Lecture 7
The Kalman filter

- Linear system driven by stochastic process
- Statistical steady-state
- Linear Gauss-Markov model
- Kalman filter
- Steady-state Kalman filter

Linear system driven by stochastic process

We consider a linear dynamical system $x(t+1) = Ax(t) + Bu(t)$, with $x(0)$ and $u(0),\ u(1), \ldots$ random variables

we'll use notation

$$\bar{x}(t) = \mathbf{E}\, x(t), \qquad \Sigma_x(t) = \mathbf{E}(x(t) - \bar{x}(t))(x(t) - \bar{x}(t))^T$$

and similarly for $\bar{u}(t)$, $\Sigma_u(t)$

taking expectation of $x(t+1) = Ax(t) + Bu(t)$ we have

$$\bar{x}(t+1) = A\bar{x}(t) + B\bar{u}(t)$$

$i.e.$, the means propagate by the same linear dynamical system now let's consider the covariance

$$x(t+1) - \bar{x}(t+1) = A(x(t) - \bar{x}(t)) + B(u(t) - \bar{u}(t))$$

and so

$$\begin{aligned}
\Sigma_x(t+1) &= \mathbf{E}\left(A(x(t) - \bar{x}(t)) + B(u(t) - \bar{u}(t))\right) \cdot \\
&\quad \cdot \left(A(x(t) - \bar{x}(t)) + B(u(t) - \bar{u}(t))\right)^T \\
&= A\Sigma_x(t)A^T + B\Sigma_u(t)B^T + A\Sigma_{xu}(t)B^T + B\Sigma_{ux}(t)A^T
\end{aligned}$$

where

$$\Sigma_{xu}(t) = \Sigma_{ux}(t)^T = \mathbf{E}(x(t) - \bar{x}(t))(u(t) - \bar{u}(t))^T$$

thus, the covariance $\Sigma_x(t)$ satisfies another, Lyapunov-like linear dynamical system, driven by Σ_{xu} and Σ_u

consider special case $\Sigma_{xu}(t) = 0$, $i.e.$, x and u are uncorrelated, so we have Lyapunov iteration

$$\Sigma_x(t+1) = A\Sigma_x(t)A^T + B\Sigma_u(t)B^T,$$

which is stable if and only if A is stable

if A is stable and $\Sigma_u(t)$ is constant, $\Sigma_x(t)$ converges to Σ_x, called the *steady-state covariance*, which satisfies Lyapunov equation

$$\Sigma_x = A\Sigma_x A^T + B\Sigma_u B^T$$

thus, we can calculate the steady-state covariance of x exactly, by solving a Lyapunov equation

(useful for starting simulations in statistical steady-state)

Question: Can you imagine situations where $\Sigma_{xu}(t) \neq 0$?

Example

we consider $x(t+1) = Ax(t) + w(t)$, with

$$A = \begin{bmatrix} 0.6 & -0.8 \\ 0.7 & 0.6 \end{bmatrix},$$

where $w(t)$ are IID $\mathcal{N}(0, I)$: i.e. white (memoryless) noise

eigenvalues of A are $0.6 \pm 0.75j$, with magnitude 0.96, so A is stable

we solve Lyapunov equation to find steady-state covariance

$$\Sigma_x = \begin{bmatrix} 13.35 & -0.03 \\ -0.03 & 11.75 \end{bmatrix}$$

covariance of $x(t)$ converges to Σ_x no matter its initial value two initial state distributions: $\Sigma_x(0) = 0$, $\Sigma_x(0) = 10^2 I$

