MATH36001 Generalized Inverses and the SVD 2015

1 Generalized Inverses of Matrices

A matrix has an inverse only if it is square and nonsingular. However there are theoretical and practical applications for which some kind of partial inverse of a matrix that is singular or even rectangular is needed. A **generalized inverse** of a matrix A is any matrix X satisfying

$$AXA = A. \tag{1}$$

Note that a generalized inverse X must be the usual inverse when A is nonsingular since multiplication by A^{-1} on both sides of (1) gives $X = A^{-1}$.

1.1 Illustration: Solvability of Linear Systems

Consider the linear system $Ax = b$, where the matrix $A \in \mathbb{C}^{n\times n}$ and the vector $b \in \mathbb{C}^n$ are given and x is an unknown vector. If A is nonsingular there is a unique solution for x given by $x = A^{-1}b$. When A is singular, there may be no solution or infinitely many.

Theorem 1 *Let $A \in \mathbb{C}^{m\times n}$. If X is any matrix satisfying $AXA = A$ then $Ax = b$ has a solution if and only if $AXb = b$, in which case the general solution is*

$$x = Xb + (I - XA)y, \tag{2}$$

where $y \in \mathbb{C}^n$ is an arbitrary vector.

Proof. Existence Let X be such that $AXA = A$.
($\Rightarrow$) Suppose there exists x such that $Ax = b$. Then by (1), $b = Ax = AXAx = AXb$.
($\Leftarrow$) $AXb = b \;\Rightarrow\; x = Xb$ solves the linear system.

General solution First show that $Xb+(I-XA)y$ is a solution. Indeed, $A(Xb+(I-XA)y) = AXb + 0 = b$.

Then if x is any solution we can write $x = XAx + (I - XA)x$, so every solution can be expressed in the form of (2). $\square$

1.2 Existence of Generalized Inverses

We now show that for every A there exist one or more generalized inverses. Let $A \in \mathbb{C}^{m\times n}$ and let $E_k, E_{k-1}, \ldots, E_1$ be elementary row operations and P a permutation matrix such that

$$EAP = \begin{bmatrix} I_r & K \\ O & O \end{bmatrix}, \quad r \le \min(n, m),$$

where $E = E_k E_{k-1} \ldots E_1$. The matrix $\left[\begin{smallmatrix} I_r & K \\ O & O \end{smallmatrix}\right]$ is the **reduced row echelon form** of A and $\mathrm{rank}(A) = r$ (see first year Linear Algebra course). Note that the two right-hand submatrices are absent when $r = n$ and the two lower submatrices are absent if $r = m$. It is easy to verify that any X given by

$$X = P \begin{bmatrix} I_r & O \\ O & L \end{bmatrix} E$$

for some $L \in \mathbb{C}^{(n-r)\times(m-r)}$ satisfies $AXA = A$. Indeed,

$$AXA = AP\begin{bmatrix} I_r & O \\ O & L \end{bmatrix} EE^{-1}\begin{bmatrix} I_r & K \\ O & O \end{bmatrix} P^{-1} = AP\begin{bmatrix} I_r & K \\ O & O \end{bmatrix} P^{-1} = E^{-1}\begin{bmatrix} I_r & K \\ O & O \end{bmatrix} P^{-1} = A.$$

Example 1 Determine a generalized inverse for

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 0 \\ 0 & 2 \end{bmatrix}. \tag{3}$$

Let us find the reduced row echelon form of A:

$$\begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} 1 & 2 \\ 2 & 0 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 0 & -4 \\ 0 & 2 \end{bmatrix}, \quad \begin{bmatrix} 1 & 1/2 & 0 \\ 0 & -1/4 & 0 \\ 0 & 1 & 2 \end{bmatrix}\begin{bmatrix} 1 & 2 \\ 0 & -4 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix},$$

so $E = \begin{bmatrix} 1 & 1/2 & 0 \\ 0 & -1/4 & 0 \\ 0 & 1 & 2 \end{bmatrix}\begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 1/2 & 0 \\ 1/2 & -1/4 & 0 \\ -2 & 1 & 2 \end{bmatrix}$, $P = I_2$. Hence,

$$X = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}\begin{bmatrix} 0 & 1/2 & 0 \\ 1/2 & -1/4 & 0 \\ -2 & 1 & 2 \end{bmatrix} = \begin{bmatrix} 0 & 1/2 & 0 \\ 1/2 & -1/4 & 0 \end{bmatrix}$$

is a generalized inverse of A. Check that $AXA = A$.

