Lecture 16

Mobile Robotics - I - Probability

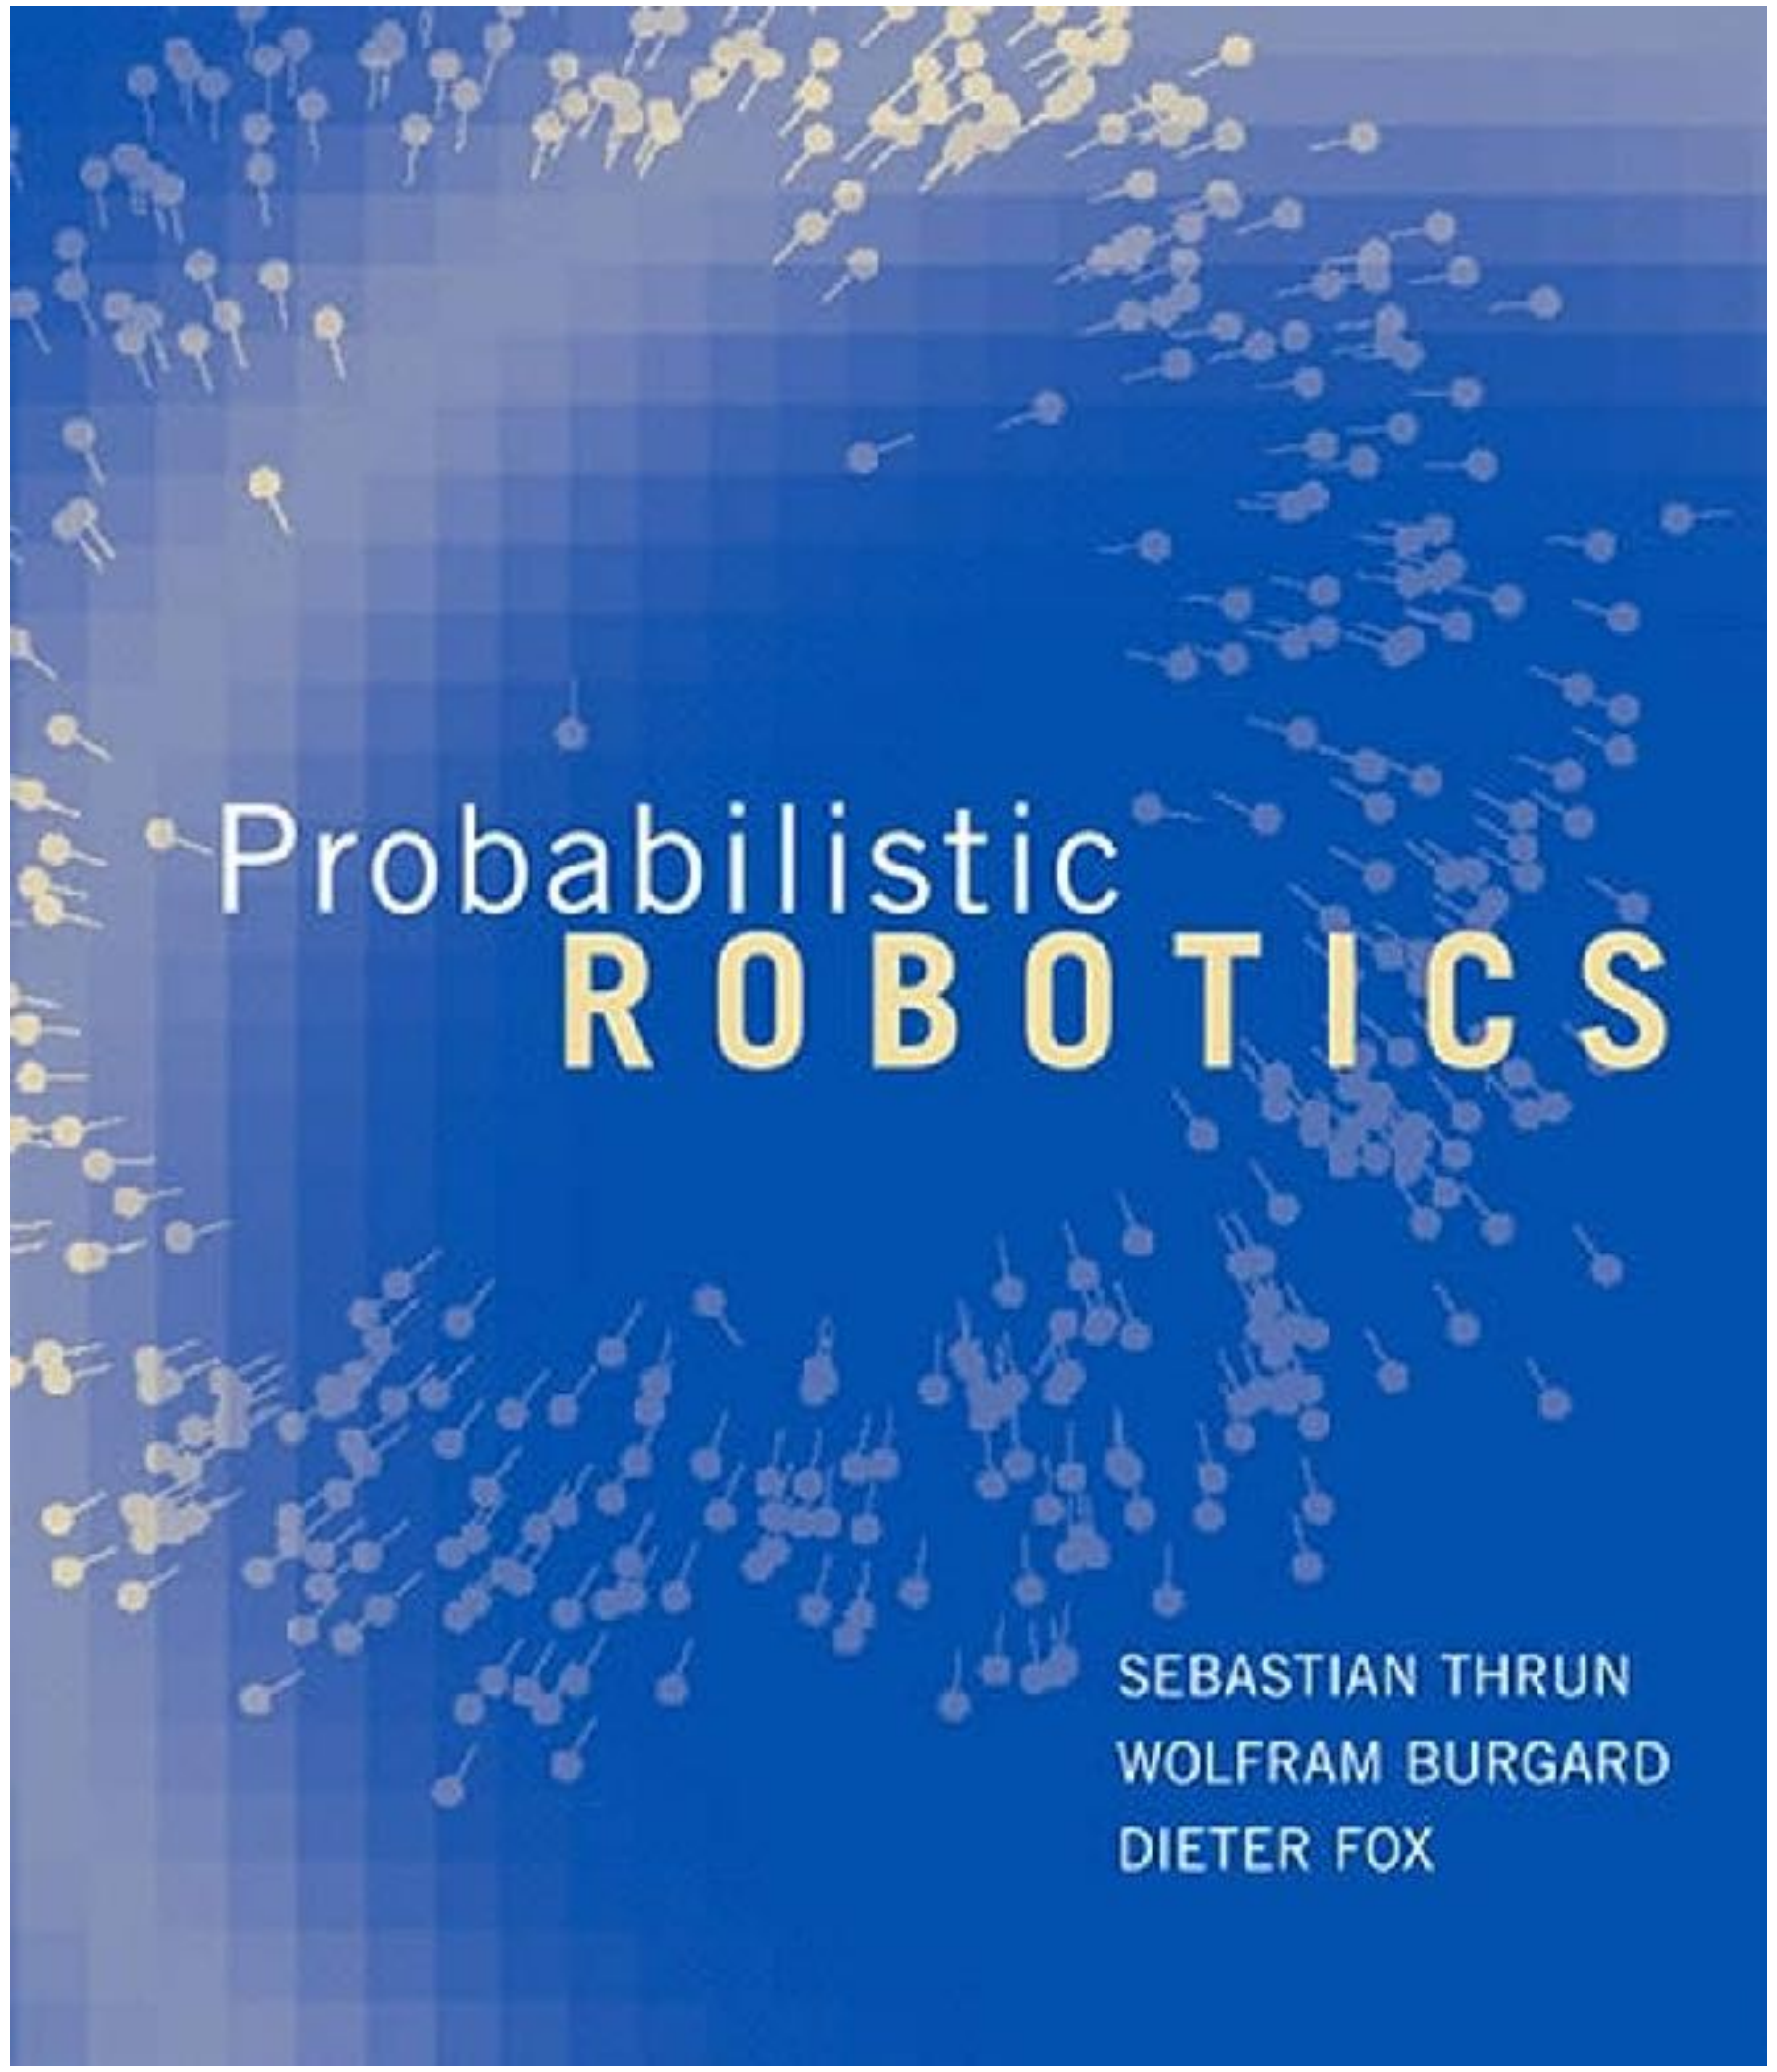

Course logistics

- Quiz 8 will be posted tomorrow noon and will be due on 03/20 noon.
- Project 5 was posted on 02/28 and is due on 03/20.
- Project 6 will be posted on 03/20 and will be on 03/27.
- Group formation for P7 and Final Project by 03/20.
  - How is that going?