plot shows $\Sigma_{11}(t)$ for the two cases

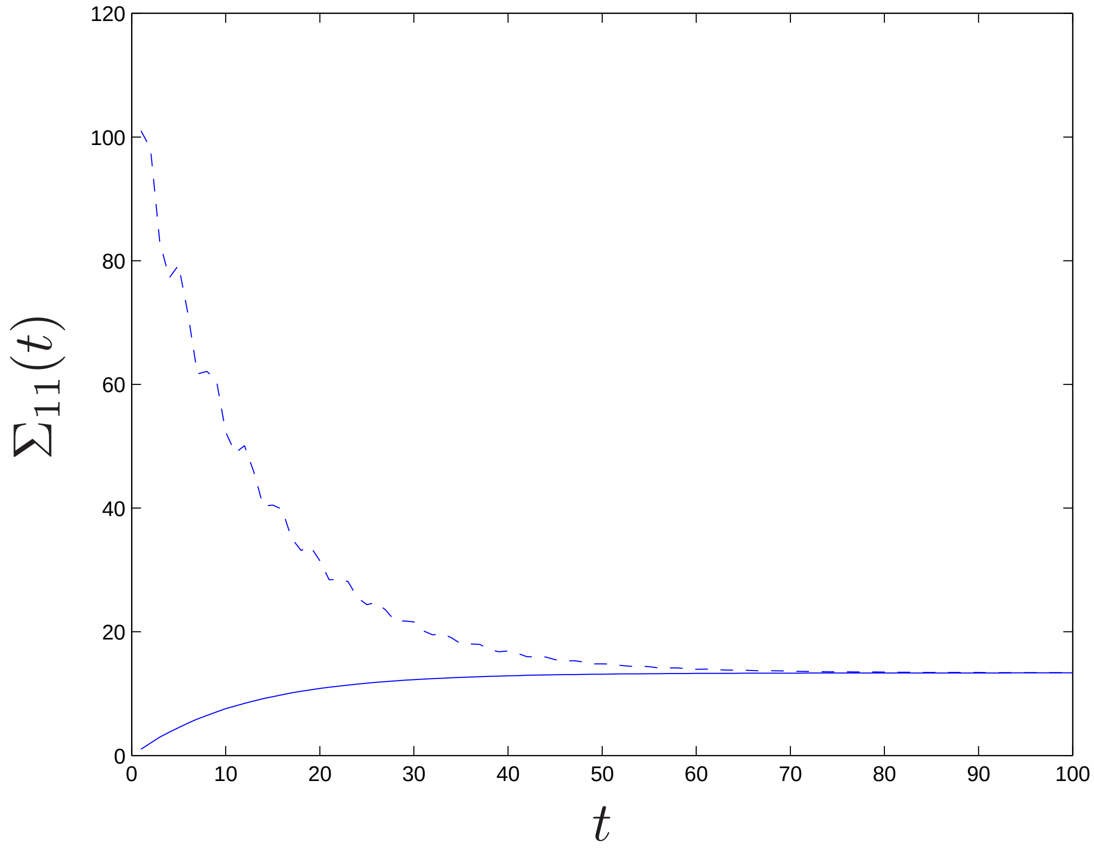

$x_1(t)$ for one realization from each case:

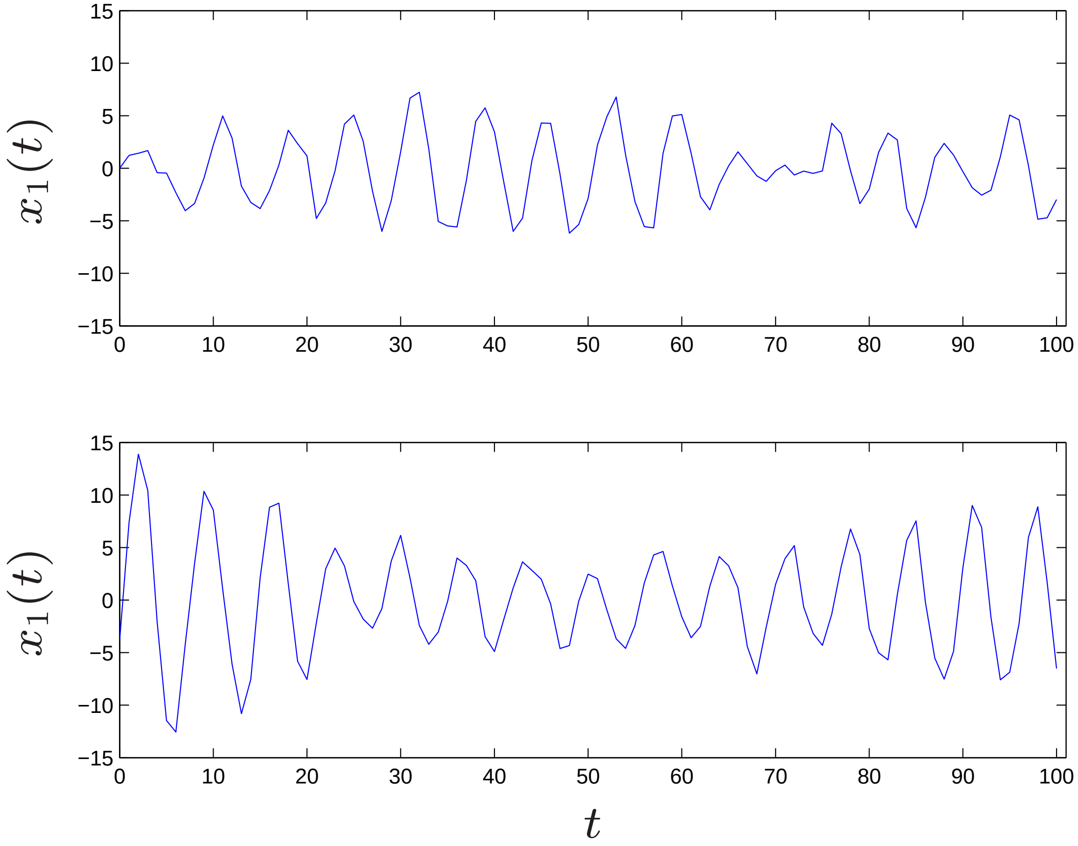

Graphical representation

Consider $x(t+1) = Ax(t) + w(t)$, and $w(t)$ is white noise.

