

Introduction to Artificial Intelligence

V22.0472-001 Fall 2009

Lecture 18: Particle & Kalman Filtering

Rob Fergus – Dept of Computer Science, Courant Institute, NYU



Slides from John DeNero, Dan Klein, Haris Baltzakis, Dieter Fox

Announcements


- Final exam will be at 7pm on Wednesday December 14th
 - Date of last class
 - 1.5 hrs long
- I won't ask anything about the last few classes.

Recap: Reasoning Over Time

- Stationary Markov models



- Hidden Markov models



X	E	P
rain	umbrella	0.9
rain	no umbrella	0.1
sun	umbrella	0.2
sun	no umbrella	0.8


Recap: Filtering

Elapse time: compute $P(X_t | e_{1:t-1})$

$$P(x_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1} | e_{1:t-1}) \cdot P(x_t | x_{t-1})$$

Observe: compute $P(X_t | e_{1:t})$

$$P(x_t | e_{1:t}) \propto P(x_t | e_{1:t-1}) \cdot P(e_t | x_t)$$



Belief: $\langle P(\text{rain}), P(\text{sun}) \rangle$

$P(X_1)$ $\langle 0.5, 0.5 \rangle$ Prior on X_1

$P(X_1 | E_1 = \text{umbrella})$ $\langle 0.82, 0.18 \rangle$ Observe

$P(X_2 | E_1 = \text{umbrella})$ $\langle 0.63, 0.37 \rangle$ Elapse time

$P(X_2 | E_1 = \text{umb}, E_2 = \text{umb})$ $\langle 0.88, 0.12 \rangle$ Observe

Particle Filtering

- Sometimes $|X|$ is too big to use exact inference


- $|X|$ may be too big to even store $B(X)$
- E.g. X is continuous
- $|X|^2$ may be too big to do updates

- Solution: approximate inference


- Track samples of X , not all values
- Samples are called particles
- Time per step is linear in the number of samples
- But: number needed may be large
- In memory: list of particles, not states

- This is how robot localization works in practice

0.0	0.1	0.0
0.0	0.0	0.2
0.0	0.2	0.5




Example: State Representations for Robot Localization



Grid Based approaches (Markov localization)

Representation: Particles

- Our representation of $P(X)$ is now a list of N particles (samples)
 - Generally, $N \ll |X|$
 - Storing map from X to counts would defeat the point
- $P(x)$ approximated by number of particles with value x
 - So, many x will have $P(x) = 0!$
 - More particles, more accuracy
- For now, all particles have a weight of 1



Particles:
 (3,3)
 (2,3)
 (3,3)
 (3,3)
 (3,2)
 (3,3)
 (3,3)
 (2,1)
 (3,3)
 (3,3)
 (2,1)


7

Particle Filtering: Elapse Time

- Each particle is moved by sampling its next position from the transition model

$$x' = \text{sample}(P(X'|x))$$

- This is like prior sampling – samples' frequencies reflect the transition probs
- Here, most samples move clockwise, but some move in another direction or stay in place
- This captures the passage of time
 - If we have enough samples, close to the exact values before and after (consistent)




Particle Filtering: Observe

- Slightly trickier:
 - Don't do rejection sampling (why not?)
 - We don't sample the observation, we fix it
 - This is similar to likelihood weighting, so we downweight our samples based on the evidence

$$w(x) = P(e|x)$$

$$B(X) \propto P(e|X)B'(X)$$

- Note that, as before, the probabilities don't sum to one, since most have been downweighted (in fact they sum to an approximation of $P(e)$)




Particle Filtering: Resample

- Rather than tracking weighted samples, we resample
- N times, we choose from our weighted sample distribution (i.e. draw with replacement)
 - This is equivalent to renormalizing the distribution
- Now the update is complete for this time step, continue with the next one

Old Particles:
 (3,3) $w=0.1$
 (2,1) $w=0.9$
 (3,1) $w=0.4$
 (3,2) $w=0.3$
 (2,2) $w=0.4$
 (1,1) $w=0.4$
 (3,1) $w=0.4$
 (2,1) $w=0.9$
 (3,2) $w=0.3$

Old Particles:
 (2,1) $w=1$
 (2,1) $w=1$
 (2,1) $w=1$
 (3,2) $w=1$
 (2,2) $w=1$
 (2,1) $w=1$
 (1,1) $w=1$
 (3,1) $w=1$
 (2,1) $w=1$
 (1,1) $w=1$



