Chapter 2

Singular Value Decomposition (SVD) and Polar Form

2.1 Polar Form

In this chapter, we assume that we are dealing with a real Euclidean space E. Let $f\colon E \to E$ be any linear map.

In general, it may not be possible to diagonalize f. However, note that $f^* \circ f$ is self-adjoint, since

$$\langle (f^* \circ f)(u), v\rangle = \langle f(u), f(v)\rangle = \langle u, (f^* \circ f)(v)\rangle.$$

Similarly, $f \circ f^*$ is self-adjoint.

The fact that $f^* \circ f$ and $f \circ f^*$ are self-adjoint is very important, because it implies that $f^* \circ f$ and $f \circ f^*$ can be diagonalized and that they have real eigenvalues.

In fact, these eigenvalues are all ≥ 0.

Thus, the eigenvalues of $f^* \circ f$ are of the form $\mu_1^2, \ldots, \mu_r^2$ or 0, where $\mu_i > 0$, and similarly for $f \circ f^*$.

The situation is even better, since we will show shortly that $f^* \circ f$ and $f \circ f^*$ have the same eigenvalues.

Remark: If $f\colon E \to F$ and $g\colon F \to E$ are linear maps, then $g \circ f$ and $f \circ g$ always have the same non-zero eigenvalues! Furthermore, if $E = F$, then 0 is an eigenvalue for $f \circ g$ iff it is an eigenvalue for $g \circ f$.

The square roots $\mu_i > 0$ of the positive eigenvalues of $f^* \circ f$ (and $f \circ f^*$) are called the *singular values of* f.

A self-adjoint linear map $f\colon E \to E$ whose eigenvalues are all ≥ 0 is called *positive semi-definite*, for short, *positive*, and if f is also invertible, *positive definite*. In the latter case, every eigenvalue is strictly positive.

We just showed that $f^* \circ f$ and $f \circ f^*$ are positive self-adjoint linear maps.

The wonderful thing about the singular value decomposition is that there exist two orthonormal bases $(u_1, \ldots, u_n)$ and $(v_1, \ldots, v_n)$ such that with respect to these bases, f is a diagonal matrix consisting of the singular values of f, or 0.

Given two Euclidean spaces E and F, where the inner product on E is denoted as $\langle -, - \rangle_1$ and the inner product on F is denoted as $\langle -, - \rangle_2$, given any linear map $f \colon E \to F$, there is a unique linear map $f^* \colon F \to E$ such that

$$\langle f(u), v \rangle_2 = \langle u, f^*(v) \rangle_1$$

for all $u \in E$ and all $v \in F$.

The linear map f^* is also called the adjoint of f.

This more general situation will be encountered when we deal with the singular value decomposition of rectangular matrices.

Recall that if $f\colon E \to F$ is a linear map, the *image* $\mathrm{Im}\, f$ *of* f is the subspace $f(E)$ of F, and the *rank of* f is the dimension $\dim(\mathrm{Im}\, f)$ of its image.

Also recall that

$$\dim\,(\mathrm{Ker}\, f) + \dim\,(\mathrm{Im}\, f) = \dim\,(E),$$

and that for every subspace W of E

$$\dim\,(W) + \dim\,(W^{\perp}) = \dim\,(E).$$

Lemma 2.1.1 *Given any two Euclidean spaces E and F, where E has dimension n and F has dimension m, for any linear map $f\colon E \to F$, we have*

$$\begin{aligned}
\mathrm{Ker}\, f &= \mathrm{Ker}\,(f^* \circ f),\\
\mathrm{Ker}\, f^* &= \mathrm{Ker}\,(f \circ f^*),\\
\mathrm{Ker}\, f &= (\mathrm{Im}\, f^*)^{\perp},\\
\mathrm{Ker}\, f^* &= (\mathrm{Im}\, f)^{\perp},\\
\dim(\mathrm{Im}\, f) &= \dim(\mathrm{Im}\, f^*),\\
\dim(\mathrm{Ker}\, f) &= \dim(\mathrm{Ker}\, f^*),
\end{aligned}$$

and f, f^*, $f^* \circ f$, and $f \circ f^*$, have the same rank.

The next Lemma shows a very useful property of positive self-adjoint linear maps.

