Lecture 29
Kalman Filter

Mark Hasegawa-Johnson

These slides are in the public domain

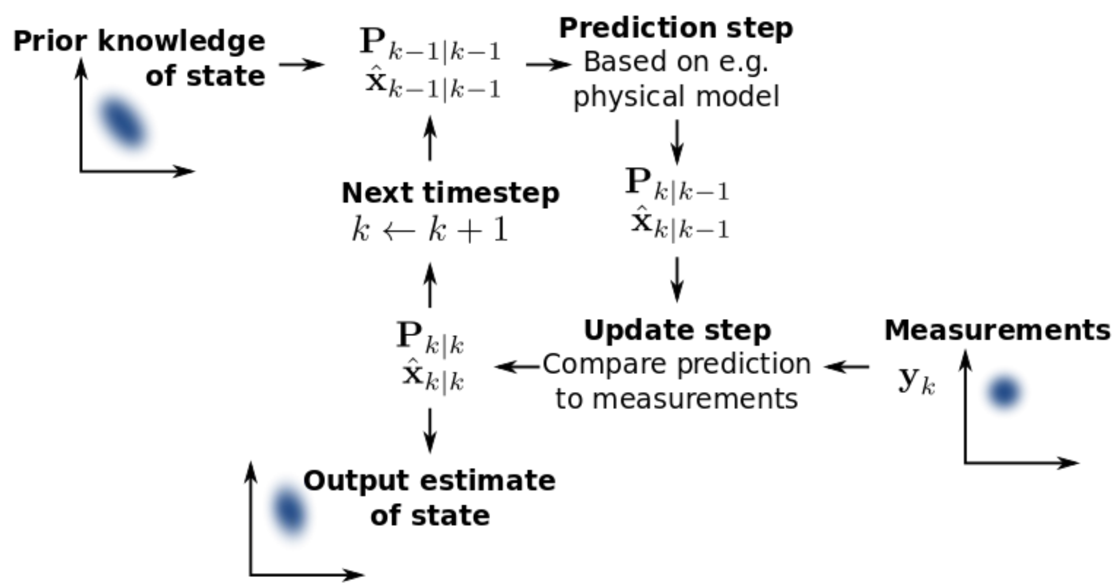

Public domain image,
https://commons.wikimedia.org/wiki/File:Basic_concept_of_Kalman_filtering.svg

