MA 5037: Optimization Methods and Applications
Least-Squares Problem

Suh-Yuh Yang (楊肅煜)

Department of Mathematics, National Central University
Jhongli District, Taoyuan City 320317, Taiwan

First version: May 21, 2018/Last updated: June 14, 2025

Solution of overdetermined systems

Consider an overdetermined linear system:

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b},$$

where $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $m \geq n$, and $\boldsymbol{b} \in \mathbb{R}^m$. *We assume that* $\boldsymbol{A}$ *has a full column rank, rank*$(\boldsymbol{A}) = n$. In this setting, the system is usually inconsistent (has no solution) and a common approach for finding an approximate solution is to

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2, \qquad (LS)$$

or equivalently, to

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \left\{ f(\boldsymbol{x}) := \boldsymbol{x}^\top(\boldsymbol{A}^\top\boldsymbol{A})\boldsymbol{x} - 2(\boldsymbol{A}^\top\boldsymbol{b})^\top\boldsymbol{x} + \|\boldsymbol{b}\|^2 \right\}. \qquad (LS)$$

Since $\boldsymbol{A}$ is of full column rank, $\nabla^2 f(\boldsymbol{x}) = 2\boldsymbol{A}^\top\boldsymbol{A} \succ \boldsymbol{0}, \forall\, \boldsymbol{x} \in \mathbb{R}^n$. Therefore, (by Lemma 2.41), the unique stationary point

$$\boldsymbol{x}_{LS} = (\boldsymbol{A}^\top\boldsymbol{A})^{-1}\boldsymbol{A}^\top\boldsymbol{b}$$

is the optimal solution of problem (LS), and $\boldsymbol{x}_{LS}$ *is called the least-squares solution of the system* $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$.

The normal system

- It is quite common not to write the explicit expression for $\boldsymbol{x}_{LS}$ but instead to write the associated system of equations that defines it:

$$(\boldsymbol{A}^\top \boldsymbol{A})\boldsymbol{x}_{LS} = \boldsymbol{A}^\top \boldsymbol{b}.$$

The above system of equations is called the normal system.

- If $m = n$ and $\boldsymbol{A}$ is of full column rank, then $\boldsymbol{A}$ is nonsingular. In this case, the least-squares solution is actually the solution of the linear system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, since

$$\boldsymbol{x}_{LS} = (\boldsymbol{A}^\top \boldsymbol{A})^{-1}\boldsymbol{A}^\top \boldsymbol{b} = \boldsymbol{A}^{-1}\boldsymbol{A}^{-\top}\boldsymbol{A}^\top \boldsymbol{b} = \boldsymbol{A}^{-1}\boldsymbol{b} = \boldsymbol{x}.$$

Example

Consider the inconsistent linear system

$$\boldsymbol{A}\boldsymbol{x} = \begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 3 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} = \boldsymbol{b}.$$

The least-squares problem can be explicitly written as

$$\min_{(x_1,x_2)^\top \in \mathbb{R}^2} \left\{ (x_1 + 2x_2)^2 + (2x_1 + x_2 - 1)^2 + (3x_1 + 2x_2 - 1)^2 \right\}.$$

We will solve the normal equations:

$$\begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 3 & 2 \end{bmatrix}^\top \begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 3 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 3 & 2 \end{bmatrix}^\top \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix},$$

which are the same as

$$\begin{bmatrix} 14 & 10 \\ 10 & 9 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 5 \\ 3 \end{bmatrix}.$$

Example (cont'd)

The solution of the above system is the least-squares estimate

$$x_{LS} = \begin{bmatrix} 15/26 \\ -8/26 \end{bmatrix}.$$

The residual vector is given by

$$r := Ax_{LS} - b = \begin{bmatrix} -0.038 \\ -0.154 \\ 0.115 \end{bmatrix},$$

and $\|r\|_2^2 = (-0.038)^2 + (-0.154)^2 + (0.115)^2 \approx 0.038.$

To find the least-squares solution in MATLAB:

```
>> A = [1, 2; 2, 1; 3, 2];
>> b = [0; 1; 1];
>> format rational;
>> A\b
ans =
15/26
-4/13
```

