

Bayesian Filtering

Dieter Fox

Probabilistic Robotics

Key idea: Explicit representation of uncertainty

(using the calculus of probability theory)

- Perception = state estimation
- Control = utility optimization

Bayes Filters: Framework

- **Given:**

- Stream of observations z and action data u :

$$d_t = \{u_1, z_2, \dots, u_{t-1}, z_t\}$$

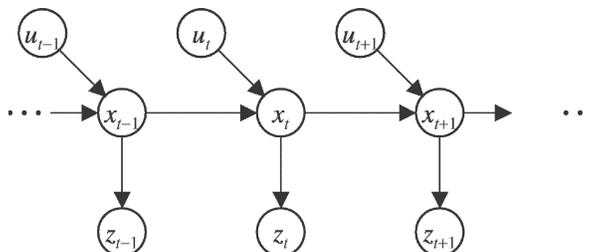
- Sensor model $P(z|x)$.
- Action model $P(x|u, x')$.
- Prior probability of the system state $P(x)$.

- **Wanted:**

- Estimate of the state X of a dynamical system.
- The posterior of the state is also called **Belief**:

$$Bel(x_t) = P(x_t | u_1, z_2, \dots, u_{t-1}, z_t)$$

Markov Assumption



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

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

Underlying Assumptions

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

Bayes Filters

z = observation
 u = action
 x = state

$$Bel(x_t) = P(x_t | u_1, z_2, \dots, u_{t-1}, z_t)$$

Bayes $= \eta P(z_t | x_t, u_1, z_2, \dots, u_{t-1}) P(x_t | u_1, z_2, \dots, u_{t-1})$

Markov $= \eta P(z_t | x_t) P(x_t | u_1, z_2, \dots, u_{t-1})$

Total prob. $= \eta P(z_t | x_t) \int P(x_t | u_1, z_2, \dots, u_{t-1}, x_{t-1})$
 $P(x_{t-1} | u_1, z_2, \dots, u_{t-1}) dx_{t-1}$

Markov $= \eta P(z_t | x_t) \int P(x_t | u_{t-1}, x_{t-1}) P(x_{t-1} | u_1, z_2, \dots, u_{t-1}) dx_{t-1}$

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

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)

Localization

“Using sensory information to locate the robot in its environment is the most fundamental problem to providing a mobile robot with autonomous capabilities.” [Cox '91]

- **Given**

- Map of the environment.
- Sequence of sensor measurements.

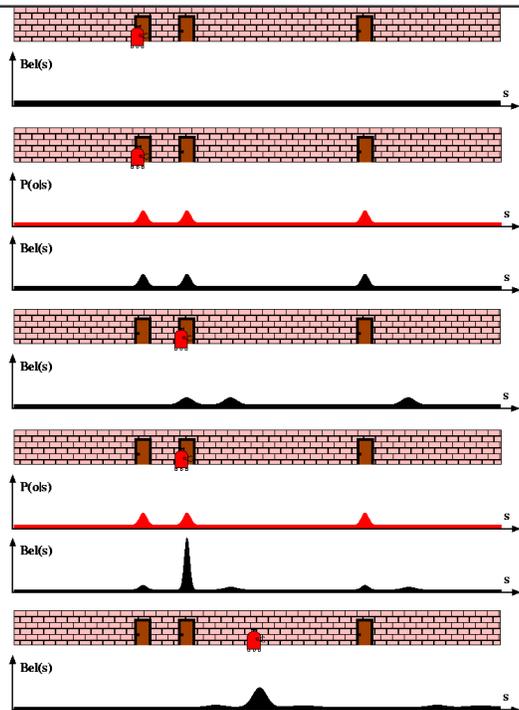
- **Wanted**

- Estimate of the robot's position.

- **Problem classes**

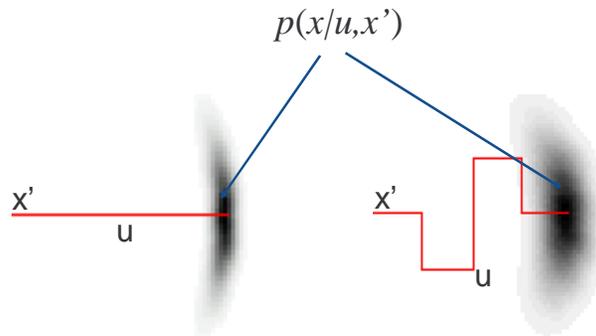
- Position tracking
- Global localization
- Kidnapped robot problem (recovery)

Bayes Filters for Robot Localization



Probabilistic Kinematics

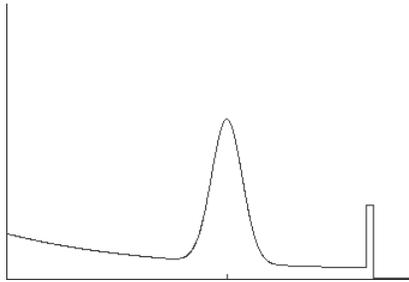
- Odometry information is inherently noisy.



Proximity Measurement

- Measurement can be caused by ...
 - a known obstacle.
 - cross-talk.
 - an unexpected obstacle (people, furniture, ...).
 - missing all obstacles (total reflection, glass, ...).
- Noise is due to uncertainty ...
 - in measuring distance to known obstacle.
 - in position of known obstacles.
 - in position of additional obstacles.
 - whether obstacle is missed.

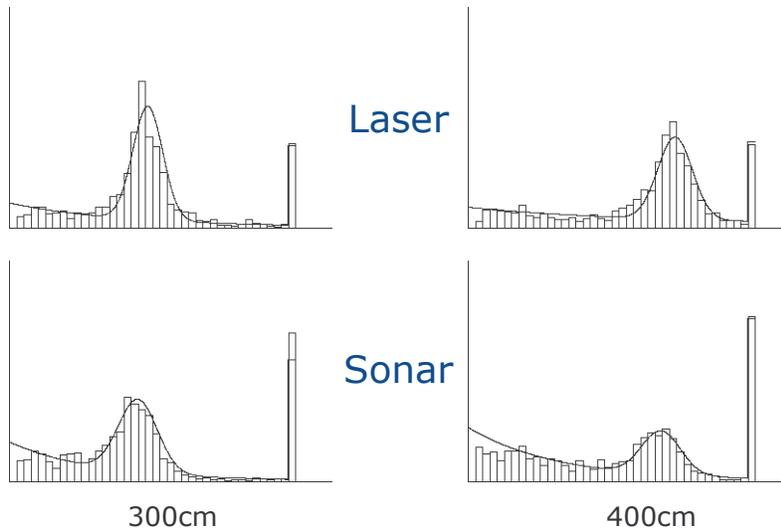
Mixture Density



$$P(z | x, m) = \begin{pmatrix} \alpha_{\text{hit}} \\ \alpha_{\text{unexp}} \\ \alpha_{\text{max}} \\ \alpha_{\text{rand}} \end{pmatrix}^T \cdot \begin{pmatrix} P_{\text{hit}}(z | x, m) \\ P_{\text{unexp}}(z | x, m) \\ P_{\text{max}}(z | x, m) \\ P_{\text{rand}}(z | x, m) \end{pmatrix}$$

Model parameters learned via EM

Approximation Results



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
- approximately linear models
- position tracking

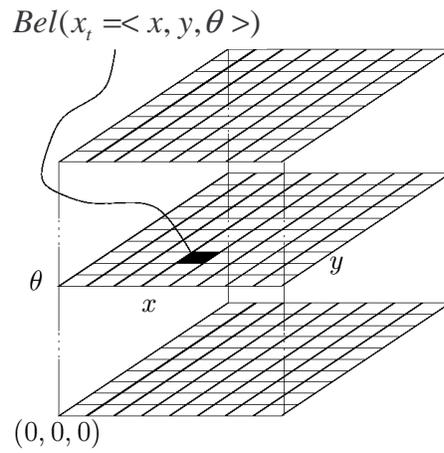
Robotics

Multi-hypothesis ('00)

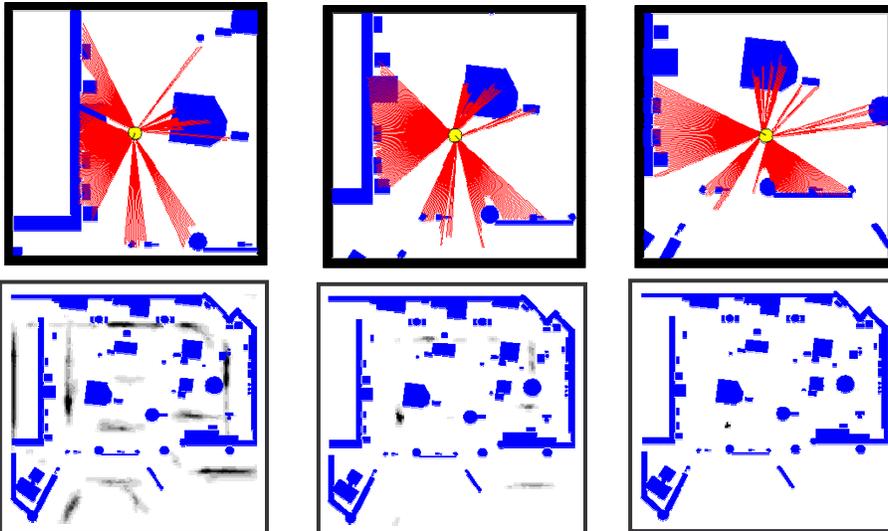
- multiple Kalman filters
- global localization, recovery

Discrete Grid Filters

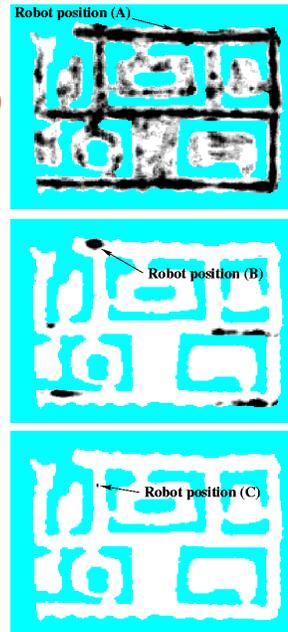
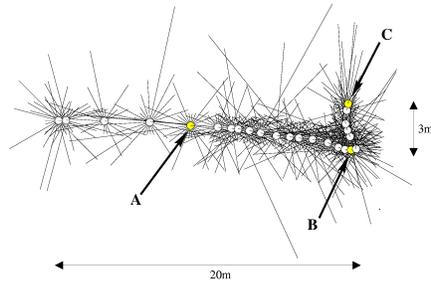
Piecewise Constant Representation



