COM521500
Math. Methods for SP I
Lecture 8: Least Squares

Least Squares (LS) Problem

In LS, we are concerned with solving

$$\min_{\mathbf{x}\in\mathbb{C}^n} \|\mathbf{A}\mathbf{x}-\mathbf{b}\|_2^2$$

for $\mathbf{x}$, given $\mathbf{A}\in\mathbb{C}^{m\times n}$, $m>n$, and $\mathbf{b}\in\mathbb{C}^m$.

In essence, $\mathbf{A}\mathbf{x}-\mathbf{b}$ represents an error vector and we seek to minimize the sum square of the error vector.

There are **so many** applications for LS.

Application I: System Identification

Let $u[n]$ be an input signal that passes through a linear time-invariant system. The output is given by

$$x[n] = \sum_{\ell=0}^{L-1} h[\ell]u[n-\ell] + \nu[n]$$

where $h[n]$ is the impulse response of the system.

Our aim is to estimate $h[n]$ from $x[n]$, given that $u[n]$ is known.

Let $\mathbf{u}[n] = [\ u[n], u[n-1], \ldots, u[n-L+1]\ ]^T$, and $\mathbf{h} = [\ h[0], h[1], \ldots, h[L-1]\ ]^T$. The output signal can be re-expressed as

$$x[n] = \mathbf{u}^T[n]\mathbf{h} + \nu[n]$$

System identification can be done by minimizing the sum squared error:

$$\min_{\mathbf{h}\in\mathbb{C}^L} \sum_{n=1}^{N} \left|\mathbf{u}^T[n]\mathbf{h} - x[n]\right|^2$$

where N is the data length.

Let $\mathbf{x} = [\ x[1], \ldots, x[N]\ ]^T$. We have

$$\mathbf{x} = \mathbf{U}\mathbf{h} + \boldsymbol{\nu}$$

where $\mathbf{U} = [\ \mathbf{u}[1], \ldots, \mathbf{u}[N]\ ]^T$.

The system identification problem can be rewritten as

$$\min_{\mathbf{h} \in \mathbb{C}^L} \|\mathbf{U}\mathbf{h} - \mathbf{x}\|_2^2$$

which is an LS.

Application II: Channel Equalization

In digital communication over a linear time-dispersive channel, the discrete signal model is generally formulated as:

$$x[n] = \sum_{\ell=0}^{L-1} h[\ell] u[n-\ell] + \nu[n]$$

where

- $u[n]$ transmitted symbol sequence
- $h[n]$ channel impulse response
- $x[n]$ received signal.

At the receiver, we apply a filtering process, called equalization

$$y[n] = \sum_{\ell=0}^{m-1} w[\ell] x[n-\ell]$$

so that $y[n] \approx u[n]$.

Let $\mathbf{x}[n] = [\ x[n], x[n-1], \ldots, x[n-m+1]\ ]^T$, and $\mathbf{w} = [\ w[0], \ldots, w[m-1]\ ]^T$. The equalizer output equation can be rewritten as

$$y[n] = \mathbf{x}^T[n]\mathbf{w}$$

Suppose that $u[n]$ is known for $n = 0, 1, \ldots, N-1$. In practice, this is made possible by having the transmitter sending signals known to the receiver, a.k.a. pilot signals.

The equalizer coefficients $w[n]$ are determined by

$$\min_{\mathbf{w} \in \mathbb{C}^m} \sum_{n=1}^{N} \left| \mathbf{x}^T[n]\mathbf{w} - u[n] \right|^2$$
$$= \min_{\mathbf{w} \in \mathbb{C}^m} \|\mathbf{X}\mathbf{w} - \mathbf{u}\|_2^2$$

where $\mathbf{X} = [\ \mathbf{x}[1], \ldots, \mathbf{x}[N]\ ]^T$, and $\mathbf{u} = [\ u[0], \ldots, u[N-1]\ ]^T$.

The problem is again an LS.

Application III: Curve Fitting

Consider that there is a collection of experimental measurements, denoted by $x(t_1), x(t_2), \ldots, x(t_N)$.

