Robotics - Localization & Bayesian Filtering

Simone Ceriani

ceriani@elet.polimi.it

Dipartimento di Elettronica e Informazione
Politecnico di Milano

10 May 2012

Outline

Outline

Mobile Robot Localization - Introduction

THE PROBLEM

- Determining *pose* of robot
- Relative to a *given* map of the environment
- a.k.a. *position estimation*

NOTES

- It's an instance of the *general localization*

  i.e., localize objects in the workspace of a manipulator

Localization - The problem

Localization Input

- Known map in a reference system
- *Perception* of the environment
- *Motion* of the robot

Localization Goal

- Determine robot position w.r.t. the map

  i.e. the relative transformation $\mathbf{T}_{WR}^{(W)}$
- Problem of *coordinate transformation*

  $\mathbf{T}_{WR}^{(W)}$ allow to express *objects* position (maps) in a local frame (w.r.t. the robot)

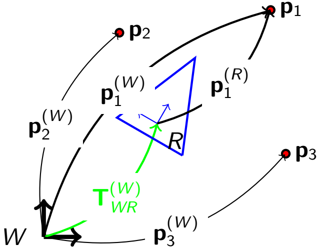

- $\mathbf{p}_i^{(W)}$: the map
- $\mathbf{p}_1^{(R)}$: robot perception
- $\mathbf{T}_{WR}^{(W)}$: localization

Localization - Issues

Direct pose sensing

- Usually impossible
- Noise corruption

Pose estimation

- Inferred from data

  Usually a single sensor measure is insufficient

- Robot need to integrate information over time

  e.g. map with two identical corridors

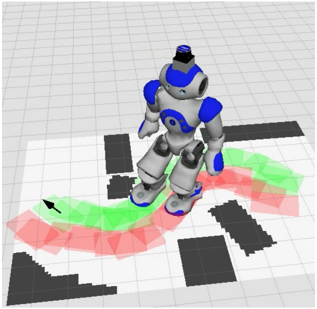

Map

- Various representations are possible

  according to the problem

- Key concept: localization *needs a precise map*

Localization - Graphical model

Graphical model of mobile robot localization

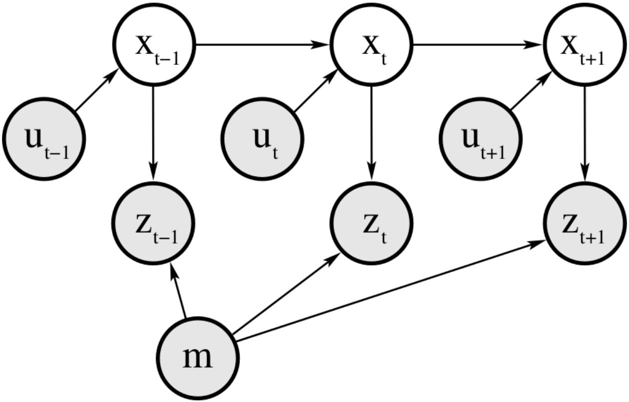

- **X**.: robot pose
- **m**: the map
- **z**.: measurements
- **u**: inputs (e.g. speed of wheels, ...)

- Values of shaded nodes are known

Localization - Maps

HAND-MADE METRIC MAP

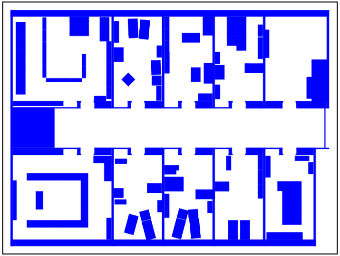

OCCUPANCY GRID MAP

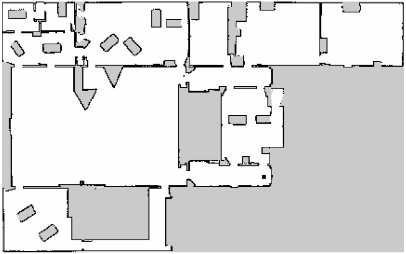

GRAPH-LIKE TOPOLOGICAL MAP

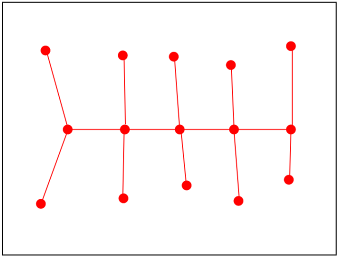

IMAGE MOSAIC OF CEILING

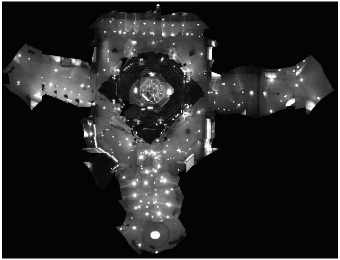

Images from *Probabilistic Robotics - S. Thrun, W. Burgard, D. Fox*

Outline

Taxonomy - Local vs Global - 1

Local vs Global localization

- Depends on information available *initially* and at *run-time*
- Three types of localization problems, with increasing degree of difficulty

1. Pose Tracking

- Assume *initial position* is known
- Localization achieved by accommodating the noise in robot motion
- Pose uncertainty often approximated by a *unimodal distribution*
- It is a *local* problem, uncertainty is confined near robot true pose

2. Global Localization

- Initial pose unknown
- The robot knows that it does not know where it is
- Approaches cannot assume bound on (initial) pose error
- It is not a local problem, estimation could be very far from true pose
- More complicated than pose tracking

Taxonomy - Local vs Global - 2

3. Kidnapped robot problem

- Variant of the *global localization problem*
- The robot can get kidnapped and teleported to some other location
- The robot might believe it knows where it is while it does not
- Even more difficult
- *Robot are not really kidnapped in practice*
- Practical importance: recover from failures in localization

In these lessons

- Markov Localization: general framework
- Pose Tracking: Extended Kalman Filtering
- Global Localization: Particle Filtering / Monte Carlo approaches

