GILBERT STRANG

INTRODUCTION TO LINEAR ALGEBRA

WELLESLEY-CAMBRIDGE PRESS

1

INTRODUCTION TO VECTORS

The heart of linear algebra is in two operations—both with vectors. We add vectors to get $\boldsymbol{v}+\boldsymbol{w}$. We multiply by numbers c and d to get $c\boldsymbol{v}$ and $d\boldsymbol{w}$. Combining those operations gives the *linear combination* $c\boldsymbol{v}+d\boldsymbol{w}$.

Chapter 1 explains these central ideas, on which everything builds. We start with two-dimensional vectors and three-dimensional vectors, which are reasonable to draw. Then we move into higher dimensions. The really impressive feature of linear algebra is how smoothly it takes that step into n-dimensional space. Your mental picture stays completely correct, even if drawing a ten-dimensional vector is impossible.

This is where the book is going (into n-dimensional space), and the first steps are the two operations in Sections 1.1 and 1.2:

1.1 ***Vector addition*** $\boldsymbol{v}+\boldsymbol{w}$ and ***linear combinations*** $c\boldsymbol{v}+d\boldsymbol{w}$.

1.2 The ***dot product*** $\boldsymbol{v}\cdot\boldsymbol{w}$ and the ***length*** $\|\boldsymbol{v}\|=\sqrt{\boldsymbol{v}\cdot\boldsymbol{v}}$.

VECTORS AND LINEAR COMBINATIONS ■ 1.1

"You can't add apples and oranges." That sentence might not be news, but it still contains some truth. In a strange way, it is the reason for vectors! If we keep the number of apples separate from the number of oranges, we have a pair of numbers. That pair is a ***two-dimensional vector*** $\boldsymbol{v}$:

$$\boldsymbol{v}=\begin{bmatrix}v_1\\v_2\end{bmatrix}\qquad\begin{aligned}v_1&=\text{number of apples}\\v_2&=\text{number of oranges.}\end{aligned}$$

We wrote $\boldsymbol{v}$ as a ***column vector***. The numbers v_1 and v_2 are its "components." The main point so far is to have a single letter $\boldsymbol{v}$ (in boldface) for this pair of numbers v_1 and v_2 (in lightface).

Even if we don't add v_1 to v_2, we do *add vectors*. The first components of $\boldsymbol{v}$ and $\boldsymbol{w}$ stay separate from the second components:

$$\boldsymbol{v}=\begin{bmatrix}v_1\\v_2\end{bmatrix}\quad\text{and}\quad\boldsymbol{w}=\begin{bmatrix}w_1\\w_2\end{bmatrix}\quad\text{add to}\quad\boldsymbol{v}+\boldsymbol{w}=\begin{bmatrix}v_1+w_1\\v_2+w_2\end{bmatrix}.$$

You see the reason. The total number of apples is $v_1 + w_1$. The number of oranges is $v_2 + w_2$. Vector addition is basic and important. Subtraction of vectors follows the same idea: *The components of* $\boldsymbol{v} - \boldsymbol{w}$ *are* $v_1 - w_1$ *and* ______.

Vectors can be multiplied by 2 or by -1 or by any number c. There are two ways to double a vector. One way is to add $\boldsymbol{v} + \boldsymbol{v}$. The other way (the usual way) is to multiply each component by 2:

$$2\boldsymbol{v} = \begin{bmatrix} 2v_1 \\ 2v_2 \end{bmatrix} \quad \text{and} \quad -\boldsymbol{v} = \begin{bmatrix} -v_1 \\ -v_2 \end{bmatrix}.$$

The components of $c\boldsymbol{v}$ are cv_1 and cv_2. The number c is called a "scalar."

Notice that the sum of $-\boldsymbol{v}$ and $\boldsymbol{v}$ is the zero vector. This is $\mathbf{0}$, which is not the same as the number zero! The vector $\mathbf{0}$ has components 0 and 0. Forgive me for hammering away at the difference between a vector and its components. Linear algebra is built on these operations $\boldsymbol{v} + \boldsymbol{w}$ and $c\boldsymbol{v}$—***adding vectors and multiplying by scalars***.

