CS 188 Summer 2024 | Introduction to Artificial Intelligence

Exam Prep 8 Solutions

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

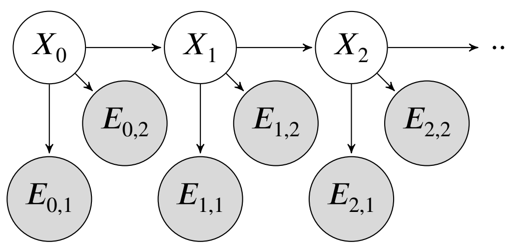

In each time step, Leland will either stay in the same region or move to an adjacent region. For example, the available actions from region 4 are (WEST, EAST, SOUTH, STAY). He chooses between all available actions with equal probability, regardless of where your robots are. Note: moving off the grid 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.

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

E	$P(E_{t,1} \mid X_t = 1)$	$P(E_{t,2} \mid X_t = 1)$
$NEAR$	0.8	0.3
FAR	0.2	0.7

(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	11
$X = 14$	0.178	9
$X = 9$	0.956	14
$X = 6$	0.790	11

(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
$X = 6$	$0.8 * 0.3$
$X = 1$	$0.8 * 0.3$
$X = 7$	$0.3 * 0.3$
$X = 8$	$0.3 * 0.8$

(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

Normalizing gives us the distribution

$$X = 1 : 0.4/0.8 = 0.5$$
$$X = 6 : 0.1/0.8 = 0.125$$
$$X = 7 : 0.1/0.8 = 0.125$$
$$X = 8 : 0.2/0.8 = 0.25$$

Use the following random numbers to resample you 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:	6	1	1	8

Q2. Hearthstone Decisions

You are playing the game Hearthstone. You are up against the famous player Alice.

On your turn, you can choose between playing 0, 1, or 2 minions. You realize Alice might be holding up an Area of Effect (AoE) card, which is more devastating the more minions you play.

- If Alice has the AoE, then your chances of winning are:
  - 60% if you play 0 minions
  - 50% if you play 1 minion
  - 20% if you play 2 minions
- If Alice does NOT have the AoE, then your chances of winning are:
  - 20% if you play 0 minions
  - 60% if you play 1 minion
  - 90% if you play 2 minions

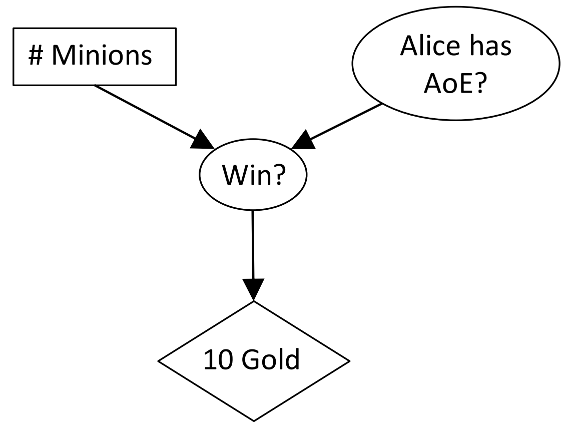

You know that there is a 50% chance that Alice has an AoE.

Winning this game is worth 10 gold and losing is worth 0.

Solution notation: A: Alice has AoE?, W: Win?, M: Number of minions

(a) How much gold would you expect to win choosing 0 minions?

$\sum_w \sum_a (P(w|Minion = 0, a)P(a)R(w) = 10 \sum_a (P(w|Minion = 0, a)P(a) = 10(.6 \cdot .5 + .2 \cdot .5) = 4$

(b) How much gold would you expect to win choosing 1 minion?

$\sum_w \sum_a (P(w|Minion = 1, a)P(a)R(w) = 10 \sum_a (P(w|Minion = 1, a)P(a) = 10(.5 \cdot .5 + .6 \cdot .5) = 5.5$

(c) How much gold would you expect to win choosing 2 minions?

$\sum_w \sum_a (P(w|Minion = 2, a)P(a)R(w) = 10 \sum_a (P(w|Minion = 2, a)P(a) = 10(.2 \cdot .5 + .9 \cdot .5) = 5.5$

(d) How much gold would you expect to win if you know the AoE is in Alice's hand?

$\max_m \sum_w P(w|m, +a)R(w) = 10 \max_m P(w|m, +a) = 10 \max\{.6, .5, .2\} = 6$

(e) How much gold would you expect to win if you know the AoE is NOT in Alice's hand?

$\max_m \sum_w P(w|m, -a)R(w) = 10 \max_m P(w|m, -a) = 10 \max\{.2, .6, .9\} = 9$

(f) How much gold would you be willing to pay for to know whether or not the AoE is in Alice's hand? (Assume your utility of gold is the same as the amount of gold.)

