General pseudo-inverse

if A has SVD $A = U\Sigma V^T$,

$$A^\dagger = V\Sigma^{-1}U^T$$

is the *pseudo-inverse* or *Moore-Penrose inverse* of A

if A is skinny and full rank,

$$A^\dagger = (A^TA)^{-1}A^T$$

gives the least-squares solution $x_{\rm ls} = A^\dagger y$

if A is fat and full rank,

$$A^\dagger = A^T(AA^T)^{-1}$$

gives the least-norm solution $x_{\rm ln} = A^\dagger y$

Full SVD

SVD of $A \in \mathbf{R}^{m\times n}$ with $\mathbf{Rank}(A) = r$:

$$A = U_1\Sigma_1 V_1^T = \begin{bmatrix} u_1 & \cdots & u_r \end{bmatrix} \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} \begin{bmatrix} v_1^T \\ \vdots \\ v_r^T \end{bmatrix}$$

- find $U_2 \in \mathbf{R}^{m\times(m-r)}$, $V_2 \in \mathbf{R}^{n\times(n-r)}$ s.t. $U = [U_1\ U_2] \in \mathbf{R}^{m\times m}$ and $V = [V_1\ V_2] \in \mathbf{R}^{n\times n}$ are orthogonal
- add zero rows/cols to Σ_1 to form $\Sigma \in \mathbf{R}^{m\times n}$:

$$\Sigma = \left[\begin{array}{c|c} \Sigma_1 & 0_{r\times(n-r)} \\ \hline 0_{(m-r)\times r} & 0_{(m-r)\times(n-r)} \end{array}\right]$$

then we have

$$A = U_1 \Sigma_1 V_1^T = \left[\begin{array}{c|c} U_1 & U_2 \end{array} \right] \left[\begin{array}{c|c} \Sigma_1 & 0_{r \times (n-r)} \\ \hline 0_{(m-r) \times r} & 0_{(m-r) \times (n-r)} \end{array} \right] \left[\begin{array}{c} V_1^T \\ \hline V_2^T \end{array} \right]$$

i.e.:

$$A = U \Sigma V^T$$

called *full SVD* of A

(SVD with positive singular values only called *compact SVD*)

Image of unit ball under linear transformation

full SVD:

$$A = U\Sigma V^T$$

gives intepretation of $y = Ax$:

- rotate (by V^T)
- stretch along axes by σ_i ($\sigma_i = 0$ for $i > r$)
- zero-pad (if $m > n$) or truncate (if $m < n$) to get m-vector
- rotate (by U)

Image of unit ball under A

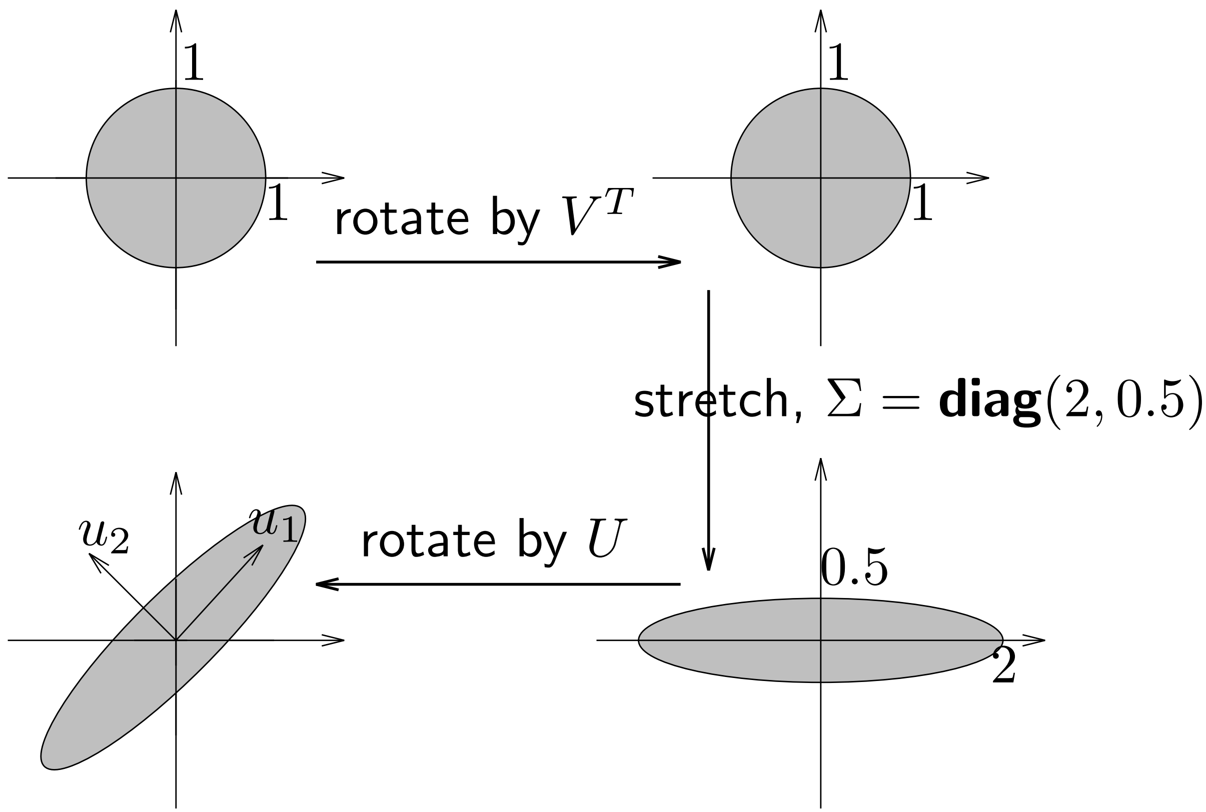

$\{Ax \mid \|x\| \leq 1\}$ is *ellipsoid* with principal axes $\sigma_i u_i$.

