

6.4 Symmetric Matrices

For projection onto a plane in \mathbf{R}^3 , the plane is full of eigenvectors (where $Px = x$). The other eigenvectors are *perpendicular* to the plane (where $Px = \mathbf{0}$). The eigenvalues $\lambda = 1, 1, 0$ are real. Three eigenvectors can be chosen perpendicular to each other. I have to write “can be chosen” because the two in the plane are not automatically perpendicular. This section makes that best possible choice for *symmetric matrices: The eigenvectors of $P = P^T$ are perpendicular unit vectors.*

Now we open up to all symmetric matrices. It is no exaggeration to say that these are the most important matrices the world will ever see—in the theory of linear algebra and also in the applications. We come immediately to the key question about symmetry. Not only the question, but also the answer.

What is special about $Ax = \lambda x$ when A is symmetric? We are looking for special properties of the eigenvalues λ and the eigenvectors x when $A = A^T$.

The diagonalization $A = S\Lambda S^{-1}$ will reflect the symmetry of A . We get some hint by transposing to $A^T = (S^{-1})^T \Lambda S^T$. Those are the same since $A = A^T$. Possibly S^{-1} in the first form equals S^T in the second form. Then $S^T S = I$. That makes each eigenvector in S orthogonal to the other eigenvectors. The key facts get first place in the Table at the end of this chapter, and here they are:

1. A symmetric matrix has only *real eigenvalues*.
2. The *eigenvectors* can be chosen *orthonormal*.

Those n orthonormal eigenvectors go into the columns of S . Every symmetric matrix can be diagonalized. *Its eigenvector matrix S becomes an orthogonal matrix Q .* Orthogonal matrices have $Q^{-1} = Q^T$ —what we suspected about S is true. To remember it we write $S = Q$, when we choose orthonormal eigenvectors.

Why do we use the word “choose”? Because the eigenvectors do not *have* to be unit vectors. Their lengths are at our disposal. We will choose unit vectors—eigenvectors of length one, which are orthonormal and not just orthogonal. Then $S\Lambda S^{-1}$ is in its special and particular form $Q\Lambda Q^T$ for symmetric matrices:

(Spectral Theorem) Every symmetric matrix has the factorization $A = Q\Lambda Q^T$ with real eigenvalues in Λ and orthonormal eigenvectors in $S = Q$:

Symmetric diagonalization $A = Q\Lambda Q^{-1} = Q\Lambda Q^T$ with $Q^{-1} = Q^T$.

It is easy to see that $Q\Lambda Q^T$ is symmetric. Take its transpose. You get $(Q^T)^T \Lambda^T Q^T$, which is $Q\Lambda Q^T$ again. The harder part is to prove that every symmetric matrix has real λ 's and orthonormal x 's. This is the “*spectral theorem*” in mathematics and the “*principal axis*

theorem” in geometry and physics. We have to prove it! No choice. I will approach the proof in three steps:

1. By an example, showing real λ 's in Λ and orthonormal x 's in Q .
2. By a proof of those facts when no eigenvalues are repeated.
3. By a proof that allows repeated eigenvalues (at the end of this section).

Example 1 Find the λ 's and x 's when $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$ and $A - \lambda I = \begin{bmatrix} 1 - \lambda & 2 \\ 2 & 4 - \lambda \end{bmatrix}$.

Solution The determinant of $A - \lambda I$ is $\lambda^2 - 5\lambda$. The eigenvalues are 0 and 5 (*both real*). We can see them directly: $\lambda = 0$ is an eigenvalue because A is singular, and $\lambda = 5$ matches the *trace* down the diagonal of A : $0 + 5$ agrees with $1 + 4$.

Two eigenvectors are $(2, -1)$ and $(1, 2)$ —orthogonal but not yet orthonormal. The eigenvector for $\lambda = 0$ is in the *nullspace* of A . The eigenvector for $\lambda = 5$ is in the *column space*. We ask ourselves, why are the nullspace and column space perpendicular? The Fundamental Theorem says that the nullspace is perpendicular to the *row space*—not the column space. But our matrix is *symmetric*! Its row and column spaces are the same. Its eigenvectors $(2, -1)$ and $(1, 2)$ must be (and are) perpendicular.

These eigenvectors have length $\sqrt{5}$. Divide them by $\sqrt{5}$ to get unit vectors. Put those into the columns of S (which is Q). Then $Q^{-1}AQ$ is Λ and $Q^{-1} = Q^T$:

$$Q^{-1}AQ = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & 1 \\ -1 & 2 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 5 \end{bmatrix} = \Lambda.$$

Now comes the n by n case. The λ 's are real when $A = A^T$ and $Ax = \lambda x$.

Real Eigenvalues All the eigenvalues of a real symmetric matrix are real.

Proof Suppose that $Ax = \lambda x$. Until we know otherwise, λ might be a complex number $a + ib$ (a and b real). Its complex conjugate is $\bar{\lambda} = a - ib$. Similarly the components of x may be complex numbers, and switching the signs of their imaginary parts gives \bar{x} . The good thing is that $\bar{\lambda}$ times \bar{x} is always the conjugate of λ times x . So we can take conjugates of $Ax = \lambda x$, remembering that A is real:

$$Ax = \lambda x \quad \text{leads to} \quad A\bar{x} = \bar{\lambda}\bar{x}. \quad \text{Transpose to} \quad \bar{x}^T A = \bar{x}^T \bar{\lambda}. \quad (1)$$

Now take the dot product of the first equation with \bar{x} and the last equation with x :

$$\bar{x}^T Ax = \bar{x}^T \lambda x \quad \text{and also} \quad \bar{x}^T Ax = \bar{x}^T \bar{\lambda} x. \quad (2)$$

The left sides are the same so the right sides are equal. One equation has λ , the other has $\bar{\lambda}$. They multiply $\bar{x}^T x = |x_1|^2 + |x_2|^2 + \dots = \text{length squared}$ which is not zero. Therefore λ must equal $\bar{\lambda}$, and $a + ib$ equals $a - ib$. The imaginary part is $b = 0$. Q.E.D.

The eigenvectors come from solving the real equation $(A - \lambda I)x = \mathbf{0}$. So the x 's are also real. The important fact is that they are perpendicular.

Orthogonal Eigenvectors Eigenvectors of a real symmetric matrix (when they correspond to different λ 's) are always perpendicular.

Proof Suppose $Ax = \lambda_1 x$ and $Ay = \lambda_2 y$. We are assuming here that $\lambda_1 \neq \lambda_2$. Take dot products of the first equation with y and the second with x :

$$\text{Use } A^T = A \quad (\lambda_1 x)^T y = (Ax)^T y = x^T A^T y = x^T A y = x^T \lambda_2 y. \quad (3)$$

The left side is $x^T \lambda_1 y$, the right side is $x^T \lambda_2 y$. Since $\lambda_1 \neq \lambda_2$, this proves that $x^T y = 0$. The eigenvector x (for λ_1) is perpendicular to the eigenvector y (for λ_2).

Example 2 The eigenvectors of a 2 by 2 symmetric matrix have a special form:

$$\text{Not widely known} \quad A = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad \text{has} \quad x_1 = \begin{bmatrix} b \\ \lambda_1 - a \end{bmatrix} \quad \text{and} \quad x_2 = \begin{bmatrix} \lambda_2 - c \\ b \end{bmatrix}. \quad (4)$$

This is in the Problem Set. The point here is that x_1 is perpendicular to x_2 :

$$x_1^T x_2 = b(\lambda_2 - c) + (\lambda_1 - a)b = b(\lambda_1 + \lambda_2 - a - c) = 0.$$

This is zero because $\lambda_1 + \lambda_2$ equals the trace $a + c$. Thus $x_1^T x_2 = 0$. Eagle eyes might notice the special case $a = c$, $b = 0$ when $x_1 = x_2 = \mathbf{0}$. This case has repeated eigenvalues, as in $A = I$. It still has perpendicular eigenvectors $(1, 0)$ and $(0, 1)$.

This example shows the main goal of this section—to *diagonalize symmetric matrices* A by *orthogonal eigenvector matrices* $S = Q$. Look again at the result:

$$\text{Symmetry} \quad A = S \Lambda S^{-1} \quad \text{becomes} \quad A = Q \Lambda Q^T \quad \text{with} \quad Q^T Q = I.$$

This says that every 2 by 2 symmetric matrix looks like

$$A = Q \Lambda Q^T = \begin{bmatrix} x_1 & x_2 \end{bmatrix} \begin{bmatrix} \lambda_1 & \\ & \lambda_2 \end{bmatrix} \begin{bmatrix} x_1^T \\ x_2^T \end{bmatrix}. \quad (5)$$

The columns x_1 and x_2 multiply the rows $\lambda_1 x_1^T$ and $\lambda_2 x_2^T$ to produce A :

$$\text{Sum of rank-one matrices} \quad A = \lambda_1 x_1 x_1^T + \lambda_2 x_2 x_2^T. \quad (6)$$

This is the great factorization $Q \Lambda Q^T$, written in terms of λ 's and x 's. When the symmetric matrix is n by n , there are n columns in Q multiplying n rows in Q^T . The n products $x_i x_i^T$ are *projection matrices*. Including the λ 's, the spectral theorem $A = Q \Lambda Q^T$ for symmetric matrices says that A is a combination of projection matrices:

$$A = \lambda_1 P_1 + \cdots + \lambda_n P_n \quad \lambda_i = \text{eigenvalue}, \quad P_i = \text{projection onto eigenspace.}$$

Complex Eigenvalues of Real Matrices

Equation (1) went from $Ax = \lambda x$ to $A\bar{x} = \bar{\lambda}\bar{x}$. In the end, λ and x were real. Those two equations were the same. But a nonsymmetric matrix can easily produce λ and x that are complex. In this case, $A\bar{x} = \bar{\lambda}\bar{x}$ is different from $Ax = \lambda x$. It gives us a new eigenvalue (which is $\bar{\lambda}$) and a new eigenvector (which is \bar{x}):

For real matrices, complex λ 's and x 's come in "conjugate pairs."

$$\text{If } Ax = \lambda x \text{ then } A\bar{x} = \bar{\lambda}\bar{x}.$$

Example 3 $A = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$ has $\lambda_1 = \cos \theta + i \sin \theta$ and $\lambda_2 = \cos \theta - i \sin \theta$.

Those eigenvalues are conjugate to each other. They are λ and $\bar{\lambda}$. The eigenvectors must be x and \bar{x} , because A is real:

$$\begin{aligned} \text{This is } \lambda x \quad Ax &= \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} 1 \\ -i \end{bmatrix} = (\cos \theta + i \sin \theta) \begin{bmatrix} 1 \\ -i \end{bmatrix} \\ \text{This is } \bar{\lambda} \bar{x} \quad A\bar{x} &= \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} 1 \\ i \end{bmatrix} = (\cos \theta - i \sin \theta) \begin{bmatrix} 1 \\ i \end{bmatrix}. \end{aligned} \quad (7)$$

Those eigenvectors $(1, -i)$ and $(1, i)$ are complex conjugates because A is real.

For this rotation matrix the absolute value is $|\lambda| = 1$, because $\cos^2 \theta + \sin^2 \theta = 1$. **This fact $|\lambda| = 1$ holds for the eigenvalues of every orthogonal matrix.**

We apologize that a touch of complex numbers slipped in. They are unavoidable even when the matrix is real. Chapter 10 goes beyond complex numbers λ and complex vectors to complex matrices A . Then you have the whole picture.

We end with two optional discussions.

Eigenvalues versus Pivots

The eigenvalues of A are very different from the pivots. For eigenvalues, we solve $\det(A - \lambda I) = 0$. For pivots, we use elimination. The only connection so far is this:

$$\text{product of pivots} = \text{determinant} = \text{product of eigenvalues.}$$

We are assuming a full set of pivots d_1, \dots, d_n . There are n real eigenvalues $\lambda_1, \dots, \lambda_n$. The d 's and λ 's are not the same, but they come from the same matrix. This paragraph is about a hidden relation. **For symmetric matrices the pivots and the eigenvalues have the same signs:**

The number of positive eigenvalues of $A = A^T$ equals the number of positive pivots.

Special case: A has all $\lambda_i > 0$ if and only if all pivots are positive.

That special case is an all-important fact for **positive definite matrices** in Section 6.5.

Example 4 This symmetric matrix A has one positive eigenvalue and one positive pivot:

$$\text{Matching signs} \quad A = \begin{bmatrix} 1 & 3 \\ 3 & 1 \end{bmatrix} \quad \begin{array}{l} \text{has pivots } 1 \text{ and } -8 \\ \text{eigenvalues } 4 \text{ and } -2. \end{array}$$

The signs of the pivots match the signs of the eigenvalues, one plus and one minus. This could be false when the matrix is not symmetric:

$$\text{Opposite signs} \quad B = \begin{bmatrix} 1 & 6 \\ -1 & -4 \end{bmatrix} \quad \begin{array}{l} \text{has pivots } 1 \text{ and } 2 \\ \text{eigenvalues } -1 \text{ and } -2. \end{array}$$

The diagonal entries are a third set of numbers and we say nothing about them.

Here is a proof that the pivots and eigenvalues have matching signs, when $A = A^T$.

You see it best when the pivots are divided out of the rows of U . Then A is LDL^T . The diagonal pivot matrix D goes between triangular matrices L and L^T :

$$\begin{bmatrix} 1 & 3 \\ 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 1 & \\ & -8 \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 0 & 1 \end{bmatrix} \quad \text{This is } A = LDL^T. \text{ It is symmetric.}$$

Watch the eigenvalues when L and L^T move toward the identity matrix: $A \rightarrow D$.

The eigenvalues of LDL^T are 4 and -2 . The eigenvalues of IDI^T are 1 and -8 (the pivots!). The eigenvalues are changing, as the “3” in L moves to zero. But to change *sign*, a real eigenvalue would have to cross zero. The matrix would at that moment be singular. Our changing matrix always has pivots 1 and -8 , so it is *never* singular. The signs cannot change, as the λ 's move to the d 's.

We repeat the proof for any $A = LDL^T$. Move L toward I , by moving the off-diagonal entries to zero. The pivots are not changing and not zero. The eigenvalues λ of LDL^T change to the eigenvalues d of IDI^T . Since these eigenvalues cannot cross zero as they move into the pivots, their signs cannot change. Q.E.D.

This connects the two halves of applied linear algebra—pivots and eigenvalues.

All Symmetric Matrices are Diagonalizable

When no eigenvalues of A are repeated, the eigenvectors are sure to be independent. Then A can be diagonalized. But a repeated eigenvalue can produce a shortage of eigenvectors. This *sometimes* happens for nonsymmetric matrices. It *never* happens for symmetric matrices. ***There are always enough eigenvectors to diagonalize $A = A^T$.***

Here is one idea for a proof. Change A slightly by a diagonal matrix $\mathbf{diag}(c, 2c, \dots, nc)$. If c is very small, the new symmetric matrix will have no repeated eigenvalues. Then we know it has a full set of orthonormal eigenvectors. As $c \rightarrow 0$ we obtain n orthonormal eigenvectors of the original A —even if some eigenvalues of that A are repeated.

Every mathematician knows that this argument is incomplete. How do we guarantee that the small diagonal matrix will separate the eigenvalues? (I am sure this is true.)

A different proof comes from a useful new factorization that applies to *all matrices*, symmetric or not. This new factorization immediately produces $A = Q\Lambda Q^T$ with a full set of real orthonormal eigenvectors when A is any symmetric matrix.

Every square matrix factors into $A = QTQ^{-1}$ where T is upper triangular and $\overline{Q}^T = Q^{-1}$. If A has real eigenvalues then Q and T can be chosen real: $Q^T Q = I$.

This is Schur's Theorem. We are looking for $AQ = QT$. The first column q_1 of Q must be a unit eigenvector of A . Then the first columns of AQ and QT are Aq_1 and $t_{11}q_1$. But the other columns of Q need not be eigenvectors when T is only triangular (not diagonal). So use any $n - 1$ columns that complete q_1 to a matrix Q_1 with orthonormal columns. At this point only the first columns of Q and T are set, where $Aq_1 = t_{11}q_1$:

$$\overline{Q_1}^T A Q_1 = \begin{bmatrix} \overline{q_1}^T \\ \vdots \\ \overline{q_n}^T \end{bmatrix} \begin{bmatrix} Aq_1 & \cdots & Aq_n \end{bmatrix} = \begin{bmatrix} t_{11} & \cdots \\ 0 & \boxed{A_2} \\ \vdots & \end{bmatrix}. \quad (8)$$

Now I will argue by "induction". Assume Schur's factorization $A_2 = Q_2 T_2 Q_2^{-1}$ is possible for that matrix A_2 of size $n - 1$. Put the orthogonal (or unitary) matrix Q_2 and the triangular T_2 into the final Q and T :

$$Q = Q_1 \begin{bmatrix} 1 & 0 \\ 0 & Q_2 \end{bmatrix} \quad \text{and} \quad T = \begin{bmatrix} t_{11} & \cdots \\ 0 & T_2 \end{bmatrix} \quad \text{and} \quad AQ = QT \quad \text{as desired.}$$

Note I had to allow q_1 and Q_1 to be complex, in case A has complex eigenvalues. But if t_{11} is a real eigenvalue, then q_1 and Q_1 can stay real. The induction step keeps everything real when A has real eigenvalues. Induction starts with 1 by 1, no problem.

Proof that T is the diagonal Λ when A is symmetric. Then we have $A = Q\Lambda Q^T$.

Every symmetric A has real eigenvalues. Schur's $A = QTQ^T$ with $Q^T Q = I$ means that $T = Q^T A Q$. This is a symmetric matrix (its transpose is $Q^T A Q$). Now the key point: *If T is triangular and also symmetric, it must be diagonal: $T = \Lambda$.*

This proves $A = Q\Lambda Q^T$. The matrix $A = A^T$ has n orthonormal eigenvectors.

■ REVIEW OF THE KEY IDEAS ■

1. A symmetric matrix has *real eigenvalues* and *perpendicular eigenvectors*.
2. Diagonalization becomes $A = Q\Lambda Q^T$ with an orthogonal matrix Q .
3. All symmetric matrices are diagonalizable, even with repeated eigenvalues.
4. The signs of the eigenvalues match the signs of the pivots, when $A = A^T$.
5. Every square matrix can be "triangularized" by $A = QTQ^{-1}$.

■ WORKED EXAMPLES ■

6.4 A What matrix A has eigenvalues $\lambda = 1, -1$ and eigenvectors $x_1 = (\cos \theta, \sin \theta)$ and $x_2 = (-\sin \theta, \cos \theta)$? Which of these properties can be predicted in advance?

$$A = A^T \quad A^2 = I \quad \det A = -1 \quad + \text{ and } - \text{ pivot} \quad A^{-1} = A$$

Solution All those properties can be predicted! With real eigenvalues in Λ and orthonormal eigenvectors in Q , the matrix $A = Q\Lambda Q^T$ must be symmetric. The eigenvalues 1 and -1 tell us that $A^2 = I$ (since $\lambda^2 = 1$) and $A^{-1} = A$ (same thing) and $\det A = -1$. The two pivots are positive and negative like the eigenvalues, since A is symmetric.

The matrix must be a reflection. Vectors in the direction of x_1 are unchanged by A (since $\lambda = 1$). Vectors in the perpendicular direction are reversed (since $\lambda = -1$). The reflection $A = Q\Lambda Q^T$ is across the “ θ -line”. Write c for $\cos \theta$, s for $\sin \theta$:

$$A = \begin{bmatrix} c & -s \\ s & c \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} c & s \\ -s & c \end{bmatrix} = \begin{bmatrix} c^2 - s^2 & 2cs \\ 2cs & s^2 - c^2 \end{bmatrix} = \begin{bmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{bmatrix}.$$

Notice that $x = (1, 0)$ goes to $Ax = (\cos 2\theta, \sin 2\theta)$ on the 2θ -line. And $(\cos 2\theta, \sin 2\theta)$ goes back across the θ -line to $x = (1, 0)$.

6.4 B Find the eigenvalues of A_3 and B_4 , and check the orthogonality of their first two eigenvectors. Graph these eigenvectors to see discrete sines and cosines:


$$A_3 = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} \quad B_4 = \begin{bmatrix} 1 & -1 & & \\ -1 & 2 & -1 & \\ & -1 & 2 & -1 \\ & & -1 & 1 \end{bmatrix}$$

The $-1, 2, -1$ pattern in both matrices is a “second difference”. Section 8.1 will explain how this is like a second derivative. Then $Ax = \lambda x$ and $Bx = \lambda x$ are like $d^2x/dt^2 = \lambda x$. This has eigenvectors $x = \sin kt$ and $x = \cos kt$ that are the bases for Fourier series. The matrices lead to “discrete sines” and “discrete cosines” that are the bases for the *Discrete Fourier Transform*. This DFT is absolutely central to all areas of digital signal processing. The favorite choice for JPEG in image processing has been B_8 of size 8.

Solution The eigenvalues of A_3 are $\lambda = 2 - \sqrt{2}$ and 2 and $2 + \sqrt{2}$. (see 6.3 B). Their sum is 6 (the trace of A_3) and their product is 4 (the determinant). The eigenvector matrix S gives the “Discrete Sine Transform” and the graph shows how the first two eigenvectors fall onto sine curves. Please draw the third eigenvector onto a third sine curve!

$$S = \begin{bmatrix} 1 & \sqrt{2} & 1 \\ \sqrt{2} & 0 & -\sqrt{2} \\ 1 & -\sqrt{2} & 1 \end{bmatrix}$$


Eigenvector matrix for A_3



The eigenvalues of B_4 are $\lambda = 2 - \sqrt{2}$ and 2 and $2 + \sqrt{2}$ and 0 (the same as for A_3 , plus the zero eigenvalue). The trace is still 6, but the determinant is now zero. The eigenvector matrix C gives the 4-point “Discrete Cosine Transform” and the graph shows how the first two eigenvectors fall onto cosine curves. (Please plot the third eigenvector.) These eigenvectors match cosines at the *halfway points* $\frac{\pi}{8}, \frac{3\pi}{8}, \frac{5\pi}{8}, \frac{7\pi}{8}$.

$$C = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & \sqrt{2}-1 & -1 & 1-\sqrt{2} \\ 1 & 1-\sqrt{2} & -1 & \sqrt{2}-1 \\ 1 & -1 & 1 & -1 \end{bmatrix}$$

Eigenvector matrix for B_4



S and C have orthogonal columns (eigenvectors of the symmetric A_3 and B_4). When we multiply a vector by S or C , that signal splits into pure frequencies—as a musical chord separates into pure notes. This is the most useful and insightful transform in all of signal processing. Here is a MATLAB code to create B_8 and its eigenvector matrix C :

```
n=8; e=ones(n-1,1); B=2*eye(n)-diag(e,-1)-diag(e,1); B(1,1)=1; B(n,n)=1;
[C, Lambda]=eig(B);
plot(C(:,1:4), 'o')
```

Problem Set 6.4

- 1 Write A as $M + N$, symmetric matrix plus skew-symmetric matrix:

$$A = \begin{bmatrix} 1 & 2 & 4 \\ 4 & 3 & 0 \\ 8 & 6 & 5 \end{bmatrix} = M + N \quad (M^T = M, N^T = -N).$$

For any square matrix, $M = \frac{A+A^T}{2}$ and $N = \frac{A-A^T}{2}$ add up to A .

- 2 If C is symmetric prove that $A^T C A$ is also symmetric. (Transpose it.) When A is 6 by 3, what are the shapes of C and $A^T C A$?

- 3 Find the eigenvalues and the unit eigenvectors of

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

- 4 Find an orthogonal matrix Q that diagonalizes $A = \begin{bmatrix} -2 & 6 \\ 6 & 7 \end{bmatrix}$. What is Λ ?

- 5 Find an orthogonal matrix Q that diagonalizes this symmetric matrix:

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

- 6 Find *all* orthogonal matrices that diagonalize $A = \begin{bmatrix} 9 & 12 \\ 12 & 16 \end{bmatrix}$.

- 7 (a) Find a symmetric matrix $\begin{bmatrix} 1 & b \\ b & 1 \end{bmatrix}$ that has a negative eigenvalue.

(b) How do you know it must have a negative pivot?

(c) How do you know it can't have two negative eigenvalues?

- 8 If $A^3 = 0$ then the eigenvalues of A must be _____. Give an example that has $A \neq 0$. But if A is symmetric, diagonalize it to prove that A must be zero.

- 9 If $\lambda = a + ib$ is an eigenvalue of a real matrix A , then its conjugate $\bar{\lambda} = a - ib$ is also an eigenvalue. (If $Ax = \lambda x$ then also $A\bar{x} = \bar{\lambda}\bar{x}$.) Prove that every real 3 by 3 matrix has at least one real eigenvalue.

- 10 Here is a quick “proof” that the eigenvalues of all real matrices are real:

$$\text{False proof } Ax = \lambda x \text{ gives } x^T Ax = \lambda x^T x \text{ so } \lambda = \frac{x^T Ax}{x^T x} \text{ is real.}$$

Find the flaw in this reasoning—a hidden assumption that is not justified. You could test those steps on the 90° rotation matrix $\begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ with $\lambda = i$ and $x = (i, 1)$.

- 11 Write A and B in the form $\lambda_1 x_1 x_1^T + \lambda_2 x_2 x_2^T$ of the spectral theorem $Q\Lambda Q^T$:

$$A = \begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix} \quad B = \begin{bmatrix} 9 & 12 \\ 12 & 16 \end{bmatrix} \quad (\text{keep } \|x_1\| = \|x_2\| = 1).$$

- 12 Every 2 by 2 symmetric matrix is $\lambda_1 x_1 x_1^T + \lambda_2 x_2 x_2^T = \lambda_1 P_1 + \lambda_2 P_2$. Explain $P_1 + P_2 = x_1 x_1^T + x_2 x_2^T = I$ from columns times rows of Q . Why is $P_1 P_2 = 0$?

- 13 What are the eigenvalues of $A = \begin{bmatrix} 0 & b \\ -b & 0 \end{bmatrix}$? Create a 4 by 4 skew-symmetric matrix ($A^T = -A$) and verify that all its eigenvalues are imaginary.

- 14** (Recommended) This matrix M is skew-symmetric and also _____. Then all its eigenvalues are pure imaginary and they also have $|\lambda| = 1$. ($\|Mx\| = \|x\|$ for every x so $\|\lambda x\| = \|x\|$ for eigenvectors.) Find all four eigenvalues from the trace of M :

$$M = \frac{1}{\sqrt{3}} \begin{bmatrix} 0 & 1 & 1 & 1 \\ -1 & 0 & -1 & 1 \\ -1 & 1 & 0 & -1 \\ -1 & -1 & 1 & 0 \end{bmatrix} \quad \text{can only have eigenvalues } i \text{ or } -i.$$

- 15** Show that A (**symmetric but complex**) has only one line of eigenvectors:

$$A = \begin{bmatrix} i & 1 \\ 1 & -i \end{bmatrix} \text{ is not even diagonalizable: eigenvalues } \lambda = 0, 0.$$

$A^T = A$ is not such a special property for complex matrices. The good property is $\overline{A}^T = A$ (Section 10.2). Then all λ 's are real and eigenvectors are orthogonal.

