Basics: Gaussian Functions

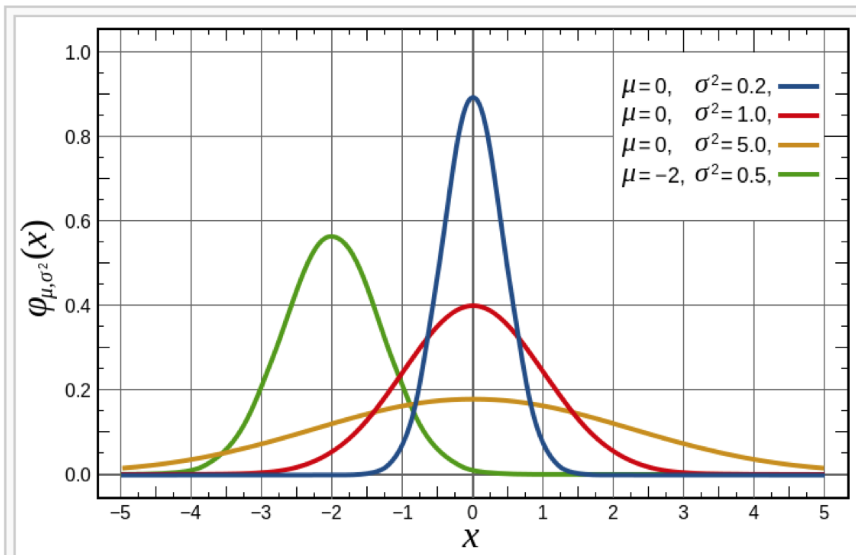

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

$$f(x_1, \ldots, x_k) = \frac{1}{\sqrt{(2\pi)^k |\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right)$$

Optimal Estimation

sensor feedback is subject to noise and imprecision

suppose that the ball is stationary at some point along the x_B axis and that stereo observations of coordinate x are drawn from the probability distribution illustrated.

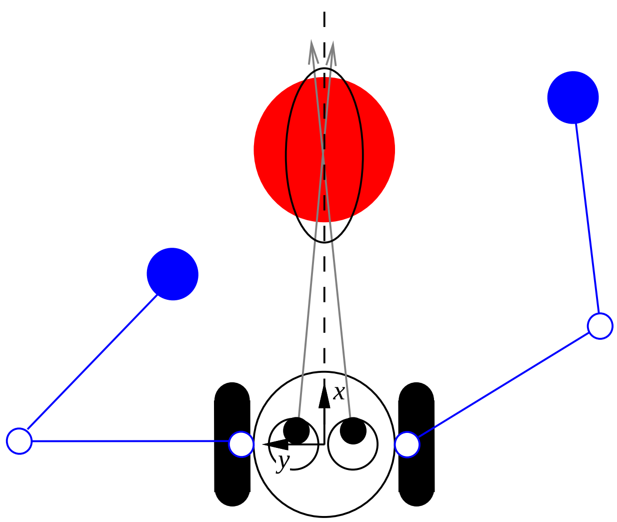

given n independent measurements,

$$\{z_i : i = 1, \ldots, n\},$$

each subject to additive zero-mean Gaussian noise

$$v_i \sim N(0, r_i),$$

- the optimal estimate, $\hat{x}$, is constructed as a weighted sum of individual (noisey) measurements $\hat{x} = k_1 z_1 + k_2 z_2$
- weights are chosen to minimize the expected squared error (variance) of the estimate.

Optimal Estimation

two independent measurements:

$$z_1 = x + v_1 \qquad z_2 = x + v_2$$

$$\hat{x} = k_1 z_1 + k_2 z_2.$$

we require that the estimator is **unbiased** $\Longleftrightarrow$ k_1 and k_2 are independent of x and the expected value of the estimation error is zero,

$$E[\tilde{x}] = E[\hat{x} - x] = 0$$

therefore $$E[(k_1(x + v_1) + k_2(x + v_2)) - x] = 0$$

and since $E[x] = x$ and $E[v_i] = 0$, this relation requires that

$$\boxed{k_2 = 1 - k_1.}$$

Optimal Estimation

the optimal filter gain (k_1) yields **minimum squared error**[1]

$$E[\tilde{x}^2] = k_1^2 r_1 + (1 - k_1)^2 r_2$$

where r_i is the observation variance for measurement i.

we find the value for k_1 that minimizes variance:

$$\frac{dE[\tilde{x}^2]}{dk_1} = 2k_1 r_1 - 2(1 - k_1) r_2 = 0$$

$$\boxed{k_1 = \frac{r_2}{r_1 + r_2}}$$

[1]which is equivalent to minimizing the estimate variance.