Outline

- Tracking an object from noisy observations
- Prediction
- Update

Tracking an object from noisy observations

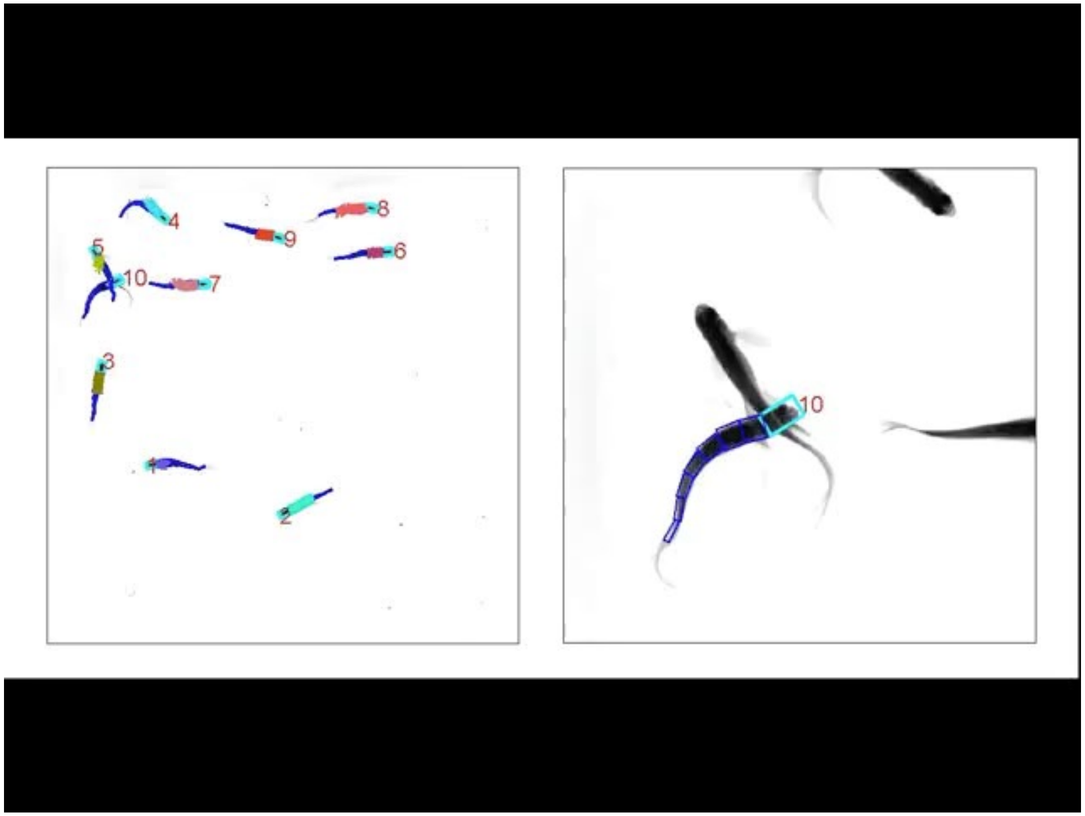

Tracking an object from noisy observations

- Y_t = current position of the object
- X_t = noisy observation of the object
- Goal: find $p(y_t|x_0, \dots, x_t)$

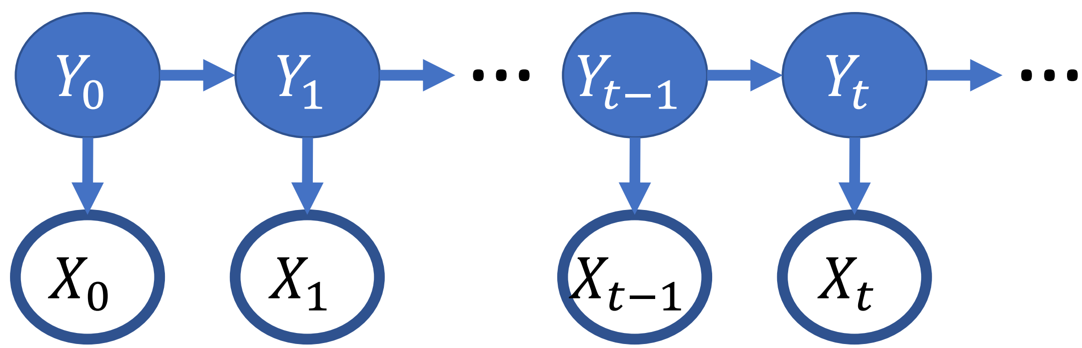

Outline

- Tracking an object from noisy observations
- Prediction
- Update

Prediction: Probability Distribution

- Suppose we already know $p(y_{t-1}|x_0, \ldots, x_{t-1})$.
- Can we find $p(y_t|x_0, \ldots, x_{t-1})$?
- Yes:

$$p(y_t|x_0, \ldots, x_{t-1}) = \sum_{y_{t-1}} p(y_{t-1}|x_0, \ldots, x_{t-1}) p(y_t|y_{t-1})$$

