CS354

Nathan Sprague

September 21, 2023

Probabilistic State Representations: Continuous

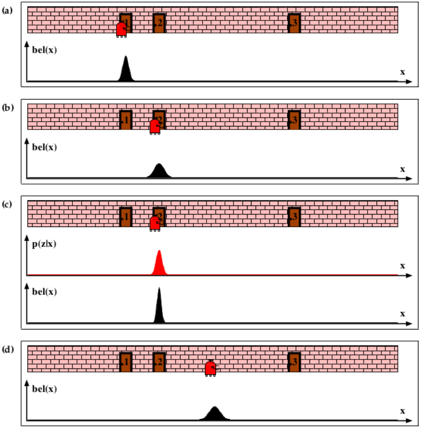

Figure 7.6 Application of the Kalman filter algorithm to mobile robot localization. All densities are represented by unimodal Gaussians.

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) (also referred to as the "mean" or "first moment")

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

- Expectation (discrete)

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

- Variance (also referred to as the "second moment")

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

Quiz 1

What is the expectation of this pdf?

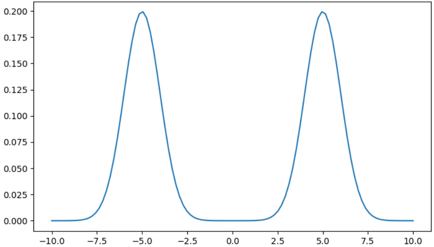

Quiz 2

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

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

Imagine we are rolling a four-sided die. The following probability distribution describes the probability for each number that we could roll:

$P(X = 1) = .7$
$P(X = 2) = .1$
$P(X = 3) = .1$
$P(X = 4) = .1$

What is the expected value of this distribution? What is the variance?

Sample Mean and Variance

Expectation and variance are properties of distributions. We can also calculate the **sample mean** and the **sample variance** for a given data set:
$\{x_1, x_2, ..., x_n\}$.

- Sample mean

$$m = \frac{1}{n}\sum_{i=1}^{n} x_i$$

- Sample variance

$$s^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - m)^2$$

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.)
All distributions

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

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

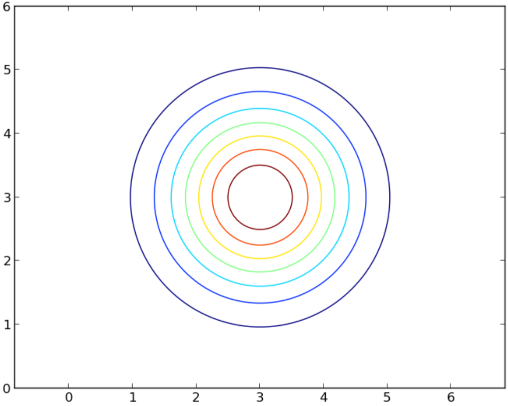

Multivariate PDF Example

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

Multivariate PDF Example

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

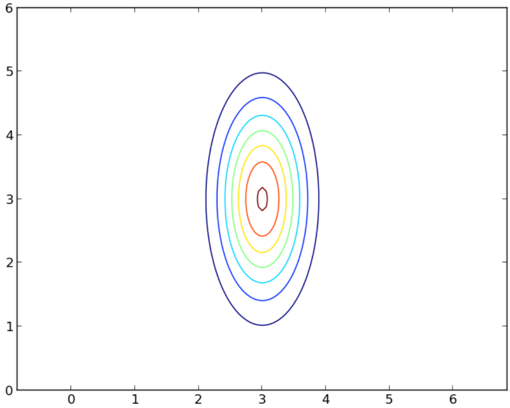

Multivariate PDF Example

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

Multivariate PDF Example

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

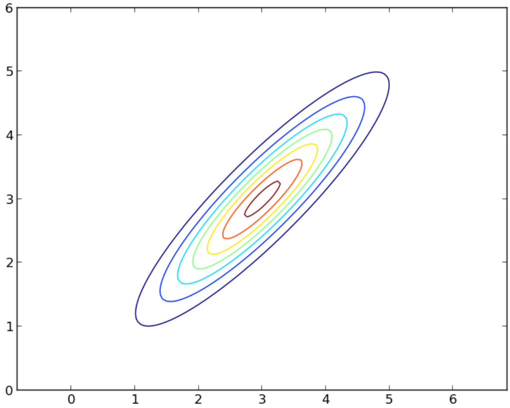

Multivariate PDF Example

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

Multivariate PDF Example

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

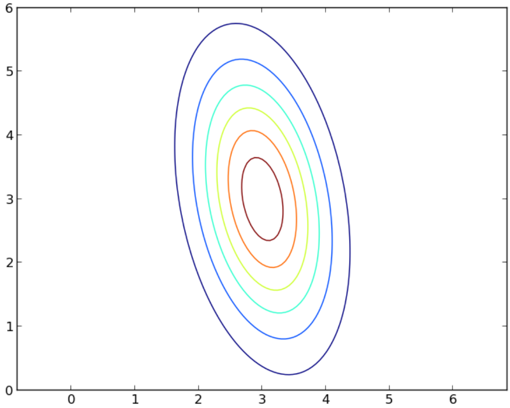

Why is this Useful For Localization?

- Memory and computation requirements grow exponentially for grid-based disributions. E.g. if we want 100 cells per dimension, we need 100^d cells.
- To approximate with a normal distribution we need $d^2 + d$ to store.

Can We Do Recursive State Estimation?

- Two Steps:
  - Prediction based on system dynamics:

$$Bel^{-}(x_t) = \int p(x_t \mid x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

  - Correction based on sensor reading:

$$Bel(x_t) = \eta p(z_t \mid x_t) Bel^{-}(x_t)$$

YES. The Kalman filter. Next time.