There is another way to see a vector, that shows all its components at once. The vector $\boldsymbol{v}$ can be represented by an arrow. When $\boldsymbol{v}$ has two components, the arrow is in two-dimensional space (a plane). If the components are v_1 and v_2, the arrow goes v_1 units to the right and v_2 units up. This vector is drawn twice in Figure 1.1. First, it starts at the origin (where the axes meet). This is the usual picture. Unless there is a special reason, our vectors will begin at $(0, 0)$. But the second arrow shows the starting point shifted over to A. The arrows $\overrightarrow{OP}$ and $\overrightarrow{AB}$ represent the same vector. One reason for allowing any starting point is to visualize the sum $\boldsymbol{v} + \boldsymbol{w}$:

Vector addition (head to tail) ***At the end of*** $\boldsymbol{v}$, ***place the start of*** $\boldsymbol{w}$.

We travel along $\boldsymbol{v}$ and then along $\boldsymbol{w}$. Or we take the shortcut along $\boldsymbol{v} + \boldsymbol{w}$. We could also go along $\boldsymbol{w}$ and then $\boldsymbol{v}$. In other words, $\boldsymbol{w} + \boldsymbol{v}$ gives the same answer as $\boldsymbol{v} + \boldsymbol{w}$. These are different ways along the parallelogram (in this example it is a rectangle). The endpoint in Figure 1.2 is the diagonal $\boldsymbol{v} + \boldsymbol{w}$ which is also $\boldsymbol{w} + \boldsymbol{v}$.

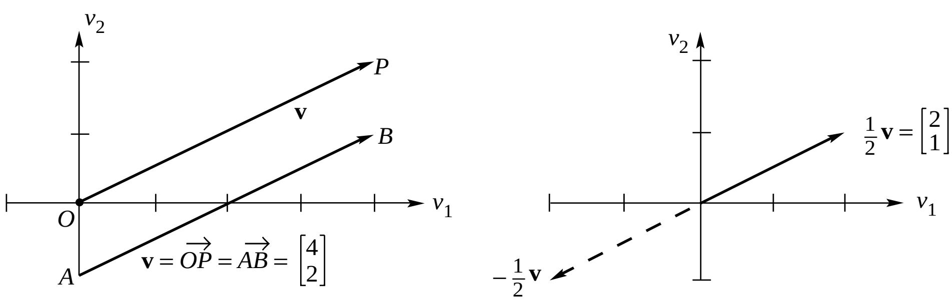

Figure 1.1 The arrow usually starts at the origin $(0, 0)$; $c\boldsymbol{v}$ is always parallel to $\boldsymbol{v}$.

LENGTHS AND DOT PRODUCTS ■ 1.2

The first section mentioned multiplication of vectors, but it backed off. Now we go forward to define the "*dot product*" of $\boldsymbol{v}$ and $\boldsymbol{w}$. This multiplication involves the separate products v_1w_1 and v_2w_2, but it doesn't stop there. Those two numbers are added to produce the single number $\boldsymbol{v} \cdot \boldsymbol{w}$.

DEFINITION The ***dot product*** or ***inner product*** of $\boldsymbol{v} = (v_1, v_2)$ and $\boldsymbol{w} = (w_1, w_2)$ is the number

$$\boldsymbol{v} \cdot \boldsymbol{w} = v_1w_1 + v_2w_2. \tag{1}$$

Example 1 The vectors $\boldsymbol{v} = (4, 2)$ and $\boldsymbol{w} = (-1, 2)$ have a *zero* dot product:

$$\begin{bmatrix} 4 \\ 2 \end{bmatrix} \cdot \begin{bmatrix} -1 \\ 2 \end{bmatrix} = -4 + 4 = 0.$$

In mathematics, zero is always a special number. For dot products, it means that *these two vectors are perpendicular*. The angle between them is 90°. When we drew them in Figure 1.2, we saw a rectangle (not just any parallelogram). The clearest example of perpendicular vectors is $\boldsymbol{i} = (1, 0)$ along the x axis and $\boldsymbol{j} = (0, 1)$ up the y axis. Again the dot product is $\boldsymbol{i} \cdot \boldsymbol{j} = 0 + 0 = 0$. Those vectors $\boldsymbol{i}$ and $\boldsymbol{j}$ form a right angle.