Lemma 2.1.2 *Given a Euclidean space E of dimension n, for any positive self-adjoint linear map $f\colon E \to E$, there is a unique positive self-adjoint linear map $h\colon E \to E$ such that $f = h^2 = h \circ h$. Furthermore,* $\mathrm{Ker}\, f = \mathrm{Ker}\, h$, *and if $\mu_1, \ldots, \mu_p$ are the distinct eigenvalues of h and E_i is the eigenspace associated with μ_i, then $\mu_1^2, \ldots, \mu_p^2$ are the distinct eigenvalues of f, and E_i is the the eigenspace associated with μ_i^2,*

To prove Lemma 2.1.2, we use the following facts:
If g is self-adjoint, positive, then

(1) $\mathrm{Ker}\ g = \mathrm{Ker}\ g^2$;

(2) If $\mu_1, \ldots, \mu_p$ are the distinct nonzero eigenvalues of g, then $\mu_1^2, \ldots, \mu_p^2$ are the distinct nonzero eigenvalues of g^2 $(= f)$;

(3) $\mathrm{Ker}\,(g - \mu_i \mathrm{id}) = \mathrm{Ker}\,(g^2 - \mu_i^2 \mathrm{id})$.

By the way, one easily checks that $g^2 = g \circ g$ is positive and self-adjoint.

In (3), if $(g^2 - \mu_i^2\mathrm{id})(u) = 0$ with $u \neq 0$, then

$$(g + \mu_i\mathrm{id}) \circ (g - \mu_i\mathrm{id})(u) = 0.$$

But, $(g + \mu_i\mathrm{id})$ must be invertible since, otherwise, $-\mu_i < 0$ would be an eigenvalue of g, which is absurd. Therefore, $(g - \mu_i\mathrm{id})(u) = 0$.

Actually, it is possible to prove that if f is self-adjoint, positive, then for every $n \geq 2$, there is a unique self-adjoint, positive, g, so that

$$f = g^n.$$

There are now two ways to proceed. We can prove directly the singular value decomposition, as Strang does [**?**, **?**], or prove the so-called *polar decomposition* theorem.

The proofs are roughly of the same difficulty. We choose the second approach since it is less common in textbook presentations, and since it also yields a little more, namely uniqueness when f is invertible.

It is somewhat disconcerting that the next two theorems are only given as an exercise in Bourbaki [**?**] (Algèbre, Chapter 9, problem 14, page 127). Yet, the SVD decomposition is of great practical importance.

This is probably typical of the attitude of "pure mathematicians". However, the proof hinted at in Bourbaki is quite elegant.

Theorem 2.1.3 *Given a Euclidean space E of dimension n, for every linear map $f\colon E \rightarrow E$, there are two positive self-adjoint linear maps $h_1\colon E \rightarrow E$ and $h_2\colon E \rightarrow E$ and an orthogonal linear map $g\colon E \rightarrow E$ such that*

$$f = g \circ h_1 = h_2 \circ g.$$

Furthermore, if f has rank r, the maps h_1 and h_2 have the same positive eigenvalues $\mu_1, \ldots, \mu_r$, which are the singular values of f, i.e., the positive square roots of the nonnull eigenvalues of both $f^* \circ f$ and $f \circ f^*$. Finally, g, h_1, h_2 are unique if f is invertible, and $h_1 = h_2$ if f is normal.

In matrix form, Theorem 2.2.1 can be stated as follows.

For every real $n \times n$ matrix A, there is some orthogonal matrix R and some positive symmetric matrix S such that

$$A = RS.$$

Furthermore, R, S are unique if A is invertible. A pair (R, S) such that $A = RS$ is called a *polar decomposition of* A.

Remark: If E is a Hermitian space, Theorem 2.1.3 also holds, but the orthogonal linear map g becomes a unitary map.

In terms of matrices, the polar decomposition states that for every complex $n \times n$ matrix A, there is some unitary matrix U and some positive Hermitian matrix H such that

$$A = UH.$$

2.2 Singular Value Decomposition (SVD)

The proof of Theorem 2.1.3 shows that there are two orthonormal bases $(u_1, \ldots, u_n)$ and $(v_1, \ldots, v_n)$, where

$$(u_1, \ldots, u_n)$$

are eigenvectors of h_1 and

$$(v_1, \ldots, v_n)$$

are eigenvectors of h_2. Furthermore,

$$(u_1, \ldots, u_r)$$

is an orthonormal basis of $\mathrm{Im}\, f^*$,

$$(u_{r+1}, \ldots, u_n)$$

is an orthonormal basis of $\mathrm{Ker}\, f$,

$$(v_1, \ldots, v_r)$$

is an orthonormal basis of $\mathrm{Im}\, f$, and

$$(v_{r+1}, \ldots, v_n)$$

is an orthonormal basis of $\mathrm{Ker}\, f^*$.