We seek to find a continuous curve that 'fits' those data.

Suppose that the curve can be parameterized as

$$y(t) = \theta_1 + \theta_2 t + \theta_3 t^2$$

and assume that $x(t_i)$ are perturbed versions of $y(t)$

$$x(t_i) = y(t_i) + \nu(t_i)$$

where $\nu(t_i)$ is noise.

COM521500 Math. Methods for Signal Processing I Lecture 8: Least Squares

Let $\mathbf{x} = [\ x(t_1), \ldots, x(t_N)\ ]^T$. We have

$$\mathbf{x} = \mathbf{H}\boldsymbol{\theta} + \boldsymbol{\nu}$$

where $\boldsymbol{\theta} = [\ \theta_1, \theta_2, \theta_3\ ]^T$, and

$$\mathbf{H} = \begin{bmatrix} 1 & t_1 & t_2^2 \\ 1 & t_2 & t_2^2 \\ & \vdots & \\ 1 & t_N & t_N^2 \end{bmatrix}$$

Again, we can use LS

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^3} \|\mathbf{H}\boldsymbol{\theta} - \mathbf{x}\|_2^2$$

to determine the curve coefficients.

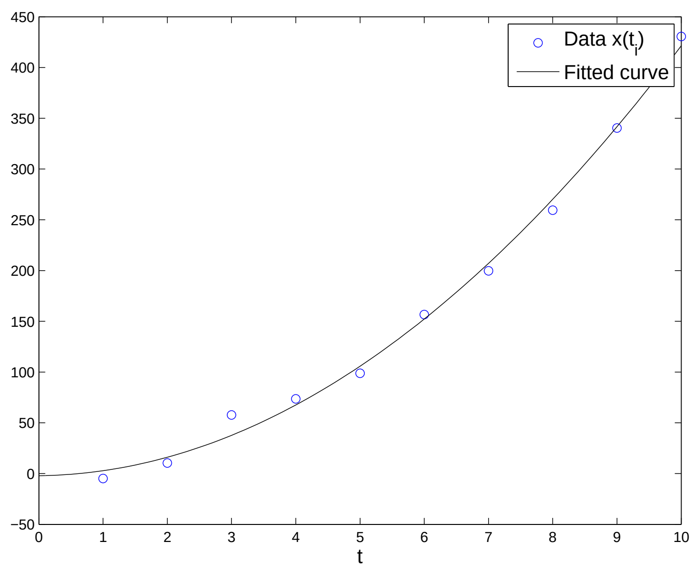

Application IV: Linear Prediction

A colored process, denoted by $y[n]$ can be modeled as

$$y[n] = \sum_{\ell=0}^{\infty} h[\ell] w[n-\ell]$$

where $w[n]$ is a zero-mean white process.

Here we are interested in the autoregressive (AR) process.

In this process $h[n]$ is an all-pole model; i.e., its z-transform is given by

$$H(z) = 1/A(z)$$

$$A(z) = 1 - \sum_{i=1}^{m} a_i z^{-i}$$

Since

$$Y(z) = H(z)W(z)$$

we have that

$$Y(z)A(z) = W(z)$$

and that

$$y[n] - \sum_{i=1}^{m} a_i y[n-i] = w[n] \qquad (*)$$

Eq. $(*)$ can be viewed as a 'prediction', where the previous samples $\{y[n-i]\}_{i=1}^{m}$ predict the present sample $y[n]$, up to a (unpredictable) perturbation $w[n]$.

Our aim is to estimate $\mathbf{a} = [\ a_1, \ldots, a_m\ ]^T$ from $y[n]$.

