

0/24 Questions Answered

HW 7 (Electronic Component)


STUDENT NAME

Q1 Decisions

21 Points

You've been job hunting, and you've narrowed your options to two companies:

[Acme](#) and Google. You already have an offer from Acme, but it expires today, and you are still waiting for a response from Google. You are faced with the dilemma of whether or not to accept the offer from Acme, which is modeled by the following decision network:



The prior probability distribution for whether Google will hire you and the utilities over possible outcomes are as follows:

Google outcome	P(Google outcome)
hired	0.25
not hired	0.75

Action	Google outcome	U
accept Acme offer	hired	2000
accept Acme offer	not hired	8000
reject Acme offer	hired	10000
reject Acme offer	not hired	0

Q1.1

5 Points

What is the expected utility of each action? (Note: throughout this problem answers will be evaluated to whole-number precision, so your answer should differ by no more than 1 from the exact answer.)

Action: accept Acme offer

Action: reject Acme offer


Which action should you take?

- reject
- accept

Q1.2

16 Points

Suddenly, the phone rings. It's your uncle, who works at Google. Your uncle tells you he has some inside information about the status of your application. Your uncle won't tell you what the information is yet, but he might be willing to divulge it for the right price. You model the new situation by adding a new node to your decision network:



You create a CPT to model the relationship between the inside information and Google's future hiring decision:

info	Google outcome	P(info Google outcome)
good news	hired	0.7
bad news	hired	0.3
good news	not hired	0.1
bad news	not hired	0.9

We'll help grind through the probabilistic inference. The resulting distributions are:

info	P(info)
good news	0.25
bad news	0.75

Google outcome	info	P(Google outcome info)
hired	good news	0.7
not hired	good news	0.3
hired	bad news	0.1
not hired	bad news	0.9

(That these are identical to $P(\text{Google outcome})$ and $P(\text{Info} | \text{Google outcome})$ is just a numerical coincidence.) Fill in the expected utilities for each action, for each possible type of information we could be given:

EU(accept Acme offer | good news)

Enter your answer here

EU(reject Acme offer | good news)

Enter your answer here

EU(accept Acme offer | bad news)

Enter your answer here

EU(reject Acme offer | bad news)

Enter your answer here

What is the maximum expected utility for each type of information we could be given?

MEU(good news)

Enter your answer here

MEU(bad news)

Enter your answer here

If we are given the inside information, what is the expected value of MEU?

Enter your answer here

What is the value of perfect information of the random variable Inside Info?

Enter your answer here

Save Answer

Q2 Value of Perfect Information

10 Points

Consider the expected value of perfect information (VPI) of observing some node in an arbitrary decision network. Which of the following are true statements?


- VPI is guaranteed to be positive (> 0).
- VPI is guaranteed to be nonnegative (≥ 0).
- VPI is guaranteed to be nonzero.
- The MEU after observing a node could potentially be less than the MEU before observing that node.
- For any two nodes X and Y , $VPI(X) + VPI(Y) \geq VPI(X, Y)$. That is, the sum of individual VPI's for two nodes is always greater than or equal to the VPI of observing both nodes.
- VPI is guaranteed to be exactly zero for any node that is conditionally independent (given the evidence so far) of all parents of the utility node.

Save Answer

Q3 Particle Filtering

60 Points

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



The robot's state is represented by an integer $X_t \in \{1 \dots 10\}$ corresponding to its location in the map at time t . We will approximate the probability distribution of 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.

Q3.1 Sensor Model

10 Points

Fill in the first two parts of the sensor model below, with probabilities rounded to three decimal places:

$$P(E_t = H \mid X_t = 1)$$

$$P(E_t = C \mid X_t = 1)$$

$$P(E_t = T \mid X_t = 1)$$

$$P(E_t = D \mid X_t = 1)$$

$$P(E_t = H \mid X_t = 2)$$

$$P(E_t = C \mid X_t = 2)$$

$$P(E_t = T \mid X_t = 2)$$


$$P(E_t = D \mid X_t = 2)$$

Q3.2 Sampling Review

8 Points

Suppose that we want to sample from a set of 4 mutually exclusive and exhaustive 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 $\{0, P_A, P_A + P_B, P_A + P_B + P_C, 1\}$. 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 $\{1 \dots 10\}$, as shown in the map. You should assume that the locations go in ascending order and that each location has equal probability initially (before any evidence is observed). The random number generated for particle i , denoted by r_i , is provided. Please fill in the locations $x_0^{(i)}$ of the eight particles.

$$r_1 = 0.914$$

$$x_0^{(1)} =$$

Enter your answer here

$$r_2 = 0.473$$

$$x_0^{(2)} =$$

$$r_3 = 0.679$$

$$x_0^{(3)} =$$

$$r_4 = 0.879$$

$$x_0^{(4)} =$$

$$r_5 = 0.212$$

$$x_0^{(5)} =$$

$$r_6 = 0.024$$

$$x_0^{(6)} =$$

$$r_7 = 0.458$$

$$x_0^{(7)} =$$

$$r_8 = 0.154$$

$$x_0^{(8)} =$$

At this point, it is *highly recommended* that you copy down the starting locations for each particle as you will need them to answer Part 3.


