

INSTRUCTIONS


- **Due: Tuesday, April 25th, 2023 at 10:00 PM EDT.** Remember that you may use up to 1 slip day for the Written Homework making the last day to submit **Wednesday, April 26th, 2023 at 10:00 PM EDT.**
- **Format:** Write your answers in the `yoursolution.tex` file and compile a pdf (preferred) or you can type directly on the blank pdf. Make sure that your answers are within the dedicated regions for each question/part. If you do not follow this format, we may deduct points. We will NOT accept handwritten solutions of any kind.
- **Images:** To insert pictures, we recommend drawing it on PowerPoint or Google Drawings, saving it as an image and including it in your latex source.
- **How to submit:** Submit a pdf with your answers on Gradescope. Log in and click on our class 15-281 and click on the submission titled HW10 and upload your pdf containing your answers.
- **Policy:** See the course website for homework policies and Academic Integrity.

Name	
Andrew ID	
Hours to complete?	<input type="radio"/> (0, 2] hours <input type="radio"/> (2, 3] hours <input type="radio"/> (3, 4] hours <input type="radio"/> (4, 5] hours <input type="radio"/> (5, 6] hours <input type="radio"/> (6, 7] hours <input type="radio"/> (7, 8] hours <input type="radio"/> > 8 hours

Q1. [30 pts] Dynamic Bayes Net and Hidden Markov Model

A professor wants to know if students are getting enough sleep. Each day, the professor observes whether the students sleep in class, and whether they have red eyes. Let S_t be the random variable of the student having enough sleep, R_t be the random variable for the student having red eyes, and C_t be the random variable of the student sleeping in class on day t . The professor has the following theory:

- The prior probability of getting enough sleep at time t , with no observations, is 0.6
- The probability of getting enough sleep on night t is 0.9 given that the student got enough sleep the previous night, and 0.2 if not
- The probability of having red eyes is 0.1 if the student got enough sleep, and 0.7 if not
- The probability of sleeping in class is 0.2 if the student got enough sleep, and 0.4 if not



S_0	$P(S_0)$	S_{t+1}	S_t	$P(S_{t+1} S_t)$	R_t	S_t	$P(R_t S_t)$	C_t	S_t	$P(C_t S_t)$
$+s$	0.6	$+s_{t+1}$	$+s_t$	0.9	$+r$	$+s$	0.1	$+c$	$+s$	0.2
$-s$	0.4	$-s_{t+1}$	$+s_t$	0.1	$-r$	$+s$	0.9	$-c$	$+s$	0.8
		$+s_{t+1}$	$-s_t$	0.2	$+r$	$-s$	0.7	$+c$	$-s$	0.4
		$-s_{t+1}$	$-s_t$	0.8	$-r$	$-s$	0.3	$-c$	$-s$	0.6

Using the DBN above and these evidence values


- $[-r_1, -c_1]$ = not red eyes, not sleeping in class
- $[+r_2, -c_2]$ = red eyes, not sleeping in class
- $[+r_3, +c_3]$ = red eyes, sleeping in class

we want to compute $P(S_t | r_{1:t}, c_{1:t})$ for each of $t = 1, 2, 3$ as well as perform smoothing to get $P(S_2 | r_{1:3}, c_{1:3})$.

In order to do so, we will compute intermediate values which will correspond to the predict and update steps of our forward algorithm as well as finding the value of α (the normalization constant) in each case.

Note: Please round your answers to 3 decimal places at the *end* of each calculation. That is, if you need to compute multiple intermediate values to get your answer, do not round until you get your final answer. Please also note we will only be able to award partial credit if work is shown.