$\Rightarrow$ we can represent the process $(x(t), w(t))$ by the following graph:

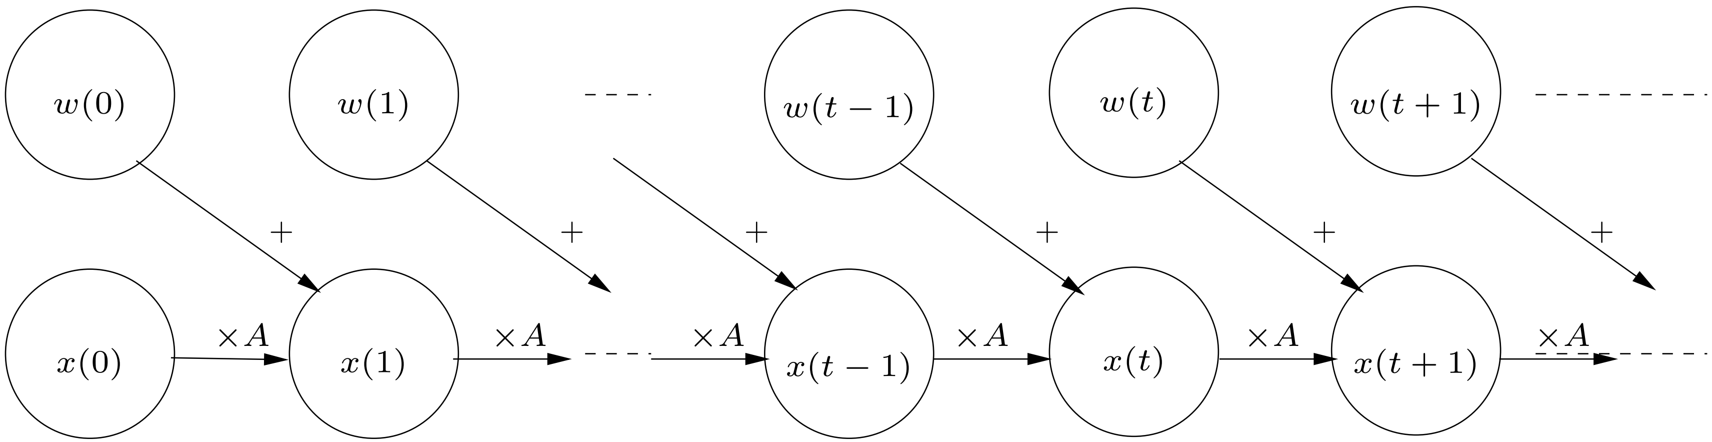

Hence, the state process $(x(t))$ is Markovian: $x(t-j) \perp x(t+k) | x(t)$

NB: The Markov property holds also if $w(t)$ and $x(0)$ are not Gaussian. It is a consequence of the assumption that the random variables $w(t)$ are independent of the previous states $x(t-j)$.

Other consequences

Under the assumption that $x(0), w(0), w(1), \ldots$ are jointly Gaussian, $x(0), x(1), x(2), \ldots$ are also jointly Gaussian.

Suppose now that the noise process is time-invariant, Gaussian and white. I.e. it is completely described by $\Sigma_w(t) = \Sigma_w$ and $\bar{w}(t) = \bar{w}$.

Suppose, also that $x(0) \sim \mathcal{N}(\bar{x}(0), \Sigma_x(0))$. Then, $\bar{x}(t+1) = A\bar{x}(t) + \bar{w}$ and $\Sigma_x(t+1) = A\Sigma_x(t)A^T + \Sigma_w$.

Consequently, the process $x(t)$ is stationary if its initial state distribution satisfies both

$$\begin{aligned} \bar{x}(0) &= A\bar{x}(0) + \bar{w} \\ \Sigma_x(0) &= A\Sigma_x(0)A^T + \Sigma_w \end{aligned} \tag{1}$$

If A is stable, the process converges over time towards stationarity, even if its initial state distribution is not 'stationary'.

