ECE 498SM: Principles of Safe Autonomy

Fall 2020 – Quiz 2

1. This is a closed book exam
2. You may use extra sheets of blank paper for answers as needed
3. You may use one sheet of pre-written notes
4. You may use a calculator
5. Absolutely no interaction between students is allowed
6. Illegible handwriting will be graded as incorrect

Name:________________________________

NetID:________________________________

Problems

1. Conditional Probability: 10 Points
2. Satisfiability: 15 Points
3. Particle Filters: 30 Points
4. Clustering: 15 Points
5. Planning and search 30 Points

1. Conditional Probability (10 points)

A robot is equipped with camera and pedestrian detection software (PD). In tests, (a) out of 100 images with people, PD successfully detected people in 75 images, and (b) out of another set of 100 images without people, PD erroneously detected people in 20.

(A) (4 points) Consider the boolean state variable X which indicates whether a person is present; and the boolean output variable Z which is the output of PD. What is a probabilistic measurement model you can write for the PD system from the above test data?

$$P(Z = 1|X = 1) = \frac{75}{100} = 0.75$$

$$P(Z = 0|X = 1) = \frac{20}{100} = 0.2$$

$$P(Z = 1|X = 0) = 0.25$$

$$P(Z = 0|X = 0) = 0.2$$

(B) (6 points) The prior probability that a person is in view is distributed uniformly. For an image acquired by the robot, PD reports that a person is present (Z=1). Given this observation, what is the probability that a person is actually present in the image P(X=1|Z=1)? Derive the formula and calculate the posterior probability.

$$P(X = 1|Z = 1) = \frac{P(Z = 1|X = 1)P(X = 1)}{P(Z = 1|X = 1)P(X = 1) + P(Z = 1|X = 0)P(X = 0)}$$

$$= \frac{0.75 * 0.5}{0.75 * 0.5 + 0.2 * 0.5} = 0.789$$

2. Satisfiability (15 points)

Consider the boolean formula:

$$\alpha = (\neg x_1 \lor x_3) \land (x_2 \lor \neg x_4 \lor \neg x_5) \land (\neg x_3 \lor \neg x_4)(x_2 \lor x_4 \lor x_5).$$

(A) (3 points) If we attempted to solve the satisfiability of this formula using the brute force search, how many assignments would have to be checked?

$2^5 = 32$

(B) (3 points) Suppose we set $x_1 = True$ what is the resulting formula α_1?

$\alpha_1 = x_3 \land (x_2 \lor \neg x_4 \lor \neg x_5) \land (\neg x_3 \lor \neg x_4) \land (x_2 \lor x_4 \lor x_5)$

(C) (3 points) Which of the three rules in the DP algorithm become applicable to α_1? What is the resulting formula α_2 when this rule is applied to α_1 and what is the corresponding variable assignment?

Unit Propagation, setting x_3 to True:

$\alpha_2 = (x_2 \lor \neg x_4 \lor \neg x_5) \land (\neg x_4) \land (x_2 \lor x_4 \lor x_5)$

Pure literal, setting x_2 to True

$\alpha_2 = x_3 \land (\neg x_3 \lor \neg x_4)$

(D) (3 points) Which of the three rules in the DP algorithm become applicable to α_2? What is the resulting formula α_3 when this rule is applied to α_2 and what is the corresponding variable assignment?

Unit Propagation, Unit Propagation, setting $\neg x_4$ to True:

$\alpha_3 = x_2 \lor x_5$

Pure literal, Unit Propagation, setting x_3 to True

$\alpha_3 = \neg x_4$

Unit Propagation, Pure literal, set x_2 to True

$\alpha_3 = \neg x_4$

Pure literal, Resolution, x_3

$\alpha_3 = \neg x_4$

(E) (3 points) Is α_3 satisfiable and what does that say about validity of $\neg\alpha$?

α_3 is satisfiable. Since there exist cases that $\neg\alpha$ is not satisfied, $\neg\alpha$ is not valid.