Figure copied from previous page for convenience:



S_0	$P(S_0)$	S_{t+1}	S_t	$P(S_{t+1} S_t)$	R_t	S_t	$P(R_t S_t)$	C_t	S_t	$P(C_t S_t)$
$+s$	0.6	$+s_{t+1}$	$+s_t$	0.9	$+r$	$+s$	0.1	$+c$	$+s$	0.2
$-s$	0.4	$-s_{t+1}$	$+s_t$	0.1	$-r$	$+s$	0.9	$-c$	$+s$	0.8
		$+s_{t+1}$	$-s_t$	0.2	$+r$	$-s$	0.7	$+c$	$-s$	0.4
		$-s_{t+1}$	$-s_t$	0.8	$-r$	$-s$	0.3	$-c$	$-s$	0.6

Round all numerical answers to 3 decimal places. Evidence values: $[-r_1, -c_1]$, $[+r_2, -c_2]$, $[+r_3, +c_3]$

(a) [7 pts] State Estimation: $t = 1$

(i) [3 pts] Predict:

Use the Bayes' Net's CPTs to write an equivalent expression for $P(+s_1)$:

$P(+s_1) =$

$P(-s_1) =$

(ii) [3 pts] Update:

Use the Bayes' Net's CPTs and α to find an equivalent expression for $P(+s_1 | -r_1, -c_1)$:

When solving for probabilities, **don't** leave α in your answer.


$P(+s_1 | -r_1, -c_1) =$

$P(-s_1 | -r_1, -c_1) =$

(iii) [1 pt] What was the value for α ?

$\alpha =$

Figure copied from previous page for convenience:



S_0	$P(S_0)$	S_{t+1}	S_t	$P(S_{t+1} S_t)$	R_t	S_t	$P(R_t S_t)$	C_t	S_t	$P(C_t S_t)$
$+s$	0.6	$+s_{t+1}$	$+s_t$	0.9	$+r$	$+s$	0.1	$+c$	$+s$	0.2
$-s$	0.4	$-s_{t+1}$	$+s_t$	0.1	$-r$	$+s$	0.9	$-c$	$+s$	0.8
		$+s_{t+1}$	$-s_t$	0.2	$+r$	$-s$	0.7	$+c$	$-s$	0.4
		$-s_{t+1}$	$-s_t$	0.8	$-r$	$-s$	0.3	$-c$	$-s$	0.6

Round all numerical answers to 3 decimal places. Evidence values: $[-r_1, -c_1]$, $[+r_2, -c_2]$, $[+r_3, +c_3]$

(b) [9 pts] State Estimation: $t = 2$

(i) [4 pts] Predict:

$$P(+s_2 | -r_1, -c_1) =$$

$$P(-s_2 | -r_1, -c_1) =$$

(ii) [4 pts] Update:

Use the Bayes' Net's CPTs, previous probabilities, and α to find an equivalent expression for $P(+s_2 | r_{1:2}, c_{1:2})$:

When solving for probabilities, **don't** leave α in your answer.


$$P(+s_2 | r_{1:2}, c_{1:2}) =$$

$$P(-s_2 | r_{1:2}, c_{1:2}) =$$

(iii) [1 pt] What was the value for α ?

$$\alpha =$$

Figure copied from previous page for convenience:



S_0	$P(S_0)$	S_{t+1}	S_t	$P(S_{t+1} S_t)$	R_t	S_t	$P(R_t S_t)$	C_t	S_t	$P(C_t S_t)$
$+s$	0.6	$+s_{t+1}$	$+s_t$	0.9	$+r$	$+s$	0.1	$+c$	$+s$	0.2
$-s$	0.4	$-s_{t+1}$	$+s_t$	0.1	$-r$	$+s$	0.9	$-c$	$+s$	0.8
		$+s_{t+1}$	$-s_t$	0.2	$+r$	$-s$	0.7	$+c$	$-s$	0.4
		$-s_{t+1}$	$-s_t$	0.8	$-r$	$-s$	0.3	$-c$	$-s$	0.6

Round all numerical answers to 3 decimal places. Evidence values: $[-r_1, -c_1]$, $[+r_2, -c_2]$, $[+r_3, +c_3]$

(c) [9 pts] State Estimation: $t = 3$

(i) [4 pts] Predict:

$$P(+s_3 | r_{1:2}, c_{1:2}) =$$

$$P(-s_3 | r_{1:2}, c_{1:2}) =$$

(ii) [4 pts] Update:

Use the Bayes' Net's CPTs, previous probabilities, and α to find an equivalent expression for $P(+s_3 | r_{1:3}, c_{1:3})$:

When solving for probabilities, **don't** leave α in your answer.


$$P(+s_3 | r_{1:3}, c_{1:3}) =$$

$$P(-s_3 | r_{1:3}, c_{1:3}) =$$

(iii) [1 pt] What was the value for α ?

$$\alpha =$$

Figure copied from previous page for convenience:



S_0	$P(S_0)$	S_{t+1}	S_t	$P(S_{t+1} S_t)$	R_t	S_t	$P(R_t S_t)$	C_t	S_t	$P(C_t S_t)$
$+s$	0.6	$+s_{t+1}$	$+s_t$	0.9	$+r$	$+s$	0.1	$+c$	$+s$	0.2
$-s$	0.4	$-s_{t+1}$	$+s_t$	0.1	$-r$	$+s$	0.9	$-c$	$+s$	0.8
		$+s_{t+1}$	$-s_t$	0.2	$+r$	$-s$	0.7	$+c$	$-s$	0.4
		$-s_{t+1}$	$-s_t$	0.8	$-r$	$-s$	0.3	$-c$	$-s$	0.6

Round all numerical answers to 3 decimal places. Evidence values: $[-r_1, -c_1]$, $[+r_2, -c_2]$, $[+r_3, +c_3]$

(d) [5 pts] We can build upon the previous three parts and use smoothing to compute $P(S_2 | r_{1:3}, c_{1:3})$.

(i) [2 pts] **Backward message:** $P(+r_3, +c_3 | S_2) = \sum_{s_3} P(+r_3, +c_3 | s_3)P(s_3 | S_2)$

$$P(+r_3, +c_3 | +s_2) =$$

$$P(+r_3, +c_3 | -s_2) =$$

(ii) [2 pts] **Smoothing:** $P(S_2 | r_{1:3}, c_{1:3}) = \alpha P(S_2 | r_{1:2}, c_{1:2})P(+r_3, +c_3 | S_2)$

When solving for probabilities, **don't** leave α in your answer.

$$P(+s_2 | r_{1:3}, c_{1:3}) =$$


$$P(-s_2 | r_{1:3}, c_{1:3}) =$$

(iii) [1 pt] What was the value for α ?

$$\alpha =$$

Q2. [12 pts] Modified HMM Updated Equations

Consider the HMM graph structure shown below.



Recall the Forward algorithm is a two step iterative algorithm used to approximate the probability distribution

$$P(X_t | e_1, \dots, e_t)$$

The two steps of the algorithm are as follows:

1. Predict:


$$P(X_t | e_{1..t-1}) = \sum_{x_{t-1}} P(X_t | x_{t-1})P(x_{t-1} | e_{1..t-1})$$

2. Observe:

$$P(X_t | e_{1..t}) = \frac{P(e_t | X_t)P(X_t | e_{1..t-1})}{\sum_{x_t} P(e_t | x_t)P(x_t | e_{1..t-1})}$$

For this problem we will consider modifying the forward algorithm as the HMM graph structure changes. Our goal will continue to be to create an iterative algorithm which is able to compute the distribution of states, X_t , given all available evidence from time 0 to time t .

- (a) [4 pts] Consider the graph below where new observed variables, Z_i , are introduced and influence the evidence.



- (i) [2 pts] What will the modified **predict** update be?

$$P(X_t | e_{1..t-1}, z_{1..t-1}) =$$


- $\sum_{x_{t-1}} P(X_t | z_{1..t-1})P(x_{t-1} | e_{1..t-1}, z_{1..t-1})$
- $\sum_{x_{t-1}} P(X_t | x_{t-1})P(x_{t-1} | e_{1..t-1}, z_{1..t-1})$
- $\sum_{x_{t-1}} P(X_t | e_{1..t-1}, z_{1..t-1})P(x_{t-1} | x_{t-1}, z_{1..t-1})$
- $\sum_{x_{t-1}} P(X_t | x_{t-1})P(x_{t-1} | e_{1..t-1})$ (no change)

(ii) [2 pts] What will the modified **observed** update be?