Particle Filter Algorithm

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





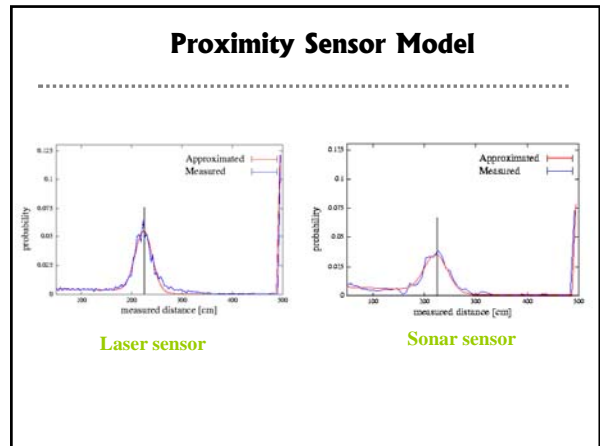
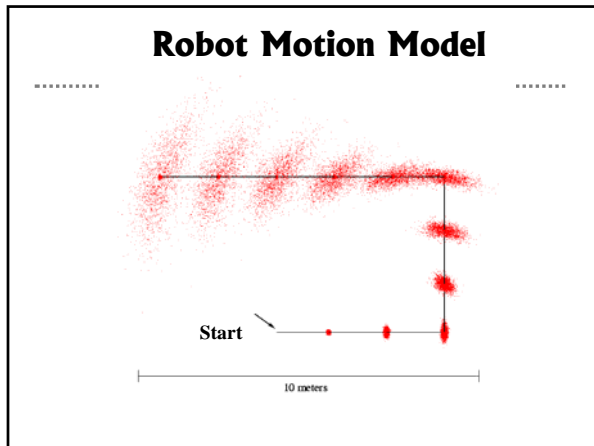
- draw x_{t-1}^j from $Bel(x_{t-1})$
- draw x_t^j from $p(x_t | x_{t-1}^j, u_{t-1})$
- Importance factor for x_t^j :






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







Robot Localization




- In robot localization:
 - We know the map, but not the robot's position
 - Observations may be vectors of range finder readings
 - State space and readings are typically continuous (works basically like a very fine grid) and so we cannot store $B(X)$
 - Particle filtering is a main technique












Robotic Cars

- DARPA Grand Challenge
- DARPA Urban Challenge
- <http://www.youtube.com/watch?v=SQFEmR50HAK>



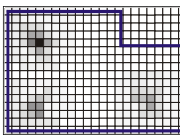

SLAM

- SLAM = Simultaneous Localization And Mapping
 - We do not know the map or our location
 - Our belief state is over maps and positions!
 - Main techniques: Kalman filtering (Gaussian HMMs) and particle methods

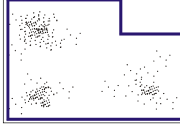


DP-SLAM, Ron Parr

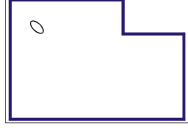
Example: State Representations for Robot Localization



Grid Based approaches
(Markov localization)



Particle Filters
(Monte Carlo localization)



Kalman
Tracking

36

Kalman Filters - Equations

Recursive filter for estimating state of linear dynamical system from noisy measurements

$$\begin{cases} P(x_t | x_{t-1}) \approx N(Ax_{t-1}, \Gamma) \\ P(y_t | x_t) \approx N(Cx_t, \Sigma) \end{cases}$$

A: State transition matrix ($n \times n$)
 C: Measurement matrix ($m \times n$)
 w_t : Process noise ($\in \mathbb{R}^n$)
 v_t : Measurement noise ($\in \mathbb{R}^m$)

$$\begin{cases} x_t = Ax_{t-1} + w_t \\ y_t = Cx_t + v_t \end{cases}$$

Process dynamics (motion model)
 measurements (observation model)

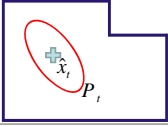
Where: $N(x; m, V) = \frac{1}{|2\pi V|^{1/2}} \exp\left(-\frac{1}{2}(x-m)^T V^{-1}(x-m)\right)$

37

Kalman Filters - Update

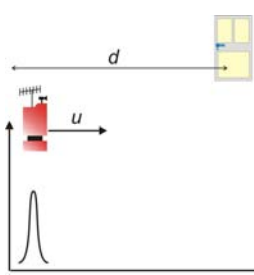
$$\begin{cases} x_t = Ax_{t-1} + w_t \\ y_t = Cx_t + v_t \\ w_t \approx N(0, \Gamma) \\ v_t \approx N(0, \Sigma) \end{cases}$$

Predict state, $\hat{x}_t^- = A\hat{x}_{t-1}$
 covariance $P_t^- = AP_{t-1}A^T + \Gamma$
 Compute Gain $K_t = P_t^- C^T (C P_t^- C^T + \Sigma)^{-1}$
 Compute Innovation $J_t = \hat{y}_t - C\hat{x}_t^-$
 Update $\hat{x}_t = \hat{x}_t^- - K_t J_t$
 $P_t = (I - K_t C) P_t^-$