Data fitting: linear fitting

Suppose that we are given a set of data points $(\boldsymbol{s}_i, t_i)$, $i = 1, 2, \cdots, m$, $\boldsymbol{s}_i \in \mathbb{R}^n$ and $t_i \in \mathbb{R}$, and assume that a linear relation of the form

$$t_i = \boldsymbol{s}_i^\top \boldsymbol{x}, \quad i = 1, 2, \cdots, m,$$

approximately holds. The objective is to find the parameters vector $\boldsymbol{x} \in \mathbb{R}^n$. The least-squares approach is to

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \sum_{i=1}^{m} (\boldsymbol{s}_i^\top \boldsymbol{x} - t_i)^2.$$

We can alternatively write the problem as

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \|S\boldsymbol{x} - \boldsymbol{t}\|^2,$$

where

$$S = \begin{bmatrix} \boldsymbol{s}_1^\top \\ \boldsymbol{s}_2^\top \\ \vdots \\ \boldsymbol{s}_m^\top \end{bmatrix}, \quad \boldsymbol{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_m \end{bmatrix}.$$

Example

Consider 30 points in $\mathbb{R}^2$, $x_i = (i-1)/29$, $y_i = 2x_i + 1 + \varepsilon_i$, for $i = 1, 2, \cdots, 30$, where ε_i is randomly generated from a standard normal distribution $\mathcal{N}(0, (0.1)^2)$. The objective is to find a line of the form $y = ax + b$ that best fits them. The corresponding linear system that needs to be "solved" is

$$\underbrace{\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \vdots \\ x_{30} & 1 \end{bmatrix}}_{\boldsymbol{X}} \begin{bmatrix} a \\ b \end{bmatrix} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{30} \end{bmatrix}}_{\boldsymbol{y}}.$$

The least squares solution is $(a, b)^\top = (\boldsymbol{X}^\top \boldsymbol{X})^{-1} \boldsymbol{X}^\top \boldsymbol{y}$.

```
randn('seed', 319);
d = linspace(0, 1, 30)';
e = 2*d + 1 + 0.1*randn(30, 1);
plot(d, e, '*')
```

Example (cont'd)

```
>> u = [d, ones(30, 1)]\e;
>> a = u(1), b = u(2)
a =
    2.0616
b =
    0.9725
```

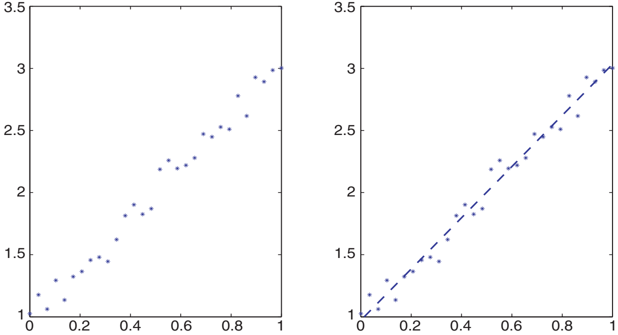

Nonlinear fitting

Suppose that we are given a set of points in $\mathbb{R}^2$, (u_i, y_i), $1 \le i \le m$, $u_i \ne u_j$ for $i \ne j$, and that we know *a priori* that these points are approximately related via a polynomial of degree at most d and $m \ge d+1$, i.e., $\exists\, a_0, a_1, \cdots, a_d$ such that

$$\sum_{j=0}^{d} a_j u_i^j \approx y_i, \quad i = 1, 2, \cdots, m.$$

The least-squares approach to this problem seeks $a_0, a_1, \cdots, a_d$ that are the least squares solution to the linear system

$$\underbrace{\begin{bmatrix} 1 & u_1 & u_1^2 & \cdots & u_1^d \\ 1 & u_2 & u_2^2 & \cdots & u_2^d \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & u_m & u_m^2 & \cdots & u_m^d \end{bmatrix}_{m\times(d+1)}}_{:=\boldsymbol{U}_{d+1}} \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_d \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix}.$$

The matrix $\boldsymbol{U}_{d+1}$ is of a full column rank since it consists of the first $d+1$ columns of the so-called $m \times m$ Vandermonde matrix which is nonsingular, $\det(\boldsymbol{U}_m) = \Pi_{1\le i<j\le m}(u_j - u_i) \ne 0$.