1.3 The Moore–Penrose Generalized Inverse

The **Moore–Penrose generalized inverse** of a matrix $A \in \mathbb{C}^{m\times n}$ is the unique matrix $X \in \mathbb{C}^{n\times m}$ satisfying the four Moore–Penrose conditions:

$$\begin{array}{ll} \text{(i)}\ \ AXA = A, & \text{(ii)}\ \ XAX = X, \\ \text{(iii)}\ \ AX = (AX)^*, & \text{(iv)}\ \ XA = (XA)^*. \end{array} \tag{4}$$

The Moore–Penrose generalized inverse of A is a generalized inverse since it satisfies (1). It is commonly called the **pseudoinverse** of A and is denoted by A^+. We show the existence of A^+ via the singular value decomposition of A in Section 2.1.

2 Singular Value Decomposition

The spectral theorem says that normal matrices can be unitarily diagonalized using a basis of eigenvectors. The singular value decomposition can be seen as a generalization of the spectral theorem to arbitrary, not necessarily square, matrices.

Theorem 2 (Singular value decomposition (SVD)) *A matrix $A \in \mathbb{C}^{m\times n}$ has a singular value decomposition (SVD)*

$$A = U\Sigma V^*,$$

where $U \in \mathbb{C}^{m\times m}$, $V \in \mathbb{C}^{n\times n}$ are unitary and $\Sigma = \operatorname{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m\times n}$, $p = \min(m, n)$, where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p \geq 0$. If A is real, U and V can be taken real orthogonal.

Proof. (*Not examinable.*)
Let $x \in \mathbb{C}^n$ and $y \in \mathbb{C}^m$ be unit 2-norm vectors ($\|x\|_2 = \|y\|_2 = 1$) such that

$$Ax = \sigma y,$$

where $\sigma = \|A\|_2$. Since any orthonormal set can be extended to form an orthonormal basis for the whole space, we can define $V_1 \in \mathbb{C}^{n\times(n-1)}$ and $U_1 \in \mathbb{C}^{m\times(m-1)}$ such that $V = [x, V_1]$ and $U = [y, U_1]$ are unitary. Then

$$U^*AV = \begin{bmatrix} y^* \\ U_1^* \end{bmatrix} A \,[\,x \quad V_1\,] = \begin{bmatrix} y^*Ax & y^*AV_1 \\ U_1^*Ax & U_1^*AV_1 \end{bmatrix} = \begin{bmatrix} \sigma & w^* \\ 0 & B \end{bmatrix} = A_1,$$

where $w^* = y^*AV_1$ and $B = U_1^*AV_1$. We now show that $w = 0$. We have

$$A_1 \begin{bmatrix} \sigma \\ w \end{bmatrix} = \begin{bmatrix} \sigma & w^* \\ 0 & B \end{bmatrix} \begin{bmatrix} \sigma \\ w \end{bmatrix} = \begin{bmatrix} \sigma^2 + w^*w \\ Bw \end{bmatrix}$$

so

$$\sigma^2 + w^*w = \left\| \begin{bmatrix} \sigma^2 + w^*w \\ 0 \end{bmatrix} \right\|_2 \le \left\| \begin{bmatrix} \sigma^2 + w^*w \\ Bw \end{bmatrix} \right\|_2 = \left\| A_1 \begin{bmatrix} \sigma \\ w \end{bmatrix} \right\|_2 \le \|A_1\|_2 (\sigma^2 + w^*w)^{1/2}.$$

Hence,

$$\|A_1\|_2 \ge (\sigma^2 + w^*w)^{1/2}.$$

But, $\|A_1\|_2 = \|U^*AV\|_2 = \|A\|_2 = \sigma$ since U and V are unitary. This implies that $w = 0$ and $U^*AV = \begin{bmatrix} \sigma & 0 \\ 0 & B \end{bmatrix}$. The proof is completed by the obvious induction. □

The σ_i are called the **singular values** of A. The nonzero singular values of A are the positive square roots of the nonzero eigenvalues of AA^* or A^*A.

The columns of $U = [u_1, \dots, u_m]$ and $V = [v_1, \dots, v_n]$ are **left and right singular vectors** of A, respectively. The left singular vectors u_i are eigenvectors of AA^* and the right singular vectors v_i are eigenvectors of A^*A.

