$\ldots x_1x_2 + 3x_1x_3 + 30x_1x_4 + 30x_2x_3 + 3x_2x_4 + x_3x_4$

$\ldots x_1x_2 + 5x_1x_3 + 7x_1x_4 + 7x_2x_3 + 5x_2x_4 + 3x_3x_4$

$\ldots 4x_1^2 - 6x_1x_2 - 10x_1x_3 - 10x_1x_4 - 6x_2x_3 - 6x_2x_4 - 2x_3x_4$

17. $-6x_1^2 - 10x_2^2 - 13x_3^2 - 13x_4^2 - 4x_1x_2 - 4x_1x_3 - 4x_1x_4 + 6x_3x_4$

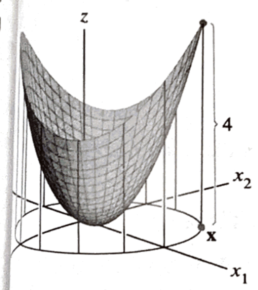

…maximum value of $Q(\mathbf{x})$ …ject to $\mathbf{x}^T\mathbf{x} = 1$ is 4.

SOLUTIONS TO PRACTICE PROBLEMS

1. The matrix of the quadratic form is $A = \begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix}$. It is easy to find the eigenvalues, 4 and 2, and corresponding unit eigenvectors, $\begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$ and $\begin{bmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$. So the desired change of variable is $\mathbf{x} = P\mathbf{y}$, where $P = \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}$. (A common error here is to forget to normalize the eigenvectors.) The new quadratic form is $\mathbf{y}^TD\mathbf{y} = 4y_1^2 + 2y_2^2$.
2. The maximum of $Q(\mathbf{x})$ for $\mathbf{x}$ a unit vector is 4, and the maximum is attained at the unit eigenvector $\begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$. [A common incorrect answer is $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$. This vector maximizes the quadratic form $\mathbf{y}^TD\mathbf{y}$ instead of $Q(\mathbf{x})$.]

7.4 THE SINGULAR VALUE DECOMPOSITION

The diagonalization theorems in Sections 5.3 and 7.1 play a part in many interesting applications. Unfortunately, as we know, not all matrices can be factored as $A = PDP^{-1}$ with D diagonal. However, a factorization $A = QDP^{-1}$ *is* possible for any $m \times n$ matrix A! A special factorization of this type, called the *singular value decomposition*, is one of the most useful matrix factorizations in applied linear algebra.

The singular value decomposition is based on the following property of the ordinary diagonalization that can be imitated for rectangular matrices: The absolute values of the eigenvalues of a symmetric matrix A measure the amounts that A stretches or shrinks certain vectors (the eigenvectors). If $A\mathbf{x} = \lambda\mathbf{x}$ and $\|\mathbf{x}\| = 1$, then

$$\|A\mathbf{x}\| = \|\lambda\mathbf{x}\| = |\lambda|\,\|\mathbf{x}\| = |\lambda| \tag{1}$$

If λ_1 is the eigenvalue with the greatest magnitude, then a corresponding unit eigenvector $\mathbf{v}_1$ identifies a direction in which the stretching effect of A is greatest. That is, the length of $A\mathbf{x}$ is maximized when $\mathbf{x} = \mathbf{v}_1$, and $\|A\mathbf{v}_1\| = |\lambda_1|$, by (1). This description of $\mathbf{v}_1$ and $|\lambda_1|$ has an analogue for rectangular matrices that will lead to the singular value decomposition.

EXAMPLE 1 If $A = \begin{bmatrix} 4 & 11 & 14 \\ 8 & 7 & -2 \end{bmatrix}$, then the linear transformation $\mathbf{x} \mapsto A\mathbf{x}$ maps the unit sphere $\{\mathbf{x} : \|\mathbf{x}\| = 1\}$ in $\mathbb{R}^3$ onto an ellipse in $\mathbb{R}^2$, shown in Fig. 1. Find a unit vector $\mathbf{x}$ at which the length $\|A\mathbf{x}\|$ is maximized, and compute this maximum length.

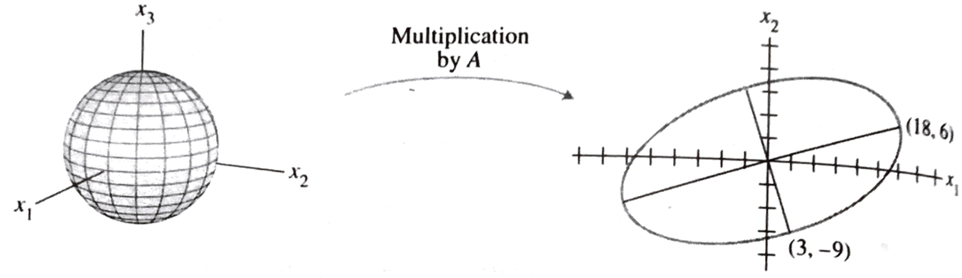

FIGURE 1 A transformation from $\mathbb{R}^3$ to $\mathbb{R}^2$.

