What is the temperature of the room I am in?

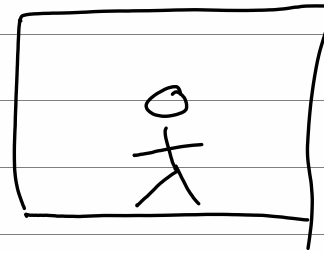

I'm going to estimate that it is 15°.

I'll call my "background" estimate x_b (for reasons which will become clear later)

After doing all the maths I know my updated, and final analysis, x_a, of what the temperature is is:

$$x_a = 15^\circ$$

I have no other information than my background guess so how can I change my guess to anything else?

Now imagine as well as my guess there is also an observation

Background (x_b) temp : 15°
Observation (y_o) temp : 20°

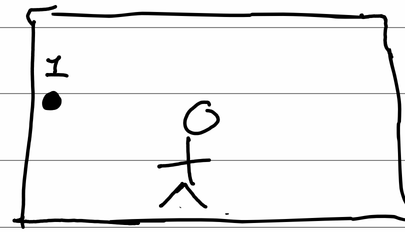

What is the temperature where I am standing?

Could just believe the observation, but I'm pretty good at estimating the temp.
"Obviously" with no more information the best estimate (x_a) is:

$$x_a = \frac{15 + 20}{2} = \underline{\underline{17.5°}}$$

Whilst this seems obvious, where does it actually come from? Why is this the best estimate?

It is derived from the least squares method (Gauss ~1795, Legendre 1805)

That is, we want:
"a value that minimises the sum of the squares of the error of each term".

We want to find x_a which minimises the sum of the squares of the error of the two observations x_b and y_o
e.g. x_a should minimise:

$$(x_a - x_b)^2 + (x_a - y_o)^2$$

There are a few ways to find the minimum, from any standard calculus course the minimum is found from the first derivative test (e.g. when it is 0)

$$f'(x_a) = 2(x_a - x_b) + 2(x_a - y_o)$$

The minimum is when this is equal to 0

$$2(x_a - x_b) + 2(x_a - y_o) = 0$$

$$(x_a - x_b) + (x_a - y_o) = 0$$

$$2x_a = x_b + y_o$$

$$x_a = \frac{x_b + y_o}{2}$$

Thus the "best guest" is the average of the 2 observations.

The problem with this least squares approach is that it is not sensitive to a change of units.

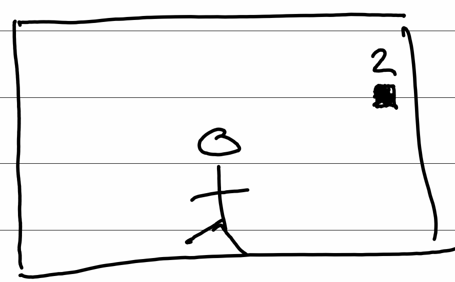

$x_b = 15^\circ$

Sensor 2 : 40° (measures 2x temp [for some reason])

Nothing has actually changed in the scenario so we should expect the same answer, but:

$$f(x_a) = (x_a - 15)^2 + (2x_a - 40)^2$$

$$f'(x_a) = 2(x_a - 15) + 2(2x_a - 40)$$

$$2(x_a - 15) + 2(2x_a - 40) = 0$$
$$(x_a - 15) + (2x_a - 40) = 0$$
$$3x_a = 55$$
$$x_a = \frac{55}{3} = 18\tfrac{1}{3}$$

The previous x_a was 17.5 so a simple change of unit (temp measurement to 2x temp measurement) changes the result.

A more fundamental weakness is that the method does not take into account measurement uncertainty.

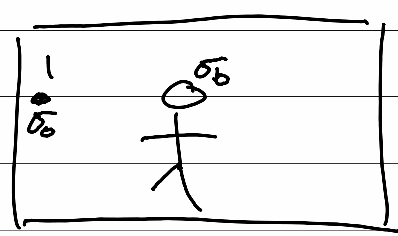

Lets assume that the uncertainty of x_b is σ_b and y_o is σ_o

e.g. $x_b = 15°$ with std. dev. $\sigma_b = 5°$
$y_o = 20°$ with std. dev $\sigma_o = 1°$

Clearly our best estimate of 17.5° is no longer the best we could do.

First we need some background assumptions.

Assuming that the sensors are unbiased
$(E[x_b] = E[x] = E[y_o])$
and the errors are uncorrelated (errors in x_b and errors in y_o are independent)

Rather than a straight average we clearly need a weighted average e.g.
$$x_a = \alpha x_b + \beta y_o.$$
and we need a method to determine α and β.
Assuming that the weights must equal 1 rewrite $x_a = \alpha x_b + \beta y_o$ without the α term

Task audience to attempt;

$$\alpha = 1 - \beta \Rightarrow \quad x_a = (1-\beta)x_b + \beta y_o$$
$$= x_b - \beta x_b + \beta y_o$$
$$= x_b + \beta(y_o - x_b)$$

This says that the optimum analysis can be found by updating the x_b estimate by some multiple of the difference between x_b and y_o.