Probabilistic Robotics

Key idea: Explicit representation of uncertainty

(using the calculus of probability theory)

- Perception = state estimation
- Action = utility optimization

Discrete Random Variables

- X denotes a random variable.
- X can take on a countable number of values in $\{x_1, x_2, \ldots, x_n\}$.
- $P(X = x_i)$, or $P(x_i)$, is the probability that the random variable X takes on value x_i.
- $P(.)$ is called probability mass function.
- E.g. $P(\text{room}) = < 0.7, 0.2, 0.08, 0.02 >$

Joint and Conditional Probability

- $P(X = x \text{ and } Y = y) = P(x, y)$

- $P(x \mid y)$ is the probability of x given y

$$P(x \mid y) = \frac{P(x, y)}{P(y)}$$

$$P(x, y) = P(x \mid y) P(y)$$

- If X and Y are independent then

$$P(x, y) = P(x) P(y)$$

- If X and Y are independent then

$$P(x \mid y) = P(x)$$

Law of Total Probability, Marginals

Discrete Case

$$\sum_x P(x) = 1$$

$$P(x) = \sum_y P(x, y)$$

$$P(x) = \sum_y P(x \mid y) P(y)$$

Continuous Case

$$\int p(x)dx = 1$$

$$p(x) = \int p(x, y)dy$$

$$p(x) = \int p(x \mid y) p(y)dy$$

Events

- P(+x, +y) ?
- P(+x) ?
- P(-y OR +x) ?
- Independent?

$$P(X,Y)$$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

Marginal Distributions

$$P(X, Y)$$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

$$P(x) = \sum_y P(x, y)$$

$$P(X)$$

X	P
+x	
-x	

$$P(y) = \sum_x P(x, y)$$

$$P(Y)$$

Y	P
+y	
-y	

Conditional Probabilities

$$P(X, Y)$$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

- P(+x | +y) ?
- P(-x | +y) ?
- P(-y | +x) ?

Bayes Formula

$$P(x,y) = P(x \mid y)P(y) = P(y \mid x)P(x)$$

$$P(x \mid y) = \frac{P(y \mid x)P(x)}{P(y)} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}$$

- Often causal knowledge is easier to obtain than diagnostic knowledge.
- Bayes rule allows us to use causal knowledge.

Simple Example of State Estimation

- Suppose a robot obtains measurement z
- What is $P(open \mid z)$?