3. Particle Filters (30 points)

Consider a world with robot that can be in three locations: $X = x_1, x_2, x_3$. Let Z be a Boolean sensor characterized by the measurement model:

$$p(Z=1|X=x_1)=0.8, \quad p(Z=0|X=x_1)=0.2$$
$$p(Z=1|X=x_2)=0.4, \quad p(Z=0|X=x_2)=0.6$$
$$p(Z=1|X=x_3)=0.1, \quad p(Z=0|X=x_3)=0.9.$$

Suppose the *prior* distribution on the locations is uniform $p(X=x_i)=[\frac{1}{3},\frac{1}{3},\frac{1}{3}]$ and we get the first measurement of Z = 1. Then (as we saw in class) the posterior distribution $p(X=x_i|Z=1)=\left[\frac{8}{13},\frac{4}{13},\frac{1}{13}\right]$. In this problem we are going to use MCL with particle filters representing these distributions.

(A) (12 points) We start with two particles $X_0 = \{x_{01}, x_{02}\}$ sampled using the prior over three locations. There are $3 \times 3 = 9$ possible combinations for X_0 (col 2) each with a probability of 1/9 (col 3). MCL uses the measurement model $p(Z|X)$ to generate particle weights. Again, the first measurement is Z = 1. In column 4, for each possible particle pair $\{x_{01}, x_{02}\}$ write the corresponding **normalized weight pair** (w_{01}, w_{02}) that should be used for resampling (importance sampling).

Suppose we only generate **one** new particle in after resampling. Write the probability of each possible value $x_1, x_2, and\ x_3$ in the last three columns, that is, $p(X_1 = \{x_{11}\}\ and\ X_0 = \{x_{01,}x_{02}\})$.

Number (1)	Sample set (2) X_0	Prob. of set (3)	Normalized weights (4)	Prob. of resampling for each $X_1 = x_{11}$		
				x_1	x_2	x_3
1	x_1, x_1	1/9	1/2, 1/2	1/9	0	0
2	x_1, x_2	1/9	2/3, 1/3	2/27	1/27	0
3	x_1, x_3	1/9	8/9, 1/9	8/81	0	1/81
4	x_2, x_1	1/9	1/3, 2/3	2/27	1/27	0
5	x_2, x_2	1/9	1/2, 1/2	0	1/9	0
6	x_2, x_3	1/9	4/5, 1/5	0	4/45	1/45
7	x_3, x_1	1/9	1/9, 8/9	8/81	0	1/81
8	x_3, x_2	1/9	1/5, 4/5	0	4/45	1/45
9	x_3, x_3	1/9	1/2, 1/2	0	0	1/9
Σ				$p(X_1=\{x_1\})$= 0.4568	$p(X_1=\{x_2\})$= 0.3630	$p(X_1=\{x_3\})$= 0.1802

(B) (3 points) How is the probability distribution $p(X_1 = x_i)$ in the last row of part (A) related to the prior distribution (uniform over x_i) and the posterior probability?

The final probability distribution is supposed to represent the true posterior. However, due to the limited number of particles $N = 2$, the result is biased towards the prior

(C) (5 points) Could the particle filter in this problem estimate the state accurately with a single particle (N=1)? Explain.

No, if using a single particle, we will always select that particle in sampling and resampling. Therefore, it's impossible to recover from a wrong initial estimation. In addition, with only 1 particle, we are not able to provide a good enough estimation to the true posterior.

(D) (5 points) How is the probability distribution $p(X_1 = x_i)$ in the last row of (A) related to the prior probability and the posterior with infinitely many particles $N = \infty$? Why do we want more particles while using particle filter?

When there's infinitely many particles $N = \infty$, the resampling probability distribution will converge to the true posterior. With more particles, the distribution of particles can better represent true posterior and therefore can provide more accurate result.

(E) (5 points) Why are noisy sensor models sometimes more preferable than very precise sensor models for particle filter-based localization?

