

Mathematical Tripos Part IB: Lent Term 2024

Numerical Analysis – Lecture 16

16 Linear least squares

16.1 Statement of the problem

Consider the problem of finding a vector $\mathbf{c} \in \mathbb{R}^n$ such that $A\mathbf{c} = \mathbf{y}$ where a matrix $A \in \mathbb{R}^{m \times n}$ and a vector $\mathbf{y} \in \mathbb{R}^m$ are given and $m > n$. When there are more equations than unknowns, the system $A\mathbf{c} = \mathbf{y}$ is called *overdetermined*. In general, an overdetermined system has no solution, but we would like to have $A\mathbf{c}$ and $\mathbf{y}$ close in a sense. Choosing the Euclidean distance $\|\mathbf{z}\| = (\sum_{i=1}^m z_i^2)^{1/2}$ as a measure of closeness, we obtain the following problem.

Problem 16.1 (Least squares in $\mathbb{R}^m$) Given $A \in \mathbb{R}^{m \times n}$ and $\mathbf{y} \in \mathbb{R}^m$, find

$$\mathbf{c}^* = \arg \min_{\mathbf{c} \in \mathbb{R}^n} \|A\mathbf{c} - \mathbf{y}\|,$$

i.e., find (the argument) $\mathbf{c}^* \in \mathbb{R}^n$ which minimizes (the functional) $\|A\mathbf{c} - \mathbf{y}\|$.

Discussion 16.2 Problem of this form occur frequently when we collect m observations (x_i, y_i) (which, typically, are prone to measurement error) and wish to exploit them to form an n -variable linear model (given a-priori by some reasons)

$$f(x_i) \approx y_i, \quad i = \overline{1, m}, \quad f(x) = c_1\phi_1(x) + c_2\phi_2(x) + \cdots + c_n\phi_n(x),$$


(e.g., trying to put some planet observations on, or near, an ellipse). That is, with f being a linear combination of ϕ_j 's (as above), we want to determine particular values $c_j = c_j^*$ such that $f(x_i)$ would fit (y_i) in a certain way:

$$A\mathbf{c} = \begin{bmatrix} \phi_1(x_1) & \cdots & \phi_n(x_1) \\ \vdots & & \vdots \\ \phi_1(x_n) & \cdots & \phi_n(x_n) \\ \vdots & & \vdots \\ \phi_1(x_m) & \cdots & \phi_n(x_m) \end{bmatrix} \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix} = \begin{bmatrix} f(x_1) \\ \vdots \\ f(x_n) \\ \vdots \\ f(x_m) \end{bmatrix} \approx \begin{bmatrix} y_1 \\ \vdots \\ y_n \\ \vdots \\ y_m \end{bmatrix} = \mathbf{y}.$$

Of course, there are many ways of doing this, but it is customary to determine c_i^* 's so as to minimize the sum of squares of the deviation, i.e., to minimize

$$\sum_{i=1}^m [f(x_i) - y_i]^2 = \|A\mathbf{c} - \mathbf{y}\|^2,$$

for this leads to a *linear* system of equations for determination of the unknowns (c_j^*) .



x_i	y_i
1	0.00
2	0.60
3	1.77
4	1.92
5	3.31
6	3.52
7	4.59
8	5.31
9	5.79
10	7.06
11	7.17

Fig 16.1 Least squares straight line data fitting

16.2 Normal equations

Theorem 16.3 The vector $\mathbf{c}^* \in \mathbb{R}^n$ is a solution to the least squares $\Leftrightarrow A^T(A\mathbf{c}^* - \mathbf{y}) = 0$.

