EE211: Robotic Perception and Intelligence

Lecture 10 Probability in Robotics

Jiankun WANG

Department of Electronic and Electrical Engineering
Southern University of Science and Technology

Undergraduate Course, Dec 2024

Outline

1. Robot Environment Interaction

2. Bayes Filters

Outline

1. Robot Environment Interaction

2. Bayes Filters

Overview

- How does the robot interact with environments?

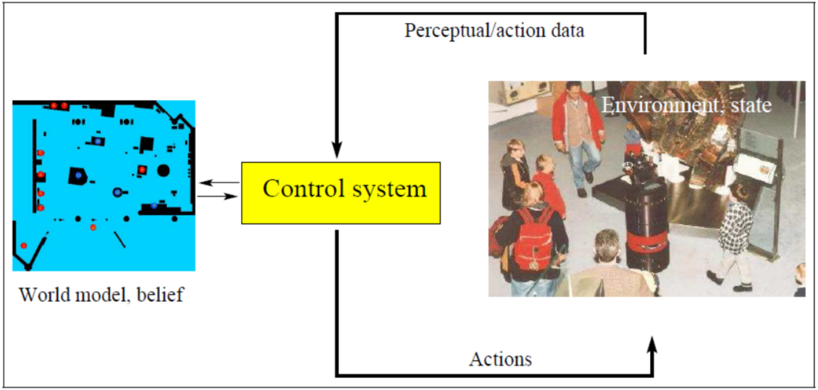

Probabilistic Robotics.

Environment Interaction

- State: x_t, robot pose, velocity, etc.
- Sensor measurements: z_t. Perception is the process by which the robot uses its sensors to obtain information about the state of its environment.
- Control actions: u_t change the state of the robot itself and the world. They do so by actively asserting forces on the robot's environment.

Probabilistic Generative Laws

- The evolution of state and measurements:

$$p(x_t | x_{0:t-1}, z_{1:t-1}, u_{1:t}).$$

- Markov chains: no variables prior to x_t may influence the stochastic evolution of future states, unless this dependence is mediated through the state x_t.
- State transition probability:

$$p(x_t | x_{0:t-1}, z_{1:t-1}, u_{1:t}) = p(x_t | x_{t-1}, u_t).$$

- Measurement probability:

$$p(z_t | x_{0:t-1}, z_{1:t-1}, u_{1:t}) = p(z_t | x_t).$$

Belief Distributions

- A belief reflects the robot's internal knowledge about the state of the environment.
- For example, a robot's pose might be $x =< 14, 12, 7 >$ in some global coordinate system, but it usually cannot know its pose, since poses are not measurable directly (not even with GPS!). Instead, the robot must infer its pose from data.
- A belief distribution assigns a probability (or density value) to each possible hypothesis with regards to the true state. Belief distributions are posterior probabilities over state variables conditioned on the available data.

$$bel(x_t) = p(x_t | z_{1:t}, u_{1:t}).$$

Belief Distributions

- Occasionally, it will prove useful to calculate a posterior before incorporating z_t, just after executing the control u_t:

$$\overline{bel}(x_t) = p(x_t | z_{1:t-1}, u_{1:t}),$$

which is the prediction of the state at time t, before incorporating the measurement at time t.

- Calculating $bel(x_t)$ from $\overline{bel}(x_t)$ is called correction or the measurement update:

$$bel(x_t) = p(x_t | z_{1:t}, u_{1:t}).$$

Outline

Bayes Filter Algorithm

- The most general algorithm for calculating beliefs.
- This algorithm calculates the belief distribution *bel* from measurement and control data.

Algorithm Bayes_filter($bel(x_{t-1}), u_t, z_t$**):**
    for all x_t do
        $\overline{bel}(x_t) = \int p(x_t \mid u_t, x_{t-1})\, bel(x_{t-1})\, dx$
        $bel(x_t) = \eta\, p(z_t \mid x_t)\, \overline{bel}(x_t)$
    endfor
    return $bel(x_t)$

Example

- A mobile robot estimating the state of a door.