Noise sensor models can be used to solve the problem of particle deprivation.

4. Clustering (15 points)

Give an example of a set of data points and a set of K cluster heads that demonstrate that the K-means clustering algorithm does not converge to the optimal clustering. That is, you need to show that:

(a) the cluster heads are indeed final, i.e., the algorithm has converged (5 points),

(b) the result gives a clustering cost J_1 (5 points), and

(c) that there is an alternative set of cluster heads that give a lower clustering cost $J_2 < J_1$.

Any proper explanations with explicit calculations of J1 and J2 are acceptable. One example:

Let k = 2, given a set of points {1,2,3,4}

C1: p1 = {1}, p2 = {2,3,4}, $J_1 = 0 + \frac{1}{3} * (1^2 + 1^2) = \frac{2}{3}$

C2: p1 = {1,2}, p2 = {3,4}, $J_2 = \frac{1}{2} * \left(\frac{1}{2}^2 + \frac{1}{2}^2\right) + \frac{1}{2} * \left(\frac{1}{2}^2 + \frac{1}{2}^2\right) = \frac{1}{2}$

5. Planning and Search (20 points)

Consider a *hybrid A* path planner* for a vehicle, similar to the one we used in MP3. The discrete coordinates are integers for each node *n* in the tree. The discrete coordinates are obtained by rounding the continuous coordinates to the nearest integers (Assume x.5 is rounded to x+1). Consider the heuristic function:

$$h(n) = \sqrt{\left(x_n - x_{goal}\right)^2 + \left(y_n - y_{goal}\right)^2}.$$

Consider the goal position $g\ at\ (0{,}0)$. The *cost-to-come* function is defined as $g(n) = g\big(previous(n)\big) + 1$. The steering angle $\delta \in \{0, -30°, 30°\}$; the speed $v = 1.0$. For heading angle θ, x-axis points to $0°$. The vehicle dynamic model is:

$$\theta_{t+1} = \theta_t + \delta$$
$$x_{t+1} = x_t + v * \cos(\theta_{t+1})$$
$$y_{t+1} = y_t + v * \sin(\theta_{t+1})$$

(A) (10 points) The *queue* in the algorithm pops out a node n_0 with continuous coordinates $(x, y, \theta) = (3.0, 5.0, 90°)$ and g(n_0) = 13. Write the continuous and discrete coordinates of all the new nodes that will be expanded from n_0 and the corresponding g, h, f values.

Node name	**Cont. coord**	**Disc. coord**	**g**	**h**	**f**
n_0	(3.0,5.0,90°)		13		
n1	(3,6,90°)	(3,6,90°)	14	6.71	20.71
n2	(2.5,5.866,120°)	(3,6,120°)	14	6.377	20.377
n3	(3.5,5.866,60°)	(4,6,60°)	14	6.832	20.832

(B) **(3 points)** Which node will be popped out from the queue after n_0? Write down its continuous coordinates.
n2 or node (2.5, 5.866, 120) should be popped out or if (a) is incorrect, the node with lowest f should be popped out

(C) **(3 points)** Is hybrid A* guaranteed to find optimal results in this environment with the heuristic function? Why or why not?
Yes, because the heuristic is admissible.

(D) **(4 points)** If we replace the heuristic function with Manhattan distance, is hybrid A* guaranteed to find optimal results in this environment with the new heuristic function? Why or why not?
No. Manhattan distance is not admissible since it could overestimate the cost. For example, suppose goal (0,0), for node (1.0, 1.0, -135°), Manhattan distance is 2 but the actual distance is 1.414.

(E) **(10 points)** Given an example graph to illustrate that best-first search is incomplete. That is, draw a graph (in the same format as in lecture slides), with the start and goal vertices, and a heuristic function, such that although there is a path from start to goal the algorithm fails to find it. Further you need to give a run of the best-first search algorithm to illustrate how it fails to find a path.
Proper explanations with graph and executions are acceptable.