Optimal Estimation

Now, the optimal estimate is:

$$\begin{aligned}\hat{x} &= \frac{r_2}{r_1 + r_2} z_1 + \frac{r_1}{r_1 + r_2} z_2 \\ &= \frac{\frac{1}{r_1}}{\frac{1}{r_1} + \frac{1}{r_2}} z_1 + \frac{\frac{1}{r_2}}{\frac{1}{r_1} + \frac{1}{r_2}} z_2,\end{aligned}$$

and the expected squared estimation error (the estimate variance) is:

$$E[\tilde{x}^2] = s = \left[\frac{1}{r_1} + \frac{1}{r_2}\right]^{-1}$$

which is generalized to k observations in a straightforward manner.

Optimal Estimation

for k observations, the optimal (least squares) estimate and the estimate variance becomes:

$$\hat{x}_k = \frac{\sum_{i=1}^{k} \frac{z_i}{r_i}}{\sum_{i=1}^{k} \frac{1}{r_i}} \qquad\qquad s_k = \left[\sum_{i=1}^{k} \frac{1}{r_i} \right]^{-1}$$

Finally, we can state the filter in the **recursive form:**

$$\hat{x}_k = \frac{\frac{1}{s_{k-1}}}{\frac{1}{s_{k-1}} + \frac{1}{r_k}} \hat{x}_{k-1} + \frac{\frac{1}{r_k}}{\frac{1}{s_{k-1}} + \frac{1}{r_k}} z_k$$

$$s_k = \left[\frac{1}{s_{k-1}} + \frac{1}{r_k} \right]^{-1}$$

Recursive Optimal Estimation - Tracking Moving Objects $\in \mathbb{R}^n$

The optimal combination of measurements, $\{\boldsymbol{z}_i : i = 1, k\}$, with associated covariance $\mathbf{R}_i$ is determined by weighting observations

$$\hat{\boldsymbol{x}}_k = \mathbf{S}_k \sum_{i=1}^{k} \mathbf{R}_i^{-1} \boldsymbol{z}_i$$

where the covariance of the estimate is computed from

$$\mathbf{S}_k = \left[\sum_{i=1}^{k} \mathbf{R}_i^{-1} \right]^{-1}$$

The recursive form of the multi-dimensional estimator becomes:

$$\hat{\boldsymbol{x}}_{k+1} = \mathbf{S}_k^{-1} \left[\mathbf{S}_k^{-1} + \mathbf{R}_k^{-1} \right]^{-1} \hat{\boldsymbol{x}}_k + \mathbf{R}_k^{-1} \left[\mathbf{S}_k^{-1} + \mathbf{R}_k^{-1} \right]^{-1} \boldsymbol{z}_{k+1}$$

Recursive Optimal Estimation - Tracking Moving Objects $\in \mathbb{R}^n$

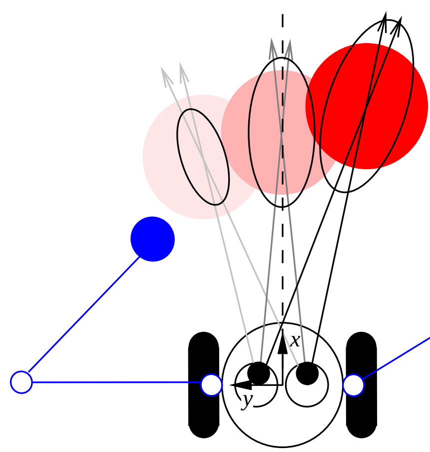

$$\hat{\boldsymbol{x}}_k^- = \begin{bmatrix} \hat{x} & \hat{y} & \dot{\hat{x}} & \dot{\hat{y}} \end{bmatrix}_k^T = \boldsymbol{\Phi}_k \hat{\boldsymbol{x}}_{k-1}^+ + \boldsymbol{w}_{k-1} \qquad \boldsymbol{w}_k \sim N(\mathbf{0}, \mathbf{Q}_k)$$

$$\boldsymbol{z}_k = \mathbf{H}_k \hat{\boldsymbol{x}}_k + \boldsymbol{v}_k \qquad \boldsymbol{v}_k \sim N(\mathbf{0}, \mathbf{R}_k)$$

The Kalman Filter

the "process" model

$$\hat{\boldsymbol{x}}_k = \boldsymbol{\Phi}_{k-1}\hat{\boldsymbol{x}}_{k-1} + \boldsymbol{w}_{k-1} \quad \text{, or}$$

$$\hat{\boldsymbol{x}}_k = \mathbf{A}_{k-1}\hat{\boldsymbol{x}}_{k-1} + \mathbf{B}_{k-1}\boldsymbol{u}_{k-1} + \boldsymbol{w}_{k-1}$$

$$\boldsymbol{w}_{k-1} \sim N(0, \mathbf{Q}_{k-1})$$

the "sensor" model

$$\boldsymbol{z}_k = \mathbf{H}_k\hat{\boldsymbol{x}}_k + \boldsymbol{v}_k$$

$$\boldsymbol{v}_k \sim N(0, \mathbf{R}_k)$$

The Kalman filter is implemented in two stages:

state prediction:

$$\hat{\boldsymbol{x}}_k^- = \mathbf{A}_{k-1}\hat{\boldsymbol{x}}_{k-1}^+ + \mathbf{B}_{k-1}\boldsymbol{u}_{k-1}$$

$$\mathbf{P}_k^- = \mathbf{A}_{k-1}\mathbf{P}_{k-1}^+\mathbf{A}_{k-1}^T + \mathbf{Q}_{k-1}$$

sensor prediction:

$$\mathbf{K}_k = \mathbf{P}_k^-\mathbf{H}_k^T\left[\mathbf{H}_k\mathbf{P}_k^-\mathbf{H}_k^T + \mathbf{R}_k\right]^{-1}$$

$$\hat{\boldsymbol{x}}_k^+ = \hat{\boldsymbol{x}}_k^- + \mathbf{K}_k\left[\boldsymbol{z}_k - \mathbf{H}_k\hat{\boldsymbol{x}}_k^-\right]$$

$$\mathbf{P}_k^+ = \left[\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right]\mathbf{P}_k^-$$

The Kalman Filter

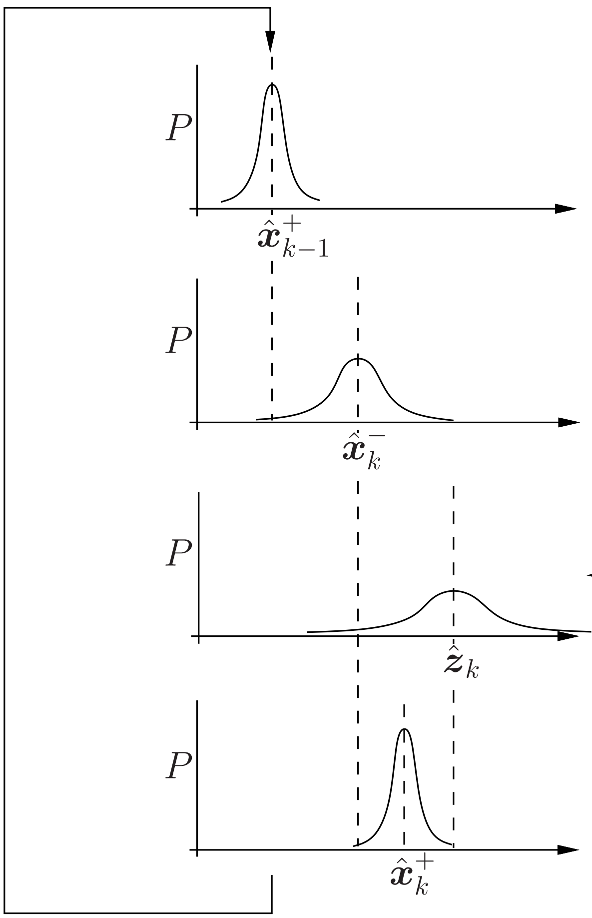

prior

forward state prediction

$\hat{\boldsymbol{x}}^-_k = \boldsymbol{\Phi}\hat{\boldsymbol{x}}^+_{k-1}$

$\mathbf{P}^-_k = \boldsymbol{\Phi}\mathbf{P}^+_{k-1}\boldsymbol{\Phi}^T + \mathbf{Q}_{k-1}$

observation

$\boldsymbol{z}_k, \; \mathbf{R}_k$

$$\mathbf{K}_k = \mathbf{P}^-_k \mathbf{H}^T_k \left[\mathbf{H}_k \mathbf{P}^-_k \mathbf{H}^T_k + \mathbf{R}_k\right]^{-1}$$

posterior

$\hat{\boldsymbol{x}}^+_k = \hat{\boldsymbol{x}}^-_k + \mathbf{K}_k \left[\boldsymbol{z}_k - \mathbf{H}_k \hat{\boldsymbol{x}}^-_k\right]$

$\mathbf{P}^+_k = \left[\mathbf{I} - \mathbf{K}_k \mathbf{H}_k\right] \mathbf{P}^-_k$

