MARKOV LOCALIZATION

Maria Isabel Ribeiro
Pedro Lima

mir@isr.ist.utl.pt pal@isr.ist.utl.pt
Instituto Superior Técnico (IST)
Instituto de Sistemas e Robótica (ISR)
Av.Rovisco Pais, 1
1049-001 Lisboa
PORTUGAL

May.2002

Localization Methods: a classification

- Position Tracking
  - A robot knows its initial position and “only” has to accommodate small errors in its odometry as it moves
- Global Localization
  - Involves a robot which is not told its initial position
  - It has to solve a much more difficult localization problem, that of estimating its position from scratch

Markov Localization family of approaches

Nourbakhsh, Powers, Birchfiled, 1995
Simmons, Koenig, 1995
Kaelbling, Cassandra, Kurien, 1996
Burgard et al, 1996

Represent the robot’s **belief** by a probability distribution over possible positions and uses Bayes rule and convolution to update the belief whenever the robot senses or moves

Localization

Handling Uncertainty in Measurements

Problems with the Kalman

- Robot kinematics and dynamics are non-linear – the extended KF must be used, but this one is not proven to converge and is not necessarily the optimal estimator
- Initial posture must be known with Gaussian uncertainty at most
- Non-Gaussian motion or measurement models cause troubles, because the noise will no longer be Gaussian
- Can not recover from tracking failures (e.g., due to “kidnapping”)
- Can not deal with multi-modal densities, typical in global localization

Markov Localization

Handling Uncertainty in Measurements

Markov Localizati

(Fox, Burgard and Thrun, 1999)

- multi-modal probability densities allowed and propagated through the motion model
- A probability density is maintained over the space of all locations of a robot in its environment
- Kalman Filter is a special case (unimodal Gaussian density, only mean and covariance need to be propagated)

Reference

- Dieter Fox, Wolfram Burgard, Sebastian Thrun, "Markov Localization for Mobile Robots in Dynamic Environments", J. of Artificial Intelligence Research 11 (1999) 391-427

See also

- Sebastian Thrun, "Probabilistic Robotics: Tutorial presented in AAAI 200" – available through S.Thrun Web page

http://www-2.cs.cmu.edu/~thrun/

Markov Localization

- What is Markov Localization ?
  - Special case of probabilistic state estimation applied to mobile robot localization
  - Initial Hypothesis:
    - Static Environment
      - **Markov assumption**
      - The robot's location is the only state in the environment which systematically affects sensor readings
  - Further Hypothesis
    - Dynamic Environment
  - Instead of maintaining a single hypothesis as to where the robot is, Markov localization maintains a probability distribution over the space of all such hypothesis
  - Uses a fine-grained and metric discretization of the state space

Markov Localization

Example

Assume the robot position is one-dimensional

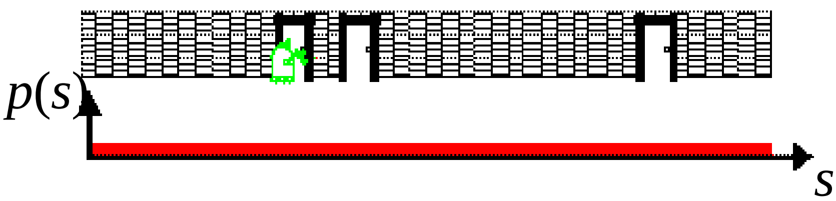

The robot is placed somewhere in the environment but it is not told its location

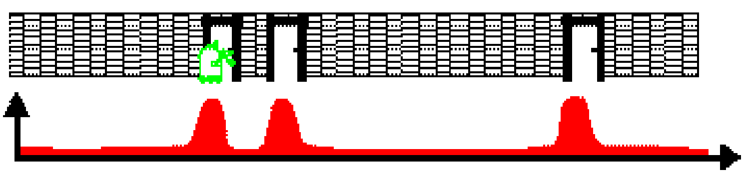

The robot queries its sensors and finds out it is next to a door

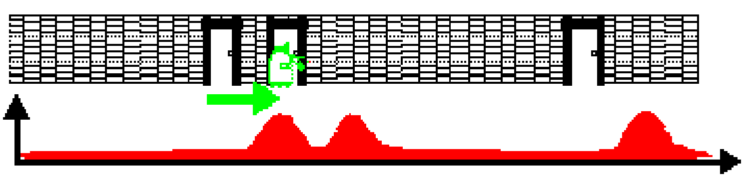

The robot moves one meter forward. To account for inherent noise in robot motion the new belief is smoother

The robot queries its sensors and again it finds itself next to a door

From S.Thrun Tutorial on Probabilistic Robotics – AAAI 2000

Markov Localization

Basic Notation

$\ell = \langle x, y, \theta \rangle$ robot location

ℓ_t robot's true location at time t

L_t Random variable that expresses the robot's location

$Bel(L_t)$ Robot's position belief at time t
Probability distribution over the space of locations