Suppose that the singular values satisfy

$$\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_r > \sigma_{r+1} = \cdots = \sigma_p = 0, \quad p = \min(m, n). \tag{5}$$

Then (see exercise 2)

$$\begin{aligned} \operatorname{rank}(A) &= r, \\ \operatorname{null}(A) &= \operatorname{span}\{v_{r+1}, \dots, v_n\}, && (6) \\ \operatorname{range}(A) &= \operatorname{span}\{u_1, u_2, \dots, u_r\}, && (7) \\ A &= \sum_{i=1}^{r} \sigma_i u_i v_i^*, \\ Av_i = \sigma_i u_i, &\quad A^*u_i = \sigma_i v_i, \quad i = 1, \dots, r. && (8) \end{aligned}$$

A geometric interpretation of the SVD: we can think of $A \in \mathbb{C}^{m\times n}$ as mapping $x \in \mathbb{C}^n$ to $y = Ax \in \mathbb{C}^m$. Then we can choose one orthogonal coordinate system for $\mathbb{C}^n$ (where the unit axes are the columns of V) and another orthogonal coordinate system for $\mathbb{C}^m$ (where the unit axes are the columns of U) such that A is diagonal (Σ), that is, maps a vector $x = \sum_{i=1}^n \beta_i v_i$ to $y = Ax = \sum_{i=1}^n \sigma_i \beta_i u_i$. In other words, "any matrix is diagonal," provided we pick up appropriate orthogonal systems for its domain and range!

Example 2 Compute the singular value decomposition of A in (3).

- The eigenvalues of $A^TA = \begin{bmatrix} 5 & 2 \\ 2 & 8 \end{bmatrix}$ are 9 and 4 so the singular values of A are 3 and 2.
- Normalized eigenvectors of A^TA are $v_1 = \frac{1}{\sqrt{5}}\left[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}\right]$, $v_2 = \frac{1}{\sqrt{5}}\left[\begin{smallmatrix} 2 \\ -1 \end{smallmatrix}\right]$.
- $u_1 = \frac{1}{3}Av_1 = \frac{1}{3\sqrt{5}}\begin{bmatrix} 5 \\ 2 \\ 4 \end{bmatrix}$, $u_2 = \frac{1}{2}Av_2 = \frac{1}{\sqrt{5}}\begin{bmatrix} 0 \\ 2 \\ -1 \end{bmatrix}$. Application of the Gram–Schmidt process to u_1, u_2 and e_1 produces $u_3 = \dfrac{e_1 - (\sum_{i=1}^2 e_1^T u_i)u_i}{\|e_1 - (\sum_{i=1}^2 e_1^T u_i)u_i\|_2} = \begin{bmatrix} 2/3 \\ -1/3 \\ -2/3 \end{bmatrix}$.
- A has the SVD $A = U\Sigma V^T$ where
$$U = \frac{1}{3\sqrt{5}}\begin{bmatrix} 5 & 0 & 2\sqrt{5} \\ 2 & 6 & -\sqrt{5} \\ 4 & -3 & -2\sqrt{5} \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 3 & 0 \\ 0 & 2 \\ 0 & 0 \end{bmatrix}, \quad V = \frac{1}{\sqrt{5}}\begin{bmatrix} 1 & 2 \\ 2 & -1 \end{bmatrix}.$$

What is the singular value decomposition of $x \in \mathbb{C}^n$?

Let $A = U\Sigma V^*$ be a singular value decomposition of $A \equiv x$ with $U \in \mathbb{C}^{n\times n}$ and $V \in \mathbb{C}^{1\times 1}$ unitary and $\Sigma = \begin{bmatrix} \sigma_1 \\ 0_{(n-1)\times 1} \end{bmatrix} \in \mathbb{R}^{n\times 1}$. Note that x^*x is a positive scalar so $\sigma_1 = (x^*x)^{1/2} = \|x\|_2$ and we can take $V = [v_1] = 1$. We have from (8) that $u_1 = \frac{1}{\sigma_1}Av_1 = \frac{1}{\|x\|_2}x$. Let $\widetilde{U}$ be any $n \times (n-1)$ matrix with orthonormal columns satisfying $u_1^*\widetilde{U} = 0$. Then $U = [\, u_1 \quad \widetilde{U}\,]$ is unitary and
$$U\Sigma V^* = \left[\, \tfrac{x}{\|x\|_2} \quad \widetilde{U}\,\right]\begin{bmatrix} \|x\|_2 \\ 0 \end{bmatrix}[\,1\,] = x.$$

2.1 Existence of the Moore–Penrose Inverse

Recall that the pseudoinverse $X \in \mathbb{C}^{n\times m}$ of $A \in \mathbb{C}^{m\times n}$ is the unique matrix satisfying the four Moore–Penrose conditions (4) and is denoted by A^+. The next theorem shows that any matrix A has a pseudoinverse.

