

Lecture Note 8: Linear Least Squares Problem

Xianyi Zeng

Department of Mathematical Sciences, UTEP

1 From Linear System to Least Squares

In previous sections we solve the linear system $A\mathbf{x} = \mathbf{b}$ when A is square and non-singular. In the more general case, the problem is not mathematically well-posed. Let A be any $n \times n$ matrix, but $\det A = 0$, then the system $A\mathbf{x} = \mathbf{b}$:

- Has no solution if $\mathbf{b} \notin \text{col}(A)$.
- Has infinite number of solutions if $\mathbf{b} \in \text{col}(A)$.

Here $\text{col}(A)$ is the column space of A :

$$\text{col}(A) = \text{span}(\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n), \quad (1.1)$$

where $\mathbf{a}_i, i = 1, \dots, n$ are the column vectors of A : $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_n]$.

In a more general case let $A \in \mathbb{R}^{m \times n}$ be any matrix, then $A\mathbf{x} = \mathbf{b}$ for $\mathbf{b} \in \mathbb{R}^m$ has at least one solution $\mathbf{x} \in \mathbb{R}^n$ if and only if $\mathbf{b} \in \text{col}(A)$. In particular, if $\mathbf{b} \in \text{col}(A)$:

- The solution is unique if and only if $n \leq m$ and $\text{rank}(A) = n$, or equivalently $\dim(\text{col}(A)) = n$.
- There are infinite number of solutions in all other situations.

We'll briefly prove the first statement:

Proof. If $n \leq m$ and $\text{rank}(A) = n$, then $A^t A \in \mathbb{R}^{n \times n}$ is non-singular (see Appendix A). Let \mathbf{x} be any vector that satisfies $A\mathbf{x} = \mathbf{b}$, then:

$$(A^t A)\mathbf{x} = A^t(A\mathbf{x}) = A^t \mathbf{b},$$

or equivalently:

$$\mathbf{x} = (A^t A)^{-1} A^t \mathbf{b}, \quad (1.2)$$

which demonstrates the uniqueness.

Conversely, we suppose $A\mathbf{x} = \mathbf{b}$ has a unique solution $\mathbf{x}_0 \in \mathbb{R}^n$. Let $\mathbf{y} \in \ker(A)$, the null space of A , then $\mathbf{x}_0 + \mathbf{y}$ is also a solution:

$$A(\mathbf{x}_0 + \mathbf{y}) = A\mathbf{x}_0 + A\mathbf{y} = \mathbf{b} + 0 = \mathbf{b}.$$

Due to the uniqueness of \mathbf{x}_0 , we know that $\ker(A)$ only contains the zero vector, i.e., $\dim(\ker(A)) = 0$. But from the *dimension theorem* of linear algebra we know that $\dim(\ker(A)) = n - \dim(\text{col}(A))$, thus:

$$n - \dim(\text{col}(A)) = 0 \quad \Rightarrow \quad \dim(\text{col}(A)) = n,$$

and consequently A has full rank n and $m \geq n$. □

In practice, however, we do not always have a well-posed system to solve (see the example in the next section). To this end, people define instead the next least squares problem for any $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{Ax} - \mathbf{b}\|_2. \quad (1.3)$$

Note that this problem always has at least one solution¹; but the solution may not be unique. Particularly, if \mathbf{x}_0 minimize the residual then $\mathbf{x}_0 + \mathbf{y}$ also minimize the residual for all $\mathbf{y} \in \ker(A)$. Thus the least squares problem (1.3) is always *solvable*, but not necessarily well-posed. A well-posed extension is presented at the end of this lecture.

Theorem 1.1. *\mathbf{x} solves the least squares problem (1.3) if and only if it solves the normal equation:*

$$A^t \mathbf{Ax} = A^t \mathbf{b}. \quad (1.4)$$

This equation always has a solution, and the solution is unique if and only if the columns of A are linearly independent (i.e., $\dim(\text{col}(A)) = \text{rank}(A) = n \leq m$).