Linear Gauss-Markov model

we consider linear dynamical system

$$x(t+1) = Ax(t) + w(t), \qquad y(t) = Cx(t) + v(t)$$

- $x(t) \in \mathbf{R}^n$ is the state; $y(t) \in \mathbf{R}^p$ is the observed output
- $w(t) \in \mathbf{R}^n$ is called *process noise* or *state noise*
- $v(t) \in \mathbf{R}^p$ is called *measurement noise*

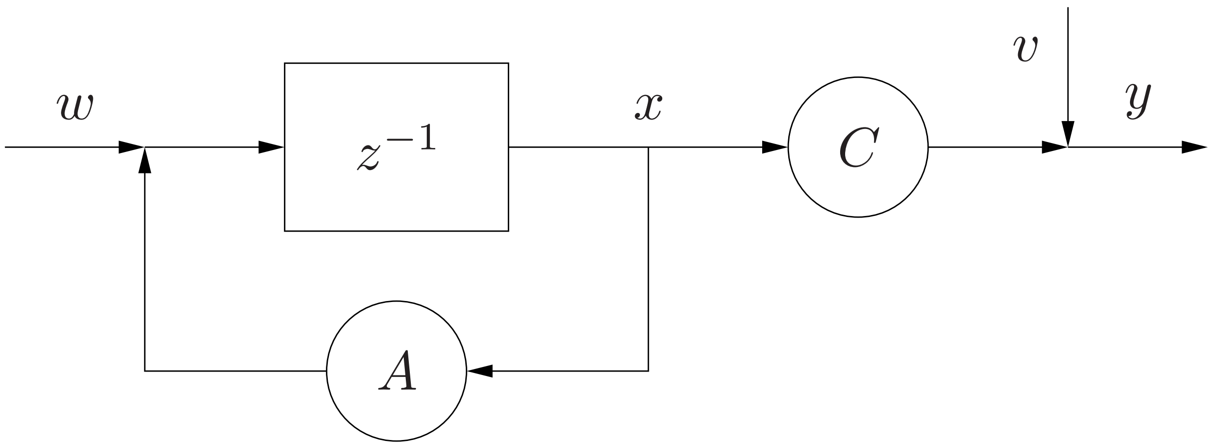

Statistical assumptions

- $x(0)$, $w(0), w(1), \ldots$, and $v(0), v(1), \ldots$ are jointly Gaussian and independent
- $w(t)$ are IID with $\mathbf{E}\, w(t) = 0$, $\mathbf{E}\, w(t)w(t)^T = W$
- $v(t)$ are IID with $\mathbf{E}\, v(t) = 0$, $\mathbf{E}\, v(t)v(t)^T = V$
- $\mathbf{E}\, x(0) = \bar{x}_0$, $\mathbf{E}(x(0) - \bar{x}_0)(x(0) - \bar{x}_0)^T = \Sigma_0$

(it's not hard to extend to case where $w(t)$, $v(t)$ are not zero mean)

we'll denote $X(t) = (x(0), \ldots, x(t))$, etc.

since $X(t)$ and $Y(t)$ are linear functions of $x(0)$, $W(t)$, and $V(t)$, we conclude they are all jointly Gaussian (*i.e.*, the process $x,\ w,\ v,\ y$ is Gaussian)

Statistical properties

- sensor noise v independent of x
- $w(t)$ is independent of $x(0), \ldots, x(t)$ and $y(0), \ldots, y(t)$
- *Markov property:* the process x is Markov, $i.e.$,

$$x(t)|x(0), \ldots, x(t-1) = x(t)|x(t-1)$$

  roughly speaking: if you know $x(t-1)$, then knowledge of $x(t-2), \ldots, x(0)$ doesn't give any more information about $x(t)$

NB: the process y is *Hidden Markov*.

Can you prove this ?

Draw factor graph of $x(0), w(0), y(0), v(0), \ldots, x(t), w(t), y(t), v(t)$.

Mean and covariance of Gauss-Markov process

mean satisfies $\bar{x}(t+1) = A\bar{x}(t)$, $\bar{x}(0) = \bar{x}_0$, so $\bar{x}(t) = A^t\bar{x}_0$

covariance satisfies

$$\Sigma_x(t+1) = A\Sigma_x(t)A^T + W$$

if A is stable, $\Sigma_x(t)$ converges to steady-state covariance Σ_x, which satisfies Lyapunov equation

$$\Sigma_x = A\Sigma_x A^T + W$$

Conditioning on observed output

we use the notation

$$\begin{aligned} \hat{x}(t|s) &= \mathbf{E}(x(t)|y(0), \ldots y(s)), \\ \Sigma_{t|s} &= \mathbf{E}(x(t) - \hat{x}(t|s))(x(t) - \hat{x}(t|s))^T \end{aligned}$$

- the random variable $x(t)|y(0), \ldots, y(s)$ is Gaussian, with mean $\hat{x}(t|s)$ and covariance $\Sigma_{t|s}$
- $\hat{x}(t|s)$ is the minimum mean-square error estimate of $x(t)$, based on $y(0), \ldots, y(s)$
- $\Sigma_{t|s}$ is the covariance of the error of the estimate $\hat{x}(t|s)$

State estimation

we focus on two state estimation problems:

- finding $\hat{x}(t|t)$, *i.e.*, estimating the current state, based on the current and past observed outputs
- finding $\hat{x}(t+1|t)$, *i.e.*, predicting the next state, based on the current and past observed outputs

since $x(t), Y(t)$ are jointly Gaussian, we can use the standard formula to find $\hat{x}(t|t)$ (and similarly for $\hat{x}(t+1|t)$)

$$\hat{x}(t|t) = \bar{x}(t) + \Sigma_{x(t)Y(t)}\Sigma_{Y(t)}^{-1}(Y(t) - \bar{Y}(t))$$

the inverse in the formula, $\Sigma_{Y(t)}^{-1}$, is size $pt \times pt$, which grows with t

the *Kalman filter* is a clever method for computing $\hat{x}(t|t)$ and $\hat{x}(t+1|t)$ recursively

Measurement update

let's find $\hat{x}(t|t)$ and $\Sigma_{t|t}$ in terms of $\hat{x}(t|t-1)$ and $\Sigma_{t|t-1}$

start with $y(t) = Cx(t) + v(t)$, and condition on $Y(t-1)$:

$$y(t)|Y(t-1) = Cx(t)|Y(t-1) + v(t)|Y(t-1) = Cx(t)|Y(t-1) + v(t)$$

since $v(t)$ and $Y(t-1)$ are independent

so $x(t)|Y(t-1)$ and $y(t)|Y(t-1)$ are jointly Gaussian with mean and covariance

$$\begin{bmatrix} \hat{x}(t|t-1) \\ C\hat{x}(t|t-1) \end{bmatrix}, \qquad \begin{bmatrix} \Sigma_{t|t-1} & \Sigma_{t|t-1}C^T \\ C\Sigma_{t|t-1} & C\Sigma_{t|t-1}C^T + V \end{bmatrix}$$

now use standard formula to get mean and covariance of

$$(x(t)|Y(t-1))|(y(t)|Y(t-1)),$$

which is exactly the same as $x(t)|Y(t)$:

$$\begin{aligned}\hat{x}(t|t) &= \hat{x}(t|t-1) + \Sigma_{t|t-1}C^T\left(C\Sigma_{t|t-1}C^T+V\right)^{-1}(y(t)-C\hat{x}(t|t-1)) \\ \Sigma_{t|t} &= \Sigma_{t|t-1} - \Sigma_{t|t-1}C^T\left(C\Sigma_{t|t-1}C^T+V\right)^{-1}C\Sigma_{t|t-1}\end{aligned}$$

this gives us $\hat{x}(t|t)$ and $\Sigma_{t|t}$ in terms of $\hat{x}(t|t-1)$ and $\Sigma_{t|t-1}$

this is called the *measurement update* since it gives our updated estimate of $x(t)$ based on the measurement $y(t)$ becoming available

Time update

now let's increment time, using $x(t+1) = Ax(t) + w(t)$

condition on $Y(t)$ to get

$$\begin{aligned} x(t+1)|Y(t) &= Ax(t)|Y(t) + w(t)|Y(t) \\ &= Ax(t)|Y(t) + w(t) \end{aligned}$$

since $w(t)$ is independent of $Y(t)$

therefore we have and

$$\begin{aligned} \hat{x}(t+1|t) &= A\hat{x}(t|t) \\ \Sigma_{t+1|t} &= A\Sigma_{t|t}A^T + W \end{aligned}$$

Kalman filter

measurement and time updates together give a recursive solution

start with prior mean and covariance, $\hat{x}(0|-1) = \bar{x}_0$, $\Sigma(0|-1) = \Sigma_0$

apply the measurement update

$$
\begin{aligned}
\hat{x}(t|t) &= \hat{x}(t|t-1) + \Sigma_{t|t-1}C^T \left(C\Sigma_{t|t-1}C^T + V\right)^{-1} (y(t) - C\hat{x}(t|t-1)) \\
\Sigma_{t|t} &= \Sigma_{t|t-1} - \Sigma_{t|t-1}C^T \left(C\Sigma_{t|t-1}C^T + V\right)^{-1} C\Sigma_{t|t-1}
\end{aligned}
$$

to get $\hat{x}(0|0)$ and $\Sigma_{0|0}$; then apply time update

$$
\hat{x}(t+1|t) = A\hat{x}(t|t), \qquad \Sigma_{t+1|t} = A\Sigma_{t|t}A^T + W
$$

to get $\hat{x}(1|0)$ and $\Sigma_{1|0}$

now, repeat measurement and time updates . . .

