CS354

Nathan Sprague

February 19, 2015

Probabilistic State Representations: Grid-Based

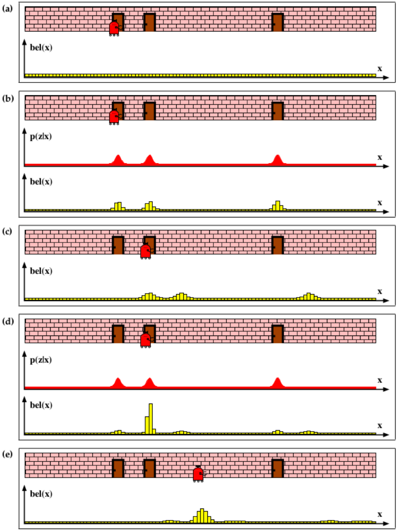

Figure 8.1 Grid localization using a fine-grained metric decomposition. Each picture depicts the position of the robot in the hallway along with its belief $bel(x_t)$, represented by a histogram over a grid.

Probabilistic Robotics. Thrun, Burgard, Fox, 2005

Probabilistic State Representations: Continuous

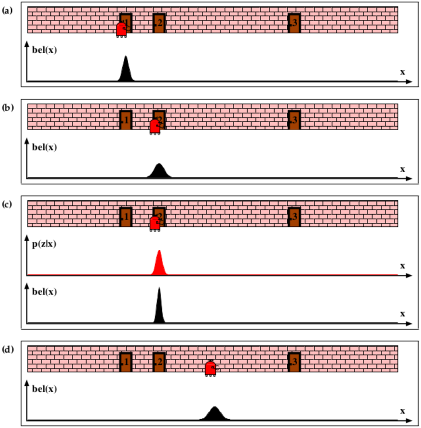

Figure 7.6 Application of the Kalman filter algorithm to mobile robot localization. All densities are represented by unimodal Gaussians.

Probabilistic Robotics. Thrun, Burgard, Fox, 2005

Probability Density Functions

Represent probability distributions over random variables:

- Properties:
  - $f(x) \geq 0$
  - $\int_{-\infty}^{\infty} f(x)dx = 1$
- Interpretation:
  - $P(a \leq x \leq b) = \int_{a}^{b} f(x)dx$

Expectation, Variance

- Expectation (continuous)

$$\mu = \mathbb{E}[x] = \int x f(x) dx$$

- Expectation (discrete)

$$\mathbb{E}[X] = \sum_{1}^{n} P(x_i) x_i$$

- Variance

$$\sigma^2 = \mathbb{E}[(x - \mathbb{E}[x])^2]$$

Normal Distribution

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

(*Normal* because of the central limit theorem.)

Combining Evidence

- Imagine two independent measurements of some unknown quantity:
  - x_1 with variance σ_1^2
  - x_2 with variance σ_2^2
- How should we combine these measumrents?

Combining Evidence

- Imagine two independent measurements of some unknown quantity:
  - x_1 with variance σ_1^2
  - x_2 with variance σ_2^2
- How should we combine these measumrents?
- We can take a weighted average:
  - $\hat{x} = \omega_1 x_1 + \omega_2 x_2$ (where $\omega_1 + \omega_2 = 1$)
- What should the weights be???

Combining Evidence

- Imagine two independent measurements of some unknown quantity:
  - x_1 with variance σ_1^2
  - x_2 with variance σ_2^2
- How should we combine these measumrents?
- We can take a weighted average:
  - $\hat{x} = \omega_1 x_1 + \omega_2 x_2$ (where $\omega_1 + \omega_2 = 1$)
- What should the weights be???
- We want to find weights that minimize variance (uncertainty) in the estimate:
  - $\sigma^2 = E[(\hat{x} - E[\hat{x}])^2]$

Combining Evidence – Solution

(Derivation not shown...)

$$\hat{x} = \frac{\sigma_2^2 x_1 + \sigma_1^2 x_2}{\sigma_2^2 + \sigma_1^2}$$

$$\sigma^2 = \frac{\sigma_1^2 \sigma_2^2}{\sigma_2^2 + \sigma_1^2}$$

Vector-Valued State

- We'll need to generalize all of this to the case where the state of the system can't be represented as a single number.
- Use a vector $\mathbf{x}$ to represent the state.

Covariance

$$cov(x, y) = \mathbb{E}[(x - \mu_x)(y - \mu_y)]$$

- Properties:
  - $cov(x, y) = cov(y, x)$
  - If x and y are independent, $cov(x, y) = 0$
  - If $cov(x, y) > 0$, y tends to increase when x increases.
  - If $cov(x, y) < 0$, y tends to decrease when x increases.

Covariance Matrix

- Covariance matrix:

$$cov(\mathbf{x}) = \Sigma_{\mathbf{x}} = \mathbb{E}[(\mathbf{x} - \hat{\mathbf{x}})(\mathbf{x} - \hat{\mathbf{x}})^T]$$

- Where $\mathbf{x}$ is a random vector and $\hat{\mathbf{x}}$ is the vector mean.
- The entry at row i, column j in the matrix is cov($\mathbf{x}_i$,$\mathbf{x}_j$)
- Multivariate normal distribution is parameterized by the mean vector and covariance matrix.