Solution The quantity $\|A\mathbf{x}\|^2$ is maximized at the same $\mathbf{x}$ that maximizes $\|A\mathbf{x}\|$, and $\|A\mathbf{x}\|^2$ is easier to study. Observe that

$$\|A\mathbf{x}\|^2 = (A\mathbf{x})^T(A\mathbf{x}) = \mathbf{x}^TA^TA\mathbf{x} = \mathbf{x}^T(A^TA)\mathbf{x}$$

Also, A^TA is a symmetric matrix, since $(A^TA)^T = A^TA^{TT} = A^TA$. So the problem now is to maximize the quadratic form $\mathbf{x}^T(A^TA)\mathbf{x}$ subject to the constraint $\|\mathbf{x}\| = 1$. That's a problem from Section 7.3, and we know the solution. By Theorem 6, the maximum value is the greatest eigenvalue λ_1 of A^TA. Also, the maximum value is attained at a unit eigenvector of A^TA corresponding to λ_1.

For the matrix A in this example,

$$A^TA = \begin{bmatrix} 4 & 8 \\ 11 & 7 \\ 14 & -2 \end{bmatrix}\begin{bmatrix} 4 & 11 & 14 \\ 8 & 7 & -2 \end{bmatrix} = \begin{bmatrix} 80 & 100 & 40 \\ 100 & 170 & 140 \\ 40 & 140 & 200 \end{bmatrix}$$

The eigenvalues of A^TA are $\lambda_1 = 360$, $\lambda_2 = 90$, and $\lambda_3 = 0$. Corresponding unit eigenvectors are, respectively,

$$\mathbf{v}_1 = \begin{bmatrix} 1/3 \\ 2/3 \\ 2/3 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} -2/3 \\ -1/3 \\ 2/3 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 2/3 \\ -2/3 \\ 1/3 \end{bmatrix}$$

The maximum value of $\|A\mathbf{x}\|^2$ is 360, attained when $\mathbf{x}$ is the unit vector $\mathbf{v}_1$. The vector $A\mathbf{v}_1$ is a point on the ellipse in Fig. 1 farthest from the origin, namely,

$$A\mathbf{v}_1 = \begin{bmatrix} 4 & 11 & 14 \\ 8 & 7 & -2 \end{bmatrix}\begin{bmatrix} 1/3 \\ 2/3 \\ 2/3 \end{bmatrix} = \begin{bmatrix} 18 \\ 6 \end{bmatrix}$$

For $\|\mathbf{x}\| = 1$, the maximum value of $\|A\mathbf{x}\|$ is $\|A\mathbf{v}_1\| = \sqrt{360} = 6\sqrt{10}$. ■

Example 1 suggests that the effect of A on the unit sphere in $\mathbb{R}^3$ is related to the quadratic form $\mathbf{x}^T(A^TA)\mathbf{x}$. In fact, the entire geometric behavior of the transformation $\mathbf{x} \mapsto A\mathbf{x}$ is captured by this quadratic form, as we shall see.

The Singular Values of an $m \times n$ Matrix

Let A be an $m \times n$ matrix. Then A^TA is symmetric and can be orthogonally diagonalized. Let $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be an orthonormal basis for $\mathbb{R}^n$ consisting of eigenvectors of A^TA, and let $\lambda_1, \dots, \lambda_n$ be the associated eigenvalues of A^TA. Then, for $1 \le i \le n$,

$$\begin{aligned} \|A\mathbf{v}_i\|^2 &= (A\mathbf{v}_i)^T A\mathbf{v}_i = \mathbf{v}_i^T A^T A\mathbf{v}_i \\ &= \mathbf{v}_i^T(\lambda_i \mathbf{v}_i) && \text{Since } \mathbf{v}_i \text{ is an eigenvector of } A^TA \\ &= \lambda_i && \text{Since } \mathbf{v}_i \text{ is a unit vector} \end{aligned} \tag{2}$$

So the eigenvalues of A^TA are all nonnegative. By renumbering, if necessary, we may assume that the eigenvalues are arranged so that

$$\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_n \ge 0$$

The **singular values** of A are the square roots of the eigenvalues of A^TA, denoted by $\sigma_1, \dots, \sigma_n$, and they are arranged in decreasing order. That is, $\sigma_i = \sqrt{\lambda_i}$ for $1 \le i \le n$. By (2), *the singular values of A are the lengths of the vectors* $A\mathbf{v}_1, \dots, A\mathbf{v}_n$.