Taxonomy - Static vs Dynamic Environment

Static Environment

- Only the robot pose change during operation
- Environment (the map) is static

Dynamic Environment

- Environment changes over time
- Environment changes affects sensor measurements
- Environment changes: temporary or permanent
- e.g.: doors, furniture, walking people, daylight
- Localization more difficult than in a static environment

Taxonomy - Passive vs Active Approaches

Passive approach

- Localization module *observes* the robot operating
- Robot motion is not aimed in facilitating localization

Active approach

- Localization module controls the robot so as to
  - minimize the localization error
  - avoid hazardous movement of a poorly localized robot
- e.g., *coastal navigation*, symmetric corridors
- Trade-off: localization performance vs ability to performs operations

Taxonomy - Single vs Multirobot

SINGLE ROBOT LOCALIZATION

- Most commonly studied
- All data is collected on the robot, no communication issue

MULTIROBOT APPROACH

- Arises in team of robot
- Could be treated as $n-$single robot localization problem
- If robots are able to detect each other, there is opportunity to do better

Outline

Uncertainty in Robotics & Probabilistic Robotics

Robotics systems

- Situated in the physical world
- Perceive information through sensors
- *manipulate* through physical forces
- Have to be able to accommodate uncertainties that exists in the physical world

Factors that contribute to robot's uncertainty

- *Real environments* are inherently unpredictable
- *Sensors* are limited in range, resolution, subject to noise
- *Actuation* involves motors; uncertainty arises from control noise, wear and tear.
- *Mathematical models* are approximation of real phenomenal

Level of uncertainty

- Depends on the application domain
  - in well known environments, like assembly lines, could be bounded
  - in the *open world* plays a key role
- Managing uncertainty is a key step towards robust real-world robot systems

Discrete Random Variables

Discrete Random Variables

- X: a random variables

  e.g., consider a die rolling experiment, X is the variable representing the outcome
- $\Pr(X = x)$ represent the probability that X has value x

  e.g., X assume one of $\{1, 2, 3, 4, 5, 6\}$
- x: a specific values that X might assume (on a discrete set)

  e.g., $\Pr(X = 1) = \Pr(X = 2) = \cdots \Pr(X = 6) = \frac{1}{6}$

Properties - 1

- $\sum_{\forall x} \Pr(X = x) = 1$: discrete probability sum to 1
- $\Pr(X = x) \geq 0$: probability is non negative

  e.g., $\Pr(X = 0) = \Pr(X = 7) = 0$, impossible event
- $\Pr(X = x) \leq 1$: probability is bounded to 1

  e.g., $\Pr(X = \{1 \text{ or } 2 \text{ or } 3 \text{ or } 4 \text{ or } 5 \text{ or } 6\})$, sure event
- Common abbreviation: $\Pr(x)$ instead of $\Pr(X = x)$

Discrete Random Variables - Properties

Properties - 2

- Consider two event A and B

  e.g., A is “die outcome is 2 or 3 ”, B is “die outcome is even”
- $\Pr(A \vee B) = \Pr(A) + \Pr(B) - \Pr(A \wedge B)$

  e.g., $A \vee B$ is “die outcome is 2 or 3 or 4 or 6”, $\Pr(A \vee B) = \frac{2}{3} = \frac{1}{3} + \frac{1}{2} - \frac{1}{6}$
- If $A \wedge B = \emptyset \quad \rightarrow \quad \Pr(A \vee B) = \Pr(A) + \Pr(B)$
- $\Pr(\overline{A}) = 1 - \Pr(A)$
- $\Pr(A \vee \overline{A}) = \Pr(A) + \Pr(\overline{A}) - \Pr(A \wedge \overline{A}) = 1$

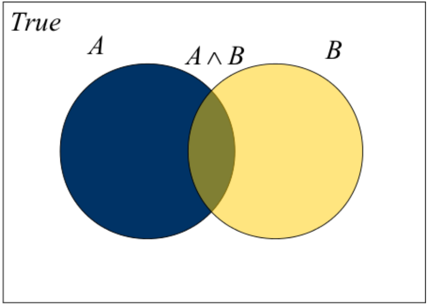

Remarks

- *Relative-frequency* (i.e., outcome of experiments) are alternative (not rigorous) ways of introduce the concept of probability.

Continuous Random Variables

Continuous Random Variables

- Allow to address continuous space
- Possess a Probability Density Function (p.d.f.)
  $f_X(x), x \in \mathbb{R}$
- $\Pr(X = x) = 0$, even though it is not impossible
- The integral is the Cumulative Density Function (c.d.f.)
  $F_X(x) = \Pr(x \leq x) = \int_{-\infty}^{x} f_X(x)dx$
- $\lim_{x \to \infty} F_X(x) = \Pr(X \leq \infty) = 1$
- $\Pr(x \in (a, b)) = \int_{a}^{b} f_X(x)dx = F_X(b) - F_X(a)$

Uniformly Distributed Continuous Random Variable

Uniformly Distributed Continuous Random Variable

- Uniformly Distributed in $[a, b]$
- p.d.f.: $f(x) = \begin{cases} \frac{1}{b-a} & x \in [a, b] \\ 0 & \text{elsewhere} \end{cases}$
- c.d.f.: $F_X(x) = \begin{cases} 0 & x < a \\ \frac{x-a}{b-a} & x \in [a, b] \\ 1 & x > b \end{cases}$

P.D.F.

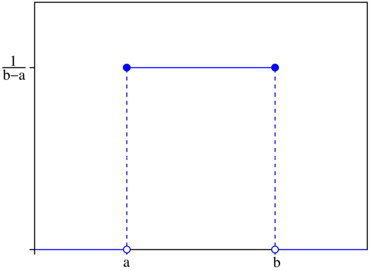

C.D.F.

