Printout

Tuesday, November 7, 2023 10:57 AM

Announcements: 11/7/2023

- **Guest lecture (11/09) - Attendance is mandatory**
- **We are doubling the F1-tenth Lab hours**
- **After Thanksgiving break F1-tenth supports (e.g. debuging the system) could be limited**
- **Utilized most of your time before/during the Thanksgiving break**
- **Most of the GEM slots are unutilized - utilize the slots**

Midter2 Review Session: 11/7/2023

Topics:

- **Filtering**
  - **Bayes**
  - **Histogram/grid**
  - **Particle**
  - **MCL**
- **Search and Planning**
  - **Uniform**
  - **Greedy**
  - **A/A*/Hybrid A***
  - **PRM**
  - **RRT/RRG**

Practice Questions will be Released:

Visit office hours

Discrete Bayes Filtering Review:

we write the probability of a random variable X taking the value x as P(X=x) or as P(x) in short.

Sequence of states $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_t$, is written in short as $\boldsymbol{x}_{1:t}$

Sequence of measurements $\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_t$, is written in short as $\boldsymbol{z}_{1:t}$

Sequence of control inputs $\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_t$, is written in short as $\boldsymbol{u}_{1:t}$

Belief: Robot's knowledge about the state of the environment

True state is unknowable / measurable typically, so, robot must infer state from data and we have to distinguish this inferred/estimated state from the actual state

We write the definition of Belief over state (x_t) in terms of conditional probability given measurements z_t and ~~control u_t~~:

$$Bel(x_t) = P(x_t \mid z_{1:t}, u_{1:t})$$

We want to write $Bel(x_t)$ **using** $\overline{Bel}(x_t)$

$$Bel(x_t) = \eta\, P(z_t \mid x_t)\, \overline{Bel}(x_t)$$

Where,

$$\overline{Bel}(x_t) = P(x_t \mid z_{1:t-1}, u_{1:t})$$

η = normalization term

$$\Rightarrow P(z_t \mid z_{1:t-1}, u_{1:t})$$

correction term :

$$P(z_t \mid x_t, z_{1:t-1}, u_{1:t})$$

Assuming that the measurement model is complete, then the correction term becomes

$$= P(z_t \mid x_t)$$

let's say we have three random variable X, Y, Z,

$$P(X \mid Y, Z) = \frac{P(X, Y, Z)}{P(Y, Z)}$$

$$P(Y \mid X Z) = \frac{P(X, Y, Z)}{P(X, Z)}$$

$$P(X \mid Y, Z) = \frac{P(Y \mid XZ)\, P(X, Z)}{P(Y, Z)}$$

$$P(X/Z) = \frac{P(X, Z)}{P(Z)}$$

$$P(Y/Z) = \frac{P(Y, Z)}{P(Z)}$$

$$P(X \mid Y, Z) = \frac{P(Y \mid X Z)\, P(X \mid Z)\, \cancel{P(Z)}}{P(Y \mid Z)\, \cancel{P(Z)}}$$

$$\boxed{P(X \mid Y, Z) = \frac{P(Y \mid X, Z)\, P(X/Z)}{P(Y/Z)}}$$

choose $X = x_t$

$Y = z_t$

$Z = z_{1:t-1}, u_{1:t}$

$$\doteq \frac{P(z_t \mid x_t, \cancel{z_{1:t-1}, u_{1:t}})\, P(x_t \mid z_{1:t-1}, u_{1:t})}{P(z_t \mid z_{1:t-1}, u_{1:t})}$$

$$= \eta \quad \_\_\_\_\_\_\_\_\_\_\_\_$$

Particle filtering

You are kidnapped in a plane with just an altitude map and you have to use particle filters to estimate your location.

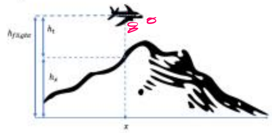

The aircraft is flying in straight line at a speed of 1 per unit time and keeping the same sea level altitude of $h_{flight} = 100$. Let the (unknown) horizonal location of the airplane at time t be X_t. Below the aircraft is a mountain. You have a map (below) that gives the exact height h_x of the mountain at any position x. A noisy sensor mounted on the airplane measures the vertical distance h_t between plane and mountain at any time t. You are going to use m particles in your calculation initialized as in the table below. For calculating the importance factor W_t^m of the m^{th} particle, use:

$W_t^m = e^{-((h_{flight} - h_{x_m}) - h_t)^2}$, where x_m is the position of the m^{th} particle. Assume we do not resample particles.

Altitude map:

Position (x)	1	2	3	4	5	6	7	8	9	10
Mountain exact altitude (h_x)	15	21	32	29	25	23	16	10	8	5

a. Suppose you saw the sensor reading $h_t = 80$. Fill in the table to update the particle filter. Writing expressions will be adequate.

Particle number m	1	2	3	4	5
X_{t-1}^m	3	6	1	5	9
X_t^m	4	7	2	6	10
W_t^m	e^{-81}	e^{-16}	e^{-1}	e^{-9}	e^{-225}

$\eta = \frac{1}{(e^{-81} + e^{-16} + e^{-1} + e^{-9} + e^{-225})}$

b. What is the estimated position $\overline{X_t}$ at time t? Just write down the equation; no need to calculate or simplify.

$\overline{X_t} = \eta(4 \times e^{-81} + 7 \times e^{-16} + 2 \times e^{-1} + 6 \times e^{-9} + 10 \times e^{-225})$

a. Consider a search algorithm like A search in which the function $f(n) = (2-w)g(n) + wh(n)$ is used to choose the path that is expanded next. Recall, $g(n)$ is the cost to arrive at n and $h(n)$ is the cost-to-go heuristic. If the heuristic h is *admissible*, what kind of search does this algorithm perform for each of these values of w and why?

	Type of search	Reasoning
$w = 0$	uniform	
$w = 1$	A / A*	
$w = 2$	greedy	

b. Draw an example graph to illustrate why best-first search is incomplete. Draw the graph with the start and goal vertices, label the edge costs and the heuristic function values.

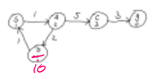

	Cost	h
<s>	0	1
<a, s>	1	2
<e, a, s>	6	3
<b, a, s>	3	1
<s, b, a, s>	4	1
<a, s, b, a, s>	5	2

c. **(5 points)** Write a run of the best-first search algorithm on the above graph to illustrate how it fails to find a path.