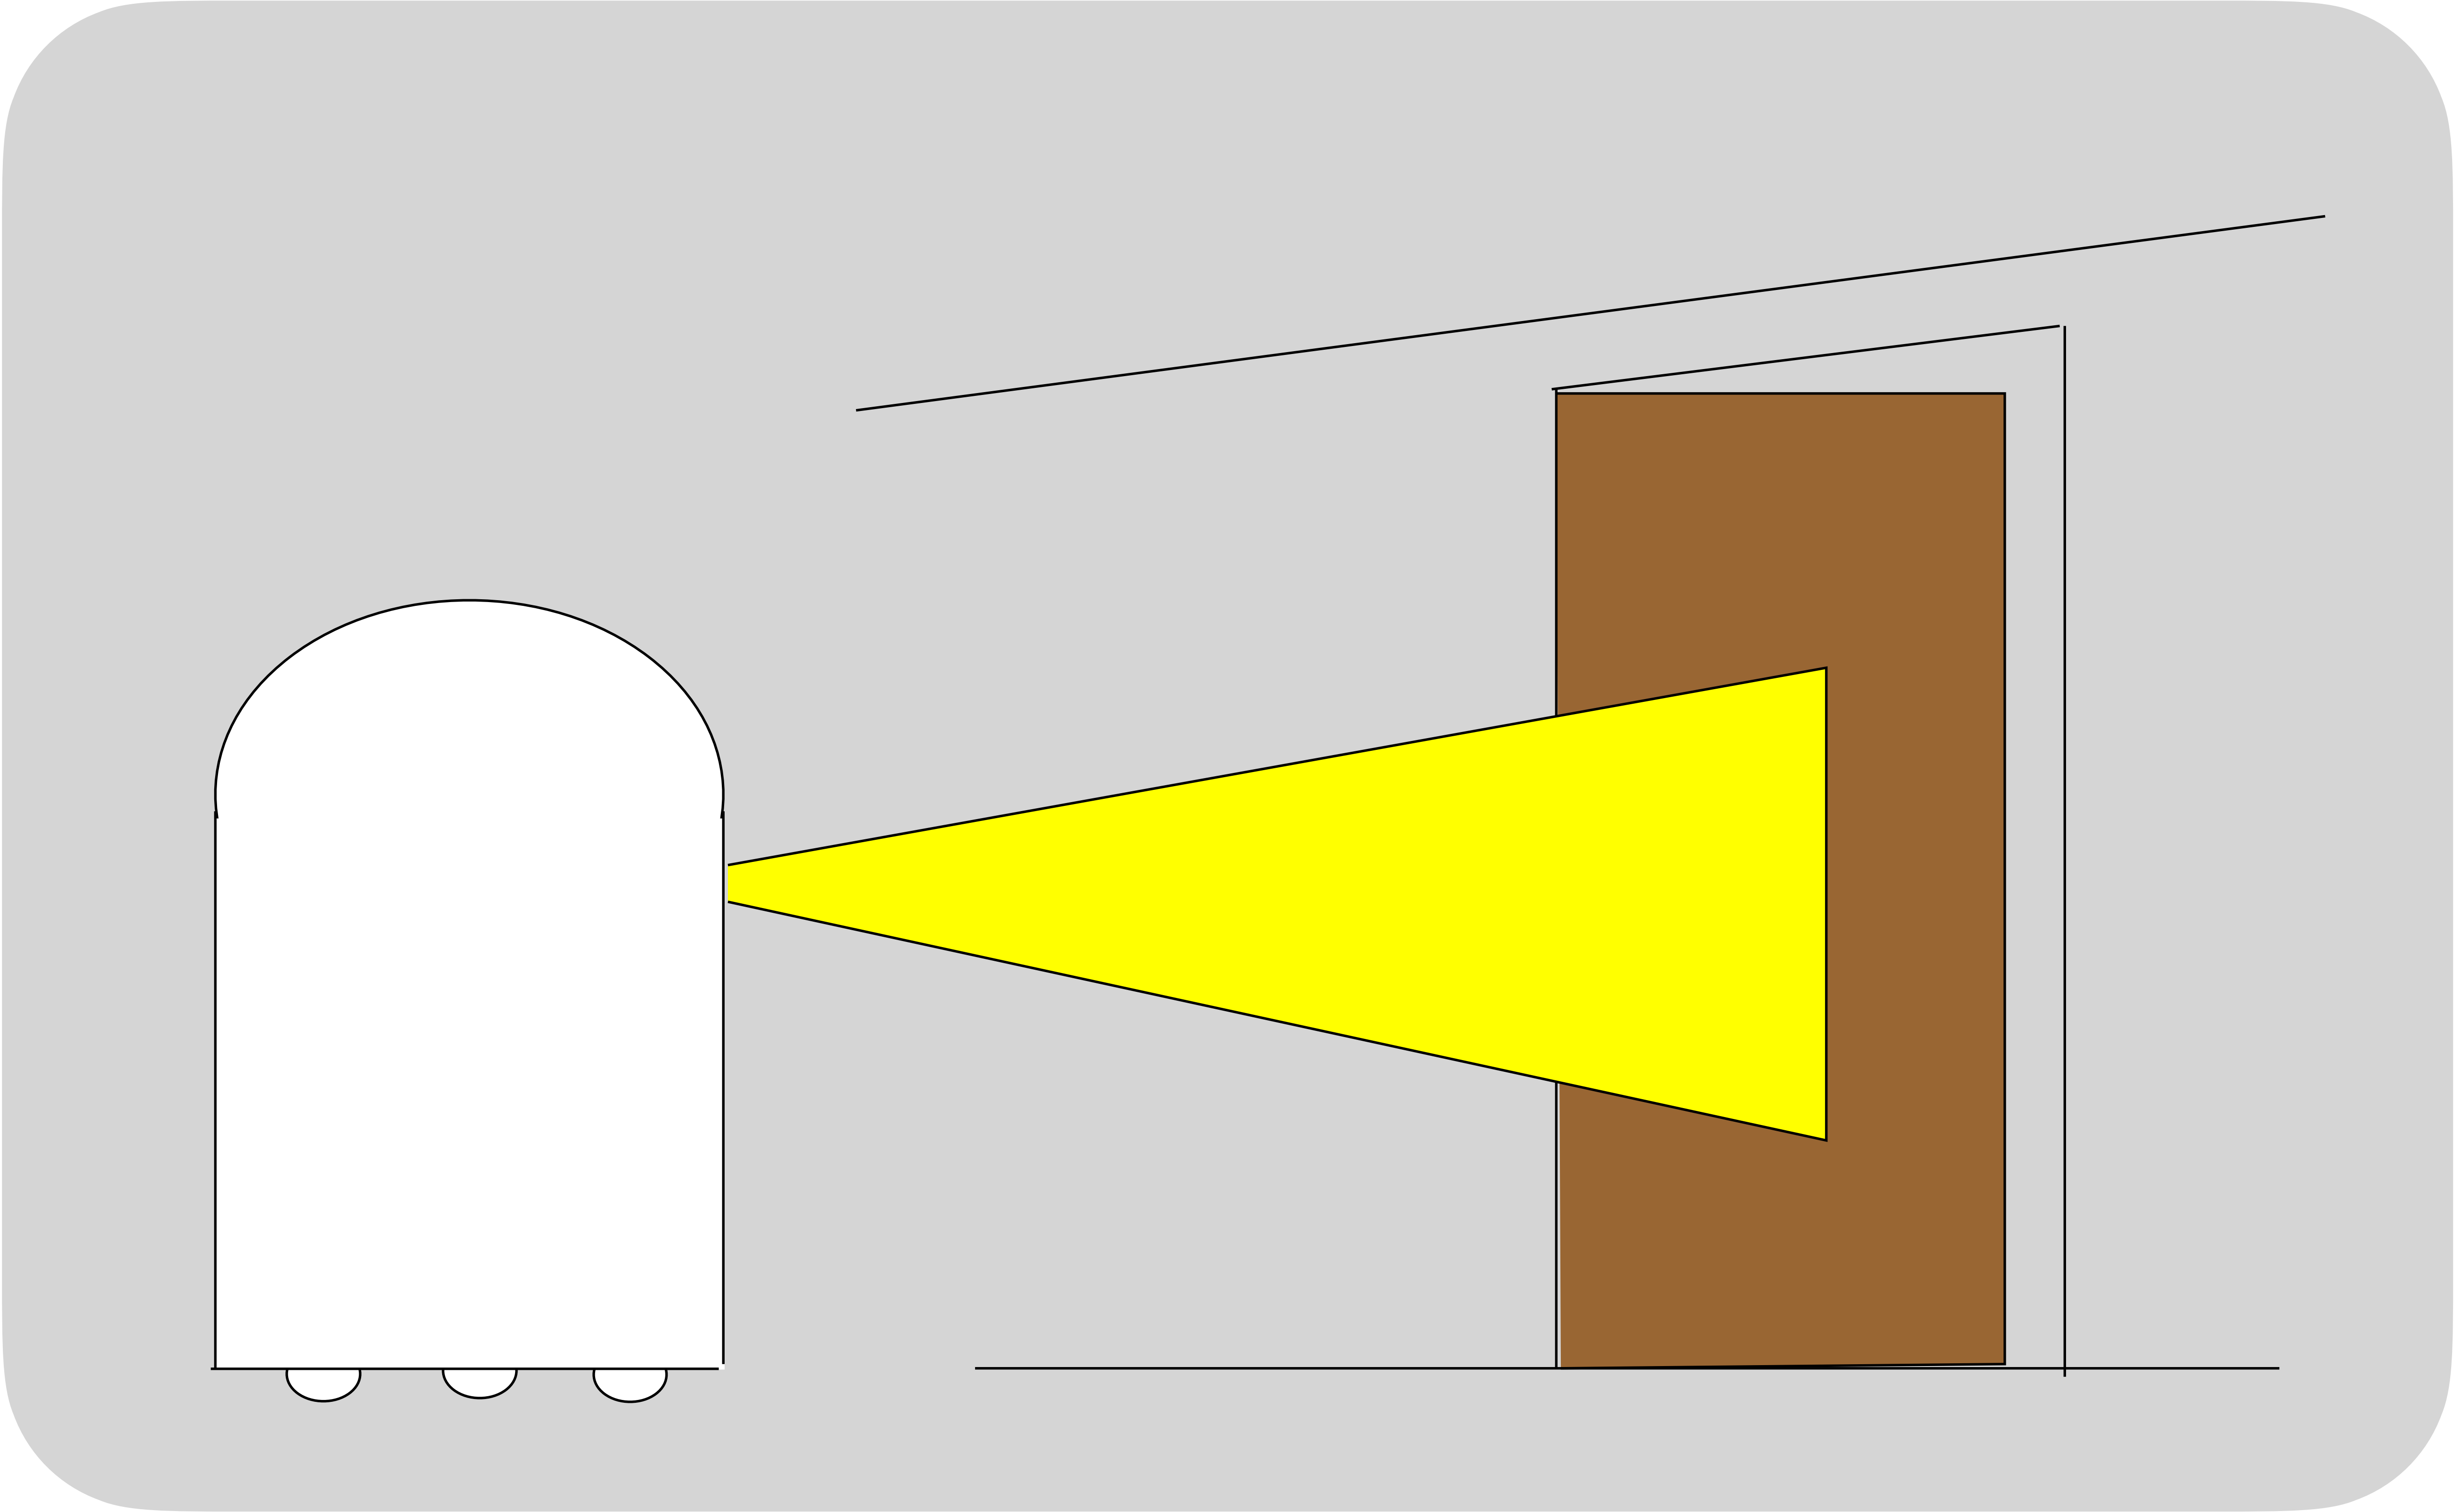

Example

$$P(z \mid \text{open}) = 0.6 \qquad P(z \mid \neg\text{open}) = 0.3$$

$$P(\text{open}) = 0.5 \qquad P(\neg\text{open}) = 0.5$$

$$P(\text{open} \mid z) =$$

$$P(\text{open} \mid z) = \frac{0.6 \times 0.5}{0.6 \times 0.5 + 0.3 \times 0.5} = \frac{2}{3} = 0.67$$

- *z* raises the probability that the door is open.

Normalization

$$P(x \mid y) = \frac{P(y \mid x)P(x)}{P(y)} = \eta P(y \mid x)P(x)$$

$$\eta = P(y)^{-1} = \frac{1}{\sum_{x'} P(y \mid x')P(x')}$$

Conditioning

- Bayes rule and background knowledge:

$$P(x \mid y, z) = \frac{P(y \mid x, z) P(x \mid z)}{P(y \mid z)}$$

$$P(x \mid y) \stackrel{?}{=} \int P(x \mid y, z) P(z) dz$$

$$\stackrel{?}{=} \int P(x \mid y, z) P(z \mid y) dz$$

$$\stackrel{?}{=} \int P(x \mid y, z) P(y \mid z) dz$$

Conditioning

- Bayes rule and background knowledge:

$$P(x \mid y, z) = \frac{P(y \mid x, z)P(x \mid z)}{P(y \mid z)}$$

$$P(x \mid y) = \int P(x \mid y, z)P(z \mid y)dz$$

Conditional Independence

$$P(x, y \mid z) = P(x \mid z) P(y \mid z)$$

- Equivalent to

$$P(x \mid z) = P(x \mid z, y)$$

and

$$P(y \mid z) = P(y \mid z, x)$$

Simple Example of State Estimation

- Suppose our robot obtains another observation z_2.
- What is $P(open \mid z_1, z_2)$?