Sensitivity of linear equations to data error

consider $y = Ax$, $A \in \mathbf{R}^{n\times n}$ invertible; of course $x = A^{-1}y$

suppose we have an error or noise in y, $i.e.$, y becomes $y + \delta y$

then x becomes $x + \delta x$ with $\delta x = A^{-1}\delta y$

hence we have $\|\delta x\| = \|A^{-1}\delta y\| \le \|A^{-1}\|\|\delta y\|$

if $\|A^{-1}\|$ is large,

- small errors in y can lead to large errors in x
- can't solve for x given y (with small errors)
- hence, A can be considered singular in practice

a more refined analysis uses *relative* instead of *absolute* errors in x and y

since $y = Ax$, we also have $\|y\| \leq \|A\|\|x\|$, hence

$$\frac{\|\delta x\|}{\|x\|} \leq \|A\|\|A^{-1}\|\frac{\|\delta y\|}{\|y\|}$$

$$\kappa(A) = \|A\|\|A^{-1}\| = \sigma_{\max}(A)/\sigma_{\min}(A)$$

is called the *condition number* of A

we have:

relative error in solution $x \leq$ condition number $\cdot$ relative error in data y

or, in terms of $\#$ bits of guaranteed accuracy:

$\#$ bits accuacy in solution $\approx$ $\#$ bits accuracy in data $-\log_2 \kappa$

we say

- A is well conditioned if κ is small
- A is poorly conditioned if κ is large

(definition of ‘small’ and ‘large’ depend on application)

same analysis holds for least-squares solutions with A nonsquare, $\kappa = \sigma_{\max}(A)/\sigma_{\min}(A)$

State estimation set up

we consider the discrete-time system

$$x(t+1) = Ax(t) + Bu(t) + w(t), \quad y(t) = Cx(t) + Du(t) + v(t)$$

- w is state *disturbance* or *noise*
- v is sensor *noise* or *error*
- A, B, C, and D are known
- u and y are observed over time interval $[0, t-1]$
- w and v are not known, but can be described statistically, or assumed small ($e.g.$, in RMS value)

State estimation problem

state estimation problem: estimate $x(s)$ from

$$u(0), \ldots, u(t-1),\ y(0), \ldots, y(t-1)$$

- $s = 0$: estimate initial state
- $s = t-1$: estimate current state
- $s = t$: estimate ($i.e.$, predict) next state

an algorithm or system that yields an estimate $\hat{x}(s)$ is called an *observer* or *state estimator*

$\hat{x}(s)$ is denoted $\hat{x}(s|t-1)$ to show what information estimate is based on (read, "$\hat{x}(s)$ given $t-1$")

Noiseless case

let's look at finding $x(0)$, with no state or measurement noise:

$$x(t+1) = Ax(t) + Bu(t), \quad y(t) = Cx(t) + Du(t)$$

with $x(t) \in \mathbf{R}^n$, $u(t) \in \mathbf{R}^m$, $y(t) \in \mathbf{R}^p$

then we have

$$\begin{bmatrix} y(0) \\ \vdots \\ y(t-1) \end{bmatrix} = \mathcal{O}_t x(0) + \mathcal{T}_t \begin{bmatrix} u(0) \\ \vdots \\ u(t-1) \end{bmatrix}$$

where

$$\mathcal{O}_t = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{t-1} \end{bmatrix}, \quad \mathcal{T}_t = \begin{bmatrix} D & 0 & \cdots & & \\ CB & D & 0 & \cdots & \\ \vdots & & & & \\ CA^{t-2}B & CA^{t-3}B & \cdots & CB & D \end{bmatrix}$$

- $\mathcal{O}_t$ maps initials state into resulting output over $[0, t-1]$
- $\mathcal{T}_t$ maps input to output over $[0, t-1]$

hence we have

$$\mathcal{O}_t x(0) = \begin{bmatrix} y(0) \\ \vdots \\ y(t-1) \end{bmatrix} - \mathcal{T}_t \begin{bmatrix} u(0) \\ \vdots \\ u(t-1) \end{bmatrix}$$

RHS is known, $x(0)$ is to be determined hence:

- can uniquely determine $x(0)$ if and only if $\mathcal{N}(\mathcal{O}_t) = \{0\}$
- $\mathcal{N}(\mathcal{O}_t)$ gives ambiguity in determining $x(0)$
- if $x(0) \in \mathcal{N}(\mathcal{O}_t)$ and $u = 0$, output is zero over interval $[0, t-1]$
- input u does not affect ability to determine $x(0)$;
  its effect can be subtracted out

Observability matrix

by C-H theorem, each A^k is linear combination of $A^0, \ldots, A^{n-1}$

hence for $t \geq n$, $\mathcal{N}(\mathcal{O}_t) = \mathcal{N}(\mathcal{O})$ where

$$\mathcal{O} = \mathcal{O}_n = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}$$

is called the *observability matrix*

if $x(0)$ can be deduced from u and y over $[0, t-1]$ for any t, then $x(0)$ can be deduced from u and y over $[0, n-1]$

$\mathcal{N}(\mathcal{O})$ is called *unobservable subspace*; describes ambiguity in determining state from input and output

system is called *observable* if $\mathcal{N}(\mathcal{O}) = \{0\}$, $i.e.$, $\mathbf{Rank}(\mathcal{O}) = n$

Observers for noiseless case

suppose $\mathbf{Rank}(\mathcal{O}_t) = n$ ($i.e.$, system is observable) and let F be any left inverse of $\mathcal{O}_t$, $i.e.$, $F\mathcal{O}_t = I$

then we have the observer

$$x(0) = F\left(\left[\begin{array}{c} y(0) \\ \vdots \\ y(t-1) \end{array}\right] - \mathcal{T}_t \left[\begin{array}{c} u(0) \\ \vdots \\ u(t-1) \end{array}\right]\right)$$

which deduces $x(0)$ (exactly) from u, y over $[0, t-1]$

in fact we have

$$x(\tau - t + 1) = F\left(\left[\begin{array}{c} y(\tau - t + 1) \\ \vdots \\ y(\tau) \end{array}\right] - \mathcal{T}_t \left[\begin{array}{c} u(\tau - t + 1) \\ \vdots \\ u(\tau) \end{array}\right]\right)$$

