Lecture 17: The Singular Value Decomposition: Theory

3.2. The singular value decomposition.

Both the normal equation and QR approaches to solving the discrete linear least squares problem assume that the matrix $\mathbf{A} \in \mathbb{C}^{m\times n}$ has full column rank, i.e., its columns are linearly independent, implying that both $\mathbf{A}^*\mathbf{A}$ and $\mathbf{R}_1$ are invertible. What if this is not the case? The discrete least squares problem still makes sense, but we need a more robust computational approach to determine the solution. What if the columns of $\mathbf{A}$ are close to being linearly dependent? What does it even mean to be 'close' to linear dependence?

To answer these questions, we shall investigate one of the most important matrix factorizations, the *singular value decomposition* (SVD). This factorization writes a matrix as the product of a unitary matrix times a diagonal matrix times another unitary matrix. It is an incredibly useful tool for proving a variety of results in matrix theory, but it also has essential computational applications: from the SVD we immediately obtain bases for the four fundamental subspaces, $\mathrm{Ran}(\mathbf{A})$, $\mathrm{Ker}(\mathbf{A})$, $\mathrm{Ran}(\mathbf{A}^*)$, and $\mathrm{Ker}(\mathbf{A}^*)$. Furthermore, the SVD facilitates the robust solution of a variety of approximation problems, including not only least squares problems with rank-deficient $\mathbf{A}$, but also other low-rank matrix approximation problems that arise throughout engineering, statistics, the physical sciences, and social science.

There are several ways to derive the singular value decomposition. We shall constructively prove the SVD based on analysis of $\mathbf{A}^*\mathbf{A}$; Trefethen and Bau follow an alternative approach somewhat different from the one we describe; see their Theorem 4.1. Before beginning, we must recall some fundamental results from linear algebra.

3.2.1. Hermitian positive definite matrices.

Theorem (Spectral Theorem). Suppose $\mathbf{H} \in \mathbb{C}^{n\times n}$ is Hermitian. Then there exist n (not necessarily distinct) eigenvalues $\lambda_1, \ldots, \lambda_n$ and corresponding unit eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ such that

$$\mathbf{H}\mathbf{v}_j = \lambda_j \mathbf{v}_j$$

and the eigenvectors form an *orthonormal* basis for $\mathbb{C}^n$.

Theorem. All eigenvalues of a Hermitian matrix are real.

Proof. Let $(\lambda_j, \mathbf{v}_j)$ be an arbitrary eigenpair of the Hermitian matrix $\mathbf{H}$, so that $\mathbf{H}\mathbf{v}_j = \lambda_j\mathbf{v}_j$. Without loss of generality, we can assume that $\mathbf{v}_j$ is scaled so that $\|\mathbf{v}_j\|_2 = 1$. Thus

$$\lambda_j = \lambda_j(\mathbf{v}_j^*\mathbf{v}_j) = \mathbf{v}_j^*(\lambda_j\mathbf{v}_j) = \mathbf{v}_j^*(\mathbf{H}\mathbf{v}_j) = \mathbf{v}_j^*\mathbf{H}^*\mathbf{v}_j = (\mathbf{H}\mathbf{v}_j)^*\mathbf{v}_j = (\lambda_j\mathbf{v}_j)^*\mathbf{v}_j = \overline{\lambda}_j\mathbf{v}_j^*\mathbf{v}_j = \overline{\lambda}_j.$$

Thus $\lambda_j = \overline{\lambda}_j$, which is only possible if λ_j is real. ■

Definition. A Hermitian matrix $\mathbf{H} \in \mathbb{C}^{n\times n}$ is *positive definite* provided $\mathbf{x}^*\mathbf{H}\mathbf{x} > 0$ for all nonzero $\mathbf{x} \in \mathbb{C}^n$; if $\mathbf{x}^*\mathbf{H}\mathbf{x} \ge 0$ for all $\mathbf{x} \in \mathbb{C}^n$, we say $\mathbf{H}$ is *positive semidefinite.*

Theorem. A Hermitian positive semidefinite matrix has nonnegative real eigenvalues.

