ECE 4960

Prof. Kirstin Hagelskjær Petersen
kirstin@cornell.edu

Vivek Thangavelu
vs353@cornell.edu

Fast Robots

Lecture Outline

- Recap
- The Bayes Filter Algorithm
- Examples (2)

In the last episode…

Robot-Environment Interaction

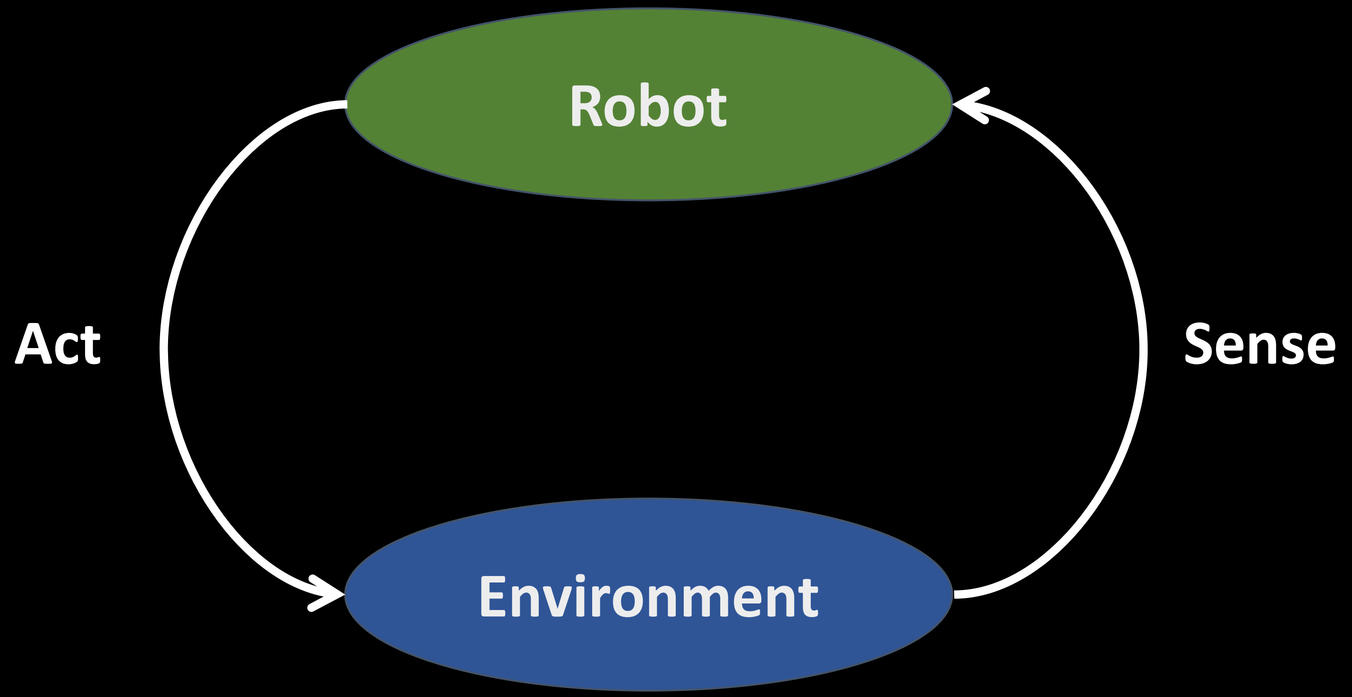

- Two fundamental types of interaction between a robot and its environment:
  - **Sensor Measurements/Observations**
  - **Control Actions**

Robot-Environment Model

- The model helps us express a robot-environment interaction using probability
- Robot environment interactions are typically modeled as a discrete time system
- The state at time t will be denoted by as x_t
- A sensor measurement at time t will be denoted as z_t
  - Provides information about the environment's state, and hence tends to increase the robot's knowledge

Conventions as per Siegwart, R., Nourbakhsh, I.R. and Scaramuzza, D., 2011. Introduction to autonomous mobile robots. MIT press.