Regularized least-squares problem

The regularized least-squares (RLS) problem has the form

$$(RLS): \quad \min_{\boldsymbol{x}\in\mathbb{R}^n}\left\{\|\boldsymbol{A}\boldsymbol{x}-\boldsymbol{b}\|^2+\lambda R(\boldsymbol{x})\right\}.$$

The positive constant λ is the regularization parameter. In many cases, the regularization is taken to be quadratic. In particular, $R(\boldsymbol{x}) = \|\boldsymbol{D}\boldsymbol{x}\|^2$, where $\boldsymbol{D} \in \mathbb{R}^{p\times n}$ is a given matrix. Then we have

$$\min_{\boldsymbol{x}\in\mathbb{R}^n}\left\{f_{RLS}(\boldsymbol{x}) := \boldsymbol{x}^\top(\boldsymbol{A}^\top\boldsymbol{A}+\lambda\boldsymbol{D}^\top\boldsymbol{D})\boldsymbol{x} - 2(\boldsymbol{A}^\top\boldsymbol{b})^\top\boldsymbol{x} + \|\boldsymbol{b}\|^2\right\}.$$

Since the Hessian of the objective function is

$$\nabla^2 f_{RLS}(\boldsymbol{x}) = 2(\boldsymbol{A}^\top\boldsymbol{A}+\lambda\boldsymbol{D}^\top\boldsymbol{D}) \succeq \boldsymbol{0},$$

any stationary point is a global minimum point (cf. Theorem 2.38). The stationary points are those satisfying $\nabla f(\boldsymbol{x}) = \boldsymbol{0}$, that is

$$(\boldsymbol{A}^\top\boldsymbol{A}+\lambda\boldsymbol{D}^\top\boldsymbol{D})\boldsymbol{x} = \boldsymbol{A}^\top\boldsymbol{b}.$$

Therefore, if $\boldsymbol{A}^\top\boldsymbol{A}+\lambda\boldsymbol{D}^\top\boldsymbol{D} \succ \boldsymbol{0}$ then then the RLS solution is given by

$$\boldsymbol{x}_{RLS} = (\boldsymbol{A}^\top\boldsymbol{A}+\lambda\boldsymbol{D}^\top\boldsymbol{D})^{-1}\boldsymbol{A}^\top\boldsymbol{b}.$$

Example of regularized least-squares solution

Let $\boldsymbol{A} \in \mathbb{R}^{3\times 3}$ be given by

$$\boldsymbol{A} = \begin{bmatrix} 2+10^{-3} & 3 & 4 \\ 3 & 5+10^{-3} & 7 \\ 4 & 7 & 10+10^{-3} \end{bmatrix}.$$

```
B = [1, 1, 1; 1, 2, 3];
A=B'*B + 0.001*eye(3);    % cond(A) ≈ 16000 is rather large!
```

The "true" vector was chosen to be $\boldsymbol{x}_{true} = (1, 2, 3)^\top$, and $\boldsymbol{b}$ is a noisy measurement of $\boldsymbol{A}\boldsymbol{x}_{true}$:

```
>> x_true = [1; 2; 3];
>> randn('seed', 315);
>> b = A*x_true + 0.01*randn(3, 1)
b =
  20.0019
  34.0004
  48.0202
```

Example of regularized least-squares solution (cont'd)

The matrix $\boldsymbol{A}$ is in fact of a full column rank since its eigenvalues are all positive (`eig(A)`). The least-squares solution $\boldsymbol{x}_{LS}$ is given by

```
>> A\b
ans =
   4.5446
  -5.1295
   6.5742
```

Note that $\boldsymbol{x}_{LS}$ is rather far from the true vector $\boldsymbol{x}_{true}$. We will add the quadratic regularization function $\|\boldsymbol{I}\boldsymbol{x}\|^2$. The regularized solution is

$$\boldsymbol{x}_{RLS} = (\boldsymbol{A}^\top \boldsymbol{A} + \lambda \boldsymbol{I})^{-1} \boldsymbol{A}^\top \boldsymbol{b}. \qquad \text{(we take } \lambda = 1 \text{ below)}$$

```
>> x_rls = (A'*A + eye(3))\(A'*b)
x_rls =
  1.1763
  2.0318
  2.8872
```

which is a much better estimate for $\boldsymbol{x}_{true}$ than $\boldsymbol{x}_{LS}$.