Grid-based Localization

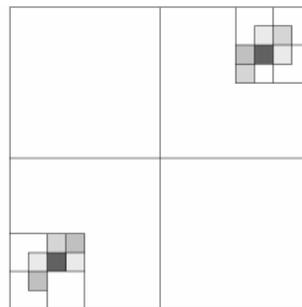
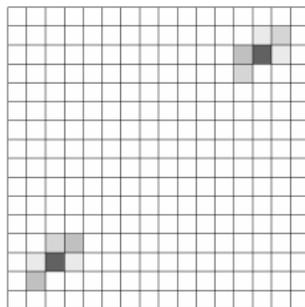


Sonars and Occupancy Grid Map



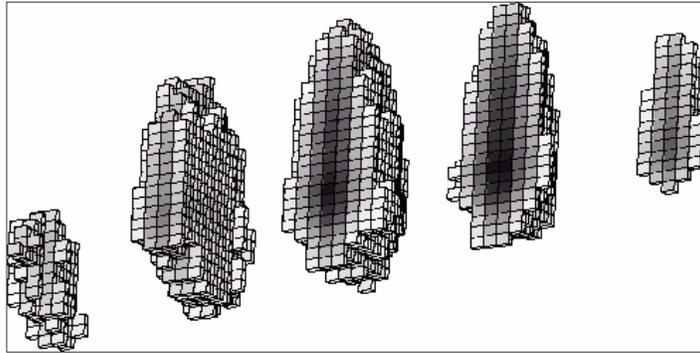
Tree-based Representation

Idea: Represent density using a variant of Octrees



Tree-based Representations

- Efficient in space and time
- Multi-resolution

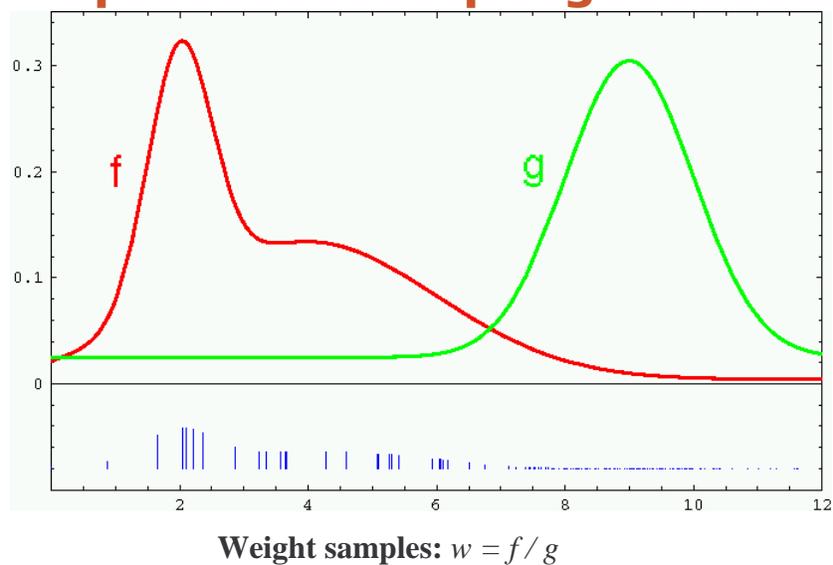


Particle Filters

Particle Filters

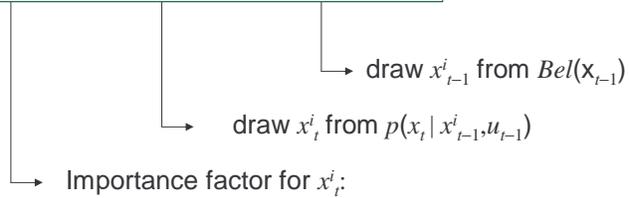
- § Represent belief by random **samples**
- § Estimation of **non-Gaussian, nonlinear** processes
- § Monte Carlo filter, Survival of the fittest, Condensation, Bootstrap filter, Particle filter
- § Filtering: [Rubin, 88], [Gordon et al., 93], [Kitagawa 96]
- § Computer vision: [Isard and Blake 96, 98]
- § Dynamic Bayesian Networks: [Kanazawa et al., 95]d

Importance Sampling



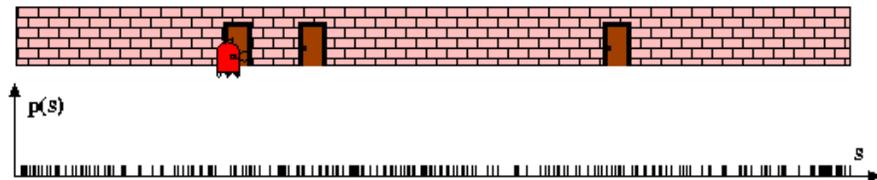
Particle Filter Algorithm

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



$$\begin{aligned}
 w_t^i &= \frac{\text{target distribution}}{\text{proposal distribution}} \\
 &= \frac{\eta p(z_t | x_t) p(x_t | x_{t-1}^i, u_{t-1}) Bel(x_{t-1}^i)}{p(x_t | x_{t-1}^i, u_{t-1}) Bel(x_{t-1}^i)} \\
 &\propto p(z_t | x_t)
 \end{aligned}$$

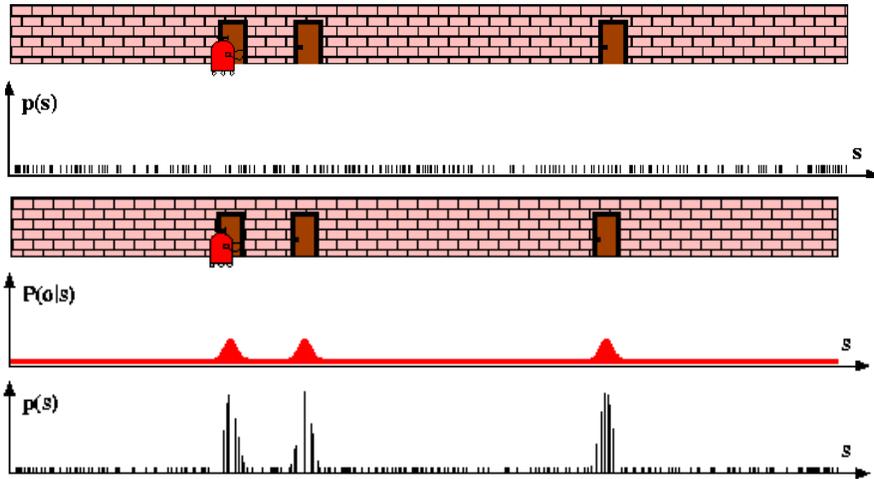
Particle Filters



Sensor Information: Importance Sampling

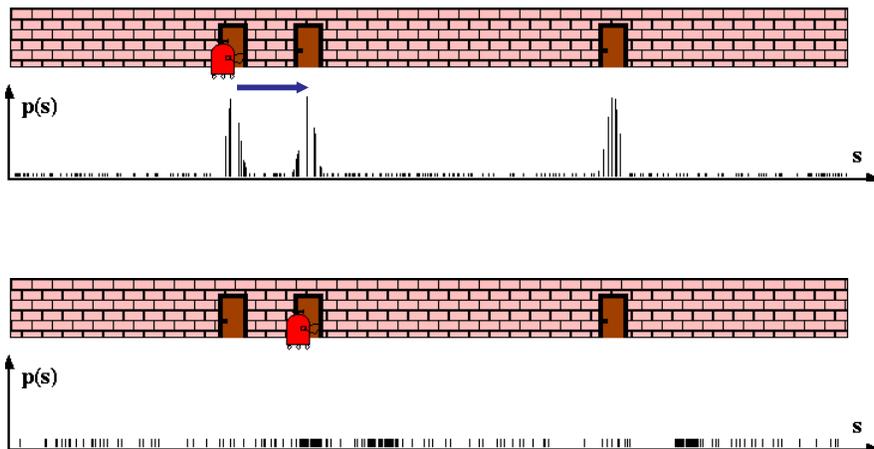
$$Bel(x) \leftarrow \alpha p(z|x) Bel^-(x)$$

$$w \leftarrow \frac{\alpha p(z|x) Bel^-(x)}{Bel^-(x)} = \alpha p(z|x)$$



Robot Motion

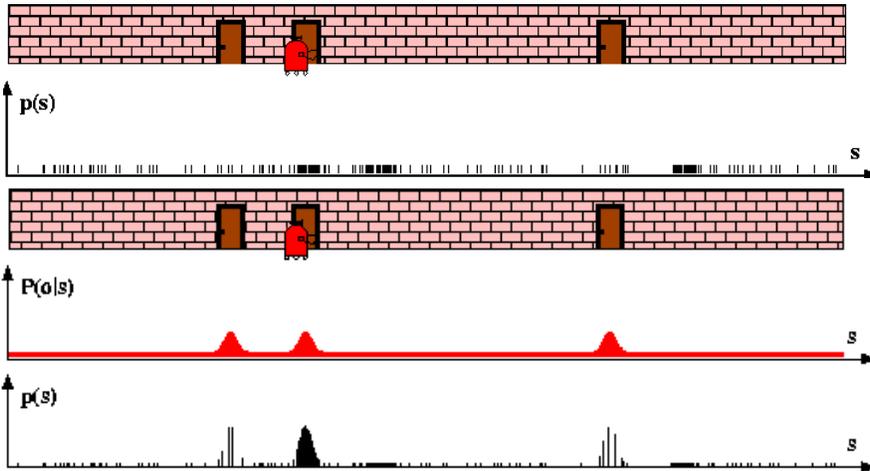
$$Bel^-(x) \leftarrow \int p(x|u, x') Bel(x') dx'$$