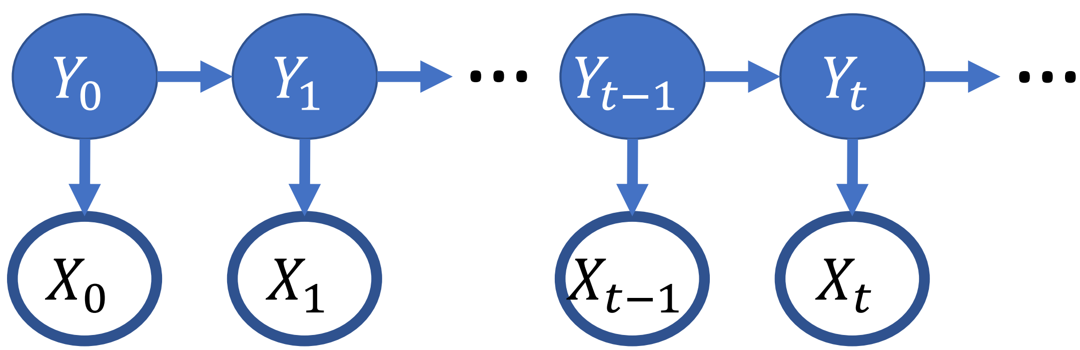

Prediction: Mean

What is its expected value?

$$E[Y_t|x_0, \ldots, x_{t-1}] = E[Y_{t-1}|x_0, \ldots, x_{t-1}] + E[\Delta]$$

… where $\Delta = Y_t - Y_{t-1}$ is the amount of movement in one second. For example, if an object is moving 10 m/s, then $E[\Delta] = 10$.

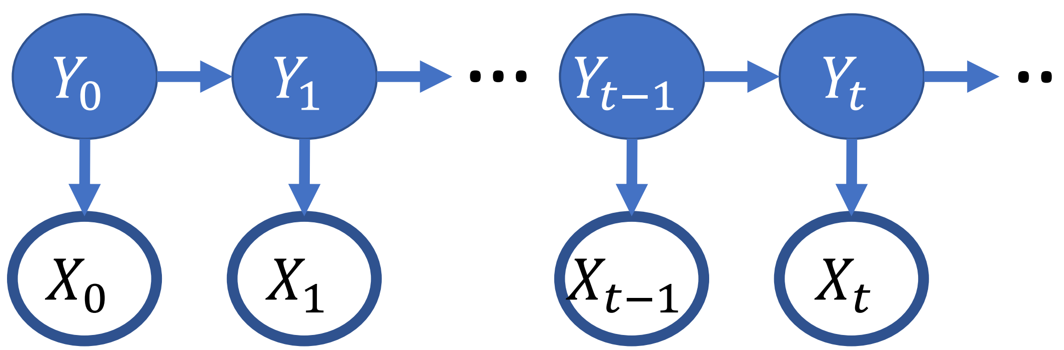

Prediction: Mean

Notation: define

$$\mu_{t|t-1} = E[Y_t | x_0, \ldots, x_{t-1}]$$
$$\mu_{t-1|t-1} = E[Y_{t-1} | x_0, \ldots, x_{t-1}]$$
$$\mu_\Delta = E[\Delta]$$

Then if we already know $\mu_{t-1|t-1}$ and μ_Δ, we can find

$$\mu_{t|t-1} = \mu_{t-1|t-1} + \mu_\Delta$$

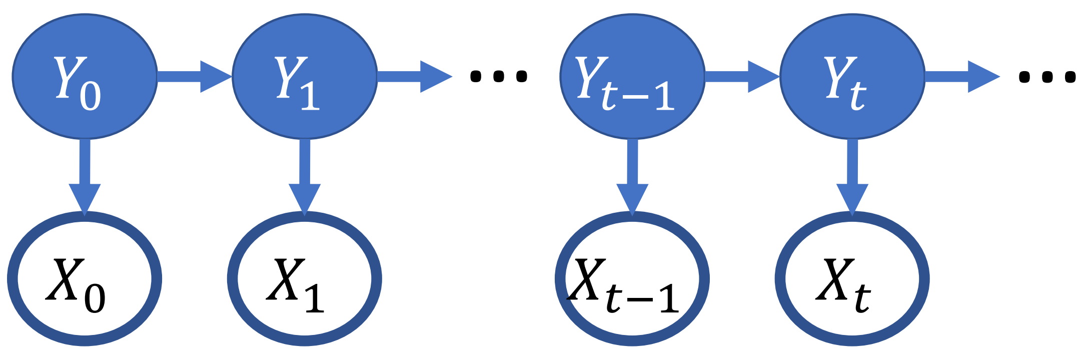

Prediction: Variance

What is its variance? If we assume that Δ and Y_{t-1} are independent, we get

$$\mathrm{Var}(Y_t|x_0,\ldots,x_{t-1}) = \mathrm{Var}(Y_{t-1}|x_0,\ldots,x_{t-1}) + \mathrm{Var}(\Delta)$$

For example, the object might be moving at 10m/s, but its velocity might have a standard deviation of 2m/s, so $\mathrm{Var}(\Delta) = 4$.