A Kalman Filter for Roger

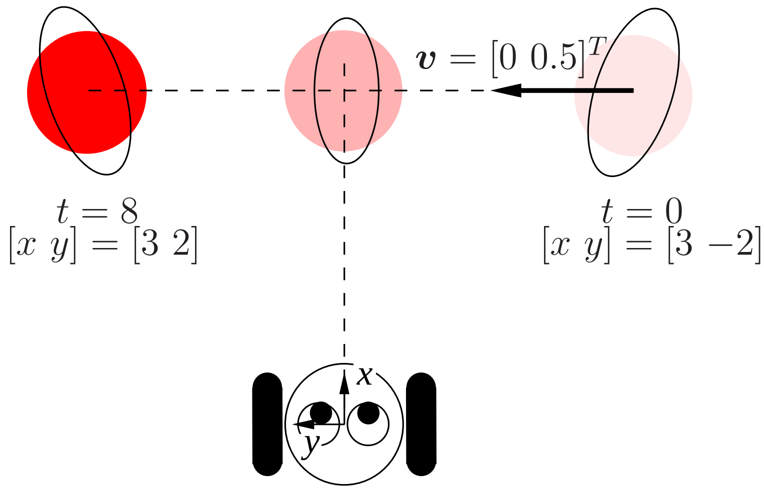

The "Process" Model

$$\hat{\boldsymbol{x}} = \begin{bmatrix} \hat{x} \\ \hat{y} \\ \dot{\hat{x}} \\ \dot{\hat{y}} \end{bmatrix}$$

$$\begin{aligned} x_k^- &= x_{k-1}^+ + \dot{x}_{k-1}^+ \Delta t + (1/2)(u_x)_{k-1} \Delta t^2 \\ y_k^- &= y_{k-1}^+ + \dot{y}_{k-1}^+ \Delta t + (1/2)(u_y)_{k-1} \Delta t^2 \\ \dot{x}_k^- &= \dot{x}_{k-1}^+ + (u_x)_{k-1} \Delta t \\ \dot{y}_k^- &= \dot{y}_{k-1}^+ + (u_y)_{k-1} \Delta t \end{aligned}$$

A Kalman Filter for Roger

The "Process" Model

$$\hat{\boldsymbol{x}} = \begin{bmatrix} \hat{x} \\ \hat{y} \\ \dot{\hat{x}} \\ \dot{\hat{y}} \end{bmatrix} \qquad \begin{aligned} x_k^- &= x_{k-1}^+ + \dot{x}_{k-1}^+ \Delta t + (1/2)(u_x)_{k-1} \Delta t^2 \\ y_k^- &= y_{k-1}^+ + \dot{y}_{k-1}^+ \Delta t + (1/2)(u_y)_{k-1} \Delta t^2 \\ \dot{x}_k^- &= \dot{x}_{k-1}^+ + (u_x)_{k-1} \Delta t \\ \dot{y}_k^- &= \dot{y}_{k-1}^+ + (u_y)_{k-1} \Delta t \end{aligned}$$

$$\hat{\boldsymbol{x}}_k^- = \mathbf{A} \quad \hat{\boldsymbol{x}}_{k-1}^+ + \mathbf{B} \quad \boldsymbol{u}_{k-1}$$

$$\begin{bmatrix} \hat{x} \\ \hat{y} \\ \dot{\hat{x}} \\ \dot{\hat{y}} \end{bmatrix}_k^- = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \hat{x} \\ \hat{y} \\ \dot{\hat{x}} \\ \dot{\hat{y}} \end{bmatrix}_{k-1}^+ + \begin{bmatrix} \Delta t^2/2 & 0 \\ 0 & \Delta t^2/2 \\ \Delta t & 0 \\ 0 & \Delta t \end{bmatrix} \begin{bmatrix} u_x \\ u_y \end{bmatrix}_{k-1}$$

consider the process with no control inputs (i.e. $\boldsymbol{u} = \mathbf{0}$)

A Kalman Filter for Roger

The "Process" Model

the process model is subject to noise $\boldsymbol{w} \sim N(0, \mathbf{Q}_k)$

how is covariance matrix $\mathbf{Q}_k$ estimated?

consider isotropic *acceleration disturbances* $\boldsymbol{u}_{dist}$

$$\mathbf{Q}_k \approx \sigma^2_{proc}\, \mathbf{B}\mathbf{B}^T = \sigma^2_{proc} \begin{bmatrix} \frac{\Delta t^4}{4} & 0 & \frac{\Delta t^3}{2} & 0 \\ 0 & \frac{\Delta t^4}{4} & 0 & \frac{\Delta t^3}{2} \\ \frac{\Delta t^3}{2} & 0 & \Delta t^2 & 0 \\ 0 & \frac{\Delta t^3}{2} & 0 & \Delta t^2 \end{bmatrix}$$

which is constant for a fixed sample rate

A Kalman Filter for Roger

The "Sensor" Model

$$\hat{\boldsymbol{z}}_k \quad = \quad \mathbf{H}_k \quad \hat{\boldsymbol{x}}_k^- \qquad\qquad \boldsymbol{v}_k \sim N(0, \mathbf{R}_k)$$

$$\begin{bmatrix} z_x \\ z_y \end{bmatrix}_k = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix}_k^- \qquad \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix} \quad \text{or, maybe } \mathbf{J}\mathbf{J}^T?$$

A Kalman Filter for Roger