- 16** Even if A is rectangular, the block matrix $B = \begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}$ is symmetric:

$$Bx = \lambda x \quad \text{is} \quad \begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix} \begin{bmatrix} y \\ z \end{bmatrix} = \lambda \begin{bmatrix} y \\ z \end{bmatrix} \quad \text{which is} \quad \begin{array}{l} Az = \lambda y \\ A^T y = \lambda z. \end{array}$$

- (a) Show that $-\lambda$ is also an eigenvalue, with the eigenvector $(y, -z)$.
 (b) Show that $A^T A z = \lambda^2 z$, so that λ^2 is an eigenvalue of $A^T A$.
 (c) If $A = I$ (2 by 2) find all four eigenvalues and eigenvectors of B .
- 17** If $A = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ in Problem 16, find all three eigenvalues and eigenvectors of B .
- 18** *Another proof that eigenvectors are perpendicular when $A = A^T$.* Two steps:
1. Suppose $Ax = \lambda x$ and $Ay = 0y$ and $\lambda \neq 0$. Then y is in the nullspace and x is in the column space. They are perpendicular because _____. Go carefully—why are these subspaces orthogonal?
 2. If $Ay = \beta y$, apply this argument to $A - \beta I$. The eigenvalue of $A - \beta I$ moves to zero and the eigenvectors stay the same—so they are perpendicular.
- 19** Find the eigenvector matrix S for A and for B . Show that S doesn't collapse at $d = 1$, even though $\lambda = 1$ is repeated. Are the eigenvectors perpendicular?

$$A = \begin{bmatrix} 0 & d & 0 \\ d & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad B = \begin{bmatrix} -d & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & d \end{bmatrix} \quad \text{have} \quad \lambda = 1, d, -d.$$

- 20** Write a 2 by 2 *complex* matrix with $\overline{A}^T = A$ (a "Hermitian matrix"). Find λ_1 and λ_2 for your complex matrix. Adjust equations (1) and (2) to show that *the eigenvalues of a Hermitian matrix are real*.

- 21 **True** (with reason) or **false** (with example). “Orthonormal” is not assumed.
- (a) A matrix with real eigenvalues and eigenvectors is symmetric.
 - (b) A matrix with real eigenvalues and orthogonal eigenvectors is symmetric.
 - (c) The inverse of a symmetric matrix is symmetric.
 - (d) The eigenvector matrix S of a symmetric matrix is symmetric.
- 22 (A paradox for instructors) If $AA^T = A^T A$ then A and A^T share the same eigenvectors (true). A and A^T always share the same eigenvalues. Find the flaw in this conclusion: They must have the same S and Λ . Therefore A equals A^T .
- 23 (Recommended) Which of these classes of matrices do A and B belong to: Invertible, orthogonal, projection, permutation, diagonalizable, Markov?

$$A = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad B = \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

Which of these factorizations are possible for A and B : LU , QR , $S\Lambda S^{-1}$, $Q\Lambda Q^T$?

- 24 What number b in $\begin{bmatrix} 2 & b \\ 1 & 0 \end{bmatrix}$ makes $A = Q\Lambda Q^T$ possible? What number makes $A = S\Lambda S^{-1}$ impossible? What number makes A^{-1} impossible?
- 25 Find all 2 by 2 matrices that are orthogonal and also symmetric. Which two numbers can be eigenvalues?
- 26 This A is nearly symmetric. But its eigenvectors are far from orthogonal:

$$A = \begin{bmatrix} 1 & 10^{-15} \\ 0 & 1 + 10^{-15} \end{bmatrix} \quad \text{has eigenvectors} \quad \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad [?]$$

What is the angle between the eigenvectors?

- 27 (MATLAB) Take two symmetric matrices with different eigenvectors, say $A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$ and $B = \begin{bmatrix} 8 & 1 \\ 1 & 0 \end{bmatrix}$. Graph the eigenvalues $\lambda_1(A + tB)$ and $\lambda_2(A + tB)$ for $-8 < t < 8$. Peter Lax says on page 113 of *Linear Algebra* that λ_1 and λ_2 appear to be on a collision course at certain values of t . “Yet at the last minute they turn aside.” How close do they come?

Challenge Problems

- 28 For complex matrices, the symmetry $A^T = A$ that produces real eigenvalues changes to $\overline{A^T} = A$. From $\det(A - \lambda I) = 0$, find the eigenvalues of the 2 by 2 “Hermitian” matrix $A = \begin{bmatrix} 4 & 2 + i \\ 2 - i & 0 \end{bmatrix} = \overline{A^T}$. To see why eigenvalues are real when $\overline{A^T} = A$, adjust equation (1) of the text to $\overline{A} \overline{x} = \overline{\lambda} \overline{x}$.

Transpose to $\overline{x^T} \overline{A^T} = \overline{x^T} \overline{\lambda}$. With $\overline{A^T} = A$, reach equation (2): $\lambda = \overline{\lambda}$.

- 29** Normal matrices have $\overline{A}^T A = A \overline{A}^T$. For real matrices, $A^T A = A A^T$ includes symmetric, skew-symmetric, and orthogonal. Those have real λ , imaginary λ , and $|\lambda| = 1$. Other normal matrices can have any complex eigenvalues λ .

Key point: *Normal matrices have n orthonormal eigenvectors.* Those vectors x_i probably will have complex components. In that complex case orthogonality means $\overline{x}_i^T x_j = 0$ as Chapter 10 explains. Inner products (dot products) become $\overline{x}^T y$.

The test for n orthonormal columns in Q becomes $\overline{Q}^T Q = I$ instead of $Q^T Q = I$.

A has n orthonormal eigenvectors ($A = Q \Lambda \overline{Q}^T$) if and only if A is normal.

- (a) Start from $A = Q \Lambda \overline{Q}^T$ with $\overline{Q}^T Q = I$. Show that $\overline{A}^T A = A \overline{A}^T$: A is normal.
 (b) Now start from $\overline{A}^T A = A \overline{A}^T$. Schur found $A = Q T \overline{Q}^T$ for every matrix A , with a triangular T . For normal matrices we must show (in 3 steps) that this T will actually be diagonal. Then $T = \Lambda$.

Step 1. Put $A = Q T \overline{Q}^T$ into $\overline{A}^T A = A \overline{A}^T$ to find $\overline{T}^T T = T \overline{T}^T$.

Step 2. Suppose $T = \begin{bmatrix} a & b \\ 0 & d \end{bmatrix}$ has $\overline{T}^T T = T \overline{T}^T$. Prove that $b = 0$.

Step 3. Extend Step 2 to size n . A normal triangular T must be diagonal.

- 30** If λ_{\max} is the largest eigenvalue of a symmetric matrix A , no diagonal entry can be larger than λ_{\max} . What is the first entry a_{11} of $A = Q \Lambda Q^T$? Show why $a_{11} \leq \lambda_{\max}$.
- 31** Suppose $A^T = -A$ (real antisymmetric matrix). Explain these facts about A :
- $x^T A x = 0$ for every real vector x .
 - The eigenvalues of A are pure imaginary.
 - The determinant of A is positive or zero (not negative).

For (a), multiply out an example of $x^T A x$ and watch terms cancel. Or reverse $x^T(Ax)$ to $(Ax)^T x$. For (b), $Az = \lambda z$ leads to $\overline{z}^T A z = \lambda \overline{z}^T z = \lambda \|z\|^2$. Part (a) shows that $\overline{z}^T A z = (x - iy)^T A(x + iy)$ has zero real part. Then (b) helps with (c).

- 32** If A is symmetric and all its eigenvalues are $\lambda = 2$, how do you know that A must be $2I$? (Key point: Symmetry guarantees that A is diagonalizable. See “Proofs of the Spectral Theorem” on web.mit.edu/18.06.)

6.5 Positive Definite Matrices

This section concentrates on *symmetric matrices that have positive eigenvalues*. If symmetry makes a matrix important, this extra property (*all $\lambda > 0$*) makes it truly special. When we say special, we don't mean rare. Symmetric matrices with positive eigenvalues are at the center of all kinds of applications. They are called *positive definite*.

The first problem is to recognize these matrices. You may say, just find the eigenvalues and test $\lambda > 0$. That is exactly what we want to avoid. Calculating eigenvalues is work. When the λ 's are needed, we can compute them. But if we just want to know that they are positive, there are faster ways. Here are two goals of this section:

- To find *quick tests* on a symmetric matrix that guarantee *positive eigenvalues*.
- To explain important applications of positive definiteness.

The λ 's are automatically real because the matrix is symmetric.

Start with 2 by 2. When does $A = \begin{bmatrix} a & b \\ b & c \end{bmatrix}$ have $\lambda_1 > 0$ and $\lambda_2 > 0$?

The eigenvalues of A are positive if and only if $a > 0$ and $ac - b^2 > 0$.

$$A_1 = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix} \text{ is not positive definite because } ac - b^2 = 1 - 4 < 0$$

$$A_2 = \begin{bmatrix} 1 & -2 \\ -2 & 6 \end{bmatrix} \text{ is positive definite because } a = 1 \text{ and } ac - b^2 = 6 - 4 > 0$$

$$A_3 = \begin{bmatrix} -1 & 2 \\ 2 & -6 \end{bmatrix} \text{ is not positive definite (even with } \det A = +2) \text{ because } a = -1$$

Notice that we didn't compute the eigenvalues 3 and -1 of A_1 . Positive trace $3 - 1 = 2$, negative determinant $(3)(-1) = -3$. And $A_3 = -A_2$ is *negative* definite. The positive eigenvalues for A_2 , two negative eigenvalues for A_3 .

Proof that the 2 by 2 test is passed when $\lambda_1 > 0$ and $\lambda_2 > 0$. Their product $\lambda_1\lambda_2$ is the determinant so $ac - b^2 > 0$. Their sum is the trace so $a + c > 0$. Then a and c are both positive (if one of them is not positive, $ac - b^2 > 0$ will fail). Problem 1 reverses the reasoning to show that the tests guarantee $\lambda_1 > 0$ and $\lambda_2 > 0$.

This test uses the 1 by 1 determinant a and the 2 by 2 determinant $ac - b^2$. When A is 3 by 3, $\det A > 0$ is the third part of the test. The next test requires *positive pivots*.

The eigenvalues of $A = A^T$ are positive if and only if the pivots are positive:

$$a > 0 \quad \text{and} \quad \frac{ac - b^2}{a} > 0.$$

$a > 0$ is required in both tests. So $ac > b^2$ is also required, for the determinant test and now the pivot. The point is to recognize that ratio as the *second pivot* of A :

$$\begin{bmatrix} a & b \\ b & c \end{bmatrix} \xrightarrow[\text{The multiplier is } b/a]{\text{The first pivot is } a} \begin{bmatrix} a & b \\ 0 & c - \frac{b}{a}b \end{bmatrix} \quad \text{The second pivot is } c - \frac{b^2}{a} = \frac{ac - b^2}{a}$$

This connects two big parts of linear algebra. **Positive eigenvalues mean positive pivots and vice versa.** We gave a proof for symmetric matrices of any size in the last section. The pivots give a quick test for $\lambda > 0$, and they are a lot faster to compute than the eigenvalues. It is very satisfying to see pivots and determinants and eigenvalues come together in this course.

$A_1 = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}$	$A_2 = \begin{bmatrix} 1 & -2 \\ -2 & 6 \end{bmatrix}$	$A_3 = \begin{bmatrix} -1 & 2 \\ 2 & -6 \end{bmatrix}$
pivots 1 and -3 (indefinite)	pivots 1 and 2 (positive definite)	pivots -1 and -2 (negative definite)

Here is a different way to look at symmetric matrices with positive eigenvalues.

Energy-based Definition

From $Ax = \lambda x$, multiply by x^T to get $x^T Ax = \lambda x^T x$. The right side is a positive λ times a positive number $x^T x = \|x\|^2$. So $x^T Ax$ is positive for any eigenvector.

The new idea is that $x^T Ax$ is **positive for all nonzero vectors** x , not just the eigenvectors. In many applications this number $x^T Ax$ (or $\frac{1}{2}x^T Ax$) is the **energy** in the system. The requirement of positive energy gives *another definition* of a positive definite matrix. I think this energy-based definition is the fundamental one.

Eigenvalues and pivots are two equivalent ways to test the new requirement $x^T Ax > 0$.

Definition A is positive definite if $x^T Ax > 0$ for every nonzero vector x :

$$x^T Ax = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = ax^2 + 2bxy + cy^2 > 0. \tag{1}$$

The four entries a, b, b, c give the four parts of $x^T Ax$. From a and c come the pure squares ax^2 and cy^2 . From b and b off the diagonal come the cross terms bxy and byx (the same). Adding those four parts gives $x^T Ax$. This energy-based definition leads to a basic fact:

If A and B are symmetric positive definite, so is $A + B$.

Reason: $x^T(A + B)x$ is simply $x^T Ax + x^T Bx$. Those two terms are positive (for $x \neq 0$) so $A + B$ is also positive definite. The pivots and eigenvalues are not easy to follow when matrices are added, but the energies just add.

$x^T Ax$ also connects with our final way to recognize a positive definite matrix. Start with any matrix R , possibly rectangular. We know that $A = R^T R$ is square and symmetric. More than that, A will be positive definite when R has independent columns:

If the columns of R are independent, then $A = R^T R$ is positive definite.

Again eigenvalues and pivots are not easy. But the number $x^T Ax$ is the same as $x^T R^T R x$. That is exactly $(Rx)^T (Rx)$ —another important proof by parenthesis! That vector Rx is not zero when $x \neq \mathbf{0}$ (this is the meaning of independent columns). Then $x^T Ax$ is the positive number $\|Rx\|^2$ and the matrix A is positive definite.

Let me collect this theory together, into five equivalent statements of positive definiteness. You will see how that key idea connects the whole subject of linear algebra: pivots, determinants, eigenvalues, and least squares (from $R^T R$). Then come the applications.

When a symmetric matrix has one of these five properties, it has them all:

1. All n pivots are positive.
2. All n upper left determinants are positive.
3. All n eigenvalues are positive.
4. $x^T Ax$ is positive except at $x = \mathbf{0}$. This is the *energy-based* definition.
5. A equals $R^T R$ for a matrix R with *independent columns*.

The “upper left determinants” are 1 by 1, 2 by 2, . . . , n by n . The last one is the determinant of the complete matrix A . This remarkable theorem ties together the whole linear algebra course—at least for symmetric matrices. We believe that two examples are more helpful than a detailed proof (we nearly have a proof already).

Example 1 Test these matrices A and B for positive definiteness:

$$A = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 2 & -1 & b \\ -1 & 2 & -1 \\ b & -1 & 2 \end{bmatrix}.$$

Solution The pivots of A are 2 and $\frac{3}{2}$ and $\frac{4}{3}$, all positive. Its upper left determinants are 2 and 3 and 4, all positive. The eigenvalues of A are $2 - \sqrt{2}$ and 2 and $2 + \sqrt{2}$, all positive. That completes tests 1, 2, and 3.

We can write $x^T Ax$ as a sum of three squares. The pivots 2, $\frac{3}{2}$, $\frac{4}{3}$ appear outside the squares. The multipliers $-\frac{1}{2}$ and $-\frac{2}{3}$ from elimination are inside the squares:

$$\begin{aligned} x^T Ax &= 2(x_1^2 - x_1 x_2 + x_2^2 - x_2 x_3 + x_3^2) && \text{Rewrite with squares} \\ &= 2(x_1 - \frac{1}{2}x_2)^2 + \frac{3}{2}(x_2 - \frac{2}{3}x_3)^2 + \frac{4}{3}(x_3)^2. && \text{This sum is positive.} \end{aligned}$$

I have two candidates to suggest for R . Either one will show that $A = R^T R$ is positive definite. R can be a rectangular first difference matrix, 4 by 3, to produce those second differences $-1, 2, -1$ in A :

$$A = R^T R \quad \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & -1 \end{bmatrix}$$

The three columns of this R are independent. A is positive definite.

Another R comes from $A = LDL^T$ (the symmetric version of $A = LU$). Elimination gives the pivots $2, \frac{3}{2}, \frac{4}{3}$ in D and the multipliers $-\frac{1}{2}, 0, -\frac{2}{3}$ in L . **Just put \sqrt{D} with L .**

$$LDL^T = \begin{bmatrix} 1 & & & \\ -\frac{1}{2} & 1 & & \\ 0 & -\frac{2}{3} & 1 & \\ & & & 1 \end{bmatrix} \begin{bmatrix} 2 & & & \\ & \frac{3}{2} & & \\ & & \frac{4}{3} & \\ & & & & 1 \end{bmatrix} \begin{bmatrix} 1 & -\frac{1}{2} & & \\ & 1 & -\frac{2}{3} & \\ & & 1 & \\ & & & 1 \end{bmatrix} = (L\sqrt{D})(L\sqrt{D})^T = R^T R. \quad (2)$$

R is the Cholesky factor

This choice of R has square roots (not so beautiful). But it is the only R that is 3 by 3 and upper triangular. It is the “Cholesky factor” of A and it is computed by MATLAB’s command $R = \text{chol}(A)$. In applications, the rectangular R is how we build A and this Cholesky R is how we break it apart.

Eigenvalues give the symmetric choice $R = Q\sqrt{\Lambda}Q^T$. This is also successful with $R^T R = Q\Lambda Q^T = A$. All these tests show that the $-1, 2, -1$ matrix A is positive definite.

Now turn to B , where the $(1, 3)$ and $(3, 1)$ entries move away from 0 to b . This b must not be too large! *The determinant test is easiest.* The 1 by 1 determinant is 2, the 2 by 2 determinant is still 3. The 3 by 3 determinant involves b :

$$\det B = 4 + 2b - 2b^2 = (1 + b)(4 - 2b) \quad \text{must be positive.}$$

At $b = -1$ and $b = 2$ we get $\det B = 0$. *Between $b = -1$ and $b = 2$ the matrix is positive definite.* The corner entry $b = 0$ in the first matrix A was safely between.

Positive Semidefinite Matrices

Often we are at the edge of positive definiteness. The determinant is zero. The smallest eigenvalue is zero. The energy in its eigenvector is $x^T Ax = x^T 0x = 0$. These matrices on the edge are called *positive semidefinite*. Here are two examples (not invertible):

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix} \quad \text{are positive semidefinite.}$$

A has eigenvalues 5 and 0. Its upper left determinants are 1 and 0. Its rank is only 1. This matrix A factors into $R^T R$ with **dependent columns** in R :

$$\begin{array}{l} \text{Dependent columns} \\ \text{Positive semidefinite} \end{array} \quad \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 2 & 0 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & 0 \end{bmatrix} = R^T R.$$

If 4 is increased by any small number, the matrix will become positive definite.

The cyclic B also has zero determinant (computed above when $b = -1$). It is singular. The eigenvector $x = (1, 1, 1)$ has $Bx = \mathbf{0}$ and $x^T Bx = 0$. Vectors x in all other directions do give positive energy. This B can be written as $R^T R$ in many ways, but R will always have *dependent* columns, with $(1, 1, 1)$ in its nullspace:

$$\begin{array}{l} \text{Second differences } A \\ \text{from first differences } R^T R \\ \text{Cyclic } A \text{ from cyclic } R \end{array} \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & -1 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix}.$$

Positive semidefinite matrices have all $\lambda \geq 0$ and all $x^T Ax \geq 0$. Those weak inequalities (\geq instead of $>$) include positive definite matrices and the singular matrices at the edge.

First Application: The Ellipse $ax^2 + 2bxy + cy^2 = 1$

Think of a tilted ellipse $x^T Ax = 1$. Its center is $(0, 0)$, as in Figure 6.7a. Turn it to line up with the coordinate axes (X and Y axes). That is Figure 6.7b. These two pictures show the geometry behind the factorization $A = Q\Lambda Q^{-1} = Q\Lambda Q^T$:

1. The tilted ellipse is associated with A . Its equation is $x^T Ax = 1$.
2. The lined-up ellipse is associated with Λ . Its equation is $X^T \Lambda X = 1$.
3. The rotation matrix that lines up the ellipse is the eigenvector matrix Q .

Example 2 Find the axes of this tilted ellipse $5x^2 + 8xy + 5y^2 = 1$.

Solution Start with the positive definite matrix that matches this equation:

The equation is $\begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 1$. The matrix is $A = \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix}$.




Figure 6.7: The tilted ellipse $5x^2 + 8xy + 5y^2 = 1$. Lined up it is $9X^2 + Y^2 = 1$.

The eigenvectors are $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$. Divide by $\sqrt{2}$ for unit vectors. Then $A = Q\Lambda Q^T$:

$$\begin{array}{l} \text{Eigenvectors in } Q \\ \text{Eigenvalues 9 and 1} \end{array} \quad \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} 9 & 0 \\ 0 & 1 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

Now multiply by $\begin{bmatrix} x & y \end{bmatrix}$ on the left and $\begin{bmatrix} x \\ y \end{bmatrix}$ on the right to get back to $x^T Ax$:

$$x^T Ax = \text{sum of squares} \quad 5x^2 + 8xy + 5y^2 = 9 \left(\frac{x+y}{\sqrt{2}} \right)^2 + 1 \left(\frac{x-y}{\sqrt{2}} \right)^2. \quad (3)$$

The coefficients are not the pivots 5 and 9/5 from D , they are the eigenvalues 9 and 1 from Λ . Inside *these* squares are the eigenvectors $(1, 1)/\sqrt{2}$ and $(1, -1)/\sqrt{2}$.

The axes of the tilted ellipse point along the eigenvectors. This explains why $A = Q\Lambda Q^T$ is called the “principal axis theorem”—it displays the axes. Not only the axis directions (from the eigenvectors) but also the axis lengths (from the eigenvalues). To see it all, use capital letters for the new coordinates that line up the ellipse:

$$\text{Lined up} \quad \frac{x+y}{\sqrt{2}} = X \quad \text{and} \quad \frac{x-y}{\sqrt{2}} = Y \quad \text{and} \quad 9X^2 + Y^2 = 1.$$

The largest value of X^2 is $1/9$. The endpoint of the shorter axis has $X = 1/3$ and $Y = 0$. Notice: The *bigger* eigenvalue λ_1 gives the *shorter* axis, of half-length $1/\sqrt{\lambda_1} = 1/3$. The smaller eigenvalue $\lambda_2 = 1$ gives the greater length $1/\sqrt{\lambda_2} = 1$.

In the xy system, the axes are along the eigenvectors of A . In the XY system, the axes are along the eigenvectors of Λ —the coordinate axes. All comes from $A = Q\Lambda Q^T$.

Suppose $A = Q\Lambda Q^T$ is positive definite, so $\lambda_i > 0$. The graph of $x^T Ax = 1$ is an ellipse:

$$\begin{bmatrix} x & y \end{bmatrix} Q\Lambda Q^T \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} X & Y \end{bmatrix} \Lambda \begin{bmatrix} X \\ Y \end{bmatrix} = \lambda_1 X^2 + \lambda_2 Y^2 = 1.$$

The axes point along eigenvectors. The half-lengths are $1/\sqrt{\lambda_1}$ and $1/\sqrt{\lambda_2}$.

$A = I$ gives the circle $x^2 + y^2 = 1$. If one eigenvalue is negative (exchange 4's and 5's in A), we don't have an ellipse. The sum of squares becomes a *difference of squares*: $9X^2 - Y^2 = 1$. This indefinite matrix gives a *hyperbola*. For a negative definite matrix like $A = -I$, with both λ 's negative, the graph of $-x^2 - y^2 = 1$ has no points at all.

■ REVIEW OF THE KEY IDEAS ■

1. Positive definite matrices have positive eigenvalues and positive pivots.
2. A quick test is given by the upper left determinants: $a > 0$ and $ac - b^2 > 0$.

3. The graph of $\mathbf{x}^T A \mathbf{x}$ is then a “bowl” going up from $\mathbf{x} = \mathbf{0}$:

$$\mathbf{x}^T A \mathbf{x} = ax^2 + 2bxy + cy^2 \text{ is positive except at } (x, y) = (0, 0).$$

4. $A = R^T R$ is automatically positive definite if R has independent columns.

5. The ellipse $\mathbf{x}^T A \mathbf{x} = 1$ has its axes along the eigenvectors of A . Lengths $1/\sqrt{\lambda}$.

■ WORKED EXAMPLES ■

6.5 A The great factorizations of a symmetric matrix are $A = LDL^T$ from pivots and multipliers, and $A = Q\Lambda Q^T$ from eigenvalues and eigenvectors. Show that $\mathbf{x}^T A \mathbf{x} > 0$ for all nonzero \mathbf{x} exactly when the pivots and eigenvalues are positive. Try these n by n tests on `pascal(6)` and `ones(6)` and `hilb(6)` and other matrices in MATLAB’s gallery.

Solution To prove $\mathbf{x}^T A \mathbf{x} > 0$, put parentheses into $\mathbf{x}^T LDL^T \mathbf{x}$ and $\mathbf{x}^T Q\Lambda Q^T \mathbf{x}$:

$$\mathbf{x}^T A \mathbf{x} = (L^T \mathbf{x})^T D (L^T \mathbf{x}) \quad \text{and} \quad \mathbf{x}^T A \mathbf{x} = (Q^T \mathbf{x})^T \Lambda (Q^T \mathbf{x}).$$

If \mathbf{x} is nonzero, then $\mathbf{y} = L^T \mathbf{x}$ and $\mathbf{z} = Q^T \mathbf{x}$ are nonzero (those matrices are invertible). So $\mathbf{x}^T A \mathbf{x} = \mathbf{y}^T D \mathbf{y} = \mathbf{z}^T \Lambda \mathbf{z}$ becomes a sum of squares and A is shown as positive definite:

$$\text{Pivots} \quad \mathbf{x}^T A \mathbf{x} = \mathbf{y}^T D \mathbf{y} = d_1 y_1^2 + \cdots + d_n y_n^2 > 0$$

$$\text{Eigenvalues} \quad \mathbf{x}^T A \mathbf{x} = \mathbf{z}^T \Lambda \mathbf{z} = \lambda_1 z_1^2 + \cdots + \lambda_n z_n^2 > 0$$

MATLAB has a gallery of unusual matrices (type `help gallery`) and here are four:

`pascal(6)` is positive definite because all its pivots are 1 (Worked Example 2.6 A).

`ones(6)` is positive *semidefinite* because its eigenvalues are 0, 0, 0, 0, 0, 6.

`H=hilb(6)` is positive definite even though `eig(H)` shows two eigenvalues very near zero.

Hilbert matrix $\mathbf{x}^T H \mathbf{x} = \int_0^1 (x_1 + x_2 s + \cdots + x_6 s^5)^2 ds > 0$, $H_{ij} = 1/(i + j + 1)$.

`rand(6)+rand(6)'` can be positive definite or not. Experiments gave only 2 in 20000.

$n = 20000$; $p = 0$; for $k = 1:n$, $A = \text{rand}(6)$; $p = p + \text{all}(\text{eig}(A + A') > 0)$; end, p/n

6.5 B When is the symmetric block matrix $M = \begin{bmatrix} A & B \\ B^T & C \end{bmatrix}$ positive definite?

Solution Multiply the first row of M by $B^T A^{-1}$ and subtract from the second row, to get a block of zeros. The *Schur complement* $S = C - B^T A^{-1} B$ appears in the corner:

$$\begin{bmatrix} I & 0 \\ -B^T A^{-1} & I \end{bmatrix} \begin{bmatrix} A & B \\ B^T & C \end{bmatrix} = \begin{bmatrix} A & B \\ 0 & C - B^T A^{-1} B \end{bmatrix} = \begin{bmatrix} A & B \\ 0 & S \end{bmatrix} \quad (4)$$

Those two blocks A and S must be positive definite. Their pivots are the pivots of M .