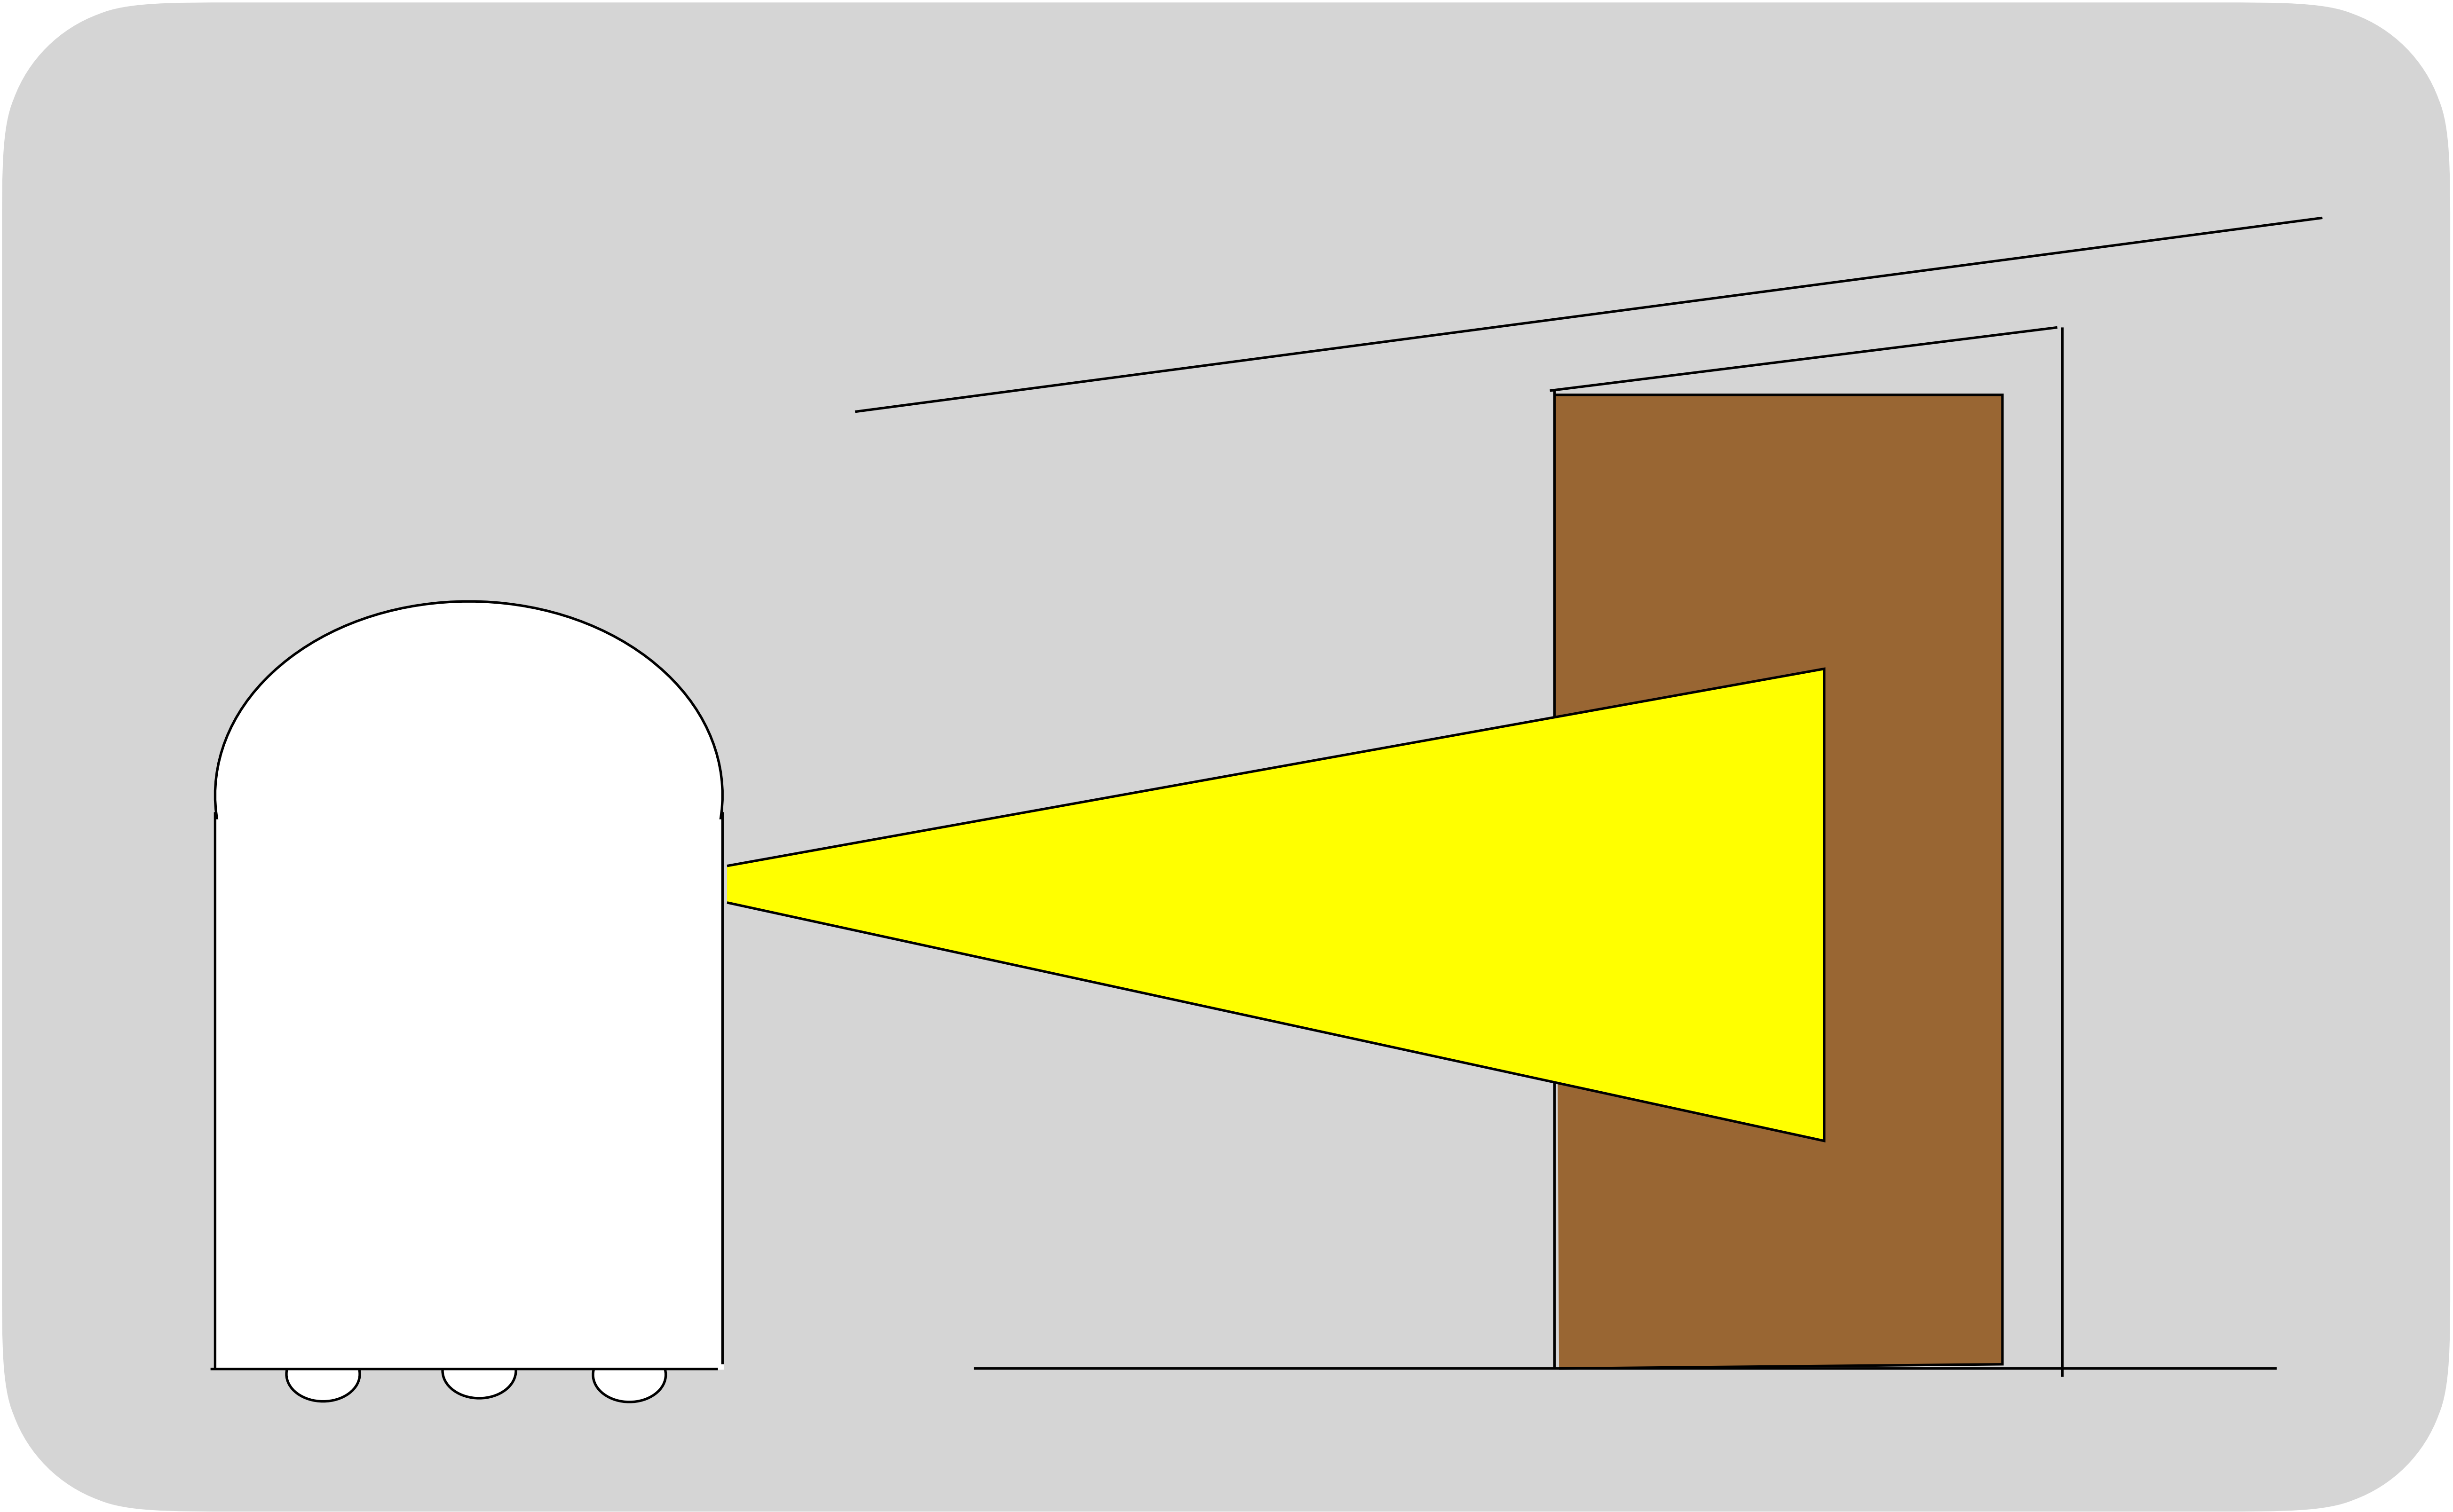

Recursive Bayesian Updating

$$P(x \mid z_1, \ldots z_n) = \frac{P(z_n \mid x, z_1, \ldots z_{n-1}) P(x \mid z_1, \ldots z_{n-1})}{P(z_n \mid z_1, \ldots z_{n-1})}$$

Markov assumption: z_n is conditionally independent of $z_1, \ldots, z_{n-1}$ given x.

$$P(x \mid z_1, \ldots z_n) = \frac{P(z_n \mid x) P(x \mid z_1, \ldots z_{n-1})}{P(z_n \mid z_1, \ldots z_{n-1})}$$

$$= \eta P(z_n \mid x) P(x \mid z_1, \ldots z_{n-1})$$

$$= \eta_{1..n} \prod_{i=1 \ldots n} P(z_i \mid x) P(x)$$

Example: Second Measurement

$$P(z_2 \mid \text{open}) = 0.5 \quad P(z_2 \mid \neg\text{open}) = 0.6$$

$$P(\text{open} \mid z_1) = 2/3 \quad P(\neg\text{open} \mid z_1) = 1/3$$

$$P(\text{open} \mid z_2, z_1) = \cdot$$

$$= \frac{1/2 \times 2/3}{1/2 \times 2/3 + 3/5 \times 1/3} = \frac{5}{8} = 0.625$$

- z_2 lowers the probability that the door is open.

Bayes Filters: Framework

- **Given:**
  - Stream of observations z and action data u:
    $$d_t = \{u_1, z_2, \ldots. u_{t-1}, z_t\}$$
  - Sensor model $P(z \mid x)$.
  - Action model $P(x \mid u, x')$.
  - Prior probability of the system state $P(x)$.
- **Wanted:**
  - Estimate of the state X of a dynamical system.
  - The posterior of the state is also called **Belief**:

$$Bel(x_t) = P(x_t \mid u_1, z_2, \ldots. u_{t-1}, z_t)$$

Bayes Filters for Robot Localization

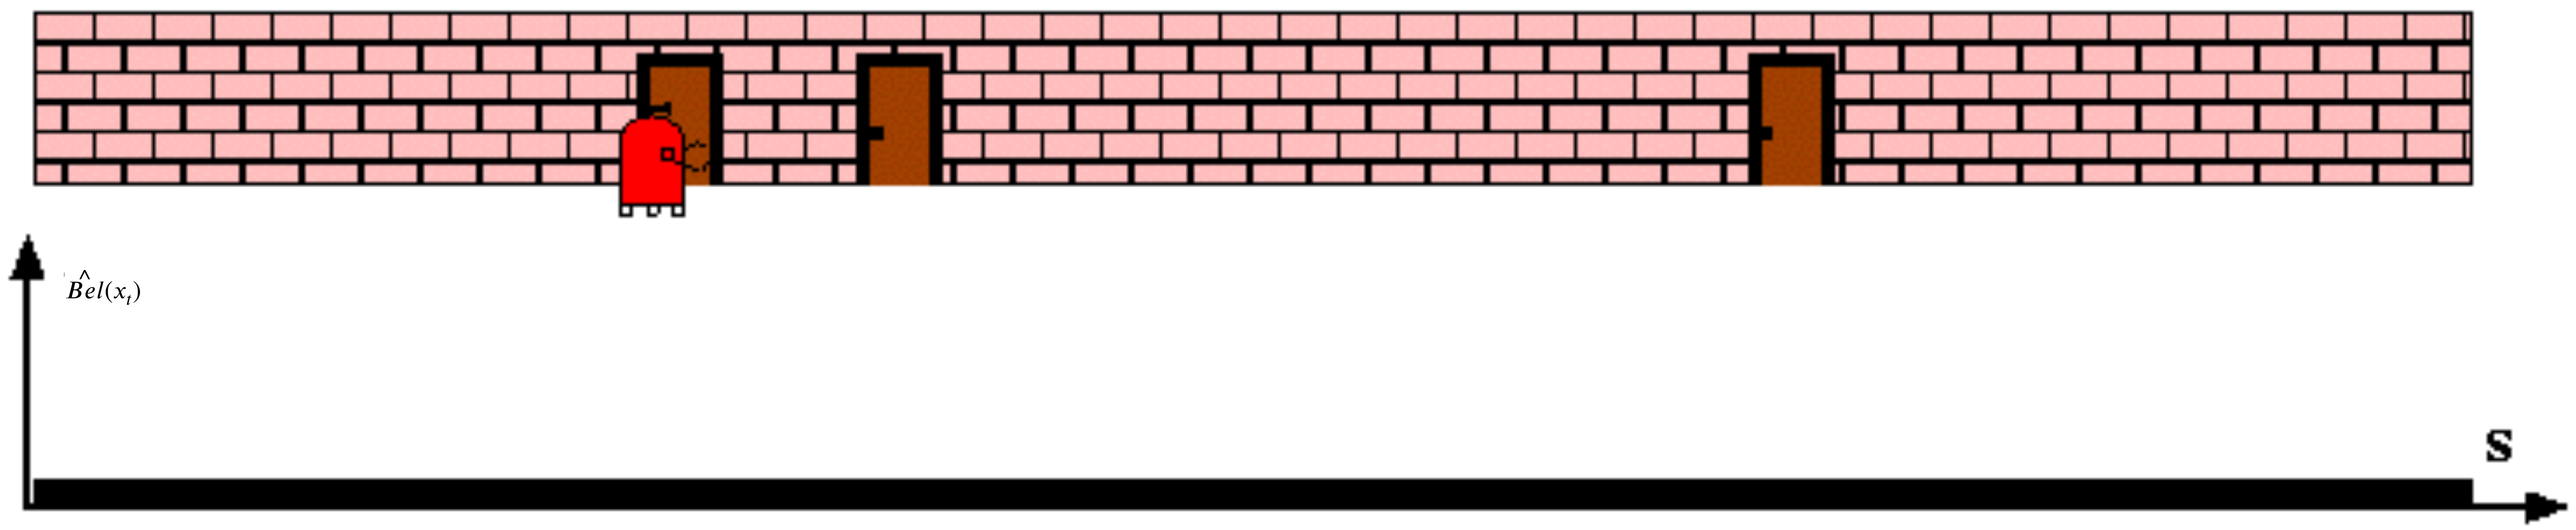

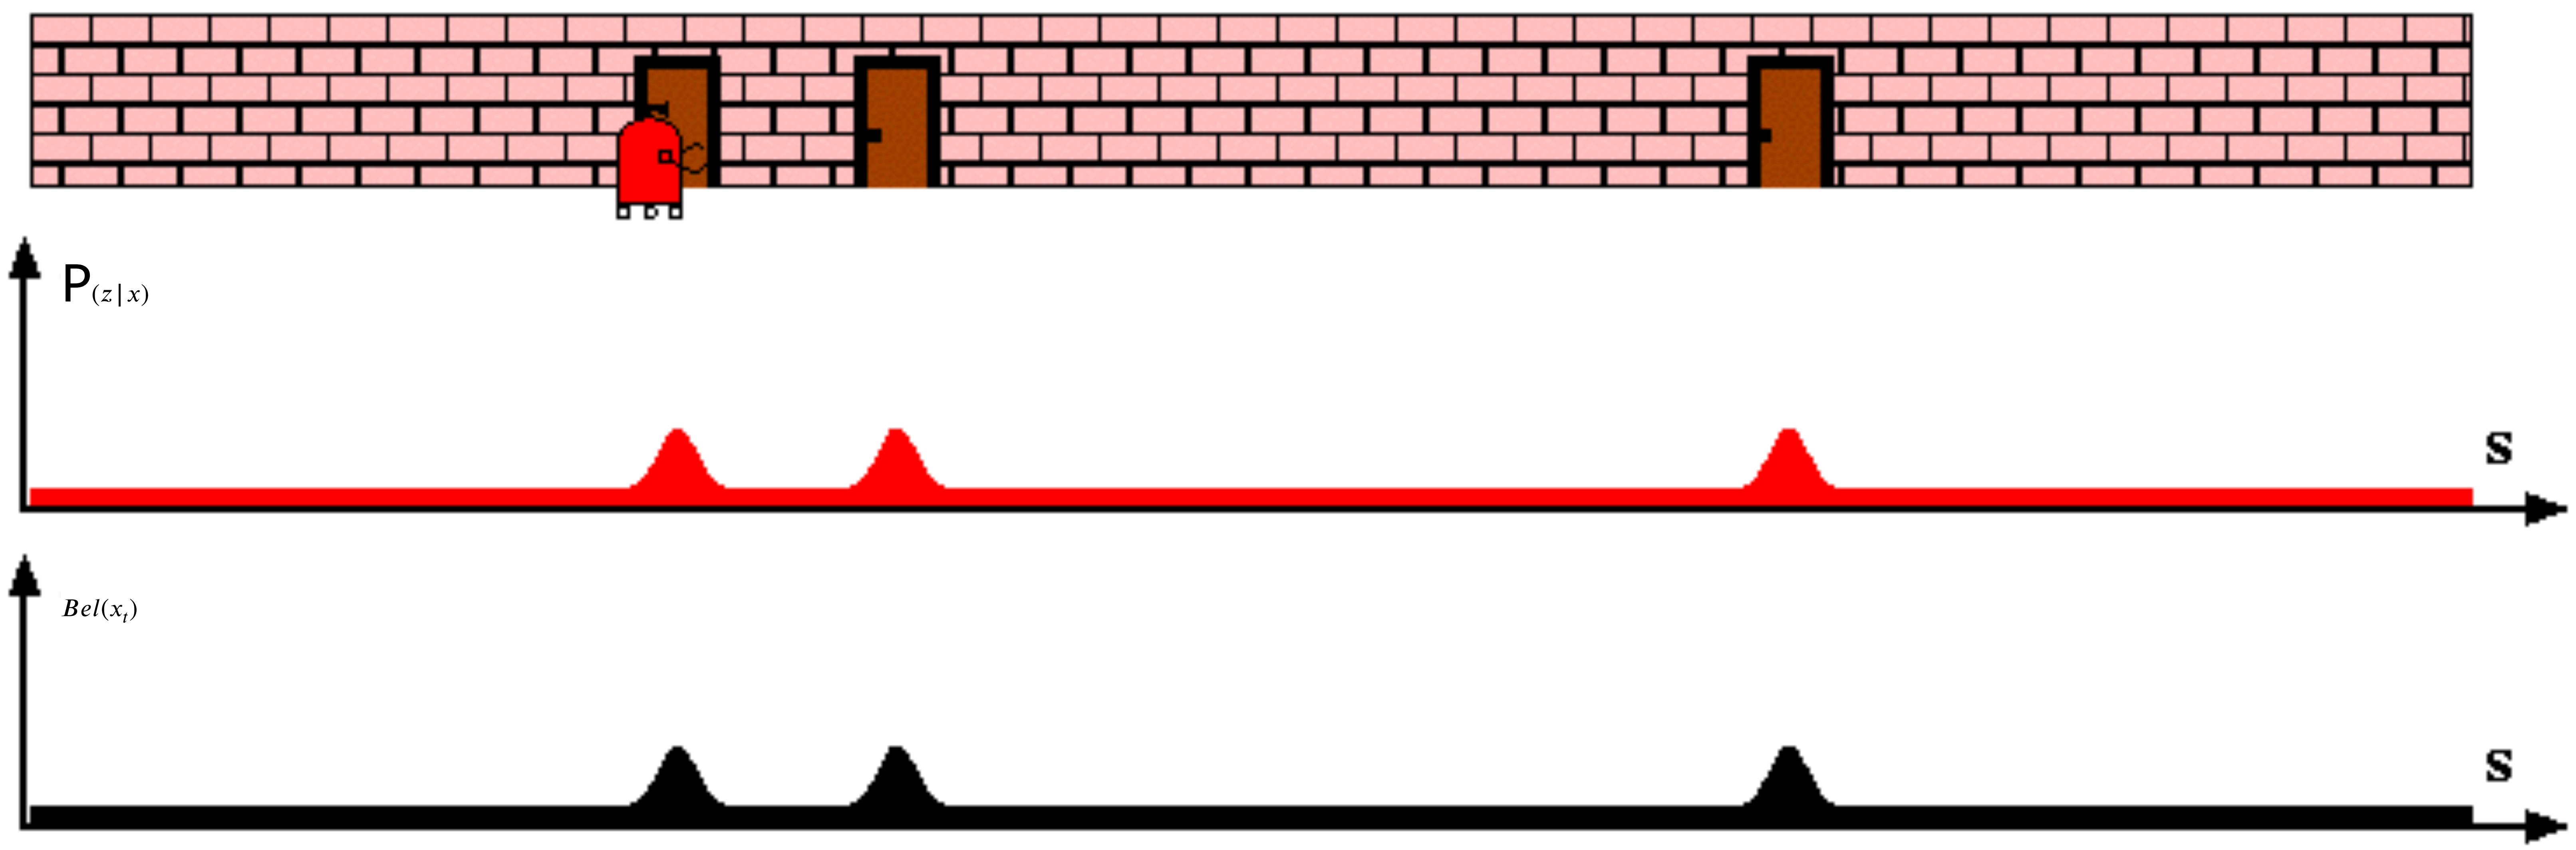

Slide borrowed from Dieter Fox

Bayes Filters for Robot Localization

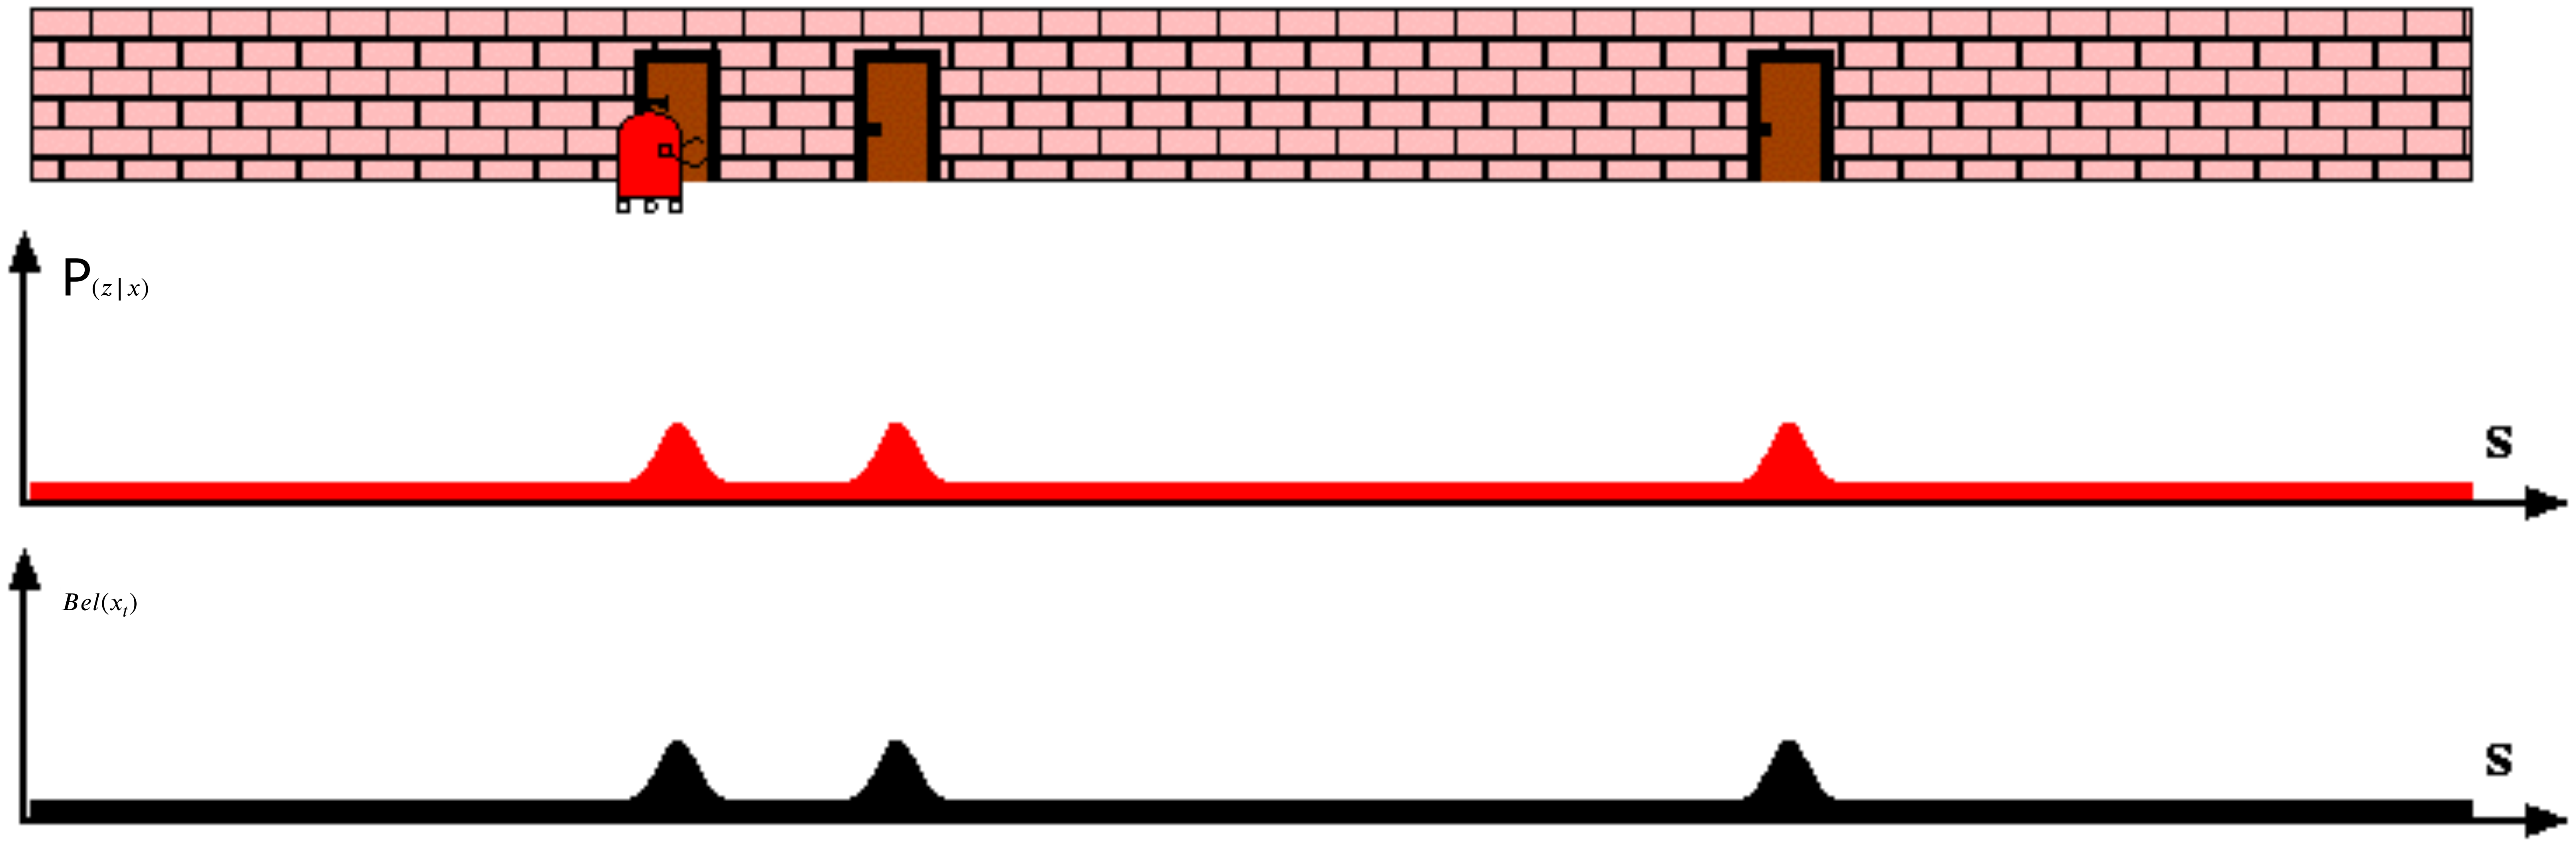

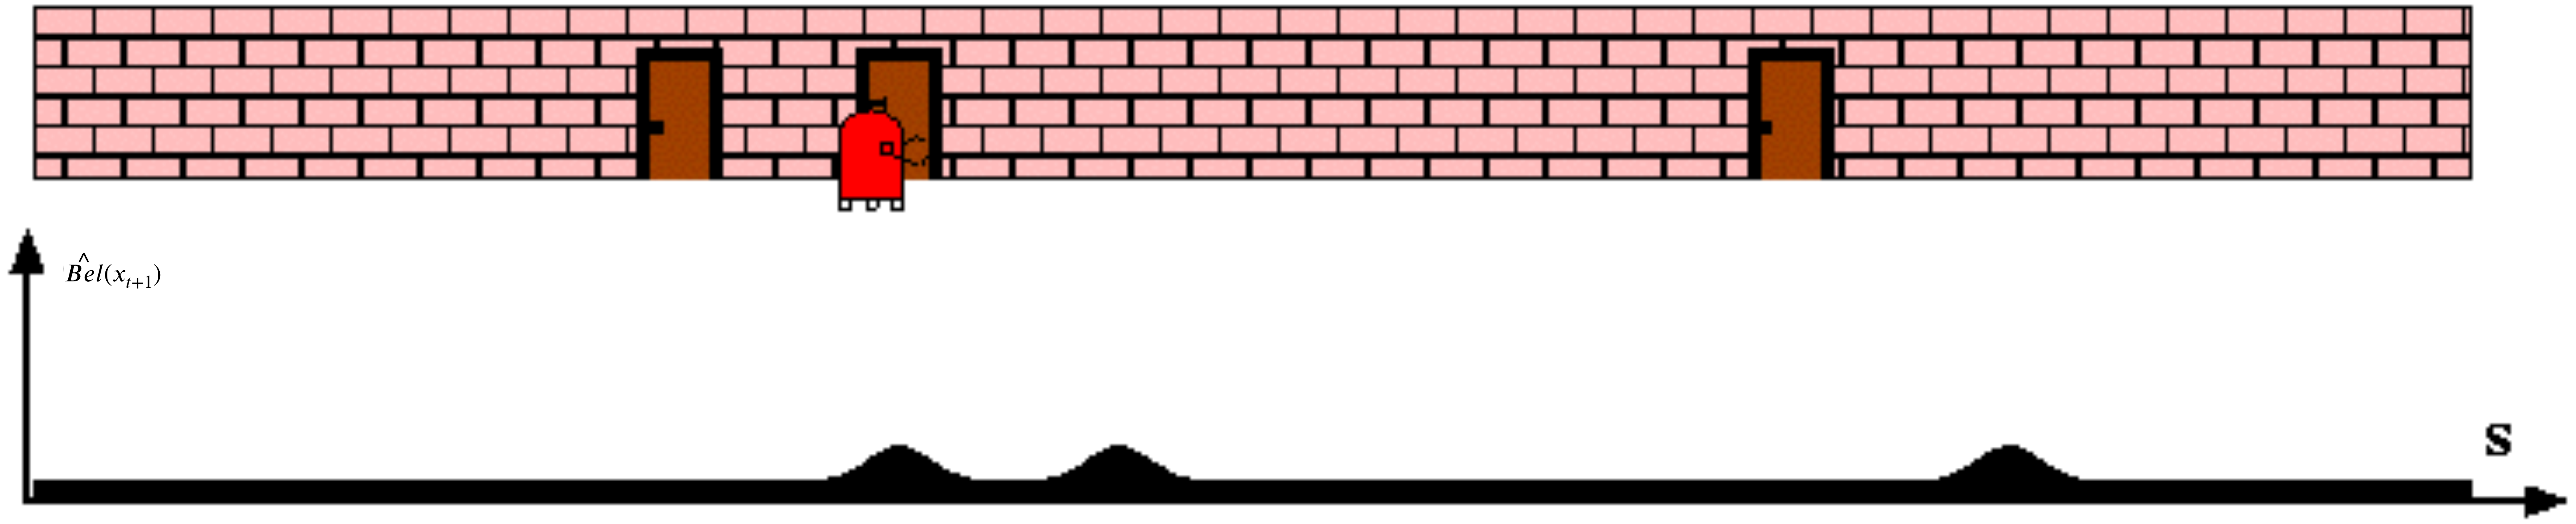

Slide borrowed from Dieter Fox

Bayes Filters for Robot Localization

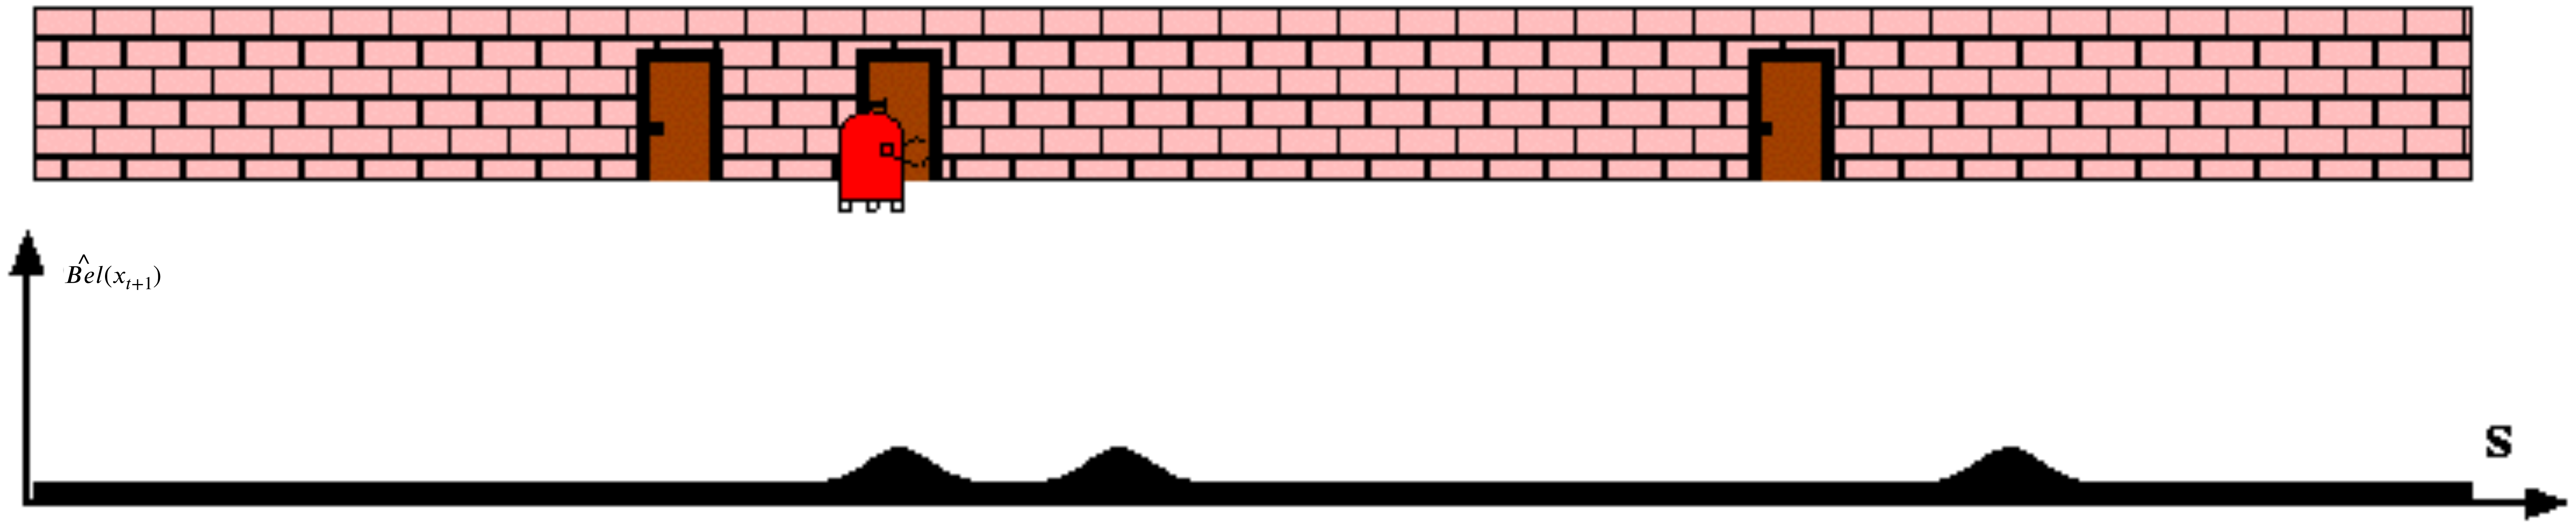

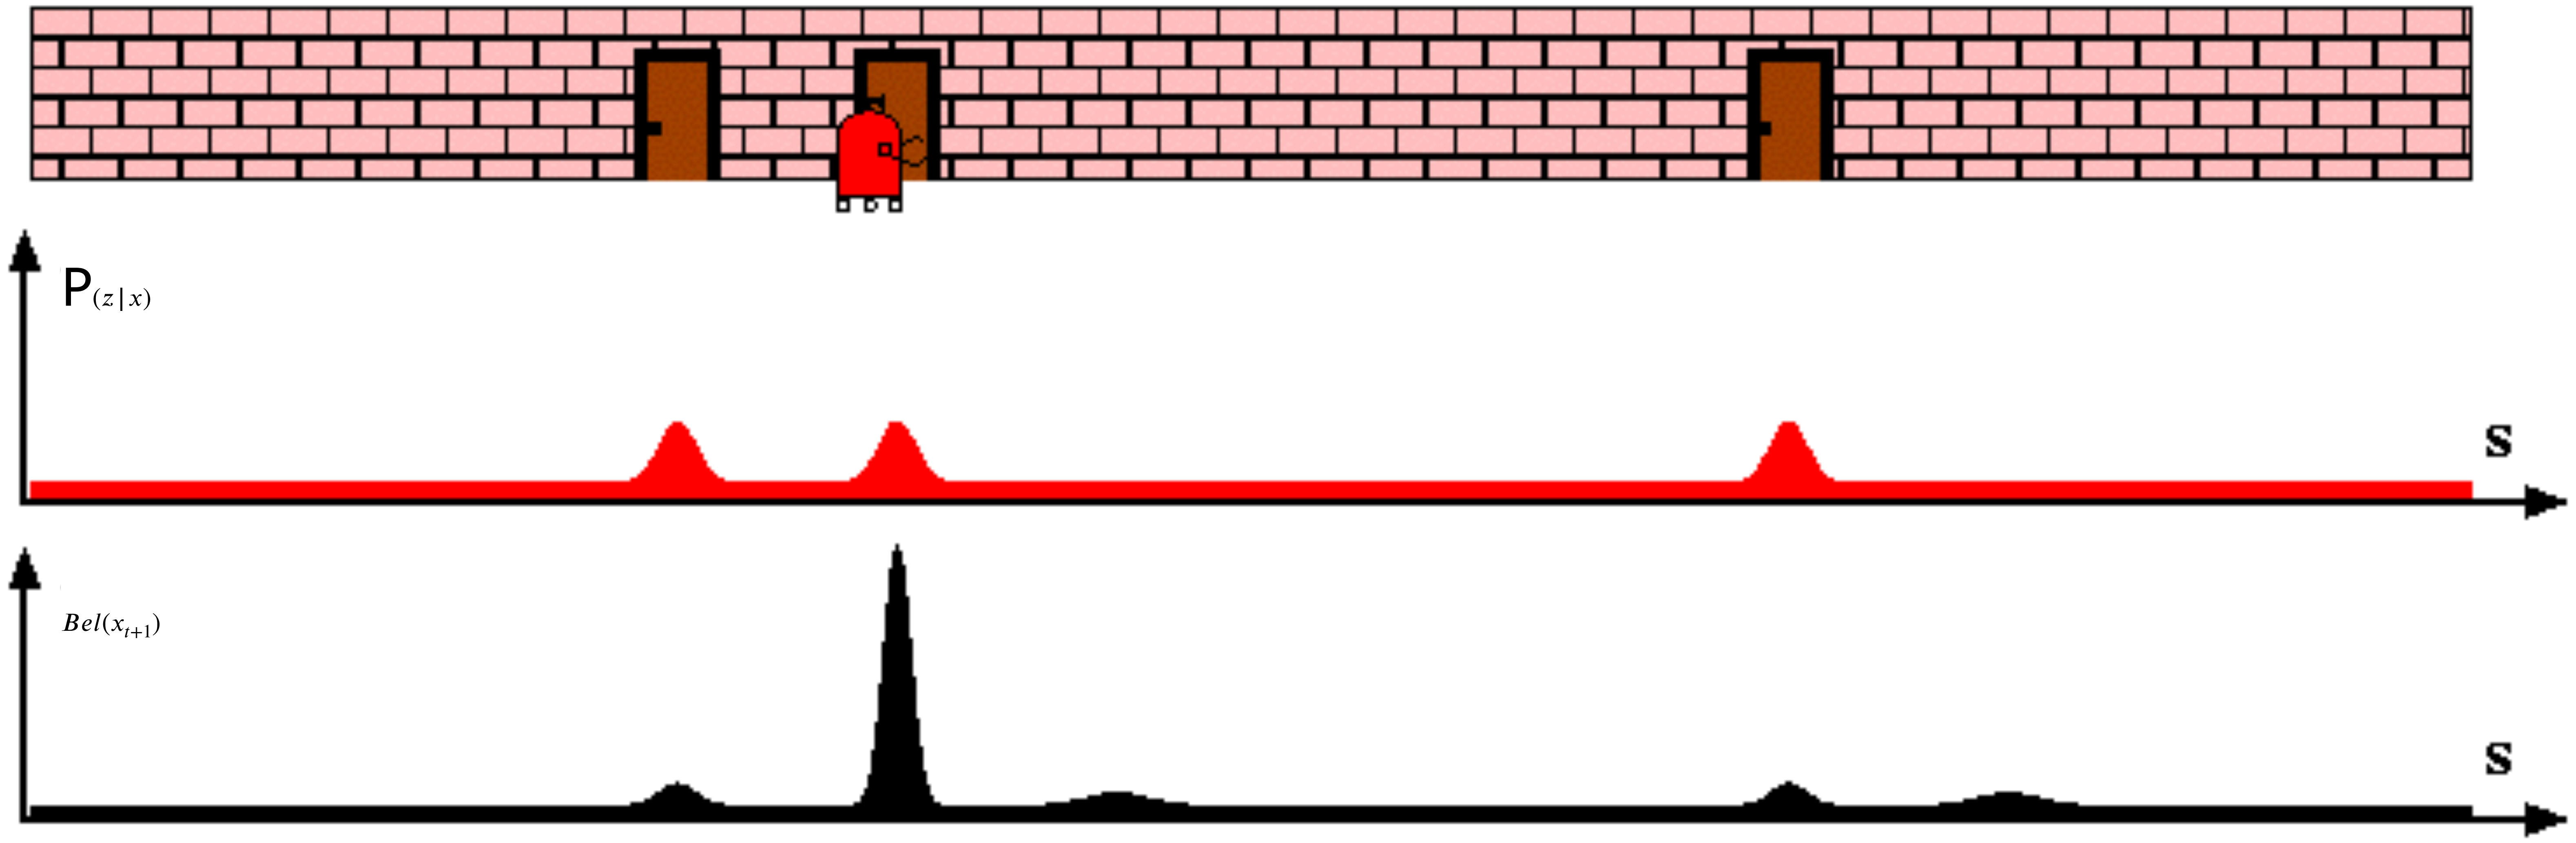

Slide borrowed from Dieter Fox

Bayes Filters for Robot Localization

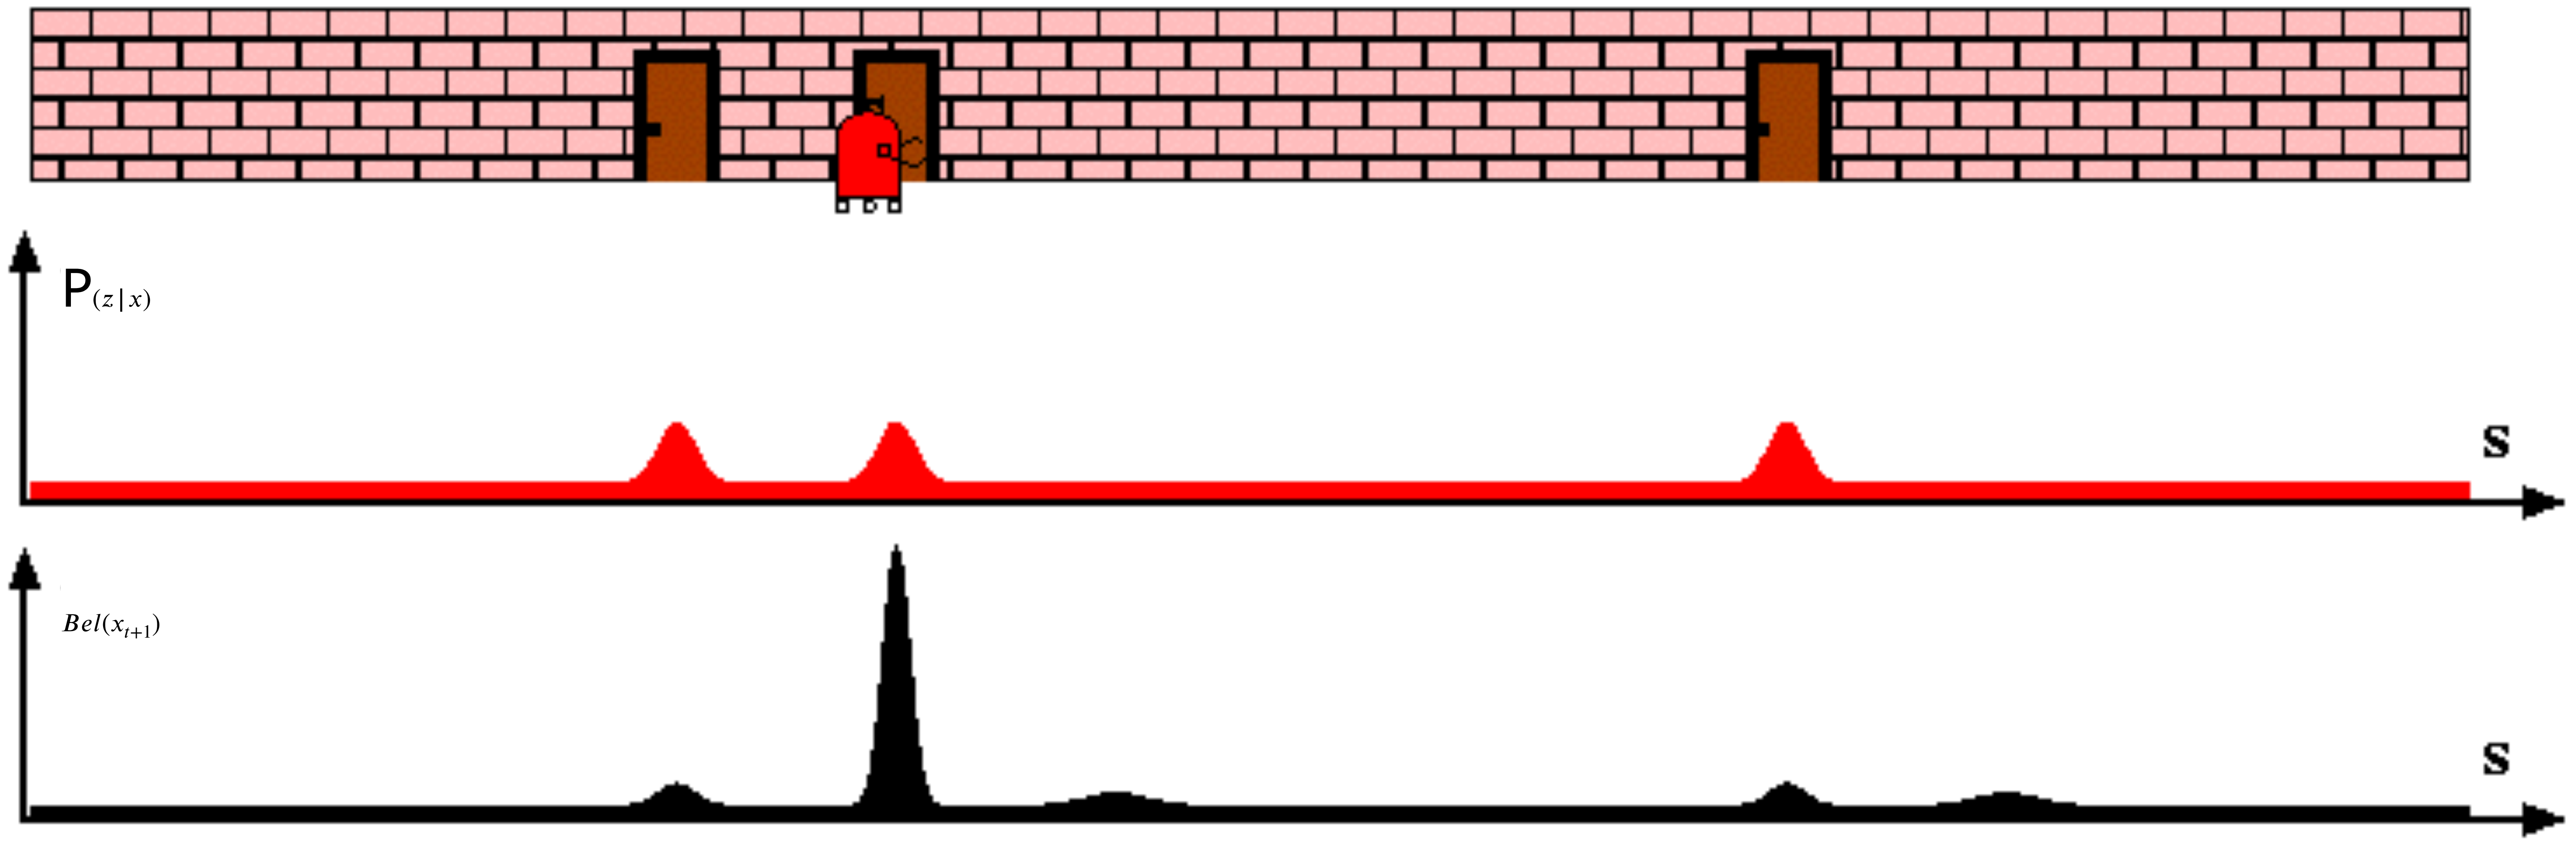

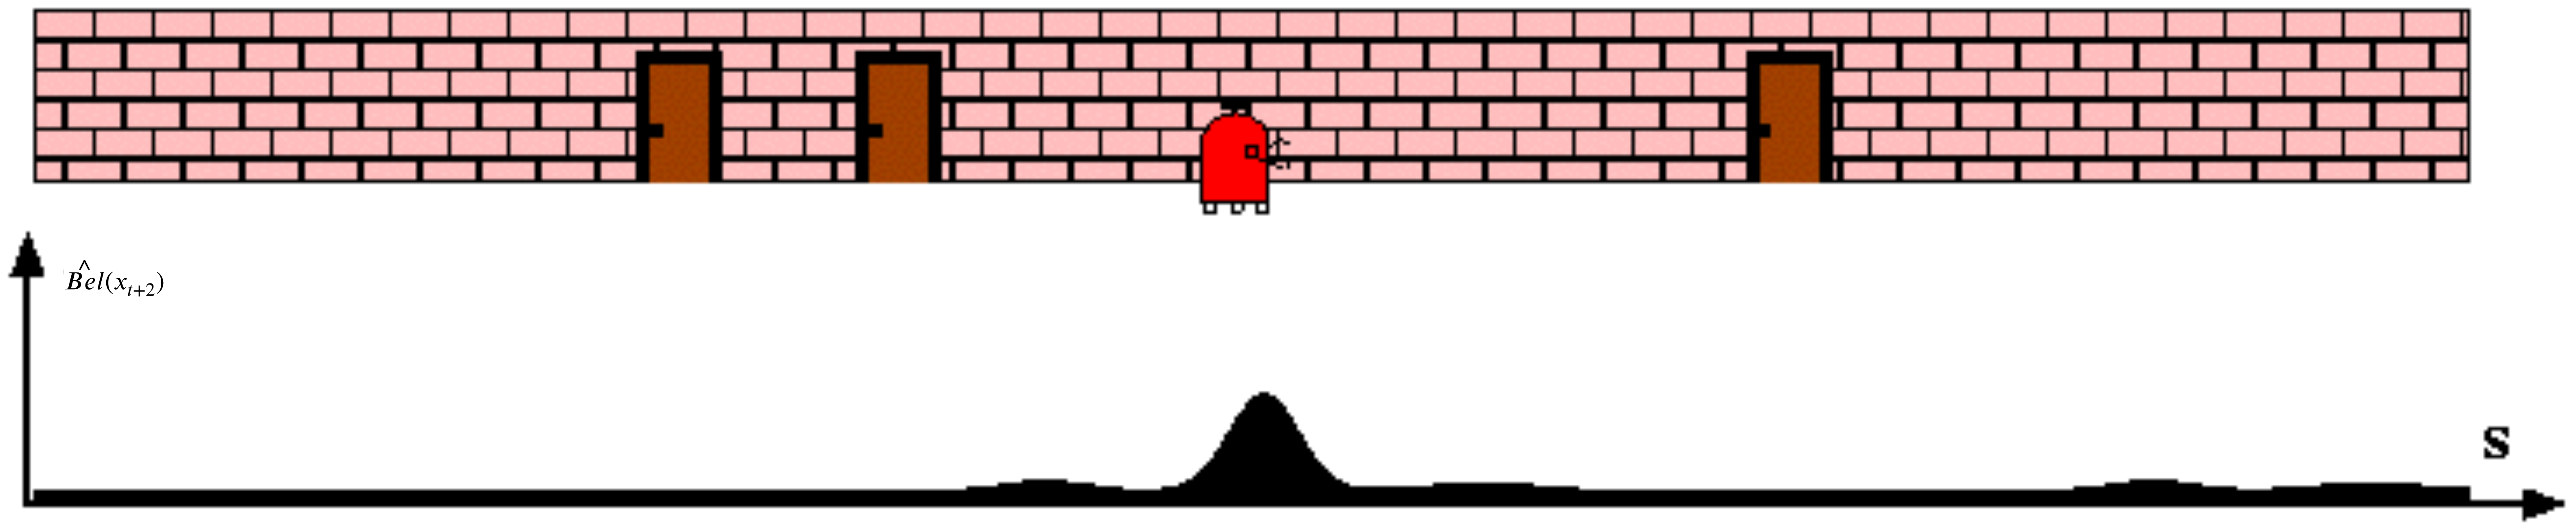