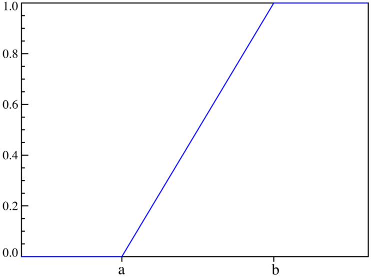

Normal Distributed Random Variable

Normal Distributed Random Variable

- a.k.a. Gaussian Random Variable
- $\mathcal{N}(\mu, \sigma^2)$
  - mean: μ
  - variance: σ^2
- p.d.f.:

$$f(x) = (2\pi\sigma^2)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right\}$$

- c.d.f.: $\nexists$ explicit formula,

  tabulated for $\mathcal{G}(\eta)$, c.d.f. of $\mathcal{N}(0,1)$
- $F(x) = \mathcal{G}(\frac{x-\mu}{\sigma})$

  $\frac{x-\mu}{\sigma}$: number of σ away from the mean

P.D.F.

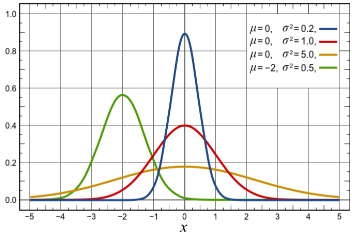

C.D.F.

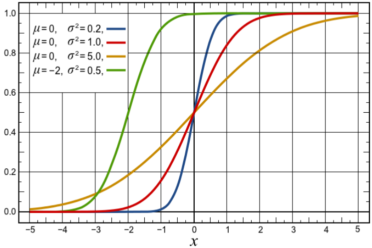

Multivariate Distributions

Multivariate

- $\mathbf{x}$ is a vector
- $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$, with
  - $\boldsymbol{\mu}$: $n \times 1$, mean vector
  - Σ: $n \times n$, covariance matrix
- $f(\mathbf{x}) =$
  $(2\pi \det(\Sigma))^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu})\right\}$
- $\mu = \begin{bmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{bmatrix}$ $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \cdots & \Sigma_{1n} \\ \Sigma_{21} & \Sigma_{22} & \cdots & \Sigma_{2n} \\ \vdots & \cdots & \ddots & \vdots \\ \Sigma_{n1} & \Sigma_{n2} & \cdots & \Sigma_{nn} \end{bmatrix}$
  - diagonal elements are the variances of the single components
  - off diagonals elements are the covariances between elements
  - Σ is symmetric and positive definite, (non singular)
  - if $\exists$ linear relation among components, Σ is positive semi-definite, (singular)

P.D.F. - $n = 2$

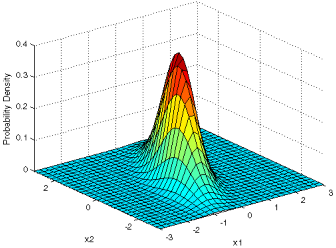

C.D.F. - $n = 2$

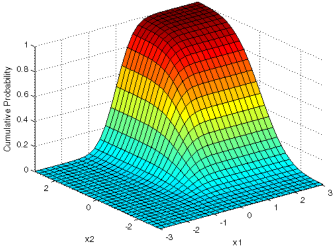

Mean, Variance, Moments

Discrete case

- Expected value: $E[X] = \sum_x xp(x) = \overline{x}$
- n-th moment: $E[X^n] = \sum_x x^n p(x)$
- Variance: $VAR(x) = E[(x-\overline{x})^2]$
  $= \sum_x (x-\overline{x})^2 p(x) = \sigma^2$
  $= E[X^2] - E[X]^2$
  a.k.a. second central moment

Continuous case

- Expected value: $E[X] = \int_x xf(x)dx$
- n-th moment: $E[X^n] = \int_x x^n p(x)$
- Variance: $VAR(x) = E[(x-\overline{x})^2]$
  $= \int_x (x-\overline{x})^2 p(x) = \sigma^2$
  $= E[X^2] - E[X]^2$
  a.k.a. second central moment

Properties

- consider $Y = aX + b$, X random variable, a, b scalar quantities
- $E[Y] = aE[X] + b$
- $VAR(Y) = a^2 VAR(Y)$

Joint Probability, Independence, Conditioning

Joint Probability

- Consider two random variables X and Y
  or a random vector $\mathbf{Z} = \left[X, Y\right]^T$
- $\Pr(X = x \wedge Y = y) = \Pr(x, y) = \Pr(\mathbf{z})$
  with $\mathbf{z} = \left[x, y\right]^T$

Independence

- X and Y are independent if and only if
- $\Pr(x, y) = \Pr(x)\Pr(y)$
  or with p.d.f. $f_{xy}(x, y) = f_x(x)f_y(y)$

Conditioning

- $\Pr(x|y)$ is the probability of x given y
- $\Pr(x, y) = \Pr(x|y)\Pr(y)$
- $\Pr(x|y) = \Pr(x)$ if x and y are independent
  if X and Y are independent, Y tell us nothing about the value of X,
  there is no advantage of knowing the value of Y if we are interested in X

Conditional Independence

Conditional Independence

- X and Y are *conditional independent* on Z if $\Pr(x, y|z) = \Pr(x|z)\Pr(y|z)$
- This is equivalent to $\Pr(x, y|z) = \frac{\Pr(x,y,z)}{p(z)} = \Pr(x|y, z)\Pr(y|z) = \frac{\Pr(x|y,z)\Pr(y,z)}{\Pr(z)}$
- Thus, X and Y are *conditional independent* on Z if $\Pr(x|z) = \Pr(x|y, z)$

  i.e., knowledge on y does not add any information to x if z is known

Total probability, marginals

- Total probability:

  $\int_{x_1=-\infty}^{\infty} \cdots \int_{x_n=-\infty}^{\infty} f_{x_1 \ldots x_n}(x_1, \ldots, x_n) dx_1 \ldots dx_n = 1$