Sensor Information: Importance Sampling

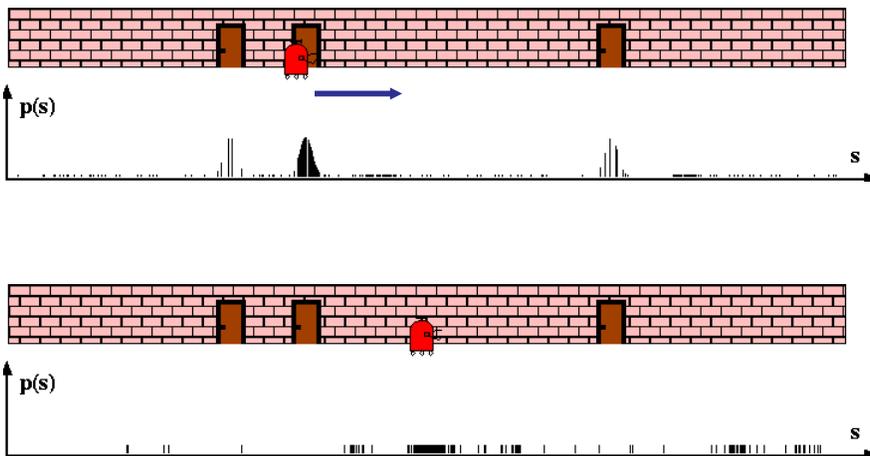
$$Bel(x) \leftarrow \alpha p(z|x) Bel^-(x)$$

$$w \leftarrow \frac{\alpha p(z|x) Bel^-(x)}{Bel^-(x)} = \alpha p(z|x)$$

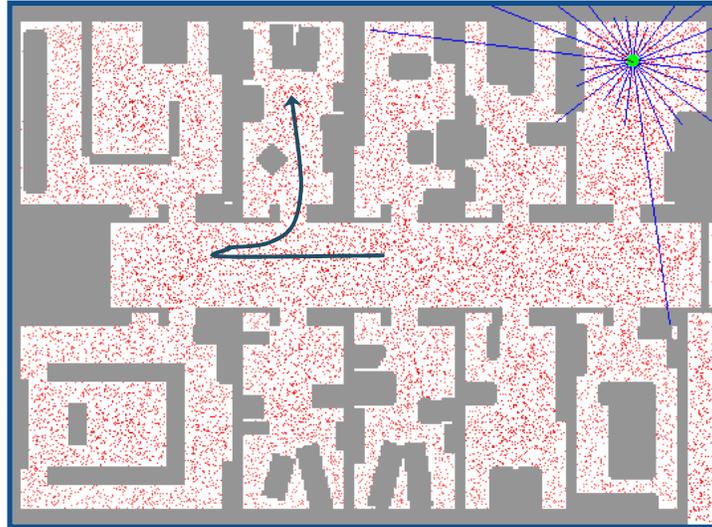


Robot Motion

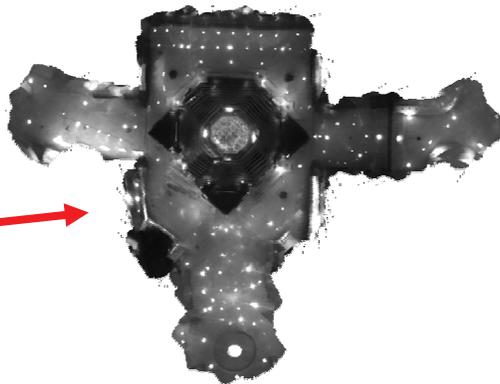
$$Bel^-(x) \leftarrow \int p(x|u, x') Bel(x') dx'$$



Sample-based Localization (sonar)

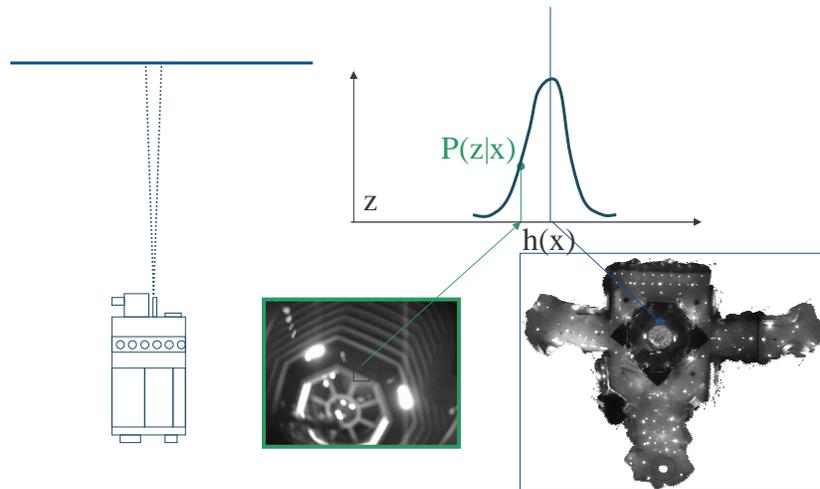


Using Ceiling Maps for Localization

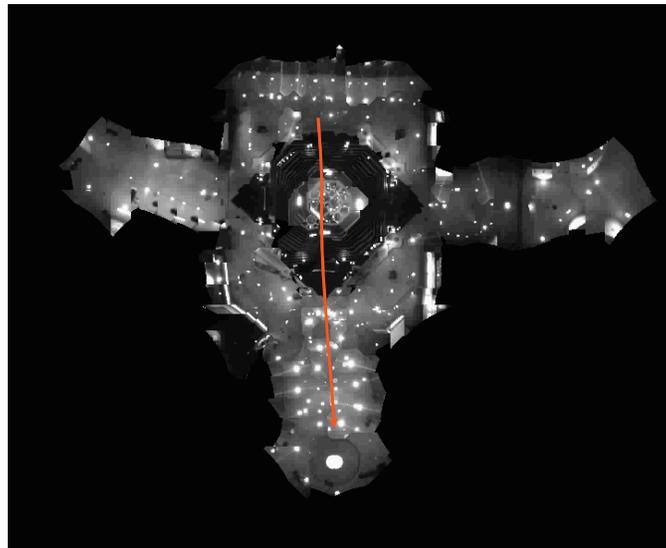


[Dellaert et al. 99]

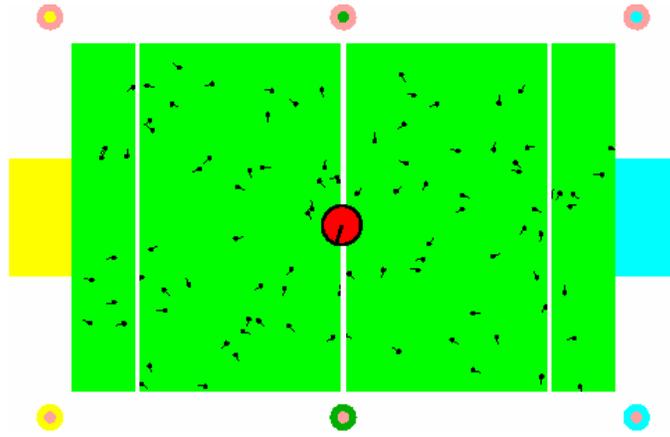
Vision-based Localization



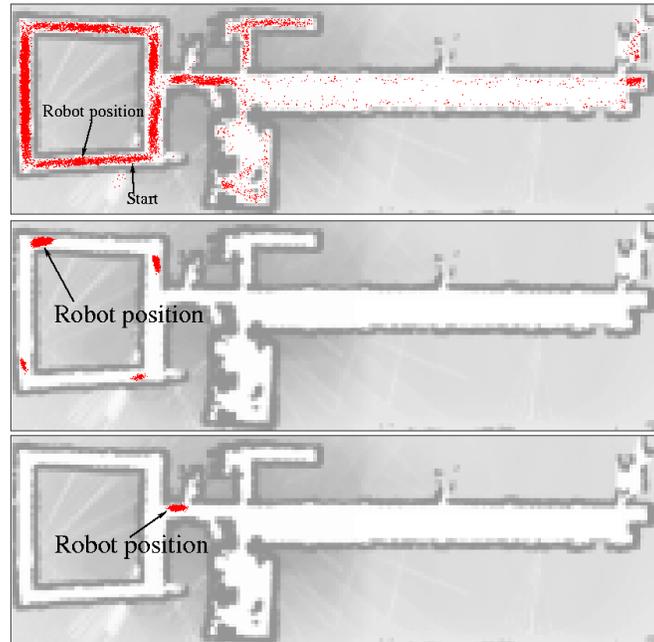
Global Localization Using Vision



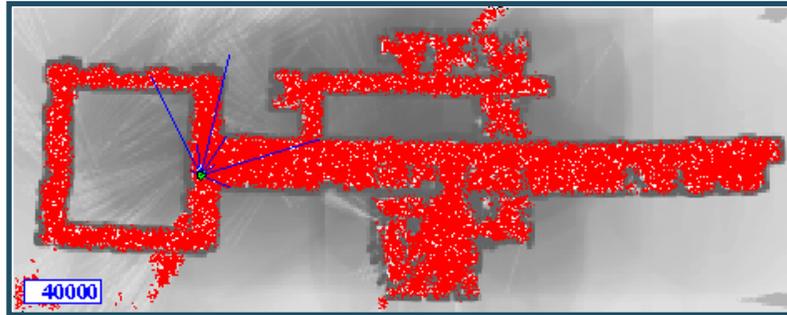
Localization for AIBO robots



Adaptive Sampling

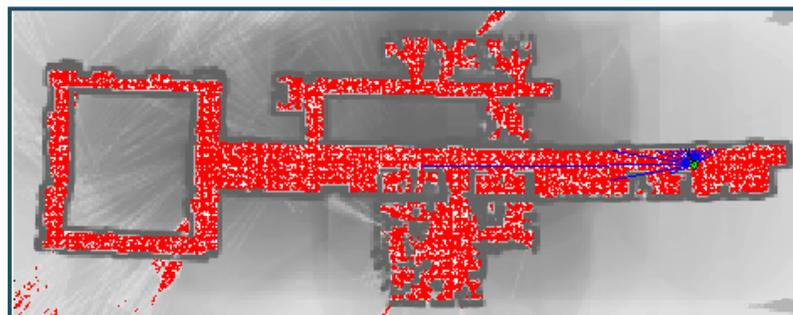


KLD-Sampling: Sonar



[NIPS-01, IJRR-03]

KLD-Sampling: Laser



Kalman Filters

Bayes Filter Reminder

- Prediction

$$\overline{bel}(x_t) = \int p(x_t | u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1}$$

- Correction

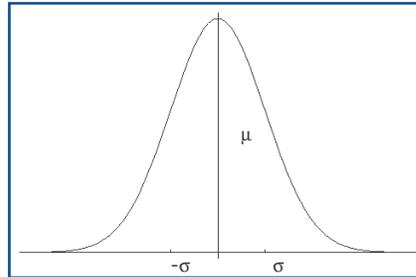
$$bel(x_t) = \eta p(z_t | x_t) \overline{bel}(x_t)$$

Gaussians

$$p(x) \sim N(\mu, \sigma^2):$$

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}}$$

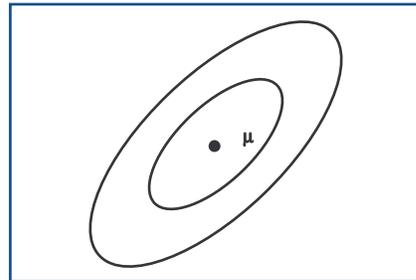
Univariate



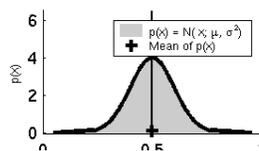
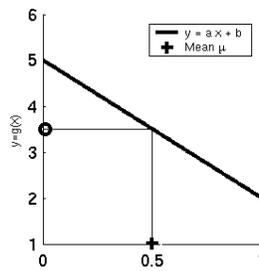
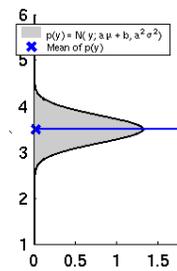
$$p(\mathbf{x}) \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}):$$