Theorem 3 *If $A = U\Sigma V^* \in \mathbb{C}^{m\times n}$ is an SVD then*
$$A^+ = V\Sigma^+U^*,$$
where $\Sigma^+ = \mathrm{diag}(\sigma_1^{-1}, \ldots, \sigma_r^{-1}, 0, \ldots, 0)$ is $n \times m$ and $r = \mathrm{rank}(A)$.

Proof. One easily checks that Σ^+ and A^+ satisfy the four Moore–Penrose conditions (4) so that they are pseudoinverses of Σ and A, respectively. □

Example 3 From the SVD of A in (3) obtained in Example 2 we have that $A^+ = \frac{1}{18}\begin{bmatrix} 2 & 8 & -2 \\ 4 & -2 & 5 \end{bmatrix}$.

In general it is *not* the case that $(AB)^+ = B^+A^+$ for $A \in \mathbb{C}^{m\times n}$, $B \in \mathbb{C}^{n\times p}$. For example, let $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$, $B = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$. Then $(AB)^+ = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}^+ = \frac{1}{2}\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}$ and $B^+A^+ = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$.

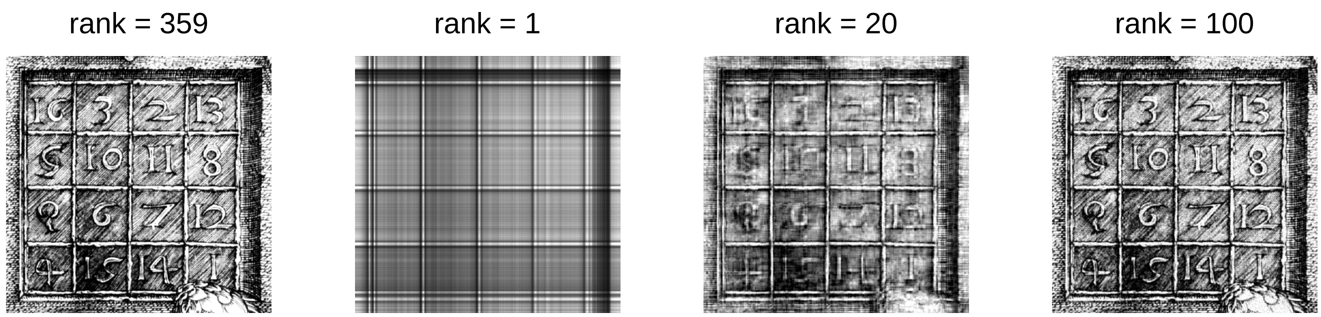

Figure 1: Low-rank approximation of Dürer's magic square.

2.2 Low-Rank Approximation

The singular values indicate how "near" a given matrix is to a matrix of low rank.

Theorem 4 *Let the SVD of* $A \in \mathbb{C}^{m\times n}$ *be given by Theorem* 2. *Write* $U = [u_1, \ldots, u_m]$ *and* $V = [v_1, \ldots, v_n]$. *If* $k < r = \operatorname{rank}(A)$ *and* $A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^*$ *then*

$$\min_{\operatorname{rank}(B)=k} \|A - B\|_2 = \|A - A_k\|_2 = \sigma_{k+1}.$$

Proof. Since $U^* A_k V = \operatorname{diag}(\sigma_1, \ldots, \sigma_k, 0 \ldots, 0)$, it follows that $\operatorname{rank}(A_k) = k$ and that $U^*(A - A_k)V = \operatorname{diag}(0, \ldots, 0, \sigma_{k+1}, \ldots, \sigma_p)$ and so $\|A - A_k\|_2 = \sigma_{k+1}$.

The theorem is proved if we can show that $\|A - B\|_2 \geq \sigma_{k+1}$ for all matrices B of rank k.