Proof. Let $(\lambda_j, \mathbf{v}_j)$ denote an eigenpair of the Hermitian positive semidefinite matrix $\mathbf{H} \in \mathbb{C}^{n\times n}$ with $\|\mathbf{v}_j\|_2^2 = \mathbf{v}_j^*\mathbf{v}_j = 1$. Since $\mathbf{H}$ is Hermitian, λ_j must be real. We conclude that

$$\lambda_j = \lambda_j\mathbf{v}_j^*\mathbf{v}_j = \mathbf{v}_j^*(\lambda_j\mathbf{v}_j) = \mathbf{v}_j^*\mathbf{H}\mathbf{v}_j \ge 0$$

since $\mathbf{H}$ is positive semidefinite. ■

3.2.2. Derivation of the singular value decomposition.

Suppose $\mathbf{A} \in \mathbb{C}^{m\times n}$ with $m \geq n$. The n-by-n matrix $\mathbf{A}^*\mathbf{A}$ is always Hermitian positive semidefinite. (Clearly $(\mathbf{A}^*\mathbf{A})^* = \mathbf{A}^*(\mathbf{A}^*)^* = \mathbf{A}^*\mathbf{A}$, so $\mathbf{A}^*\mathbf{A}$ is Hermitian. For any $\mathbf{x} \in \mathbb{C}^n$, note that $\mathbf{x}^*\mathbf{A}^*\mathbf{A}\mathbf{x} = (\mathbf{A}\mathbf{x})^*(\mathbf{A}\mathbf{x}) = \|\mathbf{A}\mathbf{x}\|_2^2 \geq 0$, so $\mathbf{A}^*\mathbf{A}$ is positive semidefinite.)

Step 1. As a consequence of results presented in §3.2.1, we can construct n eigenpairs $\{(\lambda_j, \mathbf{v}_j)\}_{j=1}^n$ of $\mathbf{A}^*\mathbf{A}$ with unit eigenvectors ($\mathbf{v}_j^*\mathbf{v}_j = 1$) that are orthogonal to one another ($\mathbf{v}_j^*\mathbf{v}_k = 0$ when $j \neq k$). We are free to pick any convenient indexing for these eigenpairs; it will be convenient to label them so that the eigenvalues are decreasing in magnitude, $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0$.

Step 2. Define $\sigma_j = \|\mathbf{A}\mathbf{v}_j\|_2$.
Note that $\sigma_j^2 = \|\mathbf{A}\mathbf{v}_j\|_2^2 = \mathbf{v}_j^*\mathbf{A}^*\mathbf{A}\mathbf{v}_j = \lambda_j$. Since the eigenvalues $\lambda_1, \ldots, \lambda_n$ are decreasing in magnitude, so are the σ_j values: $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0$.

Step 3. Next, we will build a set of related orthonormal vectors in $\mathbb{C}^m$. Suppose we have already constructed such vectors $\mathbf{u}_1, \ldots, \mathbf{u}_{j-1}$.
If $\sigma_j \neq 0$, then define $\mathbf{u}_j = \sigma_j^{-1}\mathbf{A}\mathbf{v}_j$, so that $\|\mathbf{u}_j\|_2 = \sigma_j^{-1}\|\mathbf{A}\mathbf{v}_j\|_2 = 1$.
If $\sigma_j = 0$, then pick $\mathbf{u}_j$ to be any unit vector such that

$$\mathbf{u}_j \in \operatorname{span}\{\mathbf{u}_1, \ldots, \mathbf{u}_{j-1}\}^\perp;$$

i.e., ensure $\mathbf{u}_j^*\mathbf{u}_k = 0$ for all $k < j$.[†]