$$P(X_t | e_{1...t}, z_{1...t})$$

- $\frac{P(e_t | X_t, z_t) P(X_t | e_{1...t-1}, z_{1...t-1})}{\sum_{x_t, z_t} P(e_t | x_t, z_t) P(x_t | e_{1...t-1}, z_{1...t-1})}$
 $\frac{P(e_t | X_t, z_t) P(X_t | e_{1...t-1}, z_{1...t-1})}{\sum_{x_t} P(e_t | x_t, z_t) P(x_t | e_{1...t-1}, z_{1...t-1})}$
 $\frac{P(e_t | X_t) P(X_t | e_{1...t-1})}{\sum_{z_t} P(e_t | x_t, z_t) P(x_t | e_{1...t-1}, z_{1...t-1})}$
 $\frac{P(e_t | X_t, z_t) P(X_t | e_{1...t-1}, z_{1...t-1})}{\sum_{z_t} P(e_t | x_t, z_t) P(x_t | e_{1...t-1}, z_{1...t-1})}$
 $\frac{P(e_t | X_t) P(X_t | e_{1...t-1})}{\sum_{x_t} P(e_t | x_t) P(x_t | e_{1...t-1})}$ (no change)

(b) [4 pts] Next, consider the graph below where the Z_i variables are unobserved.



(i) [2 pts] What will the modified **predict** update be?

$$P(X_t | e_{1...t-1}) =$$


- $\sum_{x_{t-1}} P(X_t | z_{1...t-1}) P(x_{t-1} | e_{1...t-1}, z_{1...t-1})$
 $\sum_{x_{t-1}} P(X_t | x_{t-1}) P(x_{t-1} | e_{1...t-1}, z_{1...t-1})$
 $\sum_{x_{t-1}} P(X_t | e_{1...t-1}, z_{1...t-1}) P(x_{t-1} | x_{t-1}, z_{1...t-1})$
 $\sum_{x_{t-1}} P(X_t | x_{t-1}) P(x_{t-1} | e_{1...t-1})$ (no change)

(ii) [2 pts] What will the modified **observed** update be?

$$P(X_t | e_{1...t}) =$$

- $\frac{P(X_t | e_{1...t-1}) P(z_t) P(e_t | X_t, z_t)}{\sum_{z_t} P(x_t | e_{1...t-1}) P(e_t | x_t, z_t) P(z_t)}$
 $\frac{P(X_t | e_{1...t-1}) \sum_{z_t} P(z_t) P(e_t | X_t, z_t)}{P(x_t | e_{1...t-1}) \sum_{z_t} P(e_t | x_t, z_t) P(z_t)}$
 $\frac{P(X_t | e_{1...t-1}) P(z_t) P(e_t | X_t, z_t)}{\sum_{x_t} P(x_t | e_{1...t-1}) P(e_t | x_t, z_t) P(z_t)}$
 $\frac{P(X_t | e_{1...t-1}) \sum_{z_t} P(z_t) P(e_t | X_t, z_t)}{\sum_{x_t} P(x_t | e_{1...t-1}) \sum_{z_t} P(e_t | x_t, z_t) P(z_t)}$
 $\frac{P(e_t | X_t) P(X_t | e_{1...t-1})}{\sum_{x_t} P(e_t | x_t) P(x_t | e_{1...t-1})}$ (no change)

- (c) [4 pts] Finally, consider a graph where the newly introduced variables are unobserved and influenced by the evidence nodes.



- (i) [2 pts] What will the modified **predict** update be?

$$P(X_t | e_{1...t-1}) =$$

- $\sum_{x_{t-1}} P(X_t | z_{1...t-1})P(x_{t-1} | e_{1...t-1}, z_{1...t-1})$
- $\sum_{x_{t-1}} P(X_t | x_{t-1})P(x_{t-1} | e_{1...t-1}, z_{1...t-1})$
- $\sum_{x_{t-1}} P(X_t | e_{1...t-1}, z_{1...t-1})P(x_{t-1} | x_{t-1}, z_{1...t-1})$
- $\sum_{x_{t-1}} P(X_t | x_{t-1})P(x_{t-1} | e_{1...t-1})$ (no change)