6.5 C *Second application: Test for a minimum.* Does $F(x, y)$ have a minimum if $\partial F/\partial x = 0$ and $\partial F/\partial y = 0$ at the point $(x, y) = (0, 0)$?

Solution For $f(x)$, the test for a minimum comes from calculus: $df/dx = 0$ and $d^2 f/dx^2 > 0$. Moving to two variables x and y produces a symmetric matrix H . It contains the four second derivatives of $F(x, y)$. Positive f'' changes to positive definite H :

$$\text{Second derivative matrix} \quad H = \begin{bmatrix} \partial^2 F/\partial x^2 & \partial^2 F/\partial x\partial y \\ \partial^2 F/\partial y\partial x & \partial^2 F/\partial y^2 \end{bmatrix}$$

$F(x, y)$ has a minimum if H is positive definite. Reason: H reveals the important terms $ax^2 + 2bxy + cy^2$ near $(x, y) = (0, 0)$. The second derivatives of F are $2a, 2b, 2b, 2c$!

6.5 D Find the eigenvalues of the $-1, 2, -1$ tridiagonal n by n matrix K (my favorite).

Solution The best way is to guess λ and x . Then check $Kx = \lambda x$. Guessing could not work for most matrices, but special cases are a big part of mathematics (pure and applied).

The key is hidden in a differential equation. The second difference matrix K is like a *second derivative*, and those eigenvalues are much easier to see:

$$\begin{array}{l} \text{Eigenvalues } \lambda_1, \lambda_2, \dots \\ \text{Eigenfunctions } y_1, y_2, \dots \end{array} \quad \frac{d^2 y}{dx^2} = \lambda y(x) \quad \text{with} \quad \begin{array}{l} y(0) = 0 \\ y(1) = 0 \end{array} \quad (5)$$

Try $y = \sin cx$. Its second derivative is $y'' = -c^2 \sin cx$. So the eigenvalue will be $\lambda = -c^2$, provided $y(x)$ satisfies the end point conditions $y(0) = 0 = y(1)$.

Certainly $\sin 0 = 0$ (this is where cosines are eliminated by $\cos 0 = 1$). At $x = 1$, we need $y(1) = \sin c = 0$. The number c must be $k\pi$, a multiple of π , and λ is $-c^2$:

$$\begin{array}{l} \text{Eigenvalues } \lambda = -k^2\pi^2 \\ \text{Eigenfunctions } y = \sin k\pi x \end{array} \quad \frac{d^2}{dx^2} \sin k\pi x = -k^2\pi^2 \sin k\pi x. \quad (6)$$

Now we go back to the matrix K and guess its eigenvectors. They come from $\sin k\pi x$ at n points $x = h, 2h, \dots, nh$, equally spaced between 0 and 1. The spacing Δx is $h = 1/(n+1)$, so the $(n+1)$ st point comes out at $(n+1)h = 1$. Multiply that sine vector s by K :

$$\begin{array}{l} \text{Eigenvector of } K = \text{sine vector } s \end{array} \quad \begin{array}{l} Ks = \lambda s = (2 - 2\cos k\pi h) s \\ s = (\sin k\pi h, \dots, \sin nk\pi h). \end{array} \quad (7)$$

I will leave that multiplication $Ks = \lambda s$ as a challenge problem. Notice what is important:

1. All eigenvalues $2 - 2\cos k\pi h$ are positive and K is positive definite.
2. The **sine matrix** S has orthogonal columns = eigenvectors s_1, \dots, s_n of K .

Discrete Sine Transform
The j, k entry is $\sin jk\pi h$

$$S = \begin{bmatrix} \sin \pi h & & \sin k\pi h & & \\ & \vdots & \dots & \vdots & \dots \\ \sin n\pi h & & \sin nk\pi h & & \end{bmatrix}$$

Those eigenvectors are orthogonal just like the eigenfunctions: $\int_0^1 \sin j\pi x \sin k\pi x dx = 0$.

Problem Set 6.5

Problems 1–13 are about tests for positive definiteness.

- 1 Suppose the 2 by 2 tests $a > 0$ and $ac - b^2 > 0$ are passed. Then $c > b^2/a$ is also positive.

- (i) λ_1 and λ_2 have the *same sign* because their product $\lambda_1\lambda_2$ equals _____.
 (i) That sign is positive because $\lambda_1 + \lambda_2$ equals _____.

Conclusion: The tests $a > 0, ac - b^2 > 0$ guarantee positive eigenvalues λ_1, λ_2 .

- 2 Which of A_1, A_2, A_3, A_4 has two positive eigenvalues? Use the test, don't compute the λ 's. Find an \mathbf{x} so that $\mathbf{x}^T A_1 \mathbf{x} < 0$, so A_1 fails the test.

$$A_1 = \begin{bmatrix} 5 & 6 \\ 6 & 7 \end{bmatrix} \quad A_2 = \begin{bmatrix} -1 & -2 \\ -2 & -5 \end{bmatrix} \quad A_3 = \begin{bmatrix} 1 & 10 \\ 10 & 100 \end{bmatrix} \quad A_4 = \begin{bmatrix} 1 & 10 \\ 10 & 101 \end{bmatrix}.$$

- 3 For which numbers b and c are these matrices positive definite?

$$A = \begin{bmatrix} 1 & b \\ b & 9 \end{bmatrix} \quad A = \begin{bmatrix} 2 & 4 \\ 4 & c \end{bmatrix} \quad A = \begin{bmatrix} c & b \\ b & c \end{bmatrix}.$$

With the pivots in D and multiplier in L , factor each A into LDL^T .

- 4 What is the quadratic $f = ax^2 + 2bxy + cy^2$ for each of these matrices? Complete the square to write f as a sum of one or two squares $d_1(\quad)^2 + d_2(\quad)^2$.

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 9 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} 1 & 3 \\ 3 & 9 \end{bmatrix}.$$

- 5 Write $f(x, y) = x^2 + 4xy + 3y^2$ as a *difference* of squares and find a point (x, y) where f is negative. The minimum is not at $(0, 0)$ even though f has positive coefficients.
- 6 The function $f(x, y) = 2xy$ certainly has a saddle point and not a minimum at $(0, 0)$. What symmetric matrix A produces this f ? What are its eigenvalues?

7 Test to see if $R^T R$ is positive definite in each case:

$$R = \begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 2 & 1 \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} 1 & 1 & 2 \\ 1 & 2 & 1 \end{bmatrix}.$$

8 The function $f(x, y) = 3(x + 2y)^2 + 4y^2$ is positive except at $(0, 0)$. What is the matrix in $f = [x \ y] A [x \ y]^T$? Check that the pivots of A are 3 and 4.

9 Find the 3 by 3 matrix A and its pivots, rank, eigenvalues, and determinant:

$$\begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix} \begin{bmatrix} A \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = 4(x_1 - x_2 + 2x_3)^2.$$

10 Which 3 by 3 symmetric matrices A and B produce these quadratics?

$$\mathbf{x}^T A \mathbf{x} = 2(x_1^2 + x_2^2 + x_3^2 - x_1 x_2 - x_2 x_3). \quad \text{Why is } A \text{ positive definite?}$$

$$\mathbf{x}^T B \mathbf{x} = 2(x_1^2 + x_2^2 + x_3^2 - x_1 x_2 - x_1 x_3 - x_2 x_3). \quad \text{Why is } B \text{ semidefinite?}$$

11 Compute the three upper left determinants of A to establish positive definiteness. Verify that their ratios give the second and third pivots.

$$\text{Pivots} = \text{ratios of determinants} \quad A = \begin{bmatrix} 2 & 2 & 0 \\ 2 & 5 & 3 \\ 0 & 3 & 8 \end{bmatrix}.$$

12 For what numbers c and d are A and B positive definite? Test the 3 determinants:

$$A = \begin{bmatrix} c & 1 & 1 \\ 1 & c & 1 \\ 1 & 1 & c \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 2 & 3 \\ 2 & d & 4 \\ 3 & 4 & 5 \end{bmatrix}.$$

13 Find a matrix with $a > 0$ and $c > 0$ and $a + c > 2b$ that has a negative eigenvalue.

Problems 14–20 are about applications of the tests.

14 If A is positive definite then A^{-1} is positive definite. Best proof: The eigenvalues of A^{-1} are positive because _____. Second proof (only for 2 by 2):

$$\text{The entries of } A^{-1} = \frac{1}{ac - b^2} \begin{bmatrix} c & -b \\ -b & a \end{bmatrix} \text{ pass the determinant tests } \underline{\hspace{2cm}}.$$

15 If A and B are positive definite, their sum $A + B$ is positive definite. Pivots and eigenvalues are not convenient for $A + B$. Better to prove $\mathbf{x}^T (A + B) \mathbf{x} > 0$. Or if $A = R^T R$ and $B = S^T S$, show that $A + B = [R \ S]^T \begin{bmatrix} R \\ S \end{bmatrix}$ with independent columns.

- 16 A positive definite matrix cannot have a zero (or even worse, a negative number) on its diagonal. Show that this matrix fails to have $\mathbf{x}^T A \mathbf{x} > 0$:

$$\begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix} \begin{bmatrix} 4 & 1 & 1 \\ 1 & 0 & 2 \\ 1 & 2 & 5 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \text{ is not positive when } (x_1, x_2, x_3) = (\quad , \quad , \quad).$$

- 17 A diagonal entry a_{jj} of a symmetric matrix cannot be smaller than all the λ 's. If it were, then $A - a_{jj}I$ would have _____ eigenvalues and would be positive definite. But $A - a_{jj}I$ has a _____ on the main diagonal.

- 18 If $A\mathbf{x} = \lambda\mathbf{x}$ then $\mathbf{x}^T A \mathbf{x} = \underline{\hspace{2cm}}$. If $\mathbf{x}^T A \mathbf{x} > 0$, prove that $\lambda > 0$.

- 19 Reverse Problem 18 to show that *if all $\lambda > 0$ then $\mathbf{x}^T A \mathbf{x} > 0$* . We must do this for *every* nonzero \mathbf{x} , not just the eigenvectors. So write \mathbf{x} as a combination of the eigenvectors and *explain why all "cross terms" are $\mathbf{x}_i^T \mathbf{x}_j = 0$* . Then $\mathbf{x}^T A \mathbf{x}$ is

$$(c_1 \mathbf{x}_1 + \cdots + c_n \mathbf{x}_n)^T (c_1 \lambda_1 \mathbf{x}_1 + \cdots + c_n \lambda_n \mathbf{x}_n) = c_1^2 \lambda_1 \mathbf{x}_1^T \mathbf{x}_1 + \cdots + c_n^2 \lambda_n \mathbf{x}_n^T \mathbf{x}_n > 0.$$

- 20 Give a quick reason why each of these statements is true:

- Every positive definite matrix is invertible.
- The only positive definite projection matrix is $P = I$.
- A diagonal matrix with positive diagonal entries is positive definite.
- A symmetric matrix with a positive determinant might not be positive definite!

Problems 21–24 use the eigenvalues; Problems 25–27 are based on pivots.

- 21 For which s and t do A and B have all $\lambda > 0$ (therefore positive definite)?

$$A = \begin{bmatrix} s & -4 & -4 \\ -4 & s & -4 \\ -4 & -4 & s \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} t & 3 & 0 \\ 3 & t & 4 \\ 0 & 4 & t \end{bmatrix}.$$

- 22 From $A = Q\Lambda Q^T$ compute the positive definite symmetric square root $Q\Lambda^{1/2}Q^T$ of each matrix. Check that this square root gives $R^2 = A$:

$$A = \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} 10 & 6 \\ 6 & 10 \end{bmatrix}.$$

- 23 You may have seen the equation for an ellipse as $x^2/a^2 + y^2/b^2 = 1$. What are a and b when the equation is written $\lambda_1 x^2 + \lambda_2 y^2 = 1$? The ellipse $9x^2 + 4y^2 = 1$ has axes with half-lengths $a = \underline{\hspace{2cm}}$ and $b = \underline{\hspace{2cm}}$.

- 24 Draw the tilted ellipse $x^2 + xy + y^2 = 1$ and find the half-lengths of its axes from the eigenvalues of the corresponding matrix A .

- 25** With positive pivots in D , the factorization $A = LDL^T$ becomes $L\sqrt{D}\sqrt{D}L^T$. (Square roots of the pivots give $D = \sqrt{D}\sqrt{D}$.) Then $C = \sqrt{D}L^T$ yields the *Cholesky factorization* $A = C^T C$ which is “symmetrized LU ”:

$$\text{From } C = \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix} \text{ find } A. \quad \text{From } A = \begin{bmatrix} 4 & 8 \\ 8 & 25 \end{bmatrix} \text{ find } C = \mathbf{chol}(A).$$

- 26** In the Cholesky factorization $A = C^T C$, with $C^T = L\sqrt{D}$, the square roots of the pivots are on the diagonal of C . Find C (upper triangular) for

$$A = \begin{bmatrix} 9 & 0 & 0 \\ 0 & 1 & 2 \\ 0 & 2 & 8 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 2 \\ 1 & 2 & 7 \end{bmatrix}.$$

- 27** The symmetric factorization $A = LDL^T$ means that $x^T A x = x^T L D L^T x$:

$$\begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & 0 \\ b/a & 1 \end{bmatrix} \begin{bmatrix} a & 0 \\ 0 & (ac - b^2)/a \end{bmatrix} \begin{bmatrix} 1 & b/a \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}.$$

The left side is $ax^2 + 2bxy + cy^2$. The right side is $a(x + \frac{b}{a}y)^2 + \text{---}y^2$. The second pivot completes the square! Test with $a = 2, b = 4, c = 10$.

- 28** Without multiplying $A = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 5 \end{bmatrix} \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$, find

- (a) the determinant of A (b) the eigenvalues of A
 (c) the eigenvectors of A (d) a reason why A is symmetric positive definite.

- 29** For $F_1(x, y) = \frac{1}{4}x^4 + x^2y + y^2$ and $F_2(x, y) = x^3 + xy - x$ find the second derivative matrices H_1 and H_2 :

$$\text{Test for minimum: } H = \begin{bmatrix} \partial^2 F / \partial x^2 & \partial^2 F / \partial x \partial y \\ \partial^2 F / \partial y \partial x & \partial^2 F / \partial y^2 \end{bmatrix} \text{ is positive definite}$$

H_1 is positive definite so F_1 is concave up (= convex). Find the minimum point of F_1 and the saddle point of F_2 (look only where first derivatives are zero).

- 30** The graph of $z = x^2 + y^2$ is a bowl opening upward. The graph of $z = x^2 - y^2$ is a saddle. The graph of $z = -x^2 - y^2$ is a bowl opening downward. What is a test on a, b, c for $z = ax^2 + 2bxy + cy^2$ to have a saddle point at $(0, 0)$?

- 31** Which values of c give a bowl and which c give a saddle point for the graph of $z = 4x^2 + 12xy + cy^2$? Describe this graph at the borderline value of c .

Challenge Problems

- 32** A *group* of nonsingular matrices includes AB and A^{-1} if it includes A and B . “Products and inverses stay in the group.” Which of these are groups (as in 2.7.37)? Invent a “subgroup” of two of these groups (not I by itself = the smallest group).
- (a) Positive definite symmetric matrices A .
 - (b) Orthogonal matrices Q .
 - (c) All exponentials e^{tA} of a fixed matrix A .
 - (d) Matrices P with positive eigenvalues.
 - (e) Matrices D with determinant 1.
- 33** When A and B are symmetric positive definite, AB might not even be symmetric. But its eigenvalues are still positive. Start from $ABx = \lambda x$ and take dot products with Bx . Then prove $\lambda > 0$.
- 34** Write down the 5 by 5 sine matrix S from Worked Example 6.5 D, containing the eigenvectors of K when $n = 5$ and $h = 1/6$. Multiply K times S to see the five positive eigenvalues. Their sum should equal the trace 10. Their product should be $\det K = 6$.
- 35** Suppose C is positive definite (so $y^T C y > 0$ whenever $y \neq \mathbf{0}$) and A has independent columns (so $Ax \neq \mathbf{0}$ whenever $x \neq \mathbf{0}$). Apply the energy test to $x^T A^T C A x$ to show that $A^T C A$ is positive definite: *the crucial matrix in engineering*.

6.6 Similar Matrices

The key step in this chapter is to diagonalize a matrix by using its eigenvectors. When S is the eigenvector matrix, the diagonal matrix $S^{-1}AS$ is Λ —the eigenvalue matrix. But diagonalization is not possible for every A . Some matrices have too few eigenvectors—we had to leave them alone. In this new section, the eigenvector matrix S remains the best choice when we can find it, *but now we allow any invertible matrix M* .

Starting from A we go to $M^{-1}AM$. This matrix may be diagonal—probably not. It still shares important properties of A . No matter which M we choose, *the eigenvalues stay the same*. The matrices A and $M^{-1}AM$ are called “similar”. A typical matrix A is similar to a whole family of other matrices because there are so many choices of M .

DEFINITION Let M be any invertible matrix. Then $B = M^{-1}AM$ is *similar* to A .

If $B = M^{-1}AM$ then immediately $A = MBM^{-1}$. That means: If B is similar to A then A is similar to B . The matrix in this reverse direction is M^{-1} —just as good as M .

A diagonalizable matrix is similar to Λ . In that special case M is S . We have $A = S\Lambda S^{-1}$ and $\Lambda = S^{-1}AS$. They certainly have the same eigenvalues! This section is opening up to other similar matrices $B = M^{-1}AM$, by allowing all invertible M .

The combination $M^{-1}AM$ appears when we change variables in a differential equation. Start with an equation for u and set $u = Mv$:

$$\frac{du}{dt} = Au \quad \text{becomes} \quad M \frac{dv}{dt} = AMv \quad \text{which is} \quad \frac{dv}{dt} = M^{-1}AMv.$$

The original coefficient matrix was A , the new one at the right is $M^{-1}AM$. Changing u to v leads to a similar matrix. When $M = S$ the new system is diagonal—the maximum in simplicity. Other choices of M could make the new system triangular and easier to solve. Since we can always go back to u , similar matrices must give the same growth or decay. More precisely, *the eigenvalues of A and B are the same*.

(No change in λ 's) Similar matrices A and $M^{-1}AM$ have the same eigenvalues. If x is an eigenvector of A , then $M^{-1}x$ is an eigenvector of $B = M^{-1}AM$.

The proof is quick, since $B = M^{-1}AM$ gives $A = MBM^{-1}$. Suppose $Ax = \lambda x$:

$$MBM^{-1}x = \lambda x \quad \text{means that} \quad B(M^{-1}x) = \lambda(M^{-1}x).$$

The eigenvalue of B is the same λ . The eigenvector has changed to $M^{-1}x$.

Two matrices can have the same *repeated* λ , and fail to be similar—as we will see.

Example 1 These matrices $M^{-1}AM$ all have the same eigenvalues 1 and 0.

The projection $A = \begin{bmatrix} .5 & .5 \\ .5 & .5 \end{bmatrix}$ is similar to $\Lambda = S^{-1}AS = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$

Now choose $M = \begin{bmatrix} 1 & 0 \\ 1 & 2 \end{bmatrix}$. The similar matrix $M^{-1}AM$ is $\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$.

Also choose $M = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$. The similar matrix $M^{-1}AM$ is $\begin{bmatrix} .5 & -.5 \\ -.5 & .5 \end{bmatrix}$.

All 2 by 2 matrices with those eigenvalues 1 and 0 are similar to each other. The eigenvectors change with M , the eigenvalues don't change.

The eigenvalues in that example are *not repeated*. This makes life easy. Repeated eigenvalues are harder. The next example has eigenvalues 0 and 0. The zero matrix shares those eigenvalues, but it is similar only to itself: $M^{-1}0M = 0$.

Example 2 A family of similar matrices with $A = 0, 0$ (repeated eigenvalue)

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \text{ is similar to } \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix} \text{ and all } B = \begin{bmatrix} cd & d^2 \\ -c^2 & -cd \end{bmatrix} \text{ except } \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

These matrices B all have zero determinant (like A). They all have rank one (like A). One eigenvalue is zero and the trace is $cd - dc = 0$, so the other must be zero. I chose any $M = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ with $ad - bc = 1$, and $B = M^{-1}AM$.

These matrices B can't be diagonalized. In fact A is as close to diagonal as possible. It is the "**Jordan form**" for the family of matrices B . This is the outstanding member (my class says "Godfather") of the family. The Jordan form $J = A$ is as near as we can come to diagonalizing these matrices, when there is only one eigenvector. In going from A to $B = M^{-1}AM$, some things change and some don't. Here is a table to show this.

<u>Not changed by M</u>	<u>Changed by M</u>
Eigenvalues	Eigenvectors
Trace and determinant	Nullspace
Rank	Column space
Number of independent eigenvectors	Row space
Jordan form	Left nullspace
	Singular values

The eigenvalues don't change for similar matrices; the eigenvectors do. The trace is the sum of the λ 's (unchanged). The determinant is the product of the same λ 's.¹ The nullspace consists of the eigenvectors for $\lambda = 0$ (if any), so it can change. Its dimension $n - r$ does not change! The *number* of eigenvectors stays the same for each λ , while the vectors themselves are multiplied by M^{-1} . The *singular values* depend on $A^T A$, which definitely changes. They come in the next section.

¹The determinant is unchanged because $\det B = (\det M^{-1})(\det A)(\det M) = \det A$.

Examples of the Jordan Form

The *Jordan form* is the serious new idea here. We lead up to it with one more example of similar matrices: *triple eigenvalue, one eigenvector*.

Example 3 This **Jordan matrix** J has $\lambda = 5, 5, 5$ on its diagonal. Its only eigenvectors are multiples of $x = (1, 0, 0)$. Algebraic multiplicity is 3, geometric multiplicity is 1:

$$\text{If } J = \begin{bmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{bmatrix} \text{ then } J - 5I = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \text{ has rank 2.}$$

Every similar matrix $B = M^{-1}JM$ has the same triple eigenvalue 5, 5, 5. Also $B - 5I$ must have the same rank 2. Its nullspace has dimension 1. So every B that is similar to this “Jordan block” J has only one independent eigenvector $M^{-1}x$.

The transpose matrix J^T has the same eigenvalues 5, 5, 5, and $J^T - 5I$ has the same rank 2. *Jordan’s theorem says that J^T is similar to J* . The matrix M that produces the similarity happens to be the *reverse identity*:

$$J^T = M^{-1}JM \text{ is } \begin{bmatrix} 5 & 0 & 0 \\ 1 & 5 & 0 \\ 0 & 1 & 5 \end{bmatrix} = \begin{bmatrix} & & 1 \\ & 1 & \\ 1 & & \end{bmatrix} \begin{bmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{bmatrix} \begin{bmatrix} & & 1 \\ & 1 & \\ 1 & & \end{bmatrix}.$$

All blank entries are zero. An eigenvector of J^T is $M^{-1}(1, 0, 0) = (0, 0, 1)$. There is one line of eigenvectors $(x_1, 0, 0)$ for J and another line $(0, 0, x_3)$ for J^T .

The key fact is that this matrix J is similar to *every* matrix A with eigenvalues 5, 5, 5 and *one line* of eigenvectors. There is an M with $M^{-1}AM = J$.

Example 4 Since J is as close to diagonal as we can get, the equation $du/dt = Ju$ cannot be simplified by changing variables. We must solve it as it stands:

$$\frac{du}{dt} = Ju = \begin{bmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \text{ is } \begin{cases} dx/dt = 5x + y \\ dy/dt = 5y + z \\ dz/dt = 5z. \end{cases}$$

The system is triangular. We think naturally of back substitution. Solve the last equation and work upwards. Main point: *All solutions contain e^{5t} since $\lambda = 5$* :

$$\begin{aligned} \text{Last equation } \frac{dz}{dt} &= 5z && \text{yields } z = z(0)e^{5t} \\ \text{Notice } te^{5t} \quad \frac{dy}{dt} &= 5y + z && \text{yields } y = (y(0) + tz(0))e^{5t} \\ \text{Notice } t^2e^{5t} \quad \frac{dx}{dt} &= 5x + y && \text{yields } x = (x(0) + ty(0) + \frac{1}{2}t^2z(0))e^{5t}. \end{aligned}$$

The two missing eigenvectors are responsible for the te^{5t} and t^2e^{5t} terms in y and x . The factors t and t^2 enter because $\lambda = 5$ is a triple eigenvalue with one eigenvector.

Note Chapter 7 will explain another approach to similar matrices. Instead of changing variables by $u = Mv$, we “*change the basis*”. In this approach, similar matrices will represent the same transformation of n -dimensional space. When we choose a basis for \mathbf{R}^n , we get a matrix. The standard basis vectors ($M = I$) lead to $I^{-1}AI$ which is A . Other bases lead to similar matrices $B = M^{-1}AM$.

The Jordan Form

For every A , we want to choose M so that $M^{-1}AM$ is as *nearly diagonal as possible*. When A has a full set of n eigenvectors, they go into the columns of M . Then $M = S$. The matrix $S^{-1}AS$ is diagonal, period. This matrix Λ is the Jordan form of A —when A can be diagonalized. In the general case, eigenvectors are missing and Λ can't be reached.

Suppose A has s independent eigenvectors. Then it is similar to a matrix with s blocks. Each block is like J in Example 3. *The eigenvalue is on the diagonal with 1's just above it.* This block accounts for one eigenvector of A . When there are n eigenvectors and n blocks, they are all 1 by 1. In that case J is Λ .

(Jordan form) If A has s independent eigenvectors, it is similar to a matrix J that has s Jordan blocks on its diagonal: Some matrix M puts A into Jordan form:

$$\text{Jordan form} \quad M^{-1}AM = \begin{bmatrix} J_1 & & \\ & \ddots & \\ & & J_s \end{bmatrix} = J. \quad (1)$$

Each block in J has one eigenvalue λ_i , one eigenvector, and 1's above the diagonal:

$$\text{Jordan block} \quad J_i = \begin{bmatrix} \lambda_i & 1 & & \\ & \cdot & \cdot & \\ & & \cdot & 1 \\ & & & \lambda_i \end{bmatrix}. \quad (2)$$

A is similar to B if they share the same Jordan form J —not otherwise.

The Jordan form J has an off-diagonal 1 for each missing eigenvector (and the 1's are next to the eigenvalues). This is the big theorem about matrix similarity. In every family of similar matrices, we are picking one outstanding member called J . It is nearly diagonal (or if possible completely diagonal). For that J , we can solve $du/dt = Ju$ as in Example 4. We can take powers J^k as in Problems 9–10. Every other matrix in the family has the form $A = MJM^{-1}$. The connection through M solves $du/dt = Au$.

The point you must see is that $MJM^{-1}MJM^{-1} = MJ^2M^{-1}$. That cancellation of $M^{-1}M$ in the middle has been used through this chapter (when M was S). We found A^{100} from $S\Lambda^{100}S^{-1}$ —by diagonalizing the matrix. Now we can't quite diagonalize A . So we use $MJ^{100}M^{-1}$ instead.

Jordan's Theorem is proved in my textbook *Linear Algebra and Its Applications*. Please refer to that book (or more advanced books) for the proof. The reasoning is rather intricate and in actual computations the Jordan form is not at all popular—its calculation is not stable. A slight change in A will separate the repeated eigenvalues and remove the off-diagonal 1's—switching to a diagonal Λ .

Proved or not, you have caught the central idea of similarity—to make A as simple as possible while preserving its essential properties.