Using this, we immediately obtain the singular value decomposition theorem.

Theorem 2.2.1 *Given a Euclidean space E of dimension n, for every linear map $f\colon E \to E$, there are two orthonormal bases $(u_1, \ldots, u_n)$ and $(v_1, \ldots, v_n)$ such that if r is the rank of f, the matrix of f w.r.t. these two bases is a diagonal matrix of the form*

$$\begin{pmatrix} \mu_1 & & \ldots & \\ & \mu_2 & \ldots & \\ \vdots & \vdots & \ddots & \vdots \\ & & \ldots & \mu_n \end{pmatrix}$$

where $\mu_1, \ldots, \mu_r$ are the singular values of f, i.e. the (positive) square roots of the nonnull eigenvalues of $f^* \circ f$ and $f \circ f^*$, and $\mu_{r+1} = \ldots = \mu_n = 0$. Furthermore, $(u_1, \ldots, u_n)$ are eigenvectors of $f^* \circ f$, $(v_1, \ldots, v_n)$ are eigenvectors of $f \circ f^*$, and $f(u_i) = \mu_i v_i$ when $1 \leq i \leq n$.

Note that $\mu_i > 0$ for all i $(1 \leq i \leq n)$ iff f is invertible.

Theorem 2.2.1 can be restated in terms of (real) matrices as follows.

Theorem 2.2.2 *For every real $n \times n$ matrix A, there are two orthogonal matrices U and V and a diagonal matrix D such that $A = V D U^\top$, where D is of the form*

$$D = \begin{pmatrix} \mu_1 & & \dots & \\ & \mu_2 & \dots & \\ \vdots & \vdots & \ddots & \vdots \\ & & \dots & \mu_n \end{pmatrix}$$

where $\mu_1, \dots, \mu_r$ are the singular values of f, i.e. the (positive) square roots of the nonnull eigenvalues of $A^\top A$ and $A\,A^\top$, and $\mu_{r+1} = \dots = \mu_n = 0$. The columns of U are eigenvectors of $A^\top A$, and the columns of V are eigenvectors of $A\,A^\top$. Furthermore, if $\det(A) \geq 0$, it is possible to choose U and V so that $\det(U) = \det(V) = +1$, i.e., U and V are rotation matrices.

A triple (U, D, V) such that $A = V D U^\top$ is called a *singular value decomposition (SVD) of A.*

Remarks. In Strang [**?**], the matrices U, V, D are denoted as $U = Q_2$, $V = Q_1$, and $D = \Sigma$, and a SVD decomposition is written as $A = Q_1 \Sigma Q_2^\top$.

This has the advantage that Q_1 comes before Q_2 in $A = Q_1 \Sigma Q_2^\top$. This has the disadvantage that A maps the columns of Q_2 (eigenvectors of $A^\top A$) to multiples of the columns of Q_1 (eigenvectors of $A\,A^\top$).

The SVD also applies to complex matrices.

In this case, for every complex $n \times n$ matrix A, there are two unitary matrices U and V and a diagonal matrix D such that

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

where D is a diagonal matrix consisting of real entries $\mu_1, \ldots, \mu_n$, where $\mu_1, \ldots, \mu_r$ are the singular values of f, i.e. the positive square roots of the nonnull eigenvalues of $A^* A$ and $A\,A^*$, and $\mu_{r+1} = \ldots = \mu_n = 0$.

It is easy to go from the polar form to the SVD, and backward.

Indeed, given a polar decomposition

$$A = R_1 S,$$

where R_1 is orthogonal and S is positive symmetric, there is an orthogonal matrix R_2 and a positive diagonal matrix D such that $S = R_2 D\, R_2^\top$, and thus

$$A = R_1 R_2 D\, R_2^\top = V D\, U^\top,$$

where $V = R_1 R_2$ and $U = R_2$ are orthogonal.

Going the other way, given an SVD decomposition

$$A = V D U^{\top},$$

let $R = V U^{\top}$ and $S = U D U^{\top}$.

It is clear that R is orthogonal and that S is positive symmetric, and

$$RS = V U^{\top} U D U^{\top} = V D U^{\top} = A.$$

Note that it is possible to require that $\det(R) = +1$ when $\det(A) \geq 0$.

Theorem 2.2.2 can be easily extended to rectangular $m \times n$ matrices (see Strang [**?**]).

As a matter of fact, both Theorems 2.1.3 and 2.2.1 can be generalized to linear maps $f \colon E \rightarrow F$ between two Euclidean spaces E and F.