EXAMPLE 2 Let A be the matrix in Example 1. Since the eigenvalues of A^TA are 360, 90, and 0, the singular values of A are

$$\sigma_1 = \sqrt{360} = 6\sqrt{10}, \quad \sigma_2 = \sqrt{90} = 3\sqrt{10}, \quad \sigma_3 = 0$$

From Example 1, the first singular value of A is the maximum of $\|A\mathbf{x}\|$ over all unit vectors, and the maximum is attained at the unit eigenvector $\mathbf{v}_1$. Theorem 7 in Section 7.3 shows that the second singular value of A is the maximum of $\|A\mathbf{x}\|$ over all unit vectors that are *orthogonal to* $\mathbf{v}_1$, and this maximum is attained at the second unit eigenvector, $\mathbf{v}_2$ (Exercise 22). For the $\mathbf{v}_2$ in Example 1,

$$A\mathbf{v}_2 = \begin{bmatrix} 4 & 11 & 14 \\ 8 & 7 & -2 \end{bmatrix} \begin{bmatrix} -2/3 \\ -1/3 \\ 2/3 \end{bmatrix} = \begin{bmatrix} 3 \\ -9 \end{bmatrix}$$

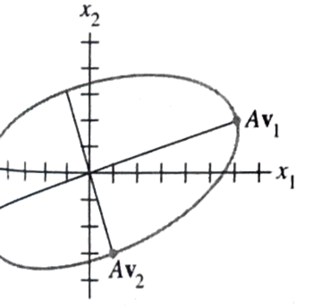

RE 2

This point is on the minor axis of the ellipse in Fig. 1, just as $A\mathbf{v}_1$ is on the major axis. (See Fig. 2.) The first two singular values of A are the lengths of the major and minor semiaxes of the ellipse.

The fact that $A\mathbf{v}_1$ and $A\mathbf{v}_2$ are orthogonal in Fig. 2 is no accident, as the next theorem shows.

THEOREM 9

Suppose $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ is an orthonormal basis of $\mathbb{R}^n$ consisting of eigenvectors of A^TA, arranged so that the corresponding eigenvalues of A^TA satisfy $\lambda_1 \ge \cdots \ge \lambda_n$, and suppose A has r nonzero singular values. Then $\{A\mathbf{v}_1, \dots, A\mathbf{v}_r\}$ is an orthogonal basis for Col A, and rank $A = r$.

PROOF Because $\mathbf{v}_i$ and $\lambda_j \mathbf{v}_j$ are orthogonal for $i \neq j$,

$$(A\mathbf{v}_i)^T (A\mathbf{v}_j) = \mathbf{v}_i^T A^T A \mathbf{v}_j = \mathbf{v}_i^T (\lambda_j \mathbf{v}_j) = 0$$

Thus $\{A\mathbf{v}_1, \dots, A\mathbf{v}_n\}$ is an orthogonal set. Furthermore, since the lengths of the vectors $A\mathbf{v}_1, \dots, A\mathbf{v}_n$ are the singular values of A, and since there are r nonzero singular values, $A\mathbf{v}_i \neq \mathbf{0}$ if and only if $1 \leq i \leq r$. So $A\mathbf{v}_1, \dots, A\mathbf{v}_r$ are linearly independent vectors, and they are in Col A. Finally, for any $\mathbf{y}$ in Col A—say, $\mathbf{y} = A\mathbf{x}$—we can write $\mathbf{x} = c_1\mathbf{v}_1 + \cdots + c_n\mathbf{v}_n$, and

$$\begin{aligned} \mathbf{y} = A\mathbf{x} &= c_1 A\mathbf{v}_1 + \cdots + c_r A\mathbf{v}_r + c_{r+1} A\mathbf{v}_{r+1} + \cdots + c_n A\mathbf{v}_n \\ &= c_1 A\mathbf{v}_1 + \cdots + c_r A\mathbf{v}_r + 0 + \cdots + 0 \end{aligned}$$

Thus $\mathbf{y}$ is in Span $\{A\mathbf{v}_1, \dots, A\mathbf{v}_r\}$, which shows that $\{A\mathbf{v}_1, \dots, A\mathbf{v}_r\}$ is an (orthogonal) basis for Col A. Hence rank $A = \dim \text{Col } A = r$. ∎

NUMERICAL NOTE

In some cases, the rank of A may be very sensitive to small changes in the entries of A. The obvious method of counting the number of pivot columns in A does not work well if A is row reduced by a computer. Roundoff error often creates an echelon form with full rank.

