

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 $P_T(X_i|X_{i-1} = s)$ 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 $M \geq 3$.

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?
- (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 $P(Z_1)$?

$$P(x_1, t_1) =$$

(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.

$$P(X_{i+1}, t_{i+1} | X_i, t_i) =$$

(iii) What is $P(Y_i|Z_i)$?

$$P(Y_i | X_i, t_i) =$$

(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:

$$P(X_t|y_1, \dots, y_{t-1}) = \underline{\hspace{2cm} \text{(i)} \hspace{2cm}} \quad \underline{\hspace{2cm} \text{(ii)} \hspace{2cm}} \quad \underline{\hspace{2cm} \text{(iii)} \hspace{2cm}}$$

(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) $P(Z_{t-1} = (s_d, i) | y_1, \dots, y_{t-1})$
 $P(X_{t-1} = s_d | y_1, \dots, y_{t-1})$

$P(Z_{t-1} = (s_i, d) | y_1, \dots, y_{t-1})$
 $P(X_{t-1} = s_i | y_1, \dots, y_{t-1})$

Q2. HMMs: Help Your House Help You

Imagine you have a smart house that wants to track your location within itself so it can turn on the lights in the room you are in and make you food in your kitchen. Your house has 4 rooms (A, B, C, D) in the floorplan below (A is connected to B and D , B is connected to A and C , C is connected to B and D , and D is connected to A and C):



At the beginning of the day ($t = 0$), your probabilities of being in each room are p_A, p_B, p_C , and p_D for rooms A, B, C , and D , respectively, and at each time t your position (following a Markovian process) is given by X_t . At each time, your probability of staying in the same room is q_0 , your probability of moving clockwise to the next room is q_1 , and your probability of moving counterclockwise to the next room is $q_{-1} = 1 - q_0 - q_1$.

(a) Initially, assume your house has no way of sensing where you are. What is the probability that you will be in room D at time $t = 1$?

- $q_0 p_D$
 $q_0 p_D + q_1 p_A + q_{-1} p_C + 2q_1 p_B$
 $q_0 p_D + q_1 p_A + q_{-1} p_C$
 $q_0 p_D + q_{-1} p_A + q_1 p_C$
 $q_1 p_A + q_1 p_C + q_0 p_D$
 None of these

Now assume your house contains a sensor M^A that detects motion ($+m$) or no motion ($-m$) in room A . However, the sensor is a bit noisy and can be tricked by movement in adjacent rooms, resulting in the conditional distributions for the sensor given in the table below. The prior distribution for the sensor's output is also given.

M^A	$P(M^A X = A)$	$P(M^A X = B)$	$P(M^A X = C)$	$P(M^A X = D)$	M^A	$P(M^A)$
$+m^A$	$1 - 2\gamma$	γ	0.0	γ	$+m^A$	0.5
$-m^A$	2γ	$1 - \gamma$	1.0	$1 - \gamma$	$-m^A$	0.5

(b) You decide to help your house to track your movements using a particle filter with three particles. At time $t = T$, the particles are at $X^0 = A, X^1 = B, X^2 = D$. What is the probability that the particles will be resampled as $X^0 = X^1 = X^2 = A$ after time elapse? Select **all terms in the product**.

- q_0
 q_0^2
 q_0^3
 q_1
 q_1^2
 q_1^3
 q_{-1}
 q_{-1}^2
 q_{-1}^3
 None of these

(c) Assume that the particles are actually resampled after time elapse as $X^0 = D, X^1 = B, X^2 = C$, and the sensor observes $M^A = -m^A$. What are the particle weights given the observation?