Q3.3 Time Update

8 Points

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 2, and perform the time update for that particle. We again sample uniformly from the range $[0, 1)$, and the bins are the possible locations **sorted in ascending numerical order** with widths proportional to their probabilities. 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 $X_{t+1} = 10$.

The map is shown again below:



$$r_1 = 0.674$$

$$x_1^{(1)} =$$

$$r_2 = 0.119$$

$$x_1^{(2)} =$$

$$r_3 = 0.748$$

$$x_1^{(3)} =$$

$$r_4 = 0.802$$

$$x_1^{(4)} =$$

$$r_5 = 0.357$$

$$x_1^{(5)} =$$

$$r_6 = 0.736$$

$$x_1^{(6)} =$$

$$r_7 = 0.425$$

$$x_1^{(7)} =$$

$$r_8 = 0.058$$

$$x_1^{(8)} =$$

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 4, Part 5, and Part 6.

Q3.4 Probability Distribution Induced by the Particles

8 Points

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 3), give the estimated probability $\hat{P}(X_1)$ that the robot is in each location.

$$\hat{P}(X_1 = 1)$$

$$\hat{P}(X_1 = 2)$$

$$\hat{P}(X_1 = 3)$$

$$\hat{P}(X_1 = 4)$$

$$\hat{P}(X_1 = 5)$$

$$\hat{P}(X_1 = 6)$$

$$\hat{P}(X_1 = 7)$$

$$\hat{P}(X_1 = 8)$$

$$\hat{P}(X_1 = 9)$$

$$\hat{P}(X_1 = 10)$$

Q3.5 Incorporating Evidence

12 Points

The sensor reading at $t = 1$ is: $E_1 = D$

Using the sensor model you specified in Part 1, incorporate the evidence by reweighting the particles. Also enter the normalized and cumulative weights for each particle. The normalized weight for a

specific particle can be calculated by taking that particle's weight and dividing by the sum of all the particle weights. The cumulative weight keeps track of a running sum of all the weights of the particles seen so far (meaning, particle i will have a cumulative weight equal to the sum of the weights of all particles j such that $j \leq i$).

Refer back to Part 3 to get the positions of your particles.

The map is shown again below:

	3	4	5	
1	2		6	7
	10	9	8	

Particle p_1 weight:

p_1 normalized weight:

p_1 cumulative normalized weight:

Particle p_2 weight:

p_2 normalized weight:

p_2 cumulative normalized weight:

Enter your answer here

Particle p_3 weight:

Enter your answer here

p_3 normalized weight:

Enter your answer here

p_3 cumulative normalized weight:

Enter your answer here

Particle p_4 weight:

Enter your answer here

p_4 normalized weight:

Enter your answer here

p_4 cumulative normalized weight:

Enter your answer here

Particle p_5 weight:

Enter your answer here

p_5 normalized weight:

Enter your answer here

p_5 cumulative normalized weight:

Particle p_6 weight: p_6 normalized weight: p_6 cumulative normalized weight:Particle p_7 weight: p_7 normalized weight: p_7 cumulative normalized weight:Particle p_8 weight: p_8 normalized weight: p_8 cumulative normalized weight:

Q3.6 Resampling

10 Points

Finally, we'll resample the particles. This reallocates resources to the most relevant parts of the state space in the next time update step.

Notice that your cumulative weights effectively tell you where the bins used in resampling the particles lie. For example, for particle 1, you calculated the cumulative weight to be some value w . Then, on a random value draw, if a value between 0 and w was chosen, you would generate a new particle where particle 1 is. Use these bounds to resample the eight particles. In the "New Particle" row, enter the id of the particle corresponding to the bin that the random value chose. In the "New Location" row, enter the location corresponding to the chosen particle. You may need to look back at Part 3 to get the locations of the particles.

$$r_1 = 0.403$$

New particle for p_1 :

New location for p_1 :

$$r_2 = 0.218$$

New particle for p_2 :

New location for p_2 :

$$r_3 = 0.217$$

New particle for p_3 :

New location for p_3 :

$$r_4 = 0.826$$

New particle for p_4 :

New location for p_4 :

$$r_5 = 0.717$$

New particle for p_5 :

New location for p_5 :

$$r_6 = 0.460$$

New particle for p_6 :

New location for p_6 :

$$r_7 = 0.794$$

New particle for p_7 :

New location for p_7 :

$$r_8 = 0.016$$

New particle for p_8 :


New location for p_8 :

Q3.7 Analysis

4 Points

The sensor provided a reading $E_1 = D$. What fraction of the particles are now on a dead end?

The map is shown again below:




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 reweighting, and resampling. We'll leave that to the computers, though.

