

Numerical Analysis – Lecture 7

5 Linear least squares

5.1 Statement of the problem

Suppose that an $m \times n$ matrix A and a vector $\mathbf{b} \in \mathbb{R}^m$ are given. The equation $A\mathbf{x} = \mathbf{b}$ has in general no solution (if $m > n$) or an infinity of solutions (if $m < n$). ‘Equations’ of this form occur frequently when we collect m observations (which, typically, are prone to measurement error) and wish to exploit them to form an n -variable linear model, whereby $n \ll m$. (In statistics, this is called *linear regression*.) Bearing in mind the likely presence of errors in A and \mathbf{b} , we seek $\mathbf{x} \in \mathbb{R}^n$ that minimizes the Euclidean length $\|A\mathbf{x} - \mathbf{b}\|$ – the *least squares problem*.

Theorem $\mathbf{x} \in \mathbb{R}^n$ is a solution of the least squares problem iff $A^T(A\mathbf{x} - \mathbf{b}) = \mathbf{0}$.

Proof. If \mathbf{x} is a solution then it minimizes

$$f(\mathbf{x}) := \|A\mathbf{x} - \mathbf{b}\|^2 = \langle A\mathbf{x} - \mathbf{b}, A\mathbf{x} - \mathbf{b} \rangle = \mathbf{x}^T A^T A \mathbf{x} - 2\mathbf{x}^T A^T \mathbf{b} + \mathbf{b}^T \mathbf{b}.$$

Hence $\nabla f(\mathbf{x}) = \mathbf{0}$. But $\frac{1}{2}\nabla f(\mathbf{x}) = A^T A \mathbf{x} - A^T \mathbf{b}$, hence $A^T(A\mathbf{x} - \mathbf{b}) = \mathbf{0}$.

Conversely, suppose that $A^T(A\mathbf{x} - \mathbf{b}) = \mathbf{0}$ and let $\mathbf{u} \in \mathbb{R}^n$. Hence, letting $\mathbf{y} = \mathbf{u} - \mathbf{x}$,

$$\begin{aligned} \|A\mathbf{u} - \mathbf{b}\|^2 &= \langle A\mathbf{x} + A\mathbf{y} - \mathbf{b}, A\mathbf{x} + A\mathbf{y} - \mathbf{b} \rangle = \langle A\mathbf{x} - \mathbf{b}, A\mathbf{x} - \mathbf{b} \rangle + 2\mathbf{y}^T A^T(A\mathbf{x} - \mathbf{b}) \\ &\quad + \langle A\mathbf{y}, A\mathbf{y} \rangle = \|A\mathbf{x} - \mathbf{b}\|^2 + \|A\mathbf{y}\|^2 \geq \|A\mathbf{x} - \mathbf{b}\|^2 \end{aligned}$$

and \mathbf{x} is indeed optimal. □

Corollary Optimality of $\mathbf{x} \Leftrightarrow A\mathbf{x} - \mathbf{b}$ is orthogonal to all columns of A .

5.2 Normal equations

One way of finding optimal \mathbf{x} is by solving the $n \times n$ linear system $A^T A \mathbf{x} = A^T \mathbf{b}$ – the *normal equations*. This approach is popular in many applications. However, there are three disadvantages. Firstly, $A^T A$ might be singular, secondly sparse A might be replaced by a dense $A^T A$ and, finally, forming $A^T A$ might lead to loss of accuracy. Thus, suppose that our computer works in the IEEE arithmetic standard (≈ 15 significant digits) and let

$$A = \begin{bmatrix} 10^8 & -10^8 \\ 1 & 1 \end{bmatrix} \quad \Rightarrow \quad A^T A = \begin{bmatrix} 10^{16} + 1 & -10^{16} + 1 \\ -10^{16} + 1 & 10^{16} + 1 \end{bmatrix} \approx 10^{16} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Given $\mathbf{b} = [0, 2]^T$ the solution of $A\mathbf{x} = \mathbf{b}$ is $[1, 1]^T$, as can be easily found by Gaussian elimination. However, our computer ‘believes’ that $A^T A$ is singular!

