

Lecture 5: Curvature, Convexity, and Separation

Lecture 4 used a linear map to approximate a differentiable function near a point. In the one-variable case, we saw that

$$f(x_0 + h) \approx f(x_0) + f'(x_0)h.$$

In this lecture we are interested in the curvature of a function, which is described by a second-order approximation. In the one-variable case, if f is twice continuously differentiable near x_0 , the second-order Taylor approximation adds a quadratic term:

$$f(x_0 + h) \approx f(x_0) + f'(x_0)h + \frac{1}{2}f''(x_0)h^2.$$

When $f'(x_0) = 0$, the linear term vanishes, so first-order information does not tell us whether the graph bends upward or downward. The sign of $f''(x_0)$ does. In several variables, the Hessian H_f generalizes $f''(x_0)$, and the quadratic term becomes $\frac{1}{2}h^\top H_f(x_0)h$.

1 Second-order approximation and curvature

1.1 Hessians and local quadratic models

Definition 1.1 (C^2 function). Let $U \subseteq \mathbb{R}^n$ be open. A function $f : U \rightarrow \mathbb{R}$ is C^2 on U , or *twice continuously differentiable on U* , if each first partial derivative $\partial f / \partial x_i$ is C^1 on U . Equivalently, all second partial derivatives of f exist and are continuous on U .

Definition 1.2 (Hessian). Let $U \subseteq \mathbb{R}^n$ be open and let $f : U \rightarrow \mathbb{R}$ be differentiable. If the gradient $\nabla f : U \rightarrow \mathbb{R}^n$ is differentiable at $x \in U$, the *Hessian* of f at x is the Jacobian of its gradient,

$$H_f(x) = D(\nabla f)(x).$$

It is the $n \times n$ matrix whose (i, j) entry is

$$(H_f(x))_{ij} = \frac{\partial}{\partial x_j} \left(\frac{\partial f}{\partial x_i} \right) (x).$$

Example 1.3 (Computing a Hessian). Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be given by

$$f(x_1, x_2) = x_1^2 x_2.$$

In Lecture 4 we computed

$$\nabla f(x) = \begin{pmatrix} 2x_1 x_2 \\ x_1^2 \end{pmatrix}.$$

The Hessian is the Jacobian of the gradient, so

$$H_f(x) = \begin{pmatrix} 2x_2 & 2x_1 \\ 2x_1 & 0 \end{pmatrix}.$$

Theorem 1.4 (Equality of mixed partials). *Let $U \subseteq \mathbb{R}^n$ be open and let $f : U \rightarrow \mathbb{R}$ be C^2 . Then*

$$\frac{\partial}{\partial x_j} \left(\frac{\partial f}{\partial x_i} \right) (x) = \frac{\partial}{\partial x_i} \left(\frac{\partial f}{\partial x_j} \right) (x)$$

for every $x \in U$ and every i, j . Thus $H_f(x)$ is symmetric.

We use this result without proof. Importantly, it requires continuity of the second partial derivatives.

Theorem 1.5 (Second-order Taylor expansion). *Let $U \subseteq \mathbb{R}^n$ be open, let $f : U \rightarrow \mathbb{R}$ be C^2 , and let $x_0 \in U$. For changes h such that $x_0 + th \in U$ for every $t \in [0, 1]$,*

$$f(x_0 + h) = f(x_0) + \nabla f(x_0) \cdot h + \frac{1}{2} h^\top H_f(x_0) h + r(h), \quad \frac{|r(h)|}{\|h\|^2} \rightarrow 0 \quad \text{as } h \rightarrow 0. \quad (1)$$

We use this result without proof. It has exactly the form of the linear approximation in Lecture 4, with one additional term. The expression

$$f(x_0) + \nabla f(x_0) \cdot h + \frac{1}{2} h^\top H_f(x_0) h$$

is the *quadratic approximation* to $f(x_0 + h)$. The first-order approximation stops after the linear term and has a remainder negligible relative to $\|h\|$. The second-order approximation keeps the quadratic term and has a remainder negligible relative to $\|h\|^2$.

1.2 Quadratic forms and their shapes

The quadratic term of the Taylor approximation has the form $x^\top A x$. We next study expressions of this form, called quadratic forms. They are essentially quadratic polynomials in several variables. Like quadratic functions in one dimension, quadratic forms have a constant Hessian. That allows us to characterize all possible types of curvature in terms of the sign of the quadratic form.

Definition 1.6 (Quadratic form). Let A be a symmetric $n \times n$ matrix. The function

$$Q_A(x) = x^\top A x, \quad x \in \mathbb{R}^n,$$

is the *quadratic form* associated with A .

The restriction to symmetric matrices is without loss of generality. That is, if a quadratic form is associated with a nonsymmetric matrix B , then the same quadratic form is associated with a symmetric matrix $