The vectors $\boldsymbol{v} = (1, 2)$ and $\boldsymbol{w} = (2, 1)$ are not perpendicular. Their dot product is 4. Soon that will reveal the angle between them (not 90°).

Example 2 Put a weight of 4 at the point $x = -1$ and a weight of 2 at the point $x = 2$. If the x axis is a see-saw, it will balance on the center point $x = 0$. The weights balance because the dot product is $(4)(-1) + (2)(2) = 0$.

This example is typical of engineering and science. The vector of weights is $(w_1, w_2) = (4, 2)$. The vector of distances from the center is $(v_1, v_2) = (-1, 2)$. The force times the distance, w_1 times v_1, gives the "moment" of the first weight. The equation for the see-saw to balance is $w_1v_1 + w_2v_2 = 0$.

The dot product $\boldsymbol{w} \cdot \boldsymbol{v}$ ***equals*** $\boldsymbol{v} \cdot \boldsymbol{w}$. The order of $\boldsymbol{v}$ and $\boldsymbol{w}$ makes no difference.

Example 3 Dot products enter in economics and business. We have five products to buy and sell. Their prices are $(p_1, p_2, p_3, p_4, p_5)$ for each unit—this is the "price vector" $\boldsymbol{p}$. The quantities we buy or sell are $(q_1, q_2, q_3, q_4, q_5)$—positive when we sell, negative when we buy. Selling q_1 units of the first product at the price p_1 brings in q_1p_1. The total income is the dot product $\boldsymbol{q} \cdot \boldsymbol{p}$:

$$\textbf{\textit{Income}} = (q_1, q_2, \ldots, q_5) \cdot (p_1, p_2, \ldots, p_5) = q_1p_1 + q_2p_2 + \cdots + q_5p_5.$$

■ REVIEW OF THE KEY IDEAS ■

1. The dot product $\boldsymbol{v} \cdot \boldsymbol{w}$ multiplies each component v_i by w_i and adds.

2. The length $\|\boldsymbol{v}\|$ is the square root of $\boldsymbol{v} \cdot \boldsymbol{v}$.

3. The vector $\boldsymbol{v}/\|\boldsymbol{v}\|$ is a ***unit vector***. Its length is 1.

4. The dot product is $\boldsymbol{v} \cdot \boldsymbol{w} = 0$ when $\boldsymbol{v}$ and $\boldsymbol{w}$ are perpendicular.

5. The cosine of θ (the angle between $\boldsymbol{v}$ and $\boldsymbol{w}$) never exceeds 1:

$$\cos\theta = \frac{\boldsymbol{v} \cdot \boldsymbol{w}}{\|\boldsymbol{v}\|\|\boldsymbol{w}\|} \quad \text{and therefore} \quad |\boldsymbol{v} \cdot \boldsymbol{w}| \leq \|\boldsymbol{v}\|\,\|\boldsymbol{w}\|.$$

Problem Set 1.2

1 Calculate the dot products $\boldsymbol{u} \cdot \boldsymbol{v}$ and $\boldsymbol{u} \cdot \boldsymbol{w}$ and $\boldsymbol{v} \cdot \boldsymbol{w}$ and $\boldsymbol{w} \cdot \boldsymbol{v}$:

$$\boldsymbol{u} = \begin{bmatrix} -.6 \\ .8 \end{bmatrix} \qquad \boldsymbol{v} = \begin{bmatrix} 3 \\ 4 \end{bmatrix} \qquad \boldsymbol{w} = \begin{bmatrix} 4 \\ 3 \end{bmatrix}.$$

2 Compute the lengths $\|\boldsymbol{u}\|$ and $\|\boldsymbol{v}\|$ and $\|\boldsymbol{w}\|$ of those vectors. Check the Schwarz inequalities $|\boldsymbol{u} \cdot \boldsymbol{v}| \leq \|\boldsymbol{u}\|\,\|\boldsymbol{v}\|$ and $|\boldsymbol{v} \cdot \boldsymbol{w}| \leq \|\boldsymbol{v}\|\,\|\boldsymbol{w}\|$.

3 Write down unit vectors in the directions of $\boldsymbol{v}$ and $\boldsymbol{w}$ in Problem 1. Find the cosine of the angle θ between them.

4 Find unit vectors $\boldsymbol{u}_1$ and $\boldsymbol{u}_2$ in the directions of $\boldsymbol{v} = (3, 1)$ and $\boldsymbol{w} = (2, 1, 2)$. Also find unit vectors $\boldsymbol{U}_1$ and $\boldsymbol{U}_2$ that are perpendicular to $\boldsymbol{v}$ and $\boldsymbol{w}$.

5 For any unit vectors $\boldsymbol{v}$ and $\boldsymbol{w}$, show that the angle is 0° or 90° or 180° between

(a) $\boldsymbol{v}$ and $\boldsymbol{v}$ (b) $\boldsymbol{w}$ and $-\boldsymbol{w}$ (c) $\boldsymbol{v} + \boldsymbol{w}$ and $\boldsymbol{v} - \boldsymbol{w}$.

6 Find the angle θ (from its cosine) between

(a) $\boldsymbol{v} = \begin{bmatrix} 1 \\ \sqrt{3} \end{bmatrix}$ and $\boldsymbol{w} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ (b) $\boldsymbol{v} = \begin{bmatrix} 2 \\ 2 \\ -1 \end{bmatrix}$ and $\boldsymbol{w} = \begin{bmatrix} 2 \\ -1 \\ 2 \end{bmatrix}$

(c) $\boldsymbol{v} = \begin{bmatrix} 1 \\ \sqrt{3} \end{bmatrix}$ and $\boldsymbol{w} = \begin{bmatrix} -1 \\ \sqrt{3} \end{bmatrix}$ (d) $\boldsymbol{v} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$ and $\boldsymbol{w} = \begin{bmatrix} -1 \\ -2 \end{bmatrix}$.

7 (a) Describe every vector (w_1, w_2) that is perpendicular to $\boldsymbol{v} = (2, -1)$.

(b) In words, describe all vectors that are perpendicular to $\boldsymbol{V} = (1, 1, 1)$.

8 True or false (give a reason if true or a counterexample if false):

TRANSPOSES AND PERMUTATIONS ■ 2.7

We need one more matrix, and fortunately it is much simpler than the inverse. It is the ***"transpose"*** of A, which is denoted by A^{T}. *Its columns are the rows of* A. When A is an m by n matrix, the transpose is n by m:

$$\text{If} \quad A = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 4 \end{bmatrix} \quad \text{then} \quad A^{\mathrm{T}} = \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 3 & 4 \end{bmatrix}.$$

You can write the rows of A into the columns of A^{T}. Or you can write the columns of A into the rows of A^{T}. The matrix "flips over" its main diagonal. The entry in row i, column j of A^{T} comes from row j, column i of the original A:

$$(A^{\mathrm{T}})_{ij} = A_{ji}.$$

The transpose of a lower triangular matrix is upper triangular. The transpose of A^{T} is—-.

Note MATLAB's symbol for the transpose is A'. To enter a column vector, type a row and transpose: $v = [\,1\ 2\ 3\,]'$. To enter a matrix with second column $w = [\,4\ 5\ 6\,]'$ you could define $M = [\,v\ \ w\,]$. Quicker to enter by rows and then transpose the whole matrix: $M = [1\ 2\ 3;\ 4\ 5\ 6]'$.

The rules for transposes are very direct. We can transpose $A+B$ to get $(A+B)^{\mathrm{T}}$. Or we can transpose A and B separately, and then add $A^{\mathrm{T}} + B^{\mathrm{T}}$—same result. The serious questions are about the transpose of a product AB and an inverse A^{-1}:

$$\text{The transpose of} \quad A+B \quad \text{is} \quad A^{\mathrm{T}} + B^{\mathrm{T}}. \tag{1}$$

$$\text{The transpose of} \quad AB \quad \text{is} \quad (AB)^{\mathrm{T}} = B^{\mathrm{T}}A^{\mathrm{T}}. \tag{2}$$

$$\text{The transpose of} \quad A^{-1} \quad \text{is} \quad (A^{-1})^{\mathrm{T}} = (A^{\mathrm{T}})^{-1}. \tag{3}$$

Notice especially how $B^{\mathrm{T}}A^{\mathrm{T}}$ comes in reverse order like $B^{-1}A^{-1}$. The proof for the inverse is quick: $B^{-1}A^{-1}$ times AB produces I. To see this reverse order for $(AB)^{\mathrm{T}}$, start with $(A\boldsymbol{x})^{\mathrm{T}}$:

$A\boldsymbol{x}$ combines the columns of A while $\boldsymbol{x}^{\mathrm{T}}A^{\mathrm{T}}$ combines the rows of A^{T}.

It is the same combination of the same vectors! In A they are columns, in A^{T} they are rows. So the transpose of the column $A\boldsymbol{x}$ is the row $\boldsymbol{x}^{\mathrm{T}}A^{\mathrm{T}}$. That fits our formula $(A\boldsymbol{x})^{\mathrm{T}} = \boldsymbol{x}^{\mathrm{T}}A^{\mathrm{T}}$. Now we can prove the formulas for $(AB)^{\mathrm{T}}$ and $(A^{-1})^{\mathrm{T}}$.

When $B = [\boldsymbol{x}_1\ \ \boldsymbol{x}_2]$ has two columns, apply the same idea to each column. The columns of AB are $A\boldsymbol{x}_1$ and $A\boldsymbol{x}_2$. Their transposes are $\boldsymbol{x}_1^{\mathrm{T}}A^{\mathrm{T}}$ and $\boldsymbol{x}_2^{\mathrm{T}}A^{\mathrm{T}}$. Those are the rows of $B^{\mathrm{T}}A^{\mathrm{T}}$:

$$\text{Transposing} \quad AB = \begin{bmatrix} A\boldsymbol{x}_1 & A\boldsymbol{x}_2 & \cdots \end{bmatrix} \quad \text{gives} \quad \begin{bmatrix} \boldsymbol{x}_1^{\mathrm{T}}A^{\mathrm{T}} \\ \boldsymbol{x}_2^{\mathrm{T}}A^{\mathrm{T}} \\ \vdots \end{bmatrix} \quad \text{which is} \quad B^{\mathrm{T}}A^{\mathrm{T}}. \tag{4}$$

INDEPENDENCE, BASIS AND DIMENSION ■ 3.5

This important section is about the true size of a subspace. There are n columns in an m by n matrix, and each column has m components. But the true "dimension" of the column space is not necessarily m or n. The dimension is measured by counting ***independent columns***—and we have to say what that means. For this particular subspace (the column space) we will see that *the true dimension is the rank* r.

The idea of independence applies to any vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$ in any vector space. Most of this section concentrates on the subspaces that we know and use—especially the column space and nullspace. In the last part we also study "vectors" that are not column vectors. These vectors can be matrices and functions; they can be linearly independent (or not). First come the key examples using column vectors, before the extra examples with matrices and functions.

The final goal is to understand a ***basis*** for a vector space. We are at the heart of our subject, and we cannot go on without a basis. The four essential ideas in this section are:

1. **Independent vectors**
2. **Spanning a space**
3. **Basis for a space**
4. **Dimension of a space.**

Linear Independence

Our first definition of independence is not so conventional, but you are ready for it.

DEFINITION The columns of A are *linearly independent* when the only solution to $A\boldsymbol{x} = \boldsymbol{0}$ is $\boldsymbol{x} = \boldsymbol{0}$. ***No other combination*** $A\boldsymbol{x}$ ***of the columns gives the zero vector***.

With linearly independent columns, the nullspace $\boldsymbol{N}(A)$ contains only the zero vector. Let me illustrate linear independence (and linear dependence) with three vectors in $\mathbf{R}^3$:

1. Figure 3.4 (left) shows three vectors *not* in the same plane. They are independent.

2. Figure 3.4 (right) shows three vectors *in the same plane*. They are dependent.

This idea of independence applies to 7 vectors in 12-dimensional space. If they are the columns of A, and independent, the nullspace only contains $\boldsymbol{x} = \boldsymbol{0}$. Now we choose different words to express the same idea. The following definition of independence will apply to any sequence of vectors in any vector space. When the vectors are the columns of A, the two definitions say exactly the same thing.

Example 5

$$A = \begin{bmatrix} 1 & 4 \\ 2 & 7 \\ 3 & 5 \end{bmatrix} \text{ and } A^{\mathrm{T}} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 7 & 5 \end{bmatrix}. \text{ Here } m = 3 \text{ and } n = 2.$$

The column space of A is spanned by the two columns of A. It is a plane in $\mathbf{R}^3$. The row space of A is spanned by the three rows of A (columns of A^{T}). It is all of $\mathbf{R}^2$. Remember: The rows are in $\mathbf{R}^n$. The columns are in $\mathbf{R}^m$. Same numbers, different vectors, different spaces.