In practice, the most reliable way to estimate the rank of a large matrix A is to count the number of nonzero singular values. In this case, extremely small nonzero singular values are assumed to be zero for all practical purposes, and the *effective rank* of the matrix is the number obtained by counting the remaining nonzero singular values.[1]

The Singular Value Decomposition

The decomposition of A involves an $m \times n$ "diagonal" matrix Σ of the form

$$\Sigma = \begin{bmatrix} D & 0 \\ 0 & 0 \end{bmatrix} \begin{matrix} \\ \leftarrow m - r \text{ rows} \end{matrix} \qquad (3)$$
$$\uparrow n - r \text{ columns}$$

where D is an $r \times r$ diagonal matrix for some r not exceeding the smaller of m and n. (If r equals m or n or both, some or all of the zero matrices do not appear.)

[1] In general, rank estimation is not a simple problem. For a discussion of the subtle issues involved, see Philip E. Gill, Walter Murray, and Margaret H. Wright, *Numerical Linear Algebra and Optimization*, vol. 1 (Redwood City, CA: Addison-Wesley, 1991), Sec. 5.8.

THEOREM 10 **The Singular Value Decomposition**

Let A be an $m \times n$ matrix with rank r. Then there exists an $m \times n$ matrix Σ as in (3) for which the diagonal entries in D are the first r singular values of A, $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$, and there exist an $m \times m$ orthogonal matrix U and an $n \times n$ orthogonal matrix V such that

$$A = U\Sigma V^T$$

Any factorization $A = U\Sigma V^T$, with U and V orthogonal, Σ as in (3), and positive diagonal entries in D, is called a **singular value decomposition** (or **SVD**) of A. The matrices U and V are not uniquely determined by A, but the diagonal entries of Σ are necessarily the singular values of A. See Exercise 19. The columns of U in such a decomposition are called **left singular vectors** of A, and the columns of V are called **right singular vectors** of A.

PROOF Let λ_i and $\mathbf{v}_i$ be as in Theorem 9, so that $\{A\mathbf{v}_1, \ldots, A\mathbf{v}_r\}$ is an orthogonal basis for Col A. Normalize each $A\mathbf{v}_i$ to obtain an orthonormal basis $\{\mathbf{u}_1, \ldots, \mathbf{u}_r\}$, where

$$\mathbf{u}_i = \frac{1}{\|A\mathbf{v}_i\|}A\mathbf{v}_i = \frac{1}{\sigma_i}A\mathbf{v}_i$$

and

$$A\mathbf{v}_i = \sigma_i\mathbf{u}_i \quad (1 \leq i \leq r) \tag{4}$$

Now extend $\{\mathbf{u}_1, \ldots, \mathbf{u}_r\}$ to an orthonormal basis $\{\mathbf{u}_1, \ldots, \mathbf{u}_m\}$ of $\mathbb{R}^m$, and let

$$U = [\,\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_m\,] \quad \text{and} \quad V = [\,\mathbf{v}_1 \quad \mathbf{v}_2 \quad \cdots \quad \mathbf{v}_n\,]$$

By construction, U and V are orthogonal matrices. Also, from (4),

$$AV = [\,A\mathbf{v}_1 \quad \cdots \quad A\mathbf{v}_r \quad \mathbf{0} \quad \cdots \quad \mathbf{0}\,] = [\,\sigma_1\mathbf{u}_1 \quad \cdots \quad \sigma_r\mathbf{u}_r \quad \mathbf{0} \quad \cdots \quad \mathbf{0}\,]$$

Let D be the diagonal matrix with diagonal entries $\sigma_1, \ldots, \sigma_r$, and let Σ be as in (3) above. Then

$$U\Sigma = [\,\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_m\,]\left[\begin{array}{cccc|c} \sigma_1 & & & 0 & \\ & \sigma_2 & & & 0 \\ & & \ddots & & \\ 0 & & & \sigma_r & \\ \hline & & 0 & & 0 \end{array}\right]$$

$$= [\,\sigma_1\mathbf{u}_1 \quad \cdots \quad \sigma_r\mathbf{u}_r \quad \mathbf{0} \quad \cdots \quad \mathbf{0}\,]$$
$$= AV$$

Since V is an orthogonal matrix, $U\Sigma V^T = AVV^T = A$. ∎

The next two examples focus attention on the internal structure of a singular value decomposition. An efficient and numerically stable algorithm for this decomposition would use a different approach. See the Numerical Note at the end of the section.