Robot-Environment Model

- A control action will be denoted by u_t, which carries information about the change of the robot state in the time interval $(t\text{-}1{:}t]$ i.e. u_t indices a transition from state x_{t-1} to x_t
  - Tends to induce a loss of knowledge due to inherent noise in robot actuation and stochasticity of robot environments
- For no particular reason, we assume the robot executes a control action u_t first and then takes a measurement z_t
- There is only one control action per time step t, and include as legal action "*do-nothing*"
- There is only one measurement per time step t

Probabilistic Generative Laws

- The evolution of state and measurements is governed by probabilistic laws
- We are interested in generative laws concerning the evolution of the:
  - State: How is x_t generated stochastically?
  - Measurements: How is z_t generated stochastically?

Markov Assumption

The Markov assumption postulates that past and future data are independent if one knows the current state

- It is assumed that future states depend only on the current state, not on the events that occurred before it
- The knowledge of past states, measurements or controls carry no additional information that would help us predict the future more accurately
- Hence, if we somehow model the state using the Markov assumption, we can recursively estimate the robot state x_t from the previous state x_{t-1} , measurement data z_t and control input u_t
- **This does not mean the future is a deterministic function of the current state**

State Generative Model

- If state x_t is modeled under the Markov Assumption, then

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

(conditional independence)

Knowledge of only the previous state x_{t-1} and control u_t is sufficient to predict x_t

Measurement Generative Model

- If x_t conforms to the Markov Assumption, then

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

(conditional independence)

- The state x_t is sufficient to predict the (potentially noisy) measurements

Knowledge of any other variable, such as past measurements, controls, or even past states, is irrelevant under the Markov Assumption

Robot-Environment Model

+

Markov Assumption

+

Bayes Theorem

=

Bayes Filter

Robot Belief

- A robot maintain an internal representation of itself and the environment
- Probabilistic robotics represents beliefs through posterior conditional probability distributions i.e. probability distributions over state variables conditioned on available data
- The belief of a robot is the posterior distribution over the state of the environment, given all past sensor measurements and all past controls
- Belief over a state variable x_t is denoted by $bel(x_t)$ which is an abbreviation for

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

- Occasionally it prove useful in our probabilistic algorithms to define a (prior) belief before incorporating the latest measurement z_t

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

Bayes Filter

- It is a recursive algorithm that calculates the belief distribution from measurements and control data
- The pseudo algorithm depicts one iteration of the Bayes Filter algorithm

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. for all x_t do
3. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | u_t, x_{t-1}) \; bel(x_{t-1})$
4. $bel(x_t) = \eta \, p(z_t | x_t) \; \overline{bel}(x_t)$
5. endfor
6. return $bel(x_t)$

Bayes Filter

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. for all x_t do
3. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | u_t, x_{t-1})\, bel(x_{t-1})$ [Prediction Step]
4. $bel(x_t) = \eta\, p(z_t | x_t)\, \overline{bel}(x_t)$ [Update/Measurement Step]
5. endfor
6. return $bel(x_t)$

Transition Probability / Action Model: $p(x_t | u_t, x_{t-1})$

Measurement Probability / Sensor Model: $p(z_t | x_t)$

Bayes Filter

- Its input is the belief $bel(x_{t-1})$ at time $t-1$ along with the most recent control input u_t and the most recent measurement z_t
- Its output is the belief $bel(x_t)$ at time t
- Control/Prediction Step:
  - In line 3, it process the control u_t by calculating a belief over the state x_t based on the belief over state x_{t-1}
- Update/Measurement Step:
  - In line 4, the posterior belief $bel(x_t)$ is calculated based on the prior belief $\overline{bel}(x_t)$ and the probability that the measurement z_t may have been observed

This is done for each hypothetical posterior state x_t

Dynamical Stochastic Model

- $p(x_t | x_{t-1}, u_t)$
  - It is known as the state transition probability
  - It specifies how the robot state evolves over time as a function of robot controls u_t
- $p(z_t | x_t)$
  - It is known as the measurement probability
  - It specifies how the measurements are generated from the robot state x_t
  - Informally, you may think of measurements as noisy projections of the state
- Remember that these predictions are stochastic and not deterministic

Bayes Filter - Initial Conditions

- To compute the posterior belief recursively, the algorithm requires an initial belief $bel(x_0)$ at time $t = 0$
- If we know the initial state with absolute certainty, we can initialize a point mass distribution that centers all probability mass on the correct value of x_0 and assign zero everywhere else
- If we are entirely ignorant of the initial state, we can initialize it with a uniform probability distribution over all the possible states