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{x}-\boldsymbol{\mu})}$$

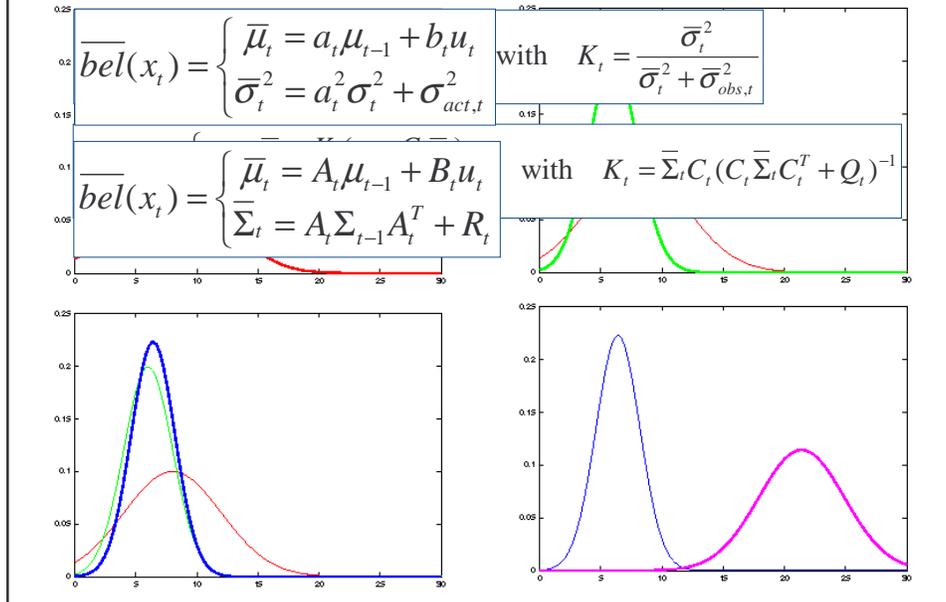
Multivariate



Gaussians and Linear Functions



Kalman Filter Updates in 1D



Kalman Filter Algorithm

1. Algorithm **Kalman_filter**(μ_{t-1} , Σ_{t-1} , u_t , z_t):
2. Prediction:
3. $\overline{\mu}_t = A_t \mu_{t-1} + B_t u_t$
4. $\overline{\Sigma}_t = A_t \Sigma_{t-1} A_t^T + R_t$
5. Correction:
6. $K_t = \overline{\Sigma}_t C_t (C_t \overline{\Sigma}_t C_t^T + Q_t)^{-1}$
7. $\mu_t = \overline{\mu}_t + K_t (z_t - C_t \overline{\mu}_t)$
8. $\Sigma_t = (I - K_t C_t) \overline{\Sigma}_t$
9. Return μ_t , Σ_t

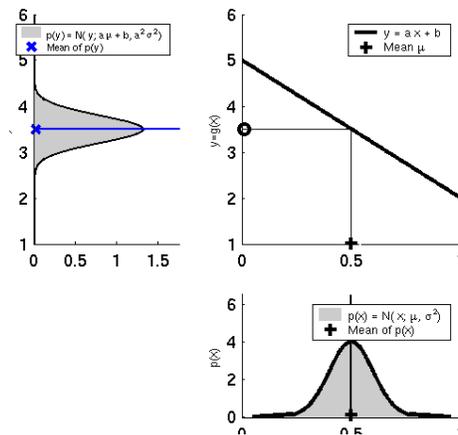
Nonlinear Dynamic Systems

- Most realistic robotic problems involve nonlinear functions

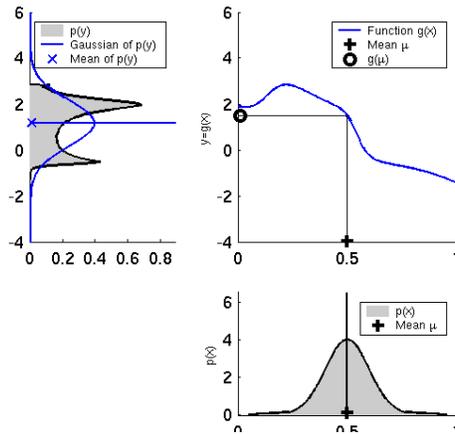
$$x_t = g(u_t, x_{t-1})$$

$$z_t = h(x_t)$$

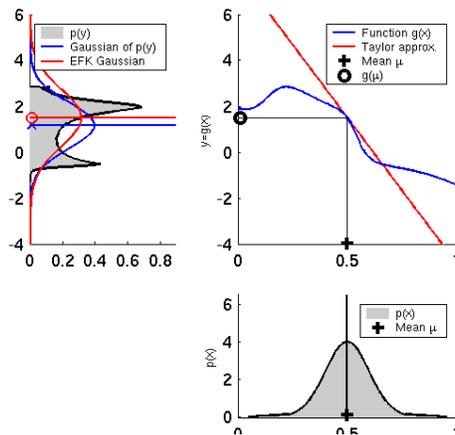
Linearity Assumption Revisited



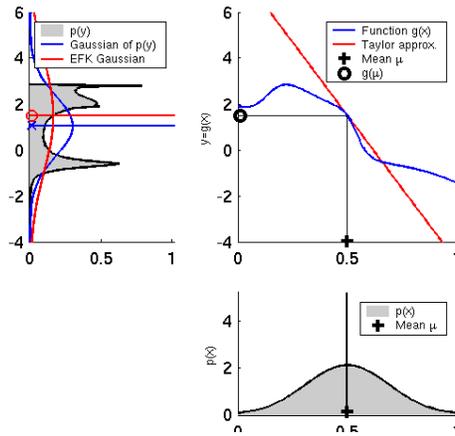
Non-linear Function



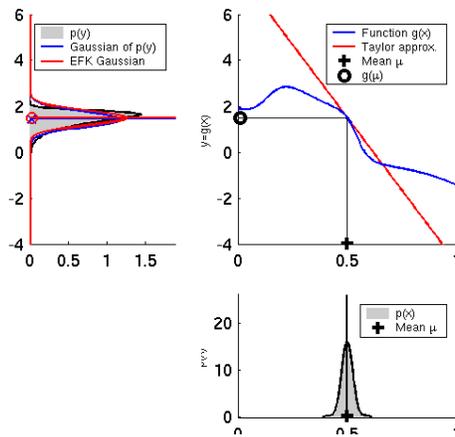
EKF Linearization (1)



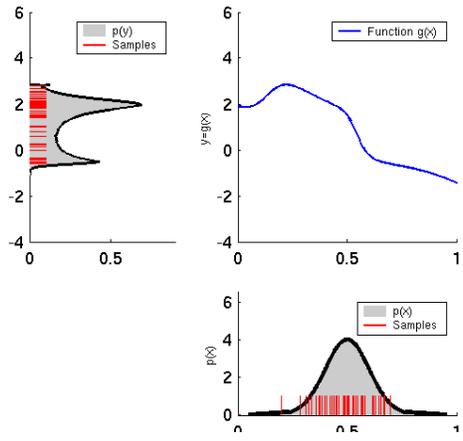
EKF Linearization (2)



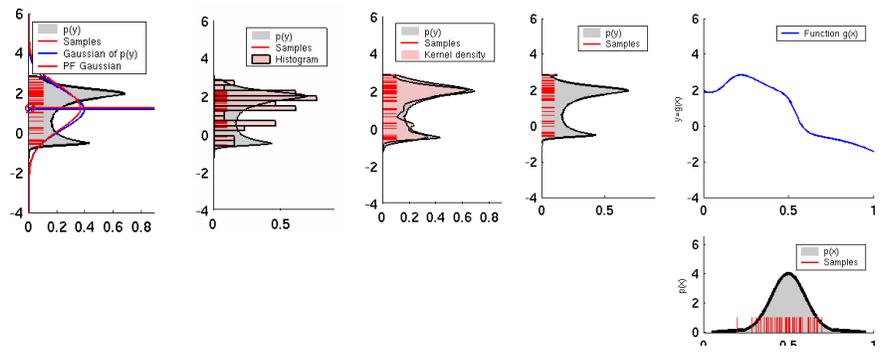
EKF Linearization (3)



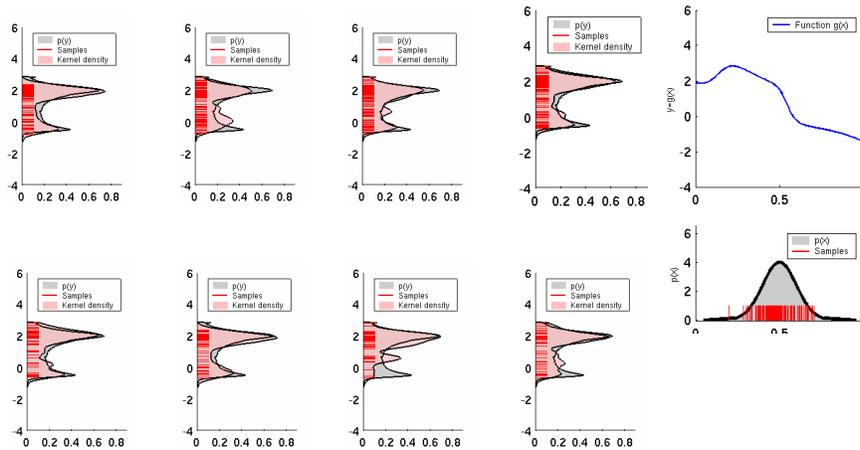
Particle Filter Projection



Density Extraction



Sampling Variance



EKF Algorithm

1. **Extended_Kalman_filter**(μ_{t-1} , Σ_{t-1} , u_t , z_t):