$i.e.$, our observer estimates what state was $t-1$ epochs ago, given past $t-1$ inputs & outputs

observer is (multi-input, multi-output) *finite impulse response* (FIR) filter, with inputs u and y, and output $\hat{x}$

Invariance of unobservable set

fact: the unobservable subspace $\mathcal{N}(\mathcal{O})$ is invariant, $i.e.$, if $z \in \mathcal{N}(\mathcal{O})$, then $Az \in \mathcal{N}(\mathcal{O})$

proof: suppose $z \in \mathcal{N}(\mathcal{O})$, $i.e.$, $CA^k z = 0$ for $k = 0, \ldots, n-1$

evidently $CA^k(Az) = 0$ for $k = 0, \ldots, n-2$;

$$CA^{n-1}(Az) = CA^n z = -\sum_{i=0}^{n-1} \alpha_i CA^i z = 0$$

(by C-H) where

$$\det(sI - A) = s^n + \alpha_{n-1}s^{n-1} + \cdots + \alpha_0$$

Continuous-time observability

continuous-time system with no sensor or state noise:

$$\dot{x} = Ax + Bu, \quad y = Cx + Du$$

can we deduce state x from u and y?

let's look at derivatives of y:

$$\begin{aligned} y &= Cx + Du \\ \dot{y} &= C\dot{x} + D\dot{u} = CAx + CBu + D\dot{u} \\ \ddot{y} &= CA^2x + CABu + CB\dot{u} + D\ddot{u} \end{aligned}$$

and so on

hence we have

$$\begin{bmatrix} y \\ \dot{y} \\ \vdots \\ y^{(n-1)} \end{bmatrix} = \mathcal{O}x + \mathcal{T} \begin{bmatrix} u \\ \dot{u} \\ \vdots \\ u^{(n-1)} \end{bmatrix}$$

where $\mathcal{O}$ is the observability matrix and

$$\mathcal{T} = \begin{bmatrix} D & 0 & \cdots & & \\ CB & D & 0 & \cdots & \\ \vdots & & & & \\ CA^{n-2}B & CA^{n-3}B & \cdots & CB & D \end{bmatrix}$$

(same matrices we encountered in discrete-time case!)

rewrite as

$$\mathcal{O}x = \begin{bmatrix} y \\ \dot{y} \\ \vdots \\ y^{(n-1)} \end{bmatrix} - \mathcal{T} \begin{bmatrix} u \\ \dot{u} \\ \vdots \\ u^{(n-1)} \end{bmatrix}$$

RHS is known; x is to be determined

hence if $\mathcal{N}(\mathcal{O}) = \{0\}$ we can deduce $x(t)$ from derivatives of $u(t)$, $y(t)$ up to order $n-1$

in this case we say system is observable

can construct an observer using any left inverse F of $\mathcal{O}$:

$$x = F \left(\begin{bmatrix} y \\ \dot{y} \\ \vdots \\ y^{(n-1)} \end{bmatrix} - \mathcal{T} \begin{bmatrix} u \\ \dot{u} \\ \vdots \\ u^{(n-1)} \end{bmatrix} \right)$$

- reconstructs $x(t)$ (exactly and instantaneously) from

$$u(t), \ldots, u^{(n-1)}(t), \; y(t), \ldots, y^{(n-1)}(t)$$

- derivative-based state reconstruction is dual of state transfer using impulsive inputs

A converse

suppose $z \in \mathcal{N}(\mathcal{O})$ (the unobservable subspace), and u is any input, with x, y the corresponding state and output, *i.e.*,

$$\dot{x} = Ax + Bu, \quad y = Cx + Du$$

then state trajectory $\tilde{x} = x + e^{tA}z$ satisfies

$$\dot{\tilde{x}} = A\tilde{x} + Bu, \quad y = C\tilde{x} + Du$$

i.e., input/output signals u, y consistent with both state trajectories x, $\tilde{x}$

hence if system is unobservable, no signal processing of any kind applied to u and y can deduce x

unobservable subspace $\mathcal{N}(\mathcal{O})$ gives fundamental ambiguity in deducing x from u, y

Least-squares observers

discrete-time system, with sensor noise:

$$x(t+1) = Ax(t) + Bu(t), \quad y(t) = Cx(t) + Du(t) + v(t)$$

we assume $\mathbf{Rank}(\mathcal{O}_t) = n$ (hence, system is observable)

least-squares observer uses pseudo-inverse:

$$\hat{x}(0) = \mathcal{O}_t^\dagger \left(\begin{bmatrix} y(0) \\ \vdots \\ y(t-1) \end{bmatrix} - \mathcal{T}_t \begin{bmatrix} u(0) \\ \vdots \\ u(t-1) \end{bmatrix} \right)$$

where $\mathcal{O}_t^\dagger = \left(\mathcal{O}_t^T \mathcal{O}_t\right)^{-1} \mathcal{O}_t^T$

interpretation: $\hat{x}_{\rm ls}(0)$ minimizes discrepancy between

- output $\hat{y}$ that *would be* observed, with input u and initial state $x(0)$ (and no sensor noise), and
- output y that *was* observed,

measured as $\sum_{\tau=0}^{t-1} \|\hat{y}(\tau) - y(\tau)\|^2$

can express least-squares initial state estimate as

$$\hat{x}_{\rm ls}(0) = \left(\sum_{\tau=0}^{t-1} (A^T)^\tau C^T C A^\tau \right)^{-1} \sum_{\tau=0}^{t-1} (A^T)^\tau C^T \tilde{y}(\tau)$$

where $\tilde{y}$ is observed output with portion due to input subtracted: $\tilde{y} = y - h * u$ where h is impulse response