Proof. The first statement can be obtained as a consequence of the stationary condition for minimizing the convex quadratic form $\phi(\mathbf{x}) = (\mathbf{Ax} - \mathbf{b})^t(\mathbf{Ax} - \mathbf{b})$, that is, $\nabla \phi(\mathbf{x}) = 0$. Here we instead use a more direct approach. Let \mathbf{x} solve the normal equation and let $\mathbf{e} \in \mathbb{R}^n$ be arbitrary, then:

$$\begin{aligned} \|A(\mathbf{x} + \mathbf{e}) - \mathbf{b}\|^2 &= (\mathbf{Ax} - \mathbf{b} - A\mathbf{e})^t(\mathbf{Ax} - \mathbf{b} - A\mathbf{e}) = \|\mathbf{Ax} - \mathbf{b}\|^2 + \|A\mathbf{e}\|^2 - 2\mathbf{e}^t(A^t \mathbf{Ax} - A^t \mathbf{b}) \\ &= \|\mathbf{Ax} - \mathbf{b}\|^2 + \|A\mathbf{e}\|^2 \geq \|\mathbf{Ax} - \mathbf{b}\|^2, \end{aligned}$$

and due to the arbitrary choice of \mathbf{e} , the same \mathbf{x} solves the least squares problem. Conversely if \mathbf{x} solves (1.3) we define $\mathbf{d} = A^t \mathbf{Ax} - A^t \mathbf{b}$, then for all $\mathbf{e} \in \mathbb{R}^n$ we have:

$$0 \leq \|A(\mathbf{x} + \mathbf{e}) - \mathbf{b}\|^2 - \|\mathbf{Ax} - \mathbf{b}\|^2 = \|A\mathbf{e}\|^2 - 2\mathbf{e}^t \mathbf{d}.$$

With the particular choice $\mathbf{e} = -\alpha \mathbf{d}$, we have for all $\alpha > 0$:

$$0 \leq \alpha^2 \|A\mathbf{d}\|^2 - 2\alpha \|\mathbf{d}\|^2 \Rightarrow 0 \leq \alpha \|A\mathbf{d}\|^2 - 2\|\mathbf{d}\|^2.$$

Letting $\alpha \rightarrow 0$, we have $2\|\mathbf{d}\|^2 \leq 0$ or $\mathbf{d} = 0$, i.e., \mathbf{x} solves the normal equation.

For the second part, clearly we just need to show that the normal equation always has a solution. Indeed, due to Appendix A the column spaces of A^t and $A^t A$ are the identical. Hence for all $\mathbf{b} \in \mathbb{R}^m$, we have:

$$A^t \mathbf{b} \in \text{col}(A^t) = \text{col}(A^t A),$$

or there exists an $\mathbf{x} \in \mathbb{R}^n$ such that $A^t \mathbf{b} = A^t \mathbf{Ax}$. □

To this end, the problem of solving the least squares problem is equivalent to solving the normal equation, which we will discuss in more detail in the next lectures.

¹An elegant proof can be obtained by the following result of real analysis: Any continuous function defined on a compact set $K \subset \mathbb{R}^n$ achieves its minimum and maximum on the set K .

2 Example: Polynomial Regression

An important application of the least squares problem is to find a polynomial fit of scattered data. Let our data set be:

$$\mathcal{D} = \{(x_i, y_i) : x_i, y_i \in \mathbb{R}, i = 1, \dots, n\}. \quad (2.1)$$

The target is to find a function $f: \mathbb{R} \mapsto \mathbb{R}$, such that:

$$y_i \approx f(x_i), \quad \forall 1 \leq i \leq n.$$

In polynomial regression, we search for a polynomial of degree m :

$$p_m(x) = a_m x^m + a_{m-1} x^{m-1} + \dots + a_1 x + a_0, \quad (2.2)$$

so that the L^2 -norm of the difference vector $\mathbf{z} = [z_i] \in \mathbb{R}^n$, $z_i = y_i - p_m(x_i)$ is minimized. Note that:

$$z_i = y_i - p_m(x_i) = y_i - \sum_{k=0}^m a_k x_i^k, \quad 1 \leq i \leq n,$$

we may write \mathbf{z} in the matrix form:

$$\mathbf{z} = \mathbf{y} - V_m(\mathbf{x})\mathbf{a}, \quad (2.3)$$

where $\mathbf{a} = [a_k] \in \mathbb{R}^{m+1}$ and $V_m(\mathbf{x}) \in \mathbb{R}^{n \times (m+1)}$ is the Vandermonde matrix:

$$V_m(\mathbf{x}) = \begin{bmatrix} 1 & x_1 & x_1^2 & \cdots & x_1^m \\ 1 & x_2 & x_2^2 & \cdots & x_2^m \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & x_n^2 & \cdots & x_n^m \end{bmatrix}_{n \times (m+1)}. \quad (2.4)$$