Dynamical Stochastic Model

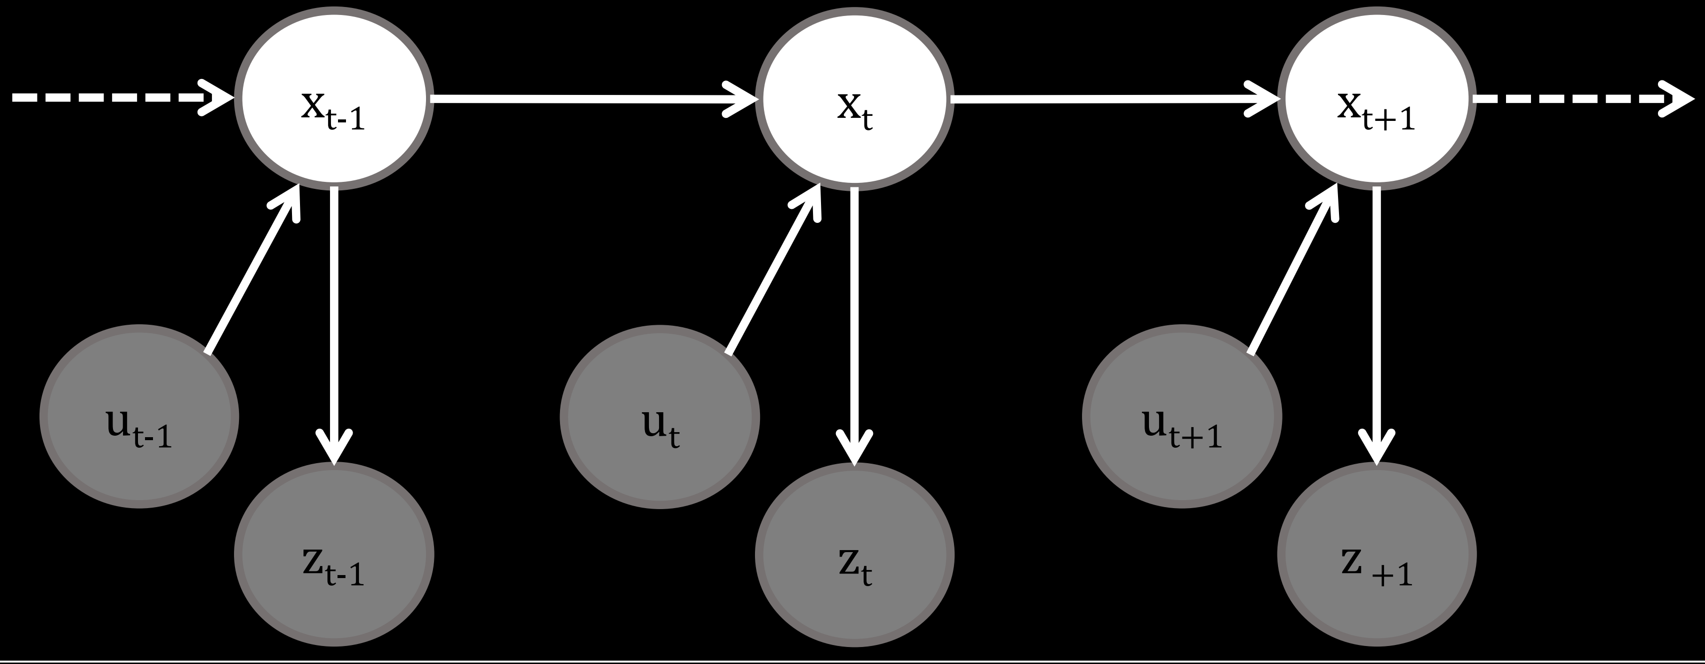

Dynamic Bayes Network that characterizes the evolution of controls, states, and measurements.