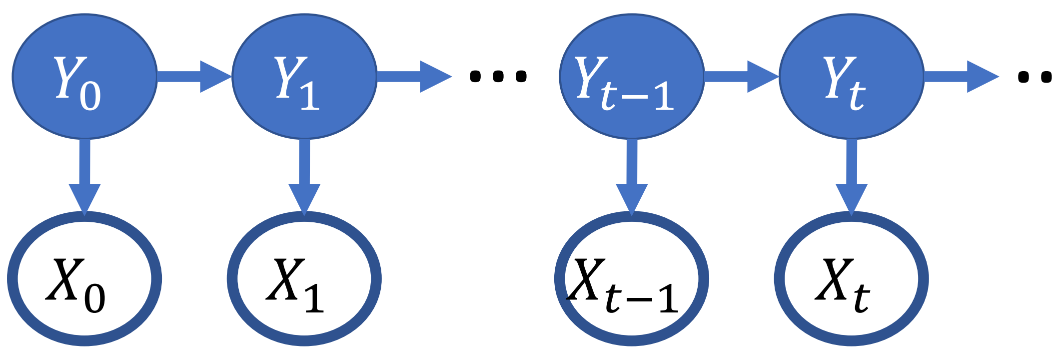

Prediction: Variance

Notation:

$$\sigma^2_{t|t-1} = \mathrm{Var}(Y_t | x_0, \ldots, x_{t-1})$$
$$\sigma^2_{t-1|t-1} = \mathrm{Var}(Y_{t-1} | x_0, \ldots, x_{t-1})$$
$$\sigma^2_{\Delta} = \mathrm{Var}(\Delta)$$

Then if we already know $\sigma^2_{t-1|t-1}$ and σ^2_{Δ}, we can find

$$\sigma^2_{t|t-1} = \sigma^2_{t-1|t-1} + \sigma^2_{\Delta}$$

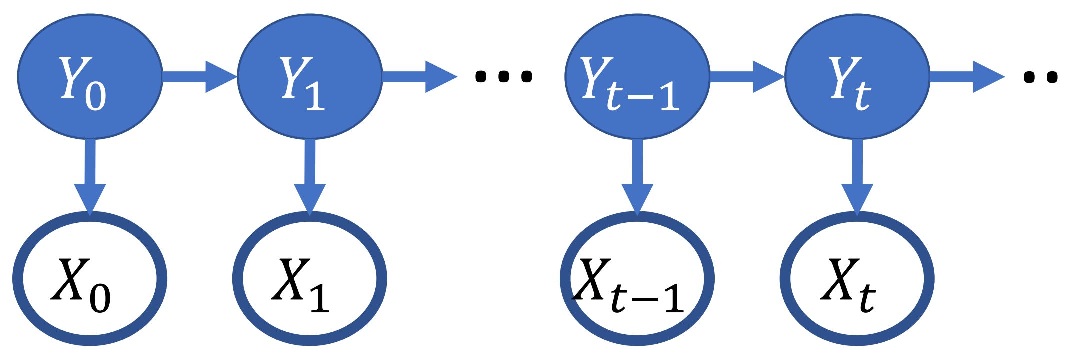

Prediction

If we know the object's location at time t-1, with some degree of uncertainty expressed by the variance $\sigma^2_{t-1|t-1}$, then we can guess where it will be at time t, with a slightly greater uncertainty caused by our uncertainty about its velocity:

$$\mu_{t|t-1} = \mu_{t-1|t-1} + \mu_\Delta$$
$$\sigma^2_{t|t-1} = \sigma^2_{t-1|t-1} + \sigma^2_\Delta$$

