

Positive Definite Matrices

A quadratic form $q(\mathbf{x})$ (in the n real variables $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$) is said to be

- **positive definite** if $q(\mathbf{x}) > 0$ for each $\mathbf{x} \neq \mathbf{0}$ in \mathbb{R}^n .
- **positive semidefinite** if $q(\mathbf{x}) \geq 0$ for each $\mathbf{x} \neq \mathbf{0}$ in \mathbb{R}^n .
- **negative definite** if $q(\mathbf{x}) < 0$ for each $\mathbf{x} \neq \mathbf{0}$ in \mathbb{R}^n .
- **negative semidefinite** if $q(\mathbf{x}) \leq 0$ for each $\mathbf{x} \neq \mathbf{0}$ in \mathbb{R}^n .
- **indefinite** if there exist $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ such that $q(\mathbf{u}) > 0$ and $q(\mathbf{v}) < 0$.

To any (real) quadratic form q there is an associated real symmetric matrix \mathbf{A} for which $q(\mathbf{x}) = \langle \mathbf{x}, \mathbf{Ax} \rangle = \langle \mathbf{Ax}, \mathbf{x} \rangle = \mathbf{x}^T \mathbf{Ax}$. We apply the same words to characterize this symmetric matrix, calling it positive/negative (semi)definite or indefinite depending on which of the above conditions hold for the quadratic form $q(\mathbf{x}) = \langle \mathbf{x}, \mathbf{Ax} \rangle$.

Note that, while it would seem, in classifying a quadratic form q , one must investigate the behavior of $q(\mathbf{x})$ over all $\mathbf{x} \in \mathbb{R}^n$, it is enough to focus on those $\mathbf{x} \in \mathbb{R}^n$ with unit length $\|\mathbf{x}\| = 1$. This is because, for $\mathbf{x} \neq \mathbf{0}$,

$$q(\mathbf{x}) = \mathbf{x}^T \mathbf{Ax} = \|\mathbf{x}\|^2 \left(\frac{\mathbf{x}}{\|\mathbf{x}\|} \right)^T \mathbf{A} \left(\frac{\mathbf{x}}{\|\mathbf{x}\|} \right) = \|\mathbf{x}\|^2 q(\mathbf{u}),$$

where $\mathbf{u} = \mathbf{x}/\|\mathbf{x}\|$ is a unit vector. This means q is positive definite, if and only if $q(\mathbf{u}) > 0$ for every unit vector $\mathbf{u} \in \mathbb{R}^n$, positive semidefinite, if and only if $q(\mathbf{u}) \geq 0$ for every unit vector $\mathbf{u} \in \mathbb{R}^n$, and so on.

Now, because the matrix \mathbf{A} associated to q is symmetric, the Spectral Theorem says there exists an orthonormal basis $\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$ of \mathbb{R}^n consisting of eigenvectors of \mathbf{A} . Without loss of generality, we may assume the vectors of this basis have been indexed so that the corresponding (real) eigenvalues are ordered from smallest to largest

$$\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n.$$

The orthogonal matrix \mathbf{S} whose j^{th} column is \mathbf{q}_j diagonalizes \mathbf{A} :

$$\mathbf{A} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^T, \quad \text{with} \quad \mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix}.$$

Thus, for any vector $\mathbf{x} \in \mathbb{R}^n$ with $\|\mathbf{x}\| = 1$,

$$q(\mathbf{x}) = \langle \mathbf{x}, \mathbf{Ax} \rangle = \langle \mathbf{x}, \mathbf{S}\mathbf{\Lambda}\mathbf{S}^T \mathbf{x} \rangle = \langle \mathbf{S}^T \mathbf{x}, \mathbf{\Lambda}\mathbf{S}^T \mathbf{x} \rangle = \langle \mathbf{\Lambda}\mathbf{S}^T \mathbf{x}, \mathbf{S}^T \mathbf{x} \rangle = \langle \mathbf{\Lambda}\mathbf{y}, \mathbf{y} \rangle = \sum_{j=1}^n \lambda_j y_j^2,$$

where $\mathbf{y} = \mathbf{S}^T \mathbf{x}$, and $\|\mathbf{y}\| = 1$, as well (since \mathbf{S}^T is orthogonal). From this, we deduce several facts:

Fact 1: For a real symmetric matrix \mathbf{A} with eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$,

- (i) $\lambda_1 \leq \langle \mathbf{x}, \mathbf{A}\mathbf{x} \rangle \leq \lambda_n$, for each unit vector \mathbf{x} (i.e., $\|\mathbf{x}\| = 1$).
- (ii) Corresponding to each eigenvalue λ_j of \mathbf{A} , there is a unit eigenvector \mathbf{x}_j (that is, $\|\mathbf{x}_j\| = 1$ with $\mathbf{A}\mathbf{x}_j = \lambda_j \mathbf{x}_j$), and for this \mathbf{x}_j ,

$$q(\mathbf{x}_j) = \langle \mathbf{x}_j, \mathbf{A}\mathbf{x}_j \rangle = \langle \lambda_j \mathbf{x}_j, \mathbf{x}_j \rangle = \lambda_j \langle \mathbf{x}_j, \mathbf{x}_j \rangle = \lambda_j.$$

- (iii) \mathbf{A} is positive semidefinite if and only if each eigenvalue $\lambda_j \geq 0$, and positive definite iff every eigenvalue is positive. Similarly, \mathbf{A} is negative semidefinite iff every eigenvalue $\lambda_j \leq 0$, and negative definite iff they are all negative.
- (iv) \mathbf{A} is indefinite if and only if $\lambda_1 < 0 < \lambda_n$.

The role of definiteness in optimization

Earlier, we said a real-valued function f of multiple real variables $\mathbf{x} = (x_1, \dots, x_n)$ which is smooth about the point $\mathbf{a} = (a_1, \dots, a_n)$ is, for small enough $\|\mathbf{h}\|$, well approximated by the second-degree Taylor polynomial

$$f(\mathbf{a} + \mathbf{h}) \approx f(\mathbf{a}) + \nabla f(\mathbf{a}) \cdot \mathbf{h} + \left\langle \mathbf{h}, \frac{1}{2} H_f(\mathbf{a}) \mathbf{h} \right\rangle. \quad (1)$$

As with the functions of a single variable studied in Calculus I, a necessary condition for differentiable f to have an extremum at \mathbf{a} is that \mathbf{a} be a critical point, a site where the 1st derivative is zero. A differentiable function $f(\mathbf{x})$ has many first *partial* derivatives, and for f to have an extremum, they must all be zero simultaneously. This means that \mathbf{a} is a critical point when $\nabla f(\mathbf{a}) = \mathbf{0}$. Thus, in the neighborhood of a critical point \mathbf{a} , (1) reduces to

$$f(\mathbf{a} + \mathbf{h}) \approx f(\mathbf{a}) + \left\langle \mathbf{h}, \frac{1}{2} H_f(\mathbf{a}) \mathbf{h} \right\rangle \quad (2)$$

for small $\|\mathbf{h}\|$. We have that $\langle \mathbf{h}, \frac{1}{2} H_f(\mathbf{a}) \mathbf{h} \rangle$ is a quadratic form. If it is the case that $\langle \mathbf{h}, H_f(\mathbf{a}) \mathbf{h} \rangle$ is positive definite, then the thing added to $f(\mathbf{a})$ on the right-hand side of (2) is positive for any nonzero vector \mathbf{h} representing the magnitude and direction we have "strayed" from \mathbf{a} ; this, in turn, means that $f(\mathbf{a})$ is a local minimum of f . By similar reasoning, $f(\mathbf{a})$ is a local maximum if $H_f(\mathbf{a})$ is a negative definite matrix. If $H_f(\mathbf{a})$ is *indefinite*, then $f(\mathbf{a})$ is neither a local max nor a local min; in that case, $\mathbf{x} = \mathbf{a}$ is called a **saddle point**. We summarize these findings:

Theorem 1 (2nd Derivative Test for Functions of Multiple Variables): Suppose $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is a smooth function in an open neighborhood of $\mathbf{a} \in \mathbb{R}^n$. If \mathbf{a} is a critical point, then

- (i) $f(\mathbf{a})$ is a local minimum if $H_f(\mathbf{a})$ is positive definite.
- (ii) $f(\mathbf{a})$ is a local maximum if $H_f(\mathbf{a})$ is negative definite.
- (iii) $f(\mathbf{a})$ is a saddle point if $H_f(\mathbf{a})$ is indefinite.