Let B be any matrix of rank k, so its null space is of dimension $n - k$. Let $\mathcal{V}_{k+1} = \operatorname{span}(v_1, \ldots, v_{k+1})$. Then $\dim \mathcal{V}_{k+1} = k + 1$ and $\operatorname{null}(B) \cap \mathcal{V}_{k+1} \neq \emptyset$ since the sum of their dimensions is $(n - k) + (k + 1) > n$. Let $z \in \operatorname{null}(B) \cap \mathcal{V}_{k+1}$ be of unit 2-norm (i.e., $\|z\|_2 = 1$) so that $Bz = 0$ and $z = \sum_{i=1}^{k+1} \zeta_i v_i$ with $\sum_{i=1}^{k+1} |\zeta_i|^2 = 1$ since z is of unit norm. Then

$$\|A-B\|_2^2 \geq \|(A-B)z\|_2^2 = \|Az\|_2^2 = \|U\Sigma V^* z\|_2^2 = \|\Sigma V^* z\|_2^2 = \sum_{i=1}^{k+1} \sigma_i^2 |\zeta_i|^2 \geq \sigma_{k+1}^2 \sum_{i=1}^{k+1} |\zeta_i|^2 = \sigma_{k+1}^2. \qquad \square$$

Example 4 (Image Compression) An $m \times n$ grayscale image is just an $m \times n$ matrix where entry (i, j) is interpreted as the brightness of pixel (i, j) ranging, say, from 0 (=black) to 1 (=white). Entries between 0 and 1 correspond to various shades of grey.

Now, let A be the matrix representing a detail from Albrecht Dürer's engraving "Melancolia I" from 1514 showing a 4×4 magic square (see first plot in Figure 1). The matrix A is of size 359×371 and of full rank. Its singular values decrease rapidly (only one $> 10^4$ and only six $> 10^3$). The three other plots in Figure 1 show the low rank approximations $A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^*$ for $k = 1, 20$, and 100. The checker board-like structure of A_1 is typical of very low-rank approximation. The individual numerals are recognizable in the $r = 20$ approximation. There is hardly any difference between the $r = 100$ approximation and the full rank image.

Although low-rank matrix approximation to images do require less computer storage and transmission time than the full image, there are more effective data compression techniques.

The primary uses of low-rank approximations in image processing involve feature recognition such as in handwritten digits, faces, and finger prints.

3 Projectors

Let $\mathcal{S}$ be a subspace of $\mathbb{C}^m$ and let $P_{\mathcal{S}} \in \mathbb{C}^{m\times m}$.

- $P_{\mathcal{S}}$ is the **projector** onto $\mathcal{S}$ if $\mathrm{range}(P_{\mathcal{S}}) = \mathcal{S}$ and $P_{\mathcal{S}}^2 = P_{\mathcal{S}}$. The projector is not unique.
- $P_{\mathcal{S}}$ is the **orthogonal projector** onto $\mathcal{S}$ if $\mathrm{range}(P_{\mathcal{S}}) = \mathcal{S}$, $P_{\mathcal{S}}^2 = P_{\mathcal{S}}$, and $P_{\mathcal{S}}^* = P_{\mathcal{S}}$. The orthogonal projector is unique (see Exercise 6). Also, $P_{\mathcal{S}^\perp} = I - P_{\mathcal{S}}$ is the orthogonal projector onto the orthogonal complement of $\mathcal{S}$ in $\mathbb{C}^m$ denoted by $\mathcal{S}^\perp$.

(Solution to Exercise 6): Let P_1 and P_2 be orthogonal projectors onto S. Since $\mathrm{range}(P_1) = \mathrm{range}(P_2)$, $P_2 = P_1 X$ for some X. Then $P_1P_2 = P_1^2 X = P_1 X = P_2$. Likewise, $P_2P_1 = P_1$. Hence, for any z,

$$\begin{aligned} \|(P_1 - P_2)z\|_2^2 &= z^*(P_1 - P_2)(P_1 - P_2)z \\ &= z^*(P_1^2 + P_2^2 - P_1P_2 - P_2P_1)z \\ &= z^*(P_1 + P_2 - P_2 - P_1)z = 0. \end{aligned}$$

Therefore $P_1 - P_2 = O$.

Example 5 Is the rank-one matrix $\dfrac{xy^*}{y^*x}$ with $x, y \in \mathbb{C}^n$, $y^*x \neq 0$, a projector? Is it an orthogonal projector?

$\dfrac{xy^*}{y^*x}$ is idempotent since $\left(\dfrac{xy^*}{y^*x}\right)\left(\dfrac{xy^*}{y^*x}\right) = \dfrac{xy^*}{y^*x}$. For any $u \in \mathbb{C}^n$, $\left(\dfrac{xy^*}{y^*x}\right)u = \dfrac{y^*u}{y^*x}x \in \mathrm{range}(x)$ so $P_{\mathrm{range}(x)} := \dfrac{xy^*}{y^*x}$ is a projector onto $\mathrm{range}(x)$.