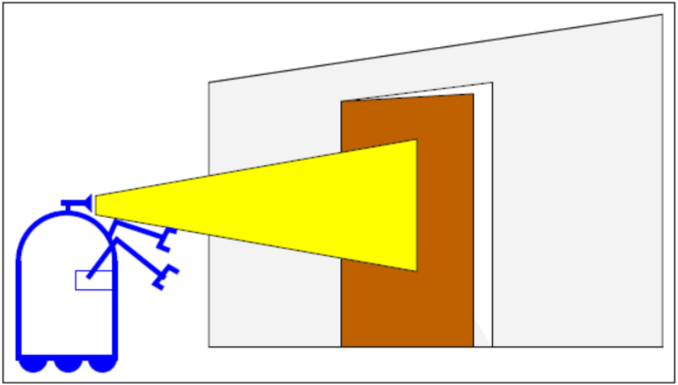

Example

$$bel(X_0 = \textbf{open}) = 0.5 \quad (2.37)$$
$$bel(X_0 = \textbf{closed}) = 0.5 \quad (2.38)$$

Let us furthermore assume the robot's sensors are noisy. The noise is characterized by the following conditional probabilities:

$$\begin{aligned} p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_open}) &= 0.6 \\ p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_open}) &= 0.4 \end{aligned} \quad (2.39)$$

and

$$\begin{aligned} p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_closed}) &= 0.2 \\ p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_closed}) &= 0.8 \end{aligned} \quad (2.40)$$

These probabilities suggest that the robot's sensors are relatively reliable in detecting a *closed* door, in that the error probability is 0.2. However, when the door is open, it has a 0.4 probability of a false measurement.

Example

These probabilities suggest that the robot's sensors are relatively reliable in detecting a *closed* door, in that the error probability is 0.2. However, when the door is open, it has a 0.4 probability of a false measurement.

Finally, let us assume the robot uses its manipulator to push the door open. If the door is already open, it will remain open. It it is closed, the robot has a 0.8 chance that it will be open afterwards:

$$\begin{aligned} p(X_t = \textbf{is\_open} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_open}) &= 1 \\ p(X_t = \textbf{is\_closed} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_open}) &= 0 \qquad (2.41) \\ p(X_t = \textbf{is\_open} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_closed}) &= 0.8 \\ p(X_t = \textbf{is\_closed} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_closed}) &= 0.2 \qquad (2.42) \end{aligned}$$

It can also choose not to use its manipulator, in which case the state of the world does not change. This is stated by the following conditional probabilities:

$$\begin{aligned} p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) &= 1 \\ p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) &= 0 \qquad (2.43) \\ p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) &= 0 \\ p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) &= 1 \qquad (2.44) \end{aligned}$$

Example

Suppose at time t, the robot takes no control action but it senses an open door. The resulting posterior belief is calculated by the Bayes filter using the prior belief $bel(X_0)$, the control $u_1 = \textbf{do\_nothing}$, and the measurement **sense_open** as input. Since the state space is finite, the integral in Line 3 turns into a finite sum:

$$
\begin{aligned}
\overline{bel}(x_1) &= \int p(x_1 \mid u_1, x_0)\; bel(x_0)\; dx_0 \\
&= \sum_{x_0} p(x_1 \mid u_1, x_0)\; bel(x_0) \\
&= p(x_1 \mid U_1 = \textbf{do\_nothing}, X_0 = \textbf{is\_open})\; bel(X_0 = \textbf{is\_open}) \\
&\quad + p(x_1 \mid U_1 = \textbf{do\_nothing}, X_0 = \textbf{is\_closed})\; bel(X_0 = \textbf{is\_closed})
\end{aligned}
\tag{2.45}
$$