- Marginal distribution:

  $\int_{-\infty}^{\infty} f_{x,y}(x, y) dx = f_y(y)$

  $\int_{-\infty}^{\infty} f_{x,y}(x, y) dy = f_x(x)$

- Marginal distribution with conditioning:

  $\int_{-\infty}^{\infty} f_{y|x}(y|x) f_x(x) dx = f_y(y)$

  $\int_{-\infty}^{\infty} f_{x|y}(x|y) f_y(y) dy = f_x(x)$

Outline

Bayes Formula

From conditioning

- $\Pr(x, y) = \Pr(x|y)\Pr(y)$
- $\Pr(x, y) = \Pr(y|x)\Pr(x)$

Bayes formula

- $\Pr(x|y) = \frac{\Pr(y|x)\Pr(x)}{\Pr(y)} = \frac{\Pr(y|x)\Pr(x)}{\int_{-\infty}^{\infty} f_{y|x}(y|x')f_x(x')dx'}$
- $\Pr(x)$ is the *prior*, the belief about x
- y is the *data*, e.g., a sensor measure
- $\Pr(y|x)$ is the *likelihood*, i.e., how much is probable to have measure y in state x
- $\Pr(x|y)$ is the *posterior*, i.e., the belief x state given the measurement y
- Bayes formula allow to infer a quantity x from data y through *inverse probability*
  i.e., through the probability of data y assuming that the state is x

Bayes Example - 1

PROBLEM

- A robot "observe" a door

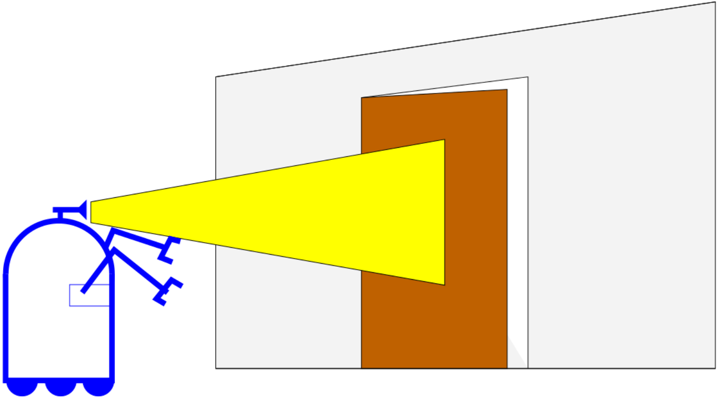

Bayes Example - 2

PROBLEM

- A robot "observe" a door
- The door could be *open* or *close*
- The sensor measure a distance as *far* or *near*
- The probability that the door is *open* is 0.4
- The probability that the sensor measure *far* when the door is *open* is 0.8
- The probability that the sensor measure *far* when the door is *close* is 0.1
- What is the probability that the door is *open* if the sensor measurement is *near*?
- What is the probability that the door is *open* if the sensor measurement is *far*?
- What is the probability that the door is *close* if the sensor measurement is *near*?
- What is the probability that the door is *close* if the sensor measurement is *far*?

Bayes Example - 3

Variable definition

- X: door state, $\{open, close\}$
- Y: sensor measure, $\{open, close\}$

P.D.F

- $\Pr(X{=}open){=}0.4$
- $\Pr(Y{=}far|X{=}open){=}0.8$
- $\Pr(Y{=}far|X{=}close){=}0.1$
- $\Pr(X{=}close){=}0.6$
- $\Pr(Y{=}near|X{=}open){=}0.2$
- $\Pr(Y{=}near|X{=}close){=}0.9$

Solution

- $\Pr(X = open|Y = near) = \frac{\Pr(Y{=}near|X{=}open)\Pr(X{=}open)}{\Pr(Y{=}near|X{=}open)\Pr(X{=}open)+\Pr(Y{=}near|X{=}close)\Pr(X{=}close)}$
  $= \frac{0.2\cdot 0.4}{0.2\cdot 0.4+0.9\cdot 0.6} = 0.13$
- $\Pr(X = open|Y = far) = \frac{\Pr(far|open)\Pr(open)}{\Pr(far|open)\Pr(open)+\Pr(far|close)\Pr(close)} = \frac{0.8\cdot 0.4}{0.8\cdot 0.4+0.1\cdot 0.6} = 0.84$
- $\Pr(X = close|Y = near) = \frac{\Pr(near|close)\Pr(close)}{\Pr(near)} = \frac{0.9\cdot 0.6}{0.62} = 0.87$
- $\Pr(X = close|Y = far) = \frac{\Pr(far|close)\Pr(close)}{\Pr(far)} = \frac{0.1\cdot 0.6}{0.9+0.2} = 0.16$

Bayes Recursive Updating - 1

New measurements

- Suppose that we get the first measurement: *near*
- Thus, $\Pr(X = \text{open}|\, Y_1 = \text{far}) = 0.84$, and $\Pr(X = \text{close}|\, Y_1 = \text{far}) = 0.16$
- A second measure Y_2 arrives: it is *far*
- $\Pr(X = \text{open}|\, Y_1 = \text{far}, Y_2 = \text{far})$?
- More generally, how to estimate $\Pr(X = \text{open}|\, Y_1 = y_1, Y_2 = y_2, \ldots Y_n = y_n)$?