2. Prediction:

3. $\bar{\mu}_t = g(u_t, \mu_{t-1})$

$\bar{\mu}_t = A_t \mu_{t-1} + B_t u_t$

4. $\bar{\Sigma}_t = G_t \Sigma_{t-1} G_t^T + R_t$

$\bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^T + R_t$

5. Correction:

6. $K_t = \bar{\Sigma}_t H_t (H_t \bar{\Sigma}_t H_t^T + Q_t)^{-1}$

$K_t = \bar{\Sigma}_t C_t (C_t \bar{\Sigma}_t C_t^T + Q_t)^{-1}$

7. $\mu_t = \bar{\mu}_t + K_t (z_t - h(\bar{\mu}_t))$

$\mu_t = \bar{\mu}_t + K_t (z_t - C_t \bar{\mu}_t)$

8. $\Sigma_t = (I - K_t H_t) \bar{\Sigma}_t$

$\Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$

9. Return μ_t , Σ_t

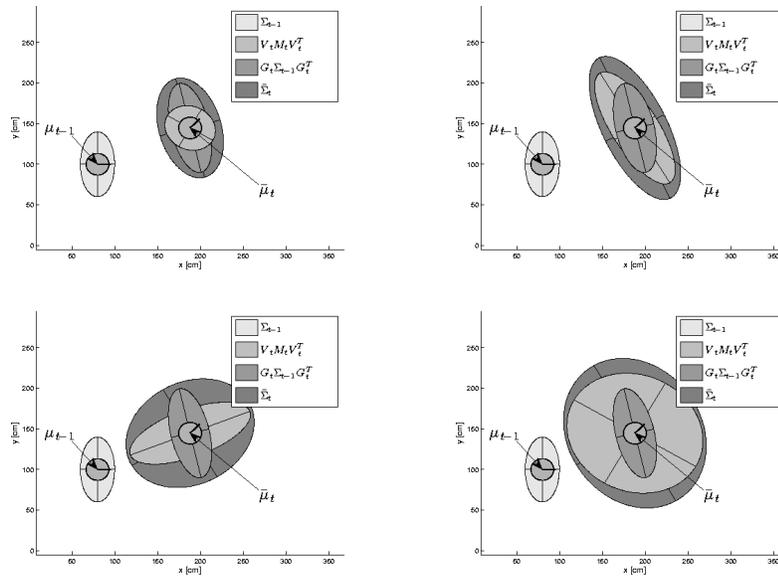
$H_t = \frac{\partial h(\bar{\mu}_t)}{\partial x_t}$

$G_t = \frac{\partial g(u_t, \mu_{t-1})}{\partial x_{t-1}}$

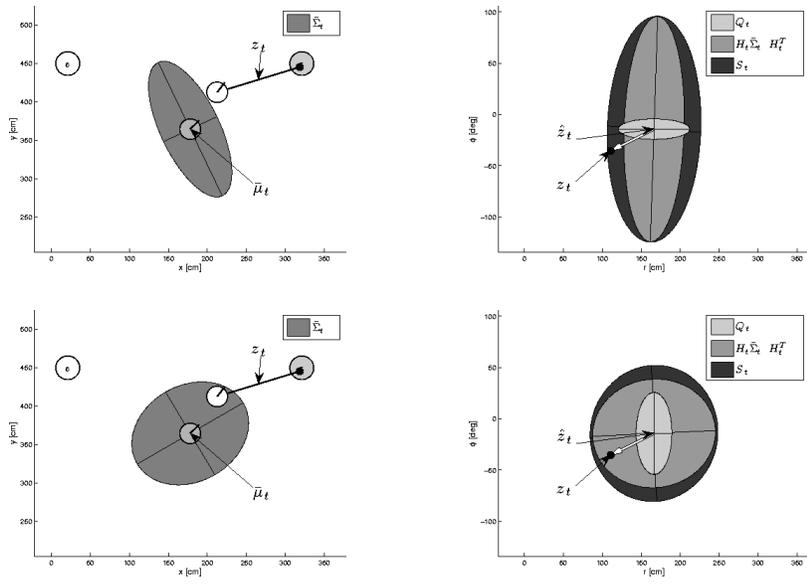
Landmark-based Localization



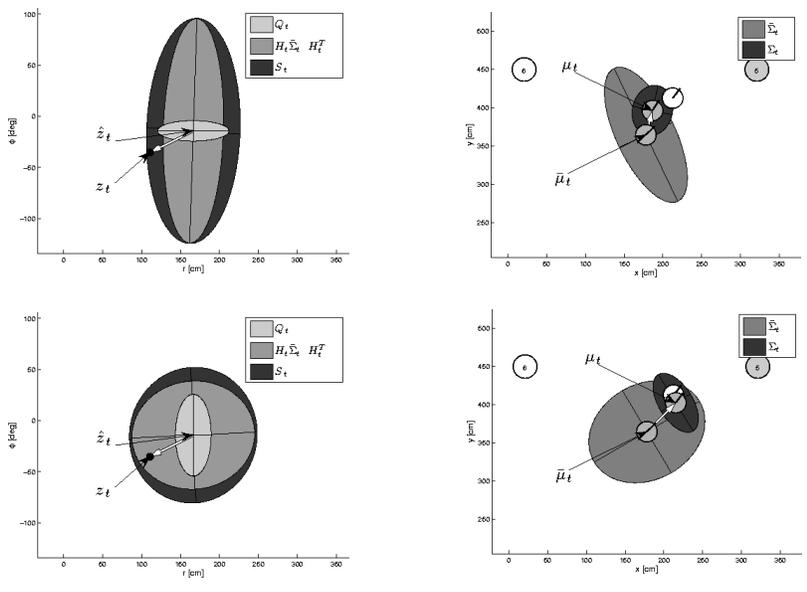
EKF Prediction Step



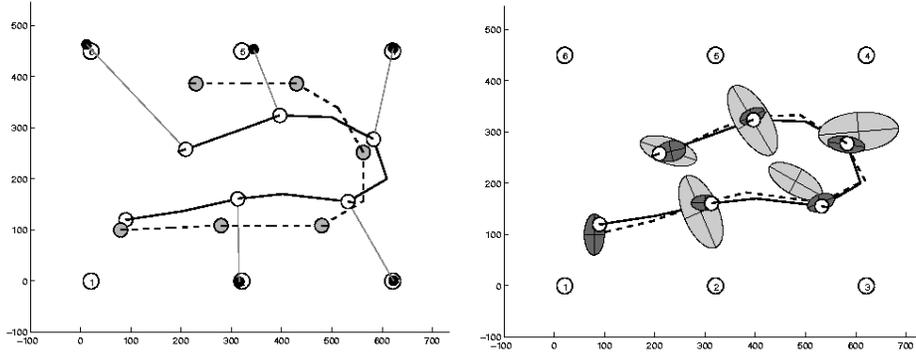
EKF Observation Prediction Step



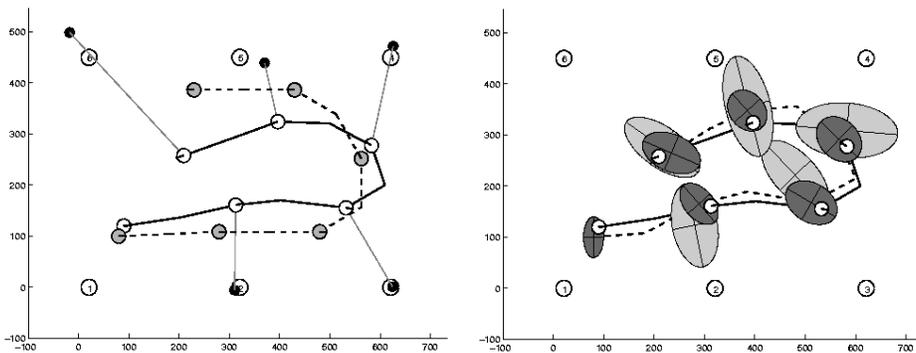
EKF Correction Step



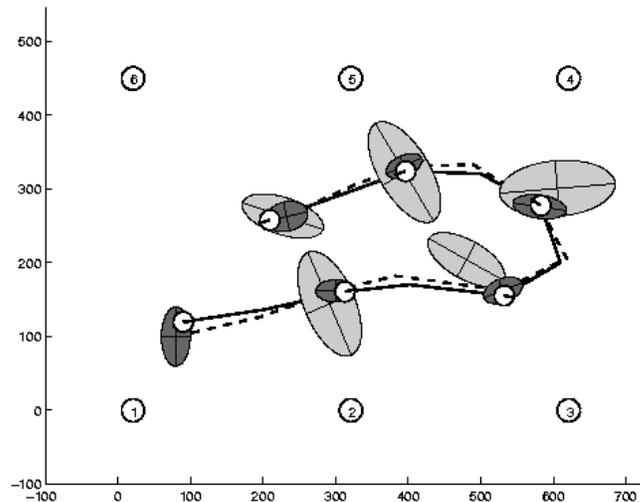
Estimation Sequence (1)



Estimation Sequence (2)



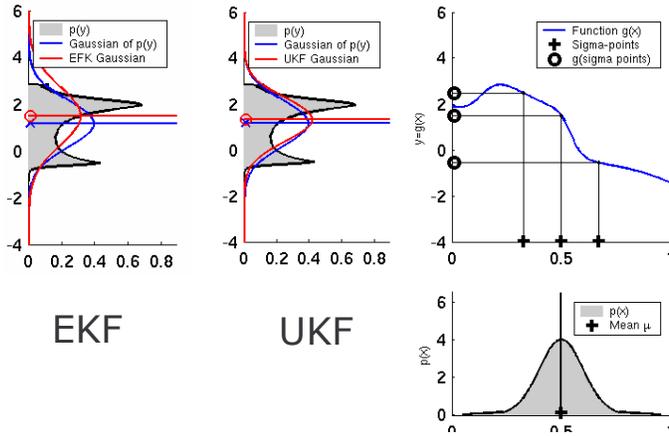
Comparison to GroundTruth



EKF Summary

- **Highly efficient:** Polynomial in measurement dimensionality k and state dimensionality n :
$$O(k^{2.376} + n^2)$$
- **Not optimal!**
- Can **diverge** if nonlinearities are large!
- Works surprisingly well even when all assumptions are violated!

