Singular value decomposition & least squares approximation

Consider the problem of fitting a model

$$m(x) = \alpha_1 \phi_1(x) + \cdots + \alpha_n \phi_n(x)$$

where $\phi_i(x)$ are the basis functions, to the data set $(x_1, y_1), (x_2, y_2) \ldots (x_m, y_m)$ where $m > n$. Notice that in general there is no solution to the linear system

$$A\alpha = y$$

where $A \in \mathbb{R}^{m \times n}$, $A_{ij} = \phi_j(x_i)$ $\begin{matrix} i=1:m \\ j=1:n \end{matrix}$

$$\alpha = \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} \in \mathbb{R}^n, \quad y = \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} \in \mathbb{R}^m$$

Least squares method:

Find $\alpha^* \in \mathbb{R}^n$: $\|A\alpha^* - y\|_2 = \min_{\alpha \in \mathbb{R}^n} \|A\alpha - y\|_2$

in general, given a matrix $A \in \mathbb{R}^{m \times n}$ and a data vector $y \in \mathbb{R}^m$, the least squares method provides an approximate solution to $Ax = b$ as the solution to the minimization problem

$$(LS_1) \qquad \min_{x \in \mathbb{R}^n} \|Ax - b\| \longrightarrow \text{Euclidean norm in } \mathbb{R}^n$$

An equivalent formulation to (LS_1) is

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

<u>Theorem</u> The vector $x^* \in \mathbb{R}^n$ is solution to (LS) if and only if

$$(*) \qquad A^T A x^* = A^T b$$

The linear system $(*)$ is called the <u>normal equations</u>

Proof Consider $f: \mathbb{R}^n \to \mathbb{R}$ defined as

$$f(x) = \|Ax-b\|^2 = (Ax-b,\ Ax-b)$$

$$= x^T A^T A x - 2(A^T b, x) + \|b\|^2 .$$

The gradient of f is

$$\nabla_x f(x) = 2(A^T A x - A^T b)$$

Thus any critical point to f

$$\nabla_x f(x) = 0$$

satisfies the normal equations $(*)$.

On another hand, if x^* is solution to the system $(*)$ then for any $x \in \mathbb{R}^n$,

$$f(x) = \|A(x-x^*) + Ax^* - b\|^2 =$$

$$= \|A(x-x^*)\|^2 + \underbrace{2(A(x-x^*),\ Ax^*-b)}_{=0}$$

$$+ \|Ax^*-b\|^2$$

$$= \|A(x-x^*)\|^2 + \underbrace{\|Ax^*-b\|^2}_{f(x^*)} \geq f(x^*)$$

Notation For a matrix $A \in \mathbb{R}^{m \times n}$ we define the null space of A as

$$\mathcal{N}(A) = \{ x \in \mathbb{R}^n : Ax = 0 \}$$

The range of A is defined as

$$\mathcal{R}(A) = \{ y \in \mathbb{R}^m : \text{there is } x \in \mathbb{R}^n : Ax = y \}$$

From linear algebra we have the following

Theorem: $\mathcal{R}(A^T) = \mathcal{N}(A)^{\perp}$

$$\mathcal{N}(A) = \mathcal{N}(A^T A)$$

and therefore, $\mathcal{R}(A^T A) = \mathcal{R}(A^T)$

Theorem (existence and uniqueness of the least squares solution)

For any $A \in \mathbb{R}^{m \times n}$ and any $b \in \mathbb{R}^m$ there is a solution to $A^T A x^* = A^T b$

The solution is unique if and only if $\mathcal{N}(A) = \{0\}$.

In particular, if $m > n$ then the solution is unique if and only if A has full column rank, $\underline{\text{rank}(A) = n}$. Henceforth, we assume $m > n$ and $\text{rank}(A) = n$. Then the unique solution to the (LS) problem is expressed as

$$\boxed{x^* = (A^T A)^{-1} A^T b}$$

The matrix $(A^T A)^{-1} A^T$ is known as the $\underline{\text{Moore-Penrose inverse of } A}$

The solution x^* may be easily evaluated if the SVD of the matrix A is available,

$$A = U S V^T .$$

$$A^TA = \left(\sum_{i=1}^{n} \sigma_i u_i v_i^T\right)^T \left(\sum_{i=1}^{n} \sigma_i u_i v_i^T\right)$$

$$= \sum_{i=1}^{n} \sigma_i^2 v_i v_i^T$$

and $(A^TA)^{-1} = \sum_{i=1}^{n} \frac{1}{\sigma_i^2} v_i v_i^T$

such that $(A^TA)^{-1}A^T = \sum_{i=1}^{n} \frac{1}{\sigma_i} v_i u_i^T$

and $x^* = (A^TA)^{-1}A^Tb$ is obtained as

$$\boxed{x^* = \sum_{i=1}^{n} \frac{u_i^T \cdot b}{\sigma_i} v_i}$$

Notice that this expression is the same as the one obtained in the case of a square matrix A and nonsingular.

in the presence of errors in data

$$\tilde{b} = b + \xi$$

the error in the (LS) solution is

$$\tilde{x}^* - x^* = \sum_{i=1}^{n} \frac{u_i^T \cdot \xi}{\sigma_i} v_i$$

such that errors in the data space along the u_i direction are amplified by small singular values $\sigma_i \ll 1$ in the v_i-direction of the state space

in practice, the discrete Picard condition is satisfied and allows filtering of noise components associated with small singular values,

$$\frac{|u_i^T \cdot b|}{\sigma_i} \ll 1 \quad \text{for high index } i$$

Regularization methods produce solutions of the form

$$x_{reg} = \sum_{i=1}^{n} f_i \frac{u_i^T \cdot b}{\sigma_i} v_i$$

where the numbers f_i are filter factors associated with a particular regularization method. The characteristic property of the filter factors is that $f_i \to 0$ for high index i such that the solution components associated with smaller σ_i, $f_i \frac{u_i^T \cdot b}{\sigma_i}$ are filtered out.

In particular, the truncated SVD (TSVD) provide a solution as

$$\boxed{x_k = \sum_{i=1}^{k} \frac{u_i^T \cdot b}{\sigma_i} v_i} \quad (TSVD)$$

The truncation parameter K should be properly selected such that all the noise-dominated SVS coefficients are eliminated. For practical applications, it is also assumed that the errors in data (noise) are of smaller magnitude than the data., thus

$$u_i^T \cdot \tilde{b} = u_i^T \cdot b + u_i^T \cdot \xi$$

satisfies $u_i^T \cdot \tilde{b} \approx u_i^T \cdot b$ for larger singular values $(\lambda_1 \cdots \lambda_K)$

and $u_i^T \cdot \tilde{b} \approx u_i^T \cdot \xi$ for smaller singular values.

The Picard plot of $\frac{u_i^T \cdot b}{\sigma_i}$, $i = 1:n$ provides insight on the selection of the truncation parameter K.