For x_a to be the best estimate of the truth (x) then the variance needs to be minimised (e.g. pick β to min $Var(x_a)$)

$$Var(x_a) = E[x_a^2] - E[x_a]^2$$

Using the fact that expectation is a linear operator : $E[aP + bQ] = aE[P] + bE[Q]$ and $\alpha + \beta = 1$ then

$$\beta = \frac{\sigma_b^2}{\sigma_b^2 + \sigma_o^2}$$

This gives an optimal update equation as:

$$x_a = x_b + \frac{\sigma_b^2}{\sigma_b^2 + \sigma_o^2}(y_o - x_b)$$

which using our example values:

$$x_a = 15 + \frac{5^2}{5^2+1^2}\left(20-15\right)$$

$$= 15 + \frac{25}{26}(5) \approx \underline{\underline{19.81}}$$

This is expected since the 20° observation has a much lower std. dev. and we should "trust" that more.

Obviously we do not usually have a 1-state (temp where you are) system and we need an approach that allows vectors to be used.

Let x be our "truth" state, a vector of values we are interested in (say the TEC at each grid point on a $5° \times 5°$ resolution map, that would mean we have 36 lat points and 72 lon points giving a state of size $36 \times 72 = 2592$ elements.

Y is going to be our observations (usually much smaller than the size of the state).

By some method we want to find $p(x|y)$ that is the posterior distribution some probability distribution of what the state of interest looks like (x) given the data we have (y).

Using Bayes' Theorem we know:

$$p(x|y) = \frac{p(y|x)\,p(x)}{p(y)}$$

$p(y|x)$ is the data distribution, it is simply the prob. distribution of the data.

Keep in mind that this is the data conditioned on x. If y has imperfect observations of x then this quantifies the distribution of measurement error, including biases.

$p(x)$ is the prior distribution, our "background guess", often in real applications this is provided from a climatology or first-principles model.

$p(y)$ can be thought of as a normalising constant and most texts miss it out entirely.

You may see: $p(x|y) \propto p(y|x)p(x)$

an extension of our

Imagine 1: temperature example, our prior (background) is Normally (Gaussian) distributed:

$$X \sim N(\mu, v^2)$$

mean μ, variance v^2

And assume we have n independent, non biased, observations of x: $y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}$ with variance σ^2

The data distribution is distributed by $N(x, \sigma^2)$

Recall for a Normal distribution:

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

So $$p(y|x) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(\frac{y_i - x}{\sigma}\right)^2}$$

$$\propto \exp\left\{-\frac{1}{2}\sum_{i}^{n}\left(\frac{y_i - x}{\sigma}\right)^2\right\}$$

change from e.

So Bayes' gives us

$$p(x|y) \propto p(y|x)\,p(x)$$

$$= \exp\left\{-\frac{1}{2}\sum_{1}^{n}\left(\frac{y_i - x}{\sigma}\right)^2\right\} \exp\left\{-\frac{1}{2}\sum_{1}^{n}\left(\frac{x-\mu}{v}\right)^2\right\}$$

$$= \exp\left\{-\frac{1}{2}\sum_{1}^{n}\left[\left(\frac{y_i - x}{\sigma}\right)^2 + \left(\frac{x-\mu}{v}\right)^2\right]\right\}$$

Maybe not entirely obv. but not hard

$$\propto \exp\left\{-\frac{1}{2}\left[x^2\left(\frac{n}{\sigma^2} + \frac{1}{v^2}\right) - 2\left(\sum_{1}\frac{y_i}{\sigma^2} + \frac{\mu}{v^2}\right)x\right]\right\}$$

This is simply the product of 2 Gaussians and so that results in a Gaussian. By completing the square we can show

$$x|y \sim N\left(\underbrace{\left[\frac{n}{\sigma^2} + \frac{1}{v^2}\right]^{-1}\left[\sum\frac{y_i}{\sigma^2} + \frac{\mu}{v^2}\right]}_{\text{mean}},\ \underbrace{\left[\frac{n}{\sigma^2} + \frac{1}{v^2}\right]^{-1}}_{\text{variance}}\right)$$

The posterior mean $E(x|y)$ is then:

$$E(x|y) = \frac{\sigma^2 v^2}{\sigma^2 + n v^2}\left(\frac{n\bar{y}}{\sigma^2} + \frac{\mu}{v^2}\right)$$

where $\bar{y}$ is the mean of y. $= \sum_{i}^{n} \frac{y_i}{n}$

$$= w_y \bar{y} + w_\mu \mu$$

"weight" of y $= \frac{n v^2}{n v^2 + \sigma^2}$

weight of μ $= \frac{\sigma^2}{n v^2 + \sigma^2}$

Worth noting that $\frac{n v^2}{n v^2 + \sigma^2} + \frac{\sigma^2}{n v^2 + \sigma^2} = 1$

$$w_y + w_\mu = 1$$

the weights add to 1. Look familiar....