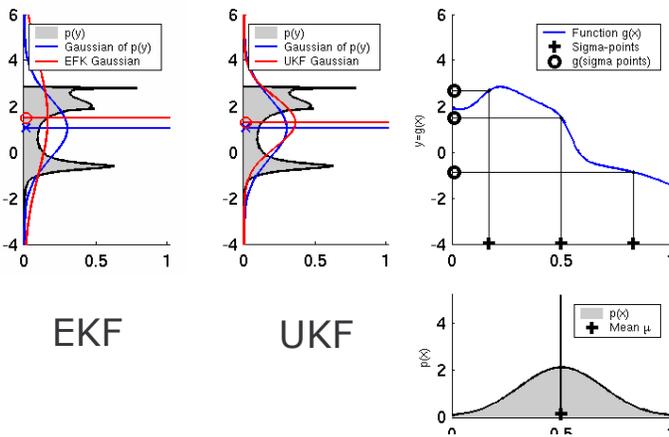
Linearization via Unscented Transform



EKF

UKF

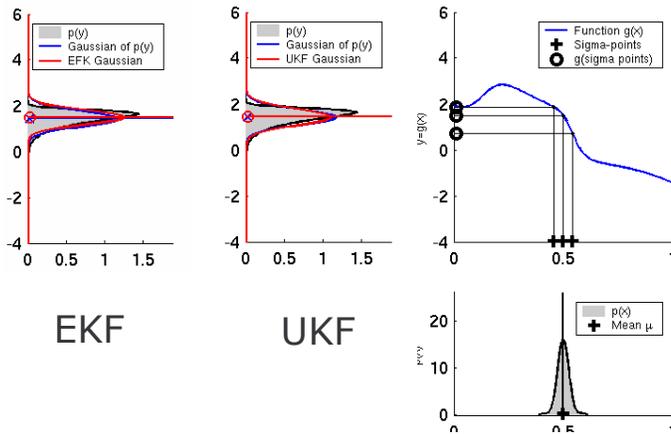
UKF Sigma-Point Estimate (2)



EKF

UKF

UKF Sigma-Point Estimate (3)



Unscented Transform

Sigma points

Weights

$$\chi^0 = \mu$$

$$w_m^0 = \frac{\lambda}{n + \lambda} \quad w_c^0 = \frac{\lambda}{n + \lambda} + (1 - \alpha^2 + \beta)$$

$$\chi^i = \mu \pm \left(\sqrt{(n + \lambda) \Sigma} \right)_i$$

$$w_m^i = w_c^i = \frac{1}{2(n + \lambda)} \quad \text{for } i = 1, \dots, 2n$$

Pass sigma points through nonlinear function

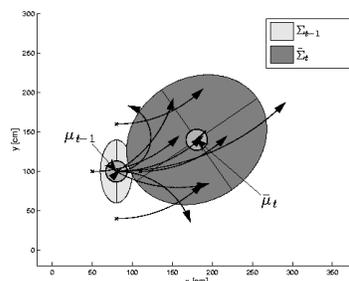
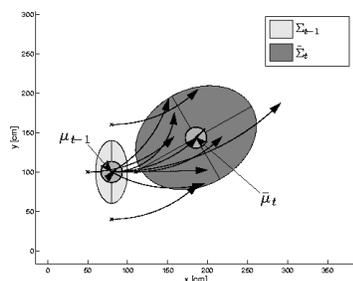
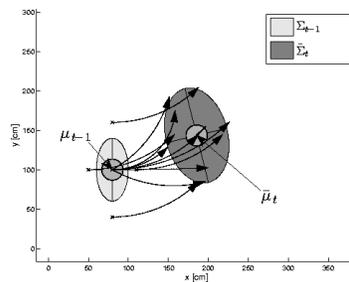
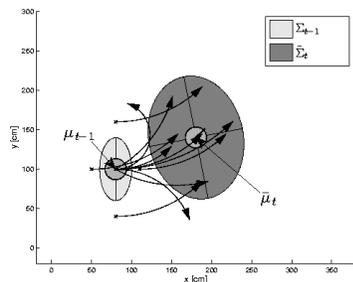
$$\psi^i = g(\chi^i)$$

Recover mean and covariance

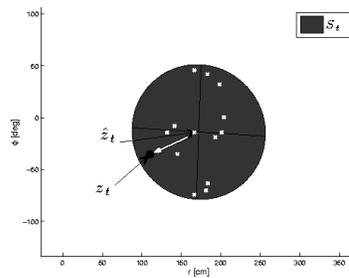
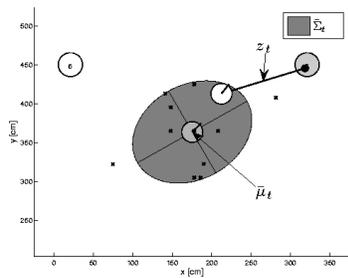
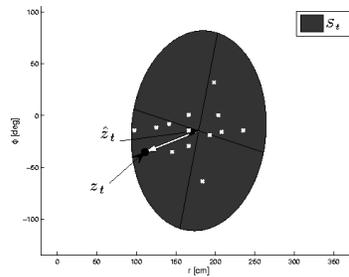
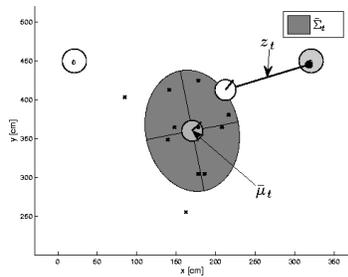
$$\mu' = \sum_{i=0}^{2n} w_m^i \psi^i$$

$$\Sigma' = \sum_{i=0}^{2n} w_c^i (\psi^i - \mu)(\psi^i - \mu)^T$$

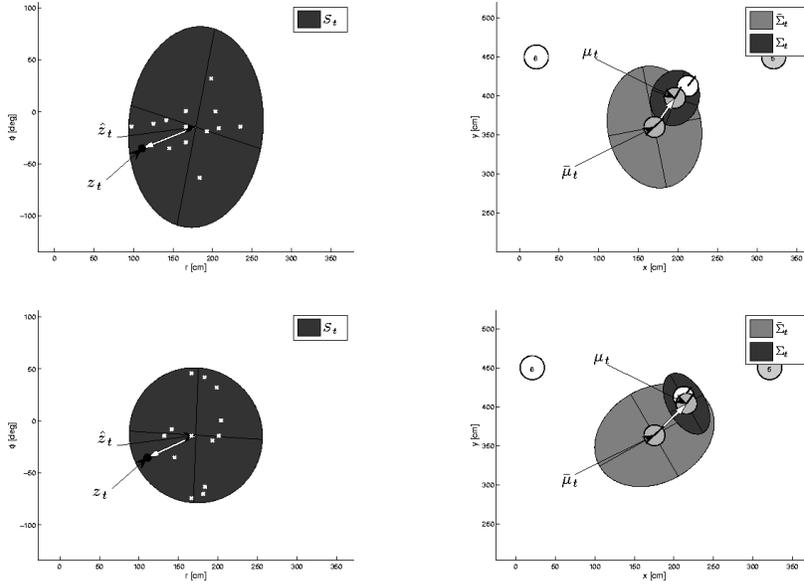
UKF Prediction Step



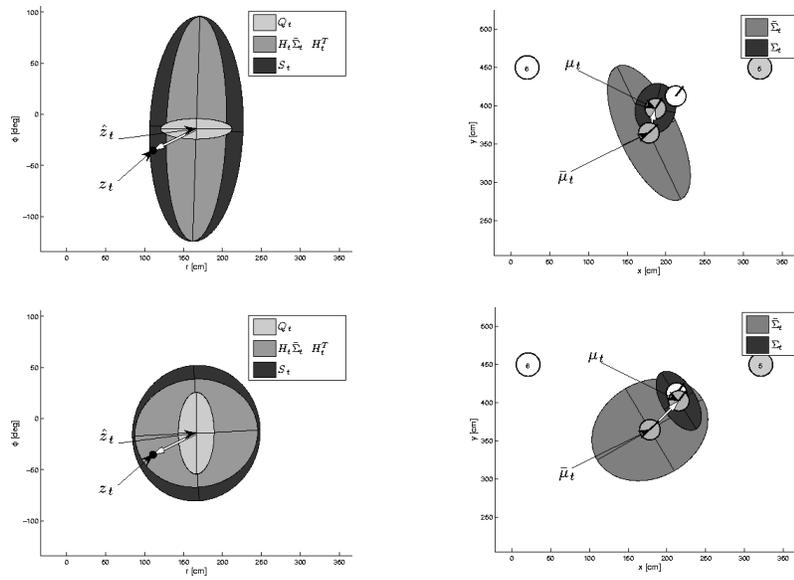
UKF Observation Prediction Step



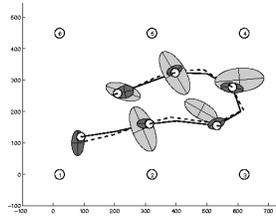
UKF Correction Step



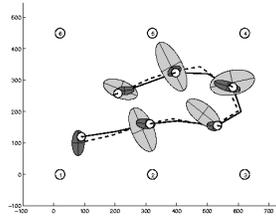
EKF Correction Step



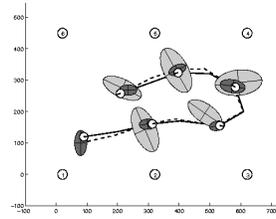
Estimation Sequence



EKF

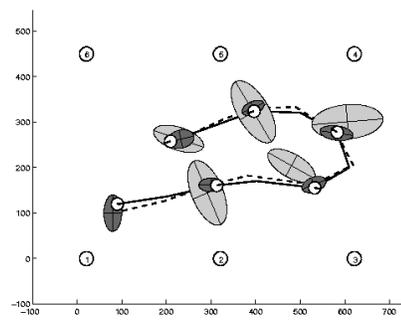


PF

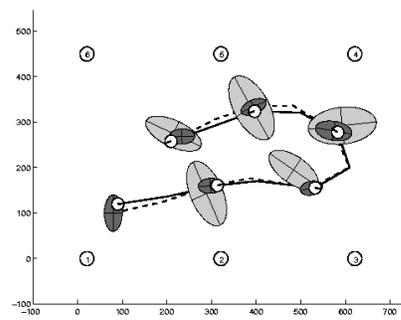


UKF

Estimation Sequence

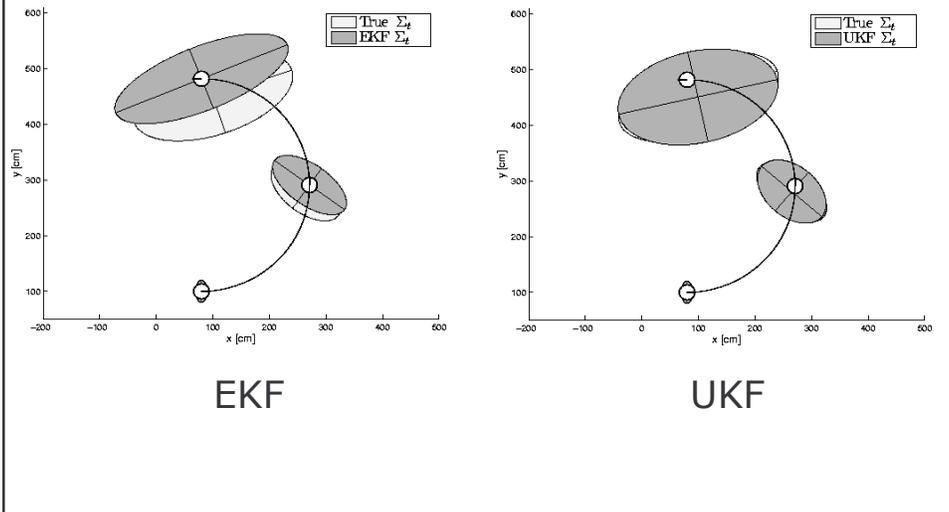


EKF



UKF

Prediction Quality



UKF Summary

- **Highly efficient:** Same complexity as EKF, with a constant factor slower in typical practical applications
- **Better linearization than EKF:** Accurate in first two terms of Taylor expansion (EKF only first term)
- **Derivative-free:** No Jacobians needed
- **Still not optimal!**