$Bel(L_t = \ell)$ Is the probability (density) that the robot assigns to the possibility that its location at time t is ℓ

The belief is updated in response to two different types of events:

- sensor readings,
- odometry data

$$d \equiv \{d_0, d_1, \dots, d_T\}$$

$$d_i = \begin{cases} a_i & \text{odometry readings} \\ s_i & \text{environment sensor readings} \end{cases}$$

Goal

Estimate the posterior distribution over L_T conditioned on all available data

$$P(L_T = \ell \mid d) = P(L_T = \ell \mid d_0, \dots, d_T)$$

Markov Localization

Markov assumption (or static world assumption)

If one knows the robot's location ℓ, future measurements are independent of past ones (and vice-versa)

$$P(d_{t+1}, d_{t+2}, \cdots \mid L_T = \ell, d_0, d_1, \cdots, d_t) = P(d_{t+1}, d_{t+2}, \cdots \mid L_T = \ell)$$

- The robot's location is the only state in the environment
- Knowing the robot state is all one needs to know about the past to predict future data.

Clearly innaccurate if the environment contains moving objects other thanthe robot itself

The basic paradigma can be extended to non-Markovian environments (see the Fox, Burgard, Thrun)

Markov Localization

- Recursive Localization

$$P(L_T = \ell \mid d)$$

  - Case 1 (**Update Phase**)
    - the most recent data item is a sensor measurement

$$d_T = s_T$$

  - Case 2 (**Prediction Phase**)
    - The most recent data item is an odometry reading

$$d_T = a_T$$

Update Phase

$$P(L_T=\ell\,|\,d)=P(L_T=\ell\,|\,d_0,d_1,\cdots,d_{T-1},s_T)$$

by Bayes' law:

$$P(a\,|\,b,c)=\frac{P(c\,|\,a,b)P(a\,|\,b)}{P(c\,|\,b)}$$

$$=\frac{P(s_T\,|\,d_0,d_1,\cdots,d_{T-1},L_T=\ell).P(L_T=\ell\,|\,d_0,d_1,\cdots,d_{T-1})}{P(s_T\,|\,d_0,d_1,\cdots,d_{T-1})}$$

by Markov assumption

$$=\frac{P(s_T\,|\,L_T=\ell).P(L_T=\ell\,|\,d_0,d_1,\cdots,d_{T-1})}{P(s_T\,|\,d_0,d_1,\cdots,d_{T-1})}$$

Does not depend on L_T

$$=\alpha_T P(s_T\,|\,L_T=\ell).P(L_T=\ell\,|\,d_0,d_1,\cdots,d_{T-1})$$

Defining

$$Bel(L_T=\ell)=P(L_T=\ell\,|\,d_0,d_1,\cdots,d_T)$$

$$Bel(L_T=\ell)=\alpha_T P(s_T\,|\,L_T=\ell)Bel(L_{T-1}=\ell)$$

independent of time

$$Bel(L_T=\ell)=\alpha_T P(s_T\,|\,\ell)Bel(L_{T-1}=\ell)$$

Incremental form

The most recently available estimate is corrected by the pdf of the expected measurement from a given location

Prediction Phase

$$P(L_T = \ell \mid d) = P(L_T = \ell \mid d_0, d_1, \cdots, d_{T-1}, a_T)$$

from odometry

Using Theorem of Total Probability

$$P(L_T = \ell \mid d) = \int P(L_T = \ell \mid d, L_{T-1} = \ell')\, P(L_{T-1} = \ell' \mid d)\, d\ell'$$

$$P(L_T = \ell \mid d, L_{T-1} = \ell') = P(L_T = \ell \mid d_0, \ldots, d_{T-1}, a_T, L_{T-1} = \ell')$$

Markov Assumption

$$= P(L_T = \ell \mid a_T, L_{T-1} = \ell')$$

$$P(L_{T-1} = \ell' \mid d) = P(L_{T-1} = \ell' \mid d_0, d_1, \ldots, d_{T-1}, a_T)$$

$$= P(L_{T-1} = \ell' \mid d_0, d_1, \ldots, d_{T-1})$$

$$P(L_T = \ell \mid d) = \int P(L_T = \ell \mid a_T, L_{T-1} = \ell')\, P(L_{T-1} = \ell' \mid d_0, \ldots, d_{T-1})\, d\ell'$$

$$Bel(L_T = \ell) = \int P(L_T = \ell \mid a_T, L_{T-1} = \ell')\, Bel(L_{T-1} = \ell')\, d\ell'$$

$$Bel(L_T = \ell) = \int P(\ell \mid a_T, \ell')\, Bel(L_{T-1} = \ell')\, d\ell'$$