In fact we can rewrite $E(x|y)$ as:

$$E(x|y) = \mu + \left(\frac{n v^2}{\sigma^2 + n v^2}\right)(\bar{y} - \mu) = \boxed{\mu + K(\bar{y} - \mu)}$$

which is exactly what we had before!

We can also write down the posterior variance:

$$Var(X|y) = (1-K)v^2$$

So the posterior variance is "updated" from the prior variance by K (the "gain").

Just checking everything is right lets re use our previous example

$$X_D \sim N(15, 5^2)$$

remember this is variance

$$Y|X \sim N(20, 1)$$

$$K = \frac{nv^2}{\sigma^2 + nv^2} = \frac{1 \cdot 5^2}{1^2 + 1 \cdot 5^2} = \frac{25}{26}$$

$$E(X|y) = \mu + K(\bar{y} - \mu)$$

$$E(X|y) = 15 + \frac{25}{26}(20 - 15) \approx \underline{\underline{19.81}}$$

The Same!

We can also find the posterior variance:

$$\mathrm{Var}(x|y) = (1-K)v^2$$
$$= \left(1 - \frac{25}{26}\right)5^2 \approx 0.96$$

Even in this simple case we are relying on two key assumptions:

1) The data is unbiased $E[y] = E[x]$

2) Data is normally distributed (product of 2 Gaussians is Gaussian)

Now go back to our state of TECs across the globe.

Assume $X \sim N(\mu, B)$ and we know μ and B.

covariance matrix

$|X| = p$.

and we have a data vector of n observations with distribution

$$Y|x \sim N(Hx, R)$$

H is an $n \times p$ matrix which maps from state to observation "space" (obs operator) again assume H and R are known.

Exactly as before $p(x|y) \propto p(y|x)p(x)$ and

$$X|y \sim N\Big(\underbrace{(H^TR^{-1}H + B^{-1})^{-1}(H^TR^{-1}y + B^{-1}\mu)}_{\text{mean}}, \underbrace{(H^TR^{-1}H + B^{-1})^{-1}}_{\text{variance}}\Big)$$

The posterior mean (with basic pre-calculus) can be written as:

$$E(x|y) = \mu + BH^T(R + HBH^T)^{-1}(y - H\mu)$$

$$= \mu + K(y - H\mu)$$

and

$$Var(x|y) = (I - KH)B$$

identity.

The key DA equations.

The only remaining concept to consider (albeit crucial) is time. The KF is a process which updates over time.

We can formally define this using Markov processes and Bayes' again, but in the interest of time the easiest way to consider this is using a linear model M which describes how the state changes from time t to t+1.

Then:

$$X_b^{t+1} = M^{t,t+1} X_a^t$$

$$B^{t+1} = M^{t,t+1} A \left(M^{t,t+1}\right)^T + Q^t$$

model error covariance.

And here lies the 3rd key assumption — linearity.

For a Gaussian, non-biased, linear system the Kalman filter is <u>optimal</u>.

Unfortunately not many real systems meet these requirements! Hence the wide variety of different KFs.

(e.g.)

Extended KF – non linear version (linearizes about a point using 1st order Taylor series expansion)

Band limited KF – only "saves" part of covariance matrix (which can be <u>very</u> big!)

Ensemble KF – most common. Multiple instances of the state are used to estimate covariance matrix and can handle non-linear model dynamics.

Recall the definition of a covariance matrix:

$$\text{cov}(X) = E\left[(X - E[X])(X - E[X])^T\right]$$

Imagine we have a set of k instances of the state vector $\{x^i\}$ for $i = 1, \ldots, k$.
(these could be with slightly different driver conditions for example)

The ensemble mean is defined as:

$$\bar{x} = \frac{1}{k}\sum_{i=1}^{k} x^k$$

If we define X_b to be a <u>perturbation matrix</u>

capital

given by:

$$X_b = \left(x^1 - \bar{x} \quad x^2 - \bar{x} \quad \ldots \quad x^k - \bar{x}\right)$$

each column is the associated ensemble member with the ensemble mean removed then $B \approx \frac{1}{k-1} X_b X_b^T$

k-1 rather than k for Bessels correction

which comes exactly from the covariance definition.

Assuming you keep adding orthogonal ensemble members then

$$\frac{1}{k-1} X_b X_b^T \longrightarrow B \quad \text{as} \quad k \to \infty.$$

This simple redefinition of B:

1) allows non-linear dynamics

2) reduces computational cost

3) lets us use the KF if we don't know B — it can be very hard to write down!

4) Opens up a world of fun....

There isn't one <u>optimal</u> way of using the ensemble KF (EnKF) so we can do lots of fun things e.g.

- localised EnKF
- transform "
- square-root EnKF
- local transform "
- etc. etc.

each have different pros and cons and always more being come up with!

Note:
- for EnKF errors in sampling decrease proportional to $1/\sqrt{k}$
- spurious relationships can be found between variables

Sean Elvidge s.elvidge@bham.ac.uk
Birmingham, UK
August 2021