CSE-571
Robotics

Probabilistic Robotics

Probabilities

Bayes rule

Bayes filters

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.
- *$P(\cdot)$* is called probability mass function.
- E.g. $P(Room) = \langle 0.7, 0.2, 0.08, 0.02 \rangle$

Joint and Conditional Probability

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

- If X and Y are independent then

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

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

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

- 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

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

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

$P(X)$

X	P
+x	
-x	

$P(Y)$

Y	P
+y	
-y	

Conditional Probabilities

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

$P(X, Y)$

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

Bayes Formula

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

$$\Rightarrow$$

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

- 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|z)?*

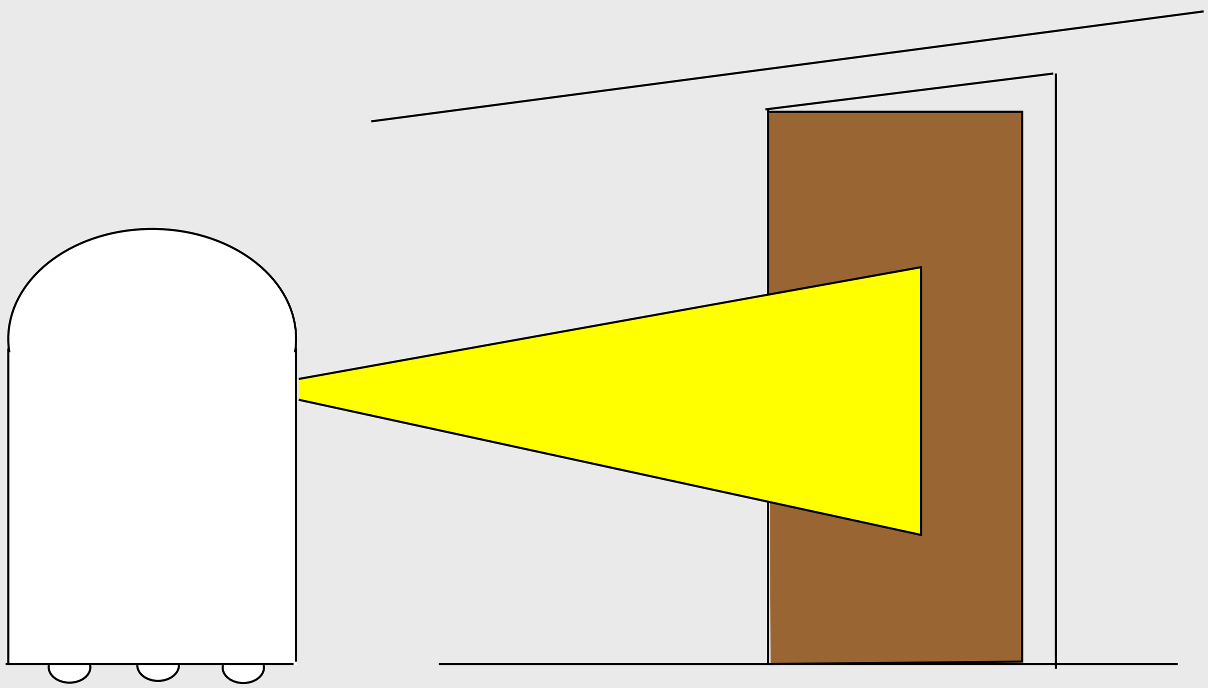

Example

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

$$P(open) = P(\neg open) = 0.5$$

$$P(open \mid z) = \frac{P(z \mid open) P(open)}{P(z \mid open) p(open) + P(z \mid \neg open) p(\neg open)}$$

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

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

Normalization

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

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

Algorithm:

$$\forall x : \text{aux}_{x|y} = P(y \mid x)\; P(x)$$

$$\eta = \frac{1}{\sum_{x} \text{aux}_{x|y}}$$

$$\forall x : P(x \mid y) = \eta\; \text{aux}_{x|y}$$

Conditioning

- Bayes rule and background knowledge:

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

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

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

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

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 $P(open|z_1, z_2)$?