Then the problem of polynomial fitting, i.e., finding the coefficient vector \mathbf{a} , reduces to the least squares problem:

$$\mathbf{a} = \arg \min_{\mathbf{a}' \in \mathbb{R}^{m+1}} \|\mathbf{V}_m(\mathbf{x})\mathbf{a}' - \mathbf{y}\|. \quad (2.5)$$

The existence of a solution (and hence p_m) is shown by Theorem 1.1, furthermore, if $x_i \neq x_j$ for all $i \neq j$ and $n \geq m+1$, the solution is unique. Due to the same theorem, we just need to show that if all x_i 's are different from each other, $V_m(\mathbf{x})$ has full rank; the latter is the consequence of a special case – when $m+1 = n$ and all x_i 's are different, $V_m(\mathbf{x})$ is non-singular. In fact, one can use induction to show that if $m+1 = n$, then:

$$\det V_m(\mathbf{x}) = \prod_{1 \leq i < j \leq n} (x_j - x_i) \neq 0. \quad (2.6)$$

Hence we deduce that when all the data points x_i are different, there is always a unique fit using the polynomial of any degree m .

3 Pseudoinverse

Lastly, we briefly discuss a well-posed problem for general matrix $A \in \mathbb{R}^{m \times n}$:

$$\mathbf{x} = \arg \min_{\mathbf{x}' \in \mathcal{S}} \|\mathbf{x}'\|, \quad \mathcal{S} = \{\mathbf{x}' \in \mathbb{R}^n : \|A\mathbf{x}' - \mathbf{b}\| = \min_{\mathbf{y} \in \mathbb{R}^n} \|A\mathbf{y} - \mathbf{b}\|\}. \quad (3.1)$$

That is to say, we want to find the least squares solution with the smallest L^2 -norm. Its existence is obtained similarly to the footnote before Theorem 1.1. To show the uniqueness, we need a lemma: if $\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{R}^n$ such that $\|\mathbf{y}_1\| = \|\mathbf{y}_2\|$, then $\|(\mathbf{y}_1 + \mathbf{y}_2)/2\| \leq \|\mathbf{y}_1\|$; and the identity holds only if $\mathbf{y}_1 = \mathbf{y}_2$. This is a direct consequence of the parallelogram identity:

$$2\|\mathbf{y}_1\|^2 + 2\|\mathbf{y}_2\|^2 = \|\mathbf{y}_1 + \mathbf{y}_2\|^2 + \|\mathbf{y}_1 - \mathbf{y}_2\|^2. \quad (3.2)$$

Now if both \mathbf{x}_1 and \mathbf{x}_2 are solutions to (3.1), we must have $\|\mathbf{x}_1\| = \|\mathbf{x}_2\|$. Furthermore, $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{S}$ indicates

$$\|A\mathbf{x}_1 - \mathbf{b}\| = \|A\mathbf{x}_2 - \mathbf{b}\|.$$

By the previous lemma

$$\|A((\mathbf{x}_1 + \mathbf{x}_2)/2) - \mathbf{b}\| \leq \|A\mathbf{x}_1 - \mathbf{b}\| \Rightarrow \frac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2) \in \mathcal{S}.$$

Invoking the lemma again, we have:

$$\|(\mathbf{x}_1 + \mathbf{x}_2)/2\| \leq \|\mathbf{x}_1\|,$$

and by the assumption we must have the identity hold, i.e., $\mathbf{x}_1 = \mathbf{x}_2$. We thusly conclude that the solution to (3.1) is unique.

Now we deviate ourselves from the numerical techniques and look at the analytical solution. Let $A = U\Sigma V$ be the singular value decomposition (SVD) of A , where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal matrices, $\Sigma \in \mathbb{R}^{m \times n}$ is diagonal in the sense that only the (i, i) -th components of Σ are (possibly) nonzero, where $1 \leq i \leq \min(m, n)$. In addition, denoting the (i, i) -th component of Σ by σ_i we have:

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k \geq 0, \quad \text{where } k = \min(m, n). \quad (3.3)$$