EXAMPLE 3 Use the results of Examples 1 and 2 to construct a singular value decomposition of $A = \begin{bmatrix} 4 & 11 & 14 \\ 8 & 7 & -2 \end{bmatrix}$.

Solution A construction can be divided into three steps.

Step 1. ***Find an orthogonal diagonalization of A^TA.*** That is, find the eigenvalues of A^TA and a corresponding orthonormal set of eigenvectors. If A had only two columns, the calculations could be done by hand. Larger matrices usually require a matrix program.[2] However, for the matrix A here, the eigendata for A^TA are provided by Example 1.

Step 2. ***Set up V and Σ.*** Arrange the eigenvalues of A^TA in decreasing order. In Example 1, the eigenvalues are already listed in decreasing order: 360, 90, and 0. The corresponding unit eigenvectors, $\mathbf{v}_1$, $\mathbf{v}_2$, and $\mathbf{v}_3$, are the right singular vectors of A. Using Example 1, construct

$$V = [\,\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3\,] = \begin{bmatrix} 1/3 & -2/3 & 2/3 \\ 2/3 & -1/3 & -2/3 \\ 2/3 & 2/3 & 1/3 \end{bmatrix}$$

The square roots of the eigenvalues are the singular values:

$$\sigma_1 = 6\sqrt{10}, \quad \sigma_2 = 3\sqrt{10}, \quad \sigma_3 = 0$$

The nonzero singular values are the diagonal entries of D. The matrix Σ is the same size as A, with D in its upper-left corner and with 0's elsewhere.

$$D = \begin{bmatrix} 6\sqrt{10} & 0 \\ 0 & 3\sqrt{10} \end{bmatrix}, \qquad \Sigma = [\,D \quad 0\,] = \begin{bmatrix} 6\sqrt{10} & 0 & 0 \\ 0 & 3\sqrt{10} & 0 \end{bmatrix}$$

Step 3. ***Construct U.*** When A has rank r, the first r columns of U are the normalized vectors obtained from $A\mathbf{v}_1, \dots, A\mathbf{v}_r$. In this example, A has two nonzero singular values, so rank $A = 2$. Recall from equation (2) and the paragraph before Example 2 that $\|A\mathbf{v}_1\| = \sigma_1$ and $\|A\mathbf{v}_2\| = \sigma_2$. Thus

$$\mathbf{u}_1 = \frac{1}{\sigma_1}A\mathbf{v}_1 = \frac{1}{6\sqrt{10}}\begin{bmatrix} 18 \\ 6 \end{bmatrix} = \begin{bmatrix} 3/\sqrt{10} \\ 1/\sqrt{10} \end{bmatrix}$$

$$\mathbf{u}_2 = \frac{1}{\sigma_2}A\mathbf{v}_2 = \frac{1}{3\sqrt{10}}\begin{bmatrix} 3 \\ -9 \end{bmatrix} = \begin{bmatrix} 1/\sqrt{10} \\ -3/\sqrt{10} \end{bmatrix}$$

Note that $\{\mathbf{u}_1, \mathbf{u}_2\}$ is already a basis for $\mathbb{R}^2$. Thus no additional vectors are needed for U, and $U = [\,\mathbf{u}_1 \quad \mathbf{u}_2\,]$. The singular value decomposition of A is

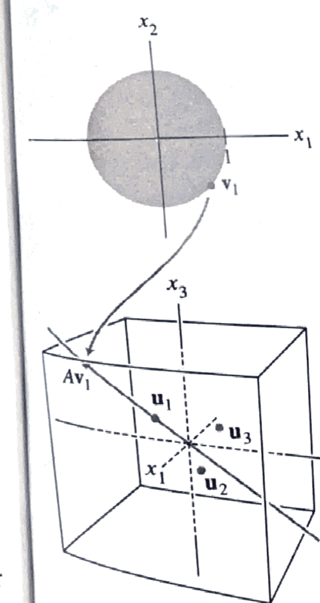

FIGURE 3

---

[2] See the *Study Guide* for software and graphing calculator commands. MATLAB, for instance, can produce both the eigenvalues and the eigenvectors with one command, `eig`.

$$A = \underbrace{\begin{bmatrix} 3/\sqrt{10} & 1/\sqrt{10} \\ 1/\sqrt{10} & -3/\sqrt{10} \end{bmatrix}}_{U} \underbrace{\begin{bmatrix} 6\sqrt{10} & 0 & 0 \\ 0 & 3\sqrt{10} & 0 \end{bmatrix}}_{\Sigma} \underbrace{\begin{bmatrix} 1/3 & 2/3 & 2/3 \\ -2/3 & -1/3 & 2/3 \\ 2/3 & -2/3 & 1/3 \end{bmatrix}}_{V^T}$$

EXAMPLE 4 Find a singular value decomposition of $A = \begin{bmatrix} 1 & -1 \\ -2 & 2 \\ 2 & -2 \end{bmatrix}$.