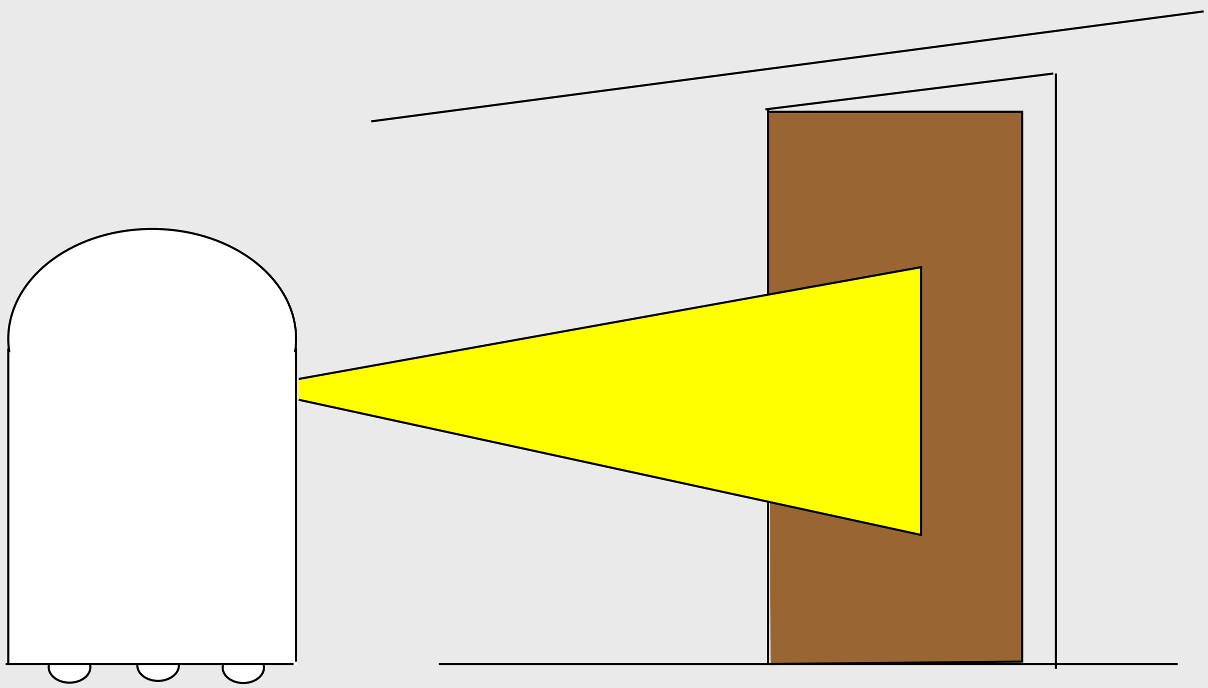

Recursive Bayesian Updating

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

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

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

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

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

Example: Second Measurement

$$P(z_2 \mid open) = 0.5 \qquad P(z_2 \mid \neg open) = 0.6$$

$$P(open \mid z_1) = 2/3 \qquad P(\neg open \mid z_1) = 1/3$$

$$P(open \mid z_2, z_1) = \frac{P(z_2 \mid open)\, P(open \mid z_1)}{P(z_2 \mid open)\, P(open \mid z_1) + P(z_2 \mid \neg open)\, P(\neg open \mid z_1)}$$

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

- 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 $P(z|x)$.
  - Action model $P(x|u,x')$.
  - 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

