15-281	Artificial Intelligence	HW12

INSTRUCTIONS

- **Due: Tuesday, April 28, 2020 at 10:00 PM EDT.** Remember that you may use up to 2 slip days for the Written Homework making the last day to submit **Thursday, April 30, 2020 at 10:00 PM EDT**.
- **Format:** Submit the answer sheet pdf containing your answers. You should solve the questions on this handout (either through a pdf annotator, or by printing, then scanning). Make sure that your answers (typed or handwritten) are within the dedicated regions for each question/part. If you do not follow this format, we may deduct points.
- **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 HW12 and upload your pdf containing your answers.
- **Policy:** See the course website for homework policies and Academic Integrity.

Name	
Andrew ID	
Hours to complete?	

For staff use only

Q1	Q2	Q3	Q4	Q5	Total
/17	/12	/21	/14	/36	/100

Q1. [17 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 players strategy is a best response to others players strategies.

(b) [6 pts] 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)

Note: Filling in this table is optional. However, it will help us understand where any misconceptions may have occurred so we can give partial credit.

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) [6 pts] *Note*: Filling in this table is optional. However, it will help us understand where any misconceptions may have occurred so we can give partial credit.

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

Q2. [12 pts] Iteratively Finding a Pure Strategy Nash Equilibrium

There are so many things that R2-D2 and BB-8 can do in Pittsburgh! Now they are each deciding between going to the movies, to the mall, or to the zoo. The following table (referred to as Table A) 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		
		Zoo	Movies	Mall
R2-D2	Zoo	(3, -1)	(2, 2)	(5, 0)
	Movies	(4, 5)	(6, 4)	(4, 3)
	Mall	(1, 2)	(2, 4)	(6, 1)

(a) [2 pts] Which, if any, of R2-D2's strategies are strictly dominated by another?

- ☐ A) Zoo
- ☐ B) Movies
- ☐ C) Mall
- ☐ D) None of them

(b) [2 pts] Which, if any, of BB-8's strategies are strictly dominated by another?

- ☐ A) Zoo
- ☐ B) Movies
- ☐ C) Mall
- ☐ D) None of them

(c) [2 pts] Now consider the table after eliminating the strategies that are strictly dominated by another, referred to as Table B.

Which, if any, of R2-D2's strategies are strictly dominated by another now? Only include strategies of R2-D2 in Table B.

- ☐ A) Zoo
- ☐ B) Movies
- ☐ C) Mall
- ☐ D) None of them

(d) [2 pts] Which, if any, of BB-8's strategies are strictly dominated by another now? Only include strategies of BB-8 in Table B.

- ☐ A) Zoo
- ☐ B) Movies
- ☐ C) Mall
- ☐ D) None of them

(e) [2 pts] Now after eliminating these strategies as well, you should be able to find the Nash equilibrium easily.

Which will be R2-D2's strategy in the Nash equilibrium?

- ☐ A) Zoo
- ☐ B) Movies
- ☐ C) Mall

(f) [2 pts] Which will be BB-8's strategy in the Nash equilibrium?

- ☐ A) Zoo
- ☐ B) Movies
- ☐ C) Mall

Q3. [21 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) [16 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:

Q4. [14 pts] Stackelberg Equilibrium

A version of the payoff matrix for the classic **Chicken** game is presented here.

		Player 2	
		Dare	Chicken
Player 1	Dare	(0,0)	(7,2)
	Chicken	(2,7)	(6,6)

(a) [5 pts] Which of the following are pure strategy Nash equilibria? The pairs denote (player 1, player 2).

☐ A) (Dare, Dare)
☐ B) (Dare, Chicken)
☐ C) (Chicken, Dare)
☐ D) (Chicken, Chicken)
☐ E) There is no pure strategy Nash equilibrium in this game.

(b) [9 pts] Please fill in the table below as it will assist in determining the Stackelberg Equilibrium. *Note:* This isn't optional.