Extend the Bayes rule

- $\Pr(X|\, Y_1, Y_2) = \dfrac{\Pr(Y_2|X, Y_1)\Pr(X|\, Y_1)}{\Pr(Y_2|\, Y_1)} = \dfrac{\Pr(Y_2|X, Y_1)\Pr(X|\, Y_1)}{\int_{-\infty}^{\infty} f_{Y_2|Y_1,x'}(Y_2|Y_1, x') f_{X|Y_1}(x'|\, Y_1) dx'}$
- $\Pr(X|\, Y_1, Y_2, \ldots, Y_n) = \dfrac{\Pr(Y_n|X, Y_1, , \ldots, Y_{n-1})\Pr(X|\, Y_1, Y_2, \ldots, Y_{n-1})}{\Pr(Y_n|\, Y_1, \ldots, Y_{n-1})} = \ldots$

Bayes Recursive Updating - 2

Markov Assumption

- *Markov assumption*: Y_2 independent of Y_1 if we know X
- Then $\Pr(Y_2|X, Y_1) = \Pr(Y_2|X)$ (see Conditional Independence formulas)
- Bayes rule

$$\Pr(X|Y_1, Y_2) = \frac{\Pr(Y_2|X, Y_1)\Pr(X|Y_1)}{\Pr(Y_2|Y_1)} = \frac{\Pr(Y_2|X)\Pr(X|Y_1)}{\Pr(Y_2|Y_1)}$$

$$\Pr(X|Y_1, Y_2, \ldots, Y_n) = \frac{\Pr(Y_n|X, Y_1, , \ldots, Y_{n-1})\Pr(X|Y_1, \ldots, Y_{n-1})}{\Pr(Y_n|Y_1, \ldots, Y_{n-1})} = \frac{\Pr(Y_n|X)\Pr(X|Y_1, \ldots, Y_{n-1})}{\Pr(Y_n|Y_1, \ldots, Y_{n-1})}$$

- with

$$\Pr(Y_2|Y_1) = \int_{-\infty}^{\infty} f_{Y_2|Y_1,x'}(Y_2|Y_1, x')f_{X|Y_1}(x'|Y_1)dx'$$

$$= \int_{-\infty}^{\infty} f_{Y_2|x'}(Y_2|x')f_{X|Y_1}(x'|Y_1)dx'$$

similar with n measurement

note: use $\sum$ in discrete case

Bayes Recursive Updating - 3

The example

- $\Pr(X = \text{open}|Y_1 = \text{far}) = 0.84$, and $\Pr(X = \text{close}|Y_1 = \text{far}) = 0.16$, $Y_2 = \textit{far}$

$$
\begin{aligned}
\Pr(Y_2|Y_1) &= \Pr(Y_2|Y_1, \text{open}) \Pr(\text{open}|Y_1) + \Pr(Y_2|Y_1, \text{close}) \Pr(\text{close}|Y_1) \\
&= \Pr(Y_2|\text{open}) \Pr(\text{open}|Y_1) + \Pr(Y_2|\text{close}) \Pr(\text{close}|Y_1) \\
&= \Pr(\text{far}|\text{open}) \Pr(\text{open}|\text{far}) + \Pr(\text{far}|\text{close}) \Pr(\text{close}|\text{far}) \\
\Pr(\text{far}|\text{far}) &= 0.8 \cdot 0.84 + 0.1 \cdot 0.16 = 0.688
\end{aligned}
$$

- $\Pr(X = \text{open}|Y_1 = \text{far}, Y_2 = \text{far}) = \frac{\Pr(\text{far}|\text{open}) \Pr(\text{open}|\text{far})}{\Pr(\text{far}|\text{far})} = \frac{0.8 \cdot 0.84}{0.688} = 0.977$

Outline

Robot environment interaction - 1

THE ENVIRONMENT

- a.k.a. *World*
- generally it's a *dynamic system*:
  - Robot can act on it
  - Changes due to time passing by

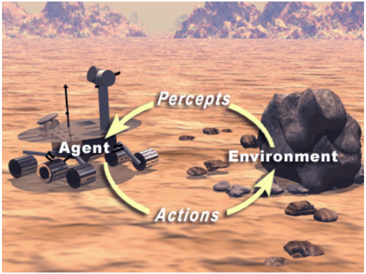

A ROBOT

- Can act on environment
  i.e., change the environment state
- Can sense environment through sensor
- Has an internal *belief* on state

STATE

- Collection of all aspects of the robot and the environment
- Generally, changes over time, some part could be static
- We will refer it with x_t

Robot environment interaction - 2

Control Actions

- Change the state of the world
  (robot and/or environment)
- $u_{t_1:t_n} = u_{t_1}, u_{t_2}, \dots u_{t_n}$

Measurements

- Information about the environment
  (distances, images, ...)
- $z_{t_1:t_n} = z_{t_1}, z_{t_2}, \dots z_{t_n}$

Complete State and Markov Chain

- x_t will be called *complete* if it is the best predictor of the future
- All past states, measurements and inputs carry no additional information
  to predict the future more accurately
- No variables prior to x_t may influence the stochastic evolution of future state
- This a *Markov Chain*

Evolution of state - 1

Evolution of state

- x_t is stochastically generated by x_{t-1}
- x_t p.d.f is conditioned on past states, inputs and measurement
- $p(x_t|x_{0:t-1}, z_{1:t-1}, u_{1:t})$
- Under Markov Chain hypothesis (or Complete State),

  thanks to conditional independence

  $p(x_t|x_{0:t-1}, z_{1:t-1}, u_{1:t}) = p(x_t|x_{t-1}, u_t)$
- $p(x_t|x_{t-1}, u_t)$ is the *state transition probability*
- State evolution is stochastic, not deterministic (i.e., is a p.d.f.)