In order to do so, we need to define the analog of the notion of an orthogonal linear map for linear maps $f: E \to F$.

By definition, the adjoint $f^*: F \to E$ of a linear map $f: E \to F$ is the unique linear map such that

$$\langle f(u), v\rangle_2 = \langle u, f^*(v)\rangle_1$$

for all $u \in E$ and all $v \in F$.

Then, we have

$$\langle f(u), f(v)\rangle_2 = \langle u, (f^* \circ f)(v)\rangle_1$$

for all $u, v \in E$.

Letting $n = \dim(E)$, $m = \dim(F)$, and $p = \min(m, n)$, if f has rank p and if for every p orthonormal vectors $(u_1, \ldots, u_p)$ in $(\mathrm{Ker}\, f)^{\perp}$, the vectors $(f(u_1), \ldots, f(u_p))$ are also orthonormal in F, then

$$f^* \circ f = \mathrm{id}$$

on $(\mathrm{Ker}\, f)^{\perp}$.

The converse is immediately proved.

Thus, we will say that a linear map $f \colon E \to F$ is *weakly orthogonal* if is has rank $p = \min(m, n)$ and if

$$f^* \circ f = \mathrm{id}$$

on $(\mathrm{Ker}\, f)^{\perp}$.

Of course, $f^* \circ f = 0$ on $\mathrm{Ker}\, f$.

In terms of matrices, we will say that a real $m \times n$ matrix A is weakly orthogonal iff its first $p = \min(m, n)$ columns are orthonormal, the remaining ones (if any) being null columns.

This is equivalent to saying that

$$A^\top A = I_n$$

if $m \geq n$, and that

$$A^\top A = \begin{pmatrix} I_m & 0_{m,n-m} \\ 0_{n-m,m} & 0_{n-m,n-m} \end{pmatrix}$$

if $n > m$.

In this latter case ($n > m$), it is immediately shown that

$$A\,A^\top = I_m,$$

and $A^\top$ is also weakly orthogonal. Weakly unitary linear maps are defined analogously.

Theorem 2.2.3 *Given any two Euclidean spaces E and F, where E has dimension n and F has dimension m, for every linear map $f\colon E \to F$, there are two positive self-adjoint linear maps $h_1\colon E \to E$ and $h_2\colon F \to F$ and a weakly orthogonal linear map $g\colon E \to F$ such that*

$$f = g \circ h_1 = h_2 \circ g.$$

Furthermore, if f has rank r, the maps h_1 and h_2 have the same positive eigenvalues $\mu_1, \ldots, \mu_r$, which are the singular values of f, i.e., the positive square roots of the nonnull eigenvalues of both $f^* \circ f$ and $f \circ f^*$. Finally, g, h_1, h_2 are unique if f is invertible, and $h_1 = h_2$ if f is normal.

In matrix form, Theorem 2.2.3 can be stated as follows.

For every real $m \times n$ matrix A, there is some weakly orthogonal $m \times n$ matrix R and some positive symmetric $n \times n$ matrix S such that

$$A = RS.$$

Remark: If E is a Hermitian space, Theorem 2.2.3 also holds, but the weakly orthogonal linear map g becomes a weakly unitary map.

In terms of matrices, the polar decomposition states that for every complex $m \times n$ matrix A, there is some weakly unitary $m \times n$ matrix U and some positive Hermitian $n \times n$ matrix H such that

$$A = UH.$$

The proof of Theorem 2.2.3 shows that there are two orthonormal bases

$$(u_1, \ldots, u_n)$$

of E and

$$(v_1, \ldots, v_m)$$

of F, where

$$(u_1, \ldots, u_n)$$

are eigenvectors of h_1 and

$$(v_1, \ldots, v_m)$$

are eigenvectors of h_2.

Furthermore,

$$(u_1, \ldots, u_r)$$

is an orthonormal basis of $\mathrm{Im}\, f^*$,

$$(u_{r+1}, \ldots, u_n)$$

is an orthonormal basis of $\mathrm{Ker}\, f$,

$$(v_1, \ldots, v_r)$$

is an orthonormal basis of $\mathrm{Im}\, f$, and

$$(v_{r+1}, \ldots, v_m)$$

is an orthonormal basis of $\mathrm{Ker}\, f^*$.

Using this, we immediately obtain the singular value decomposition theorem for linear maps $f\colon E \to F$, where E and F can have different dimensions.

