Lecture 25: Kalman filtering

Winter 2015

Ref: Hartmann, Ch. 8

25.1 Motivation

Commonly, we need to use data to help control, forecast, or estimate governing equations for a complex system, e. g. tracking a moving object, using an autopilot to control a plane, or making numerical weather forecasts. The **Kalman filter** enables the modeling system to constantly adjust to updated observations to make more skillful predictions - a form of model-data fusion.

25.2 Combining a model with sequential data; a simple example

Let's start with a very simple model system whose exact behavior at a discrete sequence of times is

$$x_n = ax_{n-1},\ n = 2, 3, \dots \qquad (25.2.1)$$

Assume we know this governing equation, but we never know the exact value of the model **state**x_n. At each time we make an unbiased observation x_n^o of x_n with variance σ_o^2 about the true value. After n times, what is our best guess $\hat{x}_n$ at x_n and what is its uncertainty, measured as an estimated variance $\hat{\sigma}_n^2$?

For $n = 1$ we have no prior information, so

$$\hat{x}_1 = x_1^o \qquad \hat{\sigma}_1^2 = \sigma_o^2.$$

For time $n > 1$ we have an unbiased estimate $\hat{x}_{n-1}$ from the previous time $n-1$, with some estimated variance $\hat{\sigma}_{n-1}^2$ about the true value x_{n-1}, and a new observation x_n^o with variance σ_o^2 about the true value x_n.

25.3 Update previous estimate

The estimate from time $n - 1$ together with the governing equation (25.2.1) implies a 'prior' estimate

$$x_n^p = a\hat{x}_{n-1} \qquad \text{with } \sigma_p^2 = a^2\hat{\sigma}_{n-1}^2. \qquad (25.3.1)$$

Now we can use the theory of optimal interpolation from the previous lecture to update the prior given the new observation. The important new wrinkle is that the model is being used not only to update the prior state but also the prior error covariance (or variance in this 1D case).

$$\hat{x}_n = x_n^p + k(x_n^o - x_n^p). \tag{25.3.2}$$

$$k = \frac{\hat{\sigma}^2}{\sigma_o^2} = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_o^2}, \tag{25.3.3}$$

$$\hat{\sigma}^2 = \frac{\sigma_p^2 \sigma_o^2}{\sigma_p^2 + \sigma_o^2} = (1-k)\sigma_p^2 \tag{25.3.4}$$

Equations (25.3.1), (25.3.2), (25.3.3), and (25.3.4) together allow us to sequentially update our best estimate of the state $\hat{x}_n$ and the variance $\hat{\sigma}_n^2$ of this estimate about the true value. An example is shown in the Matlab script **sequential_estimation_simple1D**.

If $a > 1$, after a few time steps, the estimated variance settles to an equilibrium value such that

$$\begin{aligned} \frac{1}{\hat{\sigma}^2} &= \frac{1}{\sigma_o^2} + \frac{1}{a^2 \hat{\sigma}^2} \\ &\Rightarrow \hat{\sigma}_n^2 \to \hat{\sigma}^2 = \sigma_o^2(1 - a^{-2}). \end{aligned}$$

The influence of prior observations in reducing uncertainty in the state estimates reaches a plateau, because the exponential growth of x magnifies the effects of their errors. If $a < 1$ the estimated variance will decrease monotonically with time, because the earlier observations become ever-stronger constraints due to the exponential decay of x.