38

Kalman Filter - Example

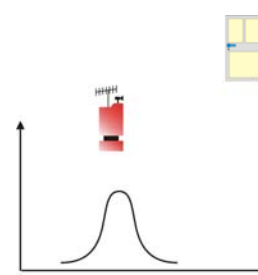


$$\begin{cases} x_t = Ax_{t-1} + B + w_t & A_t = [1] \\ y_t = Cx_t + D + v_t & B_t = [u_t] \\ w_t \approx N(0, \Gamma) & C_t = [-1] \\ v_t \approx N(0, \Sigma) & D_t = [1] \end{cases}$$

$$\begin{cases} x_t = x_{t-1} + u_t + w_t \\ y_t = d - x_t + v_t \\ w_t \approx N(0, \Gamma) \\ v_t \approx N(0, \Sigma) \end{cases}$$

39

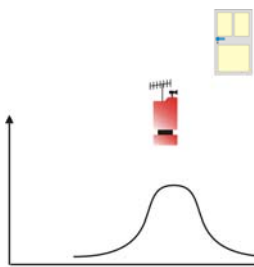
Kalman Filter - Example



Predict $\hat{x}_t^- = A\hat{x}_{t-1} + B$
 $P_t^- = AP_{t-1}A^T + \Gamma$

40

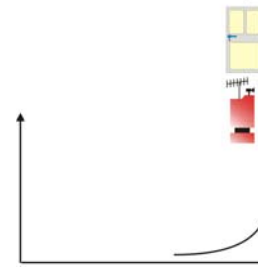
Kalman Filter - Example



Predict $\hat{x}_t^- = A\hat{x}_{t-1} + B$
 $P_t^- = AP_{t-1}A^T + \Gamma$

41


Kalman Filter - Example



Predict $\hat{x}_t^- = A\hat{x}_{t-1} + B$
 $P_t^- = AP_{t-1}A^T + \Gamma$
 Compute Innovation $J_t = \hat{y}_t - C\hat{x}_t^-$
 Compute Gain $K_t = P_t^- C^T (C P_t^- C^T + \Sigma)^{-1}$

42

Kalman Filter – Example



Predict

$$\hat{x}_i^- = A\hat{x}_{i-1} + B$$

$$P_i^- = AP_{i-1}A^T + \Gamma$$

Compute Innovation

$$J_i = \hat{y}_i - C\hat{x}_i^-$$

Compute Gain

$$K_i = P_i^- C^T (C P_i^- C^T + \Sigma)^{-1}$$


Update

$$\hat{x}_i = \hat{x}_i^- - K_i J_i$$

$$P_i = (I - K_i C) P_i^-$$

43

Kalman Filter – Example



Predict

$$\hat{x}_i^- = A\hat{x}_{i-1} + B$$

$$P_i^- = AP_{i-1}A^T + \Gamma$$

44

Kalman Filter Applications

Apollo guidance computer

Cruise missiles

Airplane autopilot

Robotics

Finance

...




Continuous State Approaches

- Perform very accurately if the inputs are precise (performance is optimal with respect to any criterion in the linear case).
- Computational efficiency.
- Requirement that the initial state is known.
- Inability to recover from catastrophic failures
- Inability to track Multiple Hypotheses the state (Gaussians have only one mode)

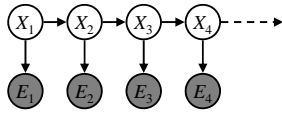
46

Discrete State Approaches

- Ability (to some degree) to operate even when its initial pose is unknown (start from uniform distribution).
- Ability to deal with noisy measurements.
- Ability to represent ambiguities (multi modal distributions).
- Computational time scales heavily with the number of possible states (dimensionality of the grid, number of samples, size of the map).
- Accuracy is limited by the size of the grid cells/number of particles-sampling method.
- Required number of particles is unknown

47

Best Explanation Queries




- Query: most likely seq:

$$\arg \max_{x_{1:t}} P(x_{1:t} | e_{1:t})$$

48

State Path Trellis


- State trellis: graph of states and transitions over time



- Each arc represents some transition $x_{t-1} \rightarrow x_t$
- Each arc has weight $P(x_t|x_{t-1})P(e_t|x_t)$
- Each path is a sequence of states
- The product of weights on a path is the seq's probability
- Can think of the Forward (and now Viterbi) algorithms as computing sums of all paths (best paths) in this graph

49

Viterbi Algorithm




$$x_{1:T}^* = \arg \max_{x_{1:T}} P(x_{1:T}|e_{1:T}) = \arg \max_{x_{1:T}} P(x_{1:T}, e_{1:T})$$

$$\begin{aligned} m_t[x_t] &= \max_{x_{1:t-1}} P(x_{1:t-1}, x_t, e_{1:t}) \\ &= \max_{x_{1:t-1}} P(x_{1:t-1}, e_{1:t-1})P(x_t|x_{t-1})P(e_t|x_t) \\ &= P(e_t|x_t) \max_{x_{t-1}} P(x_t|x_{t-1}) \max_{x_{1:t-2}} P(x_{1:t-1}, e_{1:t-1}) \\ &= P(e_t|x_t) \max_{x_{t-1}} P(x_t|x_{t-1}) m_{t-1}[x_{t-1}] \end{aligned}$$

50

Example



51

Andrew Viterbi