By construction, $\mathbf{u}_j^*\mathbf{u}_k = 0$ for $j \neq k$ if σ_j or σ_k is zero. If both σ_j and σ_k are nonzero, then

$$\mathbf{u}_j^*\mathbf{u}_k = \frac{1}{\sigma_j\sigma_k}(\mathbf{A}\mathbf{v}_j)^*(\mathbf{A}\mathbf{v}_k) = \frac{1}{\sigma_j\sigma_k}\mathbf{v}_j^*\mathbf{A}^*\mathbf{A}\mathbf{v}_k = \frac{\lambda_k}{\sigma_j\sigma_k}\mathbf{v}_j^*\mathbf{v}_k,$$

where we used the fact that $\mathbf{v}_j$ is an eigenvector of $\mathbf{A}^*\mathbf{A}$. Now if $j \neq k$, then $\mathbf{v}_j^*\mathbf{v}_k = 0$, and hence $\mathbf{u}_j^*\mathbf{u}_k = 0$. On the other hand, $j = k$ implies that $\mathbf{v}_j^*\mathbf{v}_k = 1$, so $\mathbf{u}_j^*\mathbf{u}_k = \lambda_j/\sigma_j^2 = 1$.

In conclusion, we have constructed a set of orthonormal vectors $\{\mathbf{u}_j\}_{j=1}^n$ with $\mathbf{u}_j \in \mathbb{C}^m$.

Step 4. For all $j = 1, \ldots, n$,

$$\mathbf{A}\mathbf{v}_j = \sigma_j\mathbf{u}_j,$$

regardless of whether $\sigma_j = 0$ or not. We can stack these n vector equations as columns of a single matrix equation,

$$\begin{bmatrix} | & | & & | \\ \mathbf{A}\mathbf{v}_1 & \mathbf{A}\mathbf{v}_2 & \cdots & \mathbf{A}\mathbf{v}_n \\ | & | & & | \end{bmatrix} = \begin{bmatrix} | & | & & | \\ \sigma_1\mathbf{u}_1 & \sigma_2\mathbf{u}_2 & \cdots & \sigma_n\mathbf{u}_n \\ | & | & & | \end{bmatrix}.$$

Note that both matrices in this equation can be factored into the product of simpler matrices:

$$\mathbf{A}\begin{bmatrix} | & | & & | \\ \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_n \\ | & | & & | \end{bmatrix} = \begin{bmatrix} | & | & & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_n \\ | & | & & | \end{bmatrix}\begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{bmatrix}.$$

[†] Note that $\sigma_j = 0$ implies that $\lambda_j = 0$, and so $\mathbf{A}^*\mathbf{A}$ has a zero eigenvalue; i.e., this matrix is singular. Recall from the last lecture that this case can only occur when $\mathbf{A}$ is rank-deficient: $\dim(\operatorname{Ran}(\mathbf{A})) < n$.

Denote these matrices as $\mathbf{AV} = \widehat{\mathbf{U}}\widehat{\mathbf{\Sigma}}$, where $\mathbf{A} \in \mathbb{C}^{m\times n}$, $\mathbf{V} \in \mathbb{C}^{n\times n}$, $\widehat{\mathbf{U}} \in \mathbb{C}^{m\times n}$, and $\widehat{\mathbf{\Sigma}} \in \mathbb{C}^{n\times n}$.