Denoising

Suppose that a noisy measurement of a signal $\boldsymbol{x} \in \mathbb{R}^n$ is given

$$\boldsymbol{b} = \boldsymbol{x} + \boldsymbol{w},$$

where $\boldsymbol{x}$ is an unknown signal, $\boldsymbol{w}$ is an unknown noise vector, and $\boldsymbol{b}$ is the known measurement vector. The denoising problem is to find a "good" estimate of $\boldsymbol{x}$. The associated least-squares problem is

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \|\boldsymbol{x} - \boldsymbol{b}\|^2.$$

The optimal solution of this problem is obviously $\boldsymbol{x} = \boldsymbol{b}$, which is meaningless. We will add a regularization term $\lambda \sum_{i=1}^{n-1} (x_i - x_{i+1})^2$,

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \left\{ \|\boldsymbol{I}\boldsymbol{x} - \boldsymbol{b}\|^2 + \lambda \|\boldsymbol{L}\boldsymbol{x}\|^2 \right\},$$

where parameter $\lambda > 0$ and $\boldsymbol{L} \in \mathbb{R}^{(n-1) \times n}$ is given by

$$\boldsymbol{L} := \begin{bmatrix} 1 & -1 & & & \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \end{bmatrix}.$$

The optimal solution is given by $\boldsymbol{x}_{RLS}(\lambda) = (\boldsymbol{I} + \lambda \boldsymbol{L}^\top \boldsymbol{L})^{-1} \boldsymbol{b}$.

Example

Consider the signal $\boldsymbol{x} \in \mathbb{R}^{300}$ constructed by

```
t = linspace(0, 4, 300)';
x = sin(t) + t.*(cos(t).^2);
randn('seed', 314);
b = x + 0.05*randn(300, 1);
subplot(1, 2, 1);
plot(1:300, x, 'LineWidth', 2);
subplot(1, 2, 2);
plot(1:300, b, 'LineWidth', 2);
```

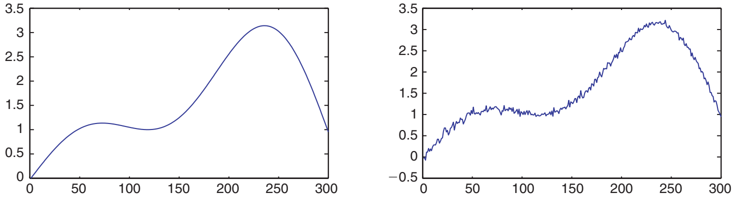

Figure 3.2. *A signal (left image) and its noisy version (right image).*

Example (cont'd): $\lambda = 1, 10, 100, 1000$

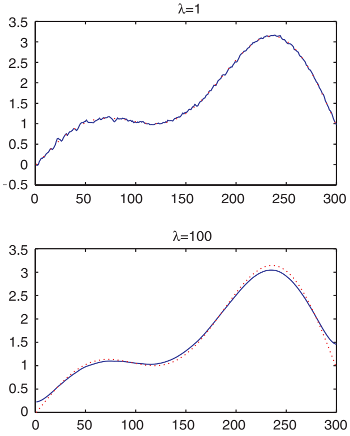

*signal **x**: marked with red dot*

Nonlinear least-squares problem

Suppose that we are given a system of nonlinear equations:

$$f_i(\boldsymbol{x}) \approx c_i, \quad i = 1, 2, \cdots, m.$$

The nonlinear least-squares (NLS) problem is formulated as

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \sum_{i=1}^{m} \left(f_i(\boldsymbol{x}) - c_i\right)^2.$$

The Gauss-Newton method is specifically devised to solve NLS problems of the form, but the method is not guaranteed to converge to the global optimal solution but rather to a stationary point (see §4.5).

Circle fitting

Suppose that we are given m points $\boldsymbol{a}_1, \boldsymbol{a}_2, \cdots, \boldsymbol{a}_m \in \mathbb{R}^n$. The circle fitting problem seeks to find a circle with center $\boldsymbol{x}$ and radius r,

$$C(\boldsymbol{x}, r) := \{\boldsymbol{y} \in \mathbb{R}^n : \|\boldsymbol{y} - \boldsymbol{x}\| = r\},$$

that best fits the m points. The nonlinear (approximate) equations associated with the problem are

$$\|\boldsymbol{x} - \boldsymbol{a}_i\| \approx r, \quad i = 1, 2, \cdots, m.$$

Since we wish to deal with differentiable functions, we will consider the squared version

$$\|\boldsymbol{x} - \boldsymbol{a}_i\|^2 \approx r^2, \quad i = 1, 2, \cdots, m.$$

The NLS problem associated with these equations is

$$\min_{x \in \mathbb{R}^n, r \geq 0} \sum_{i=1}^{m} \left(\|\boldsymbol{x} - \boldsymbol{a}_i\|^2 - r^2 \right)^2$$