Using the SVD of A and noting the fact that pre-multiplying by an orthogonal matrix preserves the L^2 -norm, (3.1) is equivalent to:

$$\mathbf{y} = \arg \min_{\mathbf{y}' \in \mathcal{S}'} \|\mathbf{y}'\|, \quad \mathcal{S}' = \{\mathbf{y}' \in \mathbb{R}^n : \|\Sigma\mathbf{y}' - U^t\mathbf{b}\| = \min_{\mathbf{z} \in \mathbb{R}^n} \|\Sigma\mathbf{z} - U^t\mathbf{b}\|\},$$

where the solution \mathbf{y} is related to the solution \mathbf{x} of (3.1) by $\mathbf{y} = V\mathbf{x}$. The members of \mathcal{S}' have a surprisingly simple structure as the m equations are decoupled from each other. In particular, $\mathbf{y}' = [y'_i] \in \mathcal{S}'$ is given by:

$$\begin{cases} y'_i = \frac{1}{\sigma_i} [U^t\mathbf{b}]_i, & \text{if } \sigma_i \neq 0; \\ y'_i \text{ is arbitrary,} & \text{if } i > k \text{ or } \sigma_i = 0. \end{cases}$$

Clearly the solution is given by setting all y'_i in the second line to be zero. The final result can be written as:

$$\mathbf{x} = V^t \mathbf{y} = V^t \Sigma^+ U^t \mathbf{b},$$

where $\Sigma^+ \in \mathbb{R}^{n \times m}$ has a similar structure as Σ^t , where every nonzero σ_i is replaced by $1/\sigma_i$. $A^+ \stackrel{\text{def}}{=} V^t \Sigma^+ U^t$ is known as the pseudoinverse of the matrix A .

Exercises

Exercise 1. Show that if we want to fit the data set (2.1) such that $x_i \neq x_j, \forall i \neq j$ by a constant function, that is, a polynomial of degree zero, this function is given by the arithmetic average of all y_i 's.

Exercise 2. Verify (2.6) for the case $m+1=n=4$. That is, show that:

$$\det \begin{bmatrix} 1 & x_1 & x_1^2 & x_1^3 \\ 1 & x_2 & x_2^2 & x_2^3 \\ 1 & x_3 & x_3^2 & x_3^3 \\ 1 & x_4 & x_4^2 & x_4^3 \end{bmatrix} = (x_4 - x_3)(x_4 - x_2)(x_4 - x_1)(x_3 - x_2)(x_3 - x_1)(x_2 - x_1).$$

Exercise 3. Compute the pseudoinverse of the following 3×2 matrix A :

$$A = \begin{bmatrix} 2\sqrt{2} & 2\sqrt{2} \\ -1 & 1 \\ 1 & -1 \end{bmatrix}.$$

You may use any software to compute the SVD decomposition of A .

A Show that $\text{col}(A^t A) = \text{col}(A)$ for any $A \in \mathbb{R}^{m \times n}$

Let $\mathcal{A}_1 = \text{col}(A)$ and $\mathcal{A}_2 = \text{col}(A^t A)$, then we have:

$$\mathcal{A}_2 = \{A^t A \mathbf{x} : \mathbf{x} \in \mathbb{R}^n\} \subseteq \{A^t \mathbf{y} : \mathbf{y} \in \mathbb{R}^m\} = \mathcal{A}_1.$$

Thus $\mathcal{A}_1^\perp \subseteq \mathcal{A}_2^\perp$. In fact, by the definition of orthogonal complement. Indeed:

$$\mathbf{u} \in \mathcal{A}_1^\perp \Rightarrow \mathbf{u} \cdot \mathbf{v} = 0, \forall \mathbf{v} \in \mathcal{A}_1 \Rightarrow \mathbf{u} \cdot \mathbf{v} = 0, \forall \mathbf{v} \in \mathcal{A}_2 \Rightarrow \mathbf{u} \in \mathcal{A}_2^\perp.$$

However, we can also show $\mathcal{A}_2^\perp \subseteq \mathcal{A}_1^\perp$ as follows:

$$\mathbf{u} \in \mathcal{A}_2^\perp \Rightarrow \mathbf{u}^t A^t A \mathbf{u} = 0 \Rightarrow \|\mathbf{A} \mathbf{u}\|^2 = \mathbf{u}^t A^t A \mathbf{u} = 0 \Rightarrow \mathbf{u}^t A^t = 0 \Rightarrow \mathbf{u} \in \mathcal{A}_1^\perp.$$

Thus $\mathcal{A}_1^\perp = \mathcal{A}_2^\perp$ and $\mathcal{A}_1 = (\mathcal{A}_1^\perp)^\perp = (\mathcal{A}_2^\perp)^\perp = \mathcal{A}_2$.

As a corollary, if A has full column rank (i.e., $\text{rank}(A) = n$) then $A^t A$ is non-singular.

²Check that for all subspace \mathcal{V} of \mathbb{R}^n , $(\mathcal{V}^\perp)^\perp = \mathcal{V}$.