Solution First, compute $A^TA = \begin{bmatrix} 9 & -9 \\ -9 & 9 \end{bmatrix}$. The eigenvalues of A^TA are 18 and 0, with corresponding unit eigenvectors

$$\mathbf{v}_1 = \begin{bmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$$

These unit vectors form the columns of V:

$$V = [\,\mathbf{v}_1 \quad \mathbf{v}_2\,] = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}$$

The singular values are $\sigma_1 = \sqrt{18} = 3\sqrt{2}$ and $\sigma_2 = 0$. Since there is only one nonzero singular value, the "matrix" D may be written as a single number. That is, $D = 3\sqrt{2}$. The matrix Σ is the same size as A, with D in its upper-left corner:

$$\Sigma = \begin{bmatrix} D & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 3\sqrt{2} & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

To construct U, first construct $A\mathbf{v}_1$ and $A\mathbf{v}_2$:

$$A\mathbf{v}_1 = \begin{bmatrix} 2/\sqrt{2} \\ -4/\sqrt{2} \\ 4/\sqrt{2} \end{bmatrix}, \quad A\mathbf{v}_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

As a check on the calculations, verify that $\|A\mathbf{v}_1\| = \sigma_1 = 3\sqrt{2}$. Of course, $A\mathbf{v}_2 = \mathbf{0}$ because $\|A\mathbf{v}_2\| = \sigma_2 = 0$. The only column found for U so far is

$$\mathbf{u}_1 = \frac{1}{3\sqrt{2}} A\mathbf{v}_1 = \begin{bmatrix} 1/3 \\ -2/3 \\ 2/3 \end{bmatrix}$$

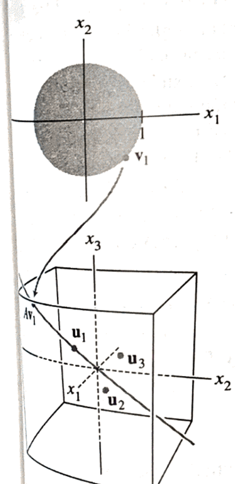

FIGURE 3

The other columns of U are found by extending the set $\{\mathbf{u}_1\}$ to an orthonormal basis for $\mathbb{R}^3$. In this case, we need two orthogonal unit vectors $\mathbf{u}_2$ and $\mathbf{u}_3$ that are orthogonal to $\mathbf{u}_1$. (See Fig. 3.) Each vector must satisfy $\mathbf{u}_1^T\mathbf{x} = 0$, which is equivalent to the equation $x_1 - 2x_2 + 2x_3 = 0$. A basis for the solution set of this equation is

$$\mathbf{w}_1 = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}, \quad \mathbf{w}_2 = \begin{bmatrix} -2 \\ 0 \\ 1 \end{bmatrix}$$

(Check that $\mathbf{w}_1$ and $\mathbf{w}_2$ are each orthogonal to $\mathbf{u}_1$.) Apply the Gram–Schmidt process (with normalizations) to $\{\mathbf{w}_1, \mathbf{w}_2\}$, and obtain

$$\mathbf{u}_2 = \begin{bmatrix} 2/\sqrt{5} \\ 1/\sqrt{5} \\ 0 \end{bmatrix}, \quad \mathbf{u}_3 = \begin{bmatrix} -2/\sqrt{45} \\ 4/\sqrt{45} \\ 5/\sqrt{45} \end{bmatrix}$$

Finally, set $U = [\,\mathbf{u}_1 \quad \mathbf{u}_2 \quad \mathbf{u}_3\,]$, take Σ and V^T from above, and write

$$A = \begin{bmatrix} 1 & -1 \\ -2 & 2 \\ 2 & -2 \end{bmatrix} = \begin{bmatrix} 1/3 & 2/\sqrt{5} & -2/\sqrt{45} \\ -2/3 & 1/\sqrt{5} & 4/\sqrt{45} \\ 2/3 & 0 & 5/\sqrt{45} \end{bmatrix} \begin{bmatrix} 3\sqrt{2} & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}$$

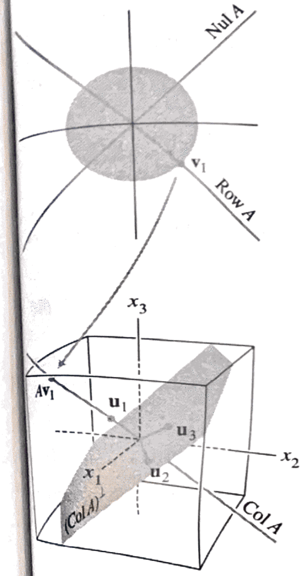

The fundamental subspaces in Example 4.