- (ii) [2 pts] What will the modified **observed** update be?

$$P(X_t | e_{1...t}) =$$

- $\frac{P(e_t | X_t, z_t)P(X_t | e_{1...t-1}, z_{1...t-1})}{\sum_{z_t} P(e_t | x_t, z_t)P(x_t | e_{1...t-1}, z_{1...t-1})}$
- $\frac{P(e_t | X_t, z_t)P(X_t | e_{1...t-1}, z_{1...t-1})}{\sum_{x_t} P(e_t | x_t, z_t)P(x_t | e_{1...t-1}, z_{1...t-1})}$
- $\frac{P(e_t | X_t)P(X_t | e_{1...t-1})}{\sum_{x_t, z_t} P(e_t | x_t, z_t)P(x_t | e_{1...t-1}, z_{1...t-1})}$
- $\frac{P(e_t | X_t)P(X_t | e_{1...t-1})}{\sum_{x_t} P(e_t | x_t)P(x_t | e_{1...t-1})}$ (no change)

Q3. [36 pts] Particle Filtering

In this question, we will use a particle filter to track the state of a robot that is lost in the small map below:



As we walk through this problem, there are many values to compute, so you may want to download and print the following worksheet and fill it in as you go.

NOTE: For your convenience, we've created a supplemental gradescope assignment to help you check your work for the first few parts. Like other online homeworks, entering in the correct answers will prompt a solution explanation. That being said, no credit is given for completing this gradescope assignment (ignore the ϵ that's required to appease gradescope), so please still put your answers in the writeup.


The robot's state is represented by an integer $1 \leq X_t \leq 10$ corresponding to its location in the map at time t . We will approximate our belief over this state with $N = 8$ particles.

You have no control over the robot's actions. At each timestep, the robot either stays in place, or moves to any one of its neighboring locations, all with equal probability. For example, if the robot starts in state $X_t = 7$, it will move to state $X_{t+1} = 6$ with probability $\frac{1}{2}$ or $X_{t+1} = 7$ with probability $\frac{1}{2}$. Similarly, if the robot starts in state $X_t = 2$, the next state X_{t+1} can be any element of $\{1, 2, 3, 10\}$, and each occurs with probability $\frac{1}{4}$.

At each time step, a sensor on the robot gives a reading $E_t \in \{H, C, T, D\}$ corresponding to the *type* of state the robot is in. The possible types are:


- **Hallway (H)** for states bordered by two parallel walls (4,9).
- **Corner (C)** for states bordered by two orthogonal walls (3,5,8,10).
- **Tee (T)** for states bordered by one wall (2,6).
- **Dead End (D)** for states bordered by three walls (1,7)

The sensor is not very reliable: it reports the correct type with probability $\frac{1}{2}$, but gives erroneous readings the rest of the time, with probability $\frac{1}{6}$ for each of the three other possible readings.



(a) [8 pts] Fill in the sensor model below:


P(Sensor Reading	State Type)	Answer
P(Sensor Reading = H	State Type = H)	
P(Sensor Reading = C	State Type = H)	
P(Sensor Reading = T	State Type = H)	
P(Sensor Reading = D	State Type = H)	
P(Sensor Reading = H	State Type = C)	
P(Sensor Reading = C	State Type = C)	
P(Sensor Reading = T	State Type = C)	
P(Sensor Reading = D	State Type = C)	
P(Sensor Reading = H	State Type = T)	
P(Sensor Reading = C	State Type = T)	
P(Sensor Reading = T	State Type = T)	
P(Sensor Reading = D	State Type = T)	
P(Sensor Reading = H	State Type = D)	
P(Sensor Reading = C	State Type = D)	
P(Sensor Reading = T	State Type = D)	
P(Sensor Reading = D	State Type = D)	



- (b) [4 pts] Suppose that we want to sample from a set of 4 events, $\{A, B, C, D\}$, which occur with corresponding probabilities P_A, P_B, P_C, P_D .