- $p(x_t \mid x_{t-1}, u_t)$ and $p(z_t \mid x_t)$ together describe the **dynamical stochastic system** of the robot and its environment
- Such a generative model is also known as a Hidden Markov Model (HMM) or Dynamic Bayes Network (DBN)

Bayes Filter

Example 1

Bayes Filter - Example 1 (1)

- A robot can "*observe*" a door through its sensor and can interact with it by "*pushing*"
- The door may be in one of two states: *open* or *close*
- The sensors and the actuators on the robot are noisy
- The robot can either *push* or *do_nothing* at any given time
- The probability that the robot can sense a **open** door is 0.6
- The probability that the robot can sense a **closed** door is 0.8
- After a *push* action, probability that a door is *open* if it was previously open is 1
- After a *push* action, probability that a door is *open* if it was previously closed is 0.8
- If the robot *does nothing*, the door continues to be in the previous state

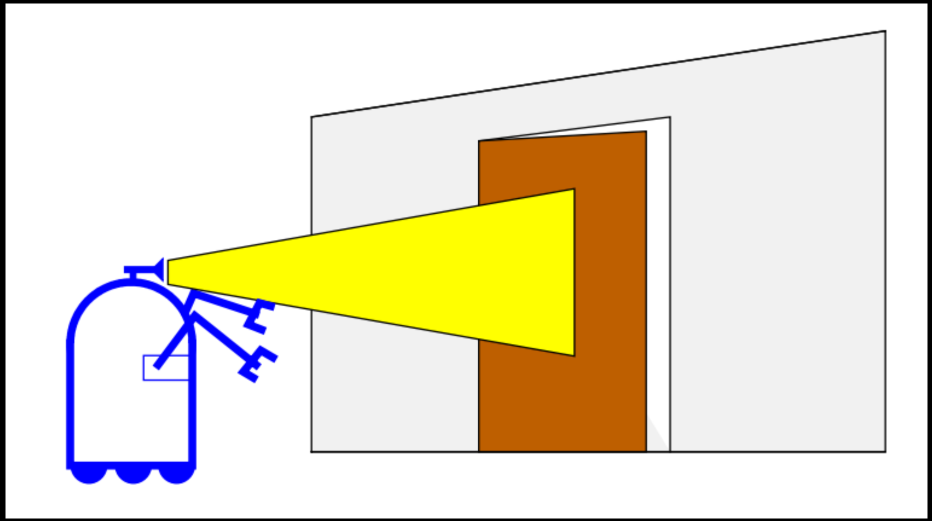

Bayes Filter - Example 1 (2)

Initial Conditions

$$bel(X_0 = is\_closed) = bel(X_0 = is\_open) = 0.5$$

Measurement Probability

$$p(Z_t = sense\_closed | X_t = is\_closed) = 0.8$$
$$p(Z_t = sense\_open | X_t = is\_closed) = 0.2$$
$$p(Z_t = sense\_closed | X_t = is\_open) = 0.4$$
$$p(Z_t = sense\_open | X_t = is\_open) = 0.6$$

Control Action/Transition Probability

$$p(X_t = is\_closed | U_t = do\_nothing, X_{t-1} = is\_closed) = 1$$
$$p(X_t = is\_open | U_t = do\_nothing, X_{t-1} = is\_closed) = 0$$
$$p(X_t = is\_closed | U_t = do\_nothing, X_{t-1} = is\_open) = 0$$
$$p(X_t = is\_open | U_t = do\_nothing, X_{t-1} = is\_open) = 1$$

$$p(X_t = is\_closed | U_t = push, X_{t-1} = is\_closed) = 0.2$$
$$p(X_t = is\_open | U_t = push, X_{t-1} = is\_closed) = 0.8$$
$$p(X_t = is\_closed | U_t = push, X_{t-1} = is\_open) = 0$$
$$p(X_t = is\_open | U_t = push, X_{t-1} = is\_open) = 1$$

Bayes Filter - Example 1 (3)

$u_1 = do\_nothing \; and \; z_1 = sense\_open$

Incorporating action

$$\overline{bel}(x_1) = \sum_{x_0} p(x_1|u_1, x_0)\, bel(x_0)$$

$$= p(x_1|U_1 = do\_nothing, X_0 = is\_open)\, bel(X_0 = is\_open)$$

$$+ p(x_1|U_1 = do\_nothing, X_0 = is\_closed)\, bel(X_0 = is\_closed)$$

We can now substitute the two possible values for the state variable X_1.