Riccati recursion

to lighten notation, we'll use $\hat{x}(t) = \hat{x}(t|t-1)$ and $\hat{\Sigma}_t = \Sigma_{t|t-1}$

we can express measurement and time updates for $\hat{\Sigma}$ as

$$\hat{\Sigma}_{t+1} = A\hat{\Sigma}_t A^T + W - A\hat{\Sigma}_t C^T (C\hat{\Sigma}_t C^T + V)^{-1} C\hat{\Sigma}_t A^T$$

which is a Riccati recursion, with initial condition $\hat{\Sigma}_0 = \Sigma_0$

- $\hat{\Sigma}_t$ can be computed *before any observations are made*
- thus, we can calculate the estimation error covariance *before* we get any observed data

Observer form

we can express KF as

$$\begin{aligned}\hat{x}(t+1) &= A\hat{x}(t) + A\hat{\Sigma}_t C^T (C\hat{\Sigma}_t C^T + V)^{-1}(y(t) - C\hat{x}(t)) \\ &= A\hat{x}(t) + L_t(y(t) - \hat{y}(t))\end{aligned}$$

where $L_t = A\hat{\Sigma}_t C^T (C\hat{\Sigma}_t C^T + V)^{-1}$ is the *observer gain*, and $\hat{y}(t)$ is $\hat{y}(t|t-1)$

- $\hat{y}(t)$ is our output prediction, $i.e.$, our estimate of $y(t)$ based on $y(0), \ldots, y(t-1)$
- $e(t) = y(t) - \hat{y}(t)$ is our output prediction error
- $A\hat{x}(t)$ is our prediction of $x(t+1)$ based on $y(0), \ldots, y(t-1)$
- our estimate of $x(t+1)$ is the prediction based on $y(0), \ldots, y(t-1)$, plus a linear function of the output prediction error

