Orthogonality

3.1 ORTHOGONAL VECTORS AND SUBSPACES

A basis is a set of independent vectors that span a space. Geometrically, it is a set of coordinate axes. A vector space is defined without those axes, but every time I think of the x-y plane or three-dimensional space or $\mathbf{R}^n$, the axes are there. They are usually perpendicular! *The coordinate axes that the imagination constructs are practically always orthogonal.* In choosing a basis, we tend to choose an orthogonal basis.

The idea of an orthogonal basis is one of the foundations of linear algebra. We need a basis to convert geometric constructions into algebraic calculations, and we need an orthogonal basis to make those calculations simple. A further specialization makes the basis just about optimal: The vectors should have *length* 1. For an ***orthonormal basis*** (orthogonal unit vectors), we will find

1. the length $\|x\|$ of a vector;
2. the test $x^{\mathrm{T}}y = 0$ for perpendicular vectors; and
3. how to create perpendicular vectors from linearly independent vectors.

More than just vectors, *subspaces* can also be perpendicular. We will discover, so beautifully and simply that it will be a delight to see, that ***the fundamental subspaces meet at right angles***. Those four subspaces are perpendicular in pairs, two in $\mathbf{R}^m$ and two in $\mathbf{R}^n$. That will complete the fundamental theorem of linear algebra.

The first step is to find the ***length of a vector***. It is denoted by $\|x\|$, and in two dimensions it comes from the hypotenuse of a right triangle (Figure 3.1a). The square of the length was given a long time ago by Pythagoras: $\|x\|^2 = x_1^2 + x_2^2$.

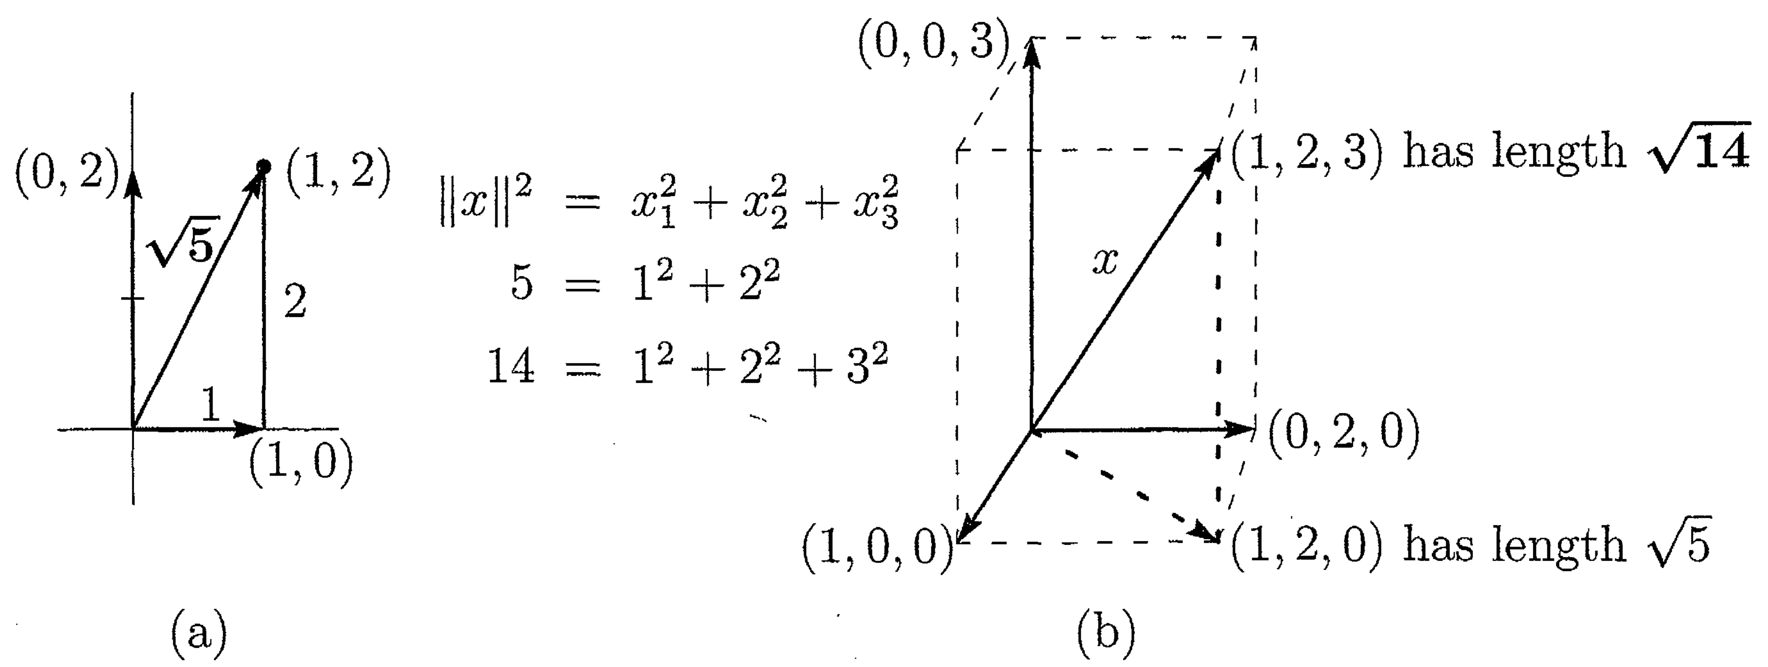

Figure 3.1 The length of vectors (x_1, x_2) and (x_1, x_2, x_3).

In three-dimensional space, $x = (x_1, x_2, x_3)$ is the diagonal of a box (Figure 3.1t
Its length comes from *two* applications of the Pythagorean formula. The two-dimension
case takes care of $(x_1, x_2, 0) = (1, 2, 0)$ across the base. This forms a right angle wi
the vertical side $(0, 0, x_3) = (0, 0, 3)$. The hypotenuse of the bold triangle (Pythagor
again) is the length $\|x\|$ we want:

$$\textbf{Length in 3D} \qquad \|x\|^2 = 1^2 + 2^2 + 3^2 \quad \text{and} \quad \|x\| = \sqrt{x_1^2 + x_2^2 + x_3^2}.$$

The extension to $x = (x_1, \ldots, x_n)$ in n dimensions is immediate. By Pythagor
$n - 1$ times, ***the length*** $\|\boldsymbol{x}\|$ ***in*** $\mathbf{R}^n$ ***is the positive square root of*** $\boldsymbol{x}^{\mathrm{T}}\boldsymbol{x}$:

$$\textbf{Length squared} \qquad \boxed{\|x\|^2 = x_1^2 + x_2^2 + \cdots + x_n^2 = x^{\mathrm{T}}x.} \qquad (\text{:}$$

The sum of squares matches $x^{\mathrm{T}}x$—and the length of $x = (1, 2, -3)$ is $\sqrt{14}$:

$$x^{\mathrm{T}}x = \begin{bmatrix} 1 & 2 & -3 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ -3 \end{bmatrix} = 1^2 + 2^2 + (-3)^2 = 14.$$

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

Orthogonal Vectors

How can we decide whether two vectors x and y are perpendicular? What is the test fc
orthogonality in Figure 3.2? In the plane spanned by x and y, those vectors are orthogonε
provided they form a *right triangle*. We go back to $a^2 + b^2 = c^2$:

$$\textbf{Sides of a right triangle} \qquad \|x\|^2 + \|y\|^2 = \|x - y\|^2. \qquad (2$$

Applying the length formula (1), this test for orthogonality in $\mathbf{R}^n$ becomes

$$\left(x_1^2 + \cdots + x_n^2\right) + \left(y_1^2 + \cdots + y_n^2\right) = (x_1 - y_1)^2 + \cdots + (x_n - y_n)^2.$$

The right-hand side has an extra $-2x_i y_i$ from each $(x_i - y_i)^2$:

$$\text{right-hand side} = \left(x_1^2 + \cdots + x_n^2\right) - 2(x_1 y_1 + \cdots + x_n y_n) + \left(y_1^2 + \cdots + y_n^2\right).$$

We have a right triangle when that sum of cross-product terms $x_i y_i$ is zero:

$$\textbf{Orthogonal vectors} \qquad \boxed{x^{\mathrm{T}}y = x_1 y_1 + \cdots + x_n y_n = 0.} \qquad (3$$

This sum is $x^{\mathrm{T}}y = \sum x_i y_i = y^{\mathrm{T}}x$, the row vector x^{T} times the column vector y:

$$\textbf{Inner product} \qquad x^{\mathrm{T}}y = \begin{bmatrix} x_1 & \ldots & x_n \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} = x_1 y_1 + \cdots + x_n y_n. \qquad (4$$

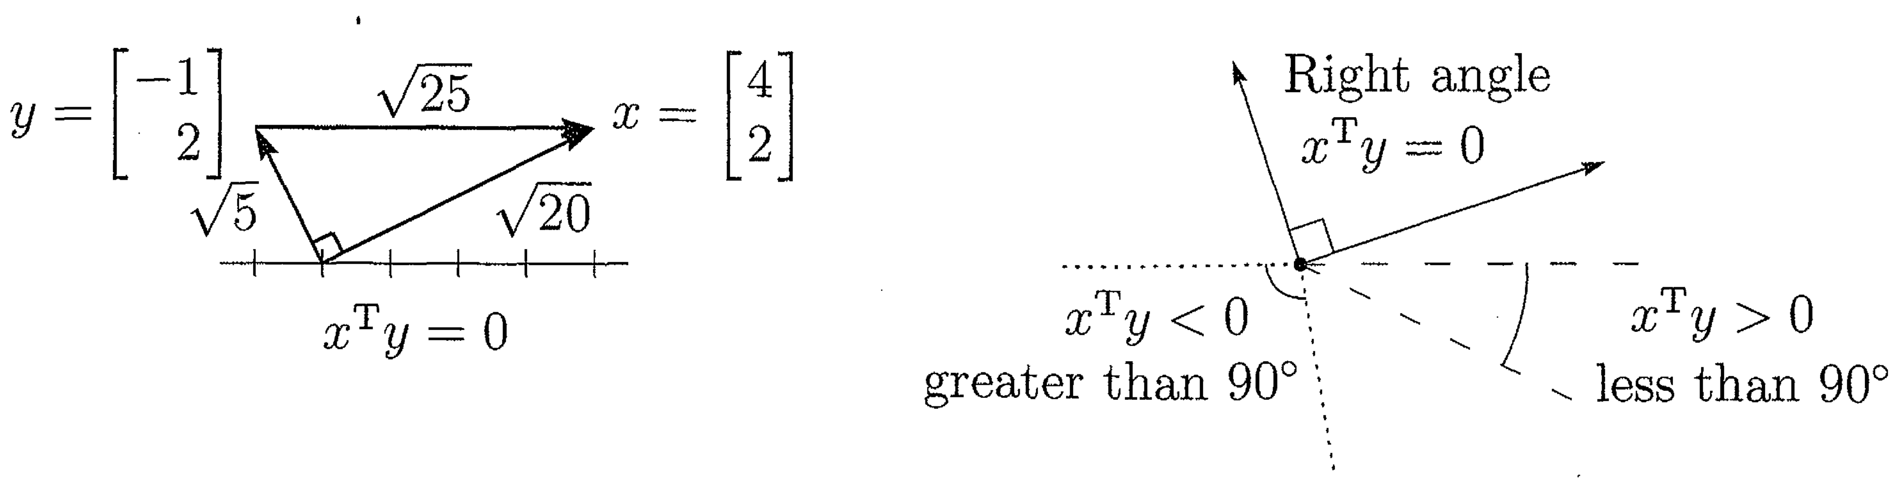

Figure 3.2 A right triangle with $5 + 20 = 25$. Dotted angle 100°, dashed angle 30°.

This number is sometimes called the scalar product or dot product, and denoted by (x, y) or $x \cdot y$. We will use the name ***inner product*** and keep the notation $x^{\mathrm{T}}y$.

3A The inner product $x^{\mathrm{T}}y$ is zero if and only if x and y are orthogonal vectors. If $x^{\mathrm{T}}y > 0$, their angle is less than 90°. If $x^{\mathrm{T}}y < 0$, their angle is greater than 90°.

The length squared is the inner product of x with itself: $x^{\mathrm{T}}x = x_1^2 + \cdots + x_n^2 = \|x\|^2$. The only vector with length zero—the only vector orthogonal to itself—is the zero vector. This vector $x = 0$ is orthogonal to every vector in $\mathbf{R}^n$.

$(2, 2, -1)$ is orthogonal to $(-1, 2, 2)$. Both have length $\sqrt{4+4+1} = 3$.

Useful fact: **If nonzero vectors $v_1, \ldots, v_k$ are mutually orthogonal** (every vector is perpendicular to every other), **then those vectors are linearly independent**.

Proof Suppose $c_1 v_1 + \cdots + c_k v_k = 0$. To show that c_1 must be zero, take the inner product of both sides with v_1. Orthogonality of the v's leaves only one term:

$$v_1^{\mathrm{T}}(c_1 v_1 + \cdots + c_k v_k) = c_1 v_1^{\mathrm{T}} v_1 = 0. \tag{5}$$

The vectors are nonzero, so $v_1^{\mathrm{T}} v_1 \neq 0$ and therefore $c_1 = 0$. The same is true of every c_i. The only combination of the v's producing zero has all $c_i = 0$: *independence*! ∎

The coordinate vectors $e_1, \ldots, e_n$ in $\mathbf{R}^n$ are the most important orthogonal vectors. Those are the columns of the identity matrix. They form the simplest basis for $\mathbf{R}^n$, and they are *unit vectors*—each has length $\|e_i\| = 1$. They point along the coordinate axes. If these axes are rotated, the result is a new **orthonormal basis**: a new system of *mutually orthogonal unit vectors*. In $\mathbf{R}^2$ we have $\cos^2\theta + \sin^2\theta = 1$:

Orthonormal vectors in $\mathbf{R}^2$ $\quad v_1 = (\cos\theta, \sin\theta)$ and $v_2 = (-\sin\theta, \cos\theta)$.

Orthogonal Subspaces

We come to the orthogonality of two subspaces. **Every vector** *in one subspace must be orthogonal to* **every vector** *in the other subspace*. Subspaces of $\mathbf{R}^3$ can have dimension 0, 1, 2, or 3. The subspaces are represented by lines or planes through the origin—and in the extreme cases, by the origin alone or the whole space. The subspace $\{0\}$ is orthogonal to all subspaces. A line can be orthogonal to another line, or it can be orthogonal to a plane, but *a plane cannot be orthogonal to a plane*.

I have to admit that the front wall and side wall of a room look like perpendicular planes in $\mathbf{R}^3$. But by our definition, that is not so! There are lines v and w in the front and side walls that do not meet at a right angle. The line along the corner is in *both* walls, and it is certainly not orthogonal to itself.

3B Two subspaces $\mathbf{V}$ and $\mathbf{W}$ of the same space $\mathbf{R}^n$ are *orthogonal* if every vector v in $\mathbf{V}$ is orthogonal to every vector w in $\mathbf{W}$: $v^{\mathrm{T}}w = 0$ for all v and w.

Suppose $\mathbf{V}$ is the plane spanned by $v_1 = (1, 0, 0, 0)$ and $v_2 = (1, 1, 0, 0)$. If $\mathbf{W}$ is the line spanned by $w = (0, 0, 4, 5)$, then w is orthogonal to both v's. The line $\mathbf{W}$ will be orthogonal to the whole plane $\mathbf{V}$.

In this case, with subspaces of dimension 2 and 1 in $\mathbf{R}^4$, there is room for a third subspace. The line $\mathbf{L}$ through $z = (0, 0, 5, -4)$ is perpendicular to $\mathbf{V}$ and $\mathbf{W}$. Then the dimensions add to $2 + 1 + 1 = 4$. What space is perpendicular to all of $\mathbf{V}$, $\mathbf{W}$, and $\mathbf{L}$?

The important orthogonal subspaces don't come by accident, and they come two at a time. In fact orthogonal subspaces are unavoidable: ***They are the fundamental subspaces!*** The first pair is the *nullspace* and *row space*. Those are subspaces of $\mathbf{R}^n$—the rows have n components and so does the vector x in $Ax = 0$. We have to show, using $Ax = 0$, that ***the rows of A are orthogonal to the nullspace vector x***.

3C Fundamental theorem of orthogonality The row space is orthogonal to the nullspace (in $\mathbf{R}^n$). The column space is orthogonal to the left nullspace (in $\mathbf{R}^m$).

First proof Suppose x is a vector in the nullspace. Then $Ax = 0$, and this system of m equations can be written out as rows of A multiplying x:

Every row is orthogonal to x

$$Ax = \begin{bmatrix} \cdots & \textbf{row 1} & \cdots \\ \cdots & \text{row } 2 & \cdots \\ & \vdots & \\ \cdots & \text{row } m & \cdots \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \tag{6}$$

The main point is already in the first equation: *row* 1 *is orthogonal to* x. Their inner product is zero; that is equation 1. Every right-hand side is zero, so x is orthogonal to every row. Therefore x is orthogonal to every *combination* of the rows. Each x in the nullspace is orthogonal to each vector in the row space, so $\boldsymbol{N}(A) \perp \boldsymbol{C}(A^{\mathrm{T}})$.

The other pair of orthogonal subspaces comes from $A^{\mathrm{T}}y = 0$, or $y^{\mathrm{T}}A = 0$:

$$y^{\mathrm{T}}A = \begin{bmatrix} y_1 & \cdots & y_m \end{bmatrix} \begin{bmatrix} \mathbf{c} & & \mathrm{c} \\ \mathbf{o} & & \mathrm{o} \\ \mathbf{l} & & \mathrm{l} \\ \mathbf{u} & \cdots & \mathrm{u} \\ \mathbf{m} & & \mathrm{m} \\ \mathbf{n} & & \mathrm{n} \\ & & \\ \mathbf{1} & & n \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \cdots & 0 \end{bmatrix}. \tag{7}$$

The vector y is orthogonal to every column. The equation says so, from the zeros on the right-hand side. Therefore y is orthogonal to every combination of the columns.

It is orthogonal to the column space, and it is a typical vector in the left nullspace: $N(A^T) \perp C(A)$. This is the same as the first half of the theorem, with A replaced by A^T. ■

Second proof The contrast with this "coordinate-free proof" should be useful to the reader. It shows a more "abstract" method of reasoning. I wish I knew which proof is clearer, and more permanently understood.

If x is in the nullspace then $Ax = 0$. If v is in the row space, it is a combination of the rows: $v = A^T z$ for some vector z. Now, in one line:

Nullspace $\perp$ Row space $$v^T x = (A^T z)^T x = z^T A x = z^T 0 = 0. \tag{8}$$ ■

Suppose A has rank 1, so its row space and column space are lines:

Rank-1 matrix $$A = \begin{bmatrix} 1 & 3 \\ 2 & 6 \\ 3 & 9 \end{bmatrix}.$$

$\begin{bmatrix}1 & 0\end{bmatrix}\begin{bmatrix}0\\-1\end{bmatrix}$ 1×2 2×1

The rows are multiples of $(1, 3)$. The nullspace contains $x = (-3, 1)$, which is orthogonal to all the rows. The nullspace and row space are perpendicular lines in $\mathbf{R}^2$:

$$\begin{bmatrix} 1 & 3 \end{bmatrix} \begin{bmatrix} 3 \\ -1 \end{bmatrix} = 0 \quad \text{and} \quad \begin{bmatrix} 2 & 6 \end{bmatrix} \begin{bmatrix} 3 \\ -1 \end{bmatrix} = 0 \quad \text{and} \quad \begin{bmatrix} 3 & 9 \end{bmatrix} \begin{bmatrix} 3 \\ -1 \end{bmatrix} = 0.$$

In contrast, the other two subspaces are in $\mathbf{R}^3$. The column space is the line through $(1, 2, 3)$. The left nullspace must be the *perpendicular plane* $y_1 + 2y_2 + 3y_3 = 0$. That equation is exactly the content of $y^T A = 0$.

The first two subspaces (the two lines) had dimensions $1 + 1 = 2$ in the space $\mathbf{R}^2$. The second pair (line and plane) had dimensions $1 + 2 = 3$ in the space $\mathbf{R}^3$. In general, *the row space and nullspace have dimensions that add to* $r + (n - r) = n$. The other pair adds to $r + (m - r) = m$. Something more than orthogonality is occurring, and I have to ask your patience about that one further point: **the dimensions**.

It is certainly true that the nullspace is perpendicular to the row space—but it is not the whole truth. $N(A)$ *contains every vector orthogonal to the row space*. The nullspace was formed from *all* solutions to $Ax = 0$.

DEFINITION Given a subspace $\mathbf{V}$ of $\mathbf{R}^n$, the space of *all* vectors orthogonal to $\mathbf{V}$ is called the **orthogonal complement** of $\mathbf{V}$. It is denoted by $\mathbf{V}^\perp$ = "$\mathbf{V}$ perp."

Using this terminology, the nullspace is the orthogonal complement of the row space: $N(A) = (C(A^T))^\perp$. At the same time, the row space contains all vectors that are orthogonal to the nullspace. A vector z can't be orthogonal to the nullspace but outside the row space. Adding z as an extra row of A would enlarge the row space, but we know that there is a fixed formula $r + (n - r) = n$:

Dimension formula $\quad \dim(\text{row space}) + \dim(\text{nullspace}) = \text{number of columns}.$

Every vector orthogonal to the nullspace is in the row space: $C(A^T) = (N(A))^\perp$.

The same reasoning applied to A^{T} produces the dual result: *The left nullspace $\boldsymbol{N}(A^{\mathrm{T}})$ and the column space $\boldsymbol{C}(A)$ are orthogonal complements.* Their dimensions add up to $(m-r)+r=m$. This completes the second half of the fundamental theorem of linear algebra. The first half gave the dimensions of the four subspaces, including the fact that row rank = column rank. Now we know that those subspaces are perpendicular. More than that, the subspaces are orthogonal complements.

3D Fundamental Theorem of Linear Algebra, Part II

The nullspace is the *orthogonal complement* of the row space in $\mathbf{R}^n$.
The left nullspace is the *orthogonal complement* of the column space in $\mathbf{R}^m$.

To repeat, the row space contains everything orthogonal to the nullspace. The column space contains everything orthogonal to the left nullspace. That is just a sentence, hidden in the middle of the book, but *it decides exactly which equations can be solved*! Looked at directly, $Ax=b$ requires b to be in the column space. Looked at indirectly, $Ax=b$ **requires b to be perpendicular to the left nullspace.**

3E $Ax=b$ is solvable if and only if $y^{\mathrm{T}}b=0$ whenever $y^{\mathrm{T}}A=0$.

The direct approach was "b must be a combination of the columns." The indirect approach is "b must be orthogonal to every vector that is orthogonal to the columns." That doesn't sound like an improvement (to put it mildly). But if only one or two vectors are orthogonal to the columns, it is much easier to check those one or two conditions $y^{\mathrm{T}}b=0$. A good example is Kirchhoff's Voltage Law in Section 2.5. Testing for zero around loops is much easier than recognizing combinations of the columns.

When the left-hand sides of $Ax=b$ add to zero, the right-hand sides must, too:

$$\begin{aligned} x_1-x_2&=b_1\\ x_2-x_3&=b_2\\ x_3-x_1&=b_3\end{aligned} \quad \text{is solvable if and only if } b_1+b_2+b_3=0. \text{ Here } A=\begin{bmatrix}1&-1&0\\0&1&-1\\-1&0&1\end{bmatrix}.$$

This test $b_1+b_2+b_3=0$ makes b orthogonal to $y=(1,1,1)$ in the left nullspace. By the Fundamental Theorem, b is a combination of the columns!

The Matrix and the Subspaces

We emphasize that $\mathbf{V}$ and $\mathbf{W}$ can be orthogonal without being complements. Their dimensions can be too small. The line $\mathbf{V}$ spanned by $(0,1,0)$ is orthogonal to the line $\mathbf{W}$ spanned by $(0,0,1)$, but $\mathbf{V}$ is not $\mathbf{W}^{\perp}$. The orthogonal complement of $\mathbf{W}$ is a two-dimensional plane, and the line is only part of $\mathbf{W}^{\perp}$. When the dimensions are right, orthogonal subspaces *are* necessarily orthogonal complements:

$$\text{If} \quad \mathbf{W}=\mathbf{V}^{\perp} \quad \text{then} \quad \mathbf{V}=\mathbf{W}^{\perp} \quad \text{and} \quad \dim\mathbf{V}+\dim\mathbf{W}=n.$$

In other words $\mathbf{V}^{\perp\perp}=\mathbf{V}$. The dimensions of $\mathbf{V}$ and $\mathbf{W}$ are right, and the whole space $\mathbf{R}^n$ is being decomposed into two perpendicular parts (Figure 3.3).

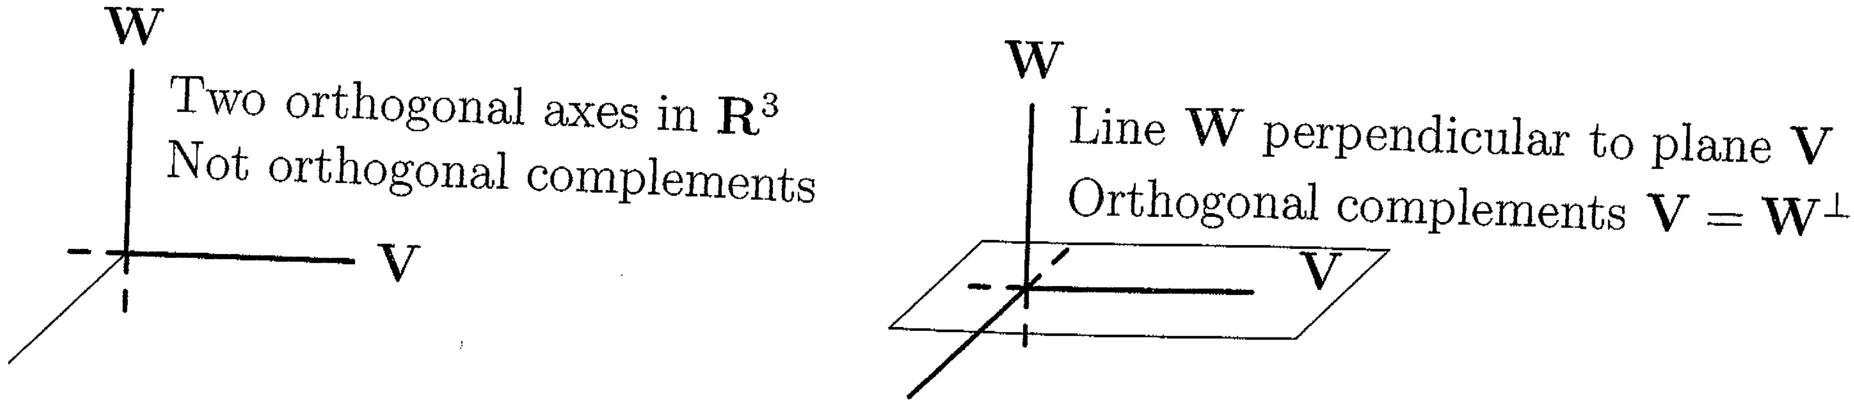

Figure 3.3 Orthogonal complements in $\mathbf{R}^3$: a plane and a line (not two lines).

Splitting $\mathbf{R}^n$ into orthogonal parts will split every vector into $x = v + w$. The vector v is the projection onto the subspace $\mathbf{V}$. The orthogonal component w is the projection of x onto $\mathbf{W}$. The next sections show how to find those projections of x. They lead to what is probably the most important figure in the book (Figure 3.4).

Figure 3.4 summarizes the fundamental theorem of linear algebra. It illustrates the true effect of a matrix—what is happening inside the multiplication Ax. The nullspace is carried to the zero vector. Every Ax is in the column space. Nothing is carried to the left nullspace. *The real action is between the row space and column space*, and you see it by looking at a typical vector x. It has a "row space component" and a "nullspace component," with $x = x_r + x_n$. When multiplied by A, this is $Ax = Ax_r + Ax_n$:

The nullspace component goes to zero: $Ax_n = 0$.
The row space component goes to the column space: $Ax_r = Ax$.

Of course everything goes to the column space—the matrix cannot do anything else. I tried to make the row and column spaces the same size, with equal dimension r.

3F From the row space to the column space, A is actually invertible. Every vector b in the column space comes from exactly one vector x_r in the row space.

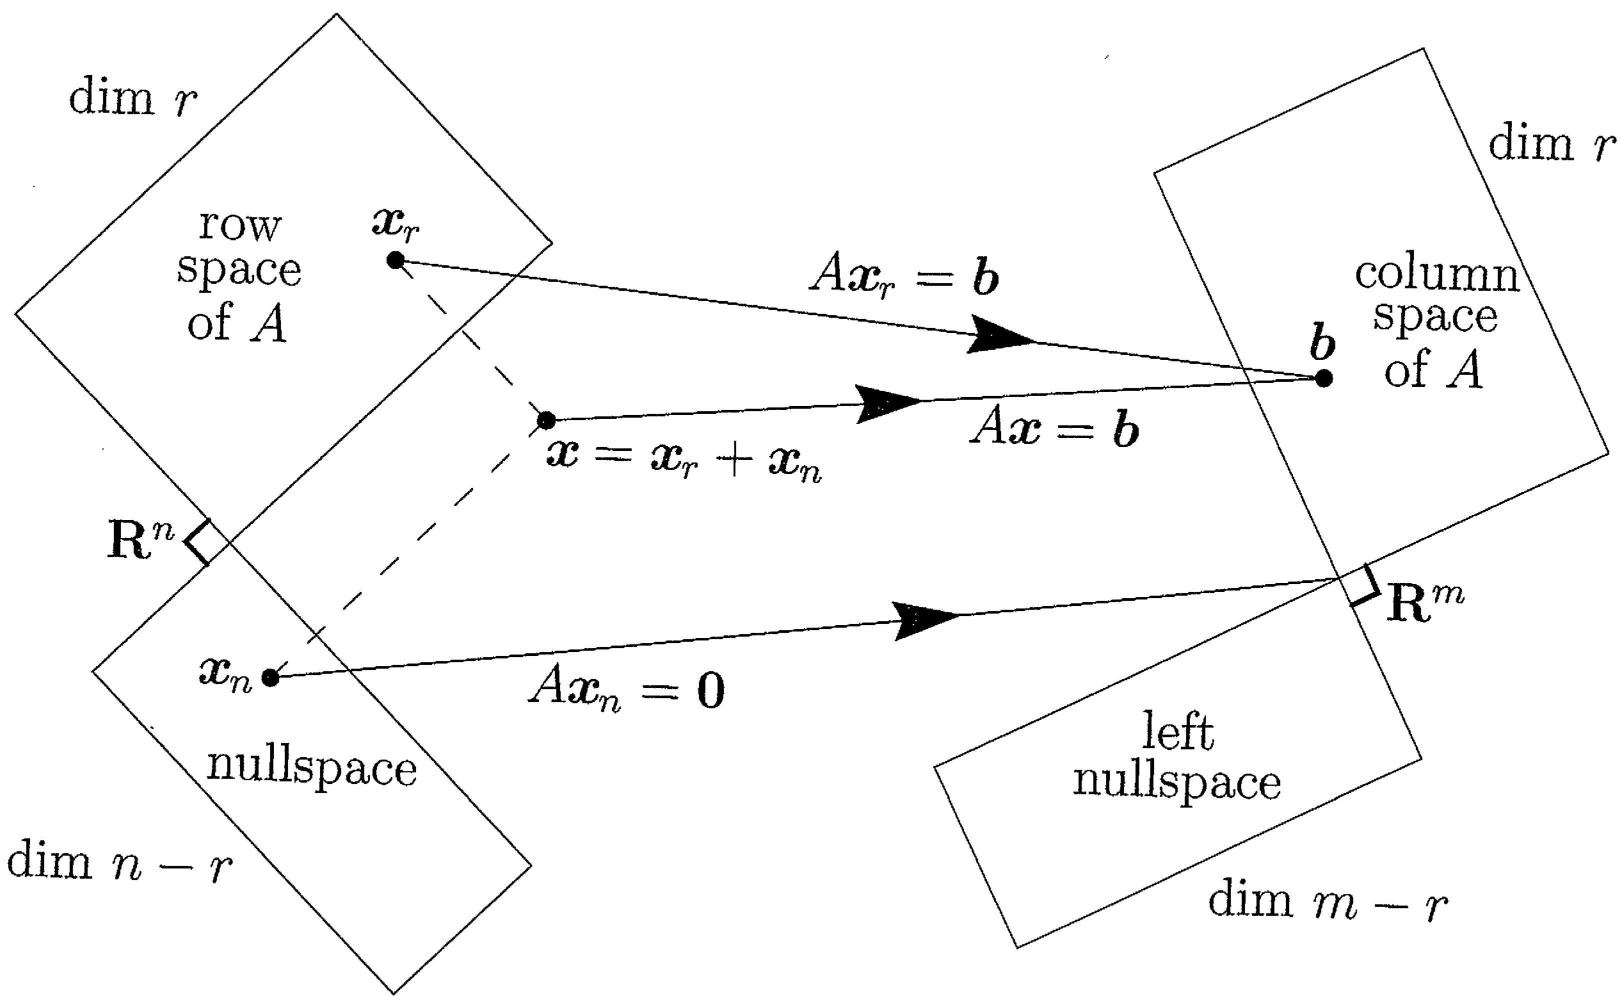

gure 3.4 The true action $Ax = A(x_{\mathbf{row}} + x_{\mathbf{null}})$ of any m by n matrix.

Proof Every b in the column space is a combination Ax of the columns. In fact, b is Ax_r, with x_r in the row space, since the nullspace component gives $Ax_n = 0$. If another vector x_r' in the row space gives $Ax_r' = b$, then $A(x_r - x_r') = b - b = 0$. This puts $x_r - x_r'$ in the nullspace and the row space, which makes it orthogonal to itself. Therefore it is zero, and $x_r = x_r'$. Exactly one vector in the row space is carried to b. ■

Every matrix transforms its row space onto its column space.

On those r-dimensional spaces A is invertible. On its nullspace A is zero. When A is diagonal, you see the invertible submatrix holding the r nonzeros.

A^T goes in the opposite direction, from $\mathbf{R}^m$ to $\mathbf{R}^n$ and from $C(A)$ back to $C(A^T)$. Of course the transpose is not the inverse! A^T moves the spaces correctly, but not the individual vectors. That honor belongs to A^{-1} if it exists—and it only exists if $r = m = n$. We cannot ask A^{-1} to bring back a whole nullspace out of the zero vector.

When A^{-1} fails to exist, the best substitute is the ***pseudoinverse*** A^+. This inverts A where that is possible: $A^+Ax = x$ for x in the row space. On the left nullspace, nothing can be done: $A^+y = 0$. Thus A^+ inverts A where it is invertible, and has the same rank r. One formula for A^+ depends on the ***singular value decomposition***—for which we first need to know about eigenvalues.

Problem Set 3.1

1. Find the lengths and the inner product of $x = (1, 4, 0, 2)$ and $y = (2, -2, 1, 3)$.

2. Give an example in $\mathbf{R}^2$ of linearly independent vectors that are not orthogonal. Also, give an example of orthogonal vectors that are not independent.

3. Two lines in the plane are perpendicular when the product of their slopes is -1. Apply this to the vectors $x = (x_1, x_2)$ and $y = (y_1, y_2)$, whose slopes are x_2/x_1 and y_2/y_1, to derive again the orthogonality condition $x^Ty = 0$.

4. How do we know that the ith row of an invertible matrix B is orthogonal to the jth column of B^{-1}, if $i \neq j$?

5. Which pairs are orthogonal among the vectors v_1, v_2, v_3, v_4?

$$v_1 = \begin{bmatrix} 1 \\ 2 \\ -2 \\ 1 \end{bmatrix}, \quad v_2 = \begin{bmatrix} 4 \\ 0 \\ 4 \\ 0 \end{bmatrix}, \quad v_3 = \begin{bmatrix} 1 \\ -1 \\ -1 \\ -1 \end{bmatrix}, \quad v_4 = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}.$$

6. Find all vectors in $\mathbf{R}^3$ that are orthogonal to $(1, 1, 1)$ and $(1, -1, 0)$. Produce an orthonormal basis from these vectors (mutually orthogonal unit vectors).

7. Find a vector x orthogonal to the row space of A, and a vector y orthogonal to the column space, and a vector z orthogonal to the nullspace:

$$A = \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 3 \\ 3 & 6 & 4 \end{bmatrix}.$$

8. If $\mathbf{V}$ and $\mathbf{W}$ are orthogonal subspaces, show that the only vector they have in common is the zero vector: $\mathbf{V} \cap \mathbf{W} = \{0\}$.

9. Find the orthogonal complement of the plane spanned by the vectors $(1, 1, 2)$ and $(1, 2, 3)$, by taking these to be the rows of A and solving $Ax = 0$. Remember that the complement is a whole line.

10. Construct a homogeneous equation in three unknowns whose solutions are the linear combinations of the vectors $(1, 1, 2)$ and $(1, 2, 3)$. This is the reverse of the previous exercise, but the two problems are really the same.

11. The fundamental theorem is often stated in the form of *Fredholm's alternative*: For any A and b, one and only one of the following systems has a solution:

(i) $Ax = b$.

(ii) $A^{\mathrm{T}}y = 0,\ y^{\mathrm{T}}b \neq 0$.

Either b is in the column space $\boldsymbol{C}(A)$ or there is a y in $\boldsymbol{N}(A^{\mathrm{T}})$ such that $y^{\mathrm{T}}b \neq 0$. Show that it is contradictory for (i) and (ii) both to have solutions.

12. Find a basis for the orthogonal complement of the row space of A:

$$A = \begin{bmatrix} 1 & 0 & 2 \\ 1 & 1 & 4 \end{bmatrix}.$$

Split $x = (3, 3, 3)$ into a row space component x_r and a nullspace component x_n.

13. Illustrate the action of A^{T} by a picture corresponding to Figure 3.4, sending $\boldsymbol{C}(A)$ back to the row space and the left nullspace to zero.

14. Show that $x - y$ is orthogonal to $x + y$ if and only if $\|x\| = \|y\|$.

15. Find a matrix whose row space contains $(1, 2, 1)$ and whose nullspace contains $(1, -2, 1)$, or prove that there is no such matrix.

16. Find all vectors that are perpendicular to $(1, 4, 4, 1)$ and $(2, 9, 8, 2)$.

17. If $\mathbf{V}$ is the orthogonal complement of $\mathbf{W}$ in $\mathbf{R}^n$, is there a matrix with row space $\mathbf{V}$ and nullspace $\mathbf{W}$? Starting with a basis for $\mathbf{V}$, construct such a matrix.

18. If $\mathbf{S} = \{0\}$ is the subspace of $\mathbf{R}^4$ containing only the zero vector, what is $\mathbf{S}^{\perp}$? If $\mathbf{S}$ is spanned by $(0, 0, 0, 1)$, what is $\mathbf{S}^{\perp}$? What is $(\mathbf{S}^{\perp})^{\perp}$?

19. *Why are these statements false?*

(a) If $\mathbf{V}$ is orthogonal to $\mathbf{W}$, then $\mathbf{V}^{\perp}$ is orthogonal to $\mathbf{W}^{\perp}$.

(b) $\mathbf{V}$ orthogonal to $\mathbf{W}$ and $\mathbf{W}$ orthogonal to $\mathbf{Z}$ makes $\mathbf{V}$ orthogonal to $\mathbf{Z}$.

20. Let $\mathbf{S}$ be a subspace of $\mathbf{R}^n$. Explain what $(\mathbf{S}^{\perp})^{\perp} = \mathbf{S}$ means and why it is true.

21. Let $\mathbf{P}$ be the plane in $\mathbf{R}^2$ with equation $x + 2y - z = 0$. Find a vector perpendicular to $\mathbf{P}$. What matrix has the plane $\mathbf{P}$ as its nullspace, and what matrix has $\mathbf{P}$ as its row space?

22. Let $\mathbf{S}$ be the subspace of $\mathbf{R}^4$ containing all vectors with $x_1 + x_2 + x_3 + x_4 = 0$. Find a basis for the space $\mathbf{S}^{\perp}$, containing all vectors orthogonal to $\mathbf{S}$.

23. Construct an unsymmetric 2 by 2 matrix of rank 1. Copy Figure 3.4 and put one vector in each subspace. Which vectors are orthogonal?

24. Redraw Figure 3.4 for a 3 by 2 matrix of rank $r = 2$. Which subspace is $\boldsymbol{Z}$ (zero vector only)? The nullspace part of any vector x in $\mathbf{R}^2$ is $x_n =$ ______.

25. Construct a matrix with the required property or say why that is impossible.

(a) Column space contains $\begin{bmatrix} 1 \\ 2 \\ -3 \end{bmatrix}$ and $\begin{bmatrix} 2 \\ -3 \\ 5 \end{bmatrix}$, nullspace contains $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

(b) Row space contains $\begin{bmatrix} 1 \\ 2 \\ -3 \end{bmatrix}$ and $\begin{bmatrix} 2 \\ -3 \\ 5 \end{bmatrix}$, nullspace contains $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

(c) $Ax = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$ has a solution and $A^{\mathrm{T}} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$.

(d) Every row is orthogonal to every column (A is not the zero matrix).

(e) The columns add up to a column of 0s, the rows add to a row of 1s.

26. If $AB = 0$ then the columns of B are in the ______ of A. The rows of A are in the ______ of B. Why can't A and B be 3 by 3 matrices of rank 2?

27. (a) If $Ax = b$ has a solution and $A^{\mathrm{T}}y = 0$, then y is perpendicular to ______.
(b) If $A^{\mathrm{T}}y = c$ has a solution and $Ax = 0$, then x is perpendicular to ______.

28. This is a system of equations $Ax = b$ with *no solution*:

$$\begin{aligned} x + 2y + 2z &= 5 \\ 2x + 2y + 3z &= 5 \\ 3x + 4y + 5z &= 9. \end{aligned}$$

Find numbers y_1, y_2, y_3 to multiply the equations so they add to $0 = 1$. You have found a vector y in which subspace? The inner product $y^{\mathrm{T}}b$ is 1.

29. In Figure 3.4, how do we know that Ax_r is equal to Ax? How do we know that this vector is in the column space? If $A = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$ and $x = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, what is x_r?

30. If Ax is in the nullspace of A^{T} then $Ax = 0$. Reason: Ax is also in the ______ of A and the spaces are ______. *Conclusion*: $A^{\mathrm{T}}A$ *has the same nullspace as* A.

31. Suppose A is a symmetric matrix ($A^{\mathrm{T}} = A$).

(a) Why is its column space perpendicular to its nullspace?

(b) If $Ax = 0$ and $Az = 5z$, which subspaces contain these "eigenvectors" x and z? **Symmetric matrices have perpendicular eigenvectors** (see Section 5.5).

32. (Recommended) Draw Figure 3.4 to show each subspace for

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 6 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 0 \\ 3 & 0 \end{bmatrix}.$$

33. Find the pieces x_r and x_n, and draw Figure 3.4 properly, if

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

Problems 34–44 are about orthogonal subspaces.

34. Put bases for the orthogonal subspaces $\mathbf{V}$ and $\mathbf{W}$ into the columns of matrices V and W. *Why does* $V^{\mathrm{T}}W =$ *zero matrix*? This matches $v^{\mathrm{T}}w = 0$ for vectors.

35. The floor and the wall are not orthogonal subspaces because they share a nonzero vector (along the line where they meet). Two planes in $\mathbf{R}^3$ cannot be orthogonal! Find a vector in both column spaces $\boldsymbol{C}(A)$ and $\boldsymbol{C}(B)$:

$$A = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 1 & 2 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 5 & 4 \\ 6 & 3 \\ 5 & 1 \end{bmatrix}.$$

This will be a vector Ax and also $B\widehat{x}$. Think 3 by 4 with the matrix $[A \;\; B]$.

36. Extend Problem 35 to a p-dimensional subspace $\mathbf{V}$ and a q-dimensional subspace $\mathbf{W}$ of $\mathbf{R}^n$. What inequality on $p + q$ guarantees that $\mathbf{V}$ intersects $\mathbf{W}$ in a nonzero vector? These subspaces cannot be orthogonal.

37. Prove that every y in $N(A^{\mathrm{T}})$ is perpendicular to every Ax in the column space, using the matrix shorthand of equation (8). Start from $A^{\mathrm{T}}y = 0$.

38. If $\mathbf{S}$ is the subspace of $\mathbf{R}^3$ containing only the zero vector, what is $\mathbf{S}^\perp$? If $\mathbf{S}$ is spanned by $(1, 1, 1)$, what is $\mathbf{S}^\perp$? If $\mathbf{S}$ is spanned by $(2, 0, 0)$ and $(0, 0, 3)$, what is $\mathbf{S}^\perp$?

39. Suppose $\mathbf{S}$ only contains $(1, 5, 1)$ and $(2, 2, 2)$ (not a subspace). Then $\mathbf{S}^\perp$ is the nullspace of the matrix $A =$ ______. $\mathbf{S}^\perp$ is a subspace even if $\mathbf{S}$ is not.

40. Suppose $\mathbf{L}$ is a one-dimensional subspace (a line) in $\mathbf{R}^3$. Its orthogonal complement $\mathbf{L}^\perp$ is the ______ perpendicular to $\mathbf{L}$. Then $(\mathbf{L}^\perp)^\perp$ is a ______ perpendicular to $\mathbf{L}^\perp$. In fact $(\mathbf{L}^\perp)^\perp$ is the same as ______.

41. Suppose $\mathbf{V}$ is the whole space $\mathbf{R}^4$. Then $\mathbf{V}^\perp$ contains only the vector ______. Then $(\mathbf{V}^\perp)^\perp$ is ______. So $(\mathbf{V}^\perp)^\perp$ is the same as ______.

42. Suppose $\mathbf{S}$ is spanned by the vectors $(1, 2, 2, 3)$ and $(1, 3, 3, 2)$. Find two vectors that span $\mathbf{S}^\perp$. This is the same as solving $Ax = 0$ for which A?

43. If $\mathbf{P}$ is the plane of vectors in $\mathbf{R}^4$ satisfying $x_1 + x_2 + x_3 + x_4 = 0$, write a basis for $\mathbf{P}^\perp$. Construct a matrix that has $\mathbf{P}$ as its nullspace.

44. If a subspace $\mathbf{S}$ is contained in a subspace $\mathbf{V}$, prove that $\mathbf{S}^\perp$ contains $\mathbf{V}^\perp$.

Problems 45–50 are about perpendicular columns and rows.

45. Suppose an n by n matrix is invertible: $AA^{-1} = I$. Then the first column of A^{-1} is orthogonal to the space spanned by which rows of A?

46. Find $A^{\mathrm{T}}A$ if the columns of A are unit vectors, all mutually perpendicular.

47. Construct a 3 by 3 matrix A with no zero entries whose columns are mutually perpendicular. Compute $A^{\mathrm{T}}A$. Why is it a diagonal matrix?

48. The lines $3x + y = b_1$ and $6x + 2y = b_2$ are ______. They are the same line if ______. In that case (b_1, b_2) is perpendicular to the vector ______. The nullspace of the matrix is the line $3x + y =$ ______. One particular vector in that nullspace is ______.

49. Why is each of these statements false?
(a) $(1, 1, 1)$ is perpendicular to $(1, 1, -2)$, so the planes $x + y + z = 0$ and $x + y - 2z = 0$ are orthogonal subspaces.
(b) The subspace spanned by $(1, 1, 0, 0, 0)$ and $(0, 0, 0, 1, 1)$ is the orthogonal complement of the subspace spanned by $(1, -1, 0, 0, 0)$ and $(2, -2, 3, 4, -4)$.
(c) Two subspaces that meet only in the zero vector are orthogonal.

50. Find a matrix with $v = (1, 2, 3)$ in the row space and column space. Find another matrix with v in the nullspace and column space. Which pairs of subspaces can v *not* be in?

51. Suppose A is 3 by 4, B is 4 by 5, and $AB = 0$. Prove $\operatorname{rank}(A) + \operatorname{rank}(B) \leq 4$.

52. The command $N = \mathsf{null}(A)$ will produce a basis for the nullspace of A. Then the command $B = \mathsf{null}(N')$ will produce a basis for the ______ of A.

3.2 COSINES AND PROJECTIONS ONTO LINES

Vectors with $x^{\mathrm{T}}y = 0$ are orthogonal. Now we allow inner products that are ***not zero***, and angles that are ***not right angles***. We want to connect inner products to angles, and also to transposes. In Chapter 1 the transpose was constructed by flipping over a matrix as if it were some kind of pancake. We have to do better than that.

One fact is unavoidable: *The orthogonal case is the most important.* Suppose we want to find the distance from a point b to the line in the direction of the vector a We are looking along that line for the point p closest to b. The key is in the geometry ***The line connecting b to p*** (the dotted line in Figure 3.5) ***is perpendicular to a***. *Thi: fact will allow us to find the projection p*. Even though a and b are not orthogonal, th‹ distance problem automatically brings in orthogonality.

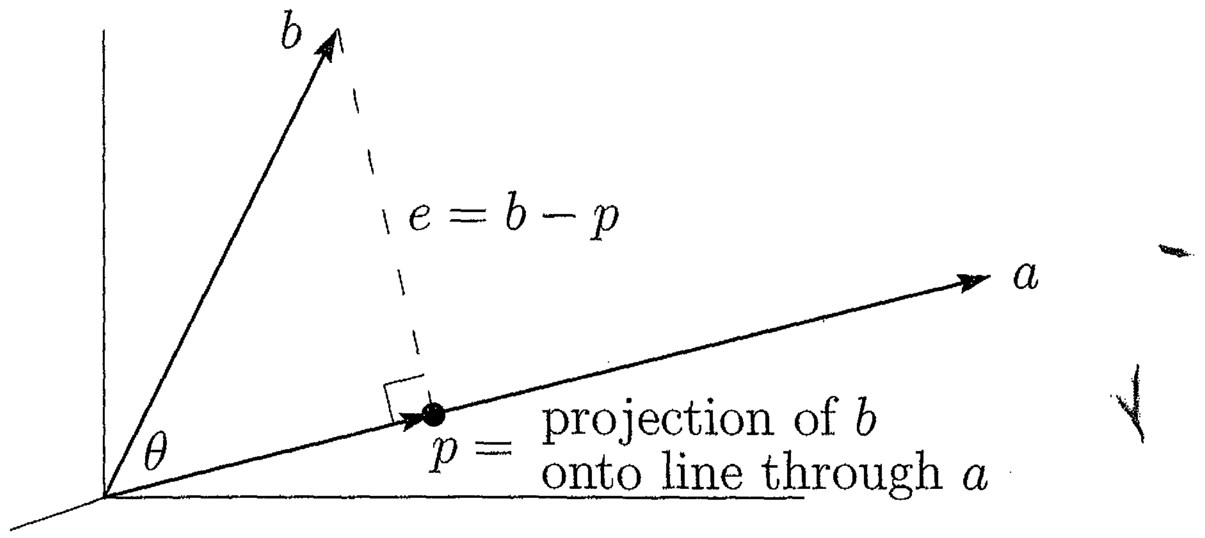

Figure 3.5 The projection p is the point (on the line through a) closest to b.

The situation is the same when we are given a plane (or any subspace $\mathbf{S}$) instea of a line. Again the problem is to find the point p on that subspace that is closest to ‹ ***This point p is the projection of b onto the subspace***. A perpendicular line from b · $\mathbf{S}$ meets the subspace at p. Geometrically, that gives the distance between points b ar subspaces **S**. But there are two questions that need to be asked:

1. Does this projection actually arise in practical applications?
2. If we have a basis for the subspace **S**, is there a formula for the projection p?

The answers are certainly yes. This is exactly the problem of the ***least-squares solution to an overdetermined system***. The vector b represents the data from experiments or questionnaires, and it contains too many errors to be found in the subspace **S**. When we try to write b as a combination of the basis vectors for **S**, it cannot be done—the equations are inconsistent, and $Ax = b$ has no solution.

The least-squares method selects p as the best choice to replace b. There can be no doubt of the importance of this application. In economics and statistics, least squares enters *regression analysis*. In geodesy, the U.S. mapping survey tackled 2.5 million equations in 400,000 unknowns.

A formula for p is easy when the subspace is a line. We will project b onto a in several different ways, and relate the projection p to inner products and angles. Projection onto a higher dimensional subspace is by far the most important case; it corresponds to a least-squares problem with several parameters, and it is solved in Section 3.3. The formulas are even simpler when we produce an orthogonal basis for **S**.

Inner Products and Cosines

We pick up the discussion of inner products and angles. You will soon see that it is not the angle, but ***the cosine of the angle***, that is directly related to inner products. We look back to trigonometry in the two-dimensional case to find that relationship. Suppose the vectors a and b make angles α and β with the x-axis (Figure 3.6).

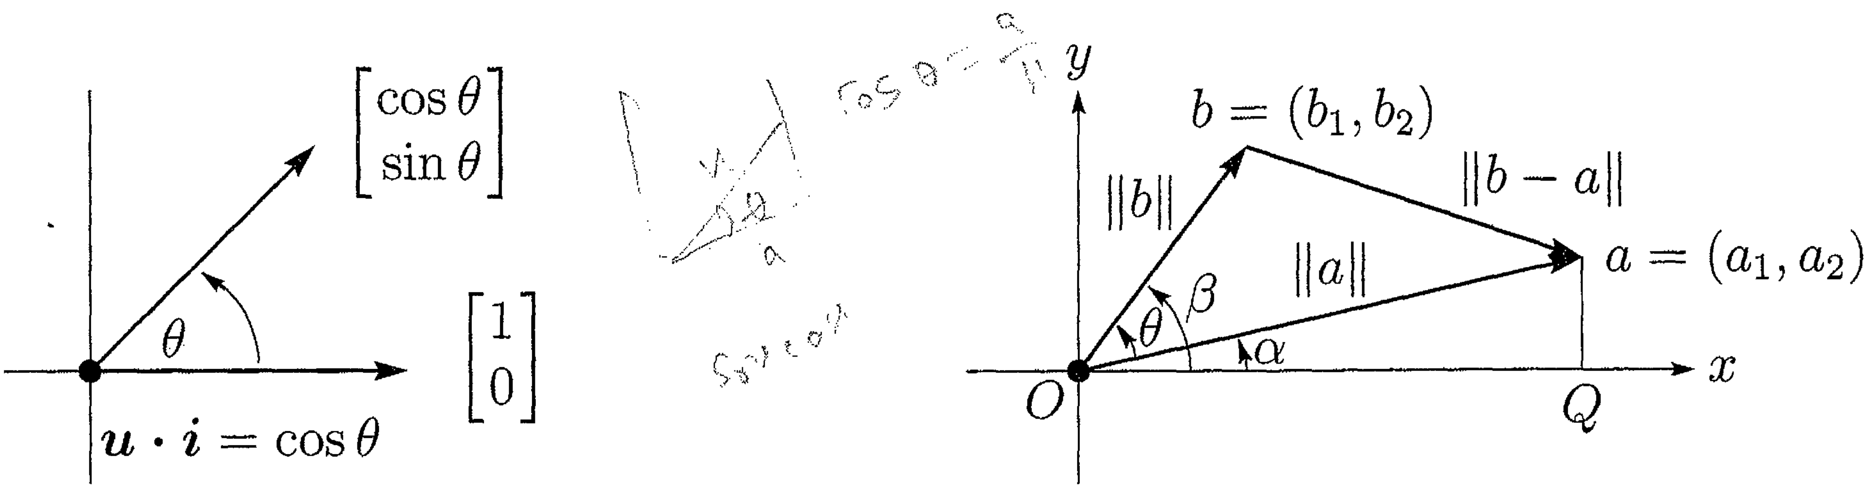

Figure 3.6 The cosine of the angle $\theta = \beta - \alpha$ using inner products.

The length $\|a\|$ is the hypotenuse in the triangle OaQ. So the sine and cosine of α are

$$\sin\alpha = \frac{a_2}{\|a\|}, \qquad \cos\alpha = \frac{a_1}{\|a\|}.$$

For the angle β, the sine is $b_2/\|b\|$ and the cosine is $b_1/\|b\|$. The cosine of $\theta = \beta - \alpha$ comes from an identity that no one could forget:

Cosine formula
$$\cos\theta = \cos\beta\cos\alpha + \sin\beta\sin\alpha = \frac{a_1b_1 + a_2b_2}{\|a\|\,\|b\|}. \qquad (1)$$

The numerator in this formula is exactly the inner product of a and b. It gives the relationship between $a^{\mathrm{T}}b$ and $\cos\theta$:

3G The cosine of the angle between any nonzero vectors a and b is

$$\textbf{Cosine of } \boldsymbol{\theta} \qquad \cos\theta = \frac{a^{\mathrm{T}}b}{\|a\|\,\|b\|}. \tag{2}$$

This formula is dimensionally correct; if we double the length of b, then both numerator and denominator are doubled, and the cosine is unchanged. Reversing the sign of b, on the other hand, reverses the sign of $\cos\theta$—and changes the angle by 180°.

There is another law of trigonometry that leads directly to the same result. It is not so unforgettable as the formula in equation (1), but it relates the lengths of the sides of any triangle:

$$\textbf{Law of Cosines} \qquad \|b-a\|^2 = \|b\|^2 + \|a\|^2 - 2\|b\|\,\|a\|\cos\theta. \tag{3}$$

When θ is a right angle, we are back to Pythagoras: $\|b-a\|^2 = \|b\|^2 + \|a\|^2$. For any angle θ, the expression $\|b-a\|^2$ is $(b-a)^{\mathrm{T}}(b-a)$, and equation (3) becomes

$$b^{\mathrm{T}}b - 2a^{\mathrm{T}}b + a^{\mathrm{T}}a = b^{\mathrm{T}}b + a^{\mathrm{T}}a - 2\|b\|\,\|a\|\cos\theta.$$

Canceling $b^{\mathrm{T}}b$ and $a^{\mathrm{T}}a$ on both sides of this equation, you recognize formula (2) for the cosine: $a^{\mathrm{T}}b = \|a\|\,\|b\|\cos\theta$. In fact, this proves the cosine formula in n dimensions, since we only have to worry about the plane triangle Oab.

Projection onto a Line

Now we want to find the projection point p. This point must be some multiple $p = \widehat{x}a$ of the given vector a—every point on the line is a multiple of a. The problem is to compute the coefficient $\widehat{x}$. All we need is the geometrical fact that ***the line from b to the closest point $p = \widehat{x}a$ is perpendicular to the vector a***:

$$(b - \widehat{x}a) \perp a, \quad \text{or} \quad a^{\mathrm{T}}(b - \widehat{x}a) = 0, \quad \text{or} \quad \widehat{x} = \frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}. \tag{4}$$

That gives the formula for the number $\widehat{x}$ and the projection p:

3H The projection of the vector b onto the line in the direction of a is $p = \widehat{x}a$:

$$\textbf{Projection onto a line} \qquad p = \widehat{x}a = \frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}a. \tag{5}$$

This allows us to redraw Figure 3.5 with a correct formula for p (Figure 3.7).

This leads to the Schwarz inequality in equation (6), which is the most important inequality in mathematics. A special case is the fact that arithmetic means $\frac{1}{2}(x+y)$ are larger than geometric means $\sqrt{xy}$. (It is also equivalent—see Problem 1 at the end of this section—to the triangle inequality for vectors.) The Schwarz inequality seems to come almost accidentally from the statement that $\|e\|^2 = \|b-p\|^2$ in Figure 3.7 cannot

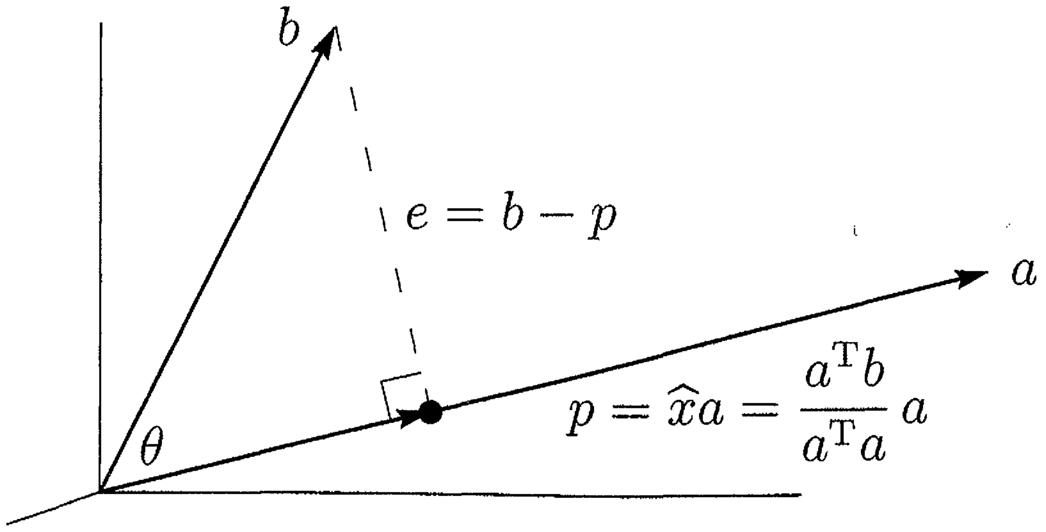

Figure 3.7 The projection p of b onto a, with $\cos\theta = \dfrac{Op}{Ob} = \dfrac{a^{\mathrm{T}}b}{\|a\|\,\|b\|}$.

be negative:

$$\left\| b - \frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}a \right\|^2 = b^{\mathrm{T}}b - 2\frac{(a^{\mathrm{T}}b)^2}{a^{\mathrm{T}}a} + \left(\frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}\right)^2 a^{\mathrm{T}}a = \frac{(b^{\mathrm{T}}b)(a^{\mathrm{T}}a) - (a^{\mathrm{T}}b)^2}{(a^{\mathrm{T}}a)} \geq 0.$$

This tells us that $(b^{\mathrm{T}}b)(a^{\mathrm{T}}a) \geq (a^{\mathrm{T}}b)^2$—and then we take square roots:

3I All vectors a and b satisfy the ***Schwarz inequality***, which is $|\cos\theta| \leq 1$ in $\mathbf{R}^n$:

$$\textbf{Schwarz inequality} \qquad |a^{\mathrm{T}}b| \leq \|a\|\,\|b\|. \tag{6}$$

According to formula (2), the ratio between $a^{\mathrm{T}}b$ and $\|a\|\,\|b\|$ is exactly $|\cos\theta|$. Since all cosines lie in the interval $-1 \leq \cos\theta \leq 1$, this gives another proof of equation (6): *the Schwarz inequality is the same as* $|\cos\theta| \leq 1$. In some ways that is a more easily understood proof, because cosines are so familiar. Either proof is all right in $\mathbf{R}^n$, but notice that ours came directly from the calculation of $\|b - p\|^2$. This stays nonnegative when we introduce new possibilities for the lengths and inner products. The name of Cauchy is also attached to this inequality $|a^{\mathrm{T}}b| \leq \|a\|\,\|b\|$, and the Russians refer to it as the Cauchy-Schwarz-Buniakowsky inequality! Mathematical historians seem to agree that Buniakowsky's claim is genuine.

One final observation about $|a^{\mathrm{T}}b| \leq \|a\|\,\|b\|$. *Equality holds if and only if b is a multiple of a.* The angle is $\theta = 0°$ or $\theta = 180°$ and the cosine is 1 or -1. In this case b is identical with its projection p, and the distance between b and the line is zero.

Project $b = (1, 2, 3)$ onto the line through $a = (1, 1, 1)$ to get $\widehat{x}$ and p:

$$\widehat{x} = \frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a} = \frac{6}{3} = 2.$$

The projection is $p = \widehat{x}a = (2, 2, 2)$. The angle between a and b has

$$\cos\theta = \frac{\|p\|}{\|b\|} = \frac{\sqrt{12}}{\sqrt{14}} \quad \text{and also} \quad \cos\theta = \frac{a^{\mathrm{T}}b}{\|a\|\,\|b\|} = \frac{6}{\sqrt{3}\sqrt{14}}.$$

The Schwarz inequality $|a^{\mathrm{T}}b| \leq \|a\|\,\|b\|$ is $6 \leq \sqrt{3}\sqrt{14}$. If we write 6 as $\sqrt{36}$, that is the same as $\sqrt{36} \leq \sqrt{42}$. The cosine is less than 1, because b is not parallel to a.

Projection Matrix of Rank 1

The projection of b onto the line through a lies at $p = a(a^{\mathrm{T}}b/a^{\mathrm{T}}a)$. That is our formula $p = \widehat{x}a$, but it is written with a slight twist: The vector a is put before the number $\widehat{x} = a^{\mathrm{T}}b/a^{\mathrm{T}}a$. There is a reason behind that apparently trivial change. Projection onto a line is carried out by a ***projection matrix*** P, and written in this new order we can see what it is. *P is the matrix that multiplies b and produces p*:

$$P = a\frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a} \quad \text{so the projection matrix is} \quad P = \frac{aa^{\mathrm{T}}}{a^{\mathrm{T}}a}. \tag{7}$$

That is a column times a row—a square matrix—divided by the number $a^{\mathrm{T}}a$.

The matrix that projects onto the line through $a = (1, 1, 1)$ is

$$P = \frac{aa^{\mathrm{T}}}{a^{\mathrm{T}}a} = \frac{1}{3}\begin{bmatrix}1\\1\\1\end{bmatrix}\begin{bmatrix}1 & 1 & 1\end{bmatrix} = \begin{bmatrix}\frac{1}{3} & \frac{1}{3} & \frac{1}{3}\\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3}\\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3}\end{bmatrix}.$$

This matrix has two properties that we will see as typical of projections:

1. **P is a symmetric matrix.**
2. **Its square is itself: $P^2 = P$.**

P^2b is the projection of Pb—and Pb is already on the line! So $P^2b = Pb$. This matrix P also gives a great example of the four fundamental subspaces:

The column space consists of the line through $a = (1, 1, 1)$.
The nullspace consists of the plane perpendicular to a.
The rank is $r = 1$.

Every column is a multiple of a, and so is $Pb = \widehat{x}a$. The vectors that project to $p = 0$ are especially important. They satisfy $a^{\mathrm{T}}b = 0$—they are perpendicular to a and their component along the line is zero. They lie in the nullspace = perpendicular plane.

Actually that example is too perfect. It has the nullspace orthogonal to the column space, which is haywire. The nullspace should be orthogonal to the *row space*. But because P is symmetric, its row and column spaces are the same.

Remark on scaling The projection matrix $aa^{\mathrm{T}}/a^{\mathrm{T}}a$ is the same if a is doubled:

$$a = \begin{bmatrix}2\\2\\2\end{bmatrix} \quad \text{gives} \quad P = \frac{1}{12}\begin{bmatrix}2\\2\\2\end{bmatrix}\begin{bmatrix}2 & 2 & 2\end{bmatrix} = \begin{bmatrix}\frac{1}{3} & \frac{1}{3} & \frac{1}{3}\\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3}\\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3}\end{bmatrix} \quad \text{as before.}$$

The line through a is the same, and that's all the projection matrix cares about. If a has unit length, the denominator is $a^{\mathrm{T}}a = 1$ and the matrix is just $P = aa^{\mathrm{T}}$.

Project onto the "θ-direction" in the x-y plane. The line goes through $a = (\cos\theta, \sin\theta)$ and the matrix is symmetric with $P^2 = P$:

$$P = \frac{aa^{\mathrm{T}}}{a^{\mathrm{T}}a} = \frac{\begin{bmatrix} c \\ s \end{bmatrix}\begin{bmatrix} c & s \end{bmatrix}}{\begin{bmatrix} c & s \end{bmatrix}\begin{bmatrix} c \\ s \end{bmatrix}} = \begin{bmatrix} c^2 & cs \\ cs & s^2 \end{bmatrix}.$$

Here c is $\cos\theta$, s is $\sin\theta$, and $c^2 + s^2 = 1$ in the denominator. This matrix P was discovered in Section 2.6 on linear transformations. Now we know P in any number of dimensions. We emphasize that it produces the projection p:

To project b onto a, multiply by the projection matrix P: $p = Pb$.

Transposes from Inner Products

Finally we connect inner products to A^{T}. Up to now, A^{T} is simply the reflection of A across its main diagonal; the rows of A become the columns of A^{T}, and vice versa. The entry in row i, column j of A^{T} is the (j, i) entry of A:

Transpose by reflection $\qquad (A^{\mathrm{T}})_{ij} = (A)_{ji}.$

There is a deeper significance to A^{T}. Its close connection to inner products gives a new and much more "abstract" definition of the transpose:

3J The transpose A^{T} can be defined by the following property: The inner product of Ax with y equals the inner product of x with $A^{\mathrm{T}}y$. Formally, this simply means that

$$(Ax)^{\mathrm{T}}y = x^{\mathrm{T}}A^{\mathrm{T}}y = x^{\mathrm{T}}(A^{\mathrm{T}}y). \tag{8}$$

This definition gives us another (better) way to verify the formula $(AB)^{\mathrm{T}} = B^{\mathrm{T}}A^{\mathrm{T}}$. Use equation (8) twice:

Move A then move B $\qquad (ABx)^{\mathrm{T}}y = (Bx)^{\mathrm{T}}(A^{\mathrm{T}}y) = x^{\mathrm{T}}(B^{\mathrm{T}}A^{\mathrm{T}}y).$

The transposes turn up in reverse order on the right side, just as the inverses do in the formula $(AB)^{-1} = B^{-1}A^{-1}$. We mention again that these two formulas meet to give the remarkable combination $(A^{-1})^{\mathrm{T}} = (A^{\mathrm{T}})^{-1}$.

Problem Set 3.2

1. (a) Given any two positive numbers x and y, choose the vector b equal to $(\sqrt{x}, \sqrt{y})$, and choose $a = (\sqrt{y}, \sqrt{x})$. Apply the Schwarz inequality to compare the arithmetic mean $\frac{1}{2}(x + y)$ with the geometric mean $\sqrt{xy}$.

(b) Suppose we start with a vector from the origin to the point x, and then add a vector of length $\|y\|$ connecting x to $x + y$. The third side of the triangle goes from the origin to $x + y$. *The triangle inequality asserts that this distance* *cannot be greater than the sum of the first two*:

$$\|x + y\| \le \|x\| + \|y\|.$$

After squaring both sides, and expanding $(x + y)^T(x + y)$, reduce this to the Schwarz inequality.

2. Verify that the length of the projection in Figure 3.7 is $\|p\| = \|b\| \cos\theta$, usin formula (5).

3. What multiple of $a = (1, 1, 1)$ is closest to the point $b = (2, 4, 4)$? Find also th point closest to a on the line through b.

4. Explain why the Schwarz inequality becomes an equality in the case that a and b lie on the same line through the origin, and only in that case. What if they lie o opposite sides of the origin?

5. In n dimensions, what angle does the vector $(1, 1, \ldots, 1)$ make with the coordinat axes? What is the projection matrix P onto that vector?

6. The Schwarz inequality has a one-line proof if a and b are normalized ahead of tim to be unit vectors:

$$|a^T b| = \left|\sum a_j b_j\right| \le \sum |a_j||b_j| \le \sum \frac{|a_j|^2 + |b_j|^2}{2} = \frac{1}{2} + \frac{1}{2} = \|a\|\,\|b\|.$$

Which previous problem justifies the middle step?

7. By choosing the correct vector b in the Schwarz inequality, prove that

$$(a_1 + \cdots + a_n)^2 \le n\left(a_1^2 + \cdots + a_n^2\right).$$

When does equality hold?

8. The methane molecule CH_4 is arranged as if the carbon atom were at the center of regular tetrahedron with four hydrogen atoms at the vertices. If vertices are place at $(0, 0, 0)$, $(1, 1, 0)$, $(1, 0, 1)$, and $(0, 1, 1)$—note that all six edges have length $\sqrt{}$ so the tetrahedron is regular—what is the cosine of the angle between the rays goi from the center $\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)$ to the vertices? (The bond angle itself is about 109.5°, old friend of chemists.)

9. Square the matrix $P = aa^T/a^Ta$, which projects onto a line, and show that $P^2 =$ (Note the number a^Ta in the middle of the matrix aa^Taa^T!)

10. Is the projection matrix P invertible? Why or why not?

11. (a) Find the projection matrix P_1 onto the line through $a = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$ and also the matr P_2 that projects onto the line perpendicular to a.
    (b) Compute $P_1 + P_2$ and P_1P_2 and explain.

12. Find the matrix that projects every point in the plane onto the line $x + 2y = 0$.

13. Prove that the *trace* of $P = aa^T/a^Ta$—which is the sum of its diagonal entries always equals 1.

14. What matrix P projects every point in $\mathbf{R}^3$ onto the line of intersection of the planes $x + y + t = 0$ and $x - t = 0$?

15. Show that the length of Ax equals the length of $A^{\mathrm{T}}x$ if $AA^{\mathrm{T}} = A^{\mathrm{T}}A$.

16. Suppose P is the projection matrix onto the line through a.

(a) Why is the inner product of x with Py equal to the inner product of Px with y?

(b) Are the two angles the same? Find their cosines if $a = (1, 1, -1)$, $x = (2, 0, 1)$, $y = (2, 1, 2)$.

(c) Why is the inner product of Px with Py again the same? What is the angle between those two?

Problems 17–26 ask for projections onto lines. Also errors $e = b - p$ and matrices P.

17. Project the vector b onto the line through a. Check that e is perpendicular to a:

$$\text{(a)}\quad b = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix} \quad \text{and} \quad a = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}. \qquad \text{(b)}\quad b = \begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix} \quad \text{and} \quad a = \begin{bmatrix} -1 \\ -3 \\ -1 \end{bmatrix}.$$

18. *Draw* the projection of b onto a and also compute it from $p = \widehat{x}a$:

$$\text{(a)}\quad b = \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix} \quad \text{and} \quad a = \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \qquad \text{(b)}\quad b = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \text{and} \quad a = \begin{bmatrix} 1 \\ -1 \end{bmatrix}.$$

19. In Problem 17, find the projection matrix $P = aa^{\mathrm{T}}/a^{\mathrm{T}}a$ onto the line through each vector a. Verify in both cases that $P^2 = P$. Multiply Pb in each case to compute the projection p.

20. Construct the projection matrices P_1 and P_2 onto the lines through the a's in Problem 18. Is it true that $(P_1 + P_2)^2 = P_1 + P_2$? This *would* be true if $P_1P_2 = 0$.

For Problems 21–26, consult the accompanying figures.

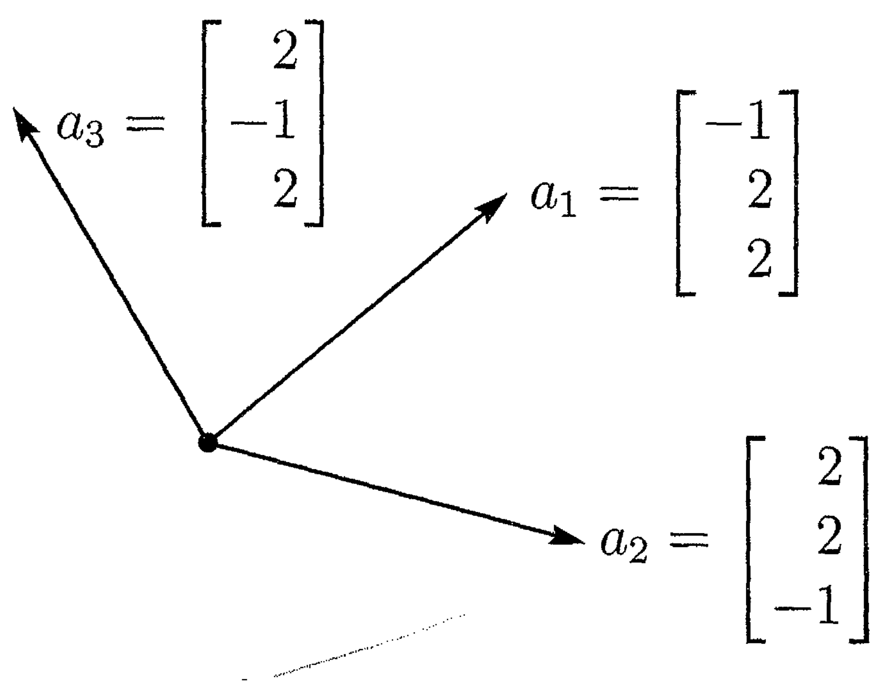

Problems 21–23

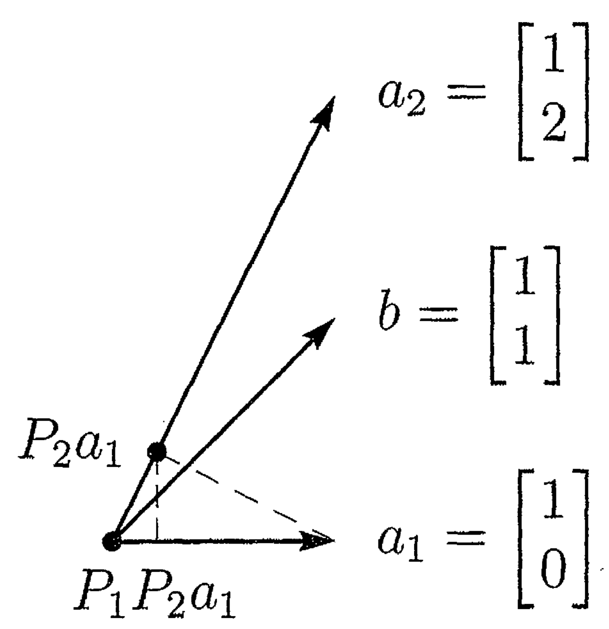

Problems 24–26

21. Compute the projection matrices $aa^{\mathrm{T}}/a^{\mathrm{T}}a$ onto the lines through $a_1 = (-1, 2, 2)$ and $a_2 = (2, 2, -1)$. Multiply those projection matrices and explain why their product P_1P_2 is what it is.

22. Project $b = (1, 0, 0)$ onto the lines through a_1 and a_2 in Problem 21 and also onto $a_3 = (2, -1, 2)$. Add the three projections $p_1 + p_2 + p_3$.

23. Continuing Problems 21–22, find the projection matrix P_3 onto $a_3 = (2, -1, 2)$. Verify that $P_1 + P_2 + P_3 = I$. The basis a_1, a_2, a_3 is orthogonal!

24. Project the vector $b = (1, 1)$ onto the lines through $a_1 = (1, 0)$ and $a_2 = (1, 2)$. Draw the projections p_1 and p_2 and add $p_1 + p_2$. The projections do not add to b because the a's are not orthogonal.

25. In Problem 24, the projection of b onto the *plane* of a_1 and a_2 will equal b. Find $P = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$ for $A = \begin{bmatrix} a_1 & a_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$.

26. Project $a_1 = (1, 0)$ onto $a_2 = (1, 2)$. Then project the result back onto a_1. Draw these projections and multiply the projection matrices $P_1 P_2$: Is this a projection?

3.3 PROJECTIONS AND LEAST SQUARES

Up to this point, $Ax = b$ either has a solution or not. If b is not in the column space $\boldsymbol{C}(A)$, the system is inconsistent and Gaussian elimination fails. This failure is almost certain when there are several equations and only one unknown:

$$\begin{array}{ll} \textbf{More equations} & 2x = b_1 \\ \textbf{than unknowns—} & 3x = b_2 \\ \textbf{no solution?} & 4x = b_3. \end{array}$$

This is solvable when b_1, b_2, b_3 are in the ratio 2:3:4. The solution x will exist only if b is on the same line as the column $a = (2, 3, 4)$.

In spite of their unsolvability, inconsistent equations arise all the time in practice. They have to be solved! One possibility is to determine x from part of the system, and ignore the rest; this is hard to justify if all m equations come from the same source. Rather than expecting no error in some equations and large errors in the others, it is much better *to choose the x that minimizes an average error E in the m equations.*

The most convenient "average" comes from the *sum of squares*:

$$\textbf{Squared error} \qquad E^2 = (2x - b_1)^2 + (3x - b_2)^2 + (4x - b_3)^2.$$

If there is an exact solution, the minimum error is $E = 0$. In the more likely case that b is not proportional to a, the graph of E^2 will be a parabola. The minimum error is at the lowest point, where the derivative is zero: √

$$\frac{dE^2}{dx} = 2[(2x - b_1)2 + (3x - b_2)3 + (4x - b_3)4] = 0.$$

Solving for x, the least-squares solution of this model system $ax = b$ is denoted by $\widehat{x}$:

$$\textbf{Least-squares solution} \qquad \widehat{x} = \frac{2b_1 + 3b_2 + 4b_3}{2^2 + 3^2 + 4^2} = \frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}.$$

You recognize $a^{\mathrm{T}}b$ in the numerator and $a^{\mathrm{T}}a$ in the denominator.

The general case is the same. We "solve" $ax = b$ by minimizing

$$E^2 = \|ax - b\|^2 = (a_1x - b_1)^2 + \cdots + (a_mx - b_m)^2.$$

The derivative of E^2 is zero at the point $\widehat{x}$, if

$$(a_1\widehat{x} - b_1)a_1 + \cdots + (a_m\widehat{x} - b_m)a_m = 0.$$

We are minimizing the distance from b to the line through a, and calculus gives the same answer, $\widehat{x} = (a_1b_1 + \cdots + a_mb_m)/(a_1^2 + \cdots + a_m^2)$, that geometry did earlier:

3K The least-squares solution to a problem $ax = b$ in one unknown is $\widehat{x} = \dfrac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}$.

You see that we keep coming back to the geometrical interpretation of a least-squares problem—to minimize a distance. By setting the derivative of E^2 to zero, calculus confirms the geometry of the previous section. *The error vector e connecting b to p must be perpendicular to a*:

$$\textbf{Orthogonality of } \boldsymbol{a} \textbf{ and } \boldsymbol{e} \qquad a^{\mathrm{T}}(b - \widehat{x}a) = a^{\mathrm{T}}b - \frac{a^{\mathrm{T}}b}{a^{\mathrm{T}}a}a^{\mathrm{T}}a = 0.$$

As a side remark, notice the degenerate case $a = 0$. All multiples of a are zero, and the line is only a point. Therefore $p = 0$ is the only candidate for the projection. But the formula for $\widehat{x}$ becomes a meaningless $0/0$, and correctly reflects the fact that $\widehat{x}$ is completely undetermined. All values of x give the same error $E = \|0x - b\|$, so E^2 is a horizontal line instead of a parabola. The "pseudoinverse" assigns the definite value $\widehat{x} = 0$, which is a more "symmetric" choice than any other number.

Least-Squares Problems with Several Variables

Now we are ready for the serious step, ***to project b onto a subspace***—rather than just onto a line. This problem arises from $Ax = b$ when A is an m by n matrix. Instead of one column and one unknown x, the matrix now has n columns. The number m of observations is still larger than the number n of unknowns, so it must be expected that $Ax = b$ will be inconsistent. *Probably, there will not exist a choice of x that perfectly fits the data b.* In other words, the vector b probably will not be a combination of the columns of A; it will be outside the column space.

Again the problem is to choose $\widehat{x}$ so as to minimize the error, and again this minimization will be done in the least-squares sense. The error is $E = \|Ax - b\|$, and ***this is exactly the distance from b to the point Ax in the column space***. Searching for the least-squares solution $\widehat{x}$, which minimizes E, is the same as locating the point $p = A\widehat{x}$ that is closer to b than any other point in the column space.

We may use geometry or calculus to determine $\widehat{x}$. In n dimensions, we prefer the appeal of geometry; p must be the "projection of b onto the column space." ***The error vector $e = b - A\widehat{x}$ must be perpendicular to that space*** (Figure 3.8). Finding $\widehat{x}$ and the projection $p = A\widehat{x}$ is so fundamental that we do it in two ways:

1. All vectors perpendicular to the column space lie in the *left nullspace*. Thus the error vector $e = b - A\widehat{x}$ must be in the nullspace of A^{T}:

$$A^{\mathrm{T}}(b - A\widehat{x}) = 0 \qquad \text{or} \qquad \boldsymbol{A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b}.$$

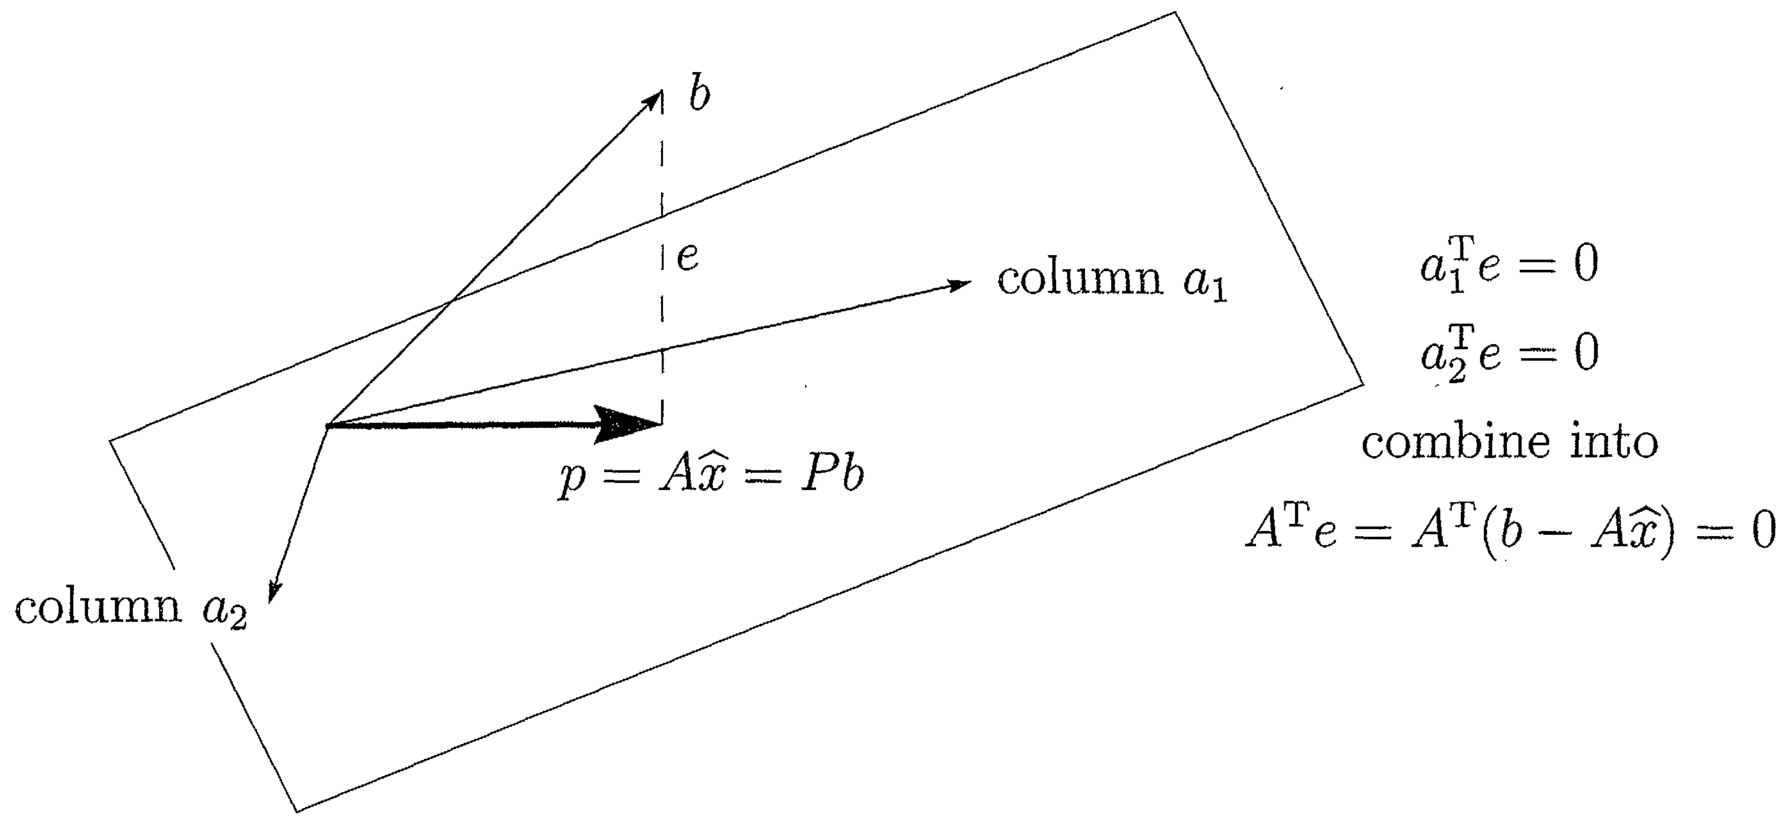

Figure 3.8 Projection onto the column space of a 3 by 2 matrix.

2. The error vector must be perpendicular to *each column* $a_1, \ldots, a_n$ of A:

$$\begin{matrix} a_1^{\mathrm{T}}(b - A\hat{x}) = 0 \\ \vdots \\ a_n^{\mathrm{T}}(b - A\hat{x}) = 0 \end{matrix} \quad \text{or} \quad \begin{bmatrix} & a_1^{\mathrm{T}} & \\ & \vdots & \\ & a_n^{\mathrm{T}} & \end{bmatrix} \begin{bmatrix} \\ b - A\hat{x} \\ \\ \end{bmatrix} = 0.$$

This is again $A^{\mathrm{T}}(b - A\hat{x}) = 0$ and $A^{\mathrm{T}}A\hat{x} = A^{\mathrm{T}}b$. The calculus way is to take partia derivatives of $E^2 = (Ax - b)^{\mathrm{T}}(Ax - b)$. That gives the same $2A^{\mathrm{T}}Ax - 2A^{\mathrm{T}}b = 0$ The fastest way is just ***to multiply the unsolvable equation $Ax = b$ by A^{T}***. Al these equivalent methods produce a square coefficient matrix $A^{\mathrm{T}}A$. It is symmetri (its transpose is not AA^{T}!) and it is the fundamental matrix of this chapter.

The equations $A^{\mathrm{T}}A\hat{x} = A^{\mathrm{T}}b$ are known in statistics as the ***normal equations***.

3L When $Ax = b$ is inconsistent, its least-squares solution minimizes $\|Ax - b\|^2$:

$$\textbf{Normal equations} \qquad A^{\mathrm{T}}A\hat{x} = A^{\mathrm{T}}b. \tag{1}$$

$A^{\mathrm{T}}A$ is invertible exactly when the columns of A are linearly independent! Then,

$$\textbf{Best estimate } \hat{x} \qquad \hat{x} = (A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}b. \tag{2}$$

The projection of b onto the column space is the nearest point $A\hat{x}$:

$$\textbf{Projection} \qquad p = A\hat{x} = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}b. \tag{3}$$

We choose an example in which our intuition is as good as the formulas:

$$A = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 0 & 0 \end{bmatrix}, \qquad b = \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix}, \qquad \begin{matrix} Ax = b & \text{has no solution} \\ A^{\mathrm{T}}A\hat{x} = A^{\mathrm{T}}b & \text{gives the best } x. \end{matrix}$$

Both columns end with a zero, so $\boldsymbol{C}(A)$ is the x-y plane within three-dimensional space. The projection of $b = (4, 5, 6)$ is $p = (4, 5, 0)$—the x and y components stay the same but $z = 6$ will disappear. That is confirmed by solving the normal equations:

$$A^{\mathrm{T}}A = \begin{bmatrix} 1 & 1 & 0 \\ 2 & 3 & 0 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 2 & 5 \\ 5 & 13 \end{bmatrix}.$$

$$\widehat{x} = (A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}b = \begin{bmatrix} 13 & -5 \\ -5 & 2 \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 \\ 2 & 3 & 0 \end{bmatrix} \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}.$$

$$\textbf{Projection} \qquad p = A\widehat{x} = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 4 \\ 5 \\ 0 \end{bmatrix}.$$

In this special case, the best we can do is to solve the first two equations of $Ax = b$. Then $\widehat{x}_1 = 2$ and $\widehat{x}_2 = 1$. The error in the equation $0x_1 + 0x_2 = 6$ is sure to be 6.

Remark 1 Suppose b is actually *in* the column space of A—it is a combination $b = Ax$ of the columns. Then the projection of b is still b:

$$\textbf{\textit{b} in column space} \qquad p = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}Ax = Ax = b.$$

The closest point p is just b itself—which is obvious.

Remark 2 At the other extreme, suppose b is *perpendicular* to every column, so $A^{\mathrm{T}}b = 0$. In this case b projects to the zero vector:

$$\textbf{\textit{b} in left nullspace} \qquad p = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}b = A(A^{\mathrm{T}}A)^{-1}0 = 0.$$

Remark 3 When A is square and invertible, the column space is the whole space. Every vector projects to itself, p equals b, and $\widehat{x} = x$:

$$\textbf{If \textit{A} is invertible} \qquad p = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}b = AA^{-1}(A^{\mathrm{T}})^{-1}A^{\mathrm{T}}b = b.$$

This is the only case when we can take apart $(A^{\mathrm{T}}A)^{-1}$, *and write it as* $A^{-1}(A^{\mathrm{T}})^{-1}$. When A is rectangular that is not possible.

Remark 4 Suppose A has only one column, containing a. Then the matrix $A^{\mathrm{T}}A$ is the number $a^{\mathrm{T}}a$ and $\widehat{x}$ is $a^{\mathrm{T}}b/a^{\mathrm{T}}a$. We return to the earlier formula.

The Cross-Product Matrix $A^{\mathrm{T}}A$

The matrix $A^{\mathrm{T}}A$ is certainly symmetric. Its transpose is $(A^{\mathrm{T}}A)^{\mathrm{T}} = A^{\mathrm{T}}A^{\mathrm{TT}}$, which is $A^{\mathrm{T}}A$ again. Its i, j entry (and j, i entry) is the inner product of column i of A with column j of A. The key question is the invertibility of $A^{\mathrm{T}}A$, and fortunately

$A^{\mathrm{T}}A$ has the same nullspace as A.

Certainly if $Ax = 0$ then $A^{\mathrm{T}}Ax = 0$. Vectors x in the nullspace of A are also in the nullspace of $A^{\mathrm{T}}A$. To go in the other direction, start by supposing that $A^{\mathrm{T}}Ax = 0$, and take the inner product with x to show that $Ax = 0$:

$$x^{\mathrm{T}} A^{\mathrm{T}} Ax = 0, \quad \text{or} \quad \|Ax\|^2 = 0, \quad \text{or} \quad Ax = 0.$$

The two nullspaces are identical. In particular, if A has independent columns (and only $x = 0$ is in its nullspace), then the same is true for $A^{\mathrm{T}}A$:

3M If A has independent columns, then $A^{\mathrm{T}}A$ is *square, symmetric*, and *invertible*.

We show later that $A^{\mathrm{T}}A$ is also positive definite (all pivots and eigenvalues are positive).

This case is by far the most common and most important. Independence is not so hard in m-dimensional space if $m > n$. We assume it in what follows.

Projection Matrices

We have shown that the closest point to b is $p = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}b$. *This formula expresses in matrix terms the construction of a perpendicular line from b to the column space of* A. The matrix that gives p is a projection matrix, denoted by P:

$$\textbf{Projection matrix} \qquad P = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}. \tag{4}$$

This matrix projects any vector b onto the column space of A.* In other words, $p = Pb$ is the component of b in the column space, and the error $e = b - Pb$ is the component in the orthogonal complement. ($I - P$ is also a projection matrix! It projects b onto the orthogonal complement, and the projection is $b - Pb$.)

In short, we have a matrix formula for splitting any b into two perpendicular components. Pb is in the column space $\boldsymbol{C}(A)$, and the other component $(I - P)b$ is in the left nullspace $\boldsymbol{N}(A^{\mathrm{T}})$—which is orthogonal to the column space.

These projection matrices can be understood geometrically and algebraically.

3N The projection matrix $P = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$ has two basic properties:

(i) It equals its square: $P^2 = P$.

(ii) It equals its transpose: $P^{\mathrm{T}} = P$.

Conversely, any symmetric matrix with $P^2 = P$ represents a projection.

Proof It is easy to see why $P^2 = P$. If we start with any b, then Pb lies in the subspace we are projecting onto. ***When we project again nothing is changed.*** The vector Pb is already in the subspace, and $P(Pb)$ is still Pb. In other words $P^2 = P$. Two or three or fifty projections give the same point p as the first projection:

$$P^2 = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}} = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}} = P.$$

---

* There may be a risk of confusion with permutation matrices, also denoted by P, but the risk should be small, and we try never to let both appear on the same page.

To prove that P is also symmetric, take its transpose. Multiply the transposes in reverse order, and use symmetry of $(A^{\mathrm{T}}A)^{-1}$, to come back to P:

$$P^{\mathrm{T}} = (A^{\mathrm{T}})^{\mathrm{T}}((A^{\mathrm{T}}A)^{-1})^{\mathrm{T}}A^{\mathrm{T}} = A((A^{\mathrm{T}}A)^{\mathrm{T}})^{-1}A^{\mathrm{T}} = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}} = P.$$

For the converse, we have to deduce from $P^2 = P$ and $P^{\mathrm{T}} = P$ that Pb **is the projection of b onto the column space of P**. The error vector $b - Pb$ is *orthogonal to the space*. For any vector Pc in the space, the inner product is zero:

$$(b - Pb)^{\mathrm{T}}Pc = b^{\mathrm{T}}(I - P)^{\mathrm{T}}Pc = b^{\mathrm{T}}(P - P^2)c = 0.$$

Thus $b - Pb$ is orthogonal to the space, and Pb is the projection onto the column space. ∎

Suppose A is actually invertible. If it is 4 by 4, then its four columns are independent and its column space is all of $\mathbf{R}^4$. What is the projection *onto the whole space*? It is the identity matrix.

$$P = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}} = AA^{-1}(A^{\mathrm{T}})^{-1}A^{\mathrm{T}} = I. \tag{5}$$

The identity matrix is symmetric, $I^2 = I$, and the error $b - Ib$ is zero.

The point of all other examples is that what happened in equation (5) is *not allowed.* To repeat: We cannot invert the separate parts A^{T} and A when those matrices are rectangular. It is the square matrix $A^{\mathrm{T}}A$ that is invertible.

Least-Squares Fitting of Data

Suppose we do a series of experiments, and expect the output b to be a linear function of the input t. We look for a ***straight line*** $b = C + Dt$. For example:

1. At different times we measure the distance to a satellite on its way to Mars. In this case t is the time and b is the distance. Unless the motor was left on or gravity is strong, the satellite should move with nearly constant velocity v: $b = b_0 + vt$.
2. We vary the load on a structure, and measure the movement it produces. In this experiment t is the load and b is the reading from the strain gauge. Unless the load is so great that the material becomes plastic, a linear relation $b = C + Dt$ is normal in the theory of elasticity.
3. The cost of producing t books like this one is nearly linear, $b = C + Dt$, with editing and typesetting in C and then printing and binding in D. C is the set-up cost and D is the cost for each additional book.

How to compute C and D? If there is no experimental error, then two measurements of b will determine the line $b = C + Dt$. But if there is error, we must be prepared to "average" the experiments and find an optimal line. That line is not to be confused with the line through a on which b was projected in the previous section! In fact, since there are two unknowns C and D to be determined, we now project onto a *two-dimensional* subspace. A perfect experiment would give a perfect C and D:

$$\begin{aligned} C + Dt_1 &= b_1 \\ C + Dt_2 &= b_2 \\ &\vdots \\ C + Dt_m &= b_m. \end{aligned} \tag{6}$$

This is an *overdetermined* system, with m equations and only two unknowns. If errors are present, it will have no solution. A has two columns, and $x = (C, D)$:

$$\begin{bmatrix} 1 & t_1 \\ 1 & t_2 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \begin{bmatrix} C \\ D \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix}, \qquad \text{or} \qquad Ax = b. \tag{7}$$

The best solution $(\widehat{C}, \widehat{D})$ is the $\widehat{x}$ that minimizes the squared error E^2:

$$\textbf{Minimize} \qquad E^2 = \|b - Ax\|^2 = (b_1 - C - Dt_1)^2 + \cdots + (b_m - C - Dt_m)^2.$$

The vector $p = A\widehat{x}$ is as close as possible to b. Of all straight lines $b = C + Dt$, we are choosing the one that best fits the data (Figure 3.9). On the graph, the errors are the ***vertical distances*** $b - C - Dt$ to the straight line (not perpendicular distances!). It is the vertical distances that are squared, summed, and minimized.

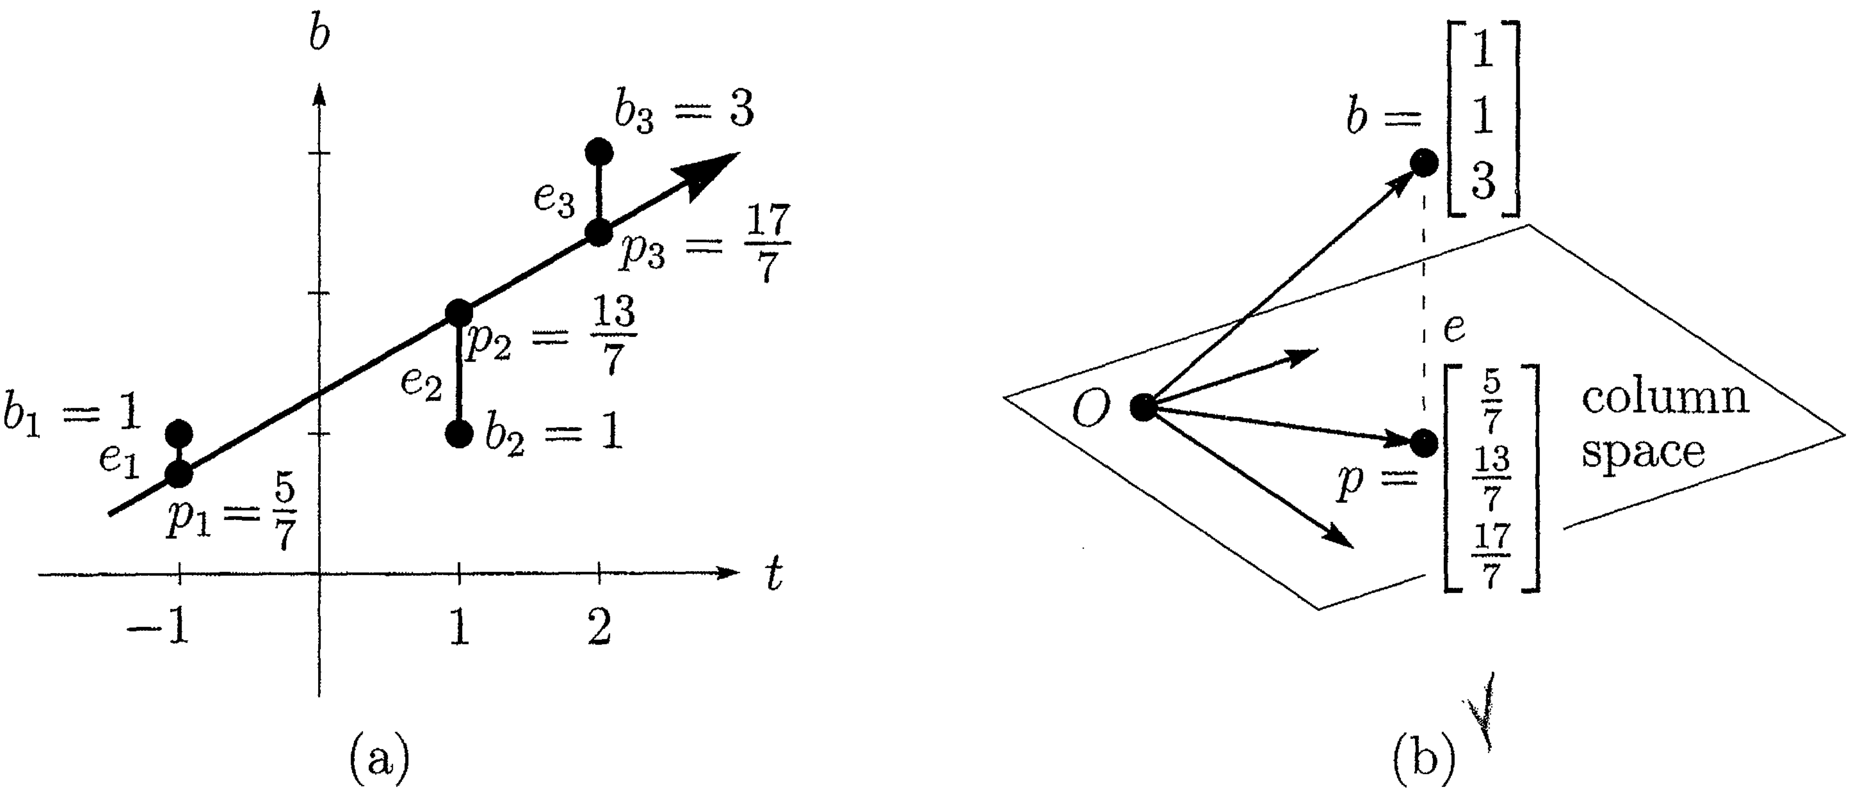

Figure 3.9 Straight-line approximation matches the projection p of b.

Three measurements b_1, b_2, b_3 are marked on Figure 3.9a:

$$b = 1 \quad \text{at} \quad t = -1, \qquad b = 1 \quad \text{at} \quad t = 1, \qquad b = 3 \quad \text{at} \quad t = 2.$$

Note that the values $t = -1, 1, 2$ are not required to be equally spaced. The first step is *to write the equations that* **would** *hold if a line could go through all three points.*

Then every $C + Dt$ would agree exactly with b:

$$Ax = b \quad \text{is} \quad \begin{aligned} C - \ \ D &= 1 \\ C + \ \ D &= 1 \\ C + 2D &= 3 \end{aligned} \quad \text{or} \quad \begin{bmatrix} 1 & -1 \\ 1 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} C \\ D \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 3 \end{bmatrix}.$$

If those equations $Ax = b$ could be solved, there would be no errors. They can't be solved because the points are not on a line. Therefore they are solved by least squares:

$$A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b \quad \text{is} \quad \begin{bmatrix} 3 & 2 \\ 2 & 6 \end{bmatrix} \begin{bmatrix} \widehat{C} \\ \widehat{D} \end{bmatrix} = \begin{bmatrix} 5 \\ 6 \end{bmatrix}.$$

The best solution is $\widehat{C} = \frac{9}{7}$, $\widehat{D} = \frac{4}{7}$ and the best line is $\frac{9}{7} + \frac{4}{7}t$.

Note the beautiful connections between the two figures. The problem is the same but the art shows it differently. In Figure 3.9b, b is not a combination of the columns $(1, 1, 1)$ and $(-1, 1, 2)$. In Figure 3.9, the three points are not on a line. Least squares replaces points b that are not on a line by points p that are! Unable to solve $Ax = b$, we solve $A\widehat{x} = p$.

The line $\frac{9}{7} + \frac{4}{7}t$ has heights $\frac{5}{7}, \frac{13}{7}, \frac{17}{7}$ at the measurement times $-1, 1, 2$. ***Those points do lie on a line.*** Therefore the vector $p = \left(\frac{5}{7}, \frac{13}{7}, \frac{17}{7}\right)$ is in the column space. *This vector is the projection.* Figure 3.9b is in three dimensions (or m dimensions if there are m points) and Figure 3.9a is in two dimensions (or n dimensions if there are n parameters).

Subtracting p from b, the errors are $e = \left(\frac{2}{7}, -\frac{6}{7}, \frac{4}{7}\right)$. Those are the vertical errors in Figure 3.9a, and they are the components of the dashed vector in Figure 3.9b. This error vector is orthogonal to the first column $(1, 1, 1)$, since $\frac{2}{7} - \frac{6}{7} + \frac{4}{7} = 0$. It is orthogonal to the second column $(-1, 1, 2)$, because $-\frac{2}{7} - \frac{6}{7} + \frac{8}{7} = 0$. It is orthogonal to the column space, and it is in the left nullspace.

Question: If the measurements $b = \left(\frac{2}{7}, -\frac{6}{7}, \frac{4}{7}\right)$ were those errors, what would be the best line and the best $\widehat{x}$? Answer: The zero line—which is the horizontal axis—and $\widehat{x} = 0$. Projection to zero.

We can quickly summarize the equations for fitting by a straight line. The first column of A contains 1s, and the second column contains the times t_i. Therefore $A^{\mathrm{T}}A$ contains the sum of the 1s and the t_i and the t_i^2:

3O The measurements $b_1, \ldots, b_m$ are given at distinct points $t_1, \ldots, t_m$. Then the straight line $\widehat{C} + \widehat{D}t$ which minimizes E^2 comes from least squares:

$$A^{\mathrm{T}}A \begin{bmatrix} \widehat{C} \\ \widehat{D} \end{bmatrix} = A^{\mathrm{T}}b \quad \text{or} \quad \begin{bmatrix} m & \sum t_i \\ \sum t_i & \sum t_i^2 \end{bmatrix} \begin{bmatrix} \widehat{C} \\ \widehat{D} \end{bmatrix} = \begin{bmatrix} \sum b_i \\ \sum t_i b_i \end{bmatrix}.$$

Remark The mathematics of least squares is not limited to fitting the data by straight lines. In many experiments there is no reason to expect a linear relationship, and it would be crazy to look for one. Suppose we are handed some radioactive material. The output b will be the reading on a Geiger counter at various times t. We may know that we are holding a mixture of two chemicals, and we may know their half-lives (or rates of decay), but we do not know how much of each is in our hands. If these two unknown amounts are C and D, then the Geiger counter readings would behave like the sum of two exponentials (and not like a straight line):

$$b = Ce^{-\lambda t} + De^{-\mu t}. \tag{8}$$

In practice, the Geiger counter is not exact. Instead, we make readings $b_1, \ldots, b_m$ at times $t_1, \ldots, t_m$, and equation (8) is approximately satisfied:

$$Ax = b \quad \text{is} \quad \begin{matrix} Ce^{-\lambda t_1} & + & De^{-\mu t_1} & \approx & b_1 \\ & & \vdots & & \\ Ce^{-\lambda t_m} & + & De^{-\mu t_m} & \approx & b_m. \end{matrix}$$

If there are more than two readings, $m > 2$, then in all likelihood we cannot solve for C and D. But the least-squares principle will give optimal values $\widehat{C}$ and $\widehat{D}$.

The situation would be completely different if we knew the amounts C and D, and were trying to discover the decay rates λ and μ. This is a problem in *nonlinear least squares*, and it is harder. We would still form E^2, the sum of the squares of the errors, and minimize it. But setting its derivatives to zero will not give linear equations for the optimal λ and μ. In the exercises, we stay with linear least squares.

Weighted Least Squares

A simple least-squares problem is the estimate $\widehat{x}$ of a patient's weight from two observations $x = b_1$ and $x = b_2$. Unless $b_1 = b_2$, we are faced with an inconsistent system of two equations in one unknown:

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} [x] = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}.$$

Up to now, we accepted b_1 and b_2 as equally reliable. We looked for the value $\widehat{x}$ that minimized $E^2 = (x - b_1)^2 + (x - b_2)^2$:

$$\frac{dE^2}{dx} = 0 \quad \text{at} \quad \widehat{x} = \frac{b_1 + b_2}{2}.$$

The optimal $\widehat{x}$ is the average. The same conclusion comes from $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$. In fact $A^{\mathrm{T}}A$ is a 1 by 1 matrix, and the normal equation is $2\widehat{x} = b_1 + b_2$.

Now suppose the two observations are not trusted to the same degree. The value $x = b_1$ may be obtained from a more accurate scale—or, in a statistical problem, from a larger sample—than $x = b_2$. Nevertheless, if b_2 contains some information, we are not willing to rely totally on b_1. The simplest compromise is to attach different weights w_1^2 and w_2^2, and choose the $\widehat{x}_W$ that minimizes the *weighted sum of squares*:

$$\textbf{Weighted error} \qquad E^2 = w_1^2(x - b_1)^2 + w_2^2(x - b_2)^2.$$

If $w_1 > w_2$, more importance is attached to b_1. The minimizing process (derivative $= 0$) tries harder to make $(x - b_1)^2$ small:

$$\frac{dE^2}{dx} = 2\left[w_1^2(x - b_1) + w_2^2(x - b_2)\right] = 0 \quad \text{at} \quad \widehat{x}_W = \frac{w_1^2 b_1 + w_2^2 b_2}{w_1^2 + w_2^2}. \tag{9}$$

Instead of the average of b_1 and b_2 (for $w_1 = w_2 = 1$), $\widehat{x}_W$ is a ***weighted average*** of the data. This average is closer to b_1 than to b_2.

The ordinary least-squares problem leading to $\widehat{x}_W$ comes from changing $Ax = b$ to the new system $WAx = Wb$. ***This changes the solution from $\widehat{x}$ to $\widehat{x}_W$***. The matrix $W^{\mathrm{T}}W$ turns up on both sides of the weighted normal equations:

The least squares solution to $WAx = Wb$ is $\widehat{x}_W$:

Weighted normal equations $\quad (A^{\mathrm{T}}W^{\mathrm{T}}WA)\widehat{x}_W = A^{\mathrm{T}}W^{\mathrm{T}}Wb.$

What happens to the picture of b projected to $A\widehat{x}$? The projection $A\widehat{x}_W$ is still the point in the column space that is closest to b. But the word "closest" has a new meaning when the length involves W. The *weighted length* of x equals the ordinary length of Wx. Perpendicularity no longer means $y^{\mathrm{T}}x = 0$; in the new system the test is $(Wy)^{\mathrm{T}}(Wx) = 0$. The matrix $W^{\mathrm{T}}W$ appears in the middle. In this new sense, the projection $A\widehat{x}_W$ and the error $b - A\widehat{x}_W$ are again perpendicular.

That last paragraph describes ***all inner products***: They come from invertible matrices W. They involve only the symmetric combination $C = W^{\mathrm{T}}W$. ***The inner product of x and y is $y^{\mathrm{T}}Cx$.*** For an orthogonal matrix $W = Q$, when this combination is $C = Q^{\mathrm{T}}Q = I$, the inner product is not new or different. Rotating the space leaves the inner product unchanged. Every other W changes the length and inner product.

For any invertible matrix W, these rules define a new inner product and length:

$$\textbf{Weighted by } \boldsymbol{W} \qquad (x, y)_W = (Wy)^{\mathrm{T}}(Wx) \qquad \textit{and} \qquad \|x\|_W = \|Wx\|. \tag{10}$$

Since W is invertible, no vector is assigned length zero (except the zero vector). All possible inner products—which depend linearly on x and y and are positive when $x = y \neq 0$—are found in this way, from some matrix $C = W^{\mathrm{T}}W$.

In practice, the important question is the choice of C. The best answer comes from statisticians, and originally from Gauss. We may know that the average error is zero. That is the "expected value" of the error in b—although the error is not really expected to be zero! We may also know the ***average of the square*** of the error; that is the ***variance***. If the errors in the b_i are independent of each other, and their variances are σ_i^2, then ***the right weights are $w_i = 1/\sigma_i$***. A more accurate measurement, which means a smaller variance, gets a heavier weight.

In addition to unequal reliability, *the observations may not be independent*. If the errors are coupled—the polls for President are not independent of those for Senator, and certainly not of those for Vice-President—then W has off-diagonal terms. The best unbiased matrix $C = W^{\mathrm{T}}W$ is the ***inverse of the covariance matrix***—whose i, j entry is the expected value of (error in b_i) times (error in b_j). Then the main diagonal of C^{-1} contains the variances σ_i^2, which are the average of (error in $b_i)^2$.

Suppose two bridge partners both guess (after the bidding) the total number of spades they hold. For each guess, the errors -1, 0, 1 might have equal probability $\frac{1}{3}$. Then the expected error is zero and the variance is $\frac{2}{3}$:

$$E(e) = \tfrac{1}{3}(-1) + \tfrac{1}{3}(0) + \tfrac{1}{3}(1) = 0$$
$$E(e^2) = \tfrac{1}{3}(-1)^2 + \tfrac{1}{3}(0)^2 + \tfrac{1}{3}(1)^2 = \tfrac{2}{3}.$$

The two guesses are dependent, because they are based on the same bidding—but n identical, because they are looking at different hands. Say the chance that they a both too high or both too low is zero, but the chance of opposite errors is $\frac{1}{3}$. The $E(e_1 e_2) = \frac{1}{3}(-1)$, and the inverse of the covariance matrix is $W^{\mathrm{T}}W$:

$$\begin{bmatrix} E(e_1^2) & E(e_1 e_2) \\ E(e_1 e_2) & E(e_2^2) \end{bmatrix}^{-1} = \begin{bmatrix} \frac{2}{3} & -\frac{1}{3} \\ -\frac{1}{3} & \frac{2}{3} \end{bmatrix}^{-1} = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} = C = W^{\mathrm{T}}W.$$

This matrix goes into the middle of the weighted normal equations.

Problem Set 3.3

1. Find the best least-squares solution $\widehat{x}$ to $3x = 10$, $4x = 5$. What error E^2 is min imized? Check that the error vector $(10 - 3\widehat{x}, 5 - 4\widehat{x})$ is perpendicular to th column $(3, 4)$.

2. Suppose the values $b_1 = 1$ and $b_2 = 7$ at times $t_1 = 1$ and $t_2 = 2$ are fitted by a lin $b = Dt$ *through the origin*. Solve $D = 1$ and $2D = 7$ by least squares, and sketcl the best line.

3. Solve $Ax = b$ by least squares, and find $p = A\widehat{x}$ if

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

Verify that the error $b - p$ is perpendicular to the columns of A.

4. Write out $E^2 = \|Ax - b\|^2$ and set to zero its derivatives with respect to u and v, if

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{bmatrix}, \qquad x = \begin{bmatrix} u \\ v \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 3 \\ 4 \end{bmatrix}.$$

Compare the resulting equations with $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$, confirming that calculus as well as geometry gives the normal equations. Find the solution $\widehat{x}$ and the projection $p = A\widehat{x}$. Why is $p = b$?

5. The following system has no solution:

$$Ax = \begin{bmatrix} 1 & -1 \\ 1 & 0 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} C \\ D \end{bmatrix} = \begin{bmatrix} 4 \\ 5 \\ 9 \end{bmatrix} = b.$$

Sketch and solve a straight-line fit that leads to the minimization of the quadratic $(C - D - 4)^2 + (C - 5)^2 + (C + D - 9)^2$. What is the projection of b onto the column space of A?

6. Find the projection of b onto the column space of A:

$$A = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ -2 & 4 \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 2 \\ 7 \end{bmatrix}.$$

Split b into $p + q$, with p in the column space and q perpendicular to that space. Which of the four subspaces contains q?

7. Find the projection matrix P onto the space spanned by $a_1 = (1, 0, 1)$ and $a_2 = (1, 1, -1)$.

8. If P is the projection matrix onto a k-dimensional subspace $\mathbf{S}$ of the whole space $\mathbf{R}^n$, what is the column space of P and what is its rank?

9. (a) If $P = P^{\mathrm{T}}P$, show that P is a projection matrix.
(b) What subspace does the matrix $P = 0$ project onto?

10. If the vectors a_1, a_2, and b are orthogonal, what are $A^{\mathrm{T}}A$ and $A^{\mathrm{T}}b$? What is the projection of b onto the plane of a_1 and a_2?

11. Suppose P is the projection matrix onto the subspace $\mathbf{S}$ and Q is the projection onto the orthogonal complement $\mathbf{S}^{\perp}$. What are $P + Q$ and PQ? Show that $P - Q$ is its own inverse.

12. If $\mathbf{V}$ is the subspace spanned by $(1, 1, 0, 1)$ and $(0, 0, 1, 0)$, find
(a) a basis for the orthogonal complement $\mathbf{V}^{\perp}$.
(b) the projection matrix P onto $\mathbf{V}$.
(c) the vector in $\mathbf{V}$ closest to the vector $b = (0, 1, 0, -1)$ in $\mathbf{V}^{\perp}$.

13. Find the best straight-line fit (least squares) to the measurements

$$\begin{array}{llll} b = 4 \text{ at } t = -2, & b = 3 \text{ at } t = -1, \\ b = 1 \text{ at } t = 0, & b = 0 \text{ at } t = 2. \end{array}$$

Then find the projection of $b = (4, 3, 1, 0)$ onto the column space of

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

14. The vectors $a_1 = (1, 1, 0)$ and $a_2 = (1, 1, 1)$ span a plane in $\mathbf{R}^3$. Find the projection matrix P onto the plane, and find a nonzero vector b that is projected to zero.

15. If P is the projection matrix onto a line in the x-y plane, draw a figure to describe the effect of the "reflection matrix" $H = I - 2P$. Explain both geometrically and algebraically why $H^2 = I$.

16. Show that if u has unit length, then the rank-1 matrix $P = uu^{\mathrm{T}}$ is a projection matrix: It has properties (i) and (ii) in 3N. By choosing $u = a/\|a\|$, P becomes the projection onto the line through a, and Pb is the point $p = \widehat{x}a$. Rank-1 projections correspond exactly to least-squares problems in one unknown.

17. What 2 by 2 matrix projects the x-y plane onto the $-45°$ line $x + y = 0$?

18. We want to fit a plane $y = C + Dt + Ez$ to the four points

$$\begin{array}{ll} y = 3 \text{ at } t = 1,\ z = 1 & y = 6 \text{ at } t = 0,\ z = 3 \\ y = 5 \text{ at } t = 2,\ z = 1 & y = 0 \text{ at } t = 0,\ z = 0. \end{array}$$

(a) Find 4 equations in 3 unknowns to pass a plane through the points (if there is such a plane).
(b) Find 3 equations in 3 unknowns for the best least-squares solution.

19. If $P_C = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$ is the projection onto the column space of A, what is the projection P_R onto the row space? (It is not P_C^{T}!)

20. If P is the projection onto the column space of A, what is the projection onto the left nullspace?

21. Suppose L_1 is the line through the origin in the direction of a_1 and L_2 is the line through b in the direction of a_2. To find the closest points x_1a_1 and $b + x_2a_2$ on the two lines, write the two equations for the x_1 and x_2 that minimize $\|x_1a_1 - x_2a_2 - b\|$. Solve for x if $a_1 = (1, 1, 0)$, $a_2 = (0, 1, 0)$, $b = (2, 1, 4)$.

22. Find the best line $C + Dt$ to fit $b = 4, 2, -1, 0, 0$ at times $t = -2, -1, 0, 1, 2$.

23. Show that the best least-squares fit to a set of measurements $y_1, \ldots, y_m$ by a *horizontal line* (a constant function $y = C$) is their average

$$C = \frac{y_1 + \cdots + y_m}{m}.$$

24. Find the best straight-line fit to the following measurements, and sketch your solution:

$$\begin{aligned} y &= 2 \quad \text{at} \quad t = -1, & y &= 0 \quad \text{at} \quad t = 0, \\ y &= -3 \quad \text{at} \quad t = 1, & y &= -5 \quad \text{at} \quad t = 2. \end{aligned}$$

25. Suppose that instead of a straight line, we fit the data in Problem 24 by a parabola: $y = C + Dt + Et^2$. In the inconsistent system $Ax = b$ that comes from the four measurements, what are the coefficient matrix A, the unknown vector x, and the data vector b? You need not compute $\widehat{x}$.

26. A Middle-Aged man was stretched on a rack to lengths $L = 5$, 6, and 7 feet under applied forces of $F = 1$, 2, and 4 tons. Assuming Hooke's law $L = a + bF$, find his normal length a by least squares.

Problems 27–31 introduce basic ideas of statistics—the foundation for least squares.

27. (Recommended) This problem projects $b = (b_1, \ldots, b_m)$ onto the line through $a = (1, \ldots, 1)$. We solve m equations $ax = b$ in 1 unknown (by least squares).
(a) Solve $a^{\mathrm{T}}a\widehat{x} = a^{\mathrm{T}}b$ to show that $\widehat{x}$ is the *mean* (the average) of the b's.
(b) Find $e = b - a\widehat{x}$, the *variance* $\|e\|^2$, and the *standard deviation* $\|e\|$.
(c) The horizontal line $\widehat{b} = 3$ is closest to $b = (1, 2, 6)$. Check that $p = (3, 3, 3)$ is perpendicular to e and find the projection matrix P.

28. First assumption behind least squares: Each measurement error has ***mean zero***. Multiply the 8 error vectors $b - Ax = (\pm 1, \pm 1, \pm 1)$ by $(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$ to show that the 8 vectors $\widehat{x} - x$ also average to zero. The estimate $\widehat{x}$ is *unbiased*.

29. Second assumption behind least squares: The m errors e_i are independent with variance σ^2, so the average of $(b - Ax)(b - Ax)^{\mathrm{T}}$ is $\sigma^2 I$. Multiply on the left by $(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$ and on the right by $A(A^{\mathrm{T}}A)^{-1}$ to show that the average of $(\widehat{x} - x)(\widehat{x} - x)^{\mathrm{T}}$ is $\sigma^2(A^{\mathrm{T}}A)^{-1}$. This is the all-important ***covariance matrix*** for the error in $\widehat{x}$.

30. A doctor takes four readings of your heart rate. The best solution to $x = b_1, \ldots, x = b_4$ is the average $\widehat{x}$ of $b_1, \ldots, b_4$. The matrix A is a column of 1s. Problem 29 gives the expected error $(\widehat{x} - x)^2$ as $\sigma^2(A^{\mathrm{T}}A)^{-1} =$ ______. By averaging, the variance drops from σ^2 to $\sigma^2/4$.

31. If you know the average $\widehat{x}_9$ of 9 numbers $b_1, \ldots, b_9$, how can you quickly find the average $\widehat{x}_{10}$ with one more number b_{10}? The idea of *recursive* least squares is to avoid adding 10 numbers. What coefficient of $\widehat{x}_9$ correctly gives $\widehat{x}_{10}$?

$$\widehat{x}_{10} = \tfrac{1}{10}b_{10} + \underline{\tfrac{9}{10}}\,\widehat{x}_9 = \tfrac{1}{10}(b_1 + \cdots + b_{10}).$$

Problems 32–37 use four points $b = (0, 8, 8, 20)$ to bring out more ideas.

32. With $b = 0, 8, 8, 20$ at $t = 0, 1, 3, 4$, set up and solve the normal equations $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$. For the best straight line as in Figure 3.9a, find its four heights p_i and four errors e_i. What is the minimum value $E^2 = e_1^2 + e_2^2 + e_3^2 + e_4^2$?

33. (Line $C + Dt$ does go through p's) With $b = 0, 8, 8, 20$ at times $t = 0, 1, 3, 4$, write the four equations $Ax = b$ (unsolvable). Change the measurements to $p = 1, 5, 13, 17$ and find an exact solution to $A\widehat{x} = p$.

34. Check that $e = b - p = (-1, 3, -5, 3)$ is perpendicular to both columns of A. What is the shortest distance $\|e\|$ from b to the column space of A?

35. For the closest parabola $b = C + Dt + Et^2$ to the same four points, write the unsolvable equations $Ax = b$ in three unknowns $x = (C, D, E)$. Set up the three normal equations $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$ (solution not required). You are now fitting a parabola to four points—what is happening in Figure 3.9b?

36. For the closest cubic $b = C + Dt + Et^2 + Ft^3$ to the same four points, write the four equations $Ax = b$. Solve them by elimination. This cubic now goes exactly through the points. What are p and e?

37. The average of the four times is $\widehat{t} = \frac{1}{4}(0 + 1 + 3 + 4) = 2$. The average of the four b's is $\widehat{b} = \frac{1}{4}(0 + 8 + 8 + 20) = 9$.

(a) Verify that the best line goes *through the center point* $(\widehat{t}, \widehat{b}) = (2, 9)$.
(b) Explain why $C + D\widehat{t} = \widehat{b}$ comes from the first equation in $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$.

38. What happens to the weighted average $\widehat{x}_W = (w_1^2 b_1 + w_2^2 b_2)/(w_1^2 + w_2^2)$ if the first weight w_1 approaches zero? The measurement b_1 is totally unreliable.

39. From m independent measurements $b_1, \ldots, b_m$ of your pulse rate, weighted by $w_1, \ldots, w_m$, what is the weighted average that replaces equation (9)? It is the best estimate when the statistical variances are $\sigma_i^2 = 1/w_i^2$.

40. If $W = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$, find the W-inner product of $x = (2, 3)$ and $y = (1, 1)$, and the W-length of x. What line of vectors is W-perpendicular to y?

41. Find the weighted least-squares solution $\widehat{x}_W$ to $Ax = b$:

$$A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix} \quad b = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \quad W = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Check that the projection $A\widehat{x}_W$ is still perpendicular (in the W-inner product!) to the error $b - A\widehat{x}_W$.

42. (a) Suppose you guess your professor's age, making errors $e = -2, -1, 5$ with probabilities $\frac{1}{2}, \frac{1}{4}, \frac{1}{4}$. Check that the expected error $E(e)$ is zero and find the variance $E(e^2)$.

(b) If the professor guesses too (or tries to remember), making errors $-1, 0, 1$ with probabilities $\frac{1}{8}, \frac{6}{8}, \frac{1}{8}$, what weights w_1 and w_2 give the reliability of your guess and the professor's guess?

3.4 ORTHOGONAL BASES AND GRAM-SCHMIDT

In an orthogonal basis, every vector is perpendicular to every other vector. The coordinate axes are mutually orthogonal. That is just about optimal, and the one possible improvement is easy: Divide each vector by its length, to make it a *unit vector.* That changes an ***orthogonal*** basis into an ***orthonormal*** basis of q's:

3P The vectors $q_1, \ldots, q_n$ are ***orthonormal*** if

$$q_i^{\mathrm{T}} q_j = \begin{cases} 0 & \text{whenever } i \neq j, \quad \text{giving the orthogonality;} \\ 1 & \text{whenever } i = j, \quad \text{giving the normalization.} \end{cases}$$

A matrix with orthonormal columns will be called Q.

The most important example is the *standard basis*. For the x-y plane, the best-known axes $e_1 = (1, 0)$ and $e_2 = (0, 1)$ are not only perpendicular but horizontal and vertical. Q is the 2 by 2 identity matrix. In n dimensions the standard basis $e_1, \ldots, e_n$ again consists of *the columns of* $Q = I$:

Standard basis
$$e_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad e_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \ldots, \quad e_n = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}.$$

That is not the only orthonormal basis! We can rotate the axes without changing the right angles at which they meet. These rotation matrices will be examples of Q.

If we have a subspace of $\mathbf{R}^n$, the standard vectors e_i might not lie in that subspace. But the subspace always has an orthonormal basis, and it can be constructed in a simple way out of any basis whatsoever. This construction, which converts a skewed set of axes into a perpendicular set, is known as ***Gram-Schmidt orthogonalization***.

To summarize, the three topics basic to this section are:

1. The definition and properties of orthogonal matrices Q.
2. The solution of $Qx = b$, either n by n or rectangular (least squares).
3. The Gram-Schmidt process and its interpretation as a new factorization $A = QR$.

$Q^T Q = I$

3Q If Q (square or rectangular) has orthonormal columns, then $Q^{\mathrm{T}}Q = I$:

Orthonormal columns
$$\begin{bmatrix} \text{---} & q_1^{\mathrm{T}} & \text{---} \\ \text{---} & q_2^{\mathrm{T}} & \text{---} \\ & & \\ \text{---} & q_n^{\mathrm{T}} & \text{---} \end{bmatrix} \begin{bmatrix} | & | & & | \\ q_1 & q_2 & & q_n \\ | & | & & | \end{bmatrix} = \begin{bmatrix} 1 & 0 & \cdot & 0 \\ 0 & 1 & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \cdot & 1 \end{bmatrix} = I. \quad (1)$$

An orthogonal matrix is a square matrix with orthonormal columns.* Then Q^{T} is Q^{-1}. For square orthogonal matrices, ***the transpose is the inverse.***

When row i of Q^{T} multiplies column j of Q, the result is $q_i^{\mathrm{T}}q_j = 0$. On the diagonal where $i = j$, we have $q_i^{\mathrm{T}}q_i = 1$. That is the normalization to unit vectors of length 1.

Note that $Q^{\mathrm{T}}Q = I$ even if Q is rectangular. But then Q^{T} is only a left-inverse.

$$Q = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}, \qquad Q^{\mathrm{T}} = Q^{-1} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}.$$

Q rotates every vector through the angle θ, and Q^{T} rotates it back through $-\theta$. The columns are clearly orthogonal, and they are orthonormal because $\sin^2\theta + \cos^2\theta = 1$. The matrix Q^{T} is just as much an orthogonal matrix as Q.

Any permutation matrix P is an orthogonal matrix. The columns are certainly unit vectors and certainly orthogonal—because the 1 appears in a different place in each column: The transpose is the inverse.

$$\text{If } P = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \text{ then } P^{-1} = P^{\mathrm{T}} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

An anti-diagonal P, with $P_{13} = P_{22} = P_{31} = 1$, takes the x-y-z axes into the z-y-x axes—a "right-handed" system into a "left-handed" system. So we were wrong if we suggested that every orthogonal Q represents a rotation. *A reflection is also allowed.* $P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ reflects every point (x, y) into (y, x), its mirror image across the 45° line. Geometrically, an orthogonal Q is the product of a rotation and a reflection.

There does remain one property that is shared by rotations and reflections, and in fact by every orthogonal matrix. It is not shared by projections, which are not orthogonal or even invertible. Projections reduce the length of a vector, whereas orthogonal matrices

---

* Orthonormal matrix would have been a better name, but it is too late to change. Also, there is no accepted word for a rectangular matrix with orthonormal columns. We still write Q, but we won't call it an "orthogonal matrix" unless it is square.

have a property that is the most important and most characteristic of all:

3R Multiplication by any Q preserves lengths:

$$\textbf{Lengths unchanged} \qquad \|Qx\| = \|x\| \quad \text{for every vector } x. \tag{2}$$

It also preserves inner products and angles, since $(Qx)^{\mathrm{T}}(Qy) = x^{\mathrm{T}}Q^{\mathrm{T}}Qy = x^{\mathrm{T}}y$.

The preservation of lengths comes directly from $Q^{\mathrm{T}}Q = I$:

$$\|Qx\|^2 = \|x\|^2 \quad \text{because} \quad (Qx)^{\mathrm{T}}(Qx) = x^{\mathrm{T}}Q^{\mathrm{T}}Qx = x^{\mathrm{T}}x. \tag{3}$$

All inner products and lengths are preserved, when the space is rotated or reflected.

We come now to the calculation that uses the special property $Q^{\mathrm{T}} = Q^{-1}$. If we have a basis, then any vector is a combination of the basis vectors. This is exceptionally simple for an orthonormal basis, which will be a key idea behind Fourier series. The problem is *to find the coefficients of the basis vectors*:

> ***Write b as a combination*** $\boldsymbol{b = x_1q_1 + x_2q_2 + \cdots + x_nq_n}$.

To compute x_1 there is a neat trick. *Multiply both sides of the equation by* q_1^{T}. On the left-hand side is $q_1^{\mathrm{T}}b$. On the right-hand side all terms disappear (because $q_1^{\mathrm{T}}q_j = 0$) except the first term. We are left with

$$q_1^{\mathrm{T}}b = x_1q_1^{\mathrm{T}}q_1.$$

Since $q_1^{\mathrm{T}}q_1 = 1$, we have found $x_1 = q_1^{\mathrm{T}}b$. Similarly the second coefficient is $x_2 = q_2^{\mathrm{T}}b$; that term survives when we multiply by q_2^{T}. The other terms die of orthogonality. Each piece of b has a simple formula, and recombining the pieces gives back b:

$$\textbf{\textit{Every vector b is equal to}} \; (q_1^{\mathrm{T}}b)q_1 + (q_2^{\mathrm{T}}b)q_2 + \cdots + (q_n^{\mathrm{T}}b)q_n. \tag{4}$$

I can't resist putting this orthonormal basis into a square matrix Q. The vector equation $x_1q_1 + \cdots + x_nq_n = b$ is identical to $Qx = b$. (The columns of Q multiply the components of x.) Its solution is $x = Q^{-1}b$. But since $Q^{-1} = Q^{\mathrm{T}}$—this is where orthonormality enters—the solution is also $x = Q^{\mathrm{T}}b$:

$$x = Q^{\mathrm{T}}b = \begin{bmatrix} \text{——} & q_1^{\mathrm{T}} & \text{——} \\ & & \\ \text{——} & q_n^{\mathrm{T}} & \text{——} \end{bmatrix} \begin{bmatrix} \\ b \\ \\ \end{bmatrix} = \begin{bmatrix} q_1^{\mathrm{T}}b \\ \\ q_n^{\mathrm{T}}b \end{bmatrix} \tag{5}$$

The components of x are the inner products $q_i^{\mathrm{T}}b$, as in equation (4).

The matrix form also shows what happens when the columns are *not* orthonormal. Expressing b as a combination $x_1a_1 + \cdots + x_na_n$ is the same as solving $Ax = b$. The basis vectors go into the columns of A. In that case we need A^{-1}, which takes work. In the orthonormal case we only need Q^{T}.

Remark 1 The ratio $a^{\mathrm{T}}b/a^{\mathrm{T}}a$ appeared earlier, when we projected b onto a line. Here a is q_1, the denominator is 1, and the projection is $(q_1^{\mathrm{T}}b)q_1$. Thus we have a new interpretation for formula (4): *Every vector b is the sum of its one-dimensional projections onto the lines through the q's.*

Since those projections are orthogonal, Pythagoras should still be correct. The square of the hypotenuse should still be the sum of squares of the components:

$$\|b\|^2 = (q_1^{\mathrm{T}}b)^2 + (q_2^{\mathrm{T}}b)^2 + \cdots + (q_n^{\mathrm{T}}b)^2 \quad \text{which is} \quad \|Q^{\mathrm{T}}b\|^2. \tag{6}$$

Remark 2 Since $Q^{\mathrm{T}} = Q^{-1}$, we also have $QQ^{\mathrm{T}} = I$. When Q comes before Q^{T}, multiplication takes the inner products of the *rows* of Q. (For $Q^{\mathrm{T}}Q$ it was the columns.) Since the result is again the identity matrix, we come to a surprising conclusion: ***The rows of a square matrix are orthonormal whenever the columns are***. The rows point in completely different directions from the columns, and I don't see geometrically why they are forced to be orthonormal—but they are.

Orthonormal columns
Orthonormal rows

$$Q = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{2} & 1/\sqrt{6} \\ 1/\sqrt{3} & 0 & -2/\sqrt{6} \\ 1/\sqrt{3} & -1/\sqrt{2} & 1/\sqrt{6} \end{bmatrix}.$$

Rectangular Matrices with Orthonormal Columns

This chapter is about $Ax = b$, when A is not necessarily square. For $Qx = b$ we now admit the same possibility—there may be more rows than columns. The n orthonormal vectors q_i in the columns of Q have $m > n$ components. Then Q is an m by n matrix and we cannot expect to solve $Qx = b$ exactly. *We solve it by least squares.*

If there is any justice, orthonormal columns should make the problem simple. It worked for square matrices, and now it will work for rectangular matrices. The key is to notice that *we still have* $Q^{\mathrm{T}}Q = I$. So Q^{T} is still the ***left-inverse*** of Q.

For least squares that is all we need. The normal equations came from multiplying $Ax = b$ by the transpose matrix, to give $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$. Now the normal equations are $Q^{\mathrm{T}}Q\widehat{x} = Q^{\mathrm{T}}b$. But $Q^{\mathrm{T}}Q$ is the identity matrix! Therefore $\widehat{x} = Q^{\mathrm{T}}b$, whether Q is square and $\widehat{x}$ is an exact solution, or Q is rectangular and we need least squares.

3S If Q has orthonormal columns, the least-squares problem becomes easy:

$Qx = b$	rectangular system with no solution for most b.
$Q^{\mathrm{T}}Q\widehat{x} = Q^{\mathrm{T}}b$	normal equation for the best $\widehat{x}$—in which $Q^{\mathrm{T}}Q = I$.
$\widehat{x} = Q^{\mathrm{T}}b$	$\widehat{x}_i$ is $q_i^{\mathrm{T}}b$.
$p = Q\widehat{x}$	the projection of b is $(q_1^{\mathrm{T}}b)q_1 + \cdots + (q_n^{\mathrm{T}}b)q_n$.
$p = QQ^{\mathrm{T}}b$	the projection matrix is $P = QQ^{\mathrm{T}}$.

The last formulas are like $p = A\widehat{x}$ and $P = A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$. When the columns are orthonormal, the "cross-product matrix" $A^{\mathrm{T}}A$ becomes $Q^{\mathrm{T}}Q = I$. The hard part of least squares disappears when vectors are orthonormal. The projections onto the axes are uncoupled, and p is the sum $p = (q_1^{\mathrm{T}}b)q_1 + \cdots + (q_n^{\mathrm{T}}b)q_n$.

We emphasize that those projections do not reconstruct b. In the square case $m = n$, they did. In the rectangular case $m > n$, they don't. They give the projection p and not the original vector b—which is all we can expect when there are more equations than unknowns, and the q's are no longer a basis. The projection matrix is usually

$A(A^{\mathrm{T}}A)^{-1}A^{\mathrm{T}}$, and here it simplifies to

$$P = Q(Q^{\mathrm{T}}Q)^{-1}Q^{\mathrm{T}} \quad \text{or} \quad P = QQ^{\mathrm{T}}.$$

Notice that $Q^{\mathrm{T}}Q$ is the n by n identity matrix, whereas QQ^{T} is an m by m projecti P. It is the identity matrix on the columns of Q (P leaves them alone). But QQ^{T} is t zero matrix on the orthogonal complement (the nullspace of Q^{T}).

The following case is simple but typical. Suppose we project a point $b = (x, y, z)$ onto tł x-y plane. Its projection is $p = (x, y, 0)$, and this is the sum of the separate projectior onto the x- and y-axes:

$$q_1 = \begin{bmatrix}1\\0\\0\end{bmatrix} \quad \text{and} \quad (q_1^{\mathrm{T}}b)q_1 = \begin{bmatrix}x\\0\\0\end{bmatrix}; \qquad q_2 = \begin{bmatrix}0\\1\\0\end{bmatrix} \quad \text{and} \quad (q_2^{\mathrm{T}}b)q_2 = \begin{bmatrix}0\\y\\0\end{bmatrix}.$$

The overall projection matrix is

$$P = q_1q_1^{\mathrm{T}} + q_2q_2^{\mathrm{T}} = \begin{bmatrix}1&0&0\\0&1&0\\0&0&0\end{bmatrix}, \quad \text{and} \quad P\begin{bmatrix}x\\y\\z\end{bmatrix} = \begin{bmatrix}x\\y\\0\end{bmatrix}.$$

Projection onto a plane = sum of projections onto orthonormal q_1 and q_2.

When the measurement times average to zero, fitting a straight line leads to orthogonal columns. Take $t_1 = -3$, $t_2 = 0$, and $t_3 = 3$. Then the attempt to fit $y = C + Dt$ leads to three equations in two unknowns:

$$\begin{aligned} C + Dt_1 &= y_1 \\ C + Dt_2 &= y_2, \\ C + Dt_3 &= y_3 \end{aligned} \quad \text{or} \quad \begin{bmatrix}1&-3\\1&0\\1&3\end{bmatrix}\begin{bmatrix}C\\D\end{bmatrix} = \begin{bmatrix}y_1\\y_2\\y_3\end{bmatrix}.$$

The columns $(1, 1, 1)$ *and* $(-3, 0, 3)$ *are orthogonal.* We can project y separately onto each column, and the best coefficients $\widehat{C}$ and $\widehat{D}$ can be found separately:

$$\widehat{C} = \frac{\begin{bmatrix}1&1&1\end{bmatrix}\begin{bmatrix}y_1&y_2&y_3\end{bmatrix}^{\mathrm{T}}}{1^2+1^2+1^2}, \qquad \widehat{D} = \frac{\begin{bmatrix}-3&0&3\end{bmatrix}\begin{bmatrix}y_1&y_2&y_3\end{bmatrix}^{\mathrm{T}}}{(-3)^2+0^2+3^2}.$$

Notice that $\widehat{C} = (y_1 + y_2 + y_3)/3$ is the *mean* of the data. $\widehat{C}$ gives the best fit by a horizontal line, whereas $\widehat{D}t$ is the best fit by a straight line through the origin. *The columns are orthogonal, so the sum of these two separate pieces is the best fit by any straight line whatsoever.* The columns are not unit vectors, so $\widehat{C}$ and $\widehat{D}$ have the length squared in the denominator.

Orthogonal columns are so much better that it is worth changing to that case. If the average of the observation times is not zero—it is $\bar{t} = (t_1 + \cdots + t_m)/m$—then the time origin can be shifted by $\bar{t}$. Instead of $y = C + Dt$ we work with $y = c + d(t - \bar{t})$. The best line is the same! As in the example, we find

$$\widehat{c} = \frac{\begin{bmatrix}1 & \cdots & 1\end{bmatrix}\begin{bmatrix}y_1 & \cdots & y_m\end{bmatrix}^{\mathrm{T}}}{1^2 + 1^2 + \cdots + 1^2} = \frac{y_1 + \cdots + y_m}{m}$$

$$\widehat{d} = \frac{\begin{bmatrix}(t_1 - \bar{t}) & \cdots & (t_m - \bar{t})\end{bmatrix}\begin{bmatrix}y_1 & \cdots & y_m\end{bmatrix}^{\mathrm{T}}}{(t_1 - \bar{t})^2 + \cdots + (t_m - \bar{t})^2} = \frac{\sum(t_i - \bar{t})y_i}{\sum(t_i - \bar{t})^2}. \tag{8}$$

The best $\widehat{c}$ is the mean, and we also get a convenient formula for $\widehat{d}$. The earlier $A^{\mathrm{T}}A$ had the off-diagonal entries $\sum t_i$, and shifting the time by $\bar{t}$ made these entries zero. This shift is an example of the Gram-Schmidt process, ***which orthogonalizes the situation in advance***.

Orthogonal matrices are crucial to numerical linear algebra, because they introduce no instability. While lengths stay the same, roundoff is under control. Orthogonalizing vectors has become an essential technique. Probably it comes second only to elimination. And it leads to a factorization $A = QR$ that is nearly as famous as $A = LU$.

The Gram-Schmidt Process

Suppose you are given three independent vectors a, b, c. If they are orthonormal, life is easy. To project a vector v onto the first one, you compute $(a^{\mathrm{T}}v)a$. To project the same vector v onto the plane of the first two, you just add $(a^{\mathrm{T}}v)a + (b^{\mathrm{T}}v)b$. To project onto the span of a, b, c, you add three projections. All calculations require only the inner products $a^{\mathrm{T}}v$, $b^{\mathrm{T}}v$, and $c^{\mathrm{T}}v$. But to make this true, we are forced to say, "***If*** they are orthonormal." Now we propose to find a way to ***make*** them orthonormal.

The method is simple. We are given a, b, c and we want q_1, q_2, q_3. There is no problem with q_1: it can go in the direction of a. We divide by the length, so that $q_1 = a/\|a\|$ is a unit vector. The real problem begins with q_2—which has to be orthogonal to q_1. If the second vector b has any component in the direction of q_1 (which is the direction of a), ***that component has to be subtracted***:

$$\textbf{Second vector} \qquad B = b - (q_1^{\mathrm{T}}b)q_1 \quad \text{and} \quad q_2 = B/\|B\|. \tag{9}$$

B is orthogonal to q_1. It is the part of b that goes in a new direction, and not in the direction of a. In Figure 3.10, B is perpendicular to q_1. It sets the direction for q_2.

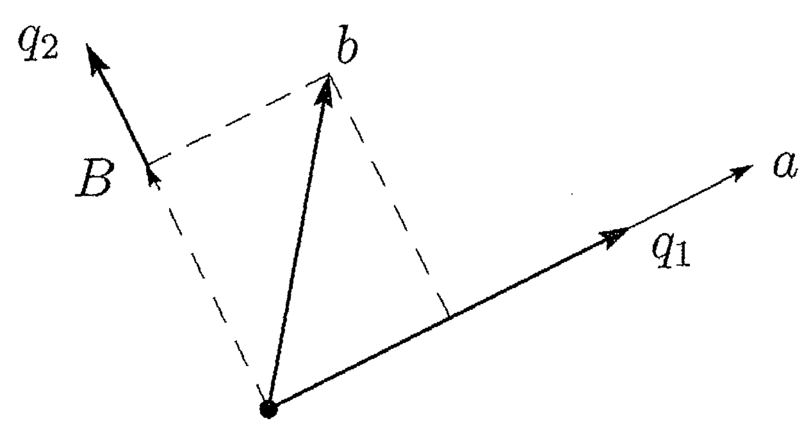

Figure 3.10 The q_1 component of b is removed; a and B normalized to q_1 and q_2.

At this point q_1 and q_2 are set. The third orthogonal direction starts with c. It will not be in the plane of q_1 and q_2, which is the plane of a and b. However, it may have a component in that plane, and that has to be subtracted. (If the result is $C = 0$, this signals that a, b, c were not independent in the first place.) What is left is the component C want, the part that is in a new direction perpendicular to the plane:

$$\textbf{Third vector} \qquad C = c - (q_1^{\mathrm{T}}c)q_1 - (\overset{b}{\cancel{q_2}}{}^{\mathrm{T}}c)\overset{B}{\cancel{q_2}} \quad \text{and} \quad q_3 = C/\|C\|. \qquad (1$$

This is the one idea of the whole Gram-Schmidt process, ***to subtract from every ne vector its components in the directions that are already settled.*** That idea is used ov and over again.* When there is a fourth vector, we subtract away its components in tl directions of q_1, q_2, q_3.

Gram-Schmidt Suppose the independent vectors are a, b, c:

$$a = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \qquad c = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}.$$

To find q_1, make the first vector into a unit vector: $q_1 = a/\sqrt{2}$. To find q_2, subtract fron the second vector its component in the first direction:

$$B = b - (q_1^{\mathrm{T}}b)q_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} - \frac{1}{\sqrt{2}} \begin{bmatrix} 1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}.$$

$a^{T}b/a^{T}a$

The normalized q_2 is B divided by its length, to produce a unit vector:

$$q_2 = \begin{bmatrix} 1/\sqrt{2} \\ 0 \\ -1/\sqrt{2} \end{bmatrix}.$$

To find q_3, subtract from c its components along q_1 and q_2:

$$\begin{aligned} C &= c - (q_1^{\mathrm{T}}c)q_1 - (q_2^{\mathrm{T}}c)q_2 \\ &= \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} - \sqrt{2} \begin{bmatrix} 1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{bmatrix} - \sqrt{2} \begin{bmatrix} 1/\sqrt{2} \\ 0 \\ -1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}. \end{aligned}$$

This is already a unit vector, so it is q_3. I went to desperate lengths to cut down the number of square roots (the painful part of Gram-Schmidt). The result is a set of orthonormal vectors q_1, q_2, q_3, which go into the columns of an orthogonal matrix Q:

$$\textbf{Orthonormal basis} \qquad Q = \begin{bmatrix} & & \\ q_1 & q_2 & q_3 \\ & & \end{bmatrix} = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} & 0 \\ 0 & 0 & 1 \\ 1/\sqrt{2} & -1/\sqrt{2} & 0 \end{bmatrix}.$$

---

* If Gram thought of it first, what was left for Schmidt?

3T The Gram-Schmidt process starts with independent vectors $a_1, \ldots, a_n$ and ends with orthonormal vectors $q_1, \ldots, q_n$. At step j it subtracts from a_j its components in the directions $q_1, \ldots, q_{j-1}$ that are already settled:

$$A_j = a_j - (q_1^{\mathrm{T}} a_j) q_1 - \cdots - (q_{j-1}^{\mathrm{T}} a_j) q_{j-1}. \tag{11}$$

Then q_j is the unit vector $A_j / \|A_j\|$.

Remark on the calculations I think it is easier to compute the orthogonal a, B, C, without forcing their lengths to equal one. Then square roots enter only at the end, when dividing by those lengths. The example above would have the same B and C, without using square roots. Notice the $\frac{1}{2}$ from $a^{\mathrm{T}}b/a^{\mathrm{T}}a$ instead of $\frac{1}{\sqrt{2}}$ from $q^{\mathrm{T}}b$:

$$B = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} - \frac{1}{2}\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \quad \text{and then} \quad C = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} - \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} - 2\begin{bmatrix} \frac{1}{2} \\ 0 \\ -\frac{1}{2} \end{bmatrix}.$$

The Factorization $A = QR$

We started with a matrix A, whose columns were a, b, c. We ended with a matrix Q, whose columns are q_1, q_2, q_3. What is the relation between those matrices? The matrices A and Q are m by n when the n vectors are in m-dimensional space, and there has to be a third matrix that connects them.

The idea is to write the a's as combinations of the q's. The vector b in Figure 3.10 is a combination of the orthonormal q_1 and q_2, and we know what combination it is:

$$b = (q_1^{\mathrm{T}} b) q_1 + (q_2^{\mathrm{T}} b) q_2.$$

Every vector in the plane is the sum of its q_1 and q_2 components. Similarly c is the sum of its q_1, q_2, q_3 components: $c = (q_1^{\mathrm{T}}c)q_1 + (q_2^{\mathrm{T}}c)q_2 + (q_3^{\mathrm{T}}c)q_3$. If we express that in matrix form we have ***the new factorization $A = QR$***:

$$\textbf{\textit{QR} factors} \qquad A = \begin{bmatrix} & & \\ a & b & c \\ & & \end{bmatrix} = \begin{bmatrix} & & \\ q_1 & q_2 & q_3 \\ & & \end{bmatrix} \begin{bmatrix} q_1^{\mathrm{T}}a & q_1^{\mathrm{T}}b & q_1^{\mathrm{T}}c \\ & q_2^{\mathrm{T}}b & q_2^{\mathrm{T}}c \\ & & q_3^{\mathrm{T}}c \end{bmatrix} = QR \tag{12}$$

Notice the zeros in the last matrix! R is *upper triangular* because of the way Gram-Schmidt was done. The first vectors a and q_1 fell on the same line. Then q_1, q_2 were in the same plane as a, b. The third vectors c and q_3 were not involved until step 3.

The QR factorization is like $A = LU$, except that the first factor Q has orthonormal columns. The second factor is called R, because the nonzeros are to the *right* of the diagonal (and the letter U is already taken). The off-diagonal entries of R are the numbers $q_1^{\mathrm{T}}b = 1/\sqrt{2}$ and $q_1^{\mathrm{T}}c = q_2^{\mathrm{T}}c = \sqrt{2}$, found above. The whole factorization is

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

You see the lengths of a, B, C on the diagonal of R. The orthonormal vectors q_1, q_2, q_3, which are the whole object of orthogonalization, are in the first factor Q.

Maybe QR is not as beautiful as LU (because of the square roots). Both factorizations are vitally important to the theory of linear algebra, and absolutely central to the calculations. If LU is Hertz, then QR is Avis.

The entries $r_{ij} = q_i^{\mathrm{T}} a_j$ appear in formula (11), when $\|A_j\| q_j$ is substituted for A_j:

$$a_j = (q_1^{\mathrm{T}} a_j) q_1 + \cdots + (q_{j-1}^{\mathrm{T}} a_j) q_{j-1} + \|A_j\| q_j = Q \text{ times column } j \text{ of } R. \qquad (13)$$

3U Every m by n matrix with independent columns can be factored into $A = QR$. The columns of Q are orthonormal, and R is upper triangular and invertible. When $m = n$ and all matrices are square, Q becomes an orthogonal matrix.

I must not forget the main point of orthogonalization. It simplifies the least-squares problem $Ax = b$. The normal equations are still correct, but $A^{\mathrm{T}}A$ becomes easier:

$$A^{\mathrm{T}}A = R^{\mathrm{T}}Q^{\mathrm{T}}QR = R^{\mathrm{T}}R. \qquad (14)$$

The fundamental equation $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$ simplifies to a triangular system:

$$R^{\mathrm{T}}R\widehat{x} = R^{\mathrm{T}}Q^{\mathrm{T}}b \qquad \text{or} \qquad R\widehat{x} = Q^{\mathrm{T}}b. \qquad (15)$$

Instead of solving $QRx = b$, which can't be done, we solve $R\widehat{x} = Q^{\mathrm{T}}b$, which is just back-substitution because R is triangular. The real cost is the mn^2 operations of Gram-Schmidt, which are needed to find Q and R in the first place.

The same idea of orthogonality applies to functions. The sines and cosines are orthogonal; the powers $1, x, x^2$ are not. When $f(x)$ is written as a combination of sines and cosines, that is a ***Fourier series***. Each term is a projection onto a line—the line in function space containing multiples of $\cos nx$ or $\sin nx$. It is completely parallel to the vector case, and very important. And finally we have a job for Schmidt: To orthogonalize the powers of x and produce the Legendre polynomials.

Function Spaces and Fourier Series

This is a brief and optional section, but it has a number of good intentions:

1. to introduce the most famous infinite-dimensional vector space (*Hilbert space*);
2. to extend the ideas of length and inner product from vectors v to functions $f(x)$;
3. to recognize the Fourier series as a sum of one-dimensional projections (the orthogonal "columns" are the sines and cosines);
4. to apply Gram-Schmidt orthogonalization to the polynomials $1, x, x^2, \ldots$; and
5. to find the best approximation to $f(x)$ by a straight line.

We will try to follow this outline, which opens up a range of new applications for linear algebra, in a systematic way.

1. Hilbert Space. After studying $\mathbf{R}^n$, it is natural to think of the space $\mathbf{R}^\infty$. It contains all vectors $v = (v_1, v_2, v_3, \ldots)$ with an infinite sequence of components. This space is actually too big when there is no control on the size of components v_j. A much better idea is to keep the familiar definition of length, using a sum of squares, and *to include* *only those vectors that have a **finite length***:

$$\textbf{Length squared} \qquad \|v\|^2 = v_1^2 + v_2^2 + v_3^2 +$$

The infinite series must converge to a finite sum. This leaves (1, 1, 1, ...). Vectors with finite length can be added ($\|v + w$ multiplied by scalars, so they form a vector space. It is the celebr

Hilbert space is the natural way to let the number of dimensions become infinite, and at the same time to keep the geometry of ordinary Euclidean space. Ellipses become infinite-dimensional ellipsoids, and perpendicular lines are recognized exactly as before. The vectors v and w are orthogonal when their inner product is zero:

$$\textbf{Orthogonality} \qquad v^{\mathrm{T}}w = v_1w_1 + v_2w_2 + v_3w_3 + \cdots = 0.$$

This sum is guaranteed to converge, and for any two vectors it still obeys the Schwarz inequality $|v^{\mathrm{T}}w| \le \|v\|\,\|w\|$. The cosine, even in Hilbert space, is never larger than 1.

There is another remarkable thing about this space: It is found under a great many different disguises. Its "vectors" can turn into functions, which is the second point.

2. Lengths and Inner Products. Suppose $f(x) = \sin x$ on the interval $0 \le x \le 2\pi$. This f is like a vector with a whole continuum of components, the values of $\sin x$ along the whole interval. To find the length of such a vector, the usual rule of adding the squares of the components becomes impossible. This summation is replaced, in a natural and inevitable way, by *integration*:

$$\textbf{Length } \|f\| \textbf{ of function} \qquad \|f\|^2 = \int_0^{2\pi} (f(x))^2\,dx = \int_0^{2\pi} (\sin x)^2\,dx = \pi. \qquad (17)$$

Our Hilbert space has become a ***function space***. The vectors are functions, we have a way to measure their length, and the space contains all those functions that have a finite length—just as in equation (16). It does not contain the function $F(x) = 1/x$, because the integral of $1/x^2$ is infinite.

The same idea of replacing summation by integration produces the ***inner product of two functions***: If $f(x) = \sin x$ and $g(x) = \cos x$, then their inner product is

$$(f, g) = \int_0^{2\pi} f(x)g(x)\,dx = \int_0^{2\pi} \sin x \cos x\,dx = 0. \qquad (18)$$

This is exactly like the vector inner product $f^{\mathrm{T}}g$. It is still related to the length by $(f, f) = \|f\|^2$. The Schwarz inequality is still satisfied: $|(f, g)| \le \|f\|\,\|g\|$. Of course, two functions like $\sin x$ and $\cos x$—whose inner product is zero—will be called orthogonal. They are even orthonormal after division by their length $\sqrt{\pi}$.

3. The ***Fourier series*** of a function is an expansion into sines and cosines:

$$f(x) = a_0 + a_1 \cos x + b_1 \sin x + a_2 \cos 2x + b_2 \sin 2x + \cdots.$$

To compute a coefficient like b_1, *multiply* both sides by the corresponding function $\sin x$ and *integrate* from 0 to 2π. (The function $f(x)$ is given on that interval.) In other words, *You see the lengths of a, B, C on the diagonal of R.* The orthonormal vectors q_1, q_2, q_3, which are the whole object of orthogonalization, are in the first factor Q.

Maybe QR is not as beautiful as LU (because of the square roots). Both factorizations are vitally important to the theory of linear algebra, and absolutely central to the calculations. If LU is Hertz, then QR is Avis.

The entries $r_{ij} = q_i^{\mathrm{T}} a_j$ appear in formula (11), when $\|A_j\| q_j$ is substituted for A_j:

$$a_j = (q_1^{\mathrm{T}} a_j) q_1 + \cdots + (q_{j-1}^{\mathrm{T}} a_j) q_{j-1} + \|A_j\| q_j = Q \text{ times column } j \text{ of } R. \tag{13}$$

3U Every m by n matrix with independent columns can be factored into $A = QR$. The columns of Q are orthonormal, and R is upper triangular and invertible. When $m = n$ and all matrices are square, Q becomes an orthogonal matrix.

I must not forget the main point of orthogonalization. It simplifies the least-squares problem $Ax = b$. The normal equations are still correct, but $A^{\mathrm{T}}A$ becomes easier:

$$A^{\mathrm{T}}A = R^{\mathrm{T}}Q^{\mathrm{T}}QR = R^{\mathrm{T}}R. \tag{14}$$

The fundamental equation $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$ simplifies to a triangular system:

$$R^{\mathrm{T}}R\widehat{x} = R^{\mathrm{T}}Q^{\mathrm{T}}b \qquad \text{or} \qquad R\widehat{x} = Q^{\mathrm{T}}b. \tag{15}$$

Instead of solving $QRx = b$, which can't be done, we solve $R\widehat{x} = Q^{\mathrm{T}}b$, which is just back-substitution because R is triangular. The real cost is the mn^2 operations of Gram-Schmidt, which are needed to find Q and R in the first place.

The same idea of orthogonality applies to functions. The sines and cosines are orthogonal; the powers $1, x, x^2$ are not. When $f(x)$ is written as a combination of sines and cosines, that is a ***Fourier series***. Each term is a projection onto a line—the line in function space containing multiples of $\cos nx$ or $\sin nx$. It is completely parallel to the vector case, and very important. And finally we have a job for Schmidt: To orthogonalize the powers of x and produce the Legendre polynomials.

Function Spaces and Fourier Series

This is a brief and optional section, but it has a number of good intentions:

1. to introduce the most famous infinite-dimensional vector space (*Hilbert space*);
2. to extend the ideas of length and inner product from vectors v to functions $f(x)$;
3. to recognize the Fourier series as a sum of one-dimensional projections (the orthogonal "columns" are the sines and cosines);
4. to apply Gram-Schmidt orthogonalization to the polynomials $1, x, x^2, \ldots$; and
5. to find the best approximation to $f(x)$ by a straight line.

We will try to follow this outline, which opens up a range of new applications fo linear algebra, in a systematic way.

1. Hilbert Space. After studying $\mathbf{R}^n$, it is natural to think of the space $\mathbf{R}^\infty$. It contain all vectors $v = (v_1, v_2, v_3, \ldots)$ with an infinite sequence of components. This space i actually too big when there is no control on the size of components v_j. A much bette idea is to keep the familiar definition of length, using a sum of squares, and *to includ* *only those vectors that have a **finite length***:

$$\textbf{Length squared} \qquad \|v\|^2 = v_1^2 + v_2^2 + v_3^2 + \cdots \tag{16}$$

The infinite series must converge to a finite sum. This leaves $\left(1, \frac{1}{2}, \frac{1}{3}, \ldots\right)$ but not $(1, 1, 1, \ldots)$. Vectors with finite length can be added ($\|v + w\| \leq \|v\| + \|w\|$) and multiplied by scalars, so they form a vector space. It is the celebrated ***Hilbert space***.

Hilbert space is the natural way to let the number of dimensions become infinite, and at the same time to keep the geometry of ordinary Euclidean space. Ellipses become infinite-dimensional ellipsoids, and perpendicular lines are recognized exactly as before. The vectors v and w are orthogonal when their inner product is zero:

$$\textbf{Orthogonality} \qquad v^{\mathrm{T}}w = v_1w_1 + v_2w_2 + v_3w_3 + \cdots = 0.$$

This sum is guaranteed to converge, and for any two vectors it still obeys the Schwarz inequality $|v^{\mathrm{T}}w| \leq \|v\|\,\|w\|$. The cosine, even in Hilbert space, is never larger than 1.

There is another remarkable thing about this space: It is found under a great many different disguises. Its "vectors" can turn into functions, which is the second point.

2. Lengths and Inner Products. Suppose $f(x) = \sin x$ on the interval $0 \leq x \leq 2\pi$. This f is like a vector with a whole continuum of components, the values of $\sin x$ along the whole interval. To find the length of such a vector, the usual rule of adding the squares of the components becomes impossible. This summation is replaced, in a natural and inevitable way, by *integration*:

$$\textbf{Length } \|f\| \textbf{ of function} \qquad \|f\|^2 = \int_0^{2\pi} (f(x))^2\,dx = \int_0^{2\pi} (\sin x)^2\,dx = \pi. \tag{17}$$

Our Hilbert space has become a ***function space***. The vectors are functions, we have a way to measure their length, and the space contains all those functions that have a finite length—just as in equation (16). It does not contain the function $F(x) = 1/x$, because the integral of $1/x^2$ is infinite.

The same idea of replacing summation by integration produces the ***inner product of two functions***: If $f(x) = \sin x$ and $g(x) = \cos x$, then their inner product is

$$(f, g) = \int_0^{2\pi} f(x)g(x)\,dx = \int_0^{2\pi} \sin x \cos x\,dx = 0. \tag{18}$$

This is exactly like the vector inner product $f^{\mathrm{T}}g$. It is still related to the length by $(f, f) = \|f\|^2$. The Schwarz inequality is still satisfied: $|(f, g)| \leq \|f\|\,\|g\|$. Of course, two functions like $\sin x$ and $\cos x$—whose inner product is zero—will be called orthogonal. They are even orthonormal after division by their length $\sqrt{\pi}$.

3. The ***Fourier series*** of a function is an expansion into sines and cosines:

$$f(x) = a_0 + a_1 \cos x + b_1 \sin x + a_2 \cos 2x + b_2 \sin 2x + \cdots.$$

To compute a coefficient like b_1, *multiply* both sides by the corresponding function $\sin x$ and *integrate* from 0 to 2π. (The function $f(x)$ is given on that interval.) In other words, take the inner product of both sides with $\sin x$:

$$\int_0^{2\pi} f(x)\sin x\,dx = a_0\int_0^{2\pi}\sin x\,dx + a_1\int_0^{2\pi}\cos x\sin x\,dx + b_1\int_0^{2\pi}(\sin x)^2\,dx + \cdots.$$

On the right-hand side, every integral is zero except one—the one in which $\sin x$ multiplies itself. *The sines and cosines are mutually orthogonal* as in equation (18). Therefore b_1 is the left-hand side divided by that one nonzero integral:

$$b_1 = \frac{\int_0^{2\pi} f(x)\sin x\,dx}{\int_0^{2\pi}(\sin x)^2\,dx} = \frac{(f, \sin x)}{(\sin x, \sin x)}.$$

The Fourier coefficient a_1 would have $\cos x$ in place of $\sin x$, and a_2 would use $\cos 2x$.

The whole point is to see the analogy with projections. The component of the vector b along the line spanned by a is $b^{\mathrm{T}}a/a^{\mathrm{T}}a$. A ***Fourier series is projecting $f(x)$ onto $\sin x$***. Its component p in this direction is exactly $b_1 \sin x$.

The coefficient b_1 is the least squares solution of the inconsistent equation $b_1 \sin x = f(x)$. This brings $b_1 \sin x$ as close as possible to $f(x)$. All the terms in the series are projections onto a sine or cosine. Since the sines and cosines are orthogonal, *the Fourier series gives the coordinates of the "vector" $f(x)$ with respect to a set of* (infinitely many) *perpendicular axes*.

4. Gram-Schmidt for Functions. There are plenty of useful functions other than sines and cosines, and they are not always orthogonal. The simplest are the powers of x, and unfortunately there is no interval on which even 1 and x^2 are perpendicular. (Their inner product is always positive, because it is the integral of x^2.) Therefore the closest parabola to $f(x)$ is *not* the sum of its projections onto 1, x, and x^2. There will be a matrix like $(A^{\mathrm{T}}A)^{-1}$, and this coupling is given by the ill-conditioned **Hilbert matrix**. On the interval $0 \le x \le 1$,

$$A^{\mathrm{T}}A = \begin{bmatrix} (1,1) & (1,x) & (1,x^2) \\ (x,1) & (x,x) & (x,x^2) \\ (x^2,1) & (x^2,x) & (x^2,x^2) \end{bmatrix} = \begin{bmatrix} \int 1 & \int x & \int x^2 \\ \int x & \int x^2 & \int x^3 \\ \int x^2 & \int x^3 & \int x^4 \end{bmatrix} = \begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \end{bmatrix}.$$

This matrix has a large inverse, because the axes $1, x, x^2$ are far from perpendicular. Th situation becomes impossible if we add a few more axes. *It is virtually hopeless to solv $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$ for the closest polynomial of degree ten.*

More precisely, it is hopeless to solve this by Gaussian elimination; every roundo error would be amplified by more than 10^{13}. On the other hand, we cannot just giv up; approximation by polynomials has to be possible. The right idea is to switch t orthogonal axes (by Gram-Schmidt). We look for combinations of 1, x, and x^2 that a orthogonal.

It is convenient to work with a symmetrically placed interval like $-1 \le x \le$ because this makes all the odd powers of x orthogonal to all the even powers:

$$(1,x) = \int_{-1}^{1} x\,dx = 0, \qquad (x,x^2) = \int_{-1}^{1} x^3\,dx = 0.$$

Therefore the Gram-Schmidt process can begin by accepting $v_1 = 1$ and $v_2 = x$ as t first two perpendicular axes. Since $(x, x^2) = 0$, it only has to correct the angle betwe

$A=S^{-1}AS$

1 and x^2. By formula (10), the third orthogonal polynomial is

$$\textbf{Orthogonalize} \qquad v_3 = x^2 - \frac{(1,x^2)}{(1,1)}1 - \frac{(x,x^2)}{(x,x)}x = x^2 - \frac{\int_{-1}^{1} x^2\,dx}{\int_{-1}^{1} 1\,dx} = x^2 - \frac{1}{3}.$$

The polynomials constructed in this way are called the ***Legendre polynomials*** and they are orthogonal to each other over the interval $-1 \le x \le 1$.

$$\textbf{Check} \qquad \left(1, x^2 - \frac{1}{3}\right) = \int_{-1}^{1}\left(x^2 - \frac{1}{3}\right)dx = \left[\frac{x^3}{3} - \frac{x}{3}\right]_{-1}^{1} = 0.$$

The closest polynomial of degree ten is now computable, without disaster, by projecting onto each of the first 10 (or 11) Legendre polynomials.

5. Best Straight Line. Suppose we want to approximate $y = x^5$ by a straight line $C + Dx$ between $x = 0$ and $x = 1$. There are at least three ways of finding that line, and if you compare them the whole chapter might become clear!

1. Solve $\begin{bmatrix} 1 & x \end{bmatrix} \begin{bmatrix} C \\ D \end{bmatrix} = x^5$ by least squares. The equation $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$ is

$$\begin{bmatrix} (1,1) & (1,x) \\ (x,1) & (x,x) \end{bmatrix} \begin{bmatrix} C \\ D \end{bmatrix} = \begin{bmatrix} (1,x^5) \\ (x,x^5) \end{bmatrix} \qquad \text{or} \qquad \begin{bmatrix} 1 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{3} \end{bmatrix} \begin{bmatrix} C \\ D \end{bmatrix} = \begin{bmatrix} \frac{1}{6} \\ \frac{1}{7} \end{bmatrix}.$$

2. Minimize $E^2 = \int_0^1 (x^5 - C - Dx)^2\,dx = \frac{1}{11} - \frac{2}{6}C - \frac{2}{7}D + C^2 + CD + \frac{1}{3}D^2$. The derivatives with respect to C and D, after dividing by 2, bring back the normal equations of method 1 (and the solution is $\widehat{C} = \frac{1}{6} - \frac{5}{14}$, $\widehat{D} = \frac{5}{7}$):

$$-\frac{1}{6} + C + \frac{1}{2}D = 0 \qquad \text{and} \qquad -\frac{1}{7} + \frac{1}{2}C + \frac{1}{3}D = 0.$$

3. Apply Gram-Schmidt to replace x by $x - (1,x)/(1,1)$. That is $x - \frac{1}{2}$, which is orthogonal to 1. Now the one-dimensional projections add to the best line:

$$C + Dx = \frac{(x^5,1)}{(1,1)}1 + \frac{\left(x^5, x - \frac{1}{2}\right)}{\left(x - \frac{1}{2}, x - \frac{1}{2}\right)}\left(x - \tfrac{1}{2}\right) = \frac{1}{6} + \frac{5}{7}\left(x - \frac{1}{2}\right).$$

Problem Set 3.4

1. (a) Write the four equations for fitting $y = C + Dt$ to the data

$$\begin{array}{llll} y = -4 & \text{at} \quad t = -2, & y = -3 & \text{at} \quad t = -1 \\ y = -1 & \text{at} \quad t = 1, & y = 0 & \text{at} \quad t = 2. \end{array}$$

Show that the columns are orthogonal.

(b) Find the optimal straight line, draw its graph, and write E^2.

(c) Interpret the zero error in terms of the original system of four equations in two unknowns: The right-hand side $(-4, -3, -1, 0)$ is in the ______ space.

2. Project $b = (0, 3, 0)$ onto each of the orthonormal vectors $a_1 = \left(\frac{2}{3}, \frac{2}{3}, -\frac{1}{3}\right)$ and $a_2 = \left(-\frac{1}{3}, \frac{2}{3}, \frac{2}{3}\right)$, and then find its projection p onto the plane of a_1 and a_2.

3. Find also the projection of $b = (0, 3, 0)$ onto $a_3 = \left(\frac{2}{3}, -\frac{1}{3}, \frac{2}{3}\right)$, and add the three projections. Why is $P = a_1a_1^{\mathrm{T}} + a_2a_2^{\mathrm{T}} + a_3a_3^{\mathrm{T}}$ equal to I?

4. If Q_1 and Q_2 are orthogonal matrices, so that $Q^{\mathrm{T}}Q = I$, show that Q_1Q_2 is also orthogonal. If Q_1 is rotation through θ, and Q_2 is rotation through ϕ, what is Q_1Q_2? Can you find the trigonometric identities for $\sin(\theta+\phi)$ and $\cos(\theta+\phi)$ in the matrix multiplication Q_1Q_2?

5. If u is a unit vector, show that $Q = I - 2uu^{\mathrm{T}}$ is a symmetric orthogonal matrix. (It is a reflection, also known as a Householder transformation.) Compute Q when $u^{\mathrm{T}} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} & -\frac{1}{2} \end{bmatrix}$.

6. Find a third column so that the matrix

$$Q = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{14} & \\ 1/\sqrt{3} & 2/\sqrt{14} & \\ 1/\sqrt{3} & -3/\sqrt{14} & \end{bmatrix}$$

is orthogonal. It must be a unit vector that is orthogonal to the other columns; how much freedom does this leave? Verify that the rows automatically become orthonormal at the same time.

7. Show, by forming $b^{\mathrm{T}}b$ directly, that Pythagoras's law holds for any combination $b = x_1q_1 + \cdots + x_nq_n$ of orthonormal vectors: $\|b\|^2 = x_1^2 + \cdots + x_n^2$. In matrix terms, $b = Qx$, so this again proves that lengths are preserved: $\|Qx\|^2 = \|x\|^2$.

8. Project the vector $b = (1, 2)$ onto two vectors that are not orthogonal, $a_1 = (1, 0)$ and $a_2 = (1, 1)$. Show that, unlike the orthogonal case, the sum of the two one-dimensional projections does not equal b.

9. If the vectors q_1, q_2, q_3 are orthonormal, what combination of q_1 and q_2 is closest to q_3?

$q_1^T q_3 = 0$ q_2

10. If q_1 and q_2 are the outputs from Gram-Schmidt, what were the possible input vectors a and b?

11. Show that an orthogonal matrix that is upper triangular must be diagonal.

12. What multiple of $a_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ should be subtracted from $a_2 = \begin{bmatrix} 4 \\ 0 \end{bmatrix}$ to make the result orthogonal to a_1? Factor $\begin{bmatrix} 1 & 4 \\ 1 & 0 \end{bmatrix}$ into QR with orthonormal vectors in Q.

13. Apply the Gram-Schmidt process to

$$a = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \qquad b = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \qquad c = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

and write the result in the form $A = QR$.

14. From the nonorthogonal a, b, c, find orthonormal vectors q_1, q_2, q_3:

$$a = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \qquad c = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}.$$

15. Find an orthonormal set q_1, q_2, q_3 for which q_1, q_2 span the column space of

$$A = \begin{bmatrix} 1 & 1 \\ 2 & -1 \\ -2 & 4 \end{bmatrix}.$$

Which fundamental subspace contains q_3? What is the least-squares solution of $Ax = b$ if $b = \begin{bmatrix} 1 & 2 & 7 \end{bmatrix}^{\mathrm{T}}$?

16. Express the Gram-Schmidt orthogonalization of a_1, a_2 as $A = QR$:

$$a_1 = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}, \qquad a_2 = \begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix}.$$

Given n vectors a_i with m components, what are the shapes of A, Q, and R?

17. With the same matrix A as in Problem 16, and with $b = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^{\mathrm{T}}$, use $A = QR$ to solve the least-squares problem $Ax = b$.

18. If $A = QR$, find a simple formula for the projection matrix P onto the column space of A.

19. Show that these *modified Gram-Schmidt* steps produce the same C as in equation (10):

$$C^* = c - (q_1^{\mathrm{T}} c) q_1 \qquad \text{and} \qquad C = C^* - (q_2^{\mathrm{T}} C^*)\, q_2.$$

This is much more stable, to subtract the projections one at a time.

20. In Hilbert space, find the length of the vector $v = (1/\sqrt{2}, 1/\sqrt{4}, 1/\sqrt{8}, \ldots)$ and the length of the function $f(x) = e^x$ (over the interval $0 \le x \le 1$). What is the inner product over this interval of e^x and e^{-x}?

21. What is the closest function $a \cos x + b \sin x$ to the function $f(x) = \sin 2x$ on the interval from $-\pi$ to π? What is the closest straight line $c + dx$?

22. By setting the derivative to zero, find the value of b_1 that minimizes

$$\|b_1 \sin x - \cos x\|^2 = \int_0^{2\pi} (b_1 \sin x - \cos x)^2\, dx.$$

Compare with the Fourier coefficient b_1.

23. Find the Fourier coefficients a_0, a_1, b_1 of the step function $y(x)$, which equals 1 on the interval $0 \le x \le \pi$ and 0 on the remaining interval $\pi < x < 2\pi$:

$$a_0 = \frac{(y, 1)}{(1, 1)} \qquad a_1 = \frac{(y, \cos x)}{(\cos x, \cos x)} \qquad b_1 = \frac{(y, \sin x)}{(\sin x, \sin x)}.$$

24. Find the fourth Legendre polynomial. It is a cubic $x^3 + ax^2 + bx + c$ that is orthogonal to 1, x, and $x^2 - \frac{1}{3}$ over the interval $-1 \le x \le 1$.

25. What is the closest straight line to the parabola $y = x^2$ over $-1 \le x \le 1$?

26. In the Gram-Schmidt formula (10), verify that C is orthogonal to q_1 and q_2.

27. Find an orthonormal basis for the subspace spanned by $a_1 = (1, -1, 0, 0)$, $a_2 = (0, 1, -1, 0)$, $a_3 = (0, 0, 1, -1)$.

28. Apply Gram-Schmidt to $(1, -1, 0)$, $(0, 1, -1)$, and $(1, 0, -1)$, to find an orthonormal basis on the plane $x_1 + x_2 + x_3 = 0$. What is the dimension of this subspace, and how many nonzero vectors come out of Gram-Schmidt?

29. (Recommended) Find orthogonal vectors A, B, C by Gram-Schmidt from a, b, c:

$$a = (1, -1, 0, 0) \qquad b = (0, 1, -1, 0) \qquad c = (0, 0, 1, -1).$$

A, B, C and a, b, c are bases for the vectors perpendicular to $d = (1, 1, 1, 1)$.

30. If $A = QR$ then $A^{\mathrm{T}}A = R^{\mathrm{T}}R =$ ______ triangular times ______ triangular. *Gram-Schmidt on A corresponds to elimination on $A^{\mathrm{T}}A$.* Compare

$$A = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & -1 \end{bmatrix} \qquad \text{with} \qquad A^{\mathrm{T}}A = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix}.$$

For $A^{\mathrm{T}}A$, the pivots are $2, \frac{3}{2}, \frac{4}{3}$ and the multipliers are $-\frac{1}{2}$ and $-\frac{2}{3}$.

(a) Using those multipliers in A, show that column 1 of A and $B =$ column 2 $-\frac{1}{2}$(column 1) and $C =$ column 3 $-\frac{2}{3}$ (column 2) are orthogonal.

(b) Check that $\|\text{column } 1\|^2 = 2$, $\|B\|^2 = \frac{3}{2}$, and $\|C\|^2 = \frac{4}{3}$, using the pivots.

31. True or false (give an example in either case):

(a) Q^{-1} is an orthogonal matrix when Q is an orthogonal matrix.

(b) If Q (3 by 2) has orthonormal columns then $\|Qx\|$ always equals $\|x\|$.

32. (a) Find a basis for the subspace $\mathbf{S}$ in $\mathbf{R}^4$ spanned by all solutions of

$$x_1 + x_2 + x_3 - x_4 = 0.$$

(b) Find a basis for the orthogonal complement $\mathbf{S}^{\perp}$.

(c) Find b_1 in $\mathbf{S}$ and b_2 in $\mathbf{S}^{\perp}$ so that $b_1 + b_2 = b = (1, 1, 1, 1)$.

3.5 THE FAST FOURIER TRANSFORM

The Fourier series is linear algebra in infinite dimensions. The "vectors" are functions $f(x)$; they are projected onto the sines and cosines; that produces the Fourier coefficients a_k and b_k. From this infinite sequence of sines and cosines, multiplied by a_k and b_k, we can reconstruct $f(x)$. That is the classical case, which Fourier dreamt about, but in actual calculations it is the ***discrete Fourier transform*** that we compute. Fourier still lives, but in finite dimensions.

This is pure linear algebra, based on orthogonality. The input is a sequence of numbers $y_0, \ldots, y_{n-1}$, instead of a function $f(x)$. The output $c_0, \ldots, c_{n-1}$ has the same length n. The relation between y and c is linear, so it must be given by a matrix. This is the ***Fourier matrix F***, and the whole technology of digital signal processing depends on it. The Fourier matrix has remarkable properties.