Slide borrowed from Dieter Fox

Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t \mid u_1, z_1 \ldots, u_t, z_t)$$

Bayes
$$= \eta\ P(z_t \mid x_t, u_1, z_1, \ldots, u_t)\ P(x_t \mid u_1, z_1, \ldots, u_t)$$

Markov
$$= \eta\ P(z_t \mid x_t)\ P(x_t \mid u_1, z_1, \ldots, u_t)$$

Total prob.
$$= \eta\ P(z_t \mid x_t) \int P(x_t \mid u_1, z_1, \ldots, u_t, x_{t-1}) P(x_{t-1} \mid u_1, z_1, \ldots, u_t)\ dx_{t-1}$$

Markov
$$= \eta\ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1})\ P(x_{t-1} \mid u_1, z_1, \ldots, u_t)\ dx_{t-1}$$

$$= \eta\ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1})\ Bel(x_{t-1})\ dx_{t-1}$$

$$Bel(x_t) = \eta \ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \ Bel(x_{t-1}) \ dx_{t-1}$$

```
Algorithm Bayes_filter( Bel(x),d ):
n=0
If d is a perceptual data item z then
    For all x do
        Bel'(x) = P(z | x) Bel(x)
        η = η + Bel'(x)
    For all x do
        Bel'(x) = η^{-1} Bel'(x)
Else if d is an action data item u then
    For all x do
        Bel'(x) = ∫ P(x | u, x') Bel(x') dx'
Return Bel'(x)
```

Markov Assumption

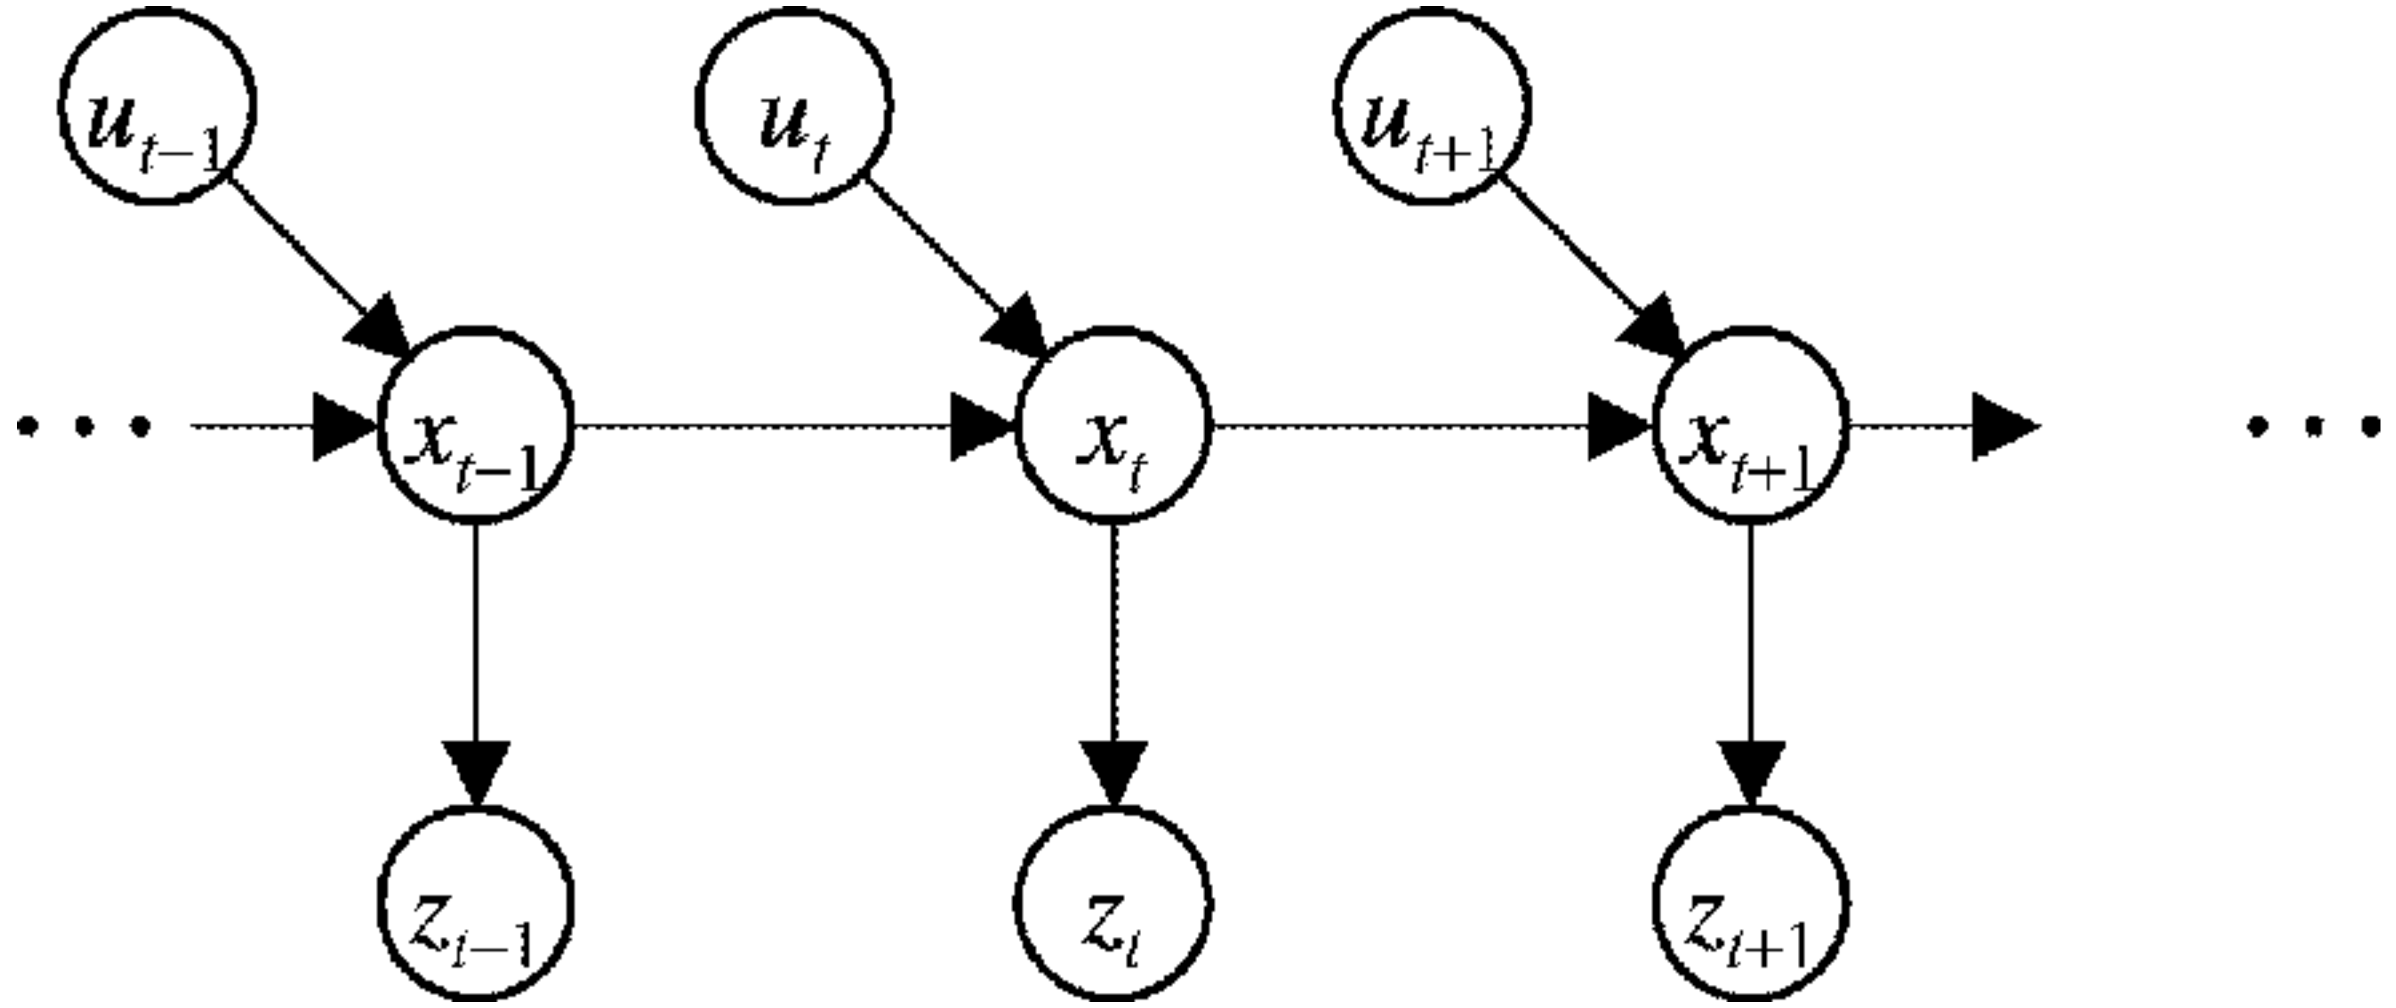

$$P(z_t \mid x_{0:t}, z_{1:t-1}, u_{1:t}) = p(z_t \mid x_t)$$

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

Underlying Assumptions

- Static world
- Independent noise
- Perfect model, no approximation errors

Bayes Filters are Familiar!

$$Bel(x_t) = \eta P(z_t \mid x_t) \int P(x_t \mid x_{t-1} u_t) Bel(x_{t-1}) dx_{t-1}$$

- Kalman filters
- Particle filters
- Hidden Markov models
- Dynamic Bayesian networks
- Partially Observable Markov Decision Processes (POMDPs)

Summary

- Bayes rule allows us to compute probabilities that are hard to assess otherwise.
- Under the Markov assumption, recursive Bayesian updating can be used to efficiently combine evidence.
- Bayes filters are a probabilistic tool for estimating the state of dynamic systems.

Next Lecture

Mobile Robotics - II - Motion & Sensor Models