Let

$$\mathbf{y}_{\rm p} = [\ y[1], \ldots, y[N]\ ]^T$$
$$\mathbf{y}[n] = [\ y[n-1], \ldots, y[n-m]\ ]^T$$
$$\mathbf{Y} = [\ \mathbf{y}[1], \ldots, \mathbf{y}[N]\ ]^T$$

AR coefficient estimation may be achieved by LS linear prediction:

$$\min_{\mathbf{a} \in \mathbb{C}^m} \|\mathbf{Y}\mathbf{a} - \mathbf{y}_{\rm p}\|_2^2$$

Solving LS

First, some remarks:

- In Lecture 4, we have learnt that for $m > n$,

$$\mathbf{A}\mathbf{x} - \mathbf{b} \neq \mathbf{0}$$

  in general, unless $\mathbf{b} \in R(\mathbf{A})$.

- If $\text{rank}(\mathbf{A}) < n$, then the solution set

$$\{ \mathbf{x}_{LS} \in \mathbb{C}^n \mid \|\mathbf{A}\mathbf{x}_{LS} - \mathbf{b}\|_2^2 = \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \}$$

  does not simply contain one element— if $\mathbf{x}_{LS}$ is a solution, then $\mathbf{x}_{LS} + \mathbf{z}$, $\mathbf{z} \in N(\mathbf{A})$ is also a solution.

Alternative I for solving LS: use Gradient

The gradient of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined to be

$$\nabla f = \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix}$$

Some useful properties for gradients:

1. The gradient of $f(\mathbf{x}) = \mathbf{x}^T\mathbf{b}$ is $\nabla f = \mathbf{b}$.
2. The gradient of $f(\mathbf{x}) = \mathbf{x}^T\mathbf{R}\mathbf{x}$ where $\mathbf{R}$ is symmetric, is $\nabla f = 2\mathbf{R}\mathbf{x}$.

For ease of exposition of ideas, assume that $\mathbf{A}$, $\mathbf{b}$, & $\mathbf{x}$ are real-valued.

Let

$$f(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$$

The LS problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$$

is an unconstrained optimization problem. Since f is convex, the sufficient & necessary condition for $\mathbf{x}_{LS}$ to be a solution is that

$$\nabla f|_{\mathbf{x}=\mathbf{x}_{LS}} = \mathbf{0}$$

We can decompose

$$f(\mathbf{x}) = \mathbf{x}^T\mathbf{A}^T\mathbf{A}\mathbf{x} - 2\mathbf{x}^T\mathbf{A}^T\mathbf{b} + \mathbf{b}^T\mathbf{b}$$

The gradient of f is

$$\nabla f = 2\mathbf{A}^T\mathbf{A}\mathbf{x} - 2\mathbf{A}^T\mathbf{b}$$

Hence, an optimal solution $\mathbf{x}_{LS}$ can be found by solving

$$\mathbf{A}^T\mathbf{A}\mathbf{x}_{LS} = \mathbf{A}^T\mathbf{b}$$

For the complex case, it can be shown (in a similar way but with more hassles) that

$$\mathbf{A}^H\mathbf{A}\mathbf{x}_{LS} = \mathbf{A}^H\mathbf{b}$$

Alternative II for solving LS: use the Orthogonal Principle

Theorem 8.1 (Orthogonal Principle) A vector $\mathbf{x}_{LS}$ is an LS solution if and only if

$$\mathbf{A}^H(\mathbf{A}\mathbf{x}_{LS} - \mathbf{b}) = \mathbf{0}$$

The equations

$$\mathbf{A}^H\mathbf{A}\mathbf{x}_{LS} = \mathbf{A}^H\mathbf{b}$$

are referred as to the **normal equations.**

If $\mathbf{A}$ is of full column rank so that $\mathbf{A}^H\mathbf{A}$ is PD, then $\mathbf{x}_{LS}$ is uniquely determined by

$$\mathbf{x}_{LS} = (\mathbf{A}^H\mathbf{A})^{-1}\mathbf{A}^H\mathbf{b}$$

Interpretations of the Normal Equations

Let $\mathbf{r}_{LS} = \mathbf{b} - \mathbf{A}\mathbf{x}_{LS}$ be the LS error vector.