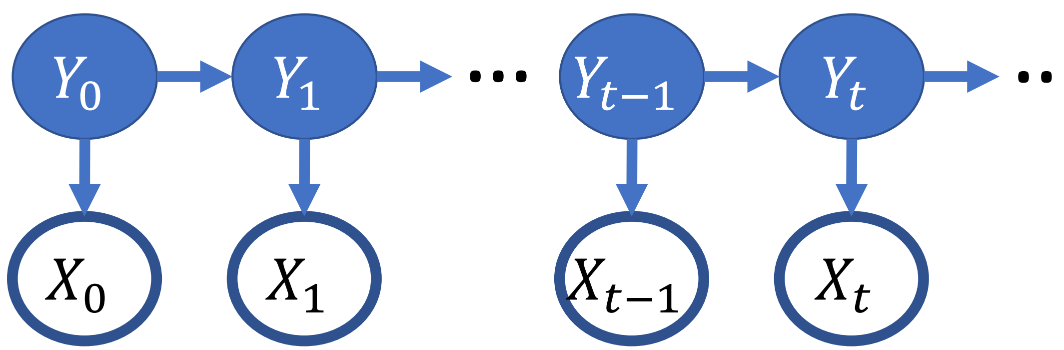

Outline

- Tracking an object from noisy observations
- Prediction
- Update

Update based on observations

The prediction step gave us $p(y_t|x_0, \ldots, x_{t-1})$. Now suppose we have a new observation, x_t. Can we use the new observation to improve our estimate of y_t?

In other words, can we find $p(y_t|x_0, \ldots, x_{t-1}, x_t)$?

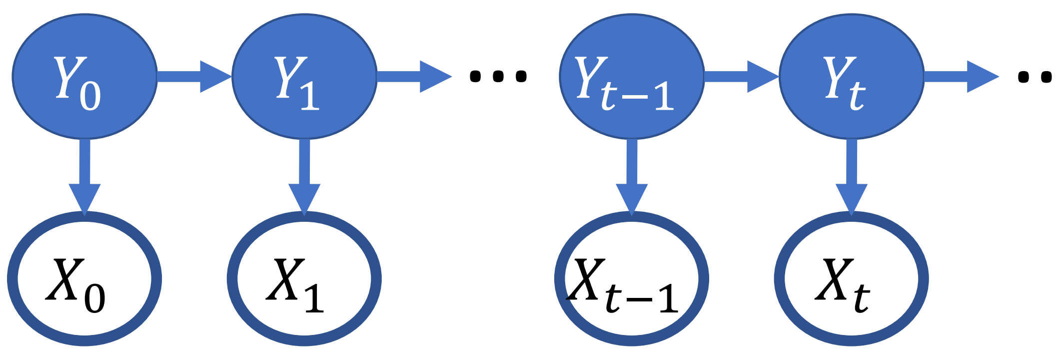

Kalman Filter: the independent noise assumption

- The Kalman filter assumes that Y_t is Gaussian, and that $X_t = Y_t + \epsilon$, where ϵ is some independent Gaussian measurement noise.

Under this assumption,

$$E[X_t | x_0, \ldots, x_{t-1}] = \mu_{t|t-1} + \mu_\epsilon$$
$$\mathrm{Var}(X_t | x_0, \ldots, x_{t-1}) = \sigma^2_{t|t-1} + \sigma^2_\epsilon$$