Applications of the Singular Value Decomposition

The SVD is often used to estimate the rank of a matrix, as noted above. Several other numerical applications are described briefly below, and an application to image processing is presented in Section 7.5.

EXAMPLE 5 (*The Condition Number*) Most numerical calculations involving an equation $A\mathbf{x} = \mathbf{b}$ are as reliable as possible when the SVD of A is used. The two orthogonal matrices U and V do not affect lengths of vectors or angles between vectors (Theorem 7 in Section 6.2). Any possible instabilities in numerical calculations are identified in Σ. If the singular values of A are extremely large or small, roundoff errors are almost inevitable, but an error analysis is aided by knowing the entries in Σ and V.

If A is an invertible $n \times n$ matrix, then the ratio σ_1/σ_n of the largest and smallest singular values gives the **condition number** of A. Exercises 41–43 in Section 2.3 showed how the condition number affects the sensitivity of a solution of $A\mathbf{x} = \mathbf{b}$ to changes (or errors) in the entries of A. (Actually, a "condition number" of A can be computed in several ways, but the definition given here is widely used for studying $A\mathbf{x} = \mathbf{b}$.)

EXAMPLE 6 (*Bases for Fundamental Subspaces*) Given an SVD for an $m \times n$ matrix A, let $\mathbf{u}_1, \ldots, \mathbf{u}_m$ be the left singular vectors, $\mathbf{v}_1, \ldots, \mathbf{v}_n$ the right singular vectors, and $\sigma_1, \ldots, \sigma_n$ the singular values, and let r be the rank of A. By Theorem 9,

$$\{\mathbf{u}_1, \ldots, \mathbf{u}_r\} \tag{5}$$

is an orthonormal basis for Col A.

Recall from Theorem 3 in Section 6.1 that $(\text{Col } A)^\perp = \text{Nul } A^T$. Hence

$$\{\mathbf{u}_{r+1}, \ldots, \mathbf{u}_m\} \tag{6}$$

is an orthonormal basis for Nul A^T.

Since $\|A\mathbf{v}_i\| = \sigma_i$ for $1 \le i \le n$, and σ_i is 0 if and only if $i > r$, the vectors $\mathbf{v}_{r+1}, \ldots, \mathbf{v}_n$ span a subspace of Nul A of dimension $n - r$. By the Rank Theorem,

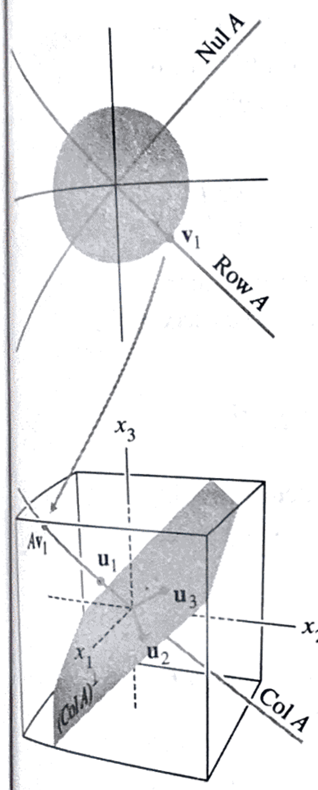

The fundamental subspaces in Example 4.

dim Nul $A = n - \text{rank } A$. It follows that

$$\{\mathbf{v}_{r+1}, \ldots, \mathbf{v}_n\} \tag{7}$$

is an orthonormal basis for Nul A, by the Basis Theorem (in Section 4.5).

From (5) and (6), the orthogonal complement of Nul A^T is Col A. Interchanging A and A^T, we have $(\text{Nul } A)^\perp = \text{Col } A^T = \text{Row } A$. Hence, from (7),

$$\{\mathbf{v}_1, \ldots, \mathbf{v}_r\} \tag{8}$$

is an orthonormal basis for Row A.