Particle	Weight							
$X^0 = D$	<input type="radio"/> γ	<input type="radio"/> $1 - \gamma$	<input type="radio"/> $1 - 2\gamma$	<input type="radio"/> 0.0	<input type="radio"/> 1.0	<input type="radio"/> 2γ	<input type="radio"/> None of these	
$X^1 = B$	<input type="radio"/> γ	<input type="radio"/> $1 - \gamma$	<input type="radio"/> $1 - 2\gamma$	<input type="radio"/> 0.0	<input type="radio"/> 1.0	<input type="radio"/> 2γ	<input type="radio"/> None of these	
$X^2 = C$	<input type="radio"/> γ	<input type="radio"/> $1 - \gamma$	<input type="radio"/> $1 - 2\gamma$	<input type="radio"/> 0.0	<input type="radio"/> 1.0	<input type="radio"/> 2γ	<input type="radio"/> None of these	

Now, assume your house also contains sensors M^B and M^D in rooms B and D , respectively, with the conditional distributions of the sensors given below and the prior equivalent to that of sensor M^A .

M^B	$P(M^B X = A)$	$P(M^B X = B)$	$P(M^B X = C)$	$P(M^B X = D)$
$+m^B$	γ	$1 - 2\gamma$	γ	0.0
$-m^B$	$1 - \gamma$	2γ	$1 - \gamma$	1.0

M^D	$P(M^D X = A)$	$P(M^D X = B)$	$P(M^D X = C)$	$P(M^D X = D)$
$+m^D$	γ	0.0	γ	$1 - 2\gamma$
$-m^D$	$1 - \gamma$	1.0	$1 - \gamma$	2γ

(d) Again, assume that the particles are actually resampled after time elapse as $X^0 = D, X^1 = B, X^2 = C$. The sensor readings are now $M^A = -m^A, M^B = -m^B, M^D = +m^D$. What are the particle weights given the observations?

Particle	Weight
$X^0 = D$	<input type="radio"/> $\gamma^2 - 2\gamma^3$ <input type="radio"/> $3 - 2\gamma$ <input type="radio"/> 0.0 <input type="radio"/> $\gamma - \gamma^2 + \gamma^3$ <input type="radio"/> $1 - 3\gamma + 2\gamma^2$ <input type="radio"/> $2 - \gamma$ <input type="radio"/> $1 - 2\gamma + \gamma^2$ <input type="radio"/> None of these
$X^1 = B$	<input type="radio"/> $\gamma^2 - 2\gamma^3$ <input type="radio"/> $3 - 2\gamma$ <input type="radio"/> 0.0 <input type="radio"/> $\gamma - \gamma^2 + \gamma^3$ <input type="radio"/> $1 - 3\gamma + 2\gamma^2$ <input type="radio"/> $2 - \gamma$ <input type="radio"/> $1 - 2\gamma + \gamma^2$ <input type="radio"/> None of these
$X^2 = C$	<input type="radio"/> $\gamma^2 - 2\gamma^3$ <input type="radio"/> $3 - 2\gamma$ <input type="radio"/> 0.0 <input type="radio"/> $\gamma - \gamma^2 + \gamma^3$ <input type="radio"/> $1 - 3\gamma + 2\gamma^2$ <input type="radio"/> $2 - \gamma$ <input type="radio"/> $1 - 2\gamma + \gamma^2$ <input type="radio"/> None of these

The sequence of observations from each sensor are expressed as the following: $m_{0:t}^A$ are all measurements $m_0^A, m_1^A, \dots, m_t^A$ from sensor M^A , $m_{0:t}^B$ are all measurements $m_0^B, m_1^B, \dots, m_t^B$ from sensor M^B , and $m_{0:t}^D$ are all measurements $m_0^D, m_1^D, \dots, m_t^D$ from sensor M^D . Your house can get an accurate estimate of where you are at a given time t using the forward algorithm. The forward algorithm update step is shown here:

$$P(X_t | m_{0:t}^A, m_{0:t}^B, m_{0:t}^D) \propto P(X_t, m_{0:t}^A, m_{0:t}^B, m_{0:t}^D) \quad (2)$$

$$= \sum_{x_{t-1}} P(X_t, x_{t-1}, m_t^A, m_t^B, m_t^D, m_{0:t-1}^A, m_{0:t-1}^B, m_{0:t-1}^D) \quad (3)$$

$$= \sum_{x_{t-1}} \frac{P(X_t | x_{t-1}) P(x_{t-1}, m_{0:t-1}^A, m_{0:t-1}^B, m_{0:t-1}^D)}{P(x_{t-1})} \quad (4)$$