Kalman filter block diagram

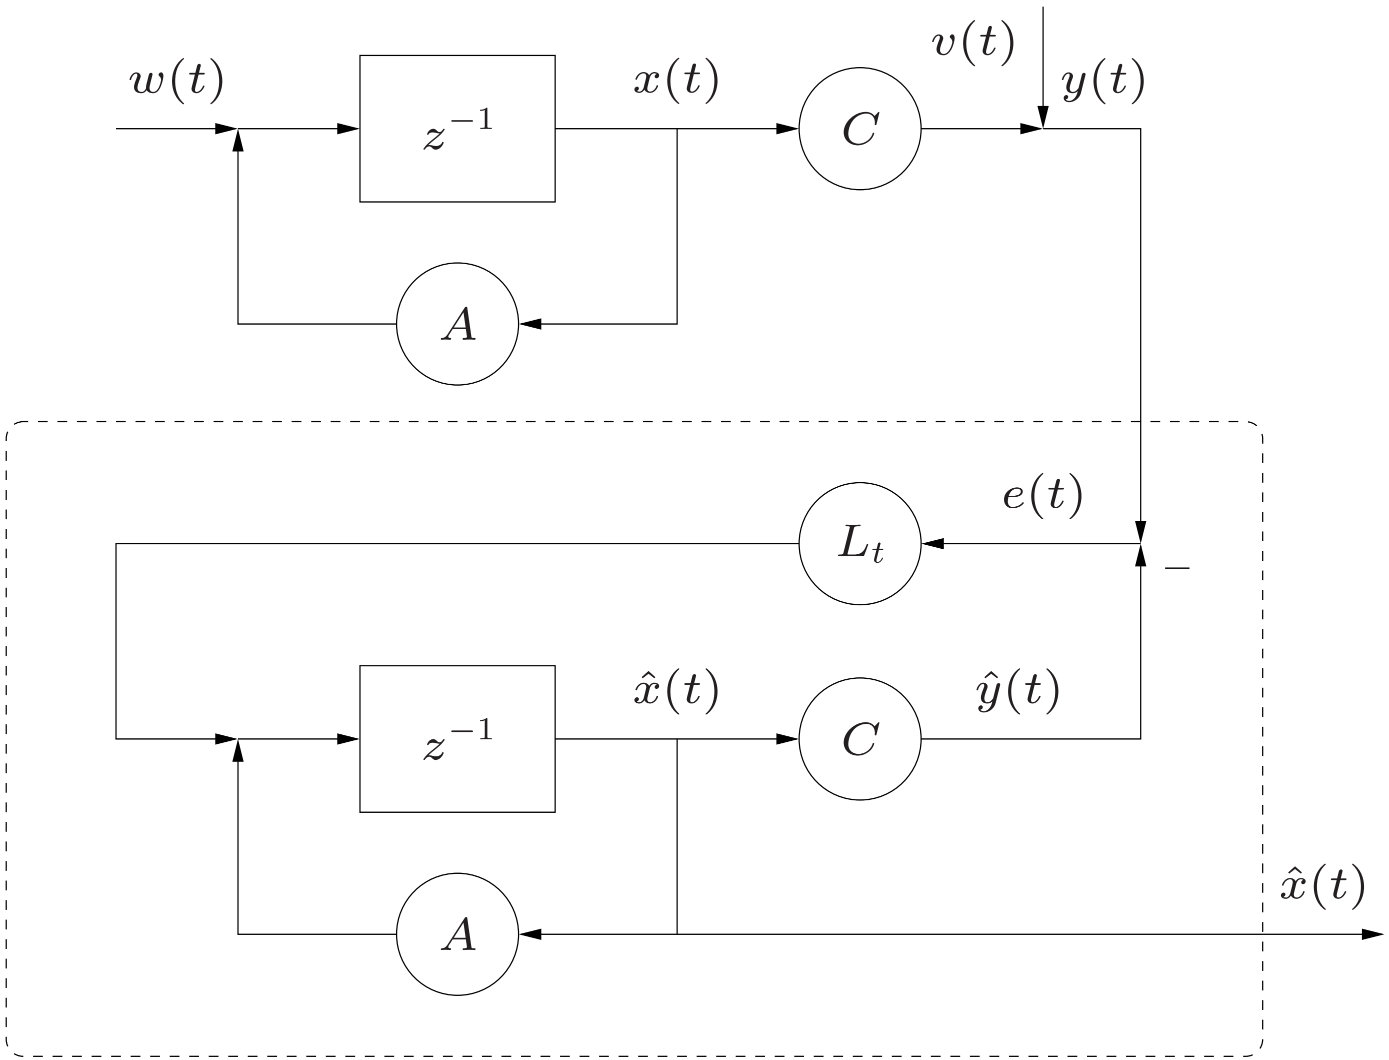

Steady-state Kalman filter

as in LQR, Riccati recursion for $\hat{\Sigma}_t$ converges to steady-state value $\hat{\Sigma}$, provided (C, A) is observable and (A, W) is controllable

$\hat{\Sigma}$ gives steady-state error covariance for estimating $x(t+1)$ given $y(0), \ldots, y(t)$

note that state prediction error covariance converges, even if system is unstable

$\hat{\Sigma}$ satisfies ARE

$$\hat{\Sigma} = A\hat{\Sigma}A^T + W - A\hat{\Sigma}C^T(C\hat{\Sigma}C^T + V)^{-1}C\hat{\Sigma}A^T$$

(which can be solved directly)

steady-state filter is a time-invariant observer:

$$\hat{x}(t+1) = A\hat{x}(t) + L(y(t) - \hat{y}(t)), \qquad \hat{y}(t) = C\hat{x}(t)$$

where $L = A\hat{\Sigma}C^T(C\hat{\Sigma}C^T + V)^{-1}$

define state estimation error $\tilde{x}(t) = x(t) - \hat{x}(t)$, so

$$y(t) - \hat{y}(t) = Cx(t) + v(t) - C\hat{x}(t) = C\tilde{x}(t) + v(t)$$

and

$$\begin{aligned}
\tilde{x}(t+1) &= x(t+1) - \hat{x}(t+1) \\
&= Ax(t) + w(t) - A\hat{x}(t) - L(C\tilde{x}(t) + v(t)) \\
&= (A - LC)\tilde{x}(t) + w(t) - Lv(t)
\end{aligned}$$

thus, the estimation error propagates according to a linear system, with closed-loop dynamics $A - LC$, driven by the process $w(t) - LCv(t)$, which is IID zero mean and covariance $W + LVL^T$

provided A, W is controllable and C, A is observable, $A - LC$ is stable

Example

system is

$$x(t+1) = Ax(t) + w(t), \qquad y(t) = Cx(t) + v(t)$$

with $x(t) \in \mathbf{R}^6$, $y(t) \in \mathbf{R}$

we'll take $\mathbf{E}\, x(0) = 0$, $\mathbf{E}\, x(0)x(0)^T = \Sigma_0 = 5^2 I$; $W = (1.5)^2 I$, $V = 1$

eigenvalues of A:

$$0.9973 \pm 0.0730j, \qquad 0.9995 \pm 0.0324j, \qquad 0.9941 \pm 0.1081j$$

(which have magnitude one)

goal: predict $y(t+1)$ based on $y(0), \dots, y(t)$

first let's find variance of $y(t)$ versus t, using Lyapunov recursion

$$\mathbf{E}\, y(t)^2 = C\Sigma_x(t)C^T + V, \qquad \Sigma_x(t+1) = A\Sigma_x(t)A^T + W, \qquad \Sigma_x(0) = \Sigma_0$$