Proof. If $\mathbf{c}^*$ is a solution then it minimizes the quadratic form

$$F(\mathbf{c}) := \|A\mathbf{c} - \mathbf{y}\|^2 = (A\mathbf{c} - \mathbf{y}, A\mathbf{c} - \mathbf{y}) = \mathbf{c}^T A^T A \mathbf{c} - 2\mathbf{c}^T A^T \mathbf{y} + \mathbf{y}^T \mathbf{y},$$

hence the gradient $\nabla F = [\frac{\partial F}{\partial c_1}, \dots, \frac{\partial F}{\partial c_n}]^T$ must vanish at $\mathbf{c} = \mathbf{c}^*$. So,

$$\frac{1}{2} \nabla F(\mathbf{c}) = A^T A \mathbf{c} - A^T \mathbf{y}, \quad \text{hence} \quad A^T(A\mathbf{c}^* - \mathbf{y}) = 0.$$

Conversely, if $A^T(A\mathbf{c}^* - \mathbf{y}) = 0$, then, for any $\mathbf{c}' \in \mathbb{R}^n$, with $\mathbf{c} := \mathbf{c}' - \mathbf{c}^*$, we find that

$$\begin{aligned} \|A\mathbf{c}' - \mathbf{y}\|^2 &= (A\mathbf{c} + [A\mathbf{c}^* - \mathbf{y}], A\mathbf{c} + [A\mathbf{c}^* - \mathbf{y}]) \\ &= (A\mathbf{c}^* - \mathbf{y}, A\mathbf{c}^* - \mathbf{y}) + 2(A^T[A\mathbf{c}^* - \mathbf{y}], \mathbf{c}) + (A\mathbf{c}, A\mathbf{c}) \\ &= \|A\mathbf{c}^* - \mathbf{y}\|^2 + \|A\mathbf{c}\|^2 \geq \|A\mathbf{c}^* - \mathbf{y}\|^2, \end{aligned}$$

i.e., $\mathbf{c}^*$ is the least squares solution. Moreover, if $A \in \mathbb{R}^{m \times n}$ has full rank n and $\mathbf{c} = \mathbf{c}' - \mathbf{c}^* \neq 0$, then $A\mathbf{c} \neq 0$, hence the last inequality is strict, thus $\mathbf{c}^*$ is unique. $\square$

Corollary 16.4 Optimality of $\mathbf{c}^* \Leftrightarrow$ the vector $A\mathbf{c}^* - \mathbf{y}$ is orthogonal to all columns of A .

Remark 16.5 This corollary has a clear geometrical visualization. If we denote the columns of A as $(\mathbf{a}_j)_{j=1}^n$, then the least squares problem is the problem of finding the value $\min_{\mathbf{c}} \|\sum_{j=1}^n c_j \mathbf{a}_j - \mathbf{y}\|$, which is the minimum of the Euclidean distance between a given vector $\mathbf{y}$ and vectors in the plane $\mathcal{A} := \text{span}(\mathbf{a}_j)$. Geometrically it is clear that this minimum is attained when $\mathbf{a}^* = \sum_{j=1}^n c_j^* \mathbf{a}_j = A\mathbf{c}^*$ is the foot of the perpendicular from $\mathbf{y}$ onto the plane $\mathcal{A}$, i.e. when $\mathbf{y} - \sum c_j \mathbf{a}_j = \mathbf{y} - A\mathbf{c}^*$ is orthogonal to all vectors $\mathbf{a}_j$ (columns of A).

Definition 16.6 Thus, we may find optimal $\mathbf{c}^*$ by solving the $n \times n$ system

$$A^T A \mathbf{c}^* = A^T \mathbf{y}.$$

This system is called the *normal equations*, $A^T A$ is the *Gram matrix*, $\mathbf{c}^*$ is the *normal solution*.

Example 16.7 The least squares approximation to the data plotted in Fig. 16.1 by a straight line

$$f(x) = c_1 + c_2 x \quad (\phi_1(x) = 1, \quad \phi_2(x) = x)$$

results in the problem $\|A\mathbf{c} - \mathbf{y}\|^2 \rightarrow \min$ which normal solution is given by

$$A = (\phi_j(x_i)) = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \\ \vdots & \vdots \\ 1 & 11 \end{bmatrix}, \quad \underbrace{\begin{bmatrix} 11 & 66 \\ 66 & 506 \end{bmatrix}}_{A^T A} \begin{bmatrix} c_1^* \\ c_2^* \end{bmatrix} = \underbrace{\begin{bmatrix} 41.04 \\ 328.05 \end{bmatrix}}_{A^T \mathbf{y}} \Rightarrow \begin{bmatrix} c_1^* \\ c_2^* \end{bmatrix} = \begin{bmatrix} -0.7314 \\ 0.7437 \end{bmatrix}.$$

Discussion 16.8 However, the approach based on the normal solution has three disadvantages. Firstly, $A^T A$ might be singular, secondly a sparse A might be replaced by a dense $A^T A$ and, finally, forming $A^T A$ might lead to loss of accuracy.

