Kalman Filter Example

COMP417

November 17, 2010

Vehicle Tracking

- Suppose you want to track the exact position of your car, while you are driving it.

Available sensors

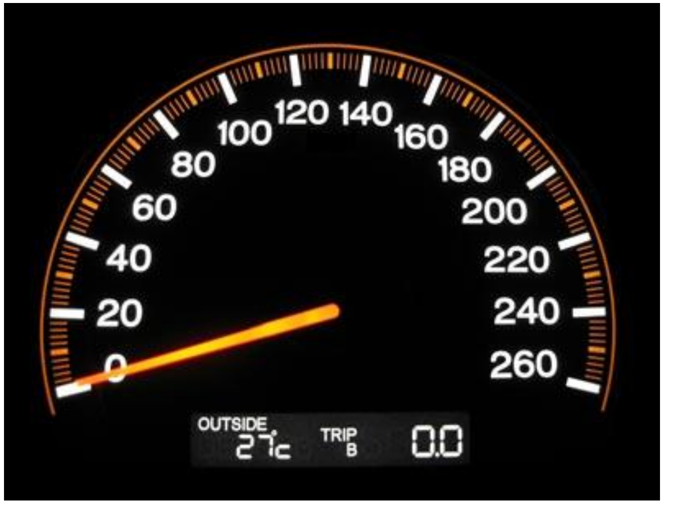

- May be intermittent
- Has error of ~5m
- Relatively slow update rate, 1Hz

- Runs continuously
- Relatively fast update rate

Representing belief (or lack thereof)

- What do you want to track?
- Position and velocity of the car

$$\mathbf{x}_t = \begin{bmatrix} x(t) \\ y(t) \\ \dot{x}(t) \\ \dot{y}(t) \end{bmatrix}$$

$$bel(\mathbf{x}_t) = p(\mathbf{x}_t \mid \mathbf{u}_{t-1}, \ldots, \mathbf{u}_0, \mathbf{z}_t, \ldots, \mathbf{z}_1)$$

All previous commands we've given the car

All previous sensor measurements

Simplifying belief

$$
\begin{aligned}
bel(\mathbf{x}_t) &= p(\mathbf{x}_t \mid \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_t,\ldots,\mathbf{z}_1) \\
&= p(\mathbf{z}_t \mid \mathbf{x}_t, \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_{t-1}\ldots,\mathbf{z}_1)\, p(\mathbf{x}_t \mid \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_{t-1}\ldots,\mathbf{z}_1) / \eta \\
&= p(\mathbf{z}_t \mid \mathbf{x}_t)\, p(\mathbf{x}_t \mid \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_{t-1}\ldots,\mathbf{z}_1) / \eta \\
&= p(\mathbf{z}_t \mid \mathbf{x}_t) \int p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_{t-1}\ldots,\mathbf{z}_1)\, p(\mathbf{x}_{t-1} \mid \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_{t-1}\ldots,\mathbf{z}_1) d\mathbf{x}_{t-1} / \eta \\
&= p(\mathbf{z}_t \mid \mathbf{x}_t) \int p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \mathbf{u}_{t-1},\ldots,\mathbf{u}_0, \mathbf{z}_{t-1}\ldots,\mathbf{z}_1) bel(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1} / \eta \\
&= p(\mathbf{z}_t \mid \mathbf{x}_t) \int p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \mathbf{u}_{t-1}) bel(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1} / \eta
\end{aligned}
$$

Bayes' Theorem

Markov Assumption

Rule of Total Probability

Measurement Model

Transition Model

The belief function is a "recursive Bayesian filter"

Kalman Filter Belief

- Kalman filters are just a special case of recursive Bayesian filters, where the belief is a Gaussian:

$$bel(\mathbf{x}_t) = p(\mathbf{z}_t \mid \mathbf{x}_t) \int p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \mathbf{u}_{t-1}) bel(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1} / \eta$$

- And the measurement and transition models are linear:

$$\mathbf{z}_t = \mathbf{H}\mathbf{x}_t + \mathbf{v} \quad \text{therefore} \quad p(\mathbf{z}_t \mid \mathbf{x}_t) = p(\mathbf{v}) \sim N(0, \mathbf{R})$$

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{B}\mathbf{u}_{t-1} + \mathbf{w} \quad \text{therefore} \quad p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \mathbf{u}_{t-1}) = p(\mathbf{w}) \sim N(0, \mathbf{Q})$$

- That's why we estimate $\hat{\mathbf{x}}_t$ and $\mathbf{P}_t = \operatorname{cov}(\mathbf{x}_t - \hat{\mathbf{x}}_t)$

We digressed…

- Back to the car example

$$\hat{\mathbf{x}}_{t|t-1} = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x(t-1) \\ y(t-1) \\ \dot{x}(t-1) \\ \dot{y}(t-1) \end{bmatrix} + \begin{bmatrix} dt^2/2 & 0 \\ 0 & dt^2/2 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \ddot{x}(t-1) \\ \ddot{y}(t-1) \end{bmatrix} + \mathbf{w}$$

F **B**

Noise of the transition / kinematic model

State estimate at time t, which takes into account measurements up to time t-1

Previous state estimate

Acceleration caused by the driver

Propagated covariance

$$\mathbf{P}_{t|t-1} = \mathbf{F}\mathbf{P}_{t-1|t-1}\mathbf{F}^T + \mathbf{Q}$$

Uncertainty grows without sensor measurements

Kalman Filter Update

- Suppose we get measurements from the GPS and the speedometer, both of which are noisy:

$$\mathbf{z}_t = \mathbf{x}_t + \mathbf{v} \quad \text{where} \quad \mathbf{v} \sim N(0, \mathbf{R})$$

- Measurement residual :

$$\mathbf{r}_t = \mathbf{z}_t - \hat{\mathbf{z}}_t = \mathbf{z}_t - \hat{\mathbf{x}}_{t|t-1}$$

- Captures the difference between what we actually observed through the sensors and what we expected to observe.

- If big then our estimate is bad, so we need to correct it significantly.

Correcting our estimates

- We can't blindly correct the state as much as the residual tells us. Why?
- Because the residual itself is noisy.
- That's why we use the "Kalman gain" $\mathbf{K}_t = \mathbf{P}_{t|t-1}(\mathbf{P}_{t|t-1} + \mathbf{R})^{-1}$ to weight the residual accordingly.
- And, finally, do the correction:

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t \mathbf{r}_t$$

$$\mathbf{P}_{t|t} = \mathbf{P}_{t|t-1} - \mathbf{K}_t \mathbf{P}_{t|t-1}$$

- Uncertainty shrinks after updates.