


## Q1. HMM: Where is the Car?

Transportation researchers are trying to improve traffic in the city but, in order to do that, they first need to estimate the location of each of the cars in the city. They need our help to model this problem as an inference problem of an HMM. For this question, assume that only *one* car is being modeled.

- (a) The structure of this modified HMM is given below, which includes  $X$ , the location of the car;  $S$ , the noisy location of the car from the signal strength at a nearby cell phone tower; and  $G$ , the noisy location of the car from GPS.



We want to perform filtering with this HMM. That is, we want to compute the belief  $P(x_t | s_{1:t}, g_{1:t})$ , the probability of a state  $x_t$  given all past and current observations.

The **dynamics update** expression has the following form:

$$P(x_t | s_{1:t-1}, g_{1:t-1}) = \underline{\hspace{2cm} \text{(i)} \hspace{2cm}} \underline{\hspace{2cm} \text{(ii)} \hspace{2cm}} \underline{\hspace{2cm} \text{(iii)} \hspace{2cm}} P(x_{t-1} | s_{1:t-1}, g_{1:t-1}).$$

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

- (i)       $P(s_{1:t-1}, g_{1:t-1})$       $P(s_{1:t}, g_{1:t})$       $P(s_{1:t-1})P(g_{1:t-1})$       $P(s_{1:t})P(g_{1:t})$      1  
 (ii)      $\sum_{x_{t-1}}$                       $\sum_{x_t}$                       $\max_{x_{t-1}}$                       $\max_{x_t}$                      1  
 (iii)      $P(x_{t-2}, x_{t-1})$       $P(x_{t-1} | x_{t-2})$       $P(x_{t-1}, x_t)$               $P(x_t | x_{t-1})$              1


The **observation update** expression has the following form:

$$P(x_t | s_{1:t}, g_{1:t}) = \underline{\hspace{2cm} \text{(iv)} \hspace{2cm}} \underline{\hspace{2cm} \text{(v)} \hspace{2cm}} \underline{\hspace{2cm} \text{(vi)} \hspace{2cm}} P(x_t | s_{1:t-1}, g_{1:t-1}).$$

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


- (iv)       $P(s_t, g_t | s_{1:t-1}, g_{1:t-1})$       $P(s_{1:t-1}, g_{1:t-1} | s_t, g_t)$       $P(s_t | s_{1:t-1})P(g_t | g_{1:t-1})$   
             $P(s_{1:t-1} | s_t)P(g_{1:t-1} | g_t)$       $\frac{1}{P(s_t, g_t | s_{1:t-1}, g_{1:t-1})}$       $\frac{1}{P(s_{1:t-1}, g_{1:t-1} | s_t, g_t)}$   
             $\frac{1}{P(s_t | s_{1:t-1})P(g_t | g_{1:t-1})}$       $\frac{1}{P(s_{1:t-1} | s_t)P(g_{1:t-1} | g_t)}$      1  
 (v)       $\sum_{x_{t-1}}$                       $\sum_{x_t}$                       $\max_{x_{t-1}}$                       $\max_{x_t}$                      1  
 (vi)      $P(x_{t-1}, s_{t-1})P(x_{t-1}, g_{t-1})$       $P(x_{t-1}, s_{t-1}, g_{t-1})$       $P(x_t | s_t)P(x_t | g_t)$   
             $P(s_{t-1} | x_{t-1})P(g_{t-1} | x_{t-1})$       $P(x_t, s_t)P(x_t, g_t)$               $P(x_t, s_t, g_t)$   
             $P(x_{t-1} | s_{t-1})P(x_{t-1} | g_{t-1})$       $P(s_t | x_t)P(g_t | x_t)$              1

- (b) It turns out that if the car moves too fast, the quality of the cell phone signal decreases. Thus, the signal-dependent location  $S_t$  not only depends on the current state  $X_t$  but it also depends on the previous state  $X_{t-1}$ . Thus, we modify our original HMM for a new more accurate one, which is given below.



Again, we want to compute the belief  $P(x_t | s_{1:t}, g_{1:t})$ . In this part we consider an update that combines the dynamics and observation update in a *single* update.