If $x = \alpha y$ for some $\alpha \in \mathbb{R}$, $P_{\mathrm{range}(x)} := \dfrac{xx^*}{x^*x}$ is an orthogonal projector onto $\mathrm{range}(x)$ since $\dfrac{xx^*}{x^*x}$ is Hermitian.

Let $A \in \mathbb{C}^{m\times n}$ be such that $\mathrm{rank}(A) = r$. In terms of the SVD of A, $A = U\Sigma V^*$ with $U = [\,U_1 \quad U_2\,]$ and $V = [\,V_1 \quad V_2\,]$ unitary and partitioned such that $U_1 \in \mathbb{C}^{m\times r}$ and $V_1 \in \mathbb{C}^{n\times r}$, the orthogonal projectors onto the four fundamental subspaces of A are given by (see Exercise 7)

$$\begin{aligned} P_{\mathrm{range}(A)} &= U_1U_1^*, \qquad & P_{\mathrm{null}(A^*)} &= U_2U_2^*, \\ P_{\mathrm{range}(A^*)} &= V_1V_1^*, \qquad & P_{\mathrm{null}(A)} &= V_2V_2^*. \end{aligned}$$

In terms of the pseudoinverse,

$$\begin{aligned} P_{\mathrm{range}(A)} &= AA^+, \qquad & P_{\mathrm{null}(A^*)} &= I - AA^+, \\ P_{\mathrm{range}(A^*)} &= A^+A, \qquad & P_{\mathrm{null}(A)} &= I - A^+A. \end{aligned}$$

4 Least Squares Problems

Consider the system of linear equations $Ax = b$, where $A \in \mathbb{C}^{m\times n}$ and $b \in \mathbb{C}^m$ are given and x is an unknown vector.

Theorem 1 says that $Ax = b$ has a solution if and only if $AA^+b = b$ or, equivalently, $b \in \mathrm{range}(A)$. In this case the general solution is $x = A^+b + (I - A^+A)y$ for any $y \in \mathbb{C}^n$.

Theorem 5 (Minimum 2-norm solution) *For given $A \in \mathbb{C}^{m\times n}$ and $b \in \mathbb{C}^m$ with $b \in \mathrm{range}(A)$, the vector $x = A^+b$ is the solution of minimum 2-norm amongst all the solutions to $Ax = b$.*

Proof. We will need the following fact

Fact 1 For $u, v \in \mathbb{C}^n$: $u^*v = 0 \Longrightarrow \|u + v\|_2^2 = \|u\|_2^2 + \|v\|_2^2$.

This follows directly from $\|u + v\|_2^2 = (u + v)^*(u + v) = \|u\|_2^2 + \|v\|_2^2 + 2\,\mathrm{Re}(u^*v)$.

Then take $u = A^+b$ and $v = (I - A^+A)y$ with $y \in \mathbb{C}^n$, arbitrary and using (4) (ii) and (iv), check that $v^*u = y^*(I - A^+A)^*A^+b = y^*(A^+ - A^+AA^+)b = y^*(A^+ - A^+)b = 0$. It follows that for the general solution $x = A^+b + (I - A^+A)y$ to $Ax = b$,

$$\|x\|_2^2 = \|A^+b\|_2^2 + \|(I - A^+A)y\|_2^2$$

and the 2-norm of x is minimal when $(I - A^+A)y = 0$, in which case $x = A^+b$. (Note that $(I - A^+A)y = 0 \Leftrightarrow y \in \mathrm{null}(A)^\perp = \mathrm{range}(A^*)$, see Exercise 7(i).)

□

Theorem 1 says that there is no solution to $Ax = b$ when $b \notin \mathrm{range}(A)$. However for some purposes we may be satisfied with a **least squares solution**, which is a vector $x \in \mathbb{C}^n$ such that $\|Ax - b\|_2$ is minimized.

Theorem 6 (Least squares solutions) *For given $A \in \mathbb{C}^{m\times n}$ and $b \in \mathbb{C}^m$ the vectors*

$$x = A^+b + (I - A^+A)y, \quad y \in \mathbb{C}^n \text{ arbitrary},$$

minimize $\|Ax - b\|_2$. Moreover $x_{LS} = A^+b$ is the least squares solution of minimum 2-norm.