Least-squares observer uncertainty ellipsoid

since $\mathcal{O}_t^\dagger \mathcal{O}_t = I$, we have

$$\tilde{x}(0) = \hat{x}_{\mathrm{ls}}(0) - x(0) = \mathcal{O}_t^\dagger \begin{bmatrix} v(0) \\ \vdots \\ v(t-1) \end{bmatrix}$$

where $\tilde{x}(0)$ is the estimation error of the initial state

in particular, $\hat{x}_{\mathrm{ls}}(0) = x(0)$ if sensor noise is zero
($i.e.$, observer recovers exact state in noiseless case)

now assume sensor noise is unknown, but has RMS value $\leq \alpha$,

$$\frac{1}{t} \sum_{\tau=0}^{t-1} \|v(\tau)\|^2 \leq \alpha^2$$

set of possible estimation errors is ellipsoid

$$\tilde{x}(0) \in \mathcal{E}_{\text{unc}} = \left\{ \mathcal{O}_t^\dagger \begin{bmatrix} v(0) \\ \vdots \\ v(t-1) \end{bmatrix} \;\middle|\; \frac{1}{t}\sum_{\tau=0}^{t-1} \|v(\tau)\|^2 \le \alpha^2 \right\}$$

$\mathcal{E}_{\text{unc}}$ is 'uncertainty ellipsoid' for $x(0)$ (least-square gives best $\mathcal{E}_{\text{unc}}$)

shape of uncertainty ellipsoid determined by matrix

$$\left(\mathcal{O}_t^T \mathcal{O}_t\right)^{-1} = \left(\sum_{\tau=0}^{t-1} (A^T)^\tau C^T C A^\tau\right)^{-1}$$

maximum norm of error is

$$\|\hat{x}_{\text{ls}}(0) - x(0)\| \le \alpha\sqrt{t}\|\mathcal{O}_t^\dagger\|$$

Infinite horizon uncertainty ellipsoid

the matrix

$$P = \lim_{t\to\infty} \left(\sum_{\tau=0}^{t-1} (A^T)^\tau C^T C A^\tau \right)^{-1}$$

always exists, and gives the limiting uncertainty in estimating $x(0)$ from u, y over longer and longer periods:

- if A is stable, $P > 0$
  $i.e.$, can't estimate initial state perfectly even with infinite number of measurements $u(t),\ y(t)$, $t = 0, \ldots$ (since memory of $x(0)$ fades . . .)
- if A is not stable, then P can have nonzero nullspace
  $i.e.$, initial state estimation error gets arbitrarily small (at least in some directions) as more and more of signals u and y are observed

Continuous-time least-squares state estimation

assume $\dot{x} = Ax + Bu$, $y = Cx + Du + v$ is observable

least-squares estimate of initial state $x(0)$, given $u(\tau)$, $y(\tau)$, $0 \leq \tau \leq t$: choose $\hat{x}_{\rm ls}(0)$ to minimize integral square residual

$$J = \int_0^t \left\| \tilde{y}(\tau) - Ce^{\tau A}x(0) \right\|^2 \, d\tau$$

where $\tilde{y} = y - h * u$ is observed output minus part due to input

let's expand as $J = x(0)^T Q x(0) + 2r^T x(0) + s$,

$$Q = \int_0^t e^{\tau A^T} C^T C e^{\tau A} \, d\tau, \quad r = \int_0^t e^{\tau A^T} C^T \tilde{y}(\tau) \, d\tau,$$

$$q = \int_0^t \tilde{y}(\tau)^T \tilde{y}(\tau) \, d\tau$$

setting $\nabla_{x(0)} J$ to zero, we obtain the least-squares observer

$$\hat{x}_{\mathrm{ls}}(0) = Q^{-1} r = \left(\int_0^t e^{\tau A^T} C^T C e^{\tau A} \, d\tau \right)^{-1} \int_0^t e^{A^T \tau} C^T \tilde{y}(\tau) \, d\tau$$

estimation error is

$$\tilde{x}(0) = \hat{x}_{\mathrm{ls}}(0) - x(0) = \left(\int_0^t e^{\tau A^T} C^T C e^{\tau A} \, d\tau \right)^{-1} \int_0^t e^{\tau A^T} C^T v(\tau) \, d\tau$$

therefore if $v = 0$ then $\hat{x}_{\mathrm{ls}}(0) = x(0)$

System Identification*

*partial, discreet time, as LTI

- Given examples $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ we want to model the relation between $\mathbf{x}_i$ and $\mathbf{y}_i$ as $\mathbf{y}_i \approx \mathbf{A}\mathbf{x}_i$. Define the estimation problem as:

$$\arg\min_{\mathbf{A}} \sum_{i=1}^{N} ||\mathbf{y}_i - \mathbf{A}\mathbf{x}_i||^2 \quad = \quad \sum_{i=1}^{N} \mathbf{x}_i'\mathbf{A}'\mathbf{A}\mathbf{x}_i - 2\mathbf{y}_i'\mathbf{A}\mathbf{x}_i + \mathbf{y}_i'\mathbf{y}_i$$

- we differentiate w.r.t. $\mathbf{A}$ and set to 0

$$\frac{\partial}{\partial \mathbf{A}} \sum_{i=1}^{N} ||\mathbf{y}_i - \mathbf{A}\mathbf{x}_i||^2 \quad = \quad 0$$

$$\sum_{i=1}^{N} 2\mathbf{A}\mathbf{x}_i\mathbf{x}_i' - 2\mathbf{y}_i\mathbf{x}_i' \quad = \quad 0$$

$$\mathbf{A}\sum_{i=1}^{N} \mathbf{x}_i\mathbf{x}_i' \quad = \quad \sum_{i=1}^{N} \mathbf{y}_i\mathbf{x}_i'$$

$$\mathbf{A}\sum_{i=1}^{N}\mathbf{x}_i\mathbf{x}_i' = \sum_{i=1}^{N}\mathbf{y}_i\mathbf{x}_i'$$

- If rank of $\sum_{i=1}^{N}\mathbf{x}_i\mathbf{x}_i'$ is full rank (requires N > n) then

$$\mathbf{A} = \left(\sum_{i=1}^{N}\mathbf{y}_i\mathbf{x}_i'\right)\left(\sum_{i=1}^{N}\mathbf{x}_i\mathbf{x}_i'\right)^{-1}$$

- E.g. we'd like to estimate **A** in system: $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \omega$ (ω is noise). To solve, simply replace $\mathbf{y}_i$ with $\mathbf{x}_{i+1}$ in above solution.