Incremental form

Given the most recent odometry measurements, all possible motions are weighted from any possible posture, by the probability of starting from that posture

Markov Localization

- Recursive Localization

$$P(L_T = \ell \mid d)$$

  - Case 1 (**Update Phase**)
    - the most recent data item is a sensor measurement

$$d_T = s_T$$

$$Bel(L_T = \ell) = \alpha_T P(s_T \mid \ell) Bel(L_{T-1} = \ell)$$

  - Case 2 (**Prediction Phase**)
    - The most recent data item is an odometry reading

$$d_T = a_T$$

$$Bel(L_T = \ell) = \int P(\ell \mid a_T, \ell') \, Bel(L_{T-1} = \ell') \, d\ell'$$

$P(\ell \mid a, \ell')$ —— motion model

$P(s \mid L)$ —— perceptual model

Markov Localization

% Initializes Belief

for each ***l*** do

$$Bel(L_0 = l) \leftarrow P(L_0 = l)$$

end for

forever do

if new ***sensor reading*** received do

$$\alpha_T \leftarrow 0$$

% Perception model

for each ***l*** do

$$\hat{B}el(L_T = l) \leftarrow P(s_T | l) \cdot Bel(L_{T-1} = l)$$

$$\alpha_T \leftarrow \alpha_T + \hat{B}el(L_T = l)$$

end for

% Normalize the Belief

for each ***l*** do

$$Bel(L_T = l) \leftarrow \alpha_T^{-1} \cdot \hat{B}el(L_{T-1} = l)$$

end for

end if

if new ***odometry reading*** received do

% Robot motion model

for each ***l*** do

$$Bel(L_T = l) \leftarrow \int P(l | l', \alpha_T) \cdot \hat{B}el(L_{T-1} = l') dl'$$

end for

end if

end forever

Questions:

- How to represent Bel(L) to reduce the computational time ?
- How to evaluate $P(\ell | a, \ell')$, $P(s | \ell)$

Markov Localization

- **Topological** (landmark-based, state space organized according to the topological structure of the environment)
- **Grid-Based** (the world is divided in cells of fixed size; resolution and precision of state estimation are fixed beforehand)
  - The latter suffers from computational overhead
- **Monte-Carlo methods**:
  - the probability density function is represented by samples randomly drawn from it
  - it is also able to represent multi-modal distributions, and thus localize the robot *globally*
  - considerably reduces the amount of memory required and can integrate measurements at a higher rate
  - state is not discretized and the method is more accurate than the grid-based methods
  - easy to implement

Monte-Carlo Localization

From
Probabilistic Algorithms in Robotics
Sebastian Thrun
Technical Report CMU CS-00-126

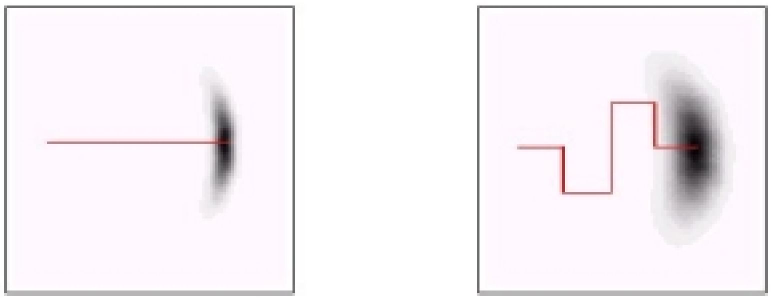

Figure 1: Probabilistic generalization of mobile robot kinematics: Each dark line illustrates a commanded robot path, and the grayly shaded shows the posterior distribution of the robot's pose. The darker an area, the more likely it is. The path in the left diagram is 40 meters and the one on the right is 80 meters long.