First, we form the set of cumulative weights, given by $\{0, P_A, P_A + P_B, P_A + P_B + P_C, 1\}$. (Note: $P_A + P_B + P_C + P_D = 1$). These weights partition the $[0, 1)$ interval into bins, as shown below. We then draw a number r uniformly at random from $[0, 1)$ and pick A, B, C , or D based on which bin r lands in.


The process is illustrated in the diagram below. If r_1 uniformly chosen from $[0, 1)$ lands in the interval $[P_A, P_A + P_B]$, then the resulting sample would be B . Similarly, if r_2 lands in $[P_A + P_B, P_A + P_B + P_C]$, the sample would be C , and r_3 landing in $[P_A + P_B, P_A + P_B + P_C]$ would also be C .



Now we will sample the starting positions for our particles at time $t = 0$. For each particle p_i , we have generated a random number r_i sampled uniformly from $[0, 1)$.

Your job is to use these numbers to sample a starting location for each particle. As a reminder, locations are integers from the range $[1, 10]$, as shown in the map. You should assume that the locations go in **ascending order and that each location has equal probability**. The random number generated for particle i , denoted by r_i is provided. Please fill in the locations of the eight particles.

r_i	p_i
$r_1 = 0.139$	
$r_2 = 0.416$	
$r_3 = 0.683$	
$r_4 = 0.825$	
$r_5 = 0.396$	
$r_6 = 0.161$	
$r_7 = 0.554$	
$r_8 = 0.013$	




- (c) [4 pts] Now we'll perform a time update from $t = 0$ to $t = 1$ using the transition model. Stated again, the transition model is as follows: At each timestep, the robot either stays in place, or moves to any one of its neighboring locations, **all with equal probability**.

For each particle, take the starting position you found in part b, and perform the time update for that particle. You should again sample from the range $[0, 1)$, where the bins are the possible locations **sorted in ascending numerical order**. As an example, if $X_t = 3$, the next state can be one of $\{2, 3, 4\}$, each with equal probability, so the $[0, 0.333)$ bin would be for $X_{t+1} = 2$, the $[0.333, 0.667)$ bin would be for $X_{t+1} = 3$, and the $[0.667, 1)$ bin would be for $X_{t+1} = 4$. Fill in the bin partitioning you are using for the p_1 update in the first table below. **Note the first row is given an example (do not use in your answer) and you may not need to use all of the rows provided.**

bin	X_{t+1}
$[0, 0.333)$	2

r_i	p_i
$r_1 = 0.822$	
$r_2 = 0.244$	
$r_3 = 0.686$	
$r_4 = 0.842$	
$r_5 = 0.113$	
$r_6 = 0.314$	
$r_7 = 0.177$	
$r_8 = 0.406$	




- (d) [5 pts] Recall that a particle filter just keeps track of a list of particles, but at any given time, we can compute a probability distribution from these particles.

Using the current newly updated set of particles (that you found in Part c), give the estimated probability that the robot is in each location.

$\hat{P}(X_i)$	Answer
$\hat{P}(X_1 = 1)$	
$\hat{P}(X_1 = 2)$	
$\hat{P}(X_1 = 3)$	
$\hat{P}(X_1 = 4)$	
$\hat{P}(X_1 = 5)$	
$\hat{P}(X_1 = 6)$	
$\hat{P}(X_1 = 7)$	
$\hat{P}(X_1 = 8)$	
$\hat{P}(X_1 = 9)$	
$\hat{P}(X_1 = 10)$	

- (e) [4 pts] The sensor reading at $t = 1$ is: $E_1 = D$. Using the sensor model you specified in Part a, incorporate the evidence by weighting the particles. Refer back to Part c to get the positions of your particles.

Particle p_i weight	Answer
p_1	
p_2	
p_3	
p_4	
p_5	
p_6	
p_7	
p_8	



- (f) [5 pts] After incorporating the evidence by weighting the particles, we can compute an updated probability distribution from these particles.

Using the set of particles (Part c, d), and the weights (Part e), give the estimated probability that the robot is in each location given the evidence, $\hat{P}(X_1 | E_1 = D)$. *Hint:* You will first want to calculate $\hat{P}(X_1, E_1 = D)$ and then normalize.