16.3 QR and least squares

An alternative to the normal equations is the QR factorization. Suppose that $A = QR$, with an orthogonal $Q \in \mathbb{R}^{m \times m}$ and an upper triangular $R \in \mathbb{R}^{m \times n}$. Then, since orthogonal mapping is length-preserving, i.e., $\|Q^T x\|^2 = \|x\|^2$ for any $x \in \mathbb{R}^m$, we have

$$\|Ac - y\| = \|Q^T(Ac - y)\| = \|Rc - Q^T y\|$$

therefore we may seek $c \in \mathbb{R}^n$ that minimises $\|Rc - Q^T y\|$.

Suppose for simplicity that $\text{rank } R = \text{rank } A = n$. Then the bottom $m - n$ rows of R are zero, thus we cannot approach corresponding components of $Q^T y$ by Rc no matter what the c_i s are. So, we should try to match all others as best as we can, therefore we find c by solving the (nonsingular) linear system given by the first n equations of

$$Rc = Q^T y.$$

Note that we don't require Q explicitly and need to evaluate only $Q^T y$.

16.4 Examples

Example 16.9 We wish to find $c \in \mathbb{R}^3$ that minimizes $\|Ac - y\|$, where

$$A = \frac{1}{2} \begin{bmatrix} 1 & 3 & 6 \\ 1 & 1 & 2 \\ 1 & 3 & 4 \\ 1 & 1 & 0 \end{bmatrix}, \quad y = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 0 \end{bmatrix}$$

We QR-factorize

$$A = \frac{1}{2} \begin{bmatrix} 1 & 3 & 6 \\ 1 & 1 & 2 \\ 1 & 3 & 4 \\ 1 & 1 & 0 \end{bmatrix} = \frac{1}{2} \underbrace{\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}}_Q \times \underbrace{\begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}}_R.$$

We thus need to solve the least-squares problem

$$\underbrace{\begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}}_R \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} \approx \frac{1}{2} \underbrace{\begin{bmatrix} 3 \\ 1 \\ 1 \\ -1 \end{bmatrix}}_{Q^T y}.$$

No matter what the c_i 's are, we cannot approach the fourth component on the right-hand side, but all others can be matched. So, we solve the 'top' square system

$$\begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 3 \\ 1 \\ 1 \end{bmatrix} \Rightarrow c^* = \frac{1}{2} \begin{bmatrix} 2 \\ -1 \\ 1 \end{bmatrix}, \quad \|Ac^* - y\| = \|Rc^* - Q^T y\| = \frac{1}{2}.$$

Example 16.10 Using the normal solution, find the least squares approximation to the data

$$\begin{array}{c|c} x_i & y_i \\ \hline -1 & 2 \\ 0 & 1 \\ 1 & 0 \end{array} \quad \text{by a function } f = c_1\phi_1 + c_2\phi_2, \text{ where } \phi_1(x) = x \text{ and } \phi_2(x) = 1 + x - x^2.$$

We solve the system of normal equations:

$$A = (\phi_j(x_i)) = \begin{bmatrix} -1 & -1 \\ 0 & 1 \\ 1 & 1 \end{bmatrix}, \quad \underbrace{\begin{bmatrix} 2 & 2 \\ 2 & 3 \end{bmatrix}}_{A^T A} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \underbrace{\begin{bmatrix} -2 \\ -1 \end{bmatrix}}_{A^T \mathbf{y}} \Rightarrow \mathbf{c}^* = \begin{bmatrix} -2 \\ 1 \end{bmatrix}.$$

The error is

$$A\mathbf{c}^* - \mathbf{y} = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix} - \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \\ -1 \end{bmatrix} \Rightarrow \|A\mathbf{c}^* - \mathbf{y}\| = \sqrt{2}.$$