- Note that this would also be the most likely **A** if ω were Gaussian noise with zero mean and unit variance.

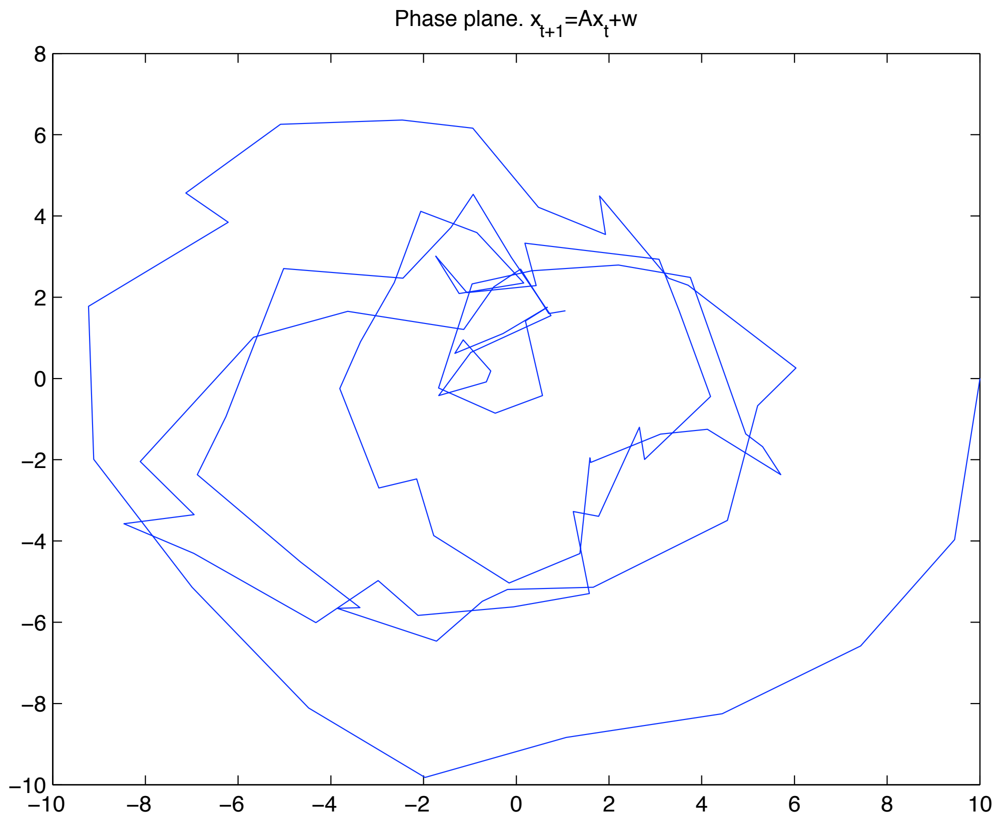

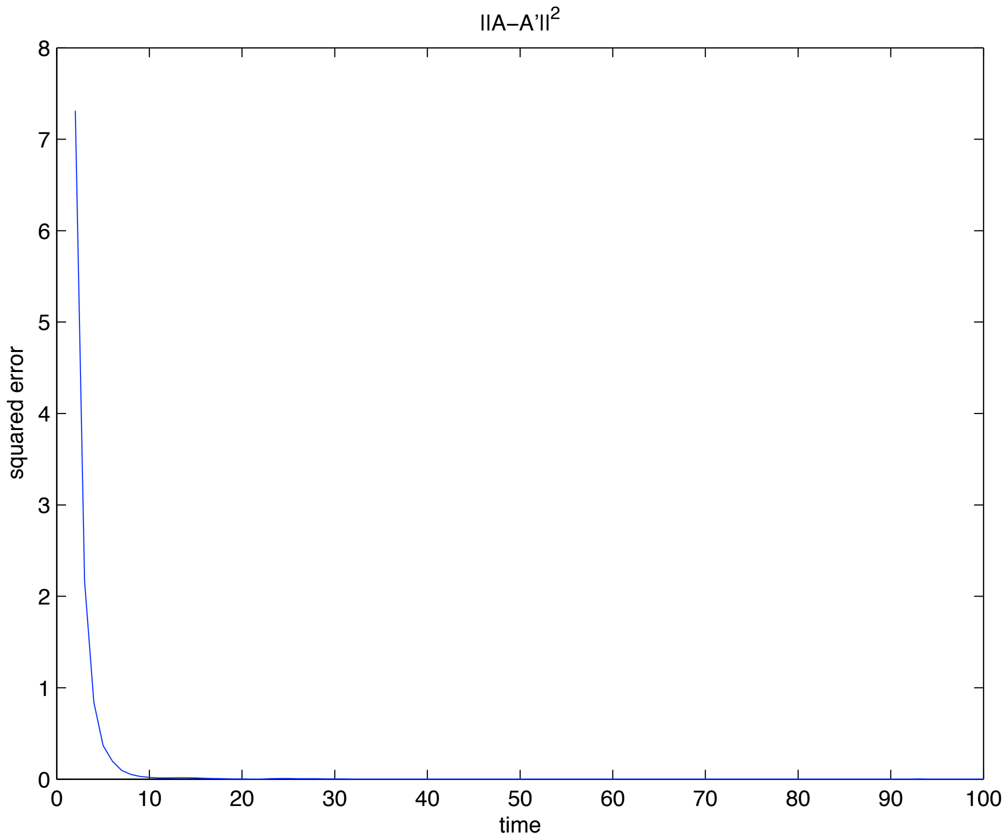

- Left: Phase plane, values of $\mathbf{x}_t$ where $\omega \sim N(\mathbf{0},\mathbf{I})$
- Right: Squared error between true and estimated **A** as function of step number. error = $\sum_{i,j}(\mathbf{A}_{ij}-\hat{\mathbf{A}}_{ij})^2$ is called the Frobenius norm.