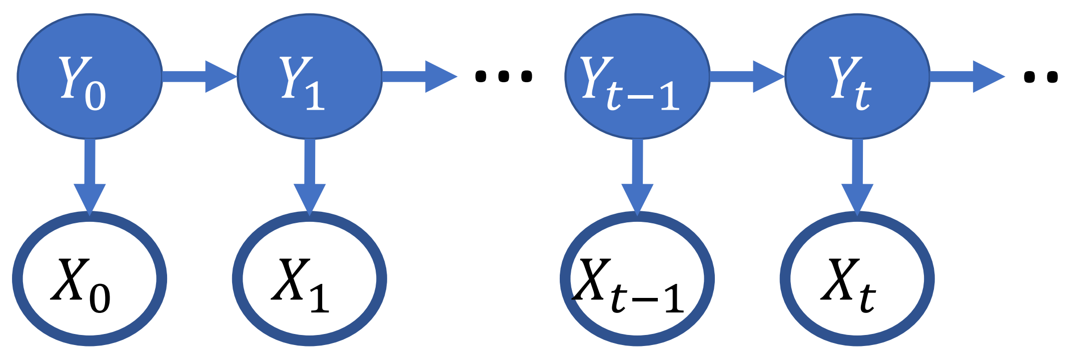

The Kalman gain

The ratio of the variances of Y_t and X_t is called the Kalman gain. It's the degree to which you trust the measurement x_t. The higher it is, the more you trust x_t:

$$k_t = \frac{\mathrm{Var}(Y_t | x_0, \ldots, x_{t-1})}{\mathrm{Var}(X_t | x_0, \ldots, x_{t-1})} = \frac{\sigma_{t|t-1}^2}{\sigma_{t|t-1}^2 + \sigma_\epsilon^2}$$

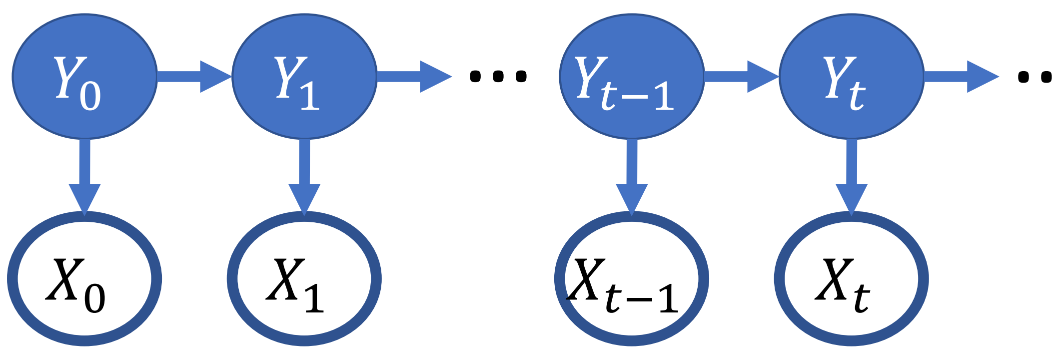

Kalman Filter: the update step

And here's the surprising result: k_t is all you need. If $X_t = Y_t + \epsilon$, and if Y_t and ϵ are Gaussian, then

$$\mu_{t|t} = E[Y_t | x_0, \dots, x_t] = \mu_{t|t-1} + k_t \left(x_t - \left(\mu_{t|t-1} + \mu_\epsilon \right) \right)$$

$$\sigma^2_{t|t} = \mathrm{Var}(Y_t | x_0, \dots, x_t) = \sigma^2_{t|t-1} (1 - k_t)$$

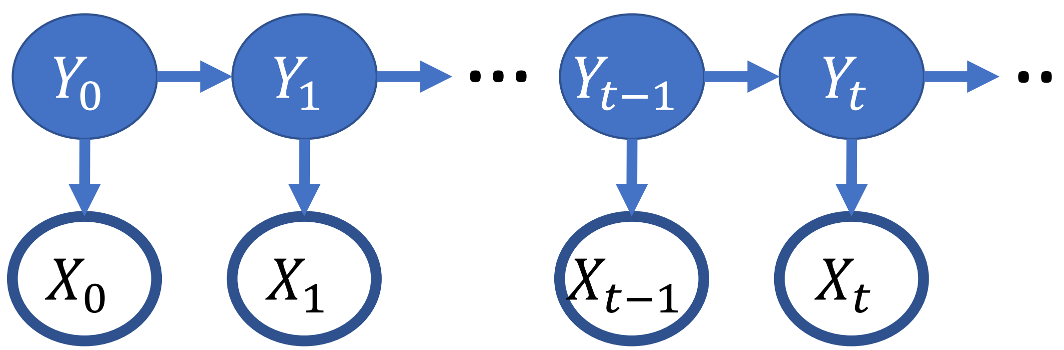

The Kalman filter

- Prediction step: given $\mu_{t-1|t-1}$ and $\sigma^2_{t-1|t-1}$, we can predict where the fish might go at time t, but with increased uncertainty:

$$\mu_{t|t-1} = \mu_{t-1|t-1} + \mu_\Delta$$
$$\sigma^2_{t|t-1} = \sigma^2_{t-1|t-1} + \sigma^2_\Delta$$