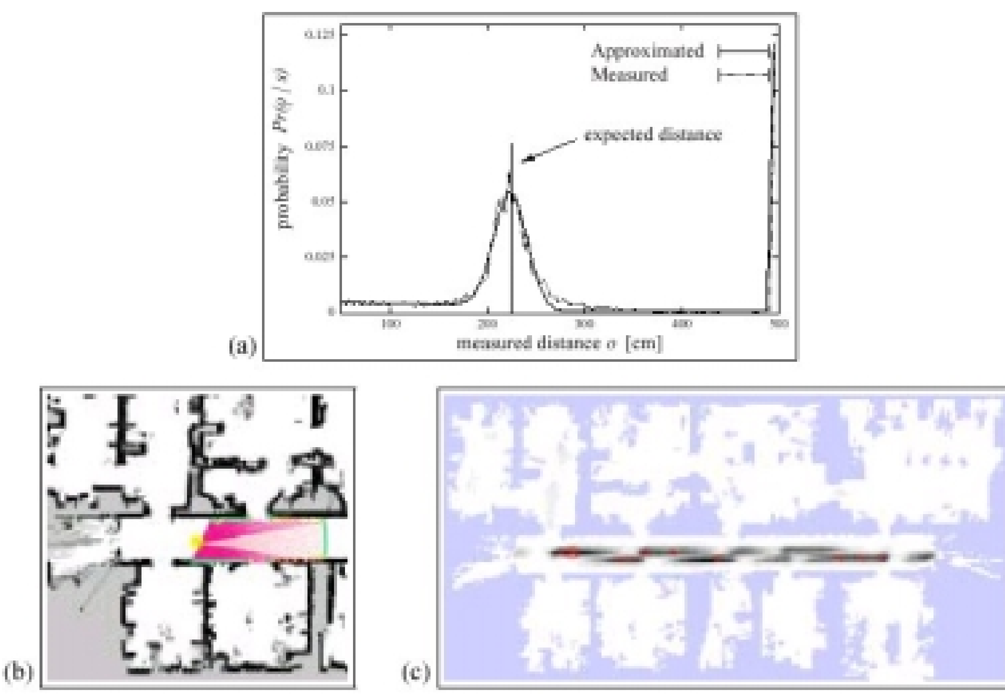

Figure 2: Probabilistic sensor model for laser range finders: (a) The density $p(o \mid s, m)$ relates the actual, measured distance of a sensor beam to its expected distance computed by ray tracing, under the assumption that the robot's pose is s. A comparison of actual data and our (learned) mixture model shows good correspondence. Diagram (b) shows a specific laser range scan o, for which diagram (c) plots the density $p(o \mid s, m)$ for different locations in the map.

Monte Carlo Localization

From
Probabilistic Algorithms in Robotics
Sebastian Thrun
Technical Report CMU CS-00-126

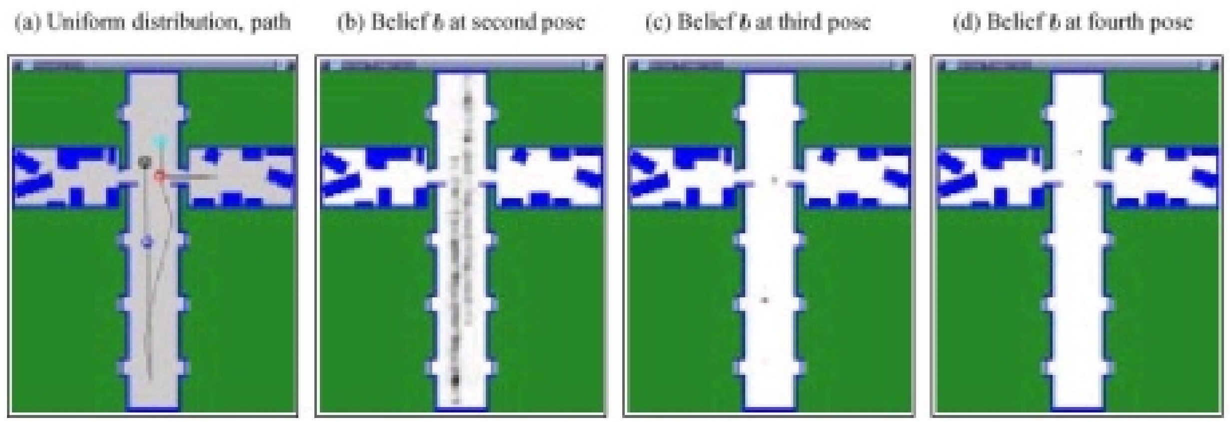

Figure 3: Example of grid-based Markov localization in a symmetric environment. (a) robot path, highlighting four robot poses, (b) to (d): belief at 2nd, 3rd, and 4th pose, respectively.

Figure 4: Global localization of a mobile robot using MCL.

References

- Dieter Fox, Wolfram Burgard, Frank Dellaert, Sebastian Thrun, “Monte Carlo Localization: Efficient Position Estimation for Mobile Robots”, Proc. 16th National Conference on Artificial Intelligence, AAAI’99, July 1999
- Dieter Fox, Wolfram Burgard, Sebastian Thrun, “Markov Localization for Mobile Robots in Dynamic Environments”, J. of Artificial Intelligence Research 11 (1999) 391-427
- Sebastian Thrun, “Probabilistic Algorithms in Robotics”, Technical Report CMU-CS-00-126, School of Computer Science, Carnegie Mellon University, Pittsburgh, USA, 2000
- Sebastian Thrun, “Probabilistic Robotics: Tutorial presented in AAAI 200” – available through S.Thrun Web page http://www-2.cs.cmu.edu/~thrun/ See also movies illustrating the methods.