

Probabilistic Robotics

Chapter 2: Recursive State Estimation

MSc course Artificial Intelligence 2017

<http://staff.fnwi.uva.nl/a.visser/education/ProbabilisticRobotics/>

Arnoud Visser & Emiel Hoogeboom

Informatics Institute

Universiteit van Amsterdam

A.Visser@uva.nl

Trends in Robotics Research

Classical Robotics (mid-70's)

- exact models
- no sensing necessary

Reactive Paradigm (mid-80's)


- no models
- relies heavily on good sensing

Hybrids (since 90's)

- model-based at higher levels
- reactive at lower levels

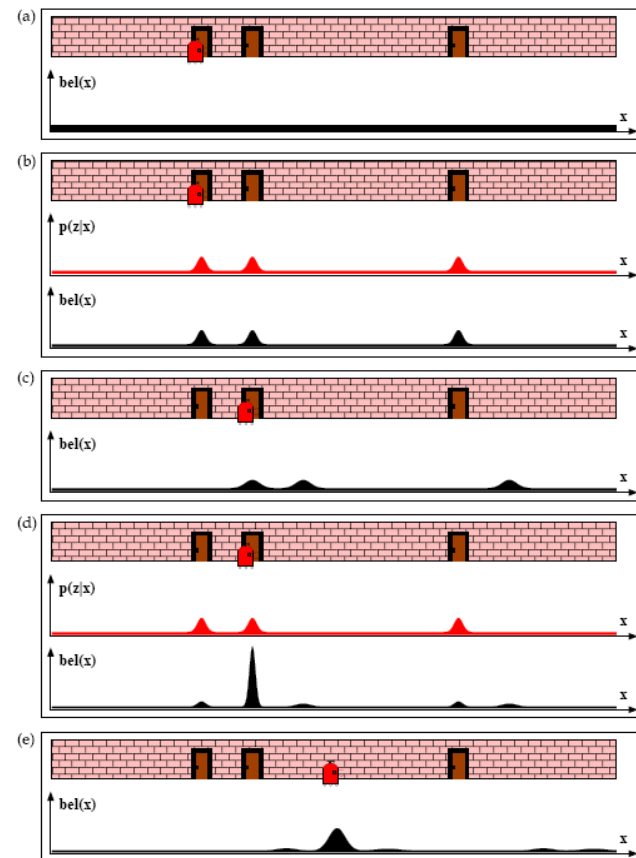

Probabilistic Robotics (since mid-90's)

- seamless integration of models and sensing
- inaccurate models, inaccurate sensors