Measurement process

- $p(z_t|x_{0:t}, z_{0:t-1}, u_{1:t}) = p(z_t|x_t)$ under Complete State
- the state x_t is sufficient to predict the (potentially noisy) measurement z_t
- $p(z_t|x_t)$ is the *measurement probability*
- Specify the probabilistic low of according to which measurement are generated

Evolution of state - 2

Evolution of state and measurements

- Describe the dynamical stochastic system of the robot and its environment
- a.k.a. as Hidden Markov Model (HMM) or Dynamic Bayes Network (DBN)

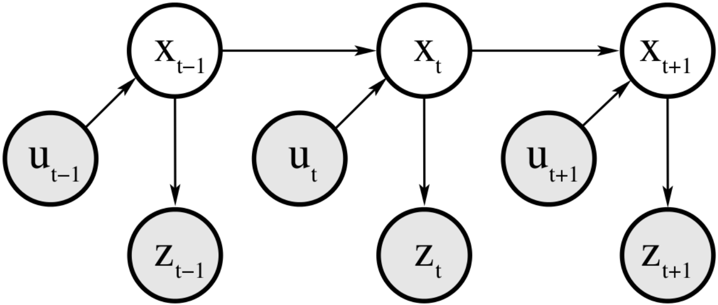

Belief Distributions

Belief

- Reflects the robot's internal knowledge about the state
- Usually, the state cannot be measure directly
- The state need to be inferred from data

Posterior Belief distribution

- Conditional probability
- Is a *posterior* probability
- Conditioned on available data
- $bel(x_t) = p(x_t | z_{1:t}, u_{1:t})$

Prior Belief distribution

- Conditional probability
- Is a *prior* probability
- Conditioned on available data *before incorporating z_t measure*
- $\overline{bel}(x_t) = p(x_t | z_{1:t-1}, u_{1:t})$
- Often referred as a *predition*

Bayes Filter Algorithm

The Algorithm

- Calculate the belief $bel(x_t)$
- Recursive: use $bel(x_{t-1})$ as input
- Use most recent measure (z_t) and input (y_t)

Algorithm

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 $bel(x_t)$

Step 1

- Calculate the prior belief $\overline{bel}(x_t)$
- Integral of the product of
  - The prior on x_t
  - The probability of the state evolution
- It is a *prediction*

Step 2

- Calculate the posterior belief $bel(x_t)$
- Product of
  - Prior distribution
  - Measurement probability
  - η: normalization factor $\eta : \sum_{\forall x_t} bel(x_t) = 1$
- It is a *measurement update*

Bayes Filter Algorithm Example - 1

THE ROBOT AND THE DOOR - V2

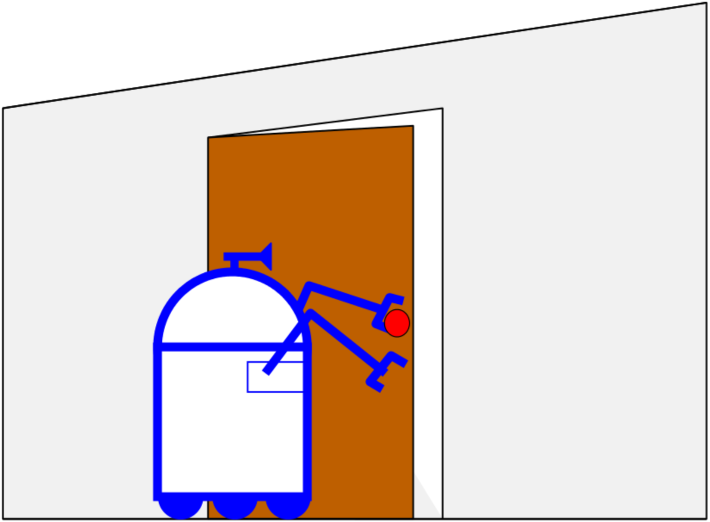

Bayes Filter Algorithm Example - 2

The door

- Can be *open* or *close*
- initial state is unknown

The robot

- Can act (stochastically) on the door:
  - *push*: try to open the door
  - *nop*: no operation
- Sense (noisly) the door presence
  - *near*: read by sensor when door is close
  - *far*: read by sensor when door is open

Initial Belief

- $bel(x_0 = \text{open}) = 0.5$
- $bel(x_0 = \text{close}) = 0.5$

Measurement Probability

- $bel(z_t = \text{far} | x_t = \text{open}) = 0.6$
- $bel(z_t = \text{near} | x_t = \text{open}) = 0.4$
- $bel(z_t = \text{far} | x_t = \text{close}) = 0.2$
- $bel(z_t = \text{near} | x_t = \text{close}) = 0.8$

Bayes Filter Algorithm Example - 3

State Transition Probability

- $bel(x_t = \text{open} | u_t = \text{push}, x_{t-1} = \text{open}) = 1$
- $bel(x_t = \text{close} | u_t = \text{push}, x_{t-1} = \text{open}) = 0$
- $bel(x_t = \text{open} | u_t = \text{push}, x_{t-1} = \text{close}) = 0.8$
- $bel(x_t = \text{close} | u_t = \text{push}, x_{t-1} = \text{close}) = 0.2$

- $bel(x_t = \text{open} | u_t = \text{nop}, x_{t-1} = \text{open}) = 1$
- $bel(x_t = \text{close} | u_t = \text{nop}, x_{t-1} = \text{open}) = 0$
- $bel(x_t = \text{open} | u_t = \text{nop}, x_{t-1} = \text{close}) = 0$
- $bel(x_t = \text{close} | u_t = \text{nop}, x_{t-1} = \text{close}) = 1$

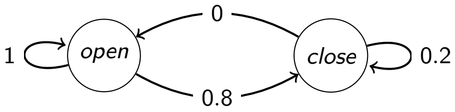

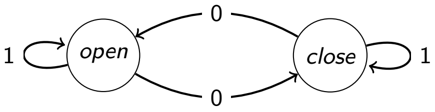

Bayes Filter Algorithm Example - 4

TIME $t = 1$, ACT

- $u_1 =$ nop, no operation performed
- $\overline{bel}(x_1) = \sum_{x_0} p(x_1|u_1, x_0)bel(x_0)$, prediction

$$\begin{aligned}
\overline{bel}(x_1) &= p(x_1|u_1 = \text{nop}, x_0 = \text{open})bel(x_0 = \text{open}) + \\
&+ p(x_1|u_1 = \text{nop}, x_0 = \text{close})bel(x_0 = \text{close}) \\
\overline{bel}(x_1 = \text{open}) &= p(x_1 = \text{open}|u_1 = \text{nop}, x_0 = \text{open})bel(x_0 = \text{open}) + \\
&+ p(x_1 = \text{open}|u_1 = \text{nop}, x_0 = \text{close})bel(x_0 = \text{close}) = \\
&= 1 \cdot 0.5 + 0 \cdot 0.5 = 0.5
\end{aligned}$$

$$\begin{aligned}
\overline{bel}(x_1 = \text{close}) &= p(x_1 = \text{close}|u_1 = \text{nop}, x_0 = \text{open})bel(x_0 = \text{open}) + \\
&+ p(x_1 = \text{close}|u_1 = \text{nop}, x_0 = \text{close})bel(x_0 = \text{close}) = \\
&= 0 \cdot 0.5 + 1 \cdot 0.5 = 0.5
\end{aligned}$$

