


## Q1. 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:				

## 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:

$$B_t(X_t) = \underline{\hspace{2cm}} \quad \text{(i)} \quad \underline{\hspace{2cm}} \quad \text{(ii)} \quad B_{t-1}(x_{t-1}).$$

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

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

$$\text{MEU}_{1:T} = \underline{\hspace{2cm}} \quad \text{(i)} \quad \underline{\hspace{2cm}} \quad \text{(ii)} \quad \underline{\hspace{2cm}} \quad \text{(iii)} \quad \underline{\hspace{2cm}} \quad \text{(iv)} \quad \underline{\hspace{2cm}} \quad \text{(v)}$$

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

(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:

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

(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 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 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 type="radio"/> 4
Particle filtering with resampling before every other time elapse	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4