- Update step: given the observation x_t, we can refine our estimate, and reduce our uncertainty:

$$k_t = \frac{\sigma^2_{t|t-1}}{\sigma^2_{t|t-1} + \sigma^2_\epsilon}$$
$$\mu_{t|t} = \mu_{t|t-1} + k_t\left(x_t - \left(\mu_{t|t-1} + \mu_\epsilon\right)\right)$$
$$\sigma^2_{t|t} = \sigma^2_{t|t-1}(1 - k_t)$$

Quiz

- Try the quiz!

https://us.prairielearn.com/pl/course_instance/129874/assessment/2340212

Conclusion

Prediction step: given $\mu_{t-1|t-1}$ and $\sigma^2_{t-1|t-1}$, we can predict where the fish might go at time t, but with increased uncertainty:

$$\mu_{t|t-1} = \mu_{t-1|t-1} + \mu_\Delta$$
$$\sigma^2_{t|t-1} = \sigma^2_{t-1|t-1} + \sigma^2_\Delta$$

Update step: given the observation x_t, we can refine our estimate, and reduce our uncertainty:

$$k_t = \frac{\sigma^2_{t|t-1}}{\sigma^2_{t|t-1} + \sigma^2_\epsilon}$$
$$\mu_{t|t} = \mu_{t|t-1} + k_t\left(x_t - \left(\mu_{t|t-1} + \mu_\epsilon\right)\right)$$
$$\sigma^2_{t|t} = \sigma^2_{t|t-1}(1 - k_t)$$

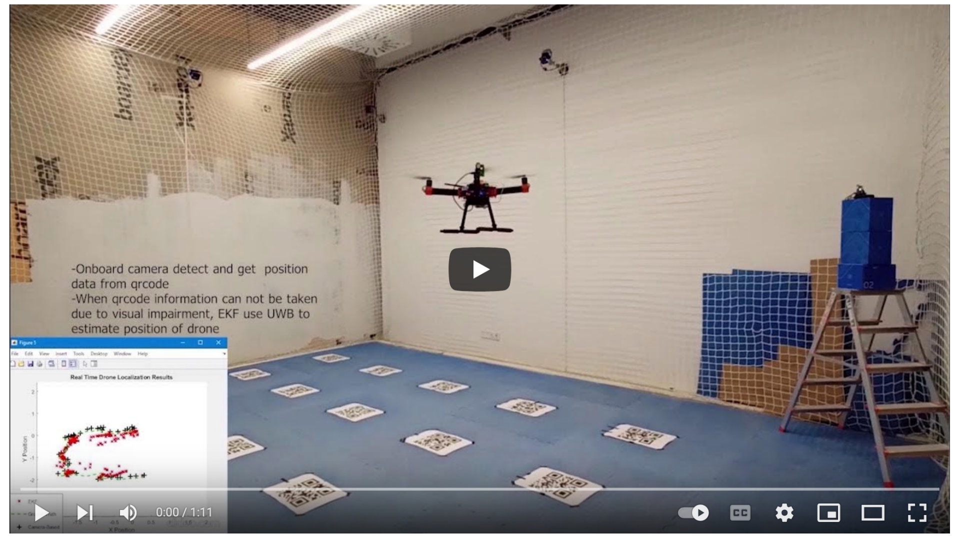

Drone Localization based on Extended Kalman Filter (EKF) with UWB sensors and camera, https://www.youtube.com/watch?v=kC8FgmhhSB8