For full rank $\mathbf{A}$,

$$\begin{aligned} \mathbf{r}_{LS} &= \mathbf{b} - \mathbf{A}(\mathbf{A}^H\mathbf{A})^{-1}\mathbf{A}^H\mathbf{b} \\ &= \mathbf{b} - \mathbf{P}\mathbf{b} = \mathbf{P}_{\perp}\mathbf{b} \end{aligned}$$

where $\mathbf{P}$ is the orthogonal projection matrix of $\mathbf{A}$, and $\mathbf{P}_{\perp}$ is the orthogonal complement.

This means that the LS error is orthogonal to any vector in $R(\mathbf{A})$.

LS for Rank Deficient $\mathbf{A}$

As we mentioned, for rank deficient $\mathbf{A}$ there are more than one LS solutions.

But we can find a unique $\mathbf{x}_{LS}$ that has its 2-norm being the smallest among all LS solutions.

Let $r = \text{rank}(\mathbf{A})$, and denote the SVD of $\mathbf{A}$ by

$$\begin{aligned} \mathbf{A} &= \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^H \\ &= [\ \mathbf{U}_1\ \mathbf{U}_2\ ] \begin{bmatrix} \tilde{\boldsymbol{\Sigma}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{V}_1^H \\ \mathbf{V}_2^H \end{bmatrix} \end{aligned}$$

where $\tilde{\boldsymbol{\Sigma}} = \text{Diag}(\sigma_1, \ldots, \sigma_r)$ contains the nonzero singular values of $\mathbf{A}$.

Define

$$\mathbf{A}^\dagger = \mathbf{V}_1 \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{U}_1^H$$

to be the **pseudo-inverse** of $\mathbf{A}$.

Theorem 8.2 The following minimum 2-norm problem

$$\begin{aligned} &\min\ \|\mathbf{x}\|_2^2 \\ &\text{s.t.}\ \ \mathbf{x} \text{ minimizes } \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \end{aligned}$$

is uniquely given by

$$\mathbf{x}_{LS} = \mathbf{A}^\dagger \mathbf{b}$$

Note that

$$\begin{aligned}\mathbf{r}_{LS} &= \mathbf{b} - \mathbf{A}\mathbf{x}_{LS} \\ &= \mathbf{b} - (\mathbf{U}_1\tilde{\mathbf{\Sigma}}\mathbf{V}_1^H)(\mathbf{V}_1\tilde{\mathbf{\Sigma}}^{-1}\mathbf{U}_1^H)\mathbf{b} \\ &= \mathbf{b} - \mathbf{U}_1\mathbf{U}_1^H\mathbf{b} \\ &= \mathbf{b} - \mathbf{P}\mathbf{b} = \mathbf{P}_\perp\mathbf{b}\end{aligned}$$

where $\mathbf{P} = \mathbf{U}_1\mathbf{U}_1^H$ is the orthogonal projection matrix of $\mathbf{A}$.

This orthogonal property is the same as that in the case of full column rank $\mathbf{A}$.

Some Relationships of the pseudo-inverse

1. For the case of full column rank $\mathbf{A}$ (i.e., $m \geq n$, $\mathrm{rank}(\mathbf{A}) = n$),

$$\mathbf{A}^\dagger = (\mathbf{A}^H\mathbf{A})^{-1}\mathbf{A}^H\mathbf{b}$$

which means that the pseudo-inverse leads to the LS in the full column rank case.

2. For the case of full row rank $\mathbf{A}$ (i.e., $m \leq n$, $\mathrm{rank}(\mathbf{A}) = m$),

$$\mathbf{A}^\dagger = \mathbf{A}^H(\mathbf{A}\mathbf{A}^H)^{-1}$$

Relationship to generalized inverse

A matrix $\mathbf{C} \in \mathbb{C}^{n \times m}$ is said to be the **Moore-Penrose generalized inverse** of $\mathbf{A}$ if the following 4 conditions hold:

1. $\mathbf{ACA} = \mathbf{A}$
2. $\mathbf{CAC} = \mathbf{C}$
3. $(\mathbf{AC})^H = \mathbf{AC}$
4. $(\mathbf{CA})^H = \mathbf{CA}$

It can be verified that $\mathbf{A}^\dagger$ is the Moore-Penrose generalized inverse.