Lecture 6
Estimation

- Gaussian random vectors
- minimum mean-square estimation (MMSE)
- MMSE with linear measurements
- relation to least-squares, pseudo-inverse

Gaussian random vectors

random vector $x \in \mathbf{R}^n$ is *Gaussian* if it has density

$$p_x(v) = (2\pi)^{-n/2}(\det \Sigma)^{-1/2} \exp\left(-\frac{1}{2}(v-\bar{x})^T \Sigma^{-1}(v-\bar{x})\right),$$

for some $\Sigma = \Sigma^T > 0$, $\bar{x} \in \mathbf{R}^n$

- denoted $x \sim \mathcal{N}(\bar{x}, \Sigma)$
- $\bar{x} \in \mathbf{R}^n$ is the *mean* or *expected* value of x, $i.e.$,

$$\bar{x} = \mathbf{E}\, x = \int v p_x(v) dv$$

- $\Sigma = \Sigma^T > 0$ is the *covariance* matrix of x, $i.e.$,

$$\Sigma \quad = \quad \mathbf{E}(x-\bar{x})(x-\bar{x})^T$$

$$= \quad \mathbf{E}\, xx^T - \bar{x}\bar{x}^T$$

$$= \quad \int (v - \bar{x})(v - \bar{x})^T p_x(v)\, dv$$

density for $x \sim \mathcal{N}(0,1)$:

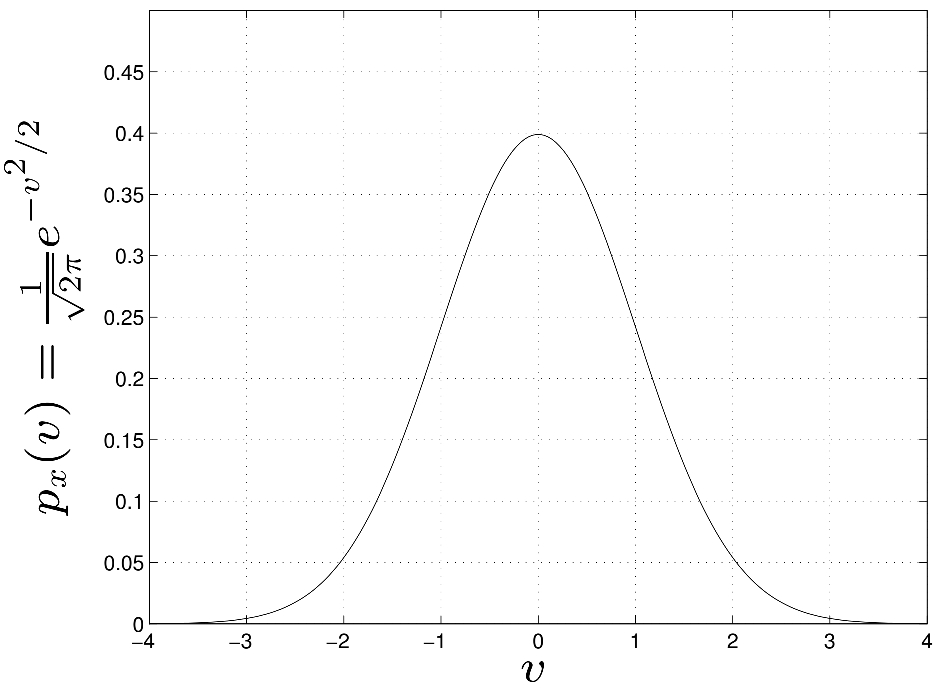

- mean and variance of scalar random variable x_i are

$$\mathbf{E}\, x_i = \bar{x}_i, \quad \mathbf{E}(x_i - \bar{x}_i)^2 = \Sigma_{ii}$$

  hence standard deviation of x_i is $\sqrt{\Sigma_{ii}}$
- covariance between x_i and x_j is $\mathbf{E}(x_i - \bar{x}_i)(x_j - \bar{x}_j) = \Sigma_{ij}$
- correlation coefficient between x_i and x_j is $\rho_{ij} = \dfrac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}$
- mean (norm) square deviation of x from $\bar{x}$ is

$$\mathbf{E}\, \|x - \bar{x}\|^2 = \mathbf{E}\, \mathbf{Tr}(x - \bar{x})(x - \bar{x})^T = \mathbf{Tr}\, \Sigma = \sum_{i=1}^{n} \Sigma_{ii}$$

  (using $\mathbf{Tr}\, AB = \mathbf{Tr}\, BA$)

example: $x \sim \mathcal{N}(0, I)$ means x_i are independent identically distributed (IID) $\mathcal{N}(0, 1)$ random variables

Confidence ellipsoids

$p_x(v)$ is constant for $(v - \bar{x})^T \Sigma^{-1} (v - \bar{x}) = \alpha$, *i.e.*, on the surface of ellipsoid

$$\mathcal{E}_\alpha = \{v \mid (v - \bar{x})^T \Sigma^{-1} (v - \bar{x}) \leq \alpha\}$$

thus $\bar{x}$ and Σ determine shape of density

can interpret $\mathcal{E}_\alpha$ as *confidence ellipsoid* for x:

the nonnegative random variable $(x - \bar{x})^T \Sigma^{-1} (x - \bar{x})$ has a χ_n^2 distribution, so $\mathbf{Prob}(x \in \mathcal{E}_\alpha) = F_{\chi_n^2}(\alpha)$ where $F_{\chi_n^2}$ is the CDF

some good approximations:

- $\mathcal{E}_n$ gives about 50% probability
- $\mathcal{E}_{n+2\sqrt{n}}$ gives about 90% probability

geometrically:

- mean $\bar{x}$ gives center of ellipsoid
- semiaxes are $\sqrt{\alpha\lambda_i}u_i$, where u_i are (orthonormal) eigenvectors of Σ with eigenvalues λ_i

example: $x \sim \mathcal{N}(\bar{x}, \Sigma)$ with $\bar{x} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$, $\Sigma = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$

- x_1 has mean 2, std. dev. $\sqrt{2}$
- x_2 has mean 1, std. dev. 1
- correlation coefficient between x_1 and x_2 is $\rho = 1/\sqrt{2}$
- $\mathbf{E}\,\|x - \bar{x}\|^2 = 3$

90% confidence ellipsoid corresponds to $\alpha = 4.6$:

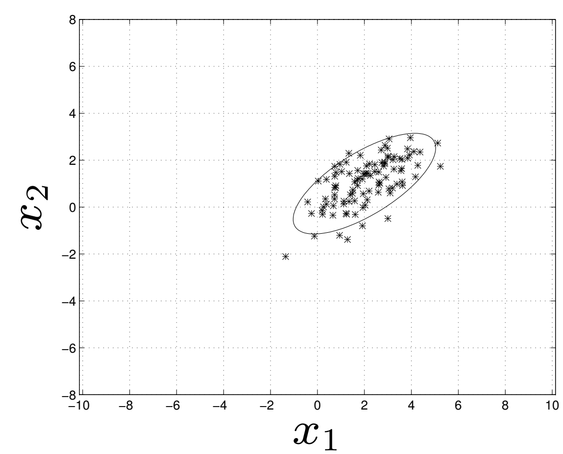

(here, 91 out of 100 fall in $\mathcal{E}_{4.6}$)

Affine transformation

suppose $x \sim \mathcal{N}(\bar{x}, \Sigma_x)$

consider affine transformation of x:

$$z = Ax + b,$$

where $A \in \mathbf{R}^{m\times n}$, $b \in \mathbf{R}^m$

then z is Gaussian, with mean

$$\mathbf{E}\, z = \mathbf{E}(Ax + b) = A\,\mathbf{E}\, x + b = A\bar{x} + b$$

and covariance

$$\begin{aligned} \Sigma_z &= \mathbf{E}(z - \bar{z})(z - \bar{z})^T \\ &= \mathbf{E}\, A(x - \bar{x})(x - \bar{x})^T A^T \\ &= A\Sigma_x A^T \end{aligned}$$

examples:

- if $w \sim \mathcal{N}(0, I)$ then $x = \Sigma^{1/2} w + \bar{x}$ is $\mathcal{N}(\bar{x}, \Sigma)$

  useful for simulating vectors with given mean and covariance

- conversely, if $x \sim \mathcal{N}(\bar{x}, \Sigma)$ then $z = \Sigma^{-1/2}(x - \bar{x})$ is $\mathcal{N}(0, I)$

  (normalizes & decorrelates)

suppose $x \sim \mathcal{N}(\bar{x}, \Sigma)$ and $c \in \mathbf{R}^n$

scalar $c^T x$ has mean $c^T \bar{x}$ and variance $c^T \Sigma c$

thus (unit length) direction of minimum variability for x is u, where

$$\Sigma u = \lambda_{\min} u, \quad \|u\| = 1$$

standard deviation of $u_n^T x$ is $\sqrt{\lambda_{\min}}$

(similarly for maximum variability)

Degenerate Gaussian vectors

it is convenient to allow Σ to be singular (but still $\Sigma = \Sigma^T \geq 0$)

(in this case density formula obviously does not hold)

meaning: in some directions x is not random at all

write Σ as

$$\Sigma = [Q_+ \; Q_0] \begin{bmatrix} \Sigma_+ & 0 \\ 0 & 0 \end{bmatrix} [Q_+ \; Q_0]^T$$

where $Q = [Q_+ \; Q_0]$ is orthogonal, $\Sigma_+ > 0$

- columns of Q_0 are orthonormal basis for $\mathcal{N}(\Sigma)$
- columns of Q_+ are orthonormal basis for $\mathrm{range}(\Sigma)$

then $Q^T x = [z^T \ w^T]^T$, where

- $z \sim \mathcal{N}(Q_+^T \bar{x}, \Sigma_+)$ is (nondegenerate) Gaussian (hence, density formula holds)
- $w = Q_0^T \bar{x} \in \mathbf{R}^n$ is not random

  ($Q_0^T x$ is called *deterministic component* of x)

Linear measurements

linear measurements with noise:

$$y = Ax + v$$

- $x \in \mathbf{R}^n$ is what we want to measure or estimate
- $y \in \mathbf{R}^m$ is measurement
- $A \in \mathbf{R}^{m \times n}$ characterizes sensors or measurements
- v is sensor noise

common assumptions:

- $x \sim \mathcal{N}(\bar{x}, \Sigma_x)$
- $v \sim \mathcal{N}(\bar{v}, \Sigma_v)$
- x and v are independent