Two. The difference between $MEU(\{\}) = 5.5$ and $MEU(\{A\}) = .5 * 6 + .5 * 9 = 7.5$ is 2.

Q3. GSI Adventures

(a) Missing Exams! The GSIs of 188 are currently looking for where all of the exams have gone! There are 5 GSIs and each one has contact with the other, and they're looking for a grand total of E exams. Imagine Berkeley as an $M \times N$ grid and each GSI starts in a different place. The E exams are spread throughout the Berkeley grid and when a GSI visits a grid space, they are able to pick up all of the exams at that space. During each timestep, a GSI can move 1 grid space. If the exams are not found in T time steps, there will not be time to grade them, and the staff will be forced to give everyone an A. The students know this, so the GSIs must always avoid S students in the grid, otherwise they will steal the exams from them.

(i) Davis and Jacob would like to model this as a search problem. The instructors know where the GSIs start, where the students start, and how they move (that is, student position is a known deterministic function of time). What is a minimal state representation to model this game? Recall that the locations of the exams are not known.

Since we don't know the location of the exams and this is an offline search problem (i.e. we run a simulation of the environment to come up with a full plan before executing it), there is no way we could know how many exams we've picked up. Therefore the only way to model this game as a search problem is to come up with a plan that visits all the squares in T time steps while avoiding the students.

We need to know the position of the GSIs, whether each square has been visited, and the current time step.

(ii) Provide the size of the state representation from above.

$2^{MN}(MN)^5T$

(iii) Which of the following are admissible heuristics for this search problem?

- ☐ The number of exams left to be found
- ☐ The number of exams left to be found divided by 5
- ■ The minimum Manhattan Distance between a GSI and an unvisited grid space
- ☐ The maximum Manhattan Distance between a GSI and an unvisited grid space
- ☐ The number of squares in the grid that have not been visited
- ■ The number of squares in the grid that have not been visited divided by 5

Note that in the correct formulation, we do not have enough information to compute the first two proposed heuristics. In any case, even if we could compute them, neither would be admissible.

(b) The exams have finally been located, and now, it's the students' turn to worry! A student's utility leading up to the exam depends on how hard they study (very hard ($+v$) or just hard ($-v$)) as well as the chance that Davis has a cold around the the exam.

If Davis has a cold ($+c$), he will be too tired to write a hard exam question. He might also be unable to hold office hours, in which case Bob (a reader) will hold office hours instead ($+b$). The decision network and the tables associated with it are shown below:

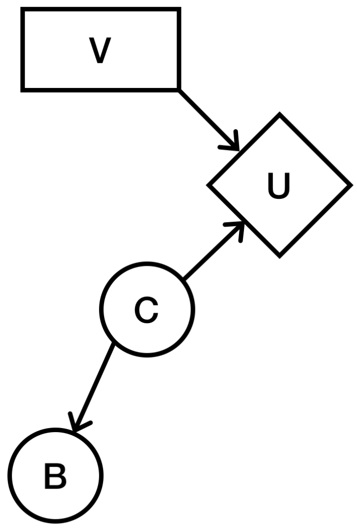

C	$P(C)$
$+c$	0.5
$-c$	0.5

B	C	$P(B\|C)$
$+b$	$+c$	0.8
$+b$	$-c$	0.1
$-b$	$+c$	0.2
$-b$	$-c$	0.9

V	C	U
$+v$	$+c$	200
$+v$	$-c$	120
$-v$	$+c$	250
$-v$	$-c$	90

Calculate the $VPI(B)$. To do this, in the calculations, calculate $MEU()$, $MEU(+b)$, and $MEU(-b)$. In order to get as much partial credit, provide these calculations, as well as any other calculations necessary, in a neat and readable order. Use the calculated tables below in order to help with the calculations. You may leave your answers as expressions in terms of probabilities in the table and your answers to previous parts.

B	$P(B)$
$+b$	0.45
$-b$	0.55

B	C	$P(C\|B)$
$+b$	$+c$	0.89
$+b$	$-c$	0.11
$-b$	$+c$	0.18
$-b$	$-c$	0.81

(i) $MEU() = \max(.5 * 200 + .5 * 120, .5 * 250 + .5 * 90) = 170$

(ii) $MEU(+b) = \max(.89 * 200 + .11 * 120, .89 * 250 + .11 * 90) = 232.4$

(iii) $MEU(-b) = \max(.18 * 200 + .81 * 120, .18 * 250 + .81 * 90) = 133.2$

(iv) $VPI(B) = (.45 * MEU(+b) + .55 * MEU(-b)) - MEU()$