LQR and Kalman Filtering
Philip Sabes
CoSMo 2012

1. **Simulate linear dynamics with Gaussian noise.** Simulate the following very simple dynamical system:

$$\ddot{x} = \frac{f}{m} + v$$

$$\mathrm{y} = x + w,$$

where $v \propto N(0, V^2)$ and $w \propto N(0, W^2)$, i.e., a simple point mass that experiences random disturbances, with noisy measurements. You'll want to express the system in the discretized form

$$x_k = Ax_{k-1} + Bu_{k-1} + v$$
$$y_k = Cx_k + w$$

What is the state of the system, x_k? Simulate, say, for 500 msec with $\Delta t = 0.01$ msec. Use the parameters $m = 4\ kg$, $V = I\ 0.05\ m$, $W = I\ 0.1\ m$, and initial condition the mass stationary at position 1 m.
If you run the simulation many times, what is the dependency of variance with time?

2. **Linear Quadratic Regulator.** Perform the backward iteration to compute the LQR control law:

Finite-horizon, discrete-time LQR

For a discrete-time linear system described by [1]

$$x_k = Ax_{k-1} + Bu_k$$

with a performance index defined as

$$J = \sum_{k=0}^{N} \left(x_k^T Q x_k + u_k^T R u_k \right)$$

the optimal control sequence minimizing the performance index is given by

$$u_k = -F_k x_{k-1}$$

where

$$F_k = (R + B^T P_k B)^{-1} B^T P_k A$$

and P_k is found iteratively backwards in time by the dynamic Riccati equation

$$P_{k-1} = Q + A^T \left(P_k - P_k B \left(R + B^T P_k B \right)^{-1} B^T P_k \right) A$$

from initial condition $P_N = Q$.

Note that both Q and R can be functions of time. Use the cost function where $R_k = 0.01$ for all k, and

$$Q_k = \begin{cases} 0 \text{ for } k \neq N \\ 1 \text{ for } k = N \end{cases}.$$

Note that in this case, the initial condition is $P_N = Q_N$.

For the LQR, you need to have access to the true state, x_k. Simulate the dynamics with no output noisy and apply the LQR. Repeat the simulation. What is the mean and variance (taken across runs) of the state as a function of time?

3. **Kalman Filter and LQG controller.** The Kalman Filter is described by the following:

Predict

Predicted (*a priori*) state estimate	$\hat{x}_{k\|k-1} = A\,\hat{x}_{k-1\|k-1} + B\,u_k$
Predicted (*a priori*) estimate covariance	$P_{k\|k-1} = A\,P_{k-1\|k-1}A^T + V_k$

Update

Innovation (measurement residual)	$\tilde{y}_k = y_k - C\,\hat{x}_{k\|k-1}$
Innovation covariance	$S_k = CP_{k\|k-1}C^T + W_k$
Optimal Kalmal gain	$K_k = P_{k\|k-1}CS_k^{-1}$
Updated (*a posteriori*) state estimate	$\hat{x}_{k\|k} = \hat{x}_{k\|k-1} + K_k\tilde{y}_k$
Updated (*a posteriori*) estimate covariance	$P_{k\|k} = (I - K_kC)P_{k\|k-1}$

Implement the Kalman Filter for this system. In practice, you can run the forward iteration to get $P_{k|k-1}$, K_k, and $P_{k|k-1}$ one time, and then use the optimal Kalman gain K_k in all the simulations with the same parameters (*N, A, B, C, V, W*). The Kalman gain is used as follows:

$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + Bu_{k-1}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k\left(y_k - C\hat{x}_{k|k-1}\right)$$

Use the Kalman Filter with the LQR to control the dynamical system from part 1, this time with output noise. If you run the simulation many times, what is the dependency of variance with time?

Now try changing the cost function to make Q_k non-zero at some point(s) other than k=N. Again, look at the mean and variance as function of time (across multiple runs). Do you see evidence for the Minimum intervention principle?