We can now substitute the two possible values for the state variable X_1. For the hypothesis $X_1 = \textbf{is\_open}$, we obtain

$$
\begin{aligned}
&\overline{bel}(X_1 = \textbf{is\_open}) \\
&= p(X_1 = \textbf{is\_open} \mid U_1 = \textbf{do\_nothing}, X_0 = \textbf{is\_open})\; bel(X_0 = \textbf{is\_open}) \\
&\quad + p(X_1 = \textbf{is\_open} \mid U_1 = \textbf{do\_nothing}, X_0 = \textbf{is\_closed})\; bel(X_0 = \textbf{is\_closed}) \\
&= 1 \cdot 0.5 + 0 \cdot 0.5 \;=\; 0.5
\end{aligned}
\tag{2.46}
$$

Example

Likewise, for $X_1 = \textbf{is\_closed}$ we get

$$\begin{aligned}
&\overline{bel}(X_1 = \textbf{is_closed}) \\
&= p(X_1 = \textbf{is_closed} \mid U_1 = \textbf{do_nothing}, X_0 = \textbf{is_open})\, bel(X_0 = \textbf{is_open}) \\
&\quad + p(X_1 = \textbf{is_closed} \mid U_1 = \textbf{do_nothing}, X_0 = \textbf{is_closed})\, bel(X_0 = \textbf{is_closed}) \\
&= 0 \cdot 0.5 + 1 \cdot 0.5 \;=\; 0.5
\end{aligned} \tag{2.47}$$

The fact that the belief $\overline{bel}(x_1)$ equals our prior belief $bel(x_0)$ should not surprise, as the action **do_nothing** does not affect the state of the world; neither does the world change over time by itself in our example.

Incorporating the measurement, however, changes the belief. Line 4 of the Bayes filter algorithm implies

$$bel(x_1) \;=\; \eta\, p(Z_1 = \textbf{sense\_open} \mid x_1)\, \overline{bel}(x_1)\,. \tag{2.48}$$

Example

For the two possible cases, $X_1 = \textbf{is\_open}$ and $X_1 = \textbf{is\_closed}$, we get

$$\begin{aligned} &bel(X_1 = \textbf{is\_open}) \\ &\quad = \quad \eta\, p(Z_1 = \textbf{sense\_open} \mid X_1 = \textbf{is\_open})\, \overline{bel}(X_1 = \textbf{is\_open}) \\ &\quad = \quad \eta\, 0.6 \cdot 0.5 \;=\; \eta\, 0.3 \end{aligned} \tag{2.49}$$

and

$$\begin{aligned} &bel(X_1 = \textbf{is\_closed}) \\ &\quad = \quad \eta\, p(Z_1 = \textbf{sense\_open} \mid X_1 = \textbf{is\_closed})\, \overline{bel}(X_1 = \textbf{is\_closed}) \\ &\quad = \quad \eta\, 0.2 \cdot 0.5 \;=\; \eta\, 0.1 \end{aligned} \tag{2.50}$$

The normalizer η is now easily calculated:

$$\eta \quad = \quad (0.3 + 0.1)^{-1} \;=\; 2.5 \tag{2.51}$$

Hence, we have

$$\begin{aligned} bel(X_1 = \textbf{is\_open}) &= 0.75 \\ bel(X_1 = \textbf{is\_closed}) &= 0.25 \end{aligned} \tag{2.52}$$

Example

This calculation is now easily iterated for the next time step. As the reader easily verifies, for $u_2 = \textbf{push}$ and $z_2 = \textbf{sense\_open}$ we get

$$\begin{aligned} \overline{bel}(X_2 = \textbf{is\_open}) &= 1 \cdot 0.75 + 0.8 \cdot 0.25 = 0.95 \\ \overline{bel}(X_2 = \textbf{is\_closed}) &= 0 \cdot 0.75 + 0.2 \cdot 0.25 = 0.05\,, \end{aligned} \tag{2.53}$$

and

$$\begin{aligned} bel(X_2 = \textbf{is\_open}) &= \eta\, 0.6 \cdot 0.95 \approx 0.983 \\ bel(X_2 = \textbf{is\_closed}) &= \eta\, 0.2 \cdot 0.05 \approx 0.017\,. \end{aligned} \tag{2.54}$$

Q & A

1. Robot Environment Interaction

2. Bayes Filters