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

        $Bel'(x) = P(z \mid x) Bel(x)$

        $\eta = \eta + Bel'(x)$

    For all *x* do

        $Bel'(x) = \eta^{-1} Bel'(x)$

Else if *d* is an action data item *u* then

    For all *x* do

        $Bel'(x) = \int P(x \mid u, x') \ Bel(x') \ dx'$

Return *Bel'(x)*

Markov Assumption

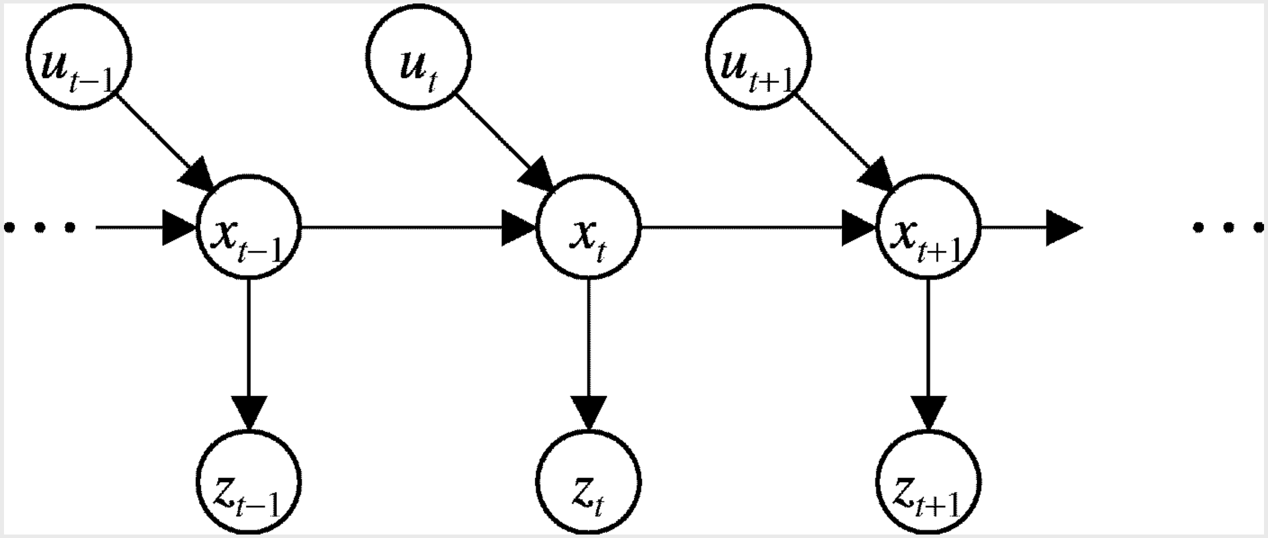

$$p(z_t \mid x_{0:t}, z_{1:t-1}, u_{1:t}) = p(z_t \mid x_t)$$

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

Underlying Assumptions

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

Dynamic Environments

- Two possible locations x_1 and x_2
- $P(x_1)=0.99$
- $P(z|x_2)=0.09$ $P(z|x_1)=0.07$