Bayes Filter Algorithm Example - 5

Time $t = 1$, sense

- $z_1 =$ far, sense open door
- $bel(x_1) = \eta p(z_1 = \text{far}|x_1)\overline{bel}(x_1)$, measurement update

$$\begin{aligned} bel(x_1 = \text{open}) &= \eta p(z_1 = \text{far}|x_1 = \text{open})\overline{bel}(x_1 = \text{open}) \\ &= \eta 0.6 \cdot 0.5 = \eta\, 0.3 \end{aligned}$$

$$\begin{aligned} bel(x_1 = \text{close}) &= \eta p(z_1 = \text{far}|x_1 = \text{close})\overline{bel}(x_1 = \text{close}) \\ &= \eta 0.2 \cdot 0.5 = \eta\, 0.1 \end{aligned}$$

- $bel(x_1 = \text{open}) + bel(x_1 = \text{close}) = 1$
- $\eta\, 0.3 + \eta\, 0.1 = 1$
- $\eta = (0.3 + 0.1)^{-1} = 2.5$

- $bel(x_1 = \text{open}) = 0.75$
- $bel(x_1 = \text{close}) = 0.25$

Bayes Filter Algorithm Example - 6

TIME $t = 2$, ACT

- $u_2 =$ push, perform *push* action
- $\overline{bel}(x_2) = \sum_{x_1} p(x_2|u_2, x_1)bel(x_1)$, prediction

$$
\begin{aligned}
\overline{bel}(x_2) &= p(x_2|u_2 = \text{push}, x_1 = \text{open})bel(x_1 = \text{open}) + \\
&+ p(x_2|u_2 = \text{push}, x_1 = \text{close})bel(x_1 = \text{close}) \\
\overline{bel}(x_2 = \text{open}) &= p(x_2 = \text{open}|u_2 = \text{push}, x_1 = \text{open})bel(x_1 = \text{open}) + \\
&+ p(x_2 = \text{open}|u_2 = \text{push}, x_1 = \text{close})bel(x_1 = \text{close}) = \\
&= 1 \cdot 0.75 + 0.8 \cdot 0.25 = 0.95
\end{aligned}
$$

$$
\begin{aligned}
\overline{bel}(x_2 = \text{close}) &= p(x_2 = \text{close}|u_2 = \text{push}, x_1 = \text{open})bel(x_1 = \text{open}) + \\
&+ p(x_2 = \text{close}|u_1 = \text{push}, x_1 = \text{close})bel(x_1 = \text{close}) = \\
&= 0 \cdot 0.75 + 0.2 \cdot 0.25 = 0.05
\end{aligned}
$$

Bayes Filter Algorithm Example - 7

TIME $t = 2$, SENSE

- $z_2 =$ far, sense open door
- $bel(x_2) = \eta p(z_2 = \text{far}|x_2)\overline{bel}(x_2)$, measurement update

$$\begin{aligned} bel(x_2 = \text{open}) &= \eta p(z_2 = \text{far}|x_2 = \text{open})\overline{bel}(x_2 = \text{open}) \\ &= \eta 0.6 \cdot 0.95 = \eta\, 0.57 \end{aligned}$$

$$\begin{aligned} bel(x_2 = \text{close}) &= \eta p(z_2 = \text{far}|x_2 = \text{close})\overline{bel}(x_2 = \text{close}) \\ &= \eta 0.2 \cdot 0.05 = \eta\, 0.01 \end{aligned}$$

- $bel(x_2 = \text{open}) + bel(x_2 = \text{close}) = 1$
- $\eta\, 0.57 + \eta\, 0.01 = 1$
- $\eta = (0.57 + 0.01)^{-1} = 1.724$

- $bel(x_1 = \text{open}) = 0.983$
- $bel(x_1 = \text{close}) = 0.017$

Outline

Markov Localization

Markov Localization Algorithm

- The state x_t is the robot pose
- Require also a map m as input
- Map m plays a key role in the measurement model

  $p(z_t|x_t, m)$

  measure the relative pose w.r.t. the map
- Can be integrated in the *motion model* (the prediction step)

  $p(x_t|u_t, x_{t-1}, m)$

  avoid prediction of impossible movement, like through a wall

Algorithm Markov_localization($bel(x_{t-1}), u_t, z_t, m$)**:**

for all x_t *do*

$\overline{bel}(x_t) = \int p(x_t \mid u_t, x_{t-1}, m)\ bel(x_{t-1})\ dx$

$bel(x_t) = \eta\ p(z_t \mid x_t, m)\ \overline{bel}(x_t)$

endfor

return $bel(x_t)$

Markov Localization - Initial belief

Pose Tracking - case 1

- The initial pose is known: $\overline{x}_0$
- The initial belief is $bel(x_0) = \begin{cases} 1 & \text{if } x_0 = \overline{x}_0 \\ 0 & \text{otherwise} \end{cases}$