For the hypothesis $X_1 = is\_open$, we obtain

$$\overline{bel}(X_1 = is\_open) = p(X_1 = is\_open|U_1 = do\_nothing, X_0 = is\_open)\, bel(X_0 = is\_open)$$

$$+ p(X_1 = is\_open|U_1 = do\_nothing, X_0 = is\_closed)\, bel(X_0 = is\_closed)$$

$$= 1 \times 0.5 + 0 \times 0.5 = 0.5$$

Bayes Filter - Example 1 (4)

$u_1 = do\_nothing \text{ and } z_1 = sense\_open$

<u>Incorporating action</u>

$$\overline{bel}(x_1) = \sum_{x_0} p(x_1|u_1, x_0)\, bel(x_0)$$

$$= p(x_1|U_1 = do\_nothing, X_0 = is\_open)\, bel(X_0 = is\_open)$$

$$+ p(x_1|U_1 = do\_nothing, X_0 = is\_closed)\, bel(X_0 = is\_closed)$$

We can now substitute the two possible values for the state variable X_1.

For the hypothesis $X_1 = is\_closed$, we obtain

$$\overline{bel}(X_1 = is\_closed) = p(X_1 = is\_closed|U_1 = do\_nothing, X_0 = is\_open)\, bel(X_0 = is\_open)$$

$$+ p(X_1 = is\_closed|U_1 = do\_nothing, X_0 = is\_closed)\, bel(X_0 = is\_closed)$$

$$= 0 \times 0.5 + 1 \times 0.5 = 0.5$$

Bayes Filter - Example 1 (5)

$u_1 = do\_nothing \ and \ z_1 = sense\_open$

Incorporating measurement

$\overline{bel}(X_1 = is\_open) = 0.5$

$\overline{bel}(X_1 = is\_closed) = 0.5$

$$bel(x_1) = \eta \, p(Z_1 = sense\_open | \, x_1)\overline{bel}(x_1)$$

For two possible cases, $X_1 = is\_open$ and $X_1 = is\_closed$, we get

$$bel(X_1 = is\_open) = \eta \, p(Z_1 = sense\_open \,|X_1 = is\_open) \, \overline{bel}(X_1 = is\_open)$$
$$= \eta \times 0.6 \times 0.5 = \eta \, 0.3$$

$$bel(X_1 = is\_closed) = \eta \, p(Z_1 = sense\_open \,|X_1 = is\_closed) \, \overline{bel}(X_1 = is\_closed)$$
$$= \eta \times 0.2 \times 0.5 = \eta \, 0.1$$

Bayes Filter - Example 1 (6)

$u_1 = do\_nothing \; and \; z_1 = sense\_open$

Incorporating measurement

$bel(X_1 = is\_open) = \eta \; 0.3$

$bel(X_1 = is\_closed) = \eta \; 0.1$

The normalizer η is now calculated as follows:

$$\eta = (0.3 + 0.1)^{-1} = 2.5$$

$$bel(X_1 = is\_closed) = \eta 0.1 = 0.25$$
$$bel(X_1 = is\_open) = \eta 0.3 = 0.75$$

Better than initial belief at time t=0!

Bayes Filter - Example 1 (7)

$$u_2 = push \text{ and } z_2 = sense\_open$$

Prediction update:

$$\overline{bel}(X_2 = is\_closed) = 0 \times 0.75 + 0.2 \times 0.25 = 0.05$$

$$\overline{bel}(X_2 = is\_open) = 1 \times 0.75 + 0.8 \times 0.25 = 0.95$$

Measurement Update:

$$bel(X_2 = is\_closed) = 0 \times 0.2 \times 0.05 \simeq 0.017$$

$$bel(X_2 = is\_open) = \eta \times 0.6 \times 0.95 \simeq 0.983$$

Waaaay better than the initial belief at time t=0!

