From Optimal Fusion to the Kalman Measurement Update

Col Juan Jurado

Goal. Derive the minimum-variance linear fusion of two *independent* Gaussian estimates. We first solve the 1D case (including the optimization for w_1, w_2), derive the fused variance, rearrange into a *prior + innovation* form with z_1 treated as the prior and z_2 as the new measurement, then upgrade to 2D/ND by replacing scalars with vectors/matrices. Finally, we show how introducing a measurement matrix $\boldsymbol{H}$ converts "fusion of two estimates" into the Kalman measurement update.

1D Optimal Fusion (Two Scalars)

Problem statement

Let z_1 and z_2 be two unbiased, independent estimates of the same scalar x:

$$z_1 = x + v_1, \quad v_1 \sim \mathcal{N}(0, \sigma_1^2), \qquad z_2 = x + v_2, \quad v_2 \sim \mathcal{N}(0, \sigma_2^2), \qquad v_1 \perp v_2.$$

Seek a linear fused estimate

$$\hat{x} = w_1 z_1 + w_2 z_2, \qquad \text{subject to } w_1 + w_2 = 1.$$

Cost function (fused variance)

$$\hat{x} - x = w_1 v_1 + w_2 v_2 \quad \Rightarrow \quad \sigma_{\hat{x}}^2 = \mathrm{Var}(\hat{x} - x) = w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2.$$

Solve for optimal weights

With $w_2 = 1 - w_1$:

$$\sigma_{\hat{x}}^2(w_1) = w_1^2 \sigma_1^2 + (1 - w_1)^2 \sigma_2^2$$

$$\frac{d}{dw_1} \sigma_{\hat{x}}^2 = 2 w_1 \sigma_1^2 - 2(1 - w_1)\sigma_2^2 = 0 \quad \Rightarrow \quad w_1^\star = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}, \quad w_2^\star = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}.$$

Fused estimate and fused variance

$$\hat{x} = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} z_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} z_2 = \frac{\frac{1}{\sigma_1^2} z_1 + \frac{1}{\sigma_2^2} z_2}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}}, \qquad \boxed{\sigma_{\hat{x}}^2 = \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}}$$

Prior + innovation form (take z_1 as prior)

$$\hat{x} = z_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}(z_2 - z_1) = \underbrace{z_1}_{\text{prior}} + \underbrace{K}_{\text{optimal weight}} \underbrace{(z_2 - z_1)}_{\text{innovation}}, \qquad K = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}.$$

Upgrade to ND (Two Vectors, Same Quantity)

Problem statement

Let $\boldsymbol{x} \in \mathbb{R}^n$ and two independent estimates:

$$\boldsymbol{z}_1 = \boldsymbol{x} + \boldsymbol{v}_1, \quad \boldsymbol{v}_1 \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_1), \qquad \boldsymbol{z}_2 = \boldsymbol{x} + \boldsymbol{v}_2, \quad \boldsymbol{v}_2 \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_2), \qquad \boldsymbol{v}_1 \perp \boldsymbol{v}_2.$$

Seek a linear unbiased fusion

$$\hat{\boldsymbol{x}} = \boldsymbol{W}_1 \boldsymbol{z}_1 + \boldsymbol{W}_2 \boldsymbol{z}_2, \qquad \boldsymbol{W}_1 + \boldsymbol{W}_2 = \boldsymbol{I}.$$

Optimal fusion result (information form)

$$\boxed{\boldsymbol{P} = (\boldsymbol{R}_1^{-1} + \boldsymbol{R}_2^{-1})^{-1}} \qquad \boxed{\hat{\boldsymbol{x}} = \boldsymbol{P}(\boldsymbol{R}_1^{-1}\boldsymbol{z}_1 + \boldsymbol{R}_2^{-1}\boldsymbol{z}_2)}$$

Prior + innovation form

$$\hat{\boldsymbol{x}} = \boldsymbol{z}_1 + \boldsymbol{P}\boldsymbol{R}_2^{-1}(\boldsymbol{z}_2 - \boldsymbol{z}_1) = \underbrace{\boldsymbol{z}_1}_{\text{prior}} + \underbrace{\boldsymbol{K}}_{\text{optimal weight}} \underbrace{(\boldsymbol{z}_2 - \boldsymbol{z}_1)}_{\text{innovation}}, \qquad \boldsymbol{K} = \boldsymbol{P}\boldsymbol{R}_2^{-1}.$$

Bring in Measurements: Turn the "Second Estimate" into an Innovation

In a Kalman filter, the prior is the predicted state estimate:

$$\boldsymbol{z}_1 \triangleq \hat{\boldsymbol{x}}^-, \qquad \boldsymbol{R}_1 \triangleq \boldsymbol{P}^-.$$

The new information is a measurement:

$$\boldsymbol{z} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{v}, \qquad \boldsymbol{v} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}).$$

Rather than fusing $\hat{\boldsymbol{x}}^-$ directly with $\boldsymbol{z}$ (different spaces), build a **predicted measurement** and a **measurement residual**:

$$\hat{\boldsymbol{z}}^- \triangleq \boldsymbol{H}\hat{\boldsymbol{x}}^-, \qquad \boldsymbol{y} \triangleq \boldsymbol{z} - \hat{\boldsymbol{z}}^- = \boldsymbol{z} - \boldsymbol{H}\hat{\boldsymbol{x}}^-.$$

Linearize the measurement about the prior:

$$\boldsymbol{z} \approx \boldsymbol{H}\hat{\boldsymbol{x}}^- + \boldsymbol{H}(\boldsymbol{x} - \hat{\boldsymbol{x}}^-) + \boldsymbol{v} \quad \Rightarrow \quad \boldsymbol{y} \approx \boldsymbol{H}(\boldsymbol{x} - \hat{\boldsymbol{x}}^-) + \boldsymbol{v}.$$

Define the error state $\delta\boldsymbol{x} \triangleq \boldsymbol{x} - \hat{\boldsymbol{x}}^-$. Then:

$$\boldsymbol{y} \approx \boldsymbol{H}\,\delta\boldsymbol{x} + \boldsymbol{v}.$$

So the "second estimate" we want to fuse with the prior is now an *estimate of the error* $\delta\boldsymbol{x}$, whose measurement is $\boldsymbol{y}$.