■ REVIEW OF THE KEY IDEAS ■

1. B is similar to A if $B = M^{-1}AM$, for some invertible matrix M .
2. Similar matrices have the same eigenvalues. Eigenvectors are multiplied by M^{-1} .
3. If A has n independent eigenvectors then A is similar to Λ (take $M = S$).
4. Every matrix is similar to a Jordan matrix J (which has Λ as its diagonal part). J has a block for each eigenvector, and 1's for missing eigenvectors.

■ WORKED EXAMPLES ■

6.6 A The 4 by 4 triangular Pascal matrix A and its inverse (alternating diagonals) are

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 1 & 3 & 3 & 1 \end{bmatrix} \quad \text{and} \quad A^{-1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 1 & -2 & 1 & 0 \\ -1 & 3 & -3 & 1 \end{bmatrix}.$$

Check that A and A^{-1} have the same eigenvalues. Find a diagonal matrix D with alternating signs that gives $A^{-1} = D^{-1}AD$. This A is similar to A^{-1} , which is unusual.

These similar matrices must have the same Jordan form J . This J has only one block because the Pascal matrix has only one line of eigenvectors.

Solution The triangular matrices A and A^{-1} both have $\lambda = 1, 1, 1, 1$ on their main diagonals. Choose D with alternating 1 and -1 on its diagonal. D equals D^{-1} :

$$D^{-1}AD = \begin{bmatrix} -1 & & & \\ & 1 & & \\ & & -1 & \\ & & & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 1 & 3 & 3 & 1 \end{bmatrix} \begin{bmatrix} -1 & & & \\ & 1 & & \\ & & -1 & \\ & & & 1 \end{bmatrix} = A^{-1}.$$

Check: Changing signs in rows 1 and 3 of A , and columns 1 and 3, produces the four negative entries in A^{-1} . We are multiplying row i by $(-1)^i$ and column j by $(-1)^j$, which gives the alternating diagonals in A^{-1} . Then AD has *columns with alternating signs*.

6.6 B *The best way to compute eigenvalues of a large matrix is not from solving $\det(A - \lambda I) = 0$. That high degree polynomial is a numerical disaster.*

Instead we compute similar matrices A_1, A_2, \dots that approach a triangular matrix. Then the eigenvalues of A (unchanged) are almost sitting on the main diagonal.

One way is to factor $A = QR$ by “Gram-Schmidt”. Reverse the order to $A_1 = RQ$. This matrix is similar to A because $RQ = Q^{-1}(QR)Q$. An example with $c = \cos \theta$ and $s = \sin \theta$ shows how a small off-diagonal s can be *cubed* in A_1 :

$$A = \begin{bmatrix} c & s \\ s & 0 \end{bmatrix} \text{ factors into } \begin{bmatrix} c & s \\ s & -c \end{bmatrix} \begin{bmatrix} 1 & cs \\ 0 & s^2 \end{bmatrix} = QR.$$

$$A_1 = RQ = \begin{bmatrix} c + cs^2 & s^3 \\ s^3 & -cs^2 \end{bmatrix} \text{ has } s^3 \text{ below the diagonal}$$

Another step can factor $A_1 = Q_1 R_1$ and reverse to $A_2 = R_1 Q_1$. This **QR method** is in Section 9.3 with a further improvement for A_1 . Add cs^2 to its diagonal (to get zero in the corner) and then subtract back from A_2 :

$$\text{Shift and factor } A_1 + cs^2 I = Q_1 R_1 \quad \text{Reverse and shift back } A_2 = R_1 Q_1 - cs^2 I$$

Shifted QR is an amazing success—just about the best way to compute eigenvalues.

Problem Set 6.6

- 1 If $C = F^{-1}AF$ and also $C = G^{-1}BG$, what matrix M gives $B = M^{-1}AM$?
Conclusion: If C is similar to A and also to B then _____.
- 2 If $A = \text{diag}(1, 3)$ and $B = \text{diag}(3, 1)$ show that A and B are similar (find an M).
- 3 Show that A and B are similar by finding M so that $B = M^{-1}AM$:

$$\begin{array}{ll} A = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} & \text{and } B = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \\ A = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} & \text{and } B = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \\ A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} & \text{and } B = \begin{bmatrix} 4 & 3 \\ 2 & 1 \end{bmatrix}. \end{array}$$

- 4 If a 2 by 2 matrix A has eigenvalues 0 and 1, why is it similar to $\Lambda = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$?
Deduce from Problem 1 that all 2 by 2 matrices with those eigenvalues are similar.
- 5 Which of these six matrices are similar? Check their eigenvalues.

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}.$$

- 6 There are sixteen 2 by 2 matrices whose entries are 0's and 1's. Similar matrices go into the same family. How many families? How many matrices (total 16) in each family?
- 7 (a) If x is in the nullspace of A show that $M^{-1}x$ is in the nullspace of $M^{-1}AM$.
 (b) The nullspaces of A and $M^{-1}AM$ have the same (vectors)(basis)(dimension).
- 8 Suppose $Ax = \lambda x$ and $Bx = \lambda x$ with the same λ 's and x 's. With n independent eigenvectors we have $A = B$: Why? Find $A \neq B$ when both have eigenvalues 0, 0 but only one line of eigenvectors $(x_1, 0)$.
- 9 By direct multiplication find A^2 and A^3 and A^5 when

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

Guess the form of A^k . Set $k = 0$ to find A^0 and $k = -1$ to find A^{-1} .

Questions 10–14 are about the Jordan form.

- 10 By direct multiplication, find J^2 and J^3 when

$$J = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}.$$

Guess the form of J^k . Set $k = 0$ to find J^0 . Set $k = -1$ to find J^{-1} .

- 11 Solve $du/dt = Ju$ for J in Problem 10, starting from $u(0) = (5, 2)$. Remember $te^{\lambda t}$.
- 12 These Jordan matrices have eigenvalues 0, 0, 0, 0. They have two eigenvectors (one from each block). But the block sizes don't match and they are *not similar*:

$$J = \left[\begin{array}{cc|cc} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{array} \right] \quad \text{and} \quad K = \left[\begin{array}{ccc|c} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 \end{array} \right].$$

For any matrix M , compare JM with MK . If they are equal show that M is not invertible. Then $M^{-1}JM = K$ is impossible: J is not similar to K .

- 13 Based on Problem 12, what are the five Jordan forms when $\lambda = 0, 0, 0, 0$?
- 14 Prove that A^T is always similar to A (we know the λ 's are the same):
1. For one Jordan block J_i : Find M_i so that $M_i^{-1}J_iM_i = J_i^T$ (see Example 3).
 2. For any J with blocks J_i : Build M_0 from blocks so that $M_0^{-1}JM_0 = J^T$.
 3. For any $A = MJM^{-1}$: Show that A^T is similar to J^T and so to J and to A .

15 Prove that $\det(A - \lambda I) = \det(M^{-1}AM - \lambda I)$. (You could write $I = M^{-1}M$ and factor out $\det M^{-1}$ and $\det M$.) Since these *characteristic polynomials* of A and $M^{-1}AM$ are the same, the eigenvalues are the same (with the same multiplicities).

16 Which pairs are similar? Choose a, b, c, d to prove that the other pairs aren't:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \begin{bmatrix} b & a \\ d & c \end{bmatrix} \quad \begin{bmatrix} c & d \\ a & b \end{bmatrix} \quad \begin{bmatrix} d & c \\ b & a \end{bmatrix}.$$

17 True or false, with a good reason:

- (a) A symmetric matrix can't be similar to a nonsymmetric matrix.
- (b) An invertible matrix can't be similar to a singular matrix.
- (c) A can't be similar to $-A$ unless $A = 0$.
- (d) A can't be similar to $A + I$.

18 If B is invertible, prove that AB is similar to BA . *They have the same eigenvalues.*

19 If A is 6 by 4 and B is 4 by 6, AB and BA have different sizes. But with blocks

$$M^{-1}FM = \begin{bmatrix} I & -A \\ 0 & I \end{bmatrix} \begin{bmatrix} AB & 0 \\ B & 0 \end{bmatrix} \begin{bmatrix} I & A \\ 0 & I \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ B & BA \end{bmatrix} = G.$$

- (a) What sizes are the four blocks (the same four sizes in each matrix)?
- (b) This equation is $M^{-1}FM = G$, so F and G have the same 10 eigenvalues. F has the 6 eigenvalues of AB plus 4 zeros; G has the 4 eigenvalues of BA plus 6 zeros. **AB has the same eigenvalues as BA plus _____ zeros.**

20 Why are these statements all true?

- (a) If A is similar to B then A^2 is similar to B^2 .
- (b) A^2 and B^2 can be similar when A and B are not similar (try $\lambda = 0, 0$).
- (c) $\begin{bmatrix} 3 & 0 \\ 0 & 4 \end{bmatrix}$ is similar to $\begin{bmatrix} 3 & 1 \\ 0 & 4 \end{bmatrix}$.
- (d) $\begin{bmatrix} 3 & 0 \\ 0 & 3 \end{bmatrix}$ is not similar to $\begin{bmatrix} 3 & 1 \\ 0 & 3 \end{bmatrix}$.
- (e) If we exchange rows 1 and 2 of A , and then exchange columns 1 and 2, **the eigenvalues stay the same.** In this case $M = \underline{\hspace{2cm}}$.

21 If J is the 5 by 5 Jordan block with $\lambda = 0$, find J^2 and count its eigenvectors and find its Jordan form (there will be two blocks).

Challenge Problems

22 If an n by n matrix A has all eigenvalues $\lambda = 0$, prove that $A^n =$ zero matrix. (Maybe prove first that $J^n =$ zero matrix, by direct multiplication. Or use the Cayley-Hamilton Theorem?)

23 For the shifted QR method in the Worked Example 6.6 B, show that A_2 is similar to A_1 . No change in eigenvalues, and the A 's quickly approach a diagonal matrix.

24 If A is similar to A^{-1} , must all the eigenvalues equal 1 or -1 ?

6.7 Singular Value Decomposition (SVD)

The Singular Value Decomposition is a highlight of linear algebra. A is any m by n matrix, square or rectangular. Its rank is r . We will diagonalize this A , but not by $S^{-1}AS$. The eigenvectors in S have three big problems: They are usually not orthogonal, there are not always enough eigenvectors, and $Ax = \lambda x$ requires A to be square. The *singular vectors* of A solve all those problems in a perfect way.

The price we pay is to have two sets of singular vectors, u 's and v 's. The u 's are eigenvectors of AA^T and the v 's are eigenvectors of $A^T A$. Since those matrices are both symmetric, their eigenvectors can be chosen orthonormal. In equation (13) below, the simple fact that A times $A^T A$ is the same as AA^T times A will lead to a remarkable property of these u 's and v 's:

$$\text{"A is diagonalized"} \quad Av_1 = \sigma_1 u_1 \quad Av_2 = \sigma_2 u_2 \quad \dots \quad Av_r = \sigma_r u_r \quad (1)$$

The singular vectors v_1, \dots, v_r are in the *row space* of A . The outputs u_1, \dots, u_r are in the *column space* of A . The *singular values* $\sigma_1, \dots, \sigma_r$ are all positive numbers. When the v 's and u 's go into the columns of V and U , orthogonality gives $V^T V = I$ and $U^T U = I$. The σ 's go into a diagonal matrix Σ .

Just as $Ax_i = \lambda_i x_i$ led to the diagonalization $AS = S\Lambda$, the equations $Av_i = \sigma_i u_i$ tell us column by column that $AV = U\Sigma$:

$$\begin{array}{l} (m \text{ by } n)(n \text{ by } r) \\ \text{equals} \\ (m \text{ by } r)(r \text{ by } r) \end{array} A \begin{bmatrix} v_1 & \cdots & v_r \end{bmatrix} = \begin{bmatrix} u_1 & \cdots & u_r \end{bmatrix} \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix}. \quad (2)$$

This is the heart of the SVD, but there is more. Those v 's and u 's account for the row space and column space of A . We need $n - r$ more v 's and $m - r$ more u 's, from the nullspace $N(A)$ and the left nullspace $N(A^T)$. They can be orthonormal bases for those two nullspaces (and then automatically orthogonal to the first r v 's and u 's). Include all the v 's and u 's in V and U , so these matrices become *square*. *We still have* $AV = U\Sigma$.

$$\begin{array}{l} (m \text{ by } n)(n \text{ by } n) \\ \text{equals} \\ (m \text{ by } m)(m \text{ by } n) \end{array} A \begin{bmatrix} v_1 & \cdots & v_r & \cdots & v_n \end{bmatrix} = \begin{bmatrix} u_1 & \cdots & u_r & \cdots & u_m \end{bmatrix} \begin{bmatrix} \sigma_1 & & & & \\ & \ddots & & & \\ & & \sigma_r & & \\ & & & \ddots & \\ & & & & 0 \end{bmatrix} \quad (3)$$

The new Σ is m by n . It is just the old r by r matrix (call that Σ_r) with $m - r$ new zero rows and $n - r$ new zero columns. The real change is in the shapes of U and V and Σ . Still $V^T V = I$ and $U^T U = I$, with sizes n and m .

V is now a square orthogonal matrix, with inverse $V^{-1} = V^T$. So $AV = U\Sigma$ can become $A = U\Sigma V^T$. This is the *Singular Value Decomposition*:

$$\text{SVD} \quad A = U\Sigma V^T = u_1 \sigma_1 v_1^T + \cdots + u_r \sigma_r v_r^T. \quad (4)$$

I would write the earlier “reduced SVD” from equation (2) as $A = U_r \Sigma_r V_r^T$. That is equally true, without the extra zeros in Σ . This reduced SVD gives the same splitting of A into a sum of r matrices, each of rank one.

We will see that $\sigma_i^2 = \lambda_i$ is an eigenvalue of $A^T A$ and also AA^T . When we put the singular values in descending order, $\sigma_1 \geq \sigma_2 \geq \dots \sigma_r > 0$, the splitting in equation (4) gives the r rank-one pieces of A *in order of importance*.

Example 1 When is $U \Sigma V^T$ (singular values) the same as $S \Lambda S^{-1}$ (eigenvalues)?

Solution We need orthonormal eigenvectors in $S = U$. We need nonnegative eigenvalues in $\Lambda = \Sigma$. So A must be a *positive semidefinite (or definite) symmetric matrix* $Q \Lambda Q^T$.

Example 2 If $A = xy^T$ with unit vectors x and y , what is the SVD of A ?

Solution The reduced SVD in (2) is exactly xy^T , with rank $r = 1$. It has $u_1 = x$ and $v_1 = y$ and $\sigma_1 = 1$. For the full SVD, complete $u_1 = x$ to an orthonormal basis of u 's, and complete $v_1 = y$ to an orthonormal basis of v 's. No new σ 's.

I will describe an application before proving that $Av_i = \sigma_i u_i$. This key equation gave the diagonalizations (2) and (3) and (4) of the SVD: $A = U \Sigma V^T$.

Image Compression

Unusually, I am going to stop the theory and describe applications. This is the century of data, and often that data is stored in a matrix. *A digital image is really a matrix of pixel values*. Each little picture element or “pixel” has a gray scale number between black and white (it has three numbers for a color picture). The picture might have $512 = 2^9$ pixels in each row and $256 = 2^8$ pixels down each column. We have a 256 by 512 pixel matrix with 2^{17} entries! To store one picture, the computer has no problem. But a CT or MR scan produces an image at every cross section—a ton of data. If the pictures are frames in a movie, 30 frames a second means 108,000 images per hour. Compression is especially needed for high definition digital TV, or the equipment could not keep up in real time.

What is compression? We want to replace those 2^{17} matrix entries by a smaller number, *without losing picture quality*. A simple way would be to use larger pixels—replace groups of four pixels by their average value. This is 4 : 1 compression. But if we carry it further, like 16 : 1, our image becomes “blocky”. We want to replace the mn entries by a smaller number, in a way that the human visual system won't notice.

Compression is a billion dollar problem and everyone has ideas. Later in this book I will describe Fourier transforms (used in **jpeg**) and wavelets (now in **JPEG2000**). Here we try an SVD approach: *Replace the 256 by 512 pixel matrix by a matrix of rank one: a column times a row*. If this is successful, the storage requirement becomes $256 + 512$ (add instead of multiply). The compression ratio $(256)(512)/(256 + 512)$ is better than 170 to 1. This is more than we hope for. We may actually use five matrices of rank one (so a matrix approximation of rank 5). The compression is still 34 : 1 and the crucial question is the picture quality.

Where does the SVD come in? *The best rank one approximation to A is the matrix $\sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$.* It uses the largest singular value σ_1 . The best rank 5 approximation includes also $\sigma_2 \mathbf{u}_2 \mathbf{v}_2^T + \cdots + \sigma_5 \mathbf{u}_5 \mathbf{v}_5^T$. **The SVD puts the pieces of A in descending order.**

A library compresses a different matrix. The rows correspond to key words. Columns correspond to titles in the library. The entry in this *word-title matrix* is $a_{ij} = 1$ if word i is in title j (otherwise $a_{ij} = 0$). We normalize the columns so long titles don't get an advantage. We might use a table of contents or an abstract. (Other books might share the title "*Introduction to Linear Algebra*".) Instead of $a_{ij} = 1$, the entries of A can include the *frequency* of the search words. See Section 8.6 for the SVD in statistics.

Once the indexing matrix is created, the search is a linear algebra problem. This giant matrix has to be compressed. The SVD approach gives an optimal low rank approximation, better for library matrices than for natural images. There is an ever-present tradeoff in the cost to compute the \mathbf{u} 's and \mathbf{v} 's. We still need a better way (with sparse matrices).

The Bases and the SVD

Start with a 2 by 2 matrix. Let its rank be $r = 2$, so A is invertible. We want \mathbf{v}_1 and \mathbf{v}_2 to be perpendicular unit vectors. **We also want $A\mathbf{v}_1$ and $A\mathbf{v}_2$ to be perpendicular.** (This is the tricky part. It is what makes the bases special.) Then the unit vectors $\mathbf{u}_1 = A\mathbf{v}_1 / \|A\mathbf{v}_1\|$ and $\mathbf{u}_2 = A\mathbf{v}_2 / \|A\mathbf{v}_2\|$ will be orthonormal. Here is a specific example:

$$\text{Unsymmetric matrix} \quad A = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix}. \quad (5)$$

No orthogonal matrix Q will make $Q^{-1}AQ$ diagonal. We need $U^{-1}AV$. The two bases will be different—one basis cannot do it. The output is $A\mathbf{v}_1 = \sigma_1 \mathbf{u}_1$ when the input is \mathbf{v}_1 . The "singular values" σ_1 and σ_2 are the lengths $\|A\mathbf{v}_1\|$ and $\|A\mathbf{v}_2\|$.

$$\begin{aligned} AV &= U\Sigma \\ A &= U\Sigma V^T \end{aligned} \quad A \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 \end{bmatrix} = \begin{bmatrix} \sigma_1 \mathbf{u}_1 & \sigma_2 \mathbf{u}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 \end{bmatrix} \begin{bmatrix} \sigma_1 & \\ & \sigma_2 \end{bmatrix}. \quad (6)$$

There is a neat way to remove U and see V by itself. *Multiply A^T times A .*

$$A^T A = (U\Sigma V^T)^T (U\Sigma V^T) = V\Sigma^T \Sigma V^T. \quad (7)$$

$U^T U$ disappears because it equals I . (We require $\mathbf{u}_1^T \mathbf{u}_1 = 1 = \mathbf{u}_2^T \mathbf{u}_2$ and $\mathbf{u}_1^T \mathbf{u}_2 = 0$.) Multiplying those diagonal Σ^T and Σ gives σ_1^2 and σ_2^2 . That leaves an ordinary diagonalization of the crucial symmetric matrix $A^T A$, whose eigenvalues are σ_1^2 and σ_2^2 :

$$\begin{aligned} \text{Eigenvalues } & \sigma_1^2, \sigma_2^2 \\ \text{Eigenvectors } & \mathbf{v}_1, \mathbf{v}_2 \end{aligned} \quad A^T A = V \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} V^T. \quad (8)$$

This is exactly like $A = Q\Lambda Q^T$. But the symmetric matrix is not A itself. Now the symmetric matrix is $A^T A$! *And the columns of V are the eigenvectors of $A^T A$.* Last is U :

Compute the eigenvectors v and eigenvalues σ^2 of $A^T A$. Then each $u = Av/\sigma$.

For large matrices LAPACK finds a special way to avoid multiplying $A^T A$ in $\text{svd}(A)$.




Figure 6.8: U and V are rotations and reflections. Σ stretches circle to ellipse.

Example 3 Find the singular value decomposition of that matrix $A = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix}$.

Solution Compute $A^T A$ and its eigenvectors. Then make them unit vectors:

$$A^T A = \begin{bmatrix} 5 & 3 \\ 3 & 5 \end{bmatrix} \text{ has unit eigenvectors } v_1 = \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} \text{ and } v_2 = \begin{bmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}.$$

The eigenvalues of $A^T A$ are 8 and 2. The v 's are perpendicular, because eigenvectors of every symmetric matrix are perpendicular—and $A^T A$ is automatically symmetric.

Now the u 's are quick to find, because Av_1 is going to be in the direction of u_1 :

$$Av_1 = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} 2\sqrt{2} \\ 0 \end{bmatrix}. \text{ The unit vector is } u_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

Clearly Av_1 is the same as $2\sqrt{2}u_1$. The first singular value is $\sigma_1 = 2\sqrt{2}$. Then $\sigma_1^2 = 8$.

$$Av_2 = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} 0 \\ \sqrt{2} \end{bmatrix}. \text{ The unit vector is } u_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Now Av_2 is $\sqrt{2}u_2$ and $\sigma_2 = \sqrt{2}$. Thus σ_2^2 agrees with the other eigenvalue 2 of $A^T A$.

$$A = U\Sigma V^T \text{ is } \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 2\sqrt{2} & \\ & \sqrt{2} \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}. \quad (9)$$

This matrix, and every invertible 2 by 2 matrix, *transforms the unit circle to an ellipse*. You can see that in the figure, which was created by Cliff Long and Tom Hern.

One final point about that example. We found the u 's from the v 's. Could we find the u 's directly? *Yes*, by multiplying AA^T instead of $A^T A$:

$$\text{Use } V^T V = I \quad AA^T = (U \Sigma V^T)(V \Sigma^T U^T) = U \Sigma \Sigma^T U^T. \quad (10)$$

Multiplying $\Sigma \Sigma^T$ gives σ_1^2 and σ_2^2 as before. *The u 's are eigenvectors of AA^T :*

Diagonal in this example $AA^T = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 2 & -1 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} 8 & 0 \\ 0 & 2 \end{bmatrix}.$

The eigenvectors $(1, 0)$ and $(0, 1)$ agree with u_1 and u_2 found earlier. Why take the first eigenvector to be $(1, 0)$ instead of $(-1, 0)$ or $(0, 1)$? Because we have to follow Av_1 (I missed that in my video lecture ...). Notice that AA^T has the same eigenvalues (8 and 2) as $A^T A$. The singular values are $\sqrt{8}$ and $\sqrt{2}$.

Example 4 Find the SVD of the singular matrix $A = \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix}$. The rank is $r = 1$.

Solution The row space has only one basis vector $v_1 = (1, 1)/\sqrt{2}$. The column space has only one basis vector $u_1 = (2, 1)/\sqrt{5}$. Then $Av_1 = (4, 2)/\sqrt{2}$ must equal $\sigma_1 u_1$. It does, with $\sigma_1 = \sqrt{10}$.




Figure 6.9: The SVD chooses orthonormal bases for 4 subspaces so that $Av_i = \sigma_i u_i$.

The SVD could stop after the row space and column space (it usually doesn't). It is customary for U and V to be square. The matrices need a second column. *The vector v_2 is in the nullspace.* It is perpendicular to v_1 in the row space. Multiply by A to get $Av_2 = \mathbf{0}$. We could say that the second singular value is $\sigma_2 = 0$, but singular values are like pivots—only the r nonzeros are counted.

$$\text{Full size } A = U \Sigma V^T \quad \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix} = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & 1 \\ 1 & -2 \end{bmatrix} \begin{bmatrix} \sqrt{10} & 0 \\ 0 & 0 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \quad (11)$$

The matrices U and V contain orthonormal bases for all four subspaces:

first	r	columns of V :	row space of A
last	$n - r$	columns of V :	nullspace of A
first	r	columns of U :	column space of A
last	$m - r$	columns of U :	nullspace of A^T

The first columns v_1, \dots, v_r and u_1, \dots, u_r are eigenvectors of $A^T A$ and $A A^T$. We now explain why $A v_i$ falls in the direction of u_i . The last v 's and u 's (in the nullspaces) are easier. As long as those are orthonormal, the SVD will be correct.

Proof of the SVD: Start from $A^T A v_i = \sigma_i^2 v_i$, which gives the v 's and σ 's. Multiplying by v_i^T leads to $\|A v_i\|^2$. To prove that $A v_i = \sigma_i u_i$, the key step is to multiply by A :

$$v_i^T A^T A v_i = \sigma_i^2 v_i^T v_i \quad \text{gives} \quad \|A v_i\|^2 = \sigma_i^2 \quad \text{so that} \quad \|A v_i\| = \sigma_i \quad (12)$$

$$A A^T A v_i = \sigma_i^2 A v_i \quad \text{gives} \quad u_i = A v_i / \sigma_i \quad \text{as a unit eigenvector of } A A^T. \quad (13)$$

Equation (12) used the small trick of placing parentheses in $(v_i^T A^T)(A v_i) = \|A v_i\|^2$. Equation (13) placed the all-important parentheses in $(A A^T)(A v_i)$. This shows that $A v_i$ is an eigenvector of $A A^T$. Divide by its length σ_i to get the unit vector $u_i = A v_i / \sigma_i$. These u 's are orthogonal because $(A v_i)^T (A v_j) = v_i^T (A^T A v_j) = v_i^T (\sigma_j^2 v_j) = 0$.

I will give my opinion directly. The SVD is the climax of this linear algebra course. I think of it as the final step in the Fundamental Theorem. First come the *dimensions* of the four subspaces. Then their *orthogonality*. Then *the orthonormal bases diagonalize A* . It is all in the formula $A = U \Sigma V^T$. You have made it to the top.

Eigshow (Part 2)

Section 6.1 described the MATLAB demo called **eigshow**. The first option is *eig*, when x moves in a circle and Ax follows on an ellipse. The second option is *svd*, when two vectors x and y stay perpendicular as they travel around a circle. Then Ax and Ay move too (not usually perpendicular). The four vectors are in the Java demo on web.mit.edu/18.06.

The SVD is seen graphically when Ax is perpendicular to Ay . Their directions at that moment give an orthonormal basis u_1, u_2 . Their lengths give the singular values σ_1, σ_2 . The vectors x and y at that same moment are the orthonormal basis v_1, v_2 .

Searching the Web

I will end with an application of the SVD to web search engines. When you google a word, you get a list of web sites in order of importance. You could try "four subspaces".

The HITS algorithm that we describe is one way to produce that ranked list. It begins with about 200 sites found from an index of key words, and after that we look only at links between pages. Search engines are link-based more than content-based.

Start with the 200 sites and all sites that link to them and all sites they link to. That is our list, to be put in order. Importance can be measured by links out and links in.

1. The site is an **authority**: *links come in from many sites*. Especially from hubs.
2. The site is a **hub**: *links go out to many sites in the list*. Especially to authorities.

We want numbers x_1, \dots, x_N to rank the authorities and y_1, \dots, y_N to rank the hubs. Start with a simple count: x_i^0 and y_i^0 count the links into and out of site i .