X_1	$\hat{P}(X_1, E_1 = D)$	Normalized Particle p_i weight
1		
2		
3		
4		
5		
6		
7		
8		
9		
10		

Use this space to show any other work if you need to

- (g) [4 pts] Finally, we'll resample the particles. This reallocates resources to the most relevant parts of the state space in the next time update step.

Use your $\hat{P}(X_1 | E_1)$ from the previous part to sample eight new particles.

First fill in the bins to partition the $[0, 1)$ interval to be sampled from in the first table. You may not need to use all of the rows provided:

bin	X_{t+1}

Using the bins above, now fill in the location of the new particles according to the given samples:

r_i	p_i	Location of new particle p_i
$r_1 = 0.803$	p_1	
$r_2 = 0.712$	p_2	
$r_3 = 0.626$	p_3	
$r_4 = 0.140$	p_4	
$r_5 = 0.559$	p_5	
$r_6 = 0.979$	p_6	
$r_7 = 0.231$	p_7	
$r_8 = 0.847$	p_8	

- (h) [2 pts] We said that the sensor provided a reading $E_1 = D$. What fraction of the particles ended up at a dead end in $t = 1$?

Answer:

This completes everything for the first time step, $t = 0 \rightarrow t = 1$. Of course, we would now continue by repeating the time update, evidence incorporation by weighting, and resampling. We'll leave that to the computers, though.

Q4. [13 pts] Game Theory Conceptual Questions and Dominant Strategy

(a) [5 pts] Which of the following statements are true? Select all that apply.

- A) It is possible to represent the rock-paper-scissors game both as a normal-form game and as an extensive form game.
- B) If a strategy's support has size 1, then it is a pure strategy.
- C) A game in which each participant's gain or loss of utility is exactly balanced by the losses or gains of the utility of the other participants is called a zero-sum game.
- D) A strategy is strictly dominant if it is never worse than any other strategy for that player.
- E) A Nash equilibrium is when every player's strategy is a best response to others players' strategies.

R2-D2 and BB-8 are each deciding between going to the park or to the museum. The following table shows the utilities for R2-D2 and BB-8 in each of the possible scenarios. The (x, y) pairs in the table denote (R2-D2's utility, BB-8's Utility).

		BB-8	
		Park	Museum
R2-D2	Park	(1, 6)	(1, 8)
	Museum	(3, 2)	(2, 5)

(b) [4 pts] Fill in the following table for R2-D2.

BB-8's Decision	R2-D2's Utility Choosing Park	R2-D2's Utility Choosing Museum	Action R2-D2 should pick
Park			
Museum			

What is R2-D2's dominant strategy?

- A) Park
- B) Museum
- C) Doesn't have one

(c) [4 pts] Fill in the following table for BB-8.

R2-D2's Decision	BB-8's Utility Choosing Park	BB-8's Utility Choosing Museum	Action BB-8 should pick
Park			
Museum			

What is BB-8's dominant strategy?

- A) Park
- B) Museum
- C) Doesn't have one

Q5. [9 pts] Equilibrium

Now, R2-D2 and BB-8 are deciding whether to attend a football game or go to a concert. The following table shows the utilities for R2-D2 and BB-8 in each of the possible scenarios. The (x, y) pairs in the table denote (R2-D2's utility, BB-8's utility).

		BB-8	
		Football	Concert
R2-D2	Football	(8, 8)	(1, 7)
	Concert	(3, 4)	(2, 9)

(a) [5 pts] Which of the following are pure strategy Nash equilibria? The pairs denote (R2-D2's choice, BB-8's choice).

- A) (football, football)
- B) (football, concert)
- C) (concert, football)
- D) (concert, concert)
- E) There is no pure strategy Nash equilibrium in this game.

(b) [4 pts] Find the mixed strategy Nash equilibrium in this game.

(i) R2-D2 will choose football with probability:

Answer:

(ii) R2-D2's expected utility is:

Answer:

(iii) BB-8 will choose football with probability:

Answer:

(iv) BB-8's expected utility is:

Answer: