Kalman Filters: Overview

- A Kalman Filter (or any estimation filter) is an algorithm/hunk_of_software which estimates some quantity of interest (the state) based on a collection of measurements.
- Like its close cousin, the oil filter, it seeks to remove the contaminating influence of unwanted stuff, and somehow return to the true/pure/ideal signal/lubricant.
- Kalman's seminal paper appeared in 1960, but it was preceded by other work and continues to inspire active research to this day. It is explained in many textbooks.
- A set of self-paced notes, complete with homework and worked solutions, is available on the unclassified system for a nominal fee. The exposition in those notes shows one approach to the subject among a large number of approaches.
- Tracking uses a KF to estimate the position/velocity of targets on the basis of measurements supplied by on-board sensors, and sometimes off-board information as well. Another widely used application for Kalman Filters is navigation error estimation.

Kalman Filters: Overview

- The KF makes certain assumptions about what is already known about the state, and about the nature and quality of the measurements. These assumptions are based on concepts from probability. Because the quantities to be estimated are often multi-dimensional, simple concepts from linear algebra are also used.

- Before looking at a full case, let's look at a very simple case: estimating a constant.

- Suppose that a constant, X, has been imperfectly estimated to be X_e, we let s_X denote its accuracy. Under "ordinary" situations the error distribution for the estimate is the famous "bell shaped curve", also known as a Gaussian or Normal distribution. We also presume that the estimation process has been unbiased, which means that the curve is centered at 0. s_X is a measure of its width. For this 1-D situation 68% of the probability lies between $\pm s_X$, 95% between $\pm 2s_X$, and 99.7% between $\pm 3s_X$.

- Also suppose that an imperfect, unbiased measurement of X has been made, call it Y, also with a Gaussian error distribution. Let s_Y be a measure of its accuracy.

- Somehow, we want to combine X_e with Y to get a new, improved estimate of X.

Kalman Filters: Overview

- The graph below indicates what might happen. The blue curve shows an estimate of $X_e=1$ for the state, while the red shows a measurement of $Y=-½$.

- The measurement indicates that an estimate smaller than 1 should be made, but the measurement is less accurate than the starting estimate. The prior estimate for X has $s=1$, while the measurement has $s=2$. How much weight should the measurement be given?

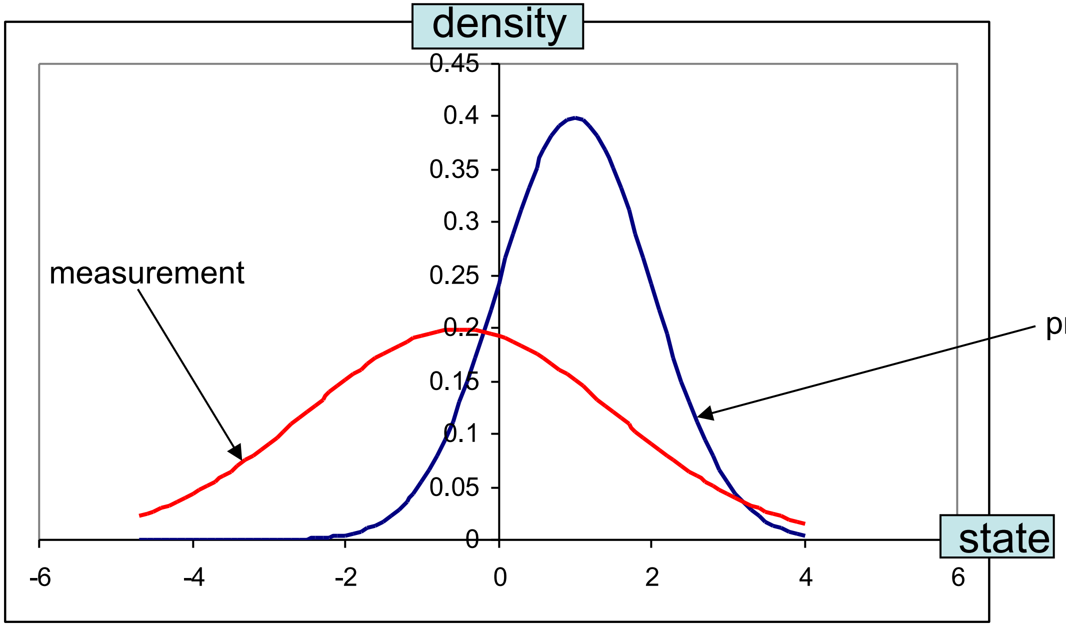

Kalman Filters: Overview

- The Kalman Filter blends the two potential values (1 and -½), taking a weighted average of the two. The weighting is proportional to the accuracy of each, as measured by the inverse of the variance ($1/s^2$). For this case, the inverse variances are $(1)^2$ (for the blue curve) and $(½)^2$ for the red curve, giving weights of 1/1.25 = 0.8 and 0.25/1.25 = 0.2. Our “new, improved” estimate is thus 0.8*1+0.2*(-1/2) = 0.7

- In addition, a new s (.89) is produced, tighter than before, because more information is behind it.

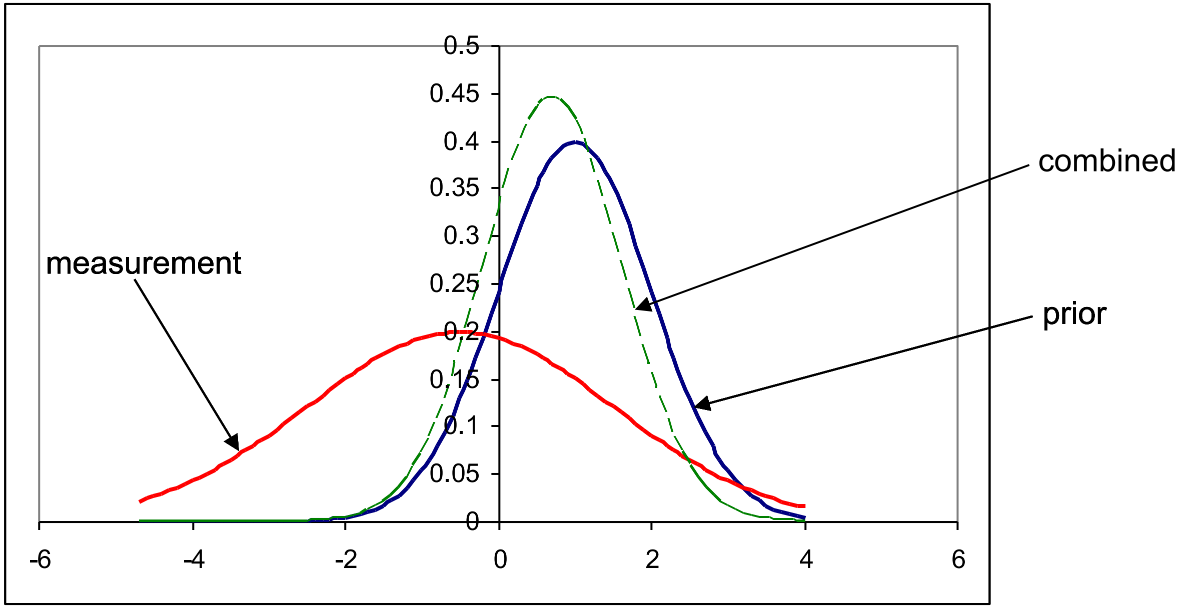

Kalman Filters: Overview

- What is the KF trying to do?

- It cannot hope to obtain unknowable "truth", so an error in an estimate is a given. The prior estimate and the measurement are both only roughly known.

- The goal is to use a method which will
  a) introduce no systematic bias in its estimate, and
  b) minimize the average error (squared) in the estimate.

- A fundamental assumption behind the derivation of the KF equations is that
  a) the noise in the measurements is Gaussian, or
  b) the measurement will be incorporated as a linear combination of the prior estimate and the measurement

  If neither of these holds then weird estimators can be constructed.

Kalman Filters: Overview

- The general model for a KF takes place with multidimensional estimates and multidimensional measurements. For example, we want to estimate the position and velocity of a target (6 dimensions), and we measure azimuth and elevation with a sensor (2 dimensions).

- The prior estimate (guess) is X, an n-dimensional state, which is presumed to have no bias in its formation, and has covariance matrix P (n-by-n). The (i,j) entry of P is the expected/mean/average value of the product of the i and j elements of X. In vector/matrix language, $P=E(X X^T)$, where X is considered to be a column vector.

- The measurement vector is Y, an m-dimensional column vector, which is a known linear function of X, contaminated by (unbiased Gaussian) noise of known magnitude. That is $\boxed{Y=HX+n}$, where Cov(n)=R. H is an m-by-n matrix, and R is an m-by-m matrix. (Note: in many books, the measurement is called Z)

- The KF forms a matrix, the "Kalman gain", as $K=PH^T(HPH^T+R)^{-1}$, and then updates both the state estimate and the covariance of the estimate as

$$\boxed{\begin{aligned} X &:= X+K(Y-HX) \\ P &:= P - KHP \end{aligned}}$$

Kalman Filters: Overview

- In general, of course, the measurement is not a linear function of the state. That is, instead of $Y=HX+n$ the actual measurement process is $Y=h(X)+n$, where h is non-linear. For this situation the equations have to be modified as

$$X := X+K(Y-h(X))$$
$$P := P - KHP$$

  where H is the m-by-n matrix of partial derivatives of the components of h relative to the elements of X, evaluated at the a-priori state estimate. This is called the "Extended Kalman Filter" (EKF).

- There is also another aspect to the estimation problem, namely the way that errors and accuracies propagate over time in the absence of new measurements. For example, if the state is a 6 dimensional estimate of an aircraft, then any errors in velocity estimation will turn into errors in position, which grow over time. Also, the starting accuracy of position needs to grow over time.

- The problem of propagation is discussed in the self study notes referenced above. It uses a model in which a differential equation is driven by white noise. (called a stochastic differential equation)

Kalman Filters: Overview

- A little more needs to be said about the Gaussian Distribution. In 1-dimension its density function is
$$f(x;\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/2\sigma^2}$$
This is the famous "bell shaped curve", centered at m with variance of s^2. The area under it is 1, for any values of m and s.

- This distribution is important because a random variable which arises as the independent combination of many other such variables, even if they are not even remotely Gaussian, will look like a Gaussian distribution. This fact is the famous "Central Limit Theorem".

- In higher dimensions the density looks like this: $f(X;\mu,P) = \frac{1}{\sqrt{(2\pi)^n |P|}} e^{-(X-\mu)^T P^{-1} (X-\mu)/2}$
where n is the dimension of X, and |P| denotes the determinant of P.

- A part of the exponent, $(X\text{-m})^T P^{-1}(X\text{-m})$, is noteworthy in its own right. It is the famous c^2 distribution, and is (amazingly) independent of P: it depends only on n, called the number of "degrees of freedom". Tracking uses this as part of its correlation logic.

- The density function is also known as the likelihood, and taking its natural logarithm gives, essentially, the c^2 value and another part mostly dependent on log|P|. This "log-likelihood" value is also used in tracking.

Kalman Filters: Overview

- In view of the fact that the EKF equations are thoroughly written out, it is natural to ask why the subject of Kalman Filtering is not just a closed subject (like solving the quadratic equation).

  1. The equations describing the state propagation over time are driven by noise whose properties are not really known (e.g. a target maneuvers)

  2. The non-linearity in the measurement equation may be severe relative to the size of the estimation error. This means that the EKF, which is just a linearized version of the KF, may be a poor approximation.

  3. Measurements (i.e. sensors) may not (do not) behave as advertised. In addition to unmodeled random errors (R matrix), there may be systematic errors (biases) in the measurements. These can come from sensor installation inaccuracies or radome distortions.

  4. Lack of modeling knowledge may lead P to become very small, which will reduce the Kalman gain and choke off the update process.

  5. For some applications (navigation) the size of the state vector may be a large computational burden. For other applications (tracking) the states may be small (usually 6 or 3 dimensions, sometimes 9), but there may be a very large number of them.

- For all these reasons, filters rarely follow the textbook formulation, but need to be adjusted/tuned/caressed by hand.