Here is the point: *A good authority has links from important sites* (like hubs). Links from universities count more heavily than links from friends. *A good hub is linked to important sites* (like authorities). A link to **amazon.com** unfortunately means more than a link to **wellesleycambridge.com**. The rankings x^0 and y^0 from counting links are updated to x^1 and y^1 by taking account of *good* links (measuring their quality by x^0 and y^0):

$$\text{Authority values } x_i^1 = \sum_{j \text{ links to } i} y_j^0 \qquad \text{Hub values } y_i^1 = \sum_{i \text{ links to } j} x_j^0 \quad (14)$$

In matrix language those are $x^1 = A^T y^0$ and $y^1 = A x^0$. The matrix A contains 1's and 0's, with $a_{ij} = 1$ when i links to j . In the language of graphs, A is an "adjacency matrix" for the World Wide Web (an enormous matrix). The new x^1 and y^1 give better rankings, but not the best. Take another step like (14), to reach x^2 and y^2 :

$$A^T A \text{ and } A A^T \text{ appear } x^2 = A^T y^1 = A^T A x^0 \quad \text{and} \quad y^2 = A^T x^1 = A A^T y^0. \quad (15)$$

In two steps we are multiplying by $A^T A$ and $A A^T$. Twenty steps will multiply by $(A^T A)^{10}$ and $(A A^T)^{10}$. When we take powers, the largest eigenvalue σ_1^2 begins to dominate. And the vectors x and y line up with the leading eigenvectors v_1 and u_1 of $A^T A$ and $A A^T$. We are computing the top terms in the SVD, by the *power method* that is discussed in Section 9.3. It is wonderful that linear algebra helps to understand the Web.

Google actually creates rankings by a random walk that follows web links. The more often this random walk goes to a site, the higher the ranking. The frequency of visits gives the leading eigenvector ($\lambda = 1$) of the normalized adjacency matrix for the Web. ***That Markov matrix has 2.7 billion rows and columns, from 2.7 billion web sites.***

This is the largest eigenvalue problem ever solved. The excellent book by Langville and Meyer, *Google's PageRank and Beyond*, explains in detail the science of search engines. See mathworks.com/company/newsletter/clevescorner/oct02_cleve.shtml

But many of the important techniques are well-kept secrets of Google. Probably Google starts with last month's eigenvector as a first approximation, and runs the random walk very fast. To get a high ranking, you want a lot of links from important sites. The HITS algorithm is described in the 1999 *Scientific American* (June 16). But I don't think the SVD is mentioned there. . .

■ REVIEW OF THE KEY IDEAS ■

1. The SVD factors A into $U \Sigma V^T$, with r singular values $\sigma_1 \geq \dots \geq \sigma_r > 0$.

2. The numbers $\sigma_1^2, \dots, \sigma_r^2$ are the nonzero eigenvalues of AA^T and $A^T A$.
3. The orthonormal columns of U and V are eigenvectors of AA^T and $A^T A$.
4. Those columns hold orthonormal bases for the four fundamental subspaces of A .
5. Those bases diagonalize the matrix: $Av_i = \sigma_i u_i$ for $i \leq r$. This is $AV = U\Sigma$.

■ WORKED EXAMPLES ■

6.7 A Identify by name these decompositions $A = c_1 b_1 + \dots + c_r b_r$ of an m by n matrix. Each term is a rank one matrix (column c times row b). The rank of A is r .

1. *Orthogonal* columns c_1, \dots, c_r and *orthogonal* rows b_1, \dots, b_r .
2. *Orthogonal* columns c_1, \dots, c_r and *triangular* rows b_1, \dots, b_r .
3. *Triangular* columns c_1, \dots, c_r and *triangular* rows b_1, \dots, b_r .

$A = CB$ is $(m \text{ by } r)(r \text{ by } n)$. Triangular vectors c_i and b_i have zeros up to component i . The matrix C with columns c_i is lower triangular, the matrix B with rows b_i is upper triangular. Where do the rank and the pivots and singular values come into this picture?

Solution These three splittings $A = CB$ are basic to linear algebra, pure or applied:

1. **Singular Value Decomposition** $A = U\Sigma V^T$ (*orthogonal* U , *orthogonal* ΣV^T)
2. **Gram-Schmidt Orthogonalization** $A = QR$ (*orthogonal* Q , *triangular* R)
3. **Gaussian Elimination** $A = LU$ (*triangular* L , *triangular* U)

You might prefer to separate out the σ_i and pivots d_i and heights h_i :

1. $A = U\Sigma V^T$ with unit vectors in U and V . The singular values are in Σ .
2. $A = QHR$ with unit vectors in Q and diagonal 1's in R . The heights h_i are in H .
3. $A = LDU$ with diagonal 1's in L and U . The pivots are in D .

Each h_i tells the height of column i above the base from earlier columns. The volume of the full n -dimensional box ($r = m = n$) comes from $A = U\Sigma V^T = LDU = QHR$:

$$|\det A| = |\text{product of } \sigma \text{'s}| = |\text{product of } d \text{'s}| = |\text{product of } h \text{'s}|.$$

6.7.B For $A = xy^T$ of rank one (2 by 2), compare $A = U\Sigma V^T$ with $A = SAS^{-1}$.

Comment This started as an exam problem in 2007. It led further and became interesting. Now there is an essay called "The Four Fundamental Subspaces: 4 Lines" on web.mit.edu/18.06. The Jordan form enters when $y^T x = 0$ and $\lambda = 0$ is repeated.

6.7.C Show that $\sigma_1 \geq |\lambda|_{\max}$. The largest singular value dominates all eigenvalues.
 Show that $\sigma_1 \geq |a_{ij}|_{\max}$. The largest singular value dominates all entries of A .

Solution Start from $A = U\Sigma V^T$. Remember that multiplying by an orthogonal matrix *does not change length*: $\|Qx\| = \|x\|$ because $\|Qx\|^2 = x^T Q^T Q x = x^T x = \|x\|^2$. This applies to $Q = U$ and $Q = V^T$. In between is the diagonal matrix Σ .

$$\|Ax\| = \|U\Sigma V^T x\| = \|\Sigma V^T x\| \leq \sigma_1 \|V^T x\| = \sigma_1 \|x\|. \quad (16)$$

An eigenvector has $\|Ax\| = |\lambda|\|x\|$. So (16) says that $|\lambda|\|x\| \leq \sigma_1 \|x\|$. Then $|\lambda| \leq \sigma_1$.

Apply also to the unit vector $x = (1, 0, \dots, 0)$. Now Ax is the first column of A . Then by inequality (16), this column has length $\leq \sigma_1$. Every entry must have magnitude $\leq \sigma_1$.

Example 5 Estimate the singular values σ_1 and σ_2 of A and A^{-1} :

$$\text{Eigenvalues} = 1 \quad A = \begin{bmatrix} 1 & 0 \\ C & 1 \end{bmatrix} \quad \text{and} \quad A^{-1} = \begin{bmatrix} 1 & 0 \\ -C & 1 \end{bmatrix}. \quad (17)$$

Solution The length of the first column is $\sqrt{1+C^2} \leq \sigma_1$, from the reasoning above. This confirms that $\sigma_1 \geq 1$ and $\sigma_1 \geq C$. Then σ_1 dominates the eigenvalues 1, 1 and the entry C . If C is very large then σ_1 is much bigger than the eigenvalues.

This matrix A has determinant = 1. $A^T A$ also has determinant = 1 and then $\sigma_1 \sigma_2 = 1$. For this matrix, $\sigma_1 \geq 1$ and $\sigma_1 \geq C$ lead to $\sigma_2 \leq 1$ and $\sigma_2 \leq 1/C$.

Conclusion: If $C = 1000$ then $\sigma_1 \geq 1000$ and $\sigma_2 \leq 1/1000$. A is *ill-conditioned*, slightly sick. Inverting A is easy by algebra, but solving $Ax = b$ by elimination could be dangerous. A is close to a singular matrix even though both eigenvalues are $\lambda = 1$. By slightly changing the 1, 2 entry from zero to $1/C = 1/1000$, the matrix becomes singular.

Section 9.2 will explain how the ratio $\sigma_{\max}/\sigma_{\min}$ governs the roundoff error in elimination. MATLAB warns you if this “*condition number*” is large. Here $\sigma_1/\sigma_2 \geq C^2$.

Problem Set 6.7

Problems 1–3 compute the SVD of a square singular matrix A .

1 Find the eigenvalues and unit eigenvectors v_1, v_2 of $A^T A$. Then find $u_1 = Av_1/\sigma_1$:

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 6 \end{bmatrix} \quad \text{and} \quad A^T A = \begin{bmatrix} 10 & 20 \\ 20 & 40 \end{bmatrix} \quad \text{and} \quad AA^T = \begin{bmatrix} 5 & 15 \\ 15 & 45 \end{bmatrix}.$$

Verify that u_1 is a unit eigenvector of AA^T . Complete the matrices U, Σ, V .

$$\text{SVD} \quad \begin{bmatrix} 1 & 2 \\ 3 & 6 \end{bmatrix} = \begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_1 & v_2 \end{bmatrix}^T.$$

- 2 Write down orthonormal bases for the four fundamental subspaces of this A .
- 3 (a) Why is the trace of $A^T A$ equal to the sum of all a_{ij}^2 ?
 (b) For every rank-one matrix, why is $\sigma_1^2 = \text{sum of all } a_{ij}^2$?

Problems 4–7 ask for the SVD of matrices of rank 2.

- 4 Find the eigenvalues and unit eigenvectors of $A^T A$ and AA^T . Keep each $Av = \sigma u$:

$$\text{Fibonacci matrix} \quad A = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$$

Construct the singular value decomposition and verify that A equals $U\Sigma V^T$.

- 5 Use the `svd` part of the MATLAB demo `eigshow` to find those v 's graphically.
- 6 Compute $A^T A$ and AA^T and their eigenvalues and unit eigenvectors for V and U .

$$\text{Rectangular matrix} \quad A = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}.$$

Check $AV = U\Sigma$ (this will decide \pm signs in U). Σ has the same shape as A .

- 7 What is the closest rank-one approximation to that 2 by 3 matrix?
- 8 A square invertible matrix has $A^{-1} = V\Sigma^{-1}U^T$. This says that the *singular values of A^{-1} are $1/\sigma(A)$* . Show that $\sigma_{\max}(A^{-1}) \sigma_{\max}(A) \geq 1$.
- 9 Suppose u_1, \dots, u_n and v_1, \dots, v_n are orthonormal bases for \mathbf{R}^n . Construct the matrix A that transforms each v_j into u_j to give $Av_1 = u_1, \dots, Av_n = u_n$.
- 10 Construct the matrix with rank one that has $Av = 12u$ for $v = \frac{1}{2}(1, 1, 1, 1)$ and $u = \frac{1}{3}(2, 2, 1)$. Its only singular value is $\sigma_1 = \underline{\hspace{2cm}}$.
- 11 Suppose A has orthogonal columns w_1, w_2, \dots, w_n of lengths $\sigma_1, \sigma_2, \dots, \sigma_n$. What are U , Σ , and V in the SVD?
- 12 Suppose A is a 2 by 2 symmetric matrix with unit eigenvectors u_1 and u_2 . If its eigenvalues are $\lambda_1 = 3$ and $\lambda_2 = -2$, what are the matrices U , Σ , V^T in its SVD?
- 13 If $A = QR$ with an orthogonal matrix Q , the SVD of A is almost the same as the SVD of R . Which of the three matrices U , Σ , V is changed because of Q ?
- 14 Suppose A is invertible (with $\sigma_1 > \sigma_2 > 0$). Change A by *as small a matrix as possible* to produce a singular matrix A_0 . Hint: U and V do not change:

$$\text{From} \quad A = \begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} \sigma_1 & \\ & \sigma_2 \end{bmatrix} \begin{bmatrix} v_1 & v_2 \end{bmatrix}^T \quad \text{find the nearest } A_0.$$

- 15 Why doesn't the SVD for $A + I$ just use $\Sigma + I$?

Challenge Problems

- 16 (Search engine) Run a **random walk** $x(2), \dots, x(n)$ starting from web site $x(1) = 1$. Count the visits to each site. At each step the code chooses the next website $x(k)$ with probabilities given by column $x(k-1)$ of A . At the end, p gives the fraction of time at each site from a histogram: count visits. **The rankings are based on p .**

Please compare p to the steady state eigenvector of the Markov matrix A :

$$A = [0 \ .1 \ .2 \ .7; \ .05 \ 0 \ .15 \ .8; \ .15 \ .25 \ 0 \ .6; \ .1 \ .3 \ .6 \ 0]'$$

$n = 100$; $x = \text{zeros}(1, n)$; $x(1) = 1$;

for $k = 2 : n$ $x(k) = \min(\text{find}(\text{rand} < \text{cumsum}(A(:, x(k-1)))))$; end

$p = \text{hist}(x, 1 : 4) / n$

- 17 The $1, -1$ first difference matrix A has $A^T A =$ second difference matrix. The singular vectors of A are *sine vectors* v and *cosine vectors* u . Then $Av = \sigma u$ is the discrete form of $d/dx(\sin cx) = c(\cos cx)$. This is the best SVD I have seen.

$$\text{SVD of } A \quad A = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & -1 \end{bmatrix} \quad A^T A = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix}$$

$$\text{Orthogonal sine matrix} \quad V = \frac{1}{\sqrt{2}} \begin{bmatrix} \sin \pi/4 & \sin 2\pi/4 & \sin 3\pi/4 \\ \sin 2\pi/4 & \sin 4\pi/4 & \sin 6\pi/4 \\ \sin 3\pi/4 & \sin 6\pi/4 & \sin 9\pi/4 \end{bmatrix}$$

- (a) Put numbers in V : The unit eigenvectors of $A^T A$ are singular vectors of A . Show that the columns of V have $A^T A v = \lambda v$ with $\lambda = 2 - \sqrt{2}, 2, 2 + \sqrt{2}$.
- (b) Multiply AV and verify that its columns are orthogonal. They are $\sigma_1 u_1$ and $\sigma_2 u_2$ and $\sigma_3 u_3$. The first columns of the cosine matrix U are u_1, u_2, u_3 .
- (c) Since A is 4 by 3, we need a fourth orthogonal vector u_4 . It comes from the nullspace of A^T . What is u_4 ?

The cosine vectors in U are eigenvectors of AA^T . The fourth cosine is $(1, 1, 1, 1)/2$.

$$AA^T = \begin{bmatrix} 1 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 1 \end{bmatrix} \quad U = \frac{1}{\sqrt{2}} \begin{bmatrix} \cos \pi/8 & \cos 2\pi/8 & \cos 3\pi/8 \\ \cos 3\pi/8 & \cos 6\pi/8 & \cos 9\pi/8 \\ \cos 5\pi/8 & \cos 10\pi/8 & \cos 15\pi/8 \\ \cos 7\pi/8 & \cos 14\pi/8 & \cos 21\pi/8 \end{bmatrix}$$

Those angles $\pi/8, 3\pi/8, 5\pi/8, 7\pi/8$ fit 4 points with spacing $\pi/4$ between 0 and π . The sine transform has three points $\pi/4, 2\pi/4, 3\pi/4$. The full cosine transform includes u_4 from the "zero frequency" or *direct current* eigenvector $(1, 1, 1, 1)$.

The 8 by 8 cosine transform in 2D is the workhorse of **jpeg compression**. Linear algebra (circulant, Toeplitz, orthogonal matrices) is at the heart of signal processing.

Table of Eigenvalues and Eigenvectors

How are the properties of a matrix reflected in its eigenvalues and eigenvectors? This question is fundamental throughout Chapter 6. A table that organizes the key facts may be helpful. Here are the special properties of the eigenvalues λ_i and the eigenvectors x_i .

Symmetric: $A^T = A$	real λ 's	orthogonal $x_i^T x_j = 0$
Orthogonal: $Q^T = Q^{-1}$	all $ \lambda = 1$	orthogonal $\bar{x}_i^T x_j = 0$
Skew-symmetric: $A^T = -A$	imaginary λ 's	orthogonal $\bar{x}_i^T x_j = 0$
Complex Hermitian: $\bar{A}^T = A$	real λ 's	orthogonal $\bar{x}_i^T x_j = 0$
Positive Definite: $x^T A x > 0$	all $\lambda > 0$	orthogonal since $A^T = A$
Markov: $m_{ij} > 0, \sum_{i=1}^n m_{ij} = 1$	$\lambda_{\max} = 1$	steady state $x > 0$
Similar: $B = M^{-1} A M$	$\lambda(B) = \lambda(A)$	$x(B) = M^{-1} x(A)$
Projection: $P = P^2 = P^T$	$\lambda = 1; 0$	column space; nullspace
Plane Rotation	$e^{i\theta}$ and $e^{-i\theta}$	$x = (1, i)$ and $(1, -i)$
Reflection: $I - 2uu^T$	$\lambda = -1; 1, \dots, 1$	u ; whole plane u^\perp
Rank One: uv^T	$\lambda = v^T u; 0, \dots, 0$	u ; whole plane v^\perp
Inverse: A^{-1}	$1/\lambda(A)$	keep eigenvectors of A
Shift: $A + cI$	$\lambda(A) + c$	keep eigenvectors of A
Stable Powers: $A^n \rightarrow 0$	all $ \lambda < 1$	any eigenvectors
Stable Exponential: $e^{At} \rightarrow 0$	all $\text{Re } \lambda < 0$	any eigenvectors
Cyclic Permutation: row 1 of I last	$\lambda_k = e^{2\pi i k/n}$	$x_k = (1, \lambda_k, \dots, \lambda_k^{n-1})$
Tridiagonal: $-1, 2, -1$ on diagonals	$\lambda_k = 2 - 2 \cos \frac{k\pi}{n+1}$	$x_k = \left(\sin \frac{k\pi}{n+1}, \sin \frac{2k\pi}{n+1}, \dots \right)$
Diagonalizable: $A = S \Lambda S^{-1}$	diagonal of Λ	columns of S are independent
Symmetric: $A = Q \Lambda Q^T$	diagonal of Λ (real)	columns of Q are orthonormal
Schur: $A = Q T Q^{-1}$	diagonal of T	columns of Q if $A^T A = A A^T$
Jordan: $J = M^{-1} A M$	diagonal of J	each block gives $x = (0, \dots, 1, \dots, 0)$
Rectangular: $A = U \Sigma V^T$	$\text{rank}(A) = \text{rank}(\Sigma)$	eigenvectors of $A^T A, A A^T$ in V, U

Chapter 7

Linear Transformations

7.1 The Idea of a Linear Transformation

When a matrix A multiplies a vector v , it “transforms” v into another vector Av . *In goes v , out comes $T(v) = Av$.* A transformation T follows the same idea as a function. In goes a number x , out comes $f(x)$. For one vector v or one number x , we multiply by the matrix or we evaluate the function. The deeper goal is to see all v 's at once. We are transforming the whole space V when we multiply every v by A .

Start again with a matrix A . It transforms v to Av . It transforms w to Aw . Then we *know* what happens to $u = v + w$. There is no doubt about Au , it has to equal $Av + Aw$. Matrix multiplication $T(v) = Av$ gives a *linear transformation*:

A *transformation* T assigns an output $T(v)$ to each input vector v in V . The transformation is *linear* if it meets these requirements for all v and w :

$$(a) \quad T(v + w) = T(v) + T(w) \qquad (b) \quad T(cv) = cT(v) \quad \text{for all } c.$$

If the input is $v = \mathbf{0}$, the output must be $T(v) = \mathbf{0}$. We combine (a) and (b) into one:

$$\text{Linear transformation} \quad T(cv + dw) \quad \text{must equal} \quad cT(v) + dT(w).$$

Again I can test matrix multiplication for linearity: $A(cv + dw) = cAv + dAw$ is *true*.

A linear transformation is highly restricted. Suppose T adds u_0 to every vector. Then $T(v) = v + u_0$ and $T(w) = w + u_0$. This isn't good, or at least *it isn't linear*. Applying T to $v + w$ produces $v + w + u_0$. That is not the same as $T(v) + T(w)$:

$$\text{Shift is not linear} \quad v + w + u_0 \quad \text{is not} \quad T(v) + T(w) = v + u_0 + w + u_0.$$

The exception is when $u_0 = \mathbf{0}$. The transformation reduces to $T(v) = v$. This is the *identity transformation* (nothing moves, as in multiplication by the identity matrix). That is certainly linear. In this case the input space V is the same as the output space W .

The linear-plus-shift transformation $T(\mathbf{v}) = A\mathbf{v} + \mathbf{u}_0$ is called “*affine*”. Straight lines stay straight although T is not linear. Computer graphics works with affine transformations in Section 8.6, because we must be able to move images.

Example 1 Choose a fixed vector $\mathbf{a} = (1, 3, 4)$, and let $T(\mathbf{v})$ be the dot product $\mathbf{a} \cdot \mathbf{v}$:

The input is $\mathbf{v} = (v_1, v_2, v_3)$. The output is $T(\mathbf{v}) = \mathbf{a} \cdot \mathbf{v} = v_1 + 3v_2 + 4v_3$.

This is linear. The inputs \mathbf{v} come from three-dimensional space, so $\mathbf{V} = \mathbf{R}^3$. The outputs are just numbers, so the output space is $\mathbf{W} = \mathbf{R}^1$. We are multiplying by the row matrix $A = [1 \ 3 \ 4]$. Then $T(\mathbf{v}) = A\mathbf{v}$.

You will get good at recognizing which transformations are linear. If the output involves squares or products or lengths, v_1^2 or $v_1 v_2$ or $\|\mathbf{v}\|$, then T is not linear.

Example 2 The length $T(\mathbf{v}) = \|\mathbf{v}\|$ is not linear. Requirement (a) for linearity would be $\|\mathbf{v} + \mathbf{w}\| = \|\mathbf{v}\| + \|\mathbf{w}\|$. Requirement (b) would be $\|c\mathbf{v}\| = c\|\mathbf{v}\|$. Both are false!

Not (a): The sides of a triangle satisfy an *inequality* $\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|$.

Not (b): The length $\|-\mathbf{v}\|$ is not $-\|\mathbf{v}\|$. For negative c , we fail.

Example 3 (Important) T is the transformation that *rotates every vector by 30°* . The “*domain*” is the xy plane (all input vectors \mathbf{v}). The “*range*” is also the xy plane (all rotated vectors $T(\mathbf{v})$). We described T without a matrix: rotate by 30° .

Is rotation linear? *Yes it is.* We can rotate two vectors and add the results. The sum of rotations $T(\mathbf{v}) + T(\mathbf{w})$ is the same as the rotation $T(\mathbf{v} + \mathbf{w})$ of the sum. The whole plane is turning together, in this linear transformation.

Lines to Lines, Triangles to Triangles

Figure 7.1 shows the line from \mathbf{v} to \mathbf{w} in the input space. It also shows the line from $T(\mathbf{v})$ to $T(\mathbf{w})$ in the output space. Linearity tells us: Every point on the input line goes onto the output line. And more than that: *Equally spaced points go to equally spaced points.* The middle point $\mathbf{u} = \frac{1}{2}\mathbf{v} + \frac{1}{2}\mathbf{w}$ goes to the middle point $T(\mathbf{u}) = \frac{1}{2}T(\mathbf{v}) + \frac{1}{2}T(\mathbf{w})$.

The second figure moves up a dimension. Now we have three corners $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$. Those inputs have three outputs $T(\mathbf{v}_1), T(\mathbf{v}_2), T(\mathbf{v}_3)$. *The input triangle goes onto the output triangle.* Equally spaced points stay equally spaced (along the edges, and then between the edges). The middle point $\mathbf{u} = \frac{1}{3}(\mathbf{v}_1 + \mathbf{v}_2 + \mathbf{v}_3)$ goes to the middle point $T(\mathbf{u}) = \frac{1}{3}(T(\mathbf{v}_1) + T(\mathbf{v}_2) + T(\mathbf{v}_3))$.

The rule of linearity extends to combinations of three vectors or n vectors:

$$\begin{array}{l} \mathbf{u} = c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + \cdots + c_n\mathbf{v}_n \quad \text{transforms to} \\ \text{Linearity} \quad T(\mathbf{u}) = c_1T(\mathbf{v}_1) + c_2T(\mathbf{v}_2) + \cdots + c_nT(\mathbf{v}_n) \end{array} \quad (1)$$




Figure 7.1: Lines to lines, equal spacing to equal spacing, $u = \mathbf{0}$ to $T(u) = \mathbf{0}$.

Note Transformations have a language of their own. Where there is no matrix, we can't talk about a column space. But the idea can be rescued. The column space consisted of all outputs Av . The nullspace consisted of all inputs for which $Av = \mathbf{0}$. Translate those into "range" and "kernel":

Range of T = set of all outputs $T(v)$: range corresponds to column space

Kernel of T = set of all inputs for which $T(v) = \mathbf{0}$: kernel corresponds to nullspace.

The range is in the output space \mathbf{W} . The kernel is in the input space \mathbf{V} . When T is multiplication by a matrix, $T(v) = Av$, you can translate to column space and nullspace.

Examples of Transformations (mostly linear)

Example 4 Project every 3-dimensional vector straight down onto the xy plane. Then $T(x, y, z) = (x, y, 0)$. The range is that plane, which contains every $T(v)$. The kernel is the z axis (which projects down to zero). This projection is linear.

Example 5 Project every 3-dimensional vector onto the horizontal plane $z = 1$. The vector $v = (x, y, z)$ is transformed to $T(v) = (x, y, 1)$. This transformation is not linear. Why not? It doesn't even transform $v = \mathbf{0}$ into $T(v) = \mathbf{0}$.

Multiply every 3-dimensional vector by a 3 by 3 matrix A . This $T(v) = Av$ is linear.

$$T(v + w) = A(v + w) \quad \text{does equal} \quad Av + Aw = T(v) + T(w).$$

Example 6 Suppose A is an *invertible matrix*. The kernel of T is the zero vector; the range \mathbf{W} equals the domain \mathbf{V} . Another linear transformation is multiplication by A^{-1} . This is the *inverse transformation* T^{-1} , which brings every vector $T(v)$ back to v :

$$T^{-1}(T(v)) = v \quad \text{matches the matrix multiplication} \quad A^{-1}(Av) = v.$$

We are reaching an unavoidable question. *Are all linear transformations from $\mathbf{V} = \mathbf{R}^n$ to $\mathbf{W} = \mathbf{R}^m$ produced by matrices?* When a linear T is described as a "rotation" or "projection" or "...", is there always a matrix hiding behind T ?

The answer is *yes*. This is an approach to linear algebra that doesn't start with matrices. The next section shows that we still end up with matrices.

Linear Transformations of the Plane

It is more interesting to *see* a transformation than to define it. When a 2 by 2 matrix A multiplies all vectors in \mathbf{R}^2 , we can watch how it acts. Start with a “house” that has eleven endpoints. Those eleven vectors v are transformed into eleven vectors Av . Straight lines between v 's become straight lines between the transformed vectors Av . (The transformation from house to house is linear!) Applying A to a standard house produces a new house—possibly stretched or rotated or otherwise unlivable.

This part of the book is visual, not theoretical. We will show four houses and the matrices that produce them. The columns of H are the eleven corners of the first house. (H is 2 by 12, so `plot2d` will connect the 11th corner to the first.) The 11 points in the house matrix H are multiplied by A to produce the corners AH of the other houses.

$$\text{House matrix } H = \begin{bmatrix} -6 & -6 & -7 & 0 & 7 & 6 & 6 & -3 & -3 & 0 & 0 & -6 \\ -7 & 2 & 1 & 8 & 1 & 2 & -7 & -7 & -2 & -2 & -7 & -7 \end{bmatrix}.$$




Figure 7.2: Linear transformations of a house drawn by `plot2d(A * H)`.

