that time step. Your robots are still in regions 1 and 9, and both report *NEAR*. What weight do we assign to each particle in order to incorporate this evidence?

Particle	Weight
$X = 6$	$0.8 * 0.3$
$X = 1$	$0.8 * 0.3$
$X = 7$	$0.3 * 0.3$
$X = 8$	$0.3 * 0.8$

- (c) To decouple this question from the previous question, let's say you just incorporated the sensor readings and found the following weights for each particle (these are not the correct answers to the previous problem!):

Particle	Weight
$X = 6$	0.1
$X = 1$	0.4
$X = 7$	0.1
$X = 8$	0.2

Normalizing gives us the distribution

$$\begin{aligned}
 X = 1 : 0.4/0.8 &= 0.5 \\
 X = 6 : 0.1/0.8 &= 0.125 \\
 X = 7 : 0.1/0.8 &= 0.125 \\
 X = 8 : 0.2/0.8 &= 0.25
 \end{aligned}$$

Use the following random numbers to resample your particles. As on the previous page, **assign region IDs to sample spaces in numerical order**.

Random number:	0.596	0.289	0.058	0.765
Particle:	6	1	1	8