Equivalent to a linear LS problem

The above NLS problem is the same as

$$\min\left\{\sum_{i=1}^{m}\left(-2\boldsymbol{a}_i^\top \boldsymbol{x}+\|\boldsymbol{x}\|^2-r^2+\|\boldsymbol{a}_i\|^2\right)^2 : \boldsymbol{x}\in\mathbb{R}^n, r\in\mathbb{R}\right\}.$$

Making the change of variables $R := \|\boldsymbol{x}\|^2 - r^2$, it reduces to

$$\min_{\boldsymbol{x}\in\mathbb{R}^n, R\in\mathbb{R}}\left\{f(\boldsymbol{x},R) := \sum_{i=1}^{m}\left(-2\boldsymbol{a}_i^\top \boldsymbol{x}+R+\|\boldsymbol{a}_i\|^2\right)^2 : \|\boldsymbol{x}\|^2 \geq R\right\}.$$

Indeed, any optimal solution $(\hat{\boldsymbol{x}}, \hat{R})$ automatically satisfies $\|\hat{\boldsymbol{x}}\|^2 \geq \hat{R}$, since otherwise, if $\|\hat{\boldsymbol{x}}\|^2 < \hat{R}$, we would have for $i = 1, 2, \cdots, m$,

$$-2\boldsymbol{a}_i^\top\hat{\boldsymbol{x}}+\hat{R}+\|\boldsymbol{a}_i\|^2 > -2\boldsymbol{a}_i^\top\hat{\boldsymbol{x}}+\|\hat{\boldsymbol{x}}\|^2+\|\boldsymbol{a}_i\|^2 = \|\hat{\boldsymbol{x}}-\boldsymbol{a}_i\|^2 \geq 0.$$

Squaring both sides and summing over i yield

$$\begin{aligned} f(\hat{\boldsymbol{x}},\hat{R}) &= \sum_{i=1}^{m}\left(-2\boldsymbol{a}_i^\top\hat{\boldsymbol{x}}+\hat{R}+\|\boldsymbol{a}_i\|^2\right)^2 \\ &> \sum_{i=1}^{m}\left(-2\boldsymbol{a}_i^\top\hat{\boldsymbol{x}}+\|\hat{\boldsymbol{x}}\|^2+\|\boldsymbol{a}_i\|^2\right)^2 = f(\hat{\boldsymbol{x}},\|\hat{\boldsymbol{x}}\|^2). \end{aligned}$$

This is a contradiction, since $(\hat{\boldsymbol{x}}, \hat{R})$ is an optimal solution.

Equivalent to a linear LS problem (cont'd)

Finally, we have the *linear* least-squares problem:

$$\min_{\boldsymbol{y}\in\mathbb{R}^{n+1}} \|\widetilde{A}\boldsymbol{y} - \boldsymbol{b}\|^2,$$

where $\boldsymbol{y} = (\boldsymbol{x}, R)^\top$ and

$$\widetilde{A} = \begin{bmatrix} 2a_1^\top & -1 \\ 2a_2^\top & -1 \\ \vdots & \vdots \\ 2a_m^\top & -1 \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} \|\boldsymbol{a}_1\|^2 \\ \|\boldsymbol{a}_2\|^2 \\ \vdots \\ \|\boldsymbol{a}_m\|^2 \end{bmatrix}.$$

If $\widetilde{\boldsymbol{A}}$ is of full column rank, then the unique solution is

$$\boldsymbol{y} = (\widetilde{\boldsymbol{A}}^\top \widetilde{A})^{-1} \widetilde{\boldsymbol{A}}^\top \boldsymbol{b},$$

and the radius r is given by $r = \sqrt{\|\boldsymbol{x}\|^2 - R}$.

Example: the best circle fitting of 10 points

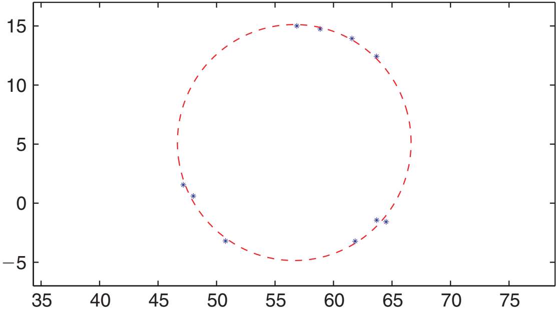

The best circle fitting of 10 points denoted by asterisks