Note that we have not stated a conclusion if $H_f(\mathbf{a})$ is only positive or negative *semidefinite*.

We must be careful not to make too much of the approximation (2) to f around a critical point \mathbf{a} . With any $\mathbf{h} \neq \mathbf{0}$, you have strayed from the point \mathbf{a} , and the two sides of (2) are (generally) unequal. But for a function f whose 2nd partial derivatives are continuous throughout an open neighborhood containing \mathbf{a} , they are enough equal for small $\|\mathbf{h}\|$ that an upswing in the term $(1/2)\langle \mathbf{h}, H_f(\mathbf{a})\mathbf{h} \rangle$ (i.e., the effects of $H_f(\mathbf{a})$ being positive definite) overrides any other effects in a small neighborhood around \mathbf{a} , and makes \mathbf{a} the site of a local minimum.

Other tests of definiteness

The eigenvalues tell us about the definiteness, or lack thereof, of a symmetric matrix \mathbf{A} . It can be time-consuming, however, to compute the eigenvalues of a matrix. We seek another method by which we may determine if a symmetric matrix is (semi)definite. This new method involves the calculation of *upper left determinants*—i.e., determinants of submatrices emanating from the upper left corner of a matrix. Given an n -by- n (symmetric) matrix $\mathbf{A} = (a_{ij})$, we define the relevant n determinants $\Delta_1, \Delta_2, \dots, \Delta_n$ to be

$$\Delta_1 = a_{11}, \quad \Delta_2 = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix}, \quad \Delta_3 = \begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix}, \quad \dots, \quad \Delta_n = \det(\mathbf{A}).$$

In the text, Strang notes that a symmetric matrix \mathbf{A} is positive definite iff each $\Delta_j > 0$. This is known as **Sylvester's Criterion**. In fact, the same upper left determinants can be used to tell that a matrix is negative definite. We state and prove the theorem below. But first we state and prove a fact relating the determinant of any square matrix to its eigenvalues.

Lemma 1: Suppose $\lambda_1, \dots, \lambda_n$ are the eigenvalues of the n -by- n matrix \mathbf{A} . The $|\mathbf{A}| = \lambda_1 \lambda_2 \cdots \lambda_n$.

Proof: Consider the n^{th} degree characteristic polynomial

$$p_A(\lambda) = \det(\mathbf{A} - \lambda\mathbf{I}) = (-1)^n(\lambda - \lambda_1) \cdots (\lambda - \lambda_n),$$

where $\lambda_1, \dots, \lambda_n$ are the (perhaps nonreal) eigenvalues of \mathbf{A} . Then

$$|\mathbf{A}| = p_A(0) = (-1)^n(0 - \lambda_1) \cdots (0 - \lambda_n) = (-1)^{2n} \prod_{j=1}^n \lambda_j = \prod_{j=1}^n \lambda_j.$$

□

Theorem 2 (Generalized(?) Sylvester's Criterion): Let $q(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x}$ be a quadratic form defined on \mathbb{R}^n with symmetric nonsingular matrix \mathbf{A} . We may conclude that q is

- (i) positive definite if and only if $\Delta_k > 0$ for each $k = 1, \dots, n$. Here, Δ_k is the determinant of the k -by- k submatrix of \mathbf{A} comprising entries a_{ij} with $1 \leq i, j \leq k$,
- (ii) negative definite if and only if $(-1)^k \Delta_k > 0$ for each $k = 1, \dots, n$.
- (iii) indefinite if neither of the previous two conditions is satisfied.

The proof, which you may read at your own behest (it continues on through the end of the next page), follows. No doubt it is the most technical proof that has been given in the course.

Proof: Each "if and only if" statement requires a proof of two statements. We begin with the "iff" statement in (i), focusing first on the assertion that $\Delta_k > 0$ for each k implies \mathbf{A} is positive definite. The proof is by induction on n , the size of the matrix. When $n = 1$ (so the matrix has just one entry, a), then trivially $\Delta_1 = a > 0$ implies

$$xax = ax^2 > 0, \quad \text{for all real } x \neq 0.$$

For our induction hypothesis, assume \mathbf{A} is n -by- n for some $n > 1$, and that the claim holds for each quadratic form defined on \mathbb{R}^{n-1} . Now, since $0 < \Delta_n = \lambda_1 \lambda_2 \cdots \lambda_n$, the eigenvalues of \mathbf{A} are all nonzero. If they are all positive, then we know $q(\mathbf{x}) = \langle \mathbf{A}\mathbf{x}, \mathbf{x} \rangle$ is positive definite. So, let us assume that some eigenvalue $\lambda_i < 0$. But for their product to be positive, there must be an even number of negative eigenvalues, so let λ_j be another negative eigenvalue, and take $\mathbf{v}_i, \mathbf{v}_j$ to be eigenvectors of unit length corresponding to λ_i, λ_j , respectively. We may assume \mathbf{v}_i and \mathbf{v}_j are orthogonal to each other, either because $\lambda_i \neq \lambda_j$, or because Gram-Schmidt may be used to select orthogonal basis vectors in an eigenspace. Let $W = \text{span}(\{\mathbf{v}_i, \mathbf{v}_j\})$, a plane in \mathbb{R}^n . Given any $\mathbf{y} \in W$, with $\mathbf{y} = y_i \mathbf{v}_i + y_j \mathbf{v}_j$, we have

$$\begin{aligned} q(\mathbf{y}) &= \langle \mathbf{A}\mathbf{y}, \mathbf{y} \rangle = \langle \lambda_i y_i \mathbf{v}_i + \lambda_j y_j \mathbf{v}_j, y_i \mathbf{v}_i + y_j \mathbf{v}_j \rangle \\ &= \lambda_i y_i^2 \langle \mathbf{v}_i, \mathbf{v}_i \rangle + \lambda_j y_i y_j \langle \mathbf{v}_j, \mathbf{v}_i \rangle + \lambda_i y_i y_j \langle \mathbf{v}_i, \mathbf{v}_j \rangle + \lambda_j y_j^2 \langle \mathbf{v}_j, \mathbf{v}_j \rangle \\ &= \lambda_i y_i^2 + \lambda_j y_j^2 < 0. \end{aligned}$$

Thus, q is negative for all nonzero $\mathbf{y} \in W$. Let $q_*: \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ be the restriction of q to \mathbb{R}^{n-1} , so that $q_*(x_1, \dots, x_{n-1}) := q(x_1, \dots, x_{n-1}, 0)$. Along with that, let $\mathbf{A}_* = (a_{ij})_{1 \leq i, j \leq n-1}$. Then

$$q_*(x_1, \dots, x_{n-1}) = q(x_1, \dots, x_{n-1}, 0) = \begin{bmatrix} x_1 & \cdots & x_{n-1} \end{bmatrix} \mathbf{A}_* \begin{bmatrix} x_1 \\ \vdots \\ x_{n-1} \end{bmatrix}.$$

By assumption, $\Delta_1, \dots, \Delta_{n-1}$ are all positive, so using the induction hypothesis, we get that q_* is positive definite on \mathbb{R}^{n-1} . But the $(n-1)$ -dimensional hyperplane in \mathbb{R}^n consisting of vectors $\mathbf{x} = (x_1, \dots, x_{n-1}, 0)$ and the plane W have at least a line in common, and we have shown that $q(\mathbf{v})$ for a nonzero vector \mathbf{v} from that line is both strictly positive and strictly negative. $\rightarrow\leftarrow$

To prove the converse statement of (i), suppose q is positive definite, and let $m > 0$ be the minimum value of q over the set $\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| = 1\}$. Similar to above, we take $\mathbf{A}_k = (a_{ij})_{1 \leq i, j \leq k}$ and define a function $q_k: \mathbb{R}^k \rightarrow \mathbb{R}$ as the "restriction of q to \mathbb{R}^k ", in the sense that

$$q_k(x_1, \dots, x_k) = q(x_1, \dots, x_k, 0, \dots, 0) = \begin{bmatrix} x_1 & \cdots & x_k \end{bmatrix} \mathbf{A}_k \begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix}.$$

Let μ_1, \dots, μ_k be the eigenvalues of \mathbf{A}_k . We know there is an $\mathbf{x} \in \mathbb{R}^k$ with $\|\mathbf{x}\| = 1$ such that $q_k(\mathbf{x}) = \mu_i$. It follows from $q_k(\mathbf{x}) = q(\mathbf{x}, \mathbf{0})$ that each $\mu_i \geq m$. Thus, by Lemma 1, $\Delta_k = |\mathbf{A}_k| = \mu_1 \cdots \mu_k \geq m^k > 0$.

To prove (ii), define $q^*(\mathbf{x}) = -q(\mathbf{x})$, so that negative definiteness of q is equivalent to positive definiteness of q^* . By part (i), q^* is positive definite if and only if each $0 < \Delta_k^* = |-\mathbf{A}_k| = (-1)^k |\mathbf{A}_k| = (-1)^k \Delta_k$.

To prove (iii), note first that $0 \neq \Delta_n = \lambda_1 \cdots \lambda_n$ means every eigenvalue is nonzero. If all were positive, then part (a) shows each $\Delta_k > 0$; if all were negative, then part (b) implies each $(-1)^k \Delta_k > 0$. Since neither of these conditions holds by supposition, it must be that \mathbf{A} has both positive and negative eigenvalues, and is nondefinite. \square

Example 1:

Classify each of the nonsingular symmetric matrices \mathbf{A} and \mathbf{B} given below according to its type of *definiteness*.

$$\mathbf{A} = \begin{bmatrix} -1 & 1/2 & 1 \\ 1/2 & -1 & 1/2 \\ 1 & 1/2 & 1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 3 & 1 & 5 & 3 \\ 1 & 5 & 2 & 0 \\ 5 & 2 & 10 & 3 \\ 3 & 0 & 3 & 14 \end{bmatrix}.$$

For \mathbf{A} , we have

$$\Delta_1 = -1, \quad \Delta_2 = \begin{vmatrix} -1 & 1/2 \\ 1/2 & -1 \end{vmatrix} = \frac{3}{4}, \quad \Delta_3 = |\mathbf{A}| = \frac{5}{2}.$$

Thus, while \mathbf{A} is nonsingular, neither of the conditions (i) nor (ii) of Theorem 2 hold, making \mathbf{A} *indefinite*.

For \mathbf{B} , we have

$$\Delta_1 = 3, \quad \Delta_2 = \begin{vmatrix} 3 & 1 \\ 1 & 5 \end{vmatrix} = 14, \quad \Delta_3 = \begin{vmatrix} 3 & 1 & 5 \\ 1 & 5 & 2 \\ 5 & 2 & 10 \end{vmatrix} = 23, \quad \Delta_4 = |\mathbf{B}| = 196.$$

Since each of these subdeterminants is positive, \mathbf{B} is *positive definite*. ■

One more result, equivalent to positive definiteness, seems pertinent.

Theorem 3: A symmetric matrix \mathbf{A} is positive definite if and only if $\mathbf{A} = \mathbf{B}^T\mathbf{B}$ for some matrix \mathbf{B} whose columns are all linearly independent.

Proof: Suppose, first, that $\mathbf{A} = \mathbf{B}^T\mathbf{B}$, where the columns of \mathbf{B} are linearly independent. Due to the linear independence of its columns, $\text{Null}(\mathbf{B}^T\mathbf{B}) = \text{Null}(\mathbf{B}) = \{\mathbf{0}\}$, showing that \mathbf{A} is nonsingular. For $\mathbf{x} \neq \mathbf{0}$, we have

$$\langle \mathbf{A}\mathbf{x}, \mathbf{x} \rangle = \langle \mathbf{B}^T\mathbf{B}\mathbf{x}, \mathbf{x} \rangle = \langle \mathbf{B}\mathbf{x}, \mathbf{B}\mathbf{x} \rangle = \|\mathbf{B}\mathbf{x}\|^2 > 0,$$

since \mathbf{x} , being nonzero, is not in $\text{Null}(\mathbf{B})$.

The converse is harder to prove. Here, I rely on another factorization, called the Cholesky factorization, which exists for any real symmetric matrix. This factorization is $\mathbf{A} = \mathbf{R}^T\mathbf{R}$, which gives the result with $\mathbf{B} = \mathbf{R}$. We note, in particular, that were the columns of \mathbf{B} *not* linearly independent, then those in $\mathbf{A} = \mathbf{R}^T\mathbf{R}$ would not be either. \square