For convenience, the HMM from part (b) from the last page is redrawn below.




$$P(x_t | s_{1:t}, g_{1:t}) = \underline{\hspace{1cm} \text{(i)} \hspace{1cm}} \underline{\hspace{1cm} \text{(ii)} \hspace{1cm}} \underline{\hspace{1cm} \text{(iii)} \hspace{1cm}} \underline{\hspace{1cm} \text{(iv)} \hspace{1cm}} P(x_{t-1} | s_{1:t-1}, g_{1:t-1}).$$

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

- (i)   $P(s_t, g_t | s_{1:t-1}, g_{1:t-1})$       $P(s_{1:t-1}, g_{1:t-1} | s_t, g_t)$       $P(s_t | s_{1:t-1})P(g_t | g_{1:t-1})$   
  $\frac{1}{P(s_{1:t-1}, g_{1:t-1} | s_t, g_t)}$       $\frac{1}{P(s_t, g_t | s_{1:t-1}, g_{1:t-1})}$       $P(s_{1:t-1} | s_t)P(g_{1:t-1} | g_t)$   
  $\frac{1}{P(s_t | s_{1:t-1})P(g_t | g_{1:t-1})}$       $\frac{1}{P(s_{1:t-1} | s_t)P(g_{1:t-1} | g_t)}$      1
- (ii)   $\sum_{x_{t-1}}$       $\sum_{x_t}$       $\max_{x_{t-1}}$       $\max_{x_t}$      1
- (iii)   $P(x_{t-2}, x_{t-1}, s_{t-1})P(x_{t-1}, g_{t-1})$       $P(x_{t-1}, x_t, s_t)P(x_t, g_t)$       $P(s_{t-1}, g_{t-1} | x_{t-1})$   
  $P(s_{t-1} | x_{t-2}, x_{t-1})P(g_{t-1} | x_{t-1})$       $P(s_t | x_{t-1}, x_t)P(g_t | x_t)$       $P(s_t, g_t | x_t)$   
  $P(x_{t-2}, x_{t-1} | s_{t-1})P(x_{t-1} | g_{t-1})$       $P(x_{t-1}, x_t | s_t)P(x_t | g_t)$      1  
  $P(x_{t-2}, x_{t-1}, s_{t-1}, g_{t-1})$       $P(x_{t-1}, x_t, s_t, g_t)$
- (iv)   $P(x_{t-1}, x_t)$       $P(x_t | x_{t-1})$       $P(x_{t-2}, x_{t-1})$       $P(x_{t-1} | x_{t-2})$      1

(c) The Viterbi algorithm finds the most probable sequence of hidden states  $X_{1:T}$ , given a sequence of observations  $s_{1:T}$ , for some time  $t = T$ . Recall the canonical HMM structure, which is shown below.



For this canonical HMM, the Viterbi algorithm performs the following dynamic programming computations:

$$m_t[x_t] = P(s_t|x_t) \max_{x_{t-1}} P(x_t|x_{t-1})m_{t-1}[x_{t-1}].$$

We consider extending the Viterbi algorithm for the modified HMM from part **(b)**. We want to find the most likely sequence of states  $X_{1:T}$  given the sequence of observations  $s_{1:T}$  and  $g_{1:T}$ . The dynamic programming update for  $t > 1$  for the modified HMM has the following form:

$$m_t[x_t] = \text{(i)} \quad \text{(ii)} \quad \text{(iii)} \quad m_{t-1}[x_{t-1}].$$


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

- (i)**        $\sum_{x_{t-1}}$         $\sum_{x_t}$         $\max_{x_{t-1}}$         $\max_{x_t}$        1
- (ii)**        $P(x_{t-2}, x_{t-1}, s_{t-1})P(x_{t-1}, g_{t-1})$         $P(x_{t-1}, x_t, s_t)P(x_t, g_t)$         $P(s_{t-1}, g_{t-1}|x_{t-1})$   
  $P(s_{t-1}|x_{t-2}, x_{t-1})P(g_{t-1}|x_{t-1})$         $P(s_t|x_{t-1}, x_t)P(g_t|x_t)$         $P(s_t, g_t|x_t)$   
  $P(x_{t-2}, x_{t-1}|s_{t-1})P(x_{t-1}|g_{t-1})$         $P(x_{t-1}, x_t|s_t)P(x_t|g_t)$        1  
  $P(x_{t-2}, x_{t-1}, s_{t-1}, g_{t-1})$         $P(x_{t-1}, x_t, s_t, g_t)$
- (iii)**        $P(x_{t-1}, x_t)$         $P(x_t|x_{t-1})$         $P(x_{t-2}, x_{t-1})$         $P(x_{t-1}|x_{t-2})$        1

## Q2. 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.

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

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

- (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	
$X = 14$	0.178	
$X = 9$	0.956	
$X = 6$	0.790	

- (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
$X = 6$	
$X = 1$	
$X = 7$	
$X = 8$	

- (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

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:				