Player 1 Action	**Player 2 Payoff Choosing Dare**	**Player 2 Payoff Choosing Chicken**	**Player 2 Action**	**Player 1 Payoff**
Dare				
Chicken				

If it is a Stackelberg game and Player 1 (row player) is the leader. You are told that the Stackelberg Equilibrium in this game is a pure strategy equilibrium. Can you tell what is the Stackelberg Equilibrium? *Note:* The choices refer to (player 1, player 2).

☐ A) (Dare, Dare)
☐ B) (Dare, Chicken)
☐ C) (Chicken, Dare)
☐ D) (Chicken, Chicken)

Q5. [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:

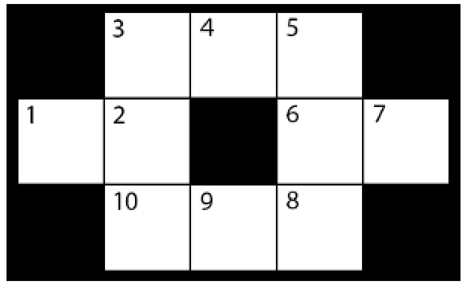

As we walk through this problem, there are many values to compute, so you may want to download and print the following worksheet and fill it in as you go.

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 more to state $X_{t+1} = 6$ with probability $\frac{1}{2}$ or $X_t = 7$ 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 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.

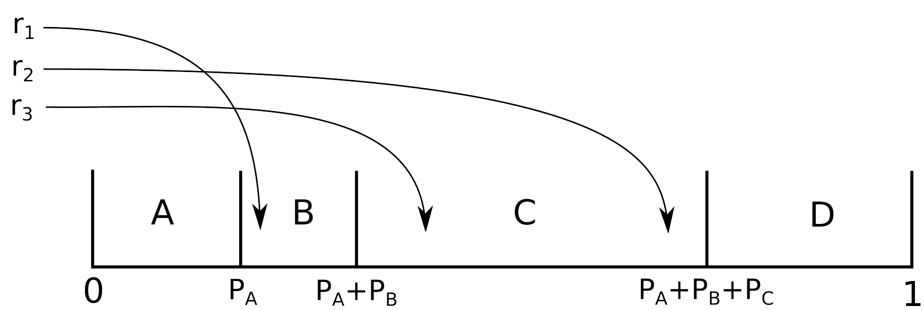

Note: P_A+P_B+P_C+P_D = 1

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 r_i.

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
$r_1 = 0.914$	
$r_2 = 0.473$	
$r_3 = 0.679$	
$r_4 = 0.879$	
$r_5 = 0.212$	
$r_6 = 0.024$	
$r_7 = 0.458$	
$r_8 = 0.154$	

(c) [4 pts] At this point, it is highly recommended that you write down the starting locations for each particle as you will need them to answer Part c.

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 $X_t = 2$, the next state can be one of $\{1, 2, 3, 10\}$, each with equal probability, so the [0,0.25) bin would be for $X_{t+1} = 1$, the $[0.25, 0.5)$ bin would be for $X_{t+1} = 2$, the $[0.5, 0.75)$ bin would be for $X_{t+1} = 3$, and the $[0.75, 1)$ bin would be for $X_{t+1} = 10$.

r_i	p_i
$r_1 = 0.647$	
$r_2 = 0.119$	
$r_3 = 0.748$	
$r_4 = 0.802$	
$r_5 = 0.357$	
$r_6 = 0.736$	
$r_7 = 0.425$	
$r_8 = 0.058$	

At this point, it is highly recommended that you copy down the new locations for each particle as you will need them to answer Part d, Part e, and Part f.

(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.

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

(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, $\hat{P}(X_1 \mid E_1 = D)$. *Hint*: You will first want to calculate $\hat{P}(X_1, E_1 = D)$ and then normalize.

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

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

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, $t = 0 \to t = 1$. 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.