Summary of Bayes Filter

- The robot is modeled as performing a series of alternating measurements and actions
- Given:
  - Sensor model $p(z|x)$
  - Action model $p(x|u, x')$
  - Initial Conditions $p(x_0)$
- To compute:
  - Estimate state x of a dynamical system
  - Posterior of the state (Belief):
    $$Bel(x_t) = p(x_t|u_1, z_1, \dots, u_t, z_t)$$

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. for all x_t do
3. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t|u_t, x_{t-1})\ bel(x_{t-1})$
4. $bel(x_t) = \eta\, p(z_t|x_t)\ \overline{bel}(x_t)$
5. endfor
6. return $bel(x_t)$

<u>Short-hand Notation</u>: x is current state and x' is previous state

Summary of Bayes Filter

- Prediction Step:
  - Incorporating action, which **increases** uncertainty
  - Compute $\overline{bel}(x_t) = p(x_t|z_{1:t-1}, u_{1:t})$
  - Requires Action Model: *p(x|u, x')*
- Measurement/Update Step:
  - Incorporating measurement, which **decreases** uncertainty
  - Compute $bel(x_t) = p(x_t|z_{1:t}, u_{1:t})$
  - Requires Sensor Model: *p(z|x)*

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. for all x_t do
3. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t|u_t, x_{t-1})\ bel(x_{t-1})$
4. $bel(x_t) = \eta\, p(z_t|x_t)\ \overline{bel}(x_t)$
5. endfor
6. return $bel(x_t)$

Bayes Filter

Example 2

Bayes Filter - Example 2 (1)

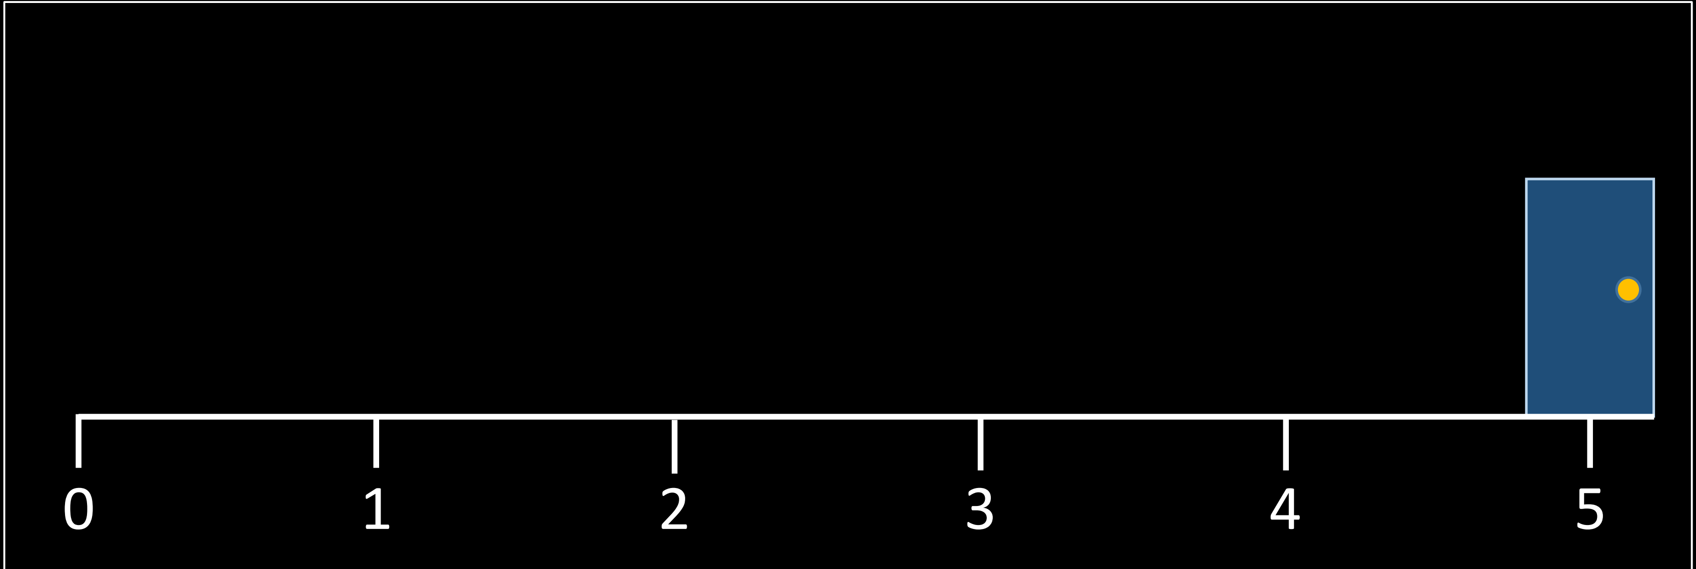

Map

Sensor Model
(for sensing a door)

$$
\begin{aligned}
p(x+1|x,u=+1) &= 0.5 \\
p(x|x,u=+1) &= 0.5 \\
\\
p(x\text{-}1|x,u=\text{-}1) &= 0.5 \\
p(x|x,u=\text{-}1) &= 0.5
\end{aligned}
$$