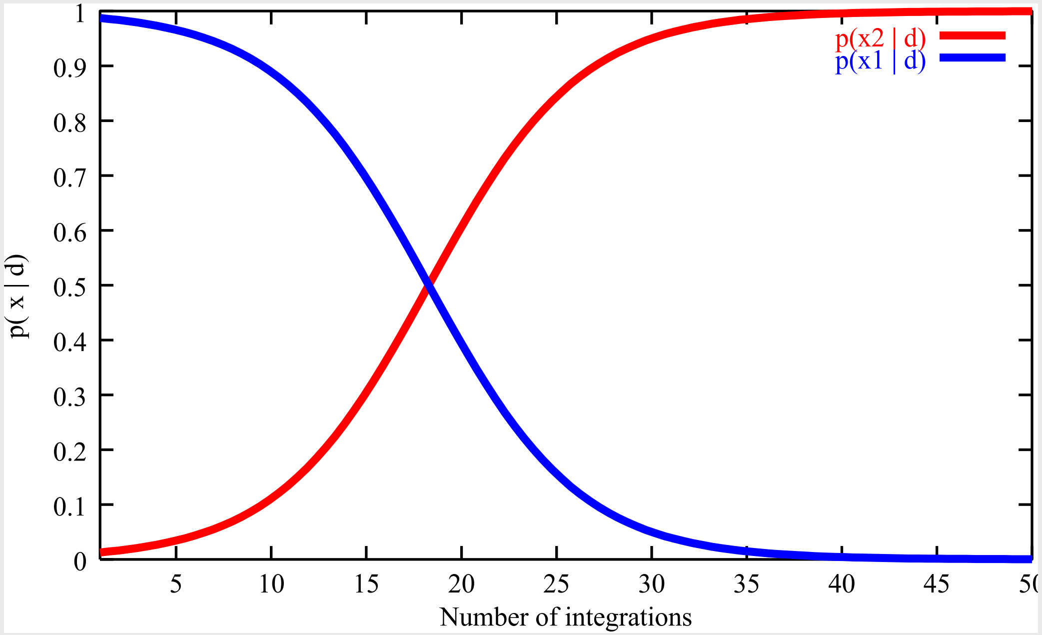

Bayes Filters for Robot Localization

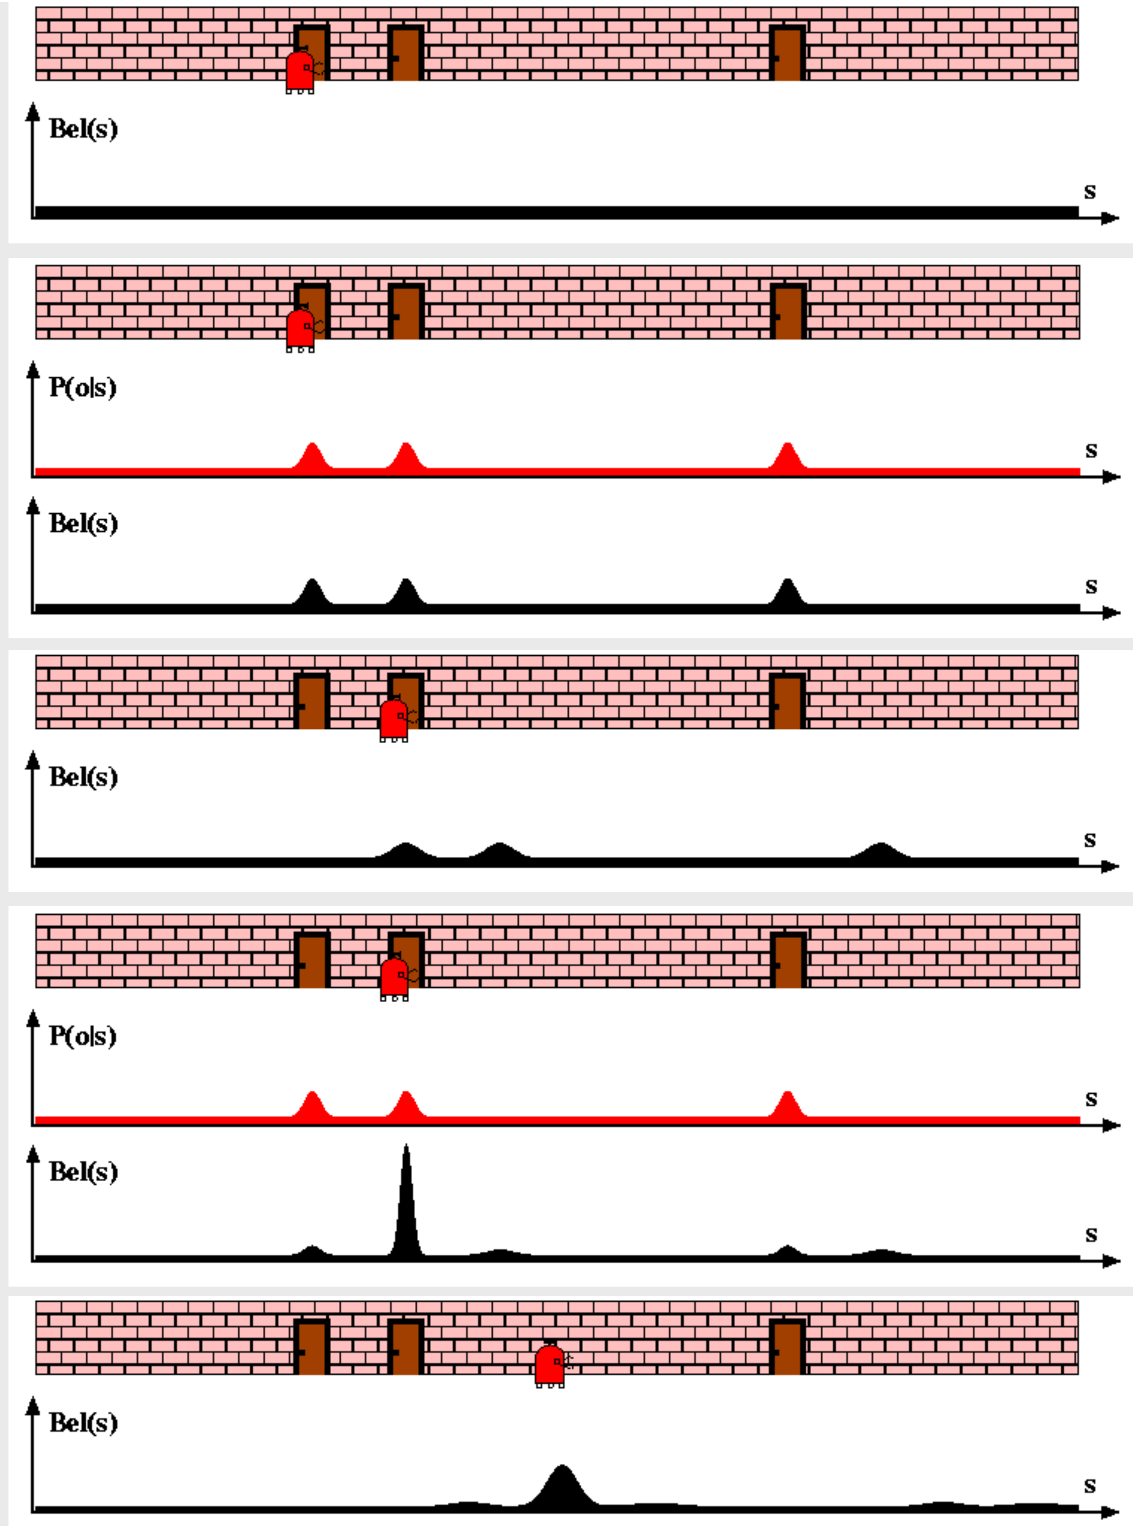

Representations for Bayesian Robot Localization

Discrete approaches (’ 95)
- Topological representation (’ 95)
  - uncertainty handling (POMDPs)
  - occas. global localization, recovery
- Grid-based, metric representation (’ 96)
  - global localization, recovery

Particle filters (’ 99)
- sample-based representation
- global localization, recovery

AI

Kalman filters (late-80s)
- Gaussians, unimodal
- approximately linear models
- position tracking

Robotics

Multi-hypothesis (’ 00)
- multiple Kalman filters
- global localization, recovery

Bayes Filters are Familiar!

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

- 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.