■ REVIEW OF THE KEY IDEAS ■

1. A transformation T takes each v in the input space to $T(v)$ in the output space.
2. T is **linear** if $T(v + w) = T(v) + T(w)$ and $T(cv) = cT(v)$: lines to lines.


3. Combinations to combinations: $T(c_1v_1 + \cdots + c_nv_n) = c_1T(v_1) + \cdots + c_nT(v_n)$.
4. The transformation $T(v) = Av + v_0$ is linear only if $v_0 = \mathbf{0}$. Then $T(v) = Av$.

■ WORKED EXAMPLES ■

7.1 A The elimination matrix $\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$ gives a *shearing transformation* from (x, y) to $T(x, y) = (x, x + y)$. Draw the xy plane and show what happens to $(1, 0)$ and $(1, 1)$. What happens to points on the vertical lines $x = 0$ and $x = a$? If the inputs fill the unit square $0 \leq x \leq 1$, $0 \leq y \leq 1$, draw the outputs (the transformed square).

Solution The points $(1, 0)$ and $(2, 0)$ on the x axis transform by T to $(1, 1)$ and $(2, 2)$. The horizontal x axis transforms to the 45° line (going through $(0, 0)$ of course). The points on the y axis are *not moved* because $T(0, y) = (0, y)$. The y axis is the line of eigenvectors of T with $\lambda = 1$. Points with $x = a$ move up by a .

Vertical lines slide up
This is the shearing
Squares to parallelograms



7.1 B A **nonlinear transformation** T is invertible if every b in the output space comes from exactly one x in the input space: $T(x) = b$ always has exactly one solution. Which of these transformations (on real numbers x) is invertible and what is T^{-1} ? *None are linear, not even T_3 .* When you solve $T(x) = b$, you are inverting T :

$$T_1(x) = x^2 \quad T_2(x) = x^3 \quad T_3(x) = x + 9 \quad T_4(x) = e^x \quad T_5(x) = \frac{1}{x} \text{ for nonzero } x\text{'s}$$

Solution T_1 is not invertible: $x^2 = 1$ has *two* solutions and $x^2 = -1$ has *no* solution. T_4 is not invertible because $e^x = -1$ has no solution. (If the output space changes to *positive* b 's then the inverse of $e^x = b$ is $x = \ln b$.)

Notice $T_5^2 = \text{identity}$. But $T_3^2(x) = x + 18$. What are $T_2^2(x)$ and $T_4^2(x)$?

T_2, T_3, T_5 are invertible. The solutions to $x^3 = b$ and $x + 9 = b$ and $\frac{1}{x} = b$ are unique:

$$x = T_2^{-1}(b) = b^{1/3} \quad x = T_3^{-1}(b) = b - 9 \quad x = T_5^{-1}(b) = 1/b$$

Problem Set 7.1

- 1 A linear transformation must leave the zero vector fixed: $T(\mathbf{0}) = \mathbf{0}$. Prove this from $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w})$ by choosing $\mathbf{w} = \underline{\hspace{2cm}}$ (and finish the proof). Prove it also from $T(c\mathbf{v}) = cT(\mathbf{v})$ by choosing $c = \underline{\hspace{2cm}}$.
- 2 Requirement (b) gives $T(c\mathbf{v}) = cT(\mathbf{v})$ and also $T(d\mathbf{w}) = dT(\mathbf{w})$. Then by addition, requirement (a) gives $T(\underline{\hspace{1cm}}) = (\underline{\hspace{1cm}})$. What is $T(c\mathbf{v} + d\mathbf{w} + e\mathbf{u})$?
- 3 Which of these transformations are not linear? The input is $\mathbf{v} = (v_1, v_2)$:
- (a) $T(\mathbf{v}) = (v_2, v_1)$ (b) $T(\mathbf{v}) = (v_1, v_1)$ (c) $T(\mathbf{v}) = (0, v_1)$
 (d) $T(\mathbf{v}) = (0, 1)$ (e) $T(\mathbf{v}) = v_1 - v_2$ (f) $T(\mathbf{v}) = v_1 v_2$.
- 4 If S and T are linear transformations, is $S(T(\mathbf{v}))$ linear or quadratic?
- (a) (Special case) If $S(\mathbf{v}) = \mathbf{v}$ and $T(\mathbf{v}) = \mathbf{v}$, then $S(T(\mathbf{v})) = \mathbf{v}$ or \mathbf{v}^2 ?
 (b) (General case) $S(\mathbf{w}_1 + \mathbf{w}_2) = S(\mathbf{w}_1) + S(\mathbf{w}_2)$ and $T(\mathbf{v}_1 + \mathbf{v}_2) = T(\mathbf{v}_1) + T(\mathbf{v}_2)$
 combine into

$$S(T(\mathbf{v}_1 + \mathbf{v}_2)) = S(\underline{\hspace{2cm}}) = \underline{\hspace{2cm}} + \underline{\hspace{2cm}}.$$

- 5 Suppose $T(\mathbf{v}) = \mathbf{v}$ except that $T(0, v_2) = (0, 0)$. Show that this transformation satisfies $T(c\mathbf{v}) = cT(\mathbf{v})$ but not $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w})$.
- 6 Which of these transformations satisfy $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w})$ and which satisfy $T(c\mathbf{v}) = cT(\mathbf{v})$?
- (a) $T(\mathbf{v}) = \mathbf{v}/\|\mathbf{v}\|$ (b) $T(\mathbf{v}) = v_1 + v_2 + v_3$ (c) $T(\mathbf{v}) = (v_1, 2v_2, 3v_3)$
 (d) $T(\mathbf{v}) =$ largest component of \mathbf{v} .
- 7 For these transformations of $\mathbf{V} = \mathbf{R}^2$ to $\mathbf{W} = \mathbf{R}^2$, find $T(T(\mathbf{v}))$. Is this transformation T^2 linear?
- (a) $T(\mathbf{v}) = -\mathbf{v}$ (b) $T(\mathbf{v}) = \mathbf{v} + (1, 1)$
 (c) $T(\mathbf{v}) = 90^\circ$ rotation $= (-v_2, v_1)$
 (d) $T(\mathbf{v}) =$ projection $= (\frac{v_1 + v_2}{2}, \frac{v_1 + v_2}{2})$.
- 8 Find the range and kernel (like the column space and nullspace) of T :
- (a) $T(v_1, v_2) = (v_1 - v_2, 0)$ (b) $T(v_1, v_2, v_3) = (v_1, v_2)$
 (c) $T(v_1, v_2) = (0, 0)$ (d) $T(v_1, v_2) = (v_1, v_1)$.
- 9 The “cyclic” transformation T is defined by $T(v_1, v_2, v_3) = (v_2, v_3, v_1)$. What is $T(T(\mathbf{v}))$? What is $T^3(\mathbf{v})$? What is $T^{100}(\mathbf{v})$? Apply T a hundred times to \mathbf{v} .

- 10 A linear transformation from \mathbf{V} to \mathbf{W} has an *inverse* from \mathbf{W} to \mathbf{V} when the range is all of \mathbf{W} and the kernel contains only $v = \mathbf{0}$. Then $T(v) = w$ has one solution v for each w in \mathbf{W} . Why are these T 's not invertible?
- (a) $T(v_1, v_2) = (v_2, v_2)$ $\mathbf{W} = \mathbf{R}^2$
 (b) $T(v_1, v_2) = (v_1, v_2, v_1 + v_2)$ $\mathbf{W} = \mathbf{R}^3$
 (c) $T(v_1, v_2) = v_1$ $\mathbf{W} = \mathbf{R}^1$
- 11 If $T(v) = Av$ and A is m by n , then T is “multiplication by A .”
- (a) What are the input and output spaces \mathbf{V} and \mathbf{W} ?
 (b) Why is range of $T =$ column space of A ?
 (c) Why is kernel of $T =$ nullspace of A ?
- 12 Suppose a linear T transforms $(1, 1)$ to $(2, 2)$ and $(2, 0)$ to $(0, 0)$. Find $T(v)$:
- (a) $v = (2, 2)$ (b) $v = (3, 1)$ (c) $v = (-1, 1)$ (d) $v = (a, b)$.

Problems 13-19 may be harder. The input space \mathbf{V} contains all 2 by 2 matrices M .

- 13 M is any 2 by 2 matrix and $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$. The transformation T is defined by $T(M) = AM$. What rules of matrix multiplication show that T is linear?
- 14 Suppose $A = \begin{bmatrix} 1 & 2 \\ 3 & 5 \end{bmatrix}$. Show that the range of T is the whole matrix space \mathbf{V} and the kernel is the zero matrix:
- (1) If $AM = 0$ prove that M must be the zero matrix.
 (2) Find a solution to $AM = B$ for any 2 by 2 matrix B .
- 15 Suppose $A = \begin{bmatrix} 1 & 2 \\ 3 & 6 \end{bmatrix}$. Show that the identity matrix I is not in the range of T . Find a nonzero matrix M such that $T(M) = AM$ is zero.
- 16 Suppose T transposes every matrix M . Try to find a matrix A which gives $AM = M^T$ for every M . Show that no matrix A will do it. *To professors:* Is this a linear transformation that doesn't come from a matrix?
- 17 The transformation T that transposes every matrix is definitely linear. Which of these extra properties are true?
- (a) $T^2 =$ identity transformation.
 (b) The kernel of T is the zero matrix.
 (c) Every matrix is in the range of T .
 (d) $T(M) = -M$ is impossible.
- 18 Suppose $T(M) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} [M] \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$. Find a matrix with $T(M) \neq 0$. Describe all matrices with $T(M) = 0$ (the kernel) and all output matrices $T(M)$ (the range).
- 19 If A and B are invertible and $T(M) = AMB$, find $T^{-1}(M)$ in the form $(\quad)M(\quad)$.

Questions 20–26 are about house transformations. The output is $T(H) = AH$.

20 How can you tell from the picture of T (house) that A is

- (a) a diagonal matrix?
- (b) a rank-one matrix?
- (c) a lower triangular matrix?

21 Draw a picture of T (house) for these matrices:

$$D = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} .7 & .7 \\ .3 & .3 \end{bmatrix} \quad \text{and} \quad U = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

22 What are the conditions on $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ to ensure that T (house) will

- (a) sit straight up?
- (b) expand the house by 3 in all directions?
- (c) rotate the house with no change in its shape?

23 Describe T (house) when $T(v) = -v + (1, 0)$. This T is “affine”.

24 Change the house matrix H to add a chimney.

25 The standard house is drawn by `plot2d(H)`. Circles from `o` and lines from `-`:

```
x = H(1, :)'; y = H(2, :)';
axis([-10 10 -10 10]), axis('square')
plot(x, y, 'o', x, y, '-');
```

Test `plot2d(A' * H)` and `plot2d(A' * A * H)` with the matrices in Figure 7.1.

26 Without a computer sketch the houses $A * H$ for these matrices A :

$$\begin{bmatrix} 1 & 0 \\ 0 & .1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} .5 & .5 \\ .5 & .5 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} .5 & .5 \\ -.5 & .5 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}.$$

27 This code creates a vector `theta` of 50 angles. It draws the unit circle and then $T(\text{circle}) = \text{ellipse}$. $T(v) = Av$ takes circles to ellipses.

```
A = [2 1; 1 2] % You can change A
theta = [0:2 * pi/50:2 * pi];
circle = [cos(theta); sin(theta)];
ellipse = A * circle;
axis([-4 4 -4 4]); axis('square')
plot(circle(1,:), circle(2,:), ellipse(1,:), ellipse(2,:))
```

28 Add two eyes and a smile to the circle in Problem 27. (If one eye is dark and the other is light, you can tell when the face is reflected across the y axis.) Multiply by matrices A to get new faces.

Challenge Problems

- 29** What conditions on $\det A = ad - bc$ ensure that the output house AH will
- (a) be squashed onto a line?
 - (b) keep its endpoints in clockwise order (not reflected)?
 - (c) have the same area as the original house?
- 30** From $A = U\Sigma V^T$ (Singular Value Decomposition) A takes **circles to ellipses**. $AV = U\Sigma$ says that the radius vectors v_1 and v_2 of the circle go to the semi-axes $\sigma_1 u_1$ and $\sigma_2 u_2$ of the ellipse. Draw the circle and the ellipse for $\theta = 30^\circ$:

$$V = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad U = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad \Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}.$$

- 31** Why does every linear transformation T from \mathbf{R}^2 to \mathbf{R}^2 take squares to parallelograms? Rectangles also go to parallelograms (squashed if T is not invertible).

7.2 The Matrix of a Linear Transformation

The next pages assign a matrix to every linear transformation T . For ordinary column vectors, the input v is in $V = \mathbf{R}^n$ and the output $T(v)$ is in $W = \mathbf{R}^m$. The matrix A for this transformation T will be m by n . Our choice of bases in V and W will decide A .

The standard basis vectors for \mathbf{R}^n and \mathbf{R}^m are the columns of I . That choice leads to a standard matrix, and $T(v) = Av$ in the normal way. But these spaces also have other bases, so the same T is represented by other matrices. A main theme of linear algebra is to choose the bases that give the best matrix for T .

When V and W are not \mathbf{R}^n and \mathbf{R}^m , they still have bases. Each choice of basis leads to a matrix for T . When the input basis is different from the output basis, the matrix for $T(v) = v$ will not be the identity I . It will be the “change of basis matrix”.

Key idea of this section

*Suppose we know $T(v_1), \dots, T(v_n)$ for the basis vectors v_1, \dots, v_n .
Then linearity produces $T(v)$ for every other input vector v .*

Reason Every v is a unique combination $c_1v_1 + \dots + c_nv_n$ of the basis vectors v_i . Since T is a linear transformation (here is the moment for linearity), $T(v)$ must be the same combination $c_1T(v_1) + \dots + c_nT(v_n)$ of the known outputs $T(v_i)$.

Our first example gives the outputs $T(v)$ for the standard basis vectors $(1, 0)$ and $(0, 1)$.

Example 1 Suppose T transforms $v_1 = (1, 0)$ to $T(v_1) = (2, 3, 4)$. Suppose the second basis vector $v_2 = (0, 1)$ goes to $T(v_2) = (5, 5, 5)$. If T is linear from \mathbf{R}^2 to \mathbf{R}^3 then its “standard matrix” is 3 by 2. Those outputs $T(v_1)$ and $T(v_2)$ go into its columns:

$$A = \begin{bmatrix} 2 & 5 \\ 3 & 5 \\ 4 & 5 \end{bmatrix}. \quad \begin{array}{l} T(v_1 + v_2) = T(v_1) + T(v_2) \\ \text{combines the columns} \end{array} \quad \begin{bmatrix} 2 & 5 \\ 3 & 5 \\ 4 & 5 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 7 \\ 8 \\ 9 \end{bmatrix}.$$

Example 2 The derivatives of the functions $1, x, x^2, x^3$ are $0, 1, 2x, 3x^2$. Those are four facts about the transformation T that “takes the derivative”. The inputs and the outputs are functions! Now add the crucial fact that the “derivative transformation” T is linear:

$$T(v) = \frac{dv}{dx} \quad \text{obeys the linearity rule} \quad \frac{d}{dx}(cv + dw) = c \frac{dv}{dx} + d \frac{dw}{dx}. \quad (1)$$

It is exactly this linearity that you use to find all other derivatives. From the derivative of each separate power $1, x, x^2, x^3$ (those are the basis vectors v_1, v_2, v_3, v_4) you find the derivative of any polynomial like $4 + x + x^2 + x^3$:

$$\frac{d}{dx}(4 + x + x^2 + x^3) = 1 + 2x + 3x^2 \quad (\text{because of linearity!})$$

This example applies T (the derivative d/dx) to the input $v = 4v_1 + v_2 + v_3 + v_4$. Here the input space \mathbf{V} contains all combinations of $1, x, x^2, x^3$. I call them vectors, you might call them functions. Those four vectors are a basis for the space \mathbf{V} of cubic polynomials (degree ≤ 3). Four derivatives tell us all derivatives in \mathbf{V} .

For the nullspace of A , we solve $Av = \mathbf{0}$. For the kernel of the derivative T , we solve $dv/dx = \mathbf{0}$. The solution is $v = \text{constant}$. The nullspace of T is one-dimensional, containing all constant functions (like the first basis function $v_1 = 1$).

To find the range (or column space), look at all outputs from $T(v) = dv/dx$. The inputs are cubic polynomials $a + bx + cx^2 + dx^3$, so the outputs are *quadratic polynomials* (degree ≤ 2). For the output space \mathbf{W} we have a choice. If $\mathbf{W} = \text{cubics}$, then the range of T (the quadratics) is a subspace. If $\mathbf{W} = \text{quadratics}$, then the range is all of \mathbf{W} .

That second choice emphasizes the difference between the domain or input space ($\mathbf{V} = \text{cubics}$) and the image or output space ($\mathbf{W} = \text{quadratics}$). \mathbf{V} has dimension $n = 4$ and \mathbf{W} has dimension $m = 3$. The “derivative matrix” below will be 3 by 4.

The range of T is a three-dimensional subspace. The matrix will have rank $r = 3$. The kernel is one-dimensional. The sum $3 + 1 = 4$ is the dimension of the input space. This was $r + (n - r) = n$ in the Fundamental Theorem of Linear Algebra. Always *(dimension of range) + (dimension of kernel) = dimension of input space*.

Example 3 The *integral* is the inverse of the derivative. That is the Fundamental Theorem of Calculus. We see it now in linear algebra. The transformation T^{-1} that “takes the integral from 0 to x ” is linear! Apply T^{-1} to $1, x, x^2$, which are w_1, w_2, w_3 :

$$\text{Integration is } T^{-1} \quad \int_0^x 1 \, dx = x, \quad \int_0^x x \, dx = \frac{1}{2}x^2, \quad \int_0^x x^2 \, dx = \frac{1}{3}x^3.$$

By linearity, the integral of $w = B + Cx + Dx^2$ is $T^{-1}(w) = Bx + \frac{1}{2}Cx^2 + \frac{1}{3}Dx^3$. The integral of a quadratic is a cubic. The input space of T^{-1} is the quadratics, the output space is the cubics. *Integration takes \mathbf{W} back to \mathbf{V}* . Its matrix will be 4 by 3.

Range of T^{-1} The outputs $Bx + \frac{1}{2}Cx^2 + \frac{1}{3}Dx^3$ are cubics with no constant term.

Kernel of T^{-1} The output is zero only if $B = C = D = 0$. The nullspace is $\mathbf{Z} = \{\mathbf{0}\}$.

Fundamental Theorem $3 + 0$ is the dimension of the input space \mathbf{W} for T^{-1} .

Matrices for the Derivative and Integral

We will show how the matrices A and A^{-1} copy the derivative T and the integral T^{-1} . This is an excellent example from calculus. (I write A^{-1} but I don't quite mean it.) Then comes the general rule—how to represent any linear transformation T by a matrix A .

The derivative transforms the space \mathbf{V} of cubics to the space \mathbf{W} of quadratics. The basis for \mathbf{V} is $1, x, x^2, x^3$. The basis for \mathbf{W} is $1, x, x^2$. *The derivative matrix is 3 by 4:*

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix} = \text{matrix form of derivative } T. \quad (2)$$

Why is A the correct matrix? Because *multiplying by A agrees with transforming by T* . The derivative of $v = a + bx + cx^2 + dx^3$ is $T(v) = b + 2cx + 3dx^2$. The same numbers b and $2c$ and $3d$ appear when we multiply by the matrix A :

$$\text{Take the derivative} \quad \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = \begin{bmatrix} b \\ 2c \\ 3d \end{bmatrix}. \quad (3)$$

Look also at T^{-1} . The integration matrix is 4 by 3. Watch how the following matrix starts with $w = B + Cx + Dx^2$ and produces its integral $0 + Bx + \frac{1}{2}Cx^2 + \frac{1}{3}Dx^3$:

$$\text{Take the integral} \quad \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & \frac{1}{3} \end{bmatrix} \begin{bmatrix} B \\ C \\ D \end{bmatrix} = \begin{bmatrix} 0 \\ B \\ \frac{1}{2}C \\ \frac{1}{3}D \end{bmatrix}. \quad (4)$$

I want to call that matrix A^{-1} , and I will. But you realize that rectangular matrices don't have inverses. At least they don't have two-sided inverses. This rectangular A has a *one-sided inverse*. The integral is a one-sided inverse of the derivative!

$$AA^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{but} \quad A^{-1}A = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

If you integrate a function and then differentiate, you get back to the start. So $AA^{-1} = I$. But if you differentiate before integrating, the constant term is lost. *The integral of the derivative of 1 is zero:*

$$T^{-1}T(1) = \text{integral of zero function} = 0.$$

This matches $A^{-1}A$, whose first column is all zero. The derivative T has a kernel (the constant functions). Its matrix A has a nullspace. Main point again: Av copies $T(v)$.

Construction of the Matrix

Now we construct a matrix for any linear transformation. Suppose T transforms the space \mathbf{V} (n -dimensional) to the space \mathbf{W} (m -dimensional). We choose a basis v_1, \dots, v_n for \mathbf{V} and we choose a basis w_1, \dots, w_m for \mathbf{W} . The matrix A will be m by n . To find the first column of A , apply T to the first basis vector v_1 . The output $T(v_1)$ is in \mathbf{W} .

$T(v_1)$ is a combination $a_{11}w_1 + \dots + a_{m1}w_m$ of the output basis for \mathbf{W} .

These numbers a_{11}, \dots, a_{m1} go into the first column of A . Transforming v_1 to $T(v_1)$ matches multiplying $(1, 0, \dots, 0)$ by A . It yields that first column of the matrix.

When T is the derivative and the first basis vector is 1, its derivative is $T(v_1) = 0$. So for the derivative matrix, the first column of A was all zero.

For the integral, the first basis function is again 1. Its integral is the second basis function x . So the first column of A^{-1} was $(0, 1, 0, 0)$. Here is the construction of A .

Key rule: The j th column of A is found by applying T to the j th basis vector v_j

$$T(v_j) = \text{combination of basis vectors of } \mathbf{W} = a_{1j}w_1 + \cdots + a_{mj}w_m. \quad (5)$$

These numbers a_{1j}, \dots, a_{mj} go into column j of A . *The matrix is constructed to get the basis vectors right. Then linearity gets all other vectors right.* Every v is a combination $c_1v_1 + \cdots + c_nv_n$, and $T(v)$ is a combination of the w 's. When A multiplies the coefficient vector $c = (c_1, \dots, c_n)$ in the v combination, Ac produces the coefficients in the $T(v)$ combination. This is because matrix multiplication (combining columns) is linear like T .

The matrix A tells us what T does. Every linear transformation from \mathbf{V} to \mathbf{W} can be converted to a matrix. This matrix depends on the bases.

Example 4 *If the bases change, T is the same but the matrix A is different.*

Suppose we reorder the basis to $x, x^2, x^3, 1$ for the cubics in \mathbf{V} . Keep the original basis $1, x, x^2$ for the quadratics in \mathbf{W} . The derivative of the first basis vector $v_1 = x$ is the first basis vector $w_1 = 1$. So the first column of A looks different:

$$A_{\text{new}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \end{bmatrix} = \begin{array}{l} \text{matrix for the derivative } T \\ \text{when the bases change to} \\ x, x^2, x^3, 1 \text{ and } 1, x, x^2. \end{array}$$

When we reorder the basis of \mathbf{V} , we reorder the columns of A . The input basis vector v_j is responsible for column j . The output basis vector w_i is responsible for row i . Soon the changes in the bases will be more than permutations.

Products AB Match Transformations TS

The examples of derivative and integral made three points. First, linear transformations T are everywhere—in calculus and differential equations and linear algebra. Second, spaces other than \mathbf{R}^n are important—we had functions in \mathbf{V} and \mathbf{W} . Third, T still boils down to a matrix A . Now we make sure that we can find this matrix.

The next examples have $\mathbf{V} = \mathbf{W}$. We choose the same basis for both spaces. Then we can compare the matrices A^2 and AB with the transformations T^2 and TS .

Example 5 T rotates every vector by the angle θ . Here $\mathbf{V} = \mathbf{W} = \mathbf{R}^2$. Find A .

Solution The standard basis is $v_1 = (1, 0)$ and $v_2 = (0, 1)$. To find A , apply T to those basis vectors. In Figure 7.3a, they are rotated by θ . *The first vector $(1, 0)$ swings around to $(\cos \theta, \sin \theta)$.* This equals $\cos \theta$ times $(1, 0)$ plus $\sin \theta$ times $(0, 1)$. Therefore those

numbers $\cos \theta$ and $\sin \theta$ go into the *first column* of A :

$$\begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} \text{ shows column 1} \quad A = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \text{ shows both columns.}$$

For the second column, transform the second vector $(0, 1)$. The figure shows it rotated to $(-\sin \theta, \cos \theta)$. *Those numbers go into the second column.* Multiplying A times $(0, 1)$ produces that column. A agrees with T on the basis, and on all v .




Figure 7.3: Two transformations: Rotation by θ and projection onto the 45° line.

Example 6 (Projection) Suppose T projects every plane vector onto the 45° line. Find its matrix for two different choices of the basis. We will find two matrices.

Solution Start with a specially chosen basis, not drawn in Figure 7.3. The basis vector v_1 is along the 45° line. *It projects to itself: $T(v_1) = v_1$.* So the first column of A contains 1 and 0. The second basis vector v_2 is along the perpendicular line (135°). *This basis vector projects to zero.* So the second column of A contains 0 and 0:

Projection $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ when V and W have the 45° and 135° basis.

Now take the standard basis $(1, 0)$ and $(0, 1)$. Figure 7.3b shows how $(1, 0)$ projects to $(\frac{1}{2}, \frac{1}{2})$. That gives the first column of A . The other basis vector $(0, 1)$ also projects to $(\frac{1}{2}, \frac{1}{2})$. So the standard matrix for this projection is A :

Same projection $A = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$ for the standard basis.

Both A 's are projection matrices. If you square A it doesn't change. Projecting twice is the same as projecting once: $T^2 = T$ so $A^2 = A$. Notice what is hidden in that statement: *The matrix for T^2 is A^2 .*

We have come to something important—the real reason for the way matrices are multiplied. *At last we discover why!* Two transformations S and T are represented by two matrices B and A . When we apply T to the output from S , we get the “composition” TS . When we apply A after B , we get the matrix product AB . **Matrix multiplication gives the correct matrix AB to represent TS .**

The transformation S is from a space U to V . Its matrix B uses a basis u_1, \dots, u_p for U and a basis v_1, \dots, v_n for V . The matrix is n by p . The transformation T is from V to W as before. *Its matrix A must use the same basis v_1, \dots, v_n for V —this is the output space for S and the input space for T . Then the matrix AB matches TS :*

Multiplication The linear transformation TS starts with any vector u in U , goes to $S(u)$ in V and then to $T(S(u))$ in W . The matrix AB starts with any x in \mathbb{R}^p , goes to Bx in \mathbb{R}^n and then to ABx in \mathbb{R}^m . The matrix AB correctly represents TS :

$$TS : U \rightarrow V \rightarrow W \quad AB : (m \text{ by } n)(n \text{ by } p) = (m \text{ by } p).$$

The input is $u = x_1u_1 + \dots + x_pu_p$. The output $T(S(u))$ matches the output ABx . **Product of transformations matches product of matrices.**

The most important cases are when the spaces U, V, W are the same and their bases are the same. With $m = n = p$ we have square matrices.