From ND Fusion to the Linear Kalman Measurement Update (Same-Space Substitution)

The ND fusion result requires two independent Gaussian estimates of the *same quantity in the same space*. In a linear Kalman filter, the prior is in *state space* ($\hat{\boldsymbol{x}}^- \in \mathbb{R}^n$), while the measurement is in *measurement space* ($\boldsymbol{z} \in \mathbb{R}^m$). We therefore perform the fusion in measurement space first.

Step 1: Convert the prior into a predicted measurement

Given the linear measurement model

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

the prior state estimate implies a **predicted measurement**

$$\hat{\boldsymbol{z}}^{-} \triangleq \boldsymbol{H}\hat{\boldsymbol{x}}^{-},$$

with covariance (propagated through $\boldsymbol{H}$)

$$\boldsymbol{R}_1 \triangleq \mathrm{Var}(\hat{\boldsymbol{z}}^{-} - \boldsymbol{H}\boldsymbol{x}) = \boldsymbol{H}\boldsymbol{P}^{-}\boldsymbol{H}^{T}.$$

The actual sensor provides

$$\boldsymbol{z} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{v}, \qquad \boldsymbol{R}_2 \triangleq \mathrm{Var}(\boldsymbol{v}) = \boldsymbol{R}.$$

Step 2: Fuse two independent estimates of the same quantity (measurement space)

Now we have two estimates of the same measurement-space quantity $\boldsymbol{H}\boldsymbol{x}$:

$$\boldsymbol{z}_1 \triangleq \hat{\boldsymbol{z}}^{-} \sim \mathcal{N}(\boldsymbol{H}\boldsymbol{x}, \boldsymbol{R}_1), \qquad \boldsymbol{z}_2 \triangleq \boldsymbol{z} \sim \mathcal{N}(\boldsymbol{H}\boldsymbol{x}, \boldsymbol{R}_2), \qquad \text{independent.}$$

Apply the ND fusion result:

$$\hat{\boldsymbol{z}}^{+} = \boldsymbol{z}_1 + \boldsymbol{K}_z(\boldsymbol{z}_2 - \boldsymbol{z}_1), \qquad \boldsymbol{K}_z = \boldsymbol{P}_z\boldsymbol{R}_2^{-1}, \qquad \boldsymbol{P}_z = (\boldsymbol{R}_1^{-1} + \boldsymbol{R}_2^{-1})^{-1}.$$

Define the **innovation**:

$$\boldsymbol{y} \triangleq \boldsymbol{z} - \hat{\boldsymbol{z}}^{-}.$$

Then

$$\hat{\boldsymbol{z}}^{+} = \hat{\boldsymbol{z}}^{-} + \boldsymbol{K}_z\,\boldsymbol{y}.$$

It is convenient to express this weighting using the **innovation covariance**:

$$\boldsymbol{S} \triangleq \mathrm{Var}(\boldsymbol{y}) = \boldsymbol{H}\boldsymbol{P}^{-}\boldsymbol{H}^{T} + \boldsymbol{R}.$$

One can show the fused-measurement gain is equivalently

$$\boldsymbol{K}_z = \boldsymbol{H}\boldsymbol{P}^{-}\boldsymbol{H}^{T}\,\boldsymbol{S}^{-1}.$$

Step 3: Map the fused measurement back into a state update

A change in state produces a change in predicted measurement:

$$\Delta\hat{\boldsymbol{z}} = \boldsymbol{H}\,\Delta\hat{\boldsymbol{x}}.$$

We choose a state correction of the form

$$\Delta\hat{\boldsymbol{x}} = \boldsymbol{K}\,\boldsymbol{y}$$

so that the corresponding measurement correction matches the fused-measurement correction in a minimum-variance sense. This yields the familiar **Kalman gain**:

$$\boxed{\boldsymbol{K} = \boldsymbol{P}^{-}\boldsymbol{H}^{T}\,\boldsymbol{S}^{-1} = \boldsymbol{P}^{-}\boldsymbol{H}^{T}\big(\boldsymbol{H}\boldsymbol{P}^{-}\boldsymbol{H}^{T} + \boldsymbol{R}\big)^{-1}.}$$

Therefore the **linear Kalman measurement update** is

$$\boxed{\hat{\boldsymbol{x}}^{+} = \hat{\boldsymbol{x}}^{-} + \boldsymbol{K}\big(\boldsymbol{z} - \boldsymbol{H}\hat{\boldsymbol{x}}^{-}\big).}$$

Updated covariance

$$\boxed{\boldsymbol{P}^{+} = (\boldsymbol{I} - \boldsymbol{K}\boldsymbol{H})\boldsymbol{P}^{-}.}$$

Equivalently,

$$\boldsymbol{P}^{+} = \boldsymbol{P}^{-} - \boldsymbol{K}\boldsymbol{H}\boldsymbol{P}^{-},$$

which highlights that the measurement update subtracts uncertainty along the directions the measurement observes.