Motion Model

Bayes Filter - Example 2 (2)

At t=0, we have no information about the robot. Therefore, we assume that it could be at any location i.e., prior is uniform.

State	0	1	2	3	4	5
$bel(x_0)$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$

Bayes Filter - Example 2 (2)

At $t = 0$

State	0	1	2	3	4	5
$bel(x_0)$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$

At $t = 1$: $U_1 = \text{do\_nothing}$, $Z_1 = \text{wall}$

State	0	1	2	3	4	5
$bel(x_1)$	0	0	0	0	$\frac{\frac{1}{6}\times\frac{1}{4}}{\frac{1}{6}\times\frac{1}{4}+\frac{1}{6}\times\frac{1}{2}}$	$\frac{\frac{1}{6}\times\frac{1}{2}}{\frac{1}{6}\times\frac{1}{4}+\frac{1}{6}\times\frac{1}{2}}$

State	0	1	2	3	4	5
$bel(x_1)$	0	0	0	0	$\frac{1}{3}$	$\frac{2}{3}$

Bayes Filter - Example 2 (3)

At $t = 1$

State	0	1	2	3	4	5
$bel(x_1)$	0	0	0	0	$\frac{1}{3}$	$\frac{2}{3}$

At $t = 2$: $U_2 =$ -1

State	0	1	2	3	4	5
$\overline{bel}(x_2)$	0	0	0	$\frac{1}{3} \times \frac{1}{2}$	$\frac{1}{3} \times \frac{1}{2} + \frac{2}{3} \times \frac{1}{2}$	$\frac{2}{3} \times \frac{1}{2}$

State	0	1	2	3	4	5
$\overline{bel}(x_2)$	0	0	0	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{3}$

Bayes Filter - Example 2 (4)

At $t = 2$

State	0	1	2	3	4	5
$\overline{bel}(x_2)$	0	0	0	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{3}$

At $t = 2$: $Z_2 = \text{wall}$

State	0	1	2	3	4	5
$bel(x_2)$	0	0	0	$\frac{1}{6} \times 0$	$\frac{\frac{1}{2} \times \frac{1}{4}}{\frac{1}{2} \times \frac{1}{4} + \frac{1}{3} \times \frac{1}{2}}$	$\frac{\frac{1}{3} \times \frac{1}{2}}{\frac{1}{2} \times \frac{1}{4} + \frac{1}{3} \times \frac{1}{2}}$

State	0	1	2	3	4	5
$bel(x_2)$	0	0	0	0	$\frac{3}{7}$	$\frac{4}{7}$

Example 2.1: Initial Conditions

At t=0, we are absolutely certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$bel(x_0)$	1	0	0	0	0	0

At t=1: $\mathrm{U}_1 = \mathrm{do\_nothing}$, $\mathrm{Z}_1 = \mathrm{wall}$

State	0	1	2	3	4	5
$bel(x_1)$	0	0	0	0	0	0

Example 2.2: Initial Conditions

At t=0, we are "absolutely" certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$bel(x_0)$	$\frac{19}{20}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$

At t=1: $\mathrm{U}_1 = \mathrm{do\_nothing}$, $\mathrm{Z}_1 = \mathrm{wall}$

State	0	1	2	3	4	5
$bel(x_1)$	0	0	0	0	$\frac{1}{3}$	$\frac{2}{3}$

References

1. Thrun, Sebastian, Wolfram Burgard, and Dieter Fox. Probabilistic robotics. MIT press, 2005
2. http://gki.informatik.uni-freiburg.de/teaching/ws0607/advanced/recordings/BayesFiltering.pdf
3. http://gki.informatik.uni-freiburg.de/teaching/ws0607/advanced/recordings/BayesFiltering.pdf