Example 7 S rotates the plane by θ and T also rotates by θ . Then TS rotates by 2θ . This transformation T^2 corresponds to the rotation matrix A^2 through 2θ :

$$T = S \quad A = B \quad T^2 = \text{rotation by } 2\theta \quad A^2 = \begin{bmatrix} \cos 2\theta & -\sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{bmatrix}. \quad (6)$$

By matching (transformation)² with (matrix)², we pick up the formulas for $\cos 2\theta$ and $\sin 2\theta$. Multiply A times A :

$$\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} = \begin{bmatrix} \cos^2 \theta - \sin^2 \theta & -2 \sin \theta \cos \theta \\ 2 \sin \theta \cos \theta & \cos^2 \theta - \sin^2 \theta \end{bmatrix}. \quad (7)$$

Comparing (6) with (7) produces $\cos 2\theta = \cos^2 \theta - \sin^2 \theta$ and $\sin 2\theta = 2 \sin \theta \cos \theta$. Trigonometry (the double angle rule) comes from linear algebra.

Example 8 S rotates by θ and T rotates by $-\theta$. Then $TS = I$ matches $AB = I$.

In this case $T(S(u))$ is u . We rotate forward and back. For the matrices to match, ABx must be x . The two matrices are inverses. Check this by putting $\cos(-\theta) = \cos \theta$ and $\sin(-\theta) = -\sin \theta$ into the backward rotation matrix:

$$AB = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} = \begin{bmatrix} \cos^2 \theta + \sin^2 \theta & 0 \\ 0 & \cos^2 \theta + \sin^2 \theta \end{bmatrix} = I.$$

Earlier T took the derivative and S took the integral. The transformation TS is the identity but not ST . Therefore AB is the identity matrix but not BA :

$$AB = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & \frac{1}{3} \end{bmatrix} = I \quad \text{but} \quad BA = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

The Identity Transformation and the Change of Basis Matrix

We now find the matrix for the special and boring transformation $T(v) = v$. This *identity transformation* does nothing to v . The matrix for $T = I$ also does nothing, *provided* the output basis is the same as the input basis. The output $T(v_1)$ is v_1 . When the bases are the same, this is w_1 . So the first column of A is $(1, 0, \dots, 0)$.

When each output $T(v_j) = v_j$ is the same as w_j , the matrix is just I .

This seems reasonable: The identity transformation is represented by the identity matrix. But suppose the bases are *different*. Then $T(v_1) = v_1$ is a combination of the w 's. That combination $m_{11}w_1 + \dots + m_{n1}w_n$ tells the first column of the matrix (call it M).

Identity transformation

When the outputs $T(v_j) = v_j$ are combinations $\sum_{i=1}^n m_{ij}w_i$, the "change of basis matrix" is M .

The basis is changing but the vectors themselves are not changing: $T(v) = v$. When the inputs have one basis and the outputs have another basis, the matrix is not I .

Example 9 The input basis is $v_1 = (3, 7)$ and $v_2 = (2, 5)$. The output basis is $w_1 = (1, 0)$ and $w_2 = (0, 1)$. Then the matrix M is easy to compute:

Change of basis The matrix for $T(v) = v$ is $M = \begin{bmatrix} 3 & 2 \\ 7 & 5 \end{bmatrix}$.

Reason The first input is the basis vector $v_1 = (3, 7)$. The output is also $(3, 7)$ which we express as $3w_1 + 7w_2$. Then the first column of M contains 3 and 7.

This seems too simple to be important. It becomes trickier when the change of basis goes the other way. We get the inverse of the previous matrix M :

Example 10 The input basis is now $v_1 = (1, 0)$ and $v_2 = (0, 1)$. The outputs are just $T(v) = v$. But the output *basis* is now $w_1 = (3, 7)$ and $w_2 = (2, 5)$.

**Reverse the bases
Invert the matrix**

The matrix for $T(v) = v$ is $\begin{bmatrix} 3 & 2 \\ 7 & 5 \end{bmatrix}^{-1} = \begin{bmatrix} 5 & -2 \\ -7 & 3 \end{bmatrix}$.

Reason The first input is $v_1 = (1, 0)$. The output is also v_1 but we express it as $5w_1 - 7w_2$. Check that $5(3, 7) - 7(2, 5)$ does produce $(1, 0)$. We are combining the columns of the previous M to get the columns of I . The matrix to do that is M^{-1} .

$$\begin{array}{l} \text{Change basis} \\ \text{Change back} \end{array} \quad [w_1 \quad w_2] \begin{bmatrix} 5 & -2 \\ -7 & 3 \end{bmatrix} = [v_1 \quad v_2] \text{ is } MM^{-1} = I.$$

A mathematician would say that the matrix MM^{-1} corresponds to the product of two identity transformations. We start and end with the same basis $(1, 0)$ and $(0, 1)$. Matrix multiplication must give I . So the two change of basis matrices are inverses.

One thing is sure. Multiplying A times $(1, 0, \dots, 0)$ gives column 1 of the matrix. The novelty of this section is that $(1, 0, \dots, 0)$ stands for the first vector v_1 , *written in the basis of v 's*. Then column 1 of the matrix is that same vector v_1 , *written in the standard basis*.

Wavelet Transform = Change to Wavelet Basis

Wavelets are little waves. They have different lengths and they are localized at different places. The first basis vector is not actually a wavelet, it is the very useful flat vector of all ones. This example shows "Haar wavelets":

$$\text{Haar basis} \quad w_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \quad w_2 = \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} \quad w_3 = \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix} \quad w_4 = \begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \end{bmatrix}. \quad (8)$$

Those vectors are *orthogonal*, which is good. You see how w_3 is localized in the first half and w_4 is localized in the second half. The **wavelet transform** finds the coefficients c_1, c_2, c_3, c_4 when the input signal $v = (v_1, v_2, v_3, v_4)$ is expressed in the wavelet basis:

$$\text{Transform } v \text{ to } c \quad v = c_1 w_1 + c_2 w_2 + c_3 w_3 + c_4 w_4 = Wc \quad (9)$$


The coefficients c_3 and c_4 tell us about details in the first half and last half of v . The coefficient c_1 is the average.

Why do want to change the basis? I think of v_1, v_2, v_3, v_4 as the intensities of a signal. In audio they are volumes of sound. In images they are pixel values on a scale of black to white. An electrocardiogram is a medical signal. Of course $n = 4$ is very short, and $n = 10,000$ is more realistic. We may need to *compress* that long signal, by keeping only the largest 5% of the coefficients. This is 20 : 1 compression and (to give only two of its applications) it makes High Definition TV and video conferencing possible.

If we keep only 5% of the *standard* basis coefficients, we lose 95% of the signal. In image processing, 95% of the image disappears. In audio, 95% of the tape goes blank. But if we choose a better basis of w 's, 5% of the basis vectors can combine to come very close to the original signal. In image processing and audio coding, you can't see or hear the difference. We don't need the other 95%!

One good basis vector is the flat $(1, 1, 1, 1)$. That part alone can represent the constant background of our image. A short wave like $(0, 0, 1, -1)$ or in higher dimensions $(0, 0, 0, 0, 0, 0, 1, -1)$ represents a detail at the end of the signal.

The three steps are the transform and compression and inverse transform.



In linear algebra, where everything is perfect, we omit the compression step. The output \hat{v} is exactly the same as the input v . The transform gives $c = W^{-1}v$ and the reconstruction brings back $v = Wc$. In true signal processing, where nothing is perfect but everything is fast, the transform (lossless) and the compression (which only loses unnecessary information) are absolutely the keys to success. The output is $\hat{v} = W\hat{c}$.

I will show those steps for a typical vector like $v = (6, 4, 5, 1)$. Its wavelet coefficients are $c = (4, 1, 1, 2)$. The reconstruction $4w_1 + w_2 + w_3 + 2w_4$ is $v = Wc$:

$$\begin{bmatrix} 6 \\ 4 \\ 5 \\ 1 \end{bmatrix} = 4 \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} + \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & -1 & 0 \\ 1 & -1 & 0 & 1 \\ 1 & -1 & 0 & -1 \end{bmatrix} \begin{bmatrix} 4 \\ 1 \\ 1 \\ 2 \end{bmatrix}. \quad (10)$$

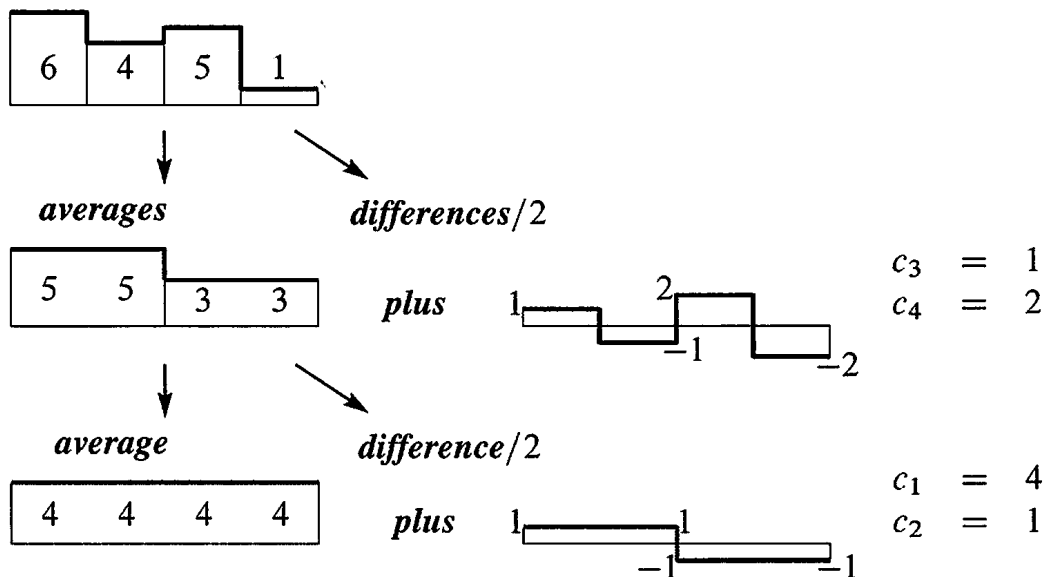
Those coefficients c are $W^{-1}v$. Inverting this basis matrix W is easy because the w 's in its columns are orthogonal. But they are not unit vectors, so rescale:

$$W^{-1} = \begin{bmatrix} \frac{1}{4} & & & \\ & \frac{1}{4} & & \\ & & \frac{1}{2} & \\ & & & \frac{1}{2} \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}.$$

The $\frac{1}{4}$'s in the first row of $c = W^{-1}v$ mean that $c_1 = 4$ is the average of 6, 4, 5, 1.

Example 11 (Same wavelet basis by recursion) I can't resist showing you a faster way to find the c 's. The special point of the wavelet basis is that you can pick off the details in c_3 and c_4 , before the coarse details in c_2 and the overall average in c_1 . A picture will explain this "multiscale" method, which is in Chapter 1 of my textbook with Nguyen on *Wavelets and Filter Banks* (Wellesley-Cambridge Press).

Split $v = (6, 4, 5, 1)$ into averages and waves at small scale and then large scale:



Fourier Transform (DFT) = Change to Fourier Basis

The first thing an electrical engineer does with a signal is to take its Fourier transform. For finite vectors we are speaking about the *Discrete Fourier Transform*. The DFT involves complex numbers (powers of $e^{2\pi i/n}$). But if we choose $n = 4$, the matrices are small and the only complex numbers are i and $i^3 = -i$. A true electrical engineer would write j instead of i for $\sqrt{-1}$.

$$\text{Fourier basis } w_1 \text{ to } w_n \text{ in the columns of } F \quad F = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & i^2 & i^3 \\ 1 & i^2 & i^4 & i^6 \\ 1 & i^3 & i^6 & i^9 \end{bmatrix}$$

The first column is the useful flat basis vector $(1, 1, 1, 1)$. It represents the average signal or the direct current (the DC term). It is a wave at zero frequency. The third column is $(1, -1, 1, -1)$, which alternates at the highest frequency. *The Fourier transform decomposes the signal into waves at equally spaced frequencies.*

The Fourier matrix F is absolutely the most important complex matrix in mathematics and science and engineering. Section 10.3 of this book explains the *Fast Fourier Transform*: it can be seen as a factorization of F into matrices with many zeros. The FFT has revolutionized entire industries, by speeding up the Fourier transform. The beautiful thing is that F^{-1} looks like F , with i changed to $-i$:

$$\begin{array}{l} \text{Fourier transform } v \text{ to } c \\ v = c_1 w_1 + \dots + c_n w_n = Fc \\ \text{Fourier coefficients } c = F^{-1}v \end{array} \quad F^{-1} = \frac{1}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & (-i) & (-i)^2 & (-i)^3 \\ 1 & (-i)^2 & (-i)^4 & (-i)^6 \\ 1 & (-i)^3 & (-i)^6 & (-i)^9 \end{bmatrix} = \frac{1}{4} \overline{F}.$$

The MATLAB command $c = \text{fft}(v)$ produces the Fourier coefficients c_1, \dots, c_n of the vector v . It multiplies v by F^{-1} (fast).

■ REVIEW OF THE KEY IDEAS ■

1. If we know $T(v_1), \dots, T(v_n)$ for a basis, linearity will determine all other $T(v)$.

2. $\left\{ \begin{array}{l} \text{Linear transformation } T \\ \text{Input basis } v_1, \dots, v_n \\ \text{Output basis } w_1, \dots, w_m \end{array} \right\} \rightarrow \begin{array}{l} \text{Matrix } A \text{ (} m \text{ by } n \text{)} \\ \text{represents } T \\ \text{in these bases} \end{array}$

3. The derivative and integral matrices are one-sided inverses: $d(\text{constant})/dx = 0$:

(Derivative) (Integral) = I is the Fundamental Theorem of Calculus.

4. If A and B represent T and S , and the output basis for S is the input basis for T , then the matrix AB represents the transformation $T(S(u))$.

5. The change of basis matrix M represents $T(v) = v$. Its columns are the coefficients of the output basis expressed in the input basis: $w_j = m_{1j}v_1 + \cdots + m_{nj}v_n$.

■ WORKED EXAMPLES ■

7.2 A Using the standard basis, find the 4 by 4 matrix P that represents a *cyclic permutation* T from $x = (x_1, x_2, x_3, x_4)$ to $T(x) = (x_4, x_1, x_2, x_3)$. Find the matrix for T^2 . What is the triple shift $T^3(x)$ and why is $T^3 = T^{-1}$?

Find two real independent eigenvectors of P . What are all the eigenvalues of P ?

Solution The first vector $(1, 0, 0, 0)$ in the standard basis transforms to $(0, 1, 0, 0)$ which is the second basis vector. So the first column of P is $(0, 1, 0, 0)$. The other three columns come from transforming the other three standard basis vectors:

$$P = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad \text{Then } P \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} x_4 \\ x_1 \\ x_2 \\ x_3 \end{bmatrix} \text{ copies } T.$$

Since we used the standard basis, T is ordinary multiplication by P . The matrix for T^2 is a “double cyclic shift” P^2 and it produces (x_3, x_4, x_1, x_2) .

The triple shift T^3 will transform $x = (x_1, x_2, x_3, x_4)$ to $T^3(x) = (x_2, x_3, x_4, x_1)$. If we apply T once more we are back to the original x . So $T^4 =$ identity transformation and $P^4 =$ identity matrix.

Two real eigenvectors of P are $(1, 1, 1, 1)$ with eigenvalue $\lambda = 1$ and $(1, -1, 1, -1)$ with eigenvalue $\lambda = -1$. The shift leaves $(1, 1, 1, 1)$ unchanged and it reverses signs in $(1, -1, 1, -1)$. The other eigenvalues are i and $-i$. The determinant is $\lambda_1\lambda_2\lambda_3\lambda_4 = -1$.

Notice that the eigenvalues $1, i, -1, -i$ add to zero (the trace of P). They are the 4th roots of 1, since $\det(P - \lambda I) = \lambda^4 - 1$. They are at angles $0^\circ, 90^\circ, 180^\circ, 270^\circ$ in the complex plane. **The Fourier matrix F is the eigenvector matrix for P .**

7.2 B The space of 2 by 2 matrices has these four “vectors” as a basis:

$$v_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad v_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \quad v_3 = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \quad v_4 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

T is the linear transformation that *transposes* every 2 by 2 matrix. What is the matrix A that represents T in this basis (output basis = input basis)? What is the inverse matrix A^{-1} ? What is the transformation T^{-1} that inverts the transpose operation?

Solution Transposing those four “basis matrices” just reverses v_2 and v_3 :

$$\begin{array}{l} T(v_1) = v_1 \\ T(v_2) = v_3 \\ T(v_3) = v_2 \\ T(v_4) = v_4 \end{array} \quad \text{gives the four columns of} \quad A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

The inverse matrix A^{-1} is the same as A . The inverse transformation T^{-1} is the same as T . If we transpose and transpose again, the final output equals the original input.

Problem Set 7.2

Questions 1–4 extend the first derivative example to higher derivatives.

- 1 The transformation S takes the *second derivative*. Keep $1, x, x^2, x^3$ as the basis v_1, v_2, v_3, v_4 and also as w_1, w_2, w_3, w_4 . Write sv_1, sv_2, sv_3, sv_4 in terms of the w 's. Find the 4 by 4 matrix B for S .
- 2 What functions have $v'' = \mathbf{0}$? They are in the kernel of the second derivative S . What vectors are in the nullspace of its matrix B in Problem 1?
- 3 B is not the square of a rectangular first derivative matrix:

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix} \text{ does not allow } A^2.$$

Add a zero row to A , so that output space = input space. Compare A^2 with B . Conclusion: For $B = A^2$ we want output basis = _____ basis. Then $m = n$.

- 4 (a) The product TS of first and second derivatives produces the *third* derivative. Add zeros to make 4 by 4 matrices, then compute AB .
(b) The matrix B^2 corresponds to $S^2 = \text{fourth derivative}$. Why is this zero?

Questions 5–9 are about a particular T and its matrix A .

- 5 With bases v_1, v_2, v_3 and w_1, w_2, w_3 , suppose $T(v_1) = w_2$ and $T(v_2) = T(v_3) = w_1 + w_3$. T is a linear transformation. Find the matrix A and multiply by the vector $(1, 1, 1)$. What is the output from T when the input is $v_1 + v_2 + v_3$?
- 6 Since $T(v_2) = T(v_3)$, the solutions to $T(v) = \mathbf{0}$ are $v = \underline{\hspace{2cm}}$. What vectors are in the nullspace of A ? Find all solutions to $T(v) = w_2$.
- 7 Find a vector that is not in the column space of A . Find a combination of w 's that is not in the range of T .
- 8 You don't have enough information to determine T^2 . Why is its matrix not necessarily A^2 ? What more information do you need?

- 9 Find the rank of A . This is not the dimension of the output space \mathbf{W} . It is the dimension of the _____ of T .

Questions 10–13 are about invertible linear transformations.

- 10 Suppose $T(v_1) = w_1 + w_2 + w_3$ and $T(v_2) = w_2 + w_3$ and $T(v_3) = w_3$. Find the matrix A for T using these basis vectors. What input vector v gives $T(v) = w_1$?
- 11 Invert the matrix A in Problem 10. Also invert the transformation T —what are $T^{-1}(w_1)$ and $T^{-1}(w_2)$ and $T^{-1}(w_3)$?
- 12 Which of these are true and why is the other one ridiculous?
- (a) $T^{-1}T = I$ (b) $T^{-1}(T(v_1)) = v_1$ (c) $T^{-1}(T(w_1)) = w_1$.
- 13 Suppose the spaces \mathbf{V} and \mathbf{W} have the same basis v_1, v_2 .
- (a) Describe a transformation T (not I) that is its own inverse.
- (b) Describe a transformation T (not I) that equals T^2 .
- (c) Why can't the same T be used for both (a) and (b)?

Questions 14–19 are about changing the basis.

- 14 (a) What matrix transforms $(1, 0)$ into $(2, 5)$ and transforms $(0, 1)$ to $(1, 3)$?
- (b) What matrix transforms $(2, 5)$ to $(1, 0)$ and $(1, 3)$ to $(0, 1)$?
- (c) Why does no matrix transform $(2, 6)$ to $(1, 0)$ and $(1, 3)$ to $(0, 1)$?
- 15 (a) What matrix M transforms $(1, 0)$ and $(0, 1)$ to (r, t) and (s, u) ?
- (b) What matrix N transforms (a, c) and (b, d) to $(1, 0)$ and $(0, 1)$?
- (c) What condition on a, b, c, d will make part (b) impossible?
- 16 (a) How do M and N in Problem 15 yield the matrix that transforms (a, c) to (r, t) and (b, d) to (s, u) ?
- (b) What matrix transforms $(2, 5)$ to $(1, 1)$ and $(1, 3)$ to $(0, 2)$?
- 17 If you keep the same basis vectors but put them in a different order, the change of basis matrix M is a _____ matrix. If you keep the basis vectors in order but change their lengths, M is a _____ matrix.
- 18 The matrix that rotates the axis vectors $(1, 0)$ and $(0, 1)$ through an angle θ is Q . What are the coordinates (a, b) of the original $(1, 0)$ using the new (rotated) axes? This *inverse* can be tricky. Draw a figure or solve for a and b :

$$Q = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad \begin{bmatrix} 1 \\ 0 \end{bmatrix} = a \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} + b \begin{bmatrix} -\sin \theta \\ \cos \theta \end{bmatrix}.$$

- 19 The matrix that transforms $(1, 0)$ and $(0, 1)$ to $(1, 4)$ and $(1, 5)$ is $M = \underline{\hspace{2cm}}$. The combination $a(1, 4) + b(1, 5)$ that equals $(1, 0)$ has $(a, b) = (\underline{\hspace{1cm}}, \underline{\hspace{1cm}})$. How are those new coordinates of $(1, 0)$ related to M or M^{-1} ?

Questions 20–23 are about the space of quadratic polynomials $A + Bx + Cx^2$.

- 20 The parabola $w_1 = \frac{1}{2}(x^2 + x)$ equals one at $x = 1$, and zero at $x = 0$ and $x = -1$. Find the parabolas w_2, w_3 , and then find $y(x)$ by linearity.
- (a) w_2 equals one at $x = 0$ and zero at $x = 1$ and $x = -1$.
- (b) w_3 equals one at $x = -1$ and zero at $x = 0$ and $x = 1$.
- (c) $y(x)$ equals 4 at $x = 1$ and 5 at $x = 0$ and 6 at $x = -1$. Use w_1, w_2, w_3 .
- 21 One basis for second-degree polynomials is $v_1 = 1$ and $v_2 = x$ and $v_3 = x^2$. Another basis is w_1, w_2, w_3 from Problem 20. Find two change of basis matrices, from the w 's to the v 's and from the v 's to the w 's.
- 22 What are the three equations for A, B, C if the parabola $Y = A + Bx + Cx^2$ equals 4 at $x = a$ and 5 at $x = b$ and 6 at $x = c$? Find the determinant of the 3 by 3 matrix. That matrix transforms values like 4, 5, 6 to parabolas—or is it the other way?
- 23 Under what condition on the numbers m_1, m_2, \dots, m_9 do these three parabolas give a basis for the space of all parabolas?
- $$v_1 = m_1 + m_2x + m_3x^2, \quad v_2 = m_4 + m_5x + m_6x^2, \quad v_3 = m_7 + m_8x + m_9x^2.$$
- 24 The Gram-Schmidt process changes a basis a_1, a_2, a_3 to an orthonormal basis q_1, q_2, q_3 . These are columns in $A = QR$. Show that R is the change of basis matrix from the a 's to the q 's (a_2 is what combination of q 's when $A = QR$?).
- 25 Elimination changes the rows of A to the rows of U with $A = LU$. Row 2 of A is what combination of the rows of U ? Writing $A^T = U^T L^T$ to work with columns, the change of basis matrix is $M = L^T$. (We have bases provided the matrices are $\underline{\hspace{2cm}}$.)
- 26 Suppose v_1, v_2, v_3 are **eigenvectors** for T . This means $T(v_i) = \lambda_i v_i$ for $i = 1, 2, 3$. What is the matrix for T when the input and output bases are the v 's?
- 27 Every invertible linear transformation can have I as its matrix! Choose any input basis v_1, \dots, v_n . For output basis choose $w_i = T(v_i)$. Why must T be invertible?
- 28 Using $v_1 = w_1$ and $v_2 = w_2$ find the standard matrix for these T 's:
- (a) $T(v_1) = \mathbf{0}$ and $T(v_2) = 3v_1$ (b) $T(v_1) = v_1$ and $T(v_1 + v_2) = v_1$.
- 29 Suppose T is reflection across the x axis and S is reflection across the y axis. The domain V is the xy plane. If $v = (x, y)$ what is $S(T(v))$? Find a simpler description of the product ST .

- 30 Suppose T is reflection across the 45° line, and S is reflection across the y axis. If $v = (2, 1)$ then $T(v) = (1, 2)$. Find $S(T(v))$ and $T(S(v))$. This shows that generally $ST \neq TS$.
- 31 Show that the product ST of two reflections is a rotation. Multiply these reflection matrices to find the rotation angle:

$$\begin{bmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{bmatrix} \begin{bmatrix} \cos 2\alpha & \sin 2\alpha \\ \sin 2\alpha & -\cos 2\alpha \end{bmatrix}.$$

- 32 True or false: If we know $T(v)$ for n different nonzero vectors in \mathbf{R}^n , then we know $T(v)$ for every vector in \mathbf{R}^n .
- 33 Express $e = (1, 0, 0, 0)$ and $v = (1, -1, 1, -1)$ in the wavelet basis, as in equations (8-10). The coefficients c_1, c_2, c_3, c_4 solve $Wc = e$ and $Wc = v$.
- 34 To represent $v = (7, 5, 3, 1)$ in the wavelet basis, start with $(6, 6, 2, 2) + (1, -1, 1, -1)$. Then write $6, 6, 2, 2$ as an overall average plus a difference, using $1, 1, 1, 1$ and $1, 1, -1, -1$.
- 35 What are the eight vectors in the wavelet basis for \mathbf{R}^8 ? They include the long wavelet $(1, 1, 1, 1, -1, -1, -1, -1)$ and the short wavelet $(1, -1, 0, 0, 0, 0, 0, 0)$.
- 36 Suppose we have two bases v_1, \dots, v_n and w_1, \dots, w_n for \mathbf{R}^n . If a vector has coefficients b_i in one basis and c_i in the other basis, what is the change of basis matrix in $b = Mc$? Start from

$$b_1v_1 + \dots + b_nv_n = Vb = c_1w_1 + \dots + c_nw_n = Wc.$$

Your answer represents $T(v) = v$ with input basis of v 's and output basis of w 's. Because of different bases, the matrix is not I .

Challenge Problems

- 37 The space \mathbf{M} of 2 by 2 matrices has the basis v_1, v_2, v_3, v_4 in Worked Example 7.2 B. Suppose T multiplies each matrix by $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$. What 4 by 4 matrix A represents this transformation T on matrix space?
- 38 Suppose A is a 3 by 4 matrix of rank $r = 2$, and $T(v) = Av$. Choose input basis vectors v_1, v_2 from the row space of A and v_3, v_4 from the nullspace. Choose output basis $w_1 = Av_1, w_2 = Av_2$ in the column space and w_3 from the nullspace of A^T . What specially simple matrix represents this T in these special bases?