Proof. Let $r = \mathrm{rank}(A)$ and $A = U\Sigma V^*$ be an SVD. For any $x \in \mathbb{C}^n$ we have

$$\begin{aligned}\|Ax - b\|_2^2 &= \|U^*AV(V^*x) - U^*b\|_2^2 \\ &= \|\Sigma y - c\|_2^2 \qquad (y = V^*x,\ c = U^*b) \\ &= \sum_{i=1}^{r} |\sigma_i y_i - c_i|^2 + \sum_{i=r+1}^{m} |c_i|^2.\end{aligned}$$

$\|Ax - b\|_2^2$ is minimized if and only if $y_i = c_i/\sigma_i = u_i^*b/\sigma_i, \quad i = 1, \ldots, r$, where $y_{r+1}, \ldots, y_n$ are arbitrary. Hence

$$x = Vy = V \begin{bmatrix} u_1^* b/\sigma_1 \\ \vdots \\ u_r^* b/\sigma_r \\ y_{r+1} \\ \vdots \\ y_n \end{bmatrix} = \sum_{i=1}^{r} \frac{u_i^* b}{\sigma_i} v_i + \sum_{i=r+1}^{n} y_i v_i, \quad \text{with } y_i \text{ arbitrary for } i = r+1, \ldots, n.$$

The formula for x in the theorem follows from $A^+ b = V\Sigma^+ U^* b = \sum_{i=1}^{r} \frac{u_i^* b}{\sigma_i} v_i$ and $\sum_{i=r+1}^{n} y_i v_i = (I - A^+ A)z$ for some $z \in \mathbb{C}^n$ since $(I - A^+ A)$ is the orthogonal projector onto $\mathrm{null}(A) = \mathrm{span}\{v_{r+1}, \ldots, v_n\}$.

Since V is unitary, $\|x\|_2 = \|y\|_2$ and we get the minimum norm solution x_{LS} by setting $y_{r+1} = \cdots = y_n = 0$. Therefore,

$$x_{LS} = Vy = \sum_{i=1}^{r} y_i v_i = \sum_{i=1}^{r} (u_i^* b/\sigma_i) v_i = A^+ b. \qquad \square$$

To summarize we have the following diagram: see Ortega p.169.

5 Polar Decomposition

The polar decomposition is the generalization to matrices of the familiar polar representation $z = re^{i\theta}$ of a complex number. It is intimately related to the singular value decomposition (SVD), as our proof of the decomposition reveals.

Theorem 7 (Polar decomposition) *Any $A \in \mathbb{C}^{m \times n}$ with $m \geq n$ can be represented in the form*

$$A = QH,$$

where $Q \in \mathbb{C}^{m \times n}$ has orthonormal columns and $H \in \mathbb{C}^{n \times n}$ is Hermitian positive semidefinite. If A is real then Q can be taken real orthogonal and H symmetric positive semidefinite.

Proof. Let $A = U\left[\begin{smallmatrix}\Sigma_r & 0\\ 0 & 0_{m-r,n-r}\end{smallmatrix}\right]V^*$ be an SVD, $r = \operatorname{rank}(A)$. Then

$$A = U\begin{bmatrix} I_r & 0\\ 0 & I_{m-r,n-r}\end{bmatrix}V^* \cdot V\begin{bmatrix}\Sigma_r & 0\\ 0 & 0_{n-r}\end{bmatrix}V^* \equiv QH. \tag{9}$$

Q has orthonormal columns since U, V are unitary and $\left[\begin{smallmatrix} I_r & 0\\ 0 & I_{m-r,n-r}\end{smallmatrix}\right]$ has orthonormal columns. H is Hermitian since $H^* = H$ and positive semidefinite since all its eigenvalues are real and nonnegative. □

Example 6 (Aligning two objects) Consider the molecule A, which we will specify by the coordinates $a_1, \dots, a_7$ of the centers of the seven spheres that represent some of its atoms, and let B be a second molecule. Define the 3×7 matrices $A = [\, a_1 \quad \dots \quad a_7 \,]$ and $B = [\, b_1 \quad \dots \quad b_7 \,]$.

Molecule A

Molecule B

How can we tell whether molecule A and molecule B are the same, that is, if B was obtained by "rotating A". In other words, is $A = QB$ for some orthogonal matrix Q? Because life is about change and imperfection, we do not expect to obtain exact equality, but we want to make the difference between B and QA as small as possible. So our task is to solve the following problem

$$\text{minimize } \|A - QB\|_F \text{ subject to } Q^TQ = I.$$

This is called an **orthogonal Procrustes problem**. It is shown in Exercise 5 that the transpose of the solution Q is the orthogonal polar factor of BA^T.