5.3 QR and least squares

Lemma Let A be any $m \times n$ matrix and let $\mathbf{b} \in \mathbb{R}^m$. The vector $\mathbf{x} \in \mathbb{R}^n$ minimizes $\|A\mathbf{x} - \mathbf{b}\|$ iff it minimizes $\|\Omega A\mathbf{x} - \Omega \mathbf{b}\|$ for an arbitrary $m \times m$ orthogonal matrix Ω .

Proof. Follows at once from the identity

$$\|\Omega A\mathbf{x} - \Omega\mathbf{b}\|^2 = (\Omega A\mathbf{x} - \Omega\mathbf{b})^T(\Omega A\mathbf{x} - \Omega\mathbf{b}) = (A\mathbf{x} - \mathbf{b})^T\Omega^T\Omega(A\mathbf{x} - \mathbf{b}) = \|A\mathbf{x} - \mathbf{b}\|^2.$$

□

An irrelevant, yet important remark The property that orthogonal matrices leave the Euclidean distance intact is called *isometry* and it has many important ramifications throughout mathematics and mathematical physics.

Method of solution Suppose that $A = QR$, a QR factorization with R in a *standard form*. Because of the lemma, we let $\Omega := Q^T$, hence seek \mathbf{x} that minimizes $\|R\mathbf{x} - Q^T\mathbf{b}\|$. In general ($m > n$) many rows of R consist of zeros. Suppose for simplicity that $\text{rank } R = \text{rank } A = n$. Then the bottom $m - n$ rows of R are zero. Therefore we find \mathbf{x} by solving the (nonsingular) linear system given by the first n equations of $R\mathbf{x} = Q^T\mathbf{b}$. Similar (although more complicated) algorithm applies when $\text{rank } R \leq n - 1$. Note, recalling our former remark, that we don't require Q explicitly and need to evaluate only $Q^T\mathbf{b}$.

6 Polynomial interpolation

6.1 The interpolation problem

Given $n + 1$ distinct real points x_0, x_1, \dots, x_n and real numbers f_0, f_1, \dots, f_n , we seek a function $p: \mathbb{R} \rightarrow \mathbb{R}$ such that $p(x_i) = f_i$, $i = 0, 1, \dots, n$. Such a function is called an *interpolant*.

We denote by $\mathbb{P}_n[x]$ the set of all real polynomials of degree at most n and observe that each $p \in \mathbb{P}_n[x]$ is uniquely defined by its $n + 1$ coefficients. This, intuitively, justifies seeking an interpolant from $\mathbb{P}_n[x]$.

6.2 The Lagrange formula

Although, in principle, we may solve a linear problem with $n + 1$ unknowns to determine a polynomial interpolant, this can be accomplished more easily by using the explicit *Lagrange formula*. We claim that

$$p(x) = \sum_{k=0}^n f_k \prod_{\substack{\ell=0 \\ \ell \neq k}}^n \frac{x - x_\ell}{x_k - x_\ell}, \quad x \in \mathbb{R}.$$

Note that $p \in \mathbb{P}_n[x]$, as required. We wish to show that it interpolates the data. Define $L_j(x) := \prod_{\ell=0, \ell \neq j}^n (x - x_\ell)/(x_j - x_\ell)$, $j = 0, 1, \dots, n$ (*Lagrange cardinal polynomials*). It is trivial to verify that $L_j(x_j) = 1$ and $L_j(x_k) = 0$ for $k \neq j$, hence $p(x_j) = \sum_{k=0}^n f_k L_k(x_j) = f_j$, $j = 0, 1, \dots, n$.

Uniqueness Suppose that both $p \in \mathbb{P}_n[x]$ and $q \in \mathbb{P}_n[x]$ interpolate to the same $n + 1$ data. Then the n th degree polynomial $p - q$ vanishes at $n + 1$ distinct points. Hence it is identically zero, $p \equiv q$ and the interpolating polynomial is unique.