The (j,k) entry of $\mathbf{V}^*\mathbf{V}$ is simply $\mathbf{v}_j^*\mathbf{v}_k$, and so $\mathbf{V}^*\mathbf{V} = \mathbf{I}$. Since $\mathbf{V}$ is a square matrix, we have just proved that it is unitary. Hence, $\mathbf{V}\mathbf{V}^* = \mathbf{I}$ as well, and we conclude that

$$\mathbf{A} = \widehat{\mathbf{U}}\widehat{\mathbf{\Sigma}}\mathbf{V}^*.$$

This matrix factorization is known as the *reduced singular value decomposition*. It can be obtained via the MATLAB command

```
[Uhat, Sighat, V] = svd(A,0);
```

While the matrix $\widehat{\mathbf{U}}$ has orthonormal columns, *it is not a unitary matrix*. In particular, we have $\widehat{\mathbf{U}}^*\widehat{\mathbf{U}} = \mathbf{I} \in \mathbb{C}^{n\times n}$, but

$$\widehat{\mathbf{U}}\widehat{\mathbf{U}}^* \in \mathbb{C}^{m\times m}$$

cannot be the identity unless $m = n$. (To see this, note that $\widehat{\mathbf{U}}\widehat{\mathbf{U}}^*$ is an orthogonal projection onto $\text{Ran}(\widehat{\mathbf{U}}) = \text{span}\{\mathbf{u}_1, \ldots, \mathbf{u}_n\}$. Since $\dim(\text{Ran}(\widehat{\mathbf{U}})) = n$, this projection cannot equal the m-by-m identity matrix when $m > n$.)

Though $\widehat{\mathbf{U}}$ is not unitary, we might call it *subunitary*.[‡] We can construct $m - n$ additional column vectors to append to $\widehat{\mathbf{U}}$ to make it unitary. Here is the recipe: For $j = n+1, \ldots, m$, pick

$$\mathbf{u}_j \in \text{span}\{\mathbf{u}_1, \ldots, \mathbf{u}_{j-1}\}^\perp$$

with $\mathbf{u}_j^*\mathbf{u}_j = 1$. Then define

$$\mathbf{U} = \left[\begin{array}{cccc} | & | & & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_m \\ | & | & & | \end{array}\right].$$

It is simple to confirm that $\mathbf{U}^*\mathbf{U} = \mathbf{U}\mathbf{U}^* = \mathbf{I} \in \mathbb{C}^{m\times m}$, so $\mathbf{U}$ is unitary.

We wish to replace the $\widehat{\mathbf{U}}$ in the reduced SVD with the unitary matrix $\mathbf{U}$. To do so, we also need to replace $\widehat{\mathbf{\Sigma}}$ by some $\mathbf{\Sigma}$ in such a way that $\widehat{\mathbf{U}}\widehat{\mathbf{\Sigma}} = \mathbf{U}\mathbf{\Sigma}$. The simplest approach is to obtain $\mathbf{\Sigma}$ by appending zeros to the end of $\widehat{\mathbf{\Sigma}}$, thus ensuring there is no contribution when the new entries of $\mathbf{U}$ multiply against the new entries of $\mathbf{\Sigma}$:

$$\mathbf{\Sigma} = \left[\begin{array}{c} \widehat{\mathbf{\Sigma}} \\ \mathbf{0} \end{array}\right] \in \mathbb{C}^{m\times n}.$$

Finally, we are prepared to state our main result, the *full singular value decomposition*.

Theorem (Singular value decomposition). Any matrix $\mathbf{A} \in \mathbb{C}^{m\times n}$ can be written in the form

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*,$$

where $\mathbf{U} \in \mathbb{C}^{m\times m}$ and $\mathbf{V} \in \mathbb{C}^{n\times n}$ are unitary matrices and $\mathbf{\Sigma} \in \mathbb{C}^{m\times n}$ is zero everywhere except for entries on the main diagonal, where the (j,j) entry is σ_j for $j = 1, \ldots, \min(m,n)$, and

$$\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_{\min(m,n)} \ge 0.$$

[‡]There is no universally accepted term for such a matrix; Gilbert Strang suggests the descriptive term *subunitary*.

We have only proved this result for $m \geq n$. The proof for $m < n$ is obtained by applying the same arguments above to $\mathbf{A}^*$ in place of $\mathbf{A}$.

The full SVD is obtained via the MATLAB command

```
[U,S,V] = svd(A).
```