Figure 4 summarizes (5)–(8), but shows the orthogonal basis $\{\sigma_1\mathbf{u}_1, \ldots, \sigma_r\mathbf{u}_r\}$ for Col A instead of the normalized basis, to remind you that $A\mathbf{v}_i = \sigma_i\mathbf{u}_i$ for $1 \le i \le r$. Explicit orthonormal bases for the four fundamental subspaces determined by A are useful in some calculations, particularly in constrained optimization problems. ■

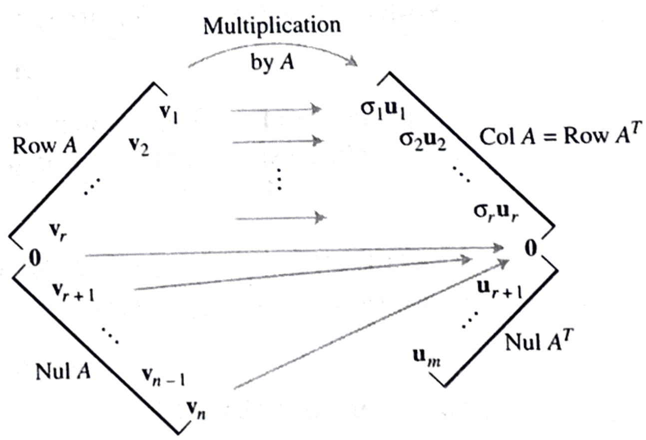

FIGURE 4 The four fundamental subspaces and the action of A.

The four fundamental subspaces and the concept of singular values provide the final statements of the Invertible Matrix Theorem. (Recall that statements about A^T have been omitted from the theorem, to avoid nearly doubling the number of statements.) The other statements were given in Sections 2.3, 2.9, 3.2, 4.6, and 5.2.

THEOREM

The Invertible Matrix Theorem (concluded)

Let A be an $n \times n$ matrix. Then the following statements are each equivalent to the statement that A is an invertible matrix.

u. $(\text{Col } A)^\perp = \{\mathbf{0}\}$.
v. $(\text{Nul } A)^\perp = \mathbb{R}^n$.
w. Row $A = \mathbb{R}^n$.
x. A has n nonzero singular values.

EXAMPLE 7 (*Reduced SVD and the Pseudoinverse of A*) When Σ contains rows columns of zeros, a more compact decomposition of A is possible. Using the notati established above, let $r = \text{rank}\,A$, and partition U and V into submatrices whose fi blocks contain r columns:

$$U = [\,U_r \quad U_{m-r}\,], \quad \text{where } U_r = [\,\mathbf{u}_1 \quad \cdots \quad \mathbf{u}_r\,]$$
$$V = [\,V_r \quad V_{n-r}\,], \quad \text{where } V_r = [\,\mathbf{v}_1 \quad \cdots \quad \mathbf{v}_r\,]$$

Then U_r is $m \times r$ and V_r is $n \times r$. (To simplify notation, we consider U_{m-r} or V_n even though one of them may have no columns.) Then partitioned matrix multiplicati shows that

$$A = [\,U_r \quad U_{m-r}\,]\begin{bmatrix} D & 0 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} V_r^T \\ V_{n-r}^T \end{bmatrix} = U_r D V_r^T \qquad (\,$$

This factorization of A is called a **reduced singular value decomposition** of A. Sin the diagonal entries in D are nonzero, we can form the following matrix, called t **pseudoinverse** (also, the **Moore–Penrose inverse**) of A:

$$A^+ = V_r D^{-1} U_r^T \qquad (1$$

Supplementary Exercises 12–14 at the end of the chapter explore some of the properti of the reduced singular value decomposition and the pseudoinverse.

EXAMPLE 8 (*Least-Squares Solution*) Given the equation $A\mathbf{x} = \mathbf{b}$, use the pseu doinverse of A in (10) to define

$$\hat{\mathbf{x}} = A^+\mathbf{b} = V_r D^{-1} U_r^T \mathbf{b}$$

Then, from the SVD in (9),

$$\begin{aligned} A\hat{\mathbf{x}} &= (U_r D V_r^T)(V_r D^{-1} U_r^T \mathbf{b}) \\ &= U_r D D^{-1} U_r^T \mathbf{b} \qquad \text{Because } V_r^T V_r = I_r \\ &= U_r U_r^T \mathbf{b} \end{aligned}$$

It follows from (5) that $U_r U_r^T \mathbf{b}$ is the orthogonal projection $\hat{\mathbf{b}}$ of $\mathbf{b}$ onto Col A. (See Theorem 10 in Section 6.3.) Thus $\hat{\mathbf{x}}$ is a least-squares solution of $A\mathbf{x} = \mathbf{b}$. In fact, this $\hat{\mathbf{x}}$ has the smallest length among all least-squares solutions of $A\mathbf{x} = \mathbf{b}$. See Supplementary Exercise 14.

NUMERICAL NOTE

Examples 1–4 and the exercises illustrate the concept of singular values and suggest how to perform calculations by hand. In practice, the computation of A^TA should be avoided, since any errors in the entries of A are squared in the entries of A^TA. There exist fast iterative methods that produce the singular values and singular vectors of A accurately to many decimal places.