7.3 Diagonalization and the Pseudoinverse

This section produces better matrices by choosing better bases. When the goal is a diagonal matrix, one way is a basis of *eigenvectors*. The other way is two bases (the input and output bases are different). Those left and right *singular vectors* are orthonormal basis vectors for the four fundamental subspaces of A . They come from the SVD.

By reversing those input and output bases, we will find the “pseudoinverse” of A . This matrix A^+ sends \mathbf{R}^m back to \mathbf{R}^n , and it sends column space back to row space.

The truth is that all our great factorizations of A can be regarded as a change of basis. But this is a short section, so we concentrate on the two outstanding examples. In both cases the good matrix is *diagonal*. It is Λ with one basis or Σ with two bases.

1. $S^{-1}AS = \Lambda$ when the input and output bases are eigenvectors of A .
2. $U^{-1}AV = \Sigma$ when those bases are eigenvectors of $A^T A$ and AA^T .

You see immediately the difference between Λ and Σ . In Λ the bases are the same. Then $m = n$ and the matrix A must be square. And some square matrices cannot be diagonalized by any S , because they don't have n independent eigenvectors.

In Σ the input and output bases are different. The matrix A can be rectangular. The bases are *orthonormal* because $A^T A$ and AA^T are symmetric. Then $U^{-1} = U^T$ and $V^{-1} = V^T$. Every matrix A is allowed, and A has the diagonal form Σ . This is the Singular Value Decomposition (SVD) of Section 6.7.

The eigenvector basis is orthonormal only when $A^T A = AA^T$ (a “normal” matrix). That includes symmetric and antisymmetric and orthogonal matrices (*special* might be a better word than normal). In this case the singular values in Σ are the absolute values $\sigma_i = |\lambda_i|$, so that $\Sigma = \mathbf{abs}(\Lambda)$. The two diagonalizations are the same when $A^T A = AA^T$, except for possible factors -1 (real) and $e^{i\theta}$ (complex).

I will just note that the Gram-Schmidt factorization $A = QR$ chooses only *one* new basis. That is the orthogonal output basis given by Q . The input uses the standard basis given by I . We don't reach a diagonal Σ , but we do reach a triangular R . The output basis matrix appears on the left and the input basis appears on the right, in $A = QRI$.

We start with input basis equal to output basis. That will produce S and S^{-1} .

Similar Matrices: A and $S^{-1}AS$ and $W^{-1}AW$

Begin with a square matrix and one basis. The input space \mathbf{V} is \mathbf{R}^n and the output space \mathbf{W} is also \mathbf{R}^n . The standard basis vectors are the columns of I . The matrix is n by n , and we call it A . The linear transformation T is “multiplication by A ”.

Most of this book has been about one fundamental problem—to *make the matrix simple*. We made it triangular in Chapter 2 (by elimination) and Chapter 4 (by Gram-Schmidt). We made it diagonal in Chapter 6 (by *eigenvectors*). Now that change from A to Λ comes from a *change of basis: Eigenvalue matrix from eigenvector basis*.

Here are the main facts in advance. When you change the basis for \mathbf{V} , the matrix changes from A to AM . Because \mathbf{V} is the input space, the matrix M goes on the right (to come first). When you change the basis for \mathbf{W} , the new matrix is $M^{-1}A$. We are working with the output space so M^{-1} is on the left (to come last).

If you change both bases in the same way, the new matrix is $M^{-1}AM$. The good basis vectors are the eigenvectors of A , when the matrix becomes $S^{-1}AS = \Lambda$.

When the basis contains the eigenvectors x_1, \dots, x_n , the matrix for T becomes Λ .

Reason To find column 1 of the matrix, input the first basis vector x_1 . The transformation multiplies by A . The output is $Ax_1 = \lambda_1 x_1$. This is λ_1 times the first basis vector plus zero times the other basis vectors. Therefore the first column of the matrix is $(\lambda_1, 0, \dots, 0)$.

In the eigenvector basis, the matrix is diagonal.

Example 1 Project onto the line $y = -x$ that goes from northwest to southeast. The vector $(1, 0)$ projects to $(.5, -.5)$ on that line. The projection of $(0, 1)$ is $(-.5, .5)$:

1. **Standard matrix: Project standard basis**
$$A = \begin{bmatrix} .5 & -.5 \\ -.5 & .5 \end{bmatrix}.$$

2. **Find the diagonal matrix Λ in the eigenvector basis.**

Solution The eigenvectors for this projection are $x_1 = (1, -1)$ and $x_2 = (1, 1)$. The first eigenvector lies on the 135° line and the second is perpendicular (on the 45° line). Their projections are x_1 and $\mathbf{0}$. The eigenvalues are $\lambda_1 = 1$ and $\lambda_2 = 0$.

2. **Diagonalized matrix: Project eigenvectors**
$$\Lambda = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}.$$

3. **Find a third matrix B using another basis $v_1 = w_1 = (2, 0)$ and $v_2 = w_2 = (1, 1)$.**

Solution w_1 is not an eigenvector, so the matrix B in this basis will not be diagonal. The first way to compute B follows the rule of Section 7.2:

Find column j of the matrix by writing the projection $T(v_j)$ as a combination of w 's.

Apply the projection T to $(2, 0)$. The result is $(1, -1)$ which is $w_1 - w_2$. So the first column of B contains 1 and -1 . The second vector $w_2 = (1, 1)$ projects to zero, so the second column of B contains 0 and 0. The eigenvalues must stay at 1 and 0:

3. **Third similar matrix: Project w_1 and w_2**
$$B = \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix}. \quad (1)$$

The second way to find the same B is more insightful. Use W^{-1} and W to change between the standard basis and the basis of w 's. Those change of basis matrices are

representing the identity transformation! The product of transformations is just ITI . The product of matrices is $B = W^{-1}AW$. This approach shows that B is *similar* to A .

For any basis w_1, \dots, w_n find the matrix B in three steps. Change the input basis to the standard basis with W . The matrix in the standard basis is A . Change the output basis back to the w 's with W^{-1} . Then $B = W^{-1}AW$ represents ITI :

$$B_{w\text{'s to } w\text{'s}} = W_{\text{standard to } w\text{'s}}^{-1} A_{\text{standard}} W_{w\text{'s to standard}} \quad (2)$$

A change of basis produces a similarity transformation to $W^{-1}AW$ in the matrix.

Example 2 (continuing with the projection) Apply this $W^{-1}AW$ rule to find B , when the basis $(2, 0)$ and $(1, 1)$ is in the columns of W :

$$W^{-1}AW = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix} \begin{bmatrix} 2 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix}.$$

The $W^{-1}AW$ rule has produced the same B as in equation (1). The matrices A and B are similar. They have the same eigenvalues (1 and 0). And Λ is similar too.

Notice that the projection matrix keeps the property $A^2 = A$ and $B^2 = B$ and $\Lambda^2 = \Lambda$. The second projection doesn't move the first projection.

The Singular Value Decomposition (SVD)

Now the input basis v_1, \dots, v_n can be different from the output basis u_1, \dots, u_m . In fact the input space \mathbf{R}^n can be different from the output space \mathbf{R}^m . Again the best matrix is diagonal (now m by n). To achieve this diagonal matrix Σ , each input vector v_j must transform into a multiple of the output vector u_j . That multiple is the *singular value* σ_j on the main diagonal of Σ :

$$\text{SVD} \quad Av_j = \begin{cases} \sigma_j u_j & \text{for } j \leq r \\ \mathbf{0} & \text{for } j > r \end{cases} \quad \text{with orthonormal bases.} \quad (3)$$

The singular values are in the order $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r$. The rank r enters because (by definition) singular values are not zero. The second part of the equation says that v_j is in the nullspace for $j = r + 1, \dots, n$. This gives the correct number $n - r$ of basis vectors for the nullspace.

Let me connect the matrices with the linear transformations they represent. A and Σ represent the *same transformation*. $A = U\Sigma V^T$ uses the standard bases for \mathbf{R}^n and \mathbf{R}^m . The diagonal Σ uses the input basis of v 's and the output basis of u 's. The orthogonal matrices V and U give the basis changes; they represent the identity transformations (in \mathbf{R}^n and \mathbf{R}^m). The product of transformations is ITI , and it is represented in the v and u bases by $U^{-1}AV$ which is Σ .

The matrix Σ in the u and v bases comes from A in the standard bases by $U^{-1}AV$:

$$\Sigma_{v\text{'s to }u\text{'s}} = U_{\text{standard to }u\text{'s}}^{-1} A_{\text{standard}} V_{v\text{'s to standard}}. \tag{4}$$

The SVD chooses orthonormal bases ($U^{-1} = U^T$ and $V^{-1} = V^T$) that diagonalize A .

The two orthonormal bases in the SVD are the eigenvector bases for $A^T A$ (the v 's) and AA^T (the u 's). Since those are symmetric matrices, their unit eigenvectors are orthonormal. Their eigenvalues are the numbers σ_j^2 . Equations (10) and (11) in Section 6.7 proved that those bases diagonalize the standard matrix A to produce Σ .

Polar Decomposition

Every complex number has the polar form $re^{i\theta}$. A nonnegative number r multiplies a number on the unit circle. (Remember that $|e^{i\theta}| = |\cos \theta + i \sin \theta| = 1$.) Thinking of these numbers as 1 by 1 matrices, $r \geq 0$ corresponds to a *positive semidefinite matrix* (call it H) and $e^{i\theta}$ corresponds to an *orthogonal matrix* Q . The *polar decomposition* extends this factorization to matrices: orthogonal times semidefinite, $A = QH$.

Every real square matrix can be factored into $A = QH$, where Q is *orthogonal* and H is *symmetric positive semidefinite*. If A is invertible, H is positive definite.

For the proof we just insert $V^T V = I$ into the middle of the SVD:

Polar decomposition $A = U \Sigma V^T = (UV^T)(V \Sigma V^T) = (Q)(H). \tag{5}$

The first factor UV^T is Q . The product of orthogonal matrices is orthogonal. The second factor $V \Sigma V^T$ is H . It is positive semidefinite because its eigenvalues are in Σ . If A is invertible then Σ and H are also invertible. **H is the symmetric positive definite square root of $A^T A$.** Equation (5) says that $H^2 = V \Sigma^2 V^T = A^T A$.

There is also a polar decomposition $A = KQ$ in the reverse order. Q is the same but now $K = U \Sigma U^T$. This is the symmetric positive definite square root of AA^T .

Example 3 Find the polar decomposition $A = QH$ from its SVD in Section 6.7:

$$A = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \sqrt{2} & \\ & 2\sqrt{2} \end{bmatrix} \begin{bmatrix} -1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix} = U \Sigma V^T.$$

Solution The orthogonal part is $Q = UV^T$. The positive definite part is $H = V \Sigma V^T$. This is also $H = Q^{-1}A$ which is $Q^T A$ because Q is orthogonal:

Orthogonal	$Q = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} -1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}$
Positive definite	$H = \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix} \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} 3/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & 3/\sqrt{2} \end{bmatrix}.$

In mechanics, the polar decomposition separates the *rotation* (in Q) from the *stretching* (in H). The eigenvalues of H are the singular values of A . They give the stretching factors. The eigenvectors of H are the eigenvectors of $A^T A$. They give the stretching directions (the principal axes). Then Q rotates those axes.

The polar decomposition just splits the key equation $Av_i = \sigma_i u_i$ into two steps. The “ H ” part multiplies v_i by σ_i . The “ Q ” part swings v_i around into u_i .

The Pseudoinverse

By choosing good bases, A multiplies v_i in the row space to give $\sigma_i u_i$ in the column space. A^{-1} must do the opposite! If $Av = \sigma u$ then $A^{-1}u = v/\sigma$. The singular values of A^{-1} are $1/\sigma$, just as the eigenvalues of A^{-1} are $1/\lambda$. The bases are reversed. The u 's are in the row space of A^{-1} , the v 's are in the column space.

Until this moment we would have added “if A^{-1} exists.” Now we don't. A matrix that multiplies u_i to produce v_i/σ_i *does* exist. It is the pseudoinverse A^+ :

$$\text{Pseudoinverse} \\ A^+ = V\Sigma^+U^T = \begin{matrix} \left[\begin{array}{ccc} v_1 & \cdots & v_r & \cdots & v_n \end{array} \right] \\ n \text{ by } n \end{matrix} \begin{matrix} \left[\begin{array}{ccc} \sigma_1^{-1} & & \\ & \ddots & \\ & & \sigma_r^{-1} \end{array} \right] \\ n \text{ by } m \end{matrix} \begin{matrix} \left[\begin{array}{ccc} u_1 & \cdots & u_r & \cdots & u_m \end{array} \right]^T \\ m \text{ by } m \end{matrix}$$

The *pseudoinverse* A^+ is an n by m matrix. If A^{-1} exists (we said it again), then A^+ is the same as A^{-1} . In that case $m = n = r$ and we are inverting $U\Sigma V^T$ to get $V\Sigma^{-1}U^T$. The new symbol A^+ is needed when $r < m$ or $r < n$. Then A has no two-sided inverse, but it has a *pseudoinverse* A^+ with that same rank r :

$$A^+u_i = \frac{1}{\sigma_i}v_i \quad \text{for } i \leq r \quad \text{and} \quad A^+u_i = \mathbf{0} \quad \text{for } i > r.$$

The vectors u_1, \dots, u_r in the column space of A go back to v_1, \dots, v_r in the row space. The other vectors u_{r+1}, \dots, u_m are in the left nullspace, and A^+ sends them to zero. When we know what happens to each basis vector u_i , we know A^+ .

Notice the pseudoinverse Σ^+ of the diagonal matrix Σ . Each σ is replaced by σ^{-1} . The product $\Sigma^+\Sigma$ is as near to the identity as we can get (it is a projection matrix, $\Sigma^+\Sigma$ is partly I and partly 0). We get r 1's. We can't do anything about the zero rows and columns. This example has $\sigma_1 = 2$ and $\sigma_2 = 3$:

$$\Sigma^+\Sigma = \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}.$$

The pseudoinverse A^+ is the n by m matrix that makes AA^+ and A^+A into projections:




Figure 7.4: Ax^+ in the column space goes back to $A^+Ax^+ = x^+$ in the row space.

Trying for
 $AA^{-1} = A^{-1}A = I$

$AA^+ =$ projection matrix onto the column space of A
 $A^+A =$ projection matrix onto the row space of A

Example 4 Find the pseudoinverse of $A = \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix}$. This matrix is not invertible. The rank is 1. The only singular value is $\sqrt{10}$. That is inverted to $1/\sqrt{10}$ in Σ^+ :

$$A^+ = V\Sigma^+U^T = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} 1/\sqrt{10} & 0 \\ 0 & 0 \end{bmatrix} \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & 1 \\ 1 & -2 \end{bmatrix} = \frac{1}{10} \begin{bmatrix} 2 & 1 \\ 2 & 1 \end{bmatrix}.$$

A^+ also has rank 1. Its column space is the row space of A . When A takes $(1, 1)$ in the row space to $(4, 2)$ in the column space, A^+ does the reverse. $A^+(4, 2) = (1, 1)$.

Every rank one matrix is a column times a row. With unit vectors u and v , that is $A = \sigma uv^T$. Then the best inverse of a rank one matrix is $A^+ = vu^T/\sigma$. The product AA^+ is uu^T , the projection onto the line through u . The product A^+A is vv^T .

Application to least squares Chapter 4 found the best solution \hat{x} to an unsolvable system $Ax = b$. The key equation is $A^T A \hat{x} = A^T b$, with the assumption that $A^T A$ is invertible. The zero vector was alone in the nullspace.

Now A may have dependent columns (rank $< n$). There can be many solutions to $A^T A \hat{x} = A^T b$. **One solution is $x^+ = A^+b$ from the pseudoinverse.** We can check that

$A^T A A^+ \mathbf{b}$, is $A^T \mathbf{b}$, because Figure 7.4 shows that $\mathbf{e} = \mathbf{b} - A A^+ \mathbf{b}$ is the part of \mathbf{b} in the nullspace of A^T . Any vector in the nullspace of A could be added to \mathbf{x}^+ , to give another solution $\hat{\mathbf{x}}$ to $A^T A \hat{\mathbf{x}} = A^T \mathbf{b}$. But \mathbf{x}^+ will be shorter than any other $\hat{\mathbf{x}}$ (Problem 16):

The shortest least squares solution to $Ax = \mathbf{b}$ is $\mathbf{x}^+ = A^+ \mathbf{b}$.

The pseudoinverse A^+ and this best solution \mathbf{x}^+ are essential in statistics, because experiments often have a matrix A with **dependent columns**.

■ REVIEW OF THE KEY IDEAS ■

1. Diagonalization $S^{-1}AS = \Lambda$ is the same as a change to the eigenvector basis.
2. The SVD chooses an input basis of \mathbf{v} 's and an output basis of \mathbf{u} 's. Those orthonormal bases diagonalize A . This is $A\mathbf{v}_i = \sigma_i \mathbf{u}_i$, and in matrix form $A = U\Sigma V^T$.
3. Polar decomposition factors A into QH , rotation UV^T times stretching $V\Sigma V^T$.
4. The pseudoinverse $A^+ = V\Sigma^+U^T$ transforms the column space of A back to its row space. A^+A is the identity on the row space (and zero on the nullspace).

■ WORKED EXAMPLES ■

7.3 A If A has rank n (full column rank) then it has a *left inverse* $C = (A^T A)^{-1} A^T$. This matrix C gives $CA = I$. Explain why the pseudoinverse is $A^+ = C$ in this case. If A has rank m (full row rank) then it has a *right inverse* B with $B = A^T (AA^T)^{-1}$. Then $AB = I$. Explain why $A^+ = B$ in this case.

Find B for A_1 and find C for A_2 . Find A^+ for all three matrices A_1, A_2, A_3 :

$$A_1 = \begin{bmatrix} 2 \\ 2 \end{bmatrix} \quad A_2 = \begin{bmatrix} 2 & 2 \end{bmatrix} \quad A_3 = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}.$$

Solution If A has rank n (independent columns) then $A^T A$ is invertible—this is a key point of Section 4.2. Certainly $C = (A^T A)^{-1} A^T$ multiplies A to give $CA = I$. In the opposite order, $AC = A(A^T A)^{-1} A^T$ is the projection matrix (Section 4.2 again) onto the column space. So C meets the requirements to be A^+ : *CA and AC are projections*.

If A has rank m (full row rank) then AA^T is invertible. Certainly A multiplies $B = A^T (AA^T)^{-1}$ to give $AB = I$. In the opposite order, $BA = A^T (AA^T)^{-1} A$ is the projection matrix onto the row space. So B is the pseudoinverse A^+ with rank m .

The example A_1 has full column rank (for C) and A_2 has full row rank (for B):

$$A_1^+ = (A_1^T A_1)^{-1} A_1^T = \frac{1}{\sqrt{8}} \begin{bmatrix} 2 & 2 \end{bmatrix} \quad A_2^+ = A_2^T (A_2 A_2^T)^{-1} = \frac{1}{\sqrt{8}} \begin{bmatrix} 2 \\ 2 \end{bmatrix}.$$

Notice $A_1^+ A_1 = [1]$ and $A_2 A_2^+ = [1]$. But A_3 (rank 1) has no left or right inverse. **Its rank is not full.** Its pseudoinverse is $A_3^+ = \sigma_1^{-1} v_1 u_1^T = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} / 4$.

Problem Set 7.3

Problems 1–4 compute and use the SVD of a particular matrix (not invertible).

- 1 (a) Compute $A^T A$ and its eigenvalues and unit eigenvectors v_1 and v_2 . Find σ_1 .

$$\text{Rank one matrix } A = \begin{bmatrix} 1 & 2 \\ 3 & 6 \end{bmatrix}$$

- (b) Compute AA^T and its eigenvalues and unit eigenvectors u_1 and u_2 .
 (c) Verify that $Av_1 = \sigma_1 u_1$. Put numbers into the SVD:

$$A = U \Sigma V^T \quad \begin{bmatrix} 1 & 2 \\ 3 & 6 \end{bmatrix} = \begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_1 & v_2 \end{bmatrix}^T.$$

- 2 (a) From the u 's and v 's in Problem 1 write down orthonormal bases for the four fundamental subspaces of this matrix A .
 (b) Describe all matrices that have those same four subspaces. Multiples of A ?
- 3 From U , V , and Σ in Problem 1 find the orthogonal matrix $Q = UV^T$ and the symmetric matrix $H = V \Sigma V^T$. Verify the polar decomposition $A = QH$. This H is only semidefinite because _____. Test $H^2 = A$.
- 4 Compute the pseudoinverse $A^+ = V \Sigma^+ U^T$. The diagonal matrix Σ^+ contains $1/\sigma_1$. Rename the four subspaces (for A) in Figure 7.4 as four subspaces for A^+ . Compute the projections $P_{\text{row}} = A^+ A$ and $P_{\text{column}} = A A^+$.

Problems 5–9 are about the SVD of an invertible matrix.

- 5 Compute $A^T A$ and its eigenvalues and unit eigenvectors v_1 and v_2 . What are the singular values σ_1 and σ_2 for this matrix A ?

$$A = \begin{bmatrix} 3 & 3 \\ -1 & 1 \end{bmatrix}.$$

- 6 AA^T has the same eigenvalues σ_1^2 and σ_2^2 as $A^T A$. Find unit eigenvectors \mathbf{u}_1 and \mathbf{u}_2 . Put numbers into the SVD:

$$A = \begin{bmatrix} 3 & 3 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 \end{bmatrix} \begin{bmatrix} \sigma_1 & \\ & \sigma_2 \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 \end{bmatrix}^T.$$

- 7 In Problem 6, multiply columns times rows to show that $A = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T$. Prove from $A = U \Sigma V^T$ that every matrix of rank r is the sum of r matrices of rank one.
- 8 From U , V , and Σ find the orthogonal matrix $Q = UV^T$ and the symmetric matrix $K = U \Sigma U^T$. Verify the polar decomposition in reverse order $A = KQ$.
- 9 The pseudoinverse of this A is the same as _____ because _____.

Problems 10–11 compute and use the SVD of a 1 by 3 rectangular matrix.

- 10 Compute $A^T A$ and AA^T and their eigenvalues and unit eigenvectors when the matrix is $A = \begin{bmatrix} 3 & 4 & 0 \end{bmatrix}$. What are the singular values of A ?
- 11 Put numbers into the singular value decomposition of A :

$$A = \begin{bmatrix} 3 & 4 & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \end{bmatrix}^T.$$

Put numbers into the pseudoinverse $V \Sigma^+ U^T$ of A . Compute AA^+ and $A^+ A$:

$$\text{Pseudoinverse } A^+ = \begin{bmatrix} & & \\ & & \\ & & \end{bmatrix} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \end{bmatrix} \begin{bmatrix} 1/\sigma_1 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \end{bmatrix}^T.$$

- 12 What is the only 2 by 3 matrix that has no pivots and no singular values? What is Σ for that matrix? A^+ is the zero matrix, but what shape?
- 13 If $\det A = 0$ why is $\det A^+ = 0$? If A has rank r , why does A^+ have rank r ?
- 14 When are the factors in $U \Sigma V^T$ the same as in $Q \Lambda Q^T$? The eigenvalues λ_i must be positive, to equal the σ_i . Then A must be _____ and positive _____.

Problems 15–18 bring out the main properties of A^+ and $\mathbf{x}^+ = A^+ \mathbf{b}$.

- 15 All matrices in this problem have rank one. The vector \mathbf{b} is (b_1, b_2) .

$$A = \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix} \quad A^T = \begin{bmatrix} .2 & .1 \\ .2 & .1 \end{bmatrix} \quad AA^T = \begin{bmatrix} .8 & .4 \\ .4 & .2 \end{bmatrix} \quad A^T A = \begin{bmatrix} .5 & .5 \\ .5 & .5 \end{bmatrix}$$

- (a) The equation $A^T A \hat{\mathbf{x}} = A^T \mathbf{b}$ has many solutions because $A^T A$ is _____.
- (b) Verify that $\mathbf{x}^+ = A^+ \mathbf{b} = (.2b_1 + .1b_2, .2b_1 + .1b_2)$ solves $A^T A \mathbf{x}^+ = A^T \mathbf{b}$.

(c) Add $(1, -1)$ to that x^+ to get another solution to $A^T A \hat{x} = A^T b$. Show that $\|\hat{x}\|^2 = \|x^+\|^2 + 2$, and x^+ is shorter.

16 *The vector $x^+ = A^+ b$ is the shortest possible solution to $A^T A \hat{x} = A^T b$. Reason: The difference $\hat{x} - x^+$ is in the nullspace of $A^T A$. This is also the nullspace of A , orthogonal to x^+ . Explain how it follows that $\|\hat{x}\|^2 = \|x^+\|^2 + \|\hat{x} - x^+\|^2$.*

17 Every b in \mathbf{R}^m is $p + e$. This is the column space part plus the left nullspace part. Every x in \mathbf{R}^n is $x_r + x_n = (\text{row space part}) + (\text{nullspace part})$. Then

$$AA^+ p = \underline{\hspace{2cm}} \quad AA^+ e = \underline{\hspace{2cm}} \quad A^+ Ax_r = \underline{\hspace{2cm}} \quad A^+ Ax_n = \underline{\hspace{2cm}}$$

18 Find A^+ and $A^+ A$ and AA^+ and x^+ for this 2 by 1 matrix and these b :

$$A = \begin{bmatrix} 3 \\ 4 \end{bmatrix} = \begin{bmatrix} .6 & -.8 \\ .8 & .6 \end{bmatrix} \begin{bmatrix} 5 \\ 0 \end{bmatrix} [1] \quad b = \begin{bmatrix} 3 \\ 4 \end{bmatrix} \text{ and } b = \begin{bmatrix} -4 \\ 3 \end{bmatrix}.$$

19 A general 2 by 2 matrix A is determined by four numbers. If triangular, it is determined by three. If diagonal, by two. If a rotation, by one. An eigenvector, by one. Check that the total count is four for each factorization of A :

$$\text{Four numbers in } LU \quad LDU \quad QR \quad U\Sigma V^T \quad SAS^{-1}.$$

20 Following Problem 19, check that LDL^T and $Q\Lambda Q^T$ are determined by *three* numbers. This is correct because the matrix A is now $\underline{\hspace{2cm}}$.

21 From $A = U\Sigma V^T$ and $A^+ = V\Sigma^+ U^T$ explain these splittings into rank 1:

$$A = \sum_1^r \sigma_i u_i v_i^T \quad A^+ = \sum_1^r \frac{v_i u_i^T}{\sigma_i} \quad A^+ A = \sum_1^r v_i v_i^T \quad AA^+ = \sum_1^r u_i u_i^T$$

Challenge Problems

22 This problem looks for all matrices A with a given column space in \mathbf{R}^m and a given row space in \mathbf{R}^n . Suppose c_1, \dots, c_r and b_1, \dots, b_r are bases for those two spaces. Make them columns of C and B . The goal is to show that $A = CMB^T$ for an r by r invertible matrix M . Hint: Start from $A = U\Sigma V^T$. A must have this form:

The first r columns of U and V must be connected to C and B by invertible matrices, because they contain bases for the same column space and row space.

23 A pair of singular vectors v and u will satisfy $Av = \sigma u$ and $A^T u = \sigma v$. This means that the double vector $x = \begin{bmatrix} u \\ v \end{bmatrix}$ is an eigenvector of what symmetric block matrix? What is the eigenvalue? The SVD of A is equivalent to the diagonalization of that symmetric block matrix.