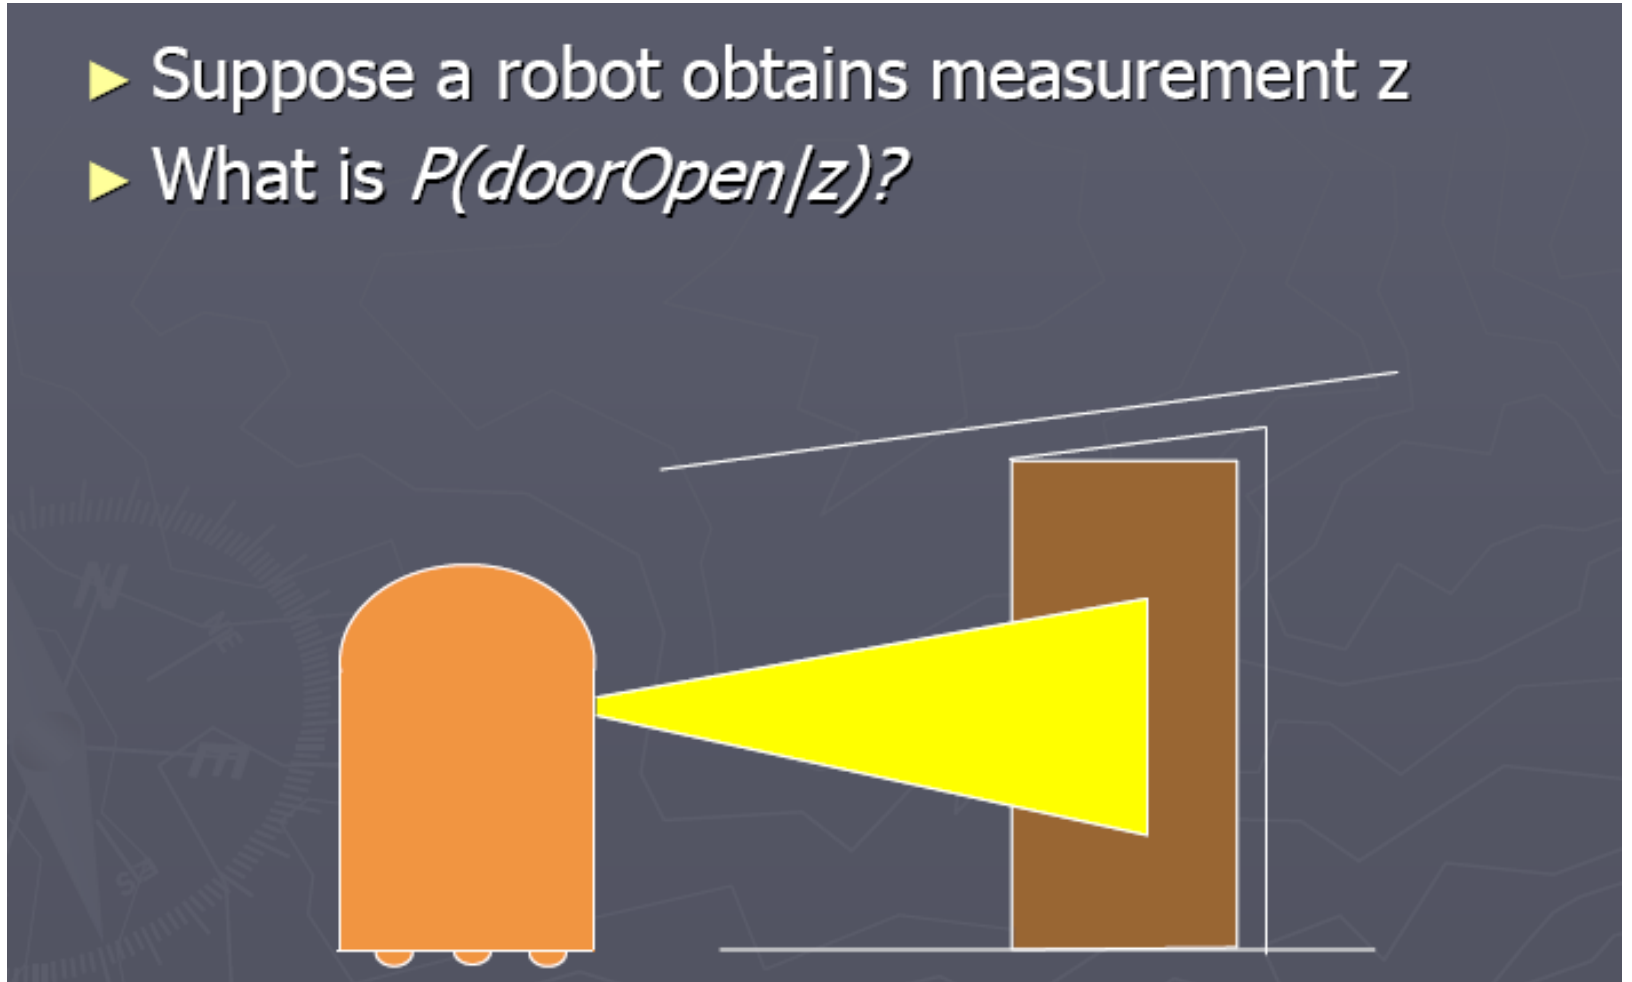
The big picture of Probabilistic Robotics

- Instead of relying on a single ‘best guess’, maintain information of probability distributions over a whole space of guesses.



Simple example of conditional probability

- ▶ Suppose a robot obtains measurement z
- ▶ What is $P(\text{doorOpen}/z)$?



Reverse the evidence: Bayesian Reasoning

- $P(open|z)$ is diagnostic.
- $P(z|open)$ is causal.
- Often causal knowledge is easier to obtain.
- Bayes rule allows us to use causal knowledge for diagnostics:

$$P(open | z) = \frac{P(z | open)P(open)}{P(z)}$$

Bayes Formula

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

\Rightarrow

$$P(x | y) = \frac{P(y | x) P(x)}{P(y)}$$

\sim

$$\text{posterior} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Thomas Bayes (1702 – 1761)



“Now that the errors arising from the imperfection of the instruments & the organs of sense shou’d be reduced to nothing or next to nothing only by multiplying the number of observations seems to me extremely incredible.

On the contrary the more observations you make with an imperfect instrument the more certain it seems to be that the error in your conclusion will be proportional to the imperfection of the instrument made use of. “

Bayes' imperfect instrument




FIG. 15. *An early billiard table without pockets.*

Modern Example

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

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

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

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

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

Combining Evidence

- Suppose our robot obtains another observation z_2 .
- How can we integrate this new information?
- More generally, how can we estimate $P(x | z_1 \dots z_n)$?