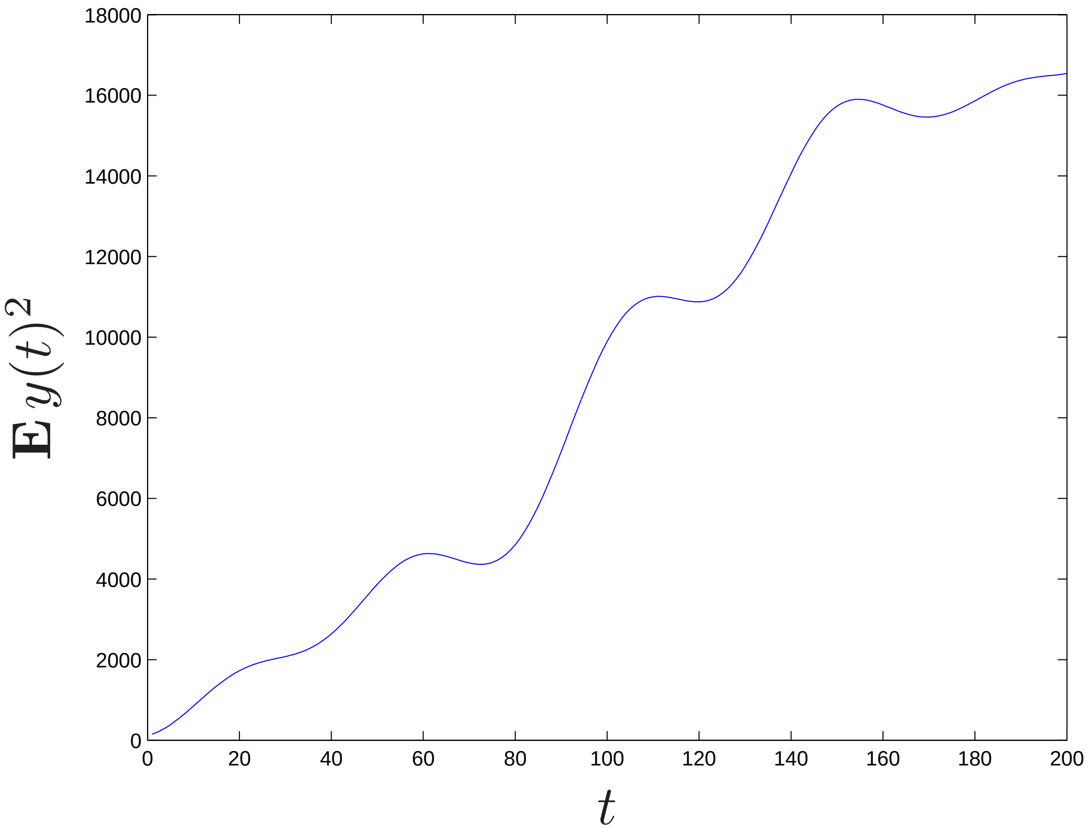

now, let's plot the prediction error variance versus t,

$$\mathbf{E}\, e(t)^2 = \mathbf{E}(\hat{y}(t) - y(t))^2 = C\hat{\Sigma}_t C^T + V,$$

where $\hat{\Sigma}_t$ satisfies Riccati recursion

$$\hat{\Sigma}_{t+1} = A\hat{\Sigma}_t A^T + W - A\hat{\Sigma}_t C^T (C\hat{\Sigma}_t C^T + V)^{-1} C\hat{\Sigma}_t A^T, \qquad \hat{\Sigma}_{-1} = \Sigma_0$$

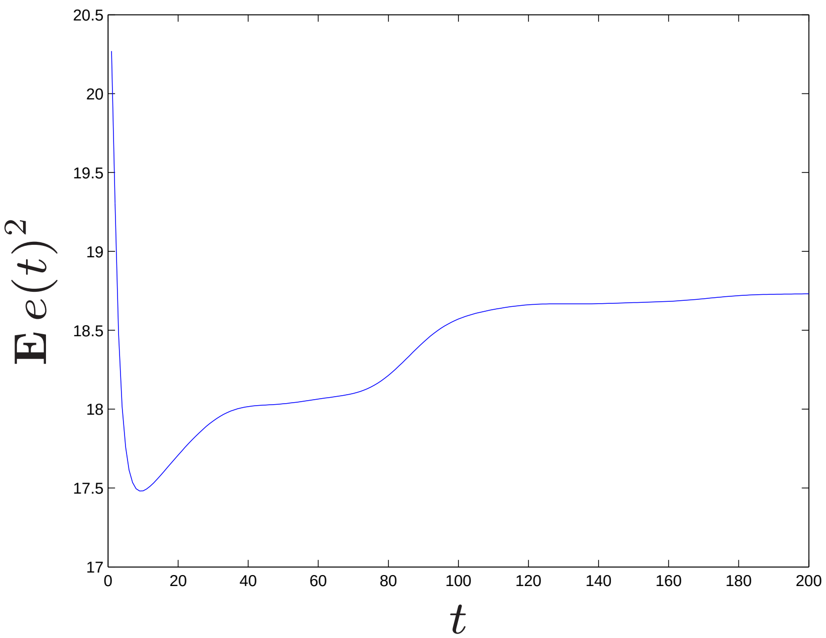

prediction error variance converges to steady-state value 18.7

now let's try the Kalman filter on a realization $y(t)$

top plot shows $y(t)$; bottom plot shows $e(t)$ (on different vertical scale)