Exercises

1. Let $A \in \mathbb{C}^{m\times n}$, $m \geq n$. If $\operatorname{rank}(A) = n$, show that A^*A is nonsingular and that $A^+ = (A^*A)^{-1}A^*$.

2. Let $A \in \mathbb{C}^{m\times n}$ with $m \geq n$ be such that $\operatorname{rank}(A) = r$. The matrix A has the singular value decomposition $A = U\Sigma V^*$, where $U = [u_1, \dots, u_m] \in \mathbb{C}^{m\times m}$, $V = [v_1, \dots, v_n] \in \mathbb{C}^{n\times n}$ are unitary, and $\Sigma = \operatorname{diag}(\sigma_1, \dots, \sigma_n) \in \mathbb{R}^{m\times n}$, where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > \sigma_{r+1} = \cdots = \sigma_n = 0$. Show that

$$\operatorname{null}(A) = \operatorname{span}\{v_{r+1}, \dots, v_n\},$$
$$\operatorname{range}(A) = \operatorname{span}\{u_1, u_2, \dots, u_r\}.$$

3. For $A \in \mathbb{C}^{n\times n}$ use the SVD to find expressions for $\|A\|_2$ and $\|A\|_F$ in terms of the singular values of A. Hence obtain a bound of the form $c_1\|A\|_2 \le \|A\|_F \le c_2\|A\|_2$, where c_1 and c_2 are constants that depend on n. When is there equality in these two inequalities?

4. Show that the pseudoinverse A^+ of $A \in \mathbb{C}^{m\times n}$ solves the problem

$$\min_{X\in\mathbb{C}^{n\times m}} \|AX - I_m\|_F.$$

[Hint: Reduce it to m standard least squares problems.] Is the solution unique?

5. Let $A, B \in \mathbb{R}^{m\times n}$. Recall that $\|A\|_F^2 = \mathrm{trace}(A^TA)$ and that for any matrix D for which the product AD is defined, $\mathrm{trace}(AD) = \mathrm{trace}(DA)$.

(i) Show that the Q that minimizes $\|A - QB\|_F$ over all choices of orthogonal Q also maximizes $\mathrm{trace}(A^TQB)$.

(ii) Suppose that the SVD of the $m\times m$ matrix BA^T is $U\Sigma V^T$, where U and V are $m\times m$ and orthogonal and Σ is diagonal with diagonal entries $\sigma_1 \ge \cdots \ge \sigma_m \ge 0$. Define $Z = V^TQU$. Use these definitions and (i) to show that

$$\mathrm{trace}(A^TQB) = \mathrm{trace}(Z\Sigma) \le \sum_{i=1}^{m} \sigma_i.$$

(iii) Identify the choice of Q that gives equality in the bound of (ii).

(iv) Carefully state a theorem summarizing the solution to the minimization of $\|A-QB\|_F$.

6. Show that the orthogonal projector onto a subspace S is unique.

7. Let $A \in \mathbb{C}^{m\times n}$.

(i) Show that $\big(\mathrm{range}(A)\big)^\perp = \mathrm{null}(A^*)$ and that $\big(\mathrm{null}(A)\big)^\perp = \mathrm{range}(A^*)$.

(ii) Show that AA^+, A^+A, $I-A^+A$ and $I-AA^+$ are the orthogonal projectors onto $\mathrm{range}(A)$, $\mathrm{range}(A^*)$, $\mathrm{null}(A)$ and $\mathrm{null}(A^*)$, respectively.

(iii) Suppose $\mathrm{rank}(A) = r$ and let $A = U\Sigma V^*$ with $U = [\,U_1 \quad U_2\,]$ and $V = [\,V_1 \quad V_2\,]$ unitary and partitioned such that $U_1 \in \mathbb{C}^{m\times r}$ and $V_1 \in \mathbb{C}^{n\times r}$. Show that $U_1U_1^*$, $V_1V_1^*$, $V_2V_2^*$ and $U_2U_2^*$ are the orthogonal projectors onto $\mathrm{range}(A)$, $\mathrm{range}(A^*)$, $\mathrm{null}(A)$ and $\mathrm{null}(A^*)$, respectively.

8. What is the pseudoinverse of a vector $x \in \mathbb{C}^n$?

9. Is it true that $(A^+)^+ = A$?