Rao-Blackwellised Particle Filters

Ball Tracking in RoboCup



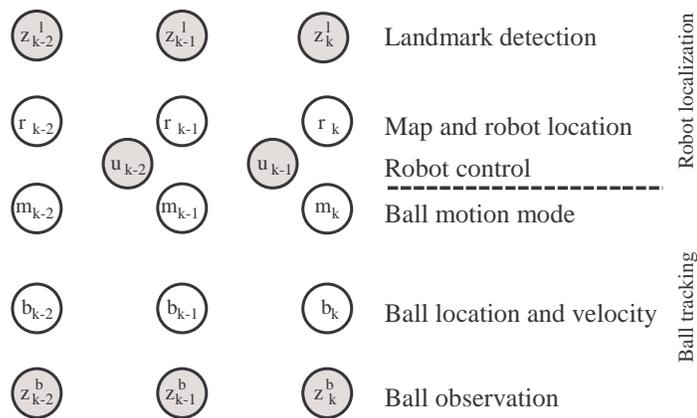
- § Extremely noisy (nonlinear) motion of observer
- § Inaccurate sensing, limited processing power
- § Interactions between target and

Goal: Unified framework for modeling the ball and its interactions. t

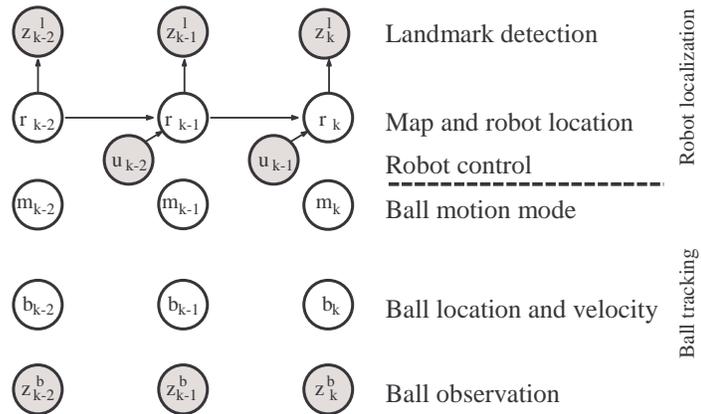
Tracking Techniques

- ⌘ Kalman Filter
 - ⌘ Highly efficient, robust (even for nonlinear)
 - ⌘ Uni-modal, limited handling of nonlinearities
- ⌘ Particle Filter
 - ⌘ Less efficient, highly robust
 - ⌘ Multi-modal, nonlinear, non-Gaussian
- ⌘ Rao-Blackwellised Particle Filter, MHT
 - ⌘ Combines PF with KF
 - ⌘ Multi-modal, highly efficient

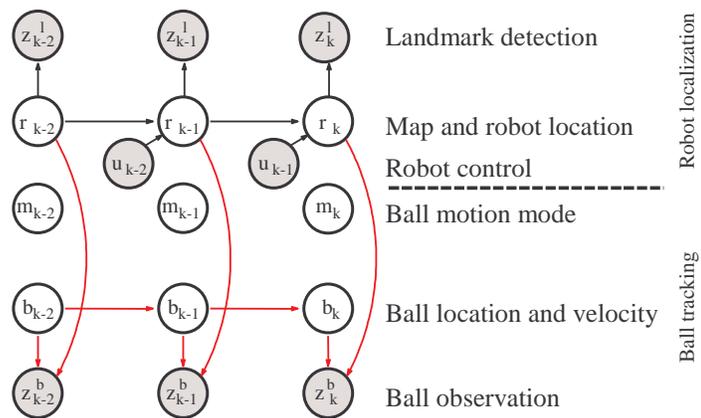
Dynamic Bayes Network for Ball Tracking



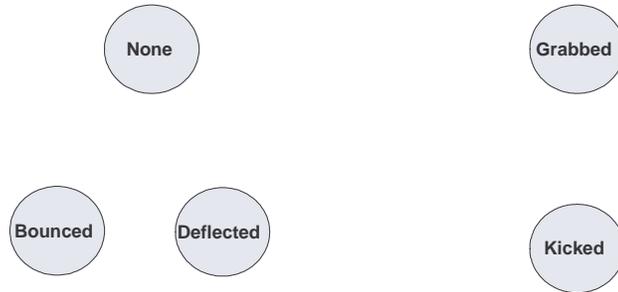
Robot Location



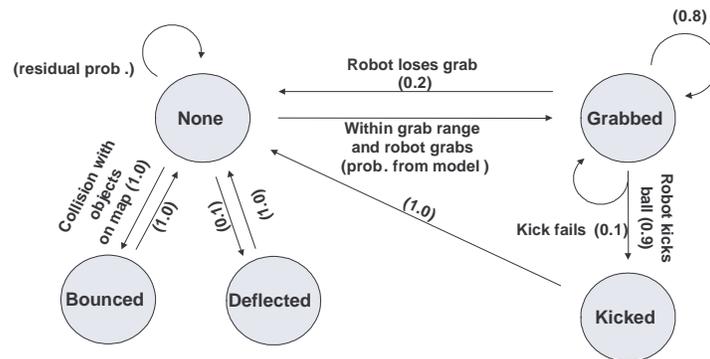
Robot and Ball Location (and velocity)



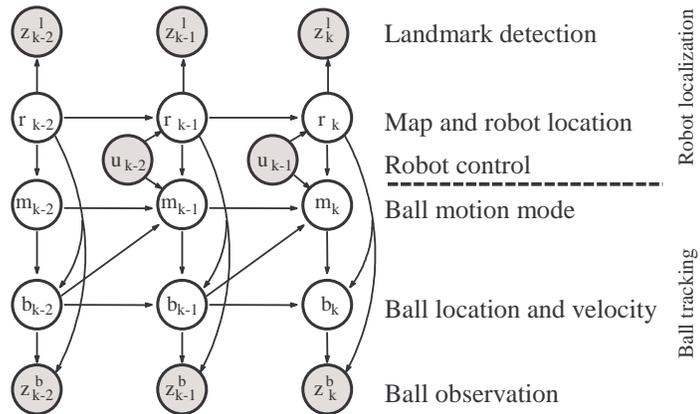
Ball-Environment Interactions



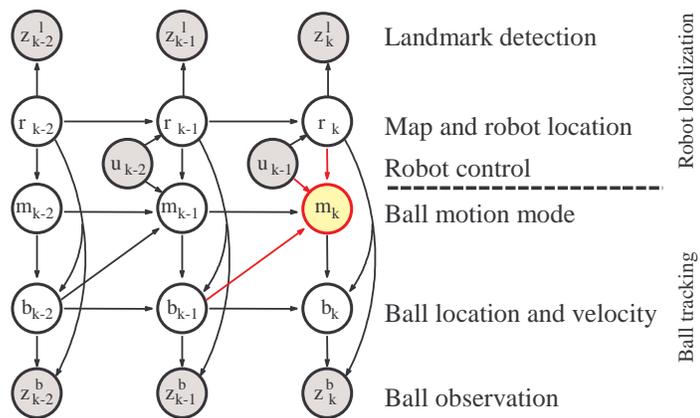
Ball-Environment Interactions



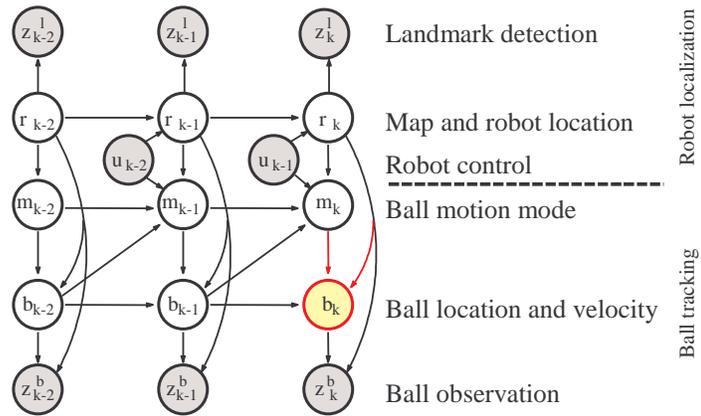
Integrating Discrete Ball Motion Mode



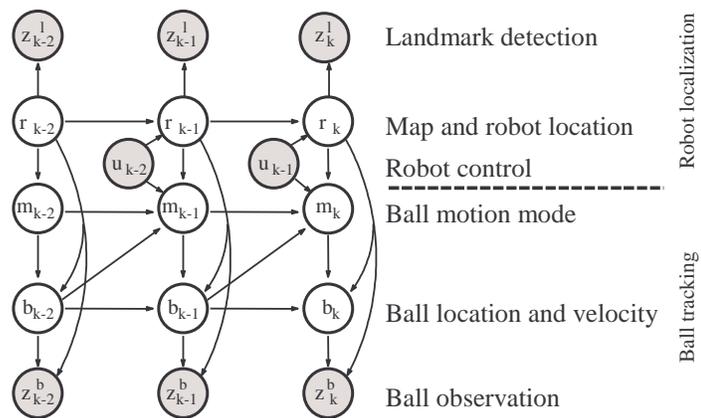
Grab Example (1)



Grab Example (2)



Inference: Posterior Estimation



$$p(b_k, m_k, r_k \mid z_{1:k}^b, z_{1:k}^l, u_{1:k-1})$$

Rao-Blackwellised PF for Inference

§ Represent posterior by random samples

§ Each sample

$$s_i = \langle r_i, m_i, b_i \rangle = \langle \langle x, y, \theta \rangle_i, m_i, \langle \mu, \Sigma \rangle_i \rangle$$

contains robot location, ball mode, ball Kalman filter

§ Generate individual components of a particle stepwise using the factorization

$$p(b_k, m_{1:k}, r_{1:k} | z_{1:k}, u_{1:k-1}) = p(b_k | m_{1:k}, r_{1:k}, z_{1:k}, u_{1:k-1}) p(m_{1:k} | r_{1:k}, z_{1:k}, u_{1:k-1}) \cdot p(r_{1:k} | z_{1:k}, u_{1:k-1})$$

Rao-Blackwellised Particle Filter for Inference

r_{k-1}

r_k

Map and robot location

m_{k-1}

m_k

Ball motion mode

b_{k-1}

b_k

Ball location and velocity

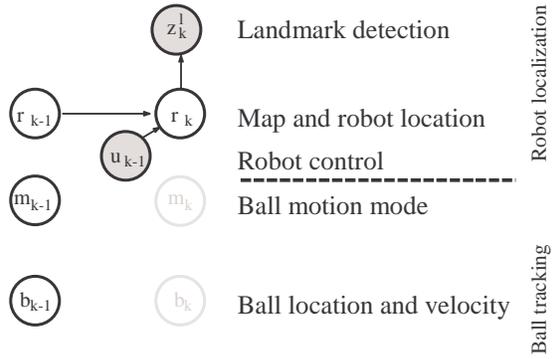
Robot localization

Ball tracking

§ Draw a sample from the previous sample set:

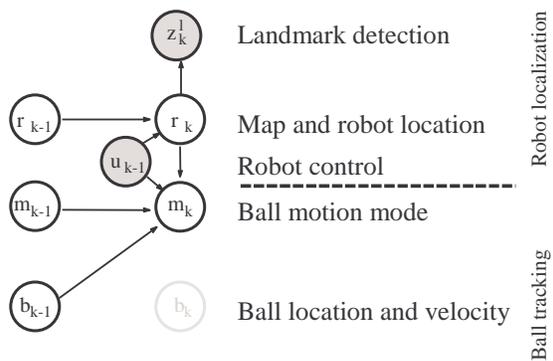
$$\langle r_{k-1}^{(i)}, m_{k-1}^{(i)}, b_{k-1}^{(i)} \rangle$$

Generate Robot Location



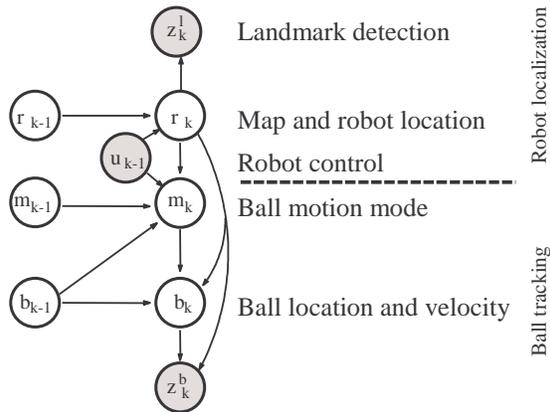
$$r_k^{(i)} \sim p(r_k | r_{k-1}^{(i)}, m_{k-1}^{(i)}, b_{k-1}^{(i)}, z_k, u_{k-1}) \Rightarrow \langle r_k^{(i)}, -, - \rangle$$

Generate Ball Motion Model



$$m_k^{(i)} \sim p(m_k | r_k^{(i)}, m_{k-1}^{(i)}, b_{k-1}^{(i)}, z_k, u_{k-1}) \Rightarrow \langle r_k^{(i)}, m_k^{(i)}, - \rangle$$

Update Ball Location and Velocity



$$b_k^{(i)} \sim p(b_k | r_k^{(i)}, m_k^{(i)}, b_{k-1}^{(i)}, z_k) \Rightarrow \langle r_k^{(i)}, m_k^{(i)}, b_k^{(i)} \rangle$$

Importance Resampling

§ Weight sample by

$$w_k^{(i)} \propto p(z_k^l | r_k^{(i)})$$

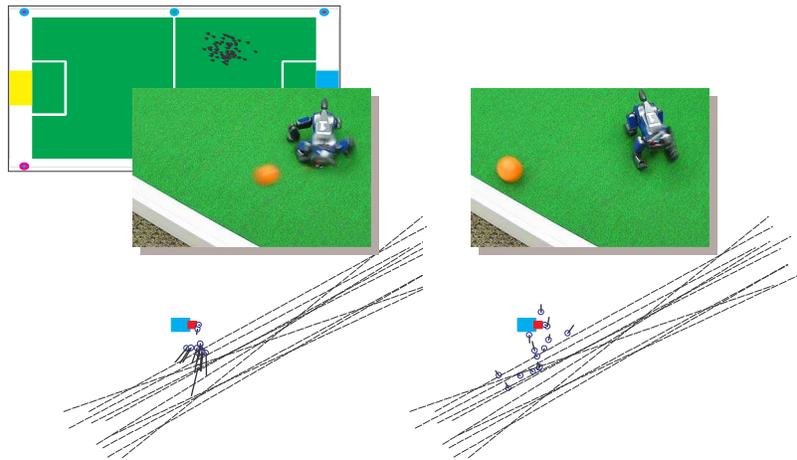
if observation is landmark detection and by

$$w_k^{(i)} \propto \sum_{M_k} p(z_k^b | M_k, r_k^{(i)}, b_{k-1}^{(i)}) p(M_k | r_k^{(i)}, m_{k-1}^{(i)}, b_{k-1}^{(i)}, u_{k-1})$$

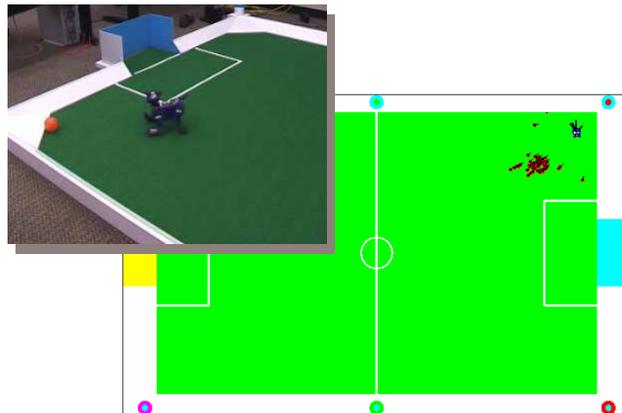
if observation is ball detection.

§ Resample

Ball-Environment Interaction

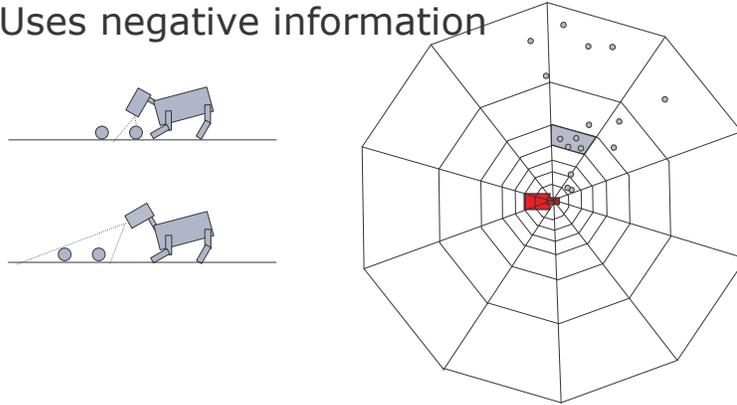


Ball-Environment Interaction



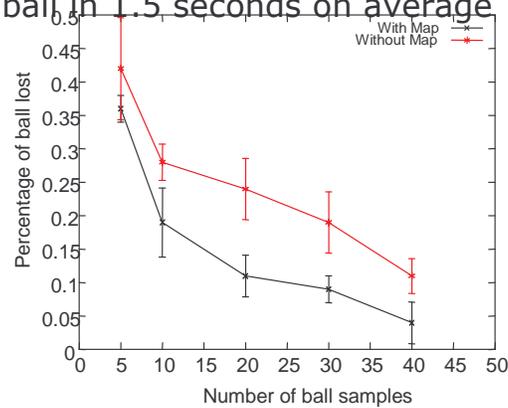
Tracking and Finding the Ball

- § Cluster ball samples by discretizing pan / tilt angles
- § Uses negative information



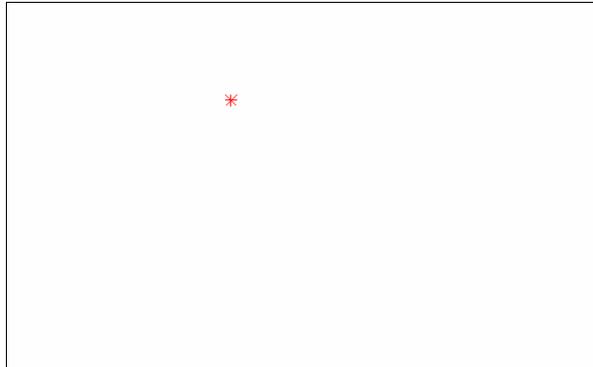
Experiment: Real Robot

- § Robot kicks ball 100 times, tries to find it afterwards
- § Finds ball in 1.5 seconds on average



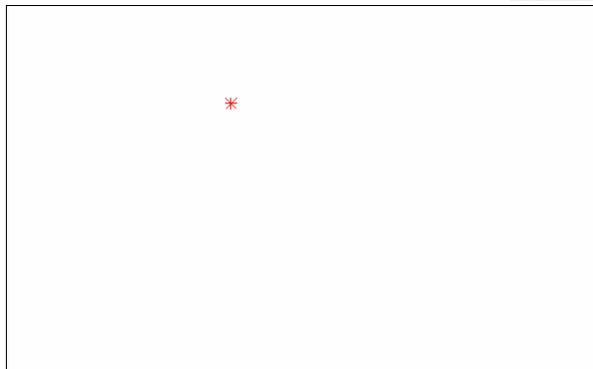
Simulation Runs

----- Reference
-*- Observations

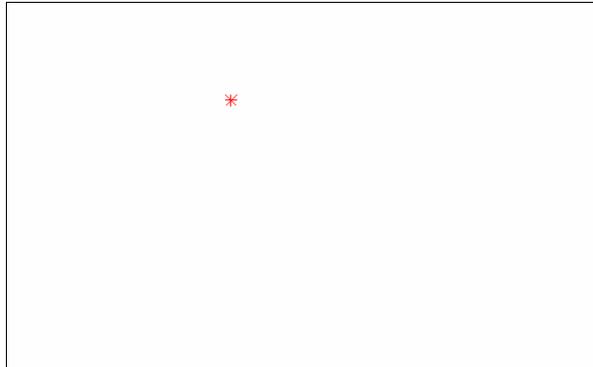
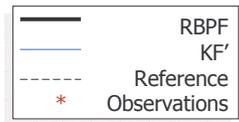


Comparison to KF* (optimized for straight motion)

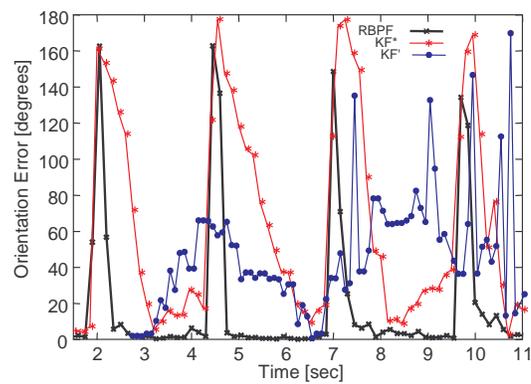
— RBPf
— KF*
----- Reference
-*- Observations



Comparison to KF' (inflated prediction noise)



Orientation Errors



Conclusions

- § Bayesian filters are the most successful technique in robotics (vision?)
- § Many instances (Kalman, particle, grid, MHT, RBPF, ...)
- § Special case of dynamic Bayesian networks
- § Recently: hierarchical models