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

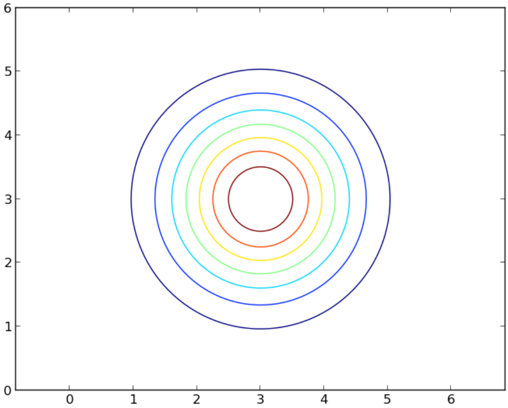

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \Sigma = \begin{bmatrix} .2 & 0 \\ 0 & 1 \end{bmatrix}$$

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \Sigma = \begin{bmatrix} .2 & 0 \\ 0 & 1 \end{bmatrix}$$

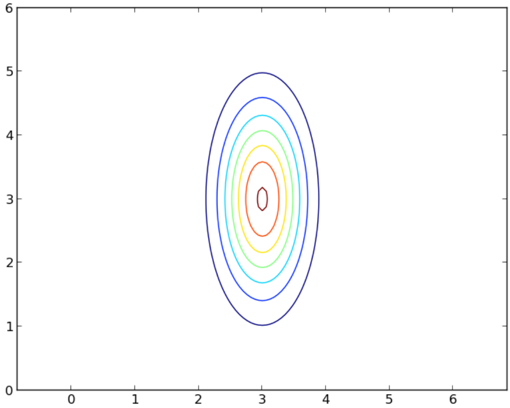

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \Sigma = \begin{bmatrix} 1 & .9 \\ .9 & 1 \end{bmatrix}$$

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \Sigma = \begin{bmatrix} 1 & .9 \\ .9 & 1 \end{bmatrix}$$

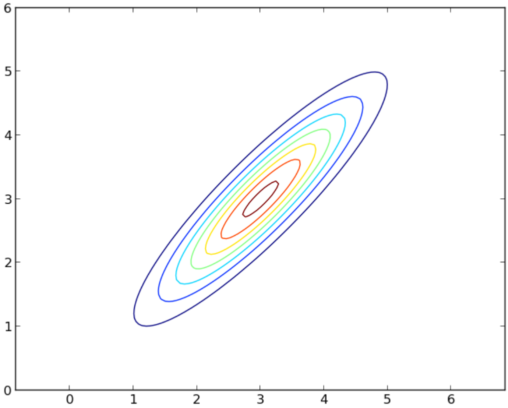

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \; \Sigma = \begin{bmatrix} .5 & -.3 \\ -.3 & 2 \end{bmatrix}$$

Multivariate PDF Example

$$\mathbf{x} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \; \Sigma = \begin{bmatrix} .5 & -.3 \\ -.3 & 2 \end{bmatrix}$$

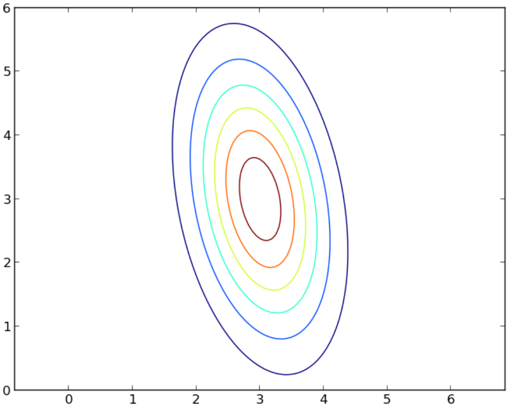

Linear System Models

- State can include information other than position. E.g. velocity.
- Linear model of an object moving with a fixed velocity in 2d:
  - $x_{t+1} = x_t + \dot{x}_t dt$
  - $y_{t+1} = y_t + \dot{y}_t dt$
  - $\dot{x}_{t+1} = \dot{x}_t$
  - $\dot{y}_{t+1} = \dot{y}_t$
- dt is time.
- $\dot{x}_t$ is velocity along the x axis.

Linear System Model in Matrix Form

This is equivalent to the last slide:

$$\mathbf{x}_t = \begin{bmatrix} x_t \\ y_t \\ \dot{x}_t \\ \dot{y}_t \end{bmatrix}$$

$$\mathbf{x}_{t+1} = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{x}_t$$

Kalman Filter

- Assumes:
  - Linear state dynamics
  - Linear sensor model
  - Normally distributed noise in the state dynamics
  - Normally distributed noise in the sensor model
- State Transition Model:
  - $\mathbf{x}_t = A\mathbf{x}_{t-1} + B\mathbf{u}_{t-1} + \mathbf{w}_{t-1}$
- Sensor Model:
  - $\mathbf{z}_t = H\mathbf{x}_t + \mathbf{v}_t$

Kalman Filter in One Slide

- Predict:
  Project the state forward:
  $$\hat{\mathbf{x}}_t^- = A\hat{\mathbf{x}}_{t-1} + B\mathbf{u}_{t-1}$$
  Project the covariance of the state estimate forward:
  $$\mathbf{P}_t^- = A\mathbf{P}_{t-1}A^T + \mathbf{Q}$$
- Correct:
  Compute the Kalman gain:
  $$\mathbf{K}_t = \mathbf{P}_t^- H^T (H\mathbf{P}_t^- H^T + \mathbf{R})^{-1}$$
  Update the estimate with the measurement:
  $$\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_t^- + \mathbf{K}_t(\mathbf{z}_t - H\hat{\mathbf{x}}_t^-)$$
  Update the estimate covariance:
  $$\mathbf{P}_t = (\mathbf{I} - \mathbf{K}_t H)\mathbf{P}_t^-$$