- $\mathcal{N}(\bar{x}, \Sigma_x)$ is the *prior distribution* of x (describes initial uncertainty about x)
- $\bar{v}$ is noise *bias* or *offset* (and is usually 0)
- Σ_v is noise *covariance*

thus

$$\begin{bmatrix} x \\ v \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \bar{x} \\ \bar{v} \end{bmatrix}, \begin{bmatrix} \Sigma_x & 0 \\ 0 & \Sigma_v \end{bmatrix} \right)$$

using

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} I & 0 \\ A & I \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix}$$

we can write

$$\mathbf{E} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \bar{x} \\ A\bar{x} + \bar{v} \end{bmatrix}$$

and

$$\begin{aligned} \mathbf{E} \begin{bmatrix} x - \bar{x} \\ y - \bar{y} \end{bmatrix} \begin{bmatrix} x - \bar{x} \\ y - \bar{y} \end{bmatrix}^T &= \begin{bmatrix} I & 0 \\ A & I \end{bmatrix} \begin{bmatrix} \Sigma_x & 0 \\ 0 & \Sigma_v \end{bmatrix} \begin{bmatrix} I & 0 \\ A & I \end{bmatrix}^T \\ &= \begin{bmatrix} \Sigma_x & \Sigma_x A^T \\ A\Sigma_x & A\Sigma_x A^T + \Sigma_v \end{bmatrix} \end{aligned}$$

covariance of measurement y is $A\Sigma_x A^T + \Sigma_v$

- $A\Sigma_x A^T$ is 'signal covariance'
- Σ_v is 'noise covariance'

Minimum mean-square estimation

suppose $x \in \mathbf{R}^n$ and $y \in \mathbf{R}^m$ are random vectors (not necessarily Gaussian)

we seek to estimate x given y

thus we seek a function $\phi : \mathbf{R}^m \to \mathbf{R}^n$ such that $\hat{x} = \phi(y)$ is near x

one common measure of nearness: mean-square error,

$$\mathbf{E}\, \|\phi(y) - x\|^2$$

minimum mean-square estimator (MMSE) $\phi_{\rm mmse}$ minimizes this quantity

general solution: $\phi_{\rm mmse}(y) = \mathbf{E}(x|y)$, *i.e.*, the conditional expectation of x given y

MMSE for Gaussian vectors

now suppose $x \in \mathbf{R}^n$ and $y \in \mathbf{R}^m$ are jointly Gaussian:

$$\left[\begin{array}{c} x \\ y \end{array}\right] \sim \mathcal{N}\left(\left[\begin{array}{c} \bar{x} \\ \bar{y} \end{array}\right], \left[\begin{array}{cc} \Sigma_x & \Sigma_{xy} \\ \Sigma_{xy}^T & \Sigma_y \end{array}\right] \right)$$

(after alot of algebra) the conditional density is

$$p_{x|y}(v|y) = (2\pi)^{-n/2}(\det \Lambda)^{-1/2} \exp\left(-\frac{1}{2}(v-w)^T \Lambda^{-1}(v-w)\right),$$

where

$$\Lambda = \Sigma_x - \Sigma_{xy}\Sigma_y^{-1}\Sigma_{xy}^T, \quad w = \bar{x} + \Sigma_{xy}\Sigma_y^{-1}(y-\bar{y})$$

hence MMSE estimator ($i.e.$, conditional expectation) is

$$\hat{x} = \phi_{\mathrm{mmse}}(y) = \mathbf{E}(x|y) = \bar{x} + \Sigma_{xy}\Sigma_y^{-1}(y-\bar{y})$$

ϕ_{mmse} is an affine function

MMSE estimation error, $\hat{x} - x$, is a Gaussian random vector

$$\hat{x} - x \;\sim\; \mathcal{N}(0, \Sigma_x - \Sigma_{xy}\Sigma_y^{-1}\Sigma_{xy}^T)$$

note that

$$\Sigma_x - \Sigma_{xy}\Sigma_y^{-1}\Sigma_{xy}^T \leq \Sigma_x$$

$i.e.$, covariance of estimation error is always less than prior covariance of x

Best linear unbiased estimator

estimator

$$\hat{x} = \phi_{\mathrm{blu}}(y) = \bar{x} + \Sigma_{xy}\Sigma_y^{-1}(y - \bar{y})$$

makes sense when x, y aren't jointly Gaussian

this estimator

- is *unbiased*, $i.e.$, $\mathbf{E}\,\hat{x} = \mathbf{E}\,x$
- often works well
- is widely used
- has minimum mean square error among all *affine* estimators

sometimes called *best linear unbiased* estimator

MMSE with linear measurements

consider specific case

$$y = Ax + v, \quad x \sim \mathcal{N}(\bar{x}, \Sigma_x), \quad v \sim \mathcal{N}(\bar{v}, \Sigma_v),$$

x, v independent

MMSE of x given y is affine function

$$\hat{x} = \bar{x} + B(y - \bar{y})$$

where $B = \Sigma_x A^T (A \Sigma_x A^T + \Sigma_v)^{-1}$, $\bar{y} = A\bar{x} + \bar{v}$

intepretation:

- $\bar{x}$ is our best prior guess of x (before measurement)
- $y - \bar{y}$ is the discrepancy between what we actually measure (y) and the expected value of what we measure ($\bar{y}$)

- estimator modifies prior guess by B times this discrepancy
- estimator blends prior information with measurement
- B gives *gain* from *observed discrepancy* to *estimate*
- B is small if noise term Σ_v in 'denominator' is large

MMSE error with linear measurements

MMSE estimation error, $\tilde{x} = \hat{x} - x$, is Gaussian with zero mean and covariance

$$\Sigma_{\rm est} = \Sigma_x - \Sigma_x A^T (A\Sigma_x A^T + \Sigma_v)^{-1} A\Sigma_x$$

- $\Sigma_{\rm est} \leq \Sigma_x$, *i.e.*, measurement always decreases uncertainty about x
- difference $\Sigma_x - \Sigma_{\rm est}$ gives *value* of measurement y in estimating x
- *e.g.*, $(\Sigma_{\rm est\ \it ii}/\Sigma_{x\ ii})^{1/2}$ gives fractional decrease in uncertainty of x_i due to measurement

note: error covariance $\Sigma_{\rm est}$ can be determined *before* measurement y is made!

to evaluate $\Sigma_{\rm est}$, only need to know

- A (which characterizes sensors)
- prior covariance of x ($i.e.$, Σ_x)
- noise covariance ($i.e.$, Σ_v)

you *do not* need to know the measurement y (or the means $\bar{x}$, $\bar{v}$)

useful for *experiment design* or *sensor selection*