Theorem 2.2.4 *Given any two Euclidean spaces E and F, where E has dimension n and F has dimension m, for every linear map $f\colon E \to F$, there are two orthonormal bases $(u_1, \ldots, u_n)$ and $(v_1, \ldots, v_m)$ such that if r is the rank of f, the matrix of f w.r.t. these two bases is a $m \times n$ matrix D of the form*

$$D = \begin{pmatrix} \mu_1 & & \ldots & \\ & \mu_2 & \ldots & \\ \vdots & \vdots & \ddots & \vdots \\ & & \ldots & \mu_n \\ 0 & \vdots & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \vdots & \ldots & 0 \end{pmatrix} \quad \textit{or}$$

$$D = \begin{pmatrix} \mu_1 & & \ldots & & 0 & \ldots & 0 \\ & \mu_2 & \ldots & & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots & 0 & \vdots & 0 \\ & & \ldots & \mu_m & 0 & \ldots & 0 \end{pmatrix}$$

where $\mu_1, \ldots, \mu_r$ are the singular values of f, i.e. the (positive) square roots of the nonnull eigenvalues of $f^* \circ f$ and $f \circ f^*$, and $\mu_{r+1} = \ldots = \mu_p = 0$, where $p = \min(m, n)$.

Furthermore, $(u_1, \ldots, u_n)$ *are eigenvectors of* $f^* \circ f$, $(v_1, \ldots, v_m)$ *are eigenvectors of* $f \circ f^*$, *and* $f(u_i) = \mu_i v_i$ *when* $1 \leq i \leq p = \min(m, n)$.

Even though the matrix D is an $m \times n$ rectangular matrix, since its only nonzero entries are on the descending diagonal, we still say that D is a diagonal matrix.

Theorem 2.2.4 can be restated in terms of (real) matrices as follows.

Theorem 2.2.5 *For every real $m \times n$ matrix A, there are two orthogonal matrices U $(n \times n)$ and V $(m \times m)$ and a diagonal $m \times n$ matrix D such that $A = V D U^\top$, where D is of the form*

$$D = \begin{pmatrix} \mu_1 & & \ldots & \\ & \mu_2 & \ldots & \\ \vdots & \vdots & \ddots & \vdots \\ & & \ldots & \mu_n \\ 0 & \vdots & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \vdots & \ldots & 0 \end{pmatrix} \quad \textit{or}$$

$$D = \begin{pmatrix} \mu_1 & & \ldots & & 0 & \ldots & 0 \\ & \mu_2 & \ldots & & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots & 0 & \vdots & 0 \\ & & \ldots & \mu_m & 0 & \ldots & 0 \end{pmatrix}$$

where $\mu_1, \ldots, \mu_r$ are the singular values of f, i.e. the (positive) square roots of the nonnull eigenvalues of $A^\top A$ and $A\,A^\top$, and $\mu_{r+1} = \ldots = \mu_p = 0$, where $p = \min(m, n)$. The columns of U are eigenvectors of $A^\top A$, and the columns of V are eigenvectors of $A\,A^\top$.

Given a (complex) $n \times n$ matrix, A, is there an interesting relationship between the eigenvalues of A and the singular values of A?

The following remarkable theorem due to Hermann Weyl shows the answer is yes!:

Theorem 2.2.6 *(Weyl's inequalities, 1949) For any (complex) $n \times n$ matrix, A, if $\lambda_1, \ldots, \lambda_n \in \mathbb{C}$ are the eigenvalues of A and $\sigma_1, \ldots, \sigma_n \in \mathbb{R}_+$ are the singular values of A, listed so that $|\lambda_1| \geq \cdots \geq |\lambda_n|$ and $\sigma_1 \geq \cdots \geq \sigma_n \geq 0$, then*

$$\begin{aligned} |\lambda_1| \cdots |\lambda_n| &= \sigma_1 \cdots \sigma_n \quad \textit{and} \\ |\lambda_1| \cdots |\lambda_k| &\leq \sigma_1 \cdots \sigma_k, \quad \textit{for} \quad k = 1, \ldots, n-1. \end{aligned}$$

Actually, Theorem 2.2.6 has a converse due to A. Horn (1954): Given two sequences $\lambda_1, \ldots, \lambda_n \in \mathbb{C}$ and $\sigma_1, \ldots, \sigma_n \in \mathbb{R}_+$, if they satisfy the Weyl inequalities, then there is some matrix, A, having $\lambda_1, \ldots, \lambda_n$ as eigenvalues and $\sigma_1, \ldots, \sigma_n$ as singular values.

The SVD decomposition of matrices can be used to define the pseudo-inverse of a rectangular matrix.