Example: Second Measurement

□ $P(z_2 | open) = 0.5$

$$P(z_2 | \neg open) = 0.6$$

□ $P(open | z_1) = 2/3$

$$\begin{aligned} P(open | z_2, z_1) &= \frac{P(z_2 | open) P(open | z_1)}{P(z_2 | open) P(open | z_1) + P(z_2 | \neg open) P(\neg open | z_1)} \\ &= \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{5} \cdot \frac{1}{3}} = \frac{5}{8} = 0.625 \end{aligned}$$

- observation z_2 lowers the probability that the door is open.

Recursive Bayesian Updating

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

This is Bayes rule with background knowledge z :

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

with $y = z_n$ and $z = z_1, \dots, z_{n-1}$

Markov Assumption

z_n is independent of z_1, \dots, z_{n-1} if we know x .

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



$$\begin{aligned} P(x | z_1, \dots, z_n) &= \frac{P(z_n | x) P(x | z_1, \dots, z_{n-1})}{P(z_n | z_1, \dots, z_{n-1})} \\ &= \eta P(z_n | x) P(x | z_1, \dots, z_{n-1}) \\ &= \eta_{1\dots n} \prod_{i=1\dots n} P(z_i | x) P(x) \end{aligned}$$

Complete state

- State x_t will be called *complete* if it is the best predictor of the future
- I.e., no additional information would help us predict the future more accurately

Typical Actions

- ❑ The robot **turns its wheels** to move
- ❑ The robot **uses its manipulator** to grasp an object
- ❑ Plants grow over **time...**

- ❑ Actions are **never carried out with absolute certainty.**
- ❑ In contrast to measurements, **actions generally increase the uncertainty.**


Modeling Actions

- To incorporate the outcome of an action u into the current “belief”, we use the state transition function

$$P(x|u,x')$$


- This term specifies the probability distribution function that **executing u changes the state from x' to x .**

Example: Closing the door



State Transitions

$P(x|u,x')$ for $u = \text{“close door”}$:



If the door is open, the action “close door” succeeds in 90% of all cases.

Integrating the Outcome of Actions

Continuous case:

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

Discrete case:

$$P(x | u) = \sum P(x | u, x') P(x')$$

Example: The Resulting Belief

$$\begin{aligned}P(\textit{closed} | u) &= \sum P(\textit{closed} | u, x')P(x') \\ &= P(\textit{closed} | u, \textit{open})P(\textit{open}) \\ &\quad + P(\textit{closed} | u, \textit{closed})P(\textit{closed}) \\ &= \frac{9}{10} * \frac{5}{8} + \frac{1}{1} * \frac{3}{8} = \frac{15}{16}\end{aligned}$$

$$\begin{aligned}P(\textit{open} | u) &= \sum P(\textit{open} | u, x')P(x') \\ &= P(\textit{open} | u, \textit{open})P(\textit{open}) \\ &\quad + P(\textit{open} | u, \textit{closed})P(\textit{closed}) \\ &= \frac{1}{10} * \frac{5}{8} + \frac{0}{1} * \frac{3}{8} = \frac{1}{16} \\ &= 1 - P(\textit{closed} | u)\end{aligned}$$

Bayes Filters: Framework

□ Given:

- Stream of observations z and action data u :

$$d_t = \{u_1, z_1 \dots, u_t, z_t\}$$





- 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**:

$$Bel(x_t) = P(x_t | u_1, z_1 \dots, u_t, z_t)$$


Bayes Filter Example



Dynamic Bayesian Network for Controls, States, and Sensations



Markov Assumption



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

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

Underlying Assumptions

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

z = observation
u = action
x = state

Bayes Filters

$$\boxed{Bel(x_t)} = P(x_t | u_1, z_1, \dots, u_t, z_t)$$

Bayes $= \eta P(z_t | x_t, u_1, z_1, \dots, u_t) P(x_t | u_1, z_1, \dots, u_t)$

Markov $= \eta P(z_t | x_t) P(x_t | u_1, z_1, \dots, u_t)$

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

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

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

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

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


1. Algorithm `Bayes_filter`($Bel(x), d$):
2. $\eta = 0$
3. If d is a **perceptual** data item z then
 4. For all x do
 5. $Bel'(x) = P(z | x) Bel(x)$
 6. $\eta = \eta + Bel'(x)$
 7. For all x do
 8. $Bel(x) = \eta^{-1} Bel'(x)$
9. Else if d is an **action** data item u then
 10. For all x do
 11. $Bel'(x) = \int P(x | u, x') Bel(x') dx'$
12. Return $Bel(x)$

Bayes Filters are Familiar!

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

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

Baves Filters in Localization



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

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.

Homework

- This Thursday: Assignment 2.8.1 – 2.8.4