The Singular Value Decomposition

Gerald Farin, Dianne Hansford

December 13, 2005

1 The Geometry of the 2x2 Case

The inverse of a matrix A is another matrix A^{-1} such that $AA^{-1} = I$, the identity matrix. Not every matrix A has an inverse; rather, A must be square and invertible, see Section 14.5. In this section, we will explore how to generalize the concept of an inverse matrix. A surprisingly simple question — what set of orthogonal vectors is taken to another set of orthogonal vectors — will be the key to our investigation.

Let A be a 2×2 matrix, nonsingular for now. Let $\mathbf{u}_1$ and $\mathbf{u}_2$ be two unit vectors which are perpendicular to each other. This means that the matrix $U = [\mathbf{u}_1, \mathbf{u}_2]$ is *orthonormal*: $U^{-1} = U^{\mathrm{T}}$. In general, A will not map two orthonormal vectors $\mathbf{u}_1, \mathbf{u}_2$ to two orthonormal image vectors $\mathbf{v}_1, \mathbf{v}_2$. But under some conditions A will map $\mathbf{u}_1$ and $\mathbf{u}_2$ to two *orthogonal* image vectors $\lambda_1\mathbf{v}_1, \lambda_2\mathbf{v}_2$.

For that to happen, we need to have $(A\mathbf{u}_i)^{\mathrm{T}}(A\mathbf{u}_i) = \lambda_i$ for some λ_i and $(A\mathbf{u}_i)^{\mathrm{T}}(A\mathbf{u}_j) = 0$ for $i \neq j$. In matrix form:

$$(AU)^{\mathrm{T}}(AU) = \Lambda, \tag{1}$$

with Λ being the diagonal matrix with λ_1 and λ_2 on the diagonal.

This leads to

$$U^{\mathrm{T}}A^{\mathrm{T}}AU = \Lambda$$

and thus to

$$A^{\mathrm{T}}A = U\Lambda U^{\mathrm{T}}.$$

Since $A^{\mathrm{T}}A$ is a symmetric matrix, this means that λ_1, λ_2 are its eigenvalues and $\mathbf{u}_1, \mathbf{u}_2$ are the corresponding eigenvectors.

Thus our question above is answered as follows: if a matrix A maps two orthonormal vectors $\mathbf{u}_1, \mathbf{u}_2$ to two orthogonal vectors $\lambda_1\mathbf{v}_1, \lambda_2\mathbf{v}_2$, then $\mathbf{u}_1, \mathbf{u}_2$ have to be eigenvectors of $A^{\mathrm{T}}A$.

We formalize this as

$$AU = V\Sigma \tag{2}$$

with an orthonormal matrix $V = [\mathbf{v}_1, \mathbf{v}_2]$ and a (so far unknown) diagonal matrix Σ. We conclude

$$A = V\Sigma U^{-1} = V\Sigma U^{\mathrm{T}}. \tag{3}$$

This is the *singular value decomposition*, SVD for short, of A.

What is the relationship between Σ and Λ? We insert (2) into (1) and obtain

$$(V\Sigma)^{\mathrm{T}}(V\Sigma) = \Lambda,$$

hence

$$\Sigma^2 = \Lambda.$$

The diagonal elements of Σ, formed by the square roots of the eigenvalues of $A^{\mathrm{T}}A$, are called the *singular values* of A.

It should be clear that our development is not limited to 2×2 matrices. Thus (5) holds for any nonsingular square matrix A.

Let us summarize: for any square nonsingular matrix, there exist orthonormal matrices U and V and a diagonal matrix Σ such that A may be decomposed as shown in (5).

2 The General Case

In our SVD development, we assumed that A was square and invertible. This had the effect that $A^{\mathrm{T}}A$ had nonzero eigenvalues and well-defined eigenvectors. However, everything still works if A is not square! From now on, A will be a rectangular matrix with m rows and n columns. We also assume that $m \geq n$, meaning that A may have more rows than columns.

Let $\mathbf{u}_1, \ldots, \mathbf{u}_n$ a set of orthonormal vectors in A's domain. In general, they will not be mapped to n orthogonal vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ in A's range.[1] But some set of $\mathbf{u}_i$ *will* be mapped to a set of orthogonal vectors $\lambda_i \mathbf{v}_i$.

Exactly as in the preceding section, this leads to

$$A^{\mathrm{T}}A = U\Lambda U^{\mathrm{T}}.$$

Since $A^{\mathrm{T}}A$ is a symmetric matrix, this means that $\lambda_1, \ldots, \lambda_n$ are its eigenvalues and $\mathbf{u}_1, \ldots, \mathbf{u}_n$ are the corresponding eigenvectors.

Proceeding exactly as before, we get

$$AU = V\Sigma \tag{4}$$

[1] Note that the $\mathbf{u}_i$ each have n elements whereas the $\mathbf{v}_i$ each have m elements.

with an orthonormal matrix $V = [\mathbf{v}_1, \ldots, \mathbf{v}_n]$ and a (so far unknown) diagonal matrix Σ. We conclude

$$A = V\Sigma U^{-1} = V\Sigma U^{\mathrm{T}}. \tag{5}$$

This is the *singular value decomposition*, SVD for short, of the rectangular matrix A. As before, the σ_i are given by $\sigma_i = \sqrt{\lambda_i}$. The matrix dimensions are illustrated in Fig. 1.

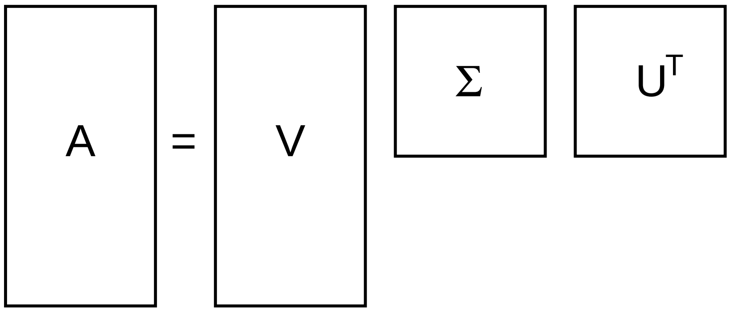

Figure 1: The SVD matrix dimensions.

An example should illuminate.

Let A and $A^{\mathrm{T}}A$ be given by

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 0 & 1 \end{bmatrix}, \quad A^{\mathrm{T}}A = \begin{bmatrix} 1 & 0 \\ 0 & 5 \end{bmatrix}.$$

We find

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{5} \end{bmatrix}, \quad U = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad V = \begin{bmatrix} 1 & 0 \\ 0 & \frac{2}{\sqrt{5}} \\ 0 & \frac{1}{\sqrt{5}} \end{bmatrix}.$$

It is now trivial to verify $AU = V\Sigma$.

In practice, the matrix Σ is arranged such that the σ_i are listed in decreasing order on its diagonal. This will necessitate permutations on the entries of V.

3 The Algorithm

Here is the SVD algorithm:
Given: an $m \times n$ matrix A.

Find: U, V, Σ such that $AU = V\Sigma$.
Do:
Find the eigenvalues $\lambda_1, \ldots, \lambda_n$ of $A^{\mathrm{T}}A$.
Create a diagonal matrix Σ having diagonal entries $\sqrt{\lambda_i}$.
Find the corresponding (normalized) eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_n$.
Create a matrix U having the $\mathbf{u}_i$ as columns.
Define $V = AU$ and normalize the columns of V.
You have now found the SVD $A = V\Sigma U^{\mathrm{T}}$.

The only "hard" task in this is finding the λ_i. But there are several highly efficient algorithms for this task, taking advantage of the fact that $A^{\mathrm{T}}A$ is symmetric. Many of these algorithms will return the corresponding eigenvectors as well.

4 Linear Systems

A linear system is given by

$$A\mathbf{x} = \mathbf{b}$$

where A may be a square matrix or a rectangular matrix. If we have the SVD

$$A = U\Sigma V^{\mathrm{T}},$$

we immediately get

$$\mathbf{x} = V\Sigma^{-1}U^{\mathrm{T}}\mathbf{b}.$$

Note we need no matrix inversions! Since Σ^{-1} is just a diagonal matrix, the elements of Σ^{-1} are simply the $1/\sigma_i$.

A word of caution: Σ may contain zeroes on its diagonal, or σ_i very close to zero. For those σ_i, one sets $\sigma_i^{-1} = 0$.

The SVD may be used to solve overdetermined linear systems. Using the matrix A from Section 2, we create a linear system as follows:

$$\begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}.$$

An approximate solution is given by using

$$\mathbf{v} = U\Sigma^{-1}V^{\mathrm{T}}\mathbf{b} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2/\sqrt{5} & 1/\sqrt{5} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 2/5 \end{bmatrix}$$

Note that this approximate solution to our overdetermined linear system is identical to that of the *normal equations*

$$A^{\mathrm{T}}A\mathbf{v} = A^{\mathrm{T}}\mathbf{b}.$$

---

5 The Pseudoinverse

What is the inverse of a nonsquare matrix? Using SVD, the answer is straightforward. We denote the "inverse" of a nonsquare matrix by $A^\dagger$. Starting from

$$A = V\Sigma U^{\mathrm{T}},$$

we immediately have

$$A^\dagger = U\Sigma^{-1}V^{\mathrm{T}}.$$

If A is square and invertible, then $A^\dagger = A^{-1}$. Otherwise, $A^\dagger$ still has some properties of an inverse:

$$A^\dagger A A^\dagger = A^\dagger,$$
$$AA^\dagger A = A,$$