A Basis for a Vector Space

In the xy plane, a set of independent vectors could be small—just one single vector. A set that spans the xy plane could be large—three vectors, or four, or infinitely many. One vector won't span the plane. Three vectors won't be independent. A "***basis***" for the plane is just right.

DEFINITION A ***basis*** for a vector space is a sequence of vectors that has two properties at once:

1. The vectors are ***linearly independent***.
2. The vectors ***span the space***.

This combination of properties is fundamental to linear algebra. Every vector $\boldsymbol{v}$ in the space is a combination of the basis vectors, because they span the space. More than that, the combination that produces $\boldsymbol{v}$ is *unique*, because the basis vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$ are independent:

There is one and only one way to write $\boldsymbol{v}$ as a combination of the basis vectors.

Reason: Suppose $\boldsymbol{v} = a_1\boldsymbol{v}_1 + \cdots + a_n\boldsymbol{v}_n$ and also $\boldsymbol{v} = b_1\boldsymbol{v}_1 + \cdots b_n\boldsymbol{v}_n$. By subtraction $(a_1 - b_1)\boldsymbol{v}_1 + \cdots + (a_n - b_n)\boldsymbol{v}_n$ is the zero vector. From the independence of the $\boldsymbol{v}$'s, each $a_i - b_i = 0$. Hence $a_i = b_i$.

Example 6 The columns of $I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ are a basis for $\mathbf{R}^2$. This is the "standard basis".

$$\text{The vectors } \boldsymbol{i} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \text{ and } \boldsymbol{j} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \text{ are independent. They span } \mathbf{R}^2.$$

We will summarize this section so far. We have four ways to recognize a positive definite matrix. Right now it is only 2 by 2.

6N When a 2 by 2 symmetric matrix has one of these four properties, it has them all:

1. Both of the eigenvalues are positive.

2. The 1 by 1 and 2 by 2 determinants are positive: $a > 0$ and $ac - b^2 > 0$.

3. The pivots are positive: $a > 0$ and $(ac - b^2)/a > 0$.

4. The function $\boldsymbol{x}^{\mathrm{T}} A \boldsymbol{x} = ax^2 + 2bxy + cy^2$ is positive except at $(0, 0)$.

When A has one (therefore all) of these four properties, it is a ***positive definite matrix***.

Note We deal only with symmetric matrices. The cross derivative $\partial^2 f/\partial x \partial y$ always equals $\partial^2 f/\partial y \partial x$. For $f(x, y, z)$ the nine second derivatives fill a symmetric 3 by 3 matrix. It is positive definite when the three pivots (and the three eigenvalues, and the three determinants) are positive.

Example 3 Is $f(x, y) = x^2 + 8xy + 3y^2$ everywhere positive—except at $(0, 0)$?

Solution The second derivatives are $f_{xx} = 2$ and $f_{xy} = f_{yx} = 8$ and $f_{yy} = 6$, all positive. But the test is *not* positive derivatives. We look for positive *definiteness*. The answer is *no*, this function is not always positive. By trial and error we locate a point $x = 1$, $y = -1$ where $f(1, -1) = 1 - 8 + 3 = -4$. Better to do linear algebra, and apply the exact tests to the matrix that produced $f(x, y)$:

$$x^2 + 8xy + 3y^2 = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & 4 \\ 4 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}.$$

The matrix has $ac - b^2 = 3 - 16$. The pivots are 1 and -13. The eigenvalues are ______ (we don't need them). The matrix is not positive definite.

Note how $8xy$ comes from $a_{12} = 4$ above the diagonal and $a_{21} = 4$ symmetrically below. That matrix multiplication in $\boldsymbol{x}^T A \boldsymbol{x}$ makes the function appear.

Main point The sign of b is not the essential thing. The cross derivative $\partial^2 f/\partial x \partial y$ can be positive or negative—it is b^2 that enters the tests. The *size* of b, compared to a and c, decides whether A is positive definite and the function has a minimum.