Pose Tracking - case 2

- The initial pose is known with some uncertainty
  e.g. Gaussian $\overline{x}_0 = (N)(\overline{x_0}, \Sigma_0)$
- The initial belief is $bel(x_0) = (2\pi \det(\Sigma))^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu})\right\}$

Global Localization

- The initial pose is unknown, uniform distribution over all the map
- The initial belief is $bel(x_0) = \frac{1}{|X|}$, where $|X|$ is the volume of the space

Markov Localization - Illustration - 1

Environment setup

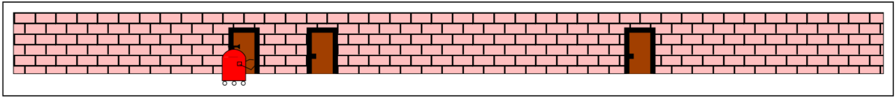

- A one-dimensional hallway
- Three indistinguishable doors
  remember that we know the map
- The robot sense (noisly) a door presence
- The robot known its direction and the relative motion performed in a time step
- The state x_t is the x robot position

Markov Localization - Illustration - 2

Initial Belief

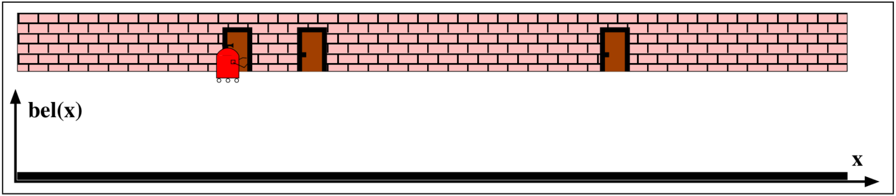

- Initial position is unknown, $bel(x_0)$ is uniformly distributed

Markov Localization - Illustration - 3

SENSE

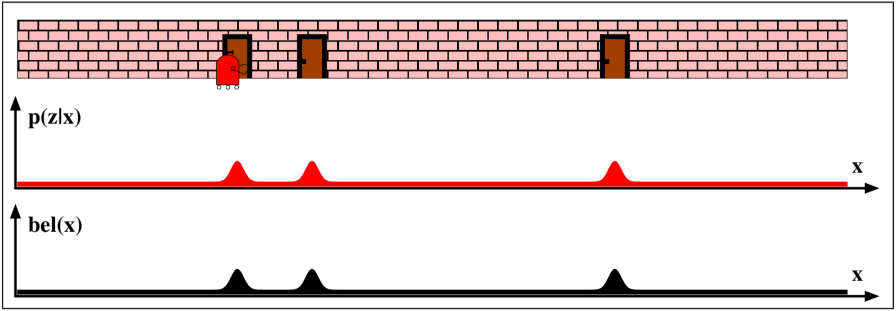

- The robot sense a door presence
- $bel(x_1)$ is higher on door locations

Markov Localization - Illustration - 4

Motion model - prediction step

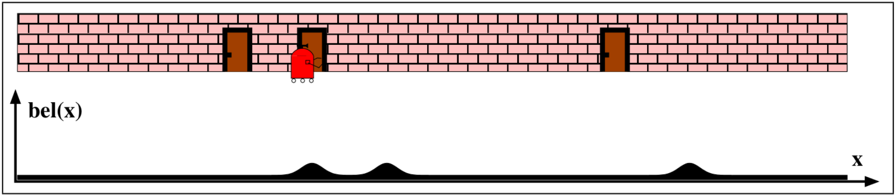

- Robot knows its movement (u, the control variable)
- $\overline{bel}(x_2)$ is shifted as result of motion
- $\overline{bel}(x_2)$ is flattened as result of uncertainty on motion

Markov Localization - Illustration - 5

Sense

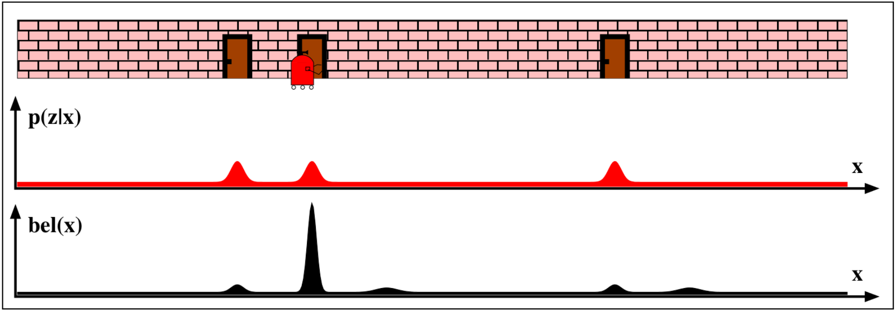

- The robot sense a door presence
- $bel(x_2)$ is focused on the correct pose

Markov Localization - Illustration - 6

Motion model - prediction step

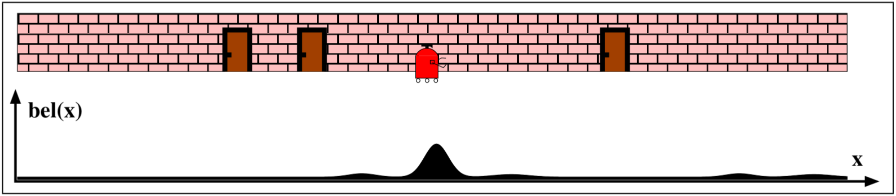

- $\overline{bel}(x_3)$ is focused on the correct pose
- $\overline{bel}(x_3)$ is flattened

References

Reference

"Probabilistic Robotics"

(Intelligent Robotics and Autonomous Agents series)

The MIT Press (2005)

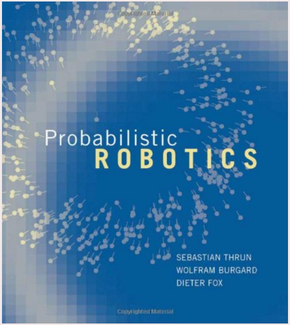

Chapters 2, 3, 7.