Q4 Modified HMM Update Equations

24 Points

Consider the HMM graph structure shown below.



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

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

The two steps of the algorithm are as follows:

Predict:

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

Update:


$$P(X_t | e_{1:t}) = \alpha P(e_t | X_t)P(X_t | e_{1:t-1})$$

For this problem we will consider modifying the forward algorithm for slightly different HMM graph structures. 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 .

Q4.1

8 Points

Consider the graph below where new observed variables, Z_i , are introduced and influence the evidence.



What will the modified prediction step 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)

What will the modified update step be?

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


- $\alpha P(e_t, z_t | X_t)P(X_t | e_{1:t-1}, z_{1:t-1})$
- $\alpha P(e_t | X_t, z_t)P(X_t | e_{1:t-1}, z_{1:t-1})$
- $\alpha P(e_t | X_t, z_t)P(X_t | e_{1:t-1})$
- $\alpha P(e_t | X_t)P(X_t | e_{1:t-1})$ (no change)

Save Answer

Q4.2

8 Points

Next, consider the graph below where the Z_i variables are unobserved.



What will the modified prediction step 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)

What will the modified update step be?

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


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

Save Answer

Q4.3

8 Points

Finally, consider a graph where the newly introduced variables are unobserved and influenced by the evidence nodes.



What will the modified prediction step 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)

What will the modified update step be?

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

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

Save Answer

Q5 Rationality of Utilities

9 Points

Q5.1

3 Points

Consider a lottery $L = [0.2, A; 0.3, B; 0.4, C; 0.1, D]$,

where the utility values of each of the outcomes are $U(A) = 1$,

$U(B) = 3$,

$U(C) = 5, U(D) = 2$. What is the utility of this lottery, $U(L)$?

Q5.2

3 Points

Consider a lottery $L1 = [0.5, A; 0.5, L2]$, where $U(A) = 4$, and $L2 = [0.5, X; 0.5, Y]$ is a lottery, and $U(X) = 4, U(Y) = 8$. What is the utility of

the the first lottery, $U(L1)$?

Q5.3

3 Points

Assume $A \succ B, B \succ L$, where $L = [0.5, C; 0.5, D]$, and $D \succ A$.

Assuming

rational preferences, which of the following statements are guaranteed to be


true?

$A \succ L$
 $A \succ C$
 $A \succ D$
 $B \succ C$
 $B \succ D$

Q6 Certainty Equivalent Values

6 Points

Consider the utility function shown below.



Under the above utility function, what is the certainty equivalent monetary value in dollars (\$) of the lottery $[0.6, \$0; 0.4, \$100]$?

I.e., what is X such that $U(\$X) = U([0.6, \$0; 0.4, \$100])$?

Hint:

Keep in mind that $U([p, A; 1 - p, B])$ is not equal to $U(pA + (1 - p)B)$.

Enter your answer here

Save Answer

Q7 Preferences and Utilities

14 Points

Our Pacman board now has food pellets of 3 different sizes - pellet P_1 of radius 1, P_2 of radius 2 and P_3 of radius 3. In different moods, Pacman has different preferences among these pellets. In each of the following questions, you are given Pacman's preference for the different pellets. From among the options pick the utility functions that are consistent with Pacman's preferences, where each utility function $U(r)$ is given as a function of the pellet radius r , and is defined over non-negative values of r .

Q7.1

2 Points

$$P_1 \sim P_2 \sim P_3$$

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

 Irrational preferences!

Save Answer

Q7.2

2 Points

$$P_1 \succ P_2 \succ P_3$$

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

 Irrational preferences!

Save Answer

Q7.3

2 Points

$$P_1 \succ P_2 \succ P_3$$

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

 Irrational preferences!**Q7.4**

2 Points

 $(P_1 \succ P_2 \succ P_3)$ and $(P_2 \succ (50\text{-}50 \text{ lottery among } P_1 \text{ and } P_3))$

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

 Irrational preferences!**Q7.5**

2 Points

 $(P_1 \succ P_2 \succ P_3)$ and $(P_2 \succ (50\text{-}50 \text{ lottery among } P_1 \text{ and } P_3))$

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

 Irrational preferences!**Q7.6**

2 Points

$(P_1 \succ P_2)$ and $(P_2 \succ P_3)$ and ((50-50 lottery among P_2 and P_3) \succ (5

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

 Irrational preferences!**Q7.7**

2 Points

Which of the following would be a utility function for a risk-seeking preference? That is, for which utility(s) would Pacman prefer entering a lottery for a random food pellet, with expected size s , over receiving a pellet of size s ?

$U(r) = 0$

$U(r) = 3$

$U(r) = r$

$U(r) = 2r + 4$

$U(r) = -r$

$U(r) = r^2$

$U(r) = -r^2$

$U(r) = \sqrt{r}$

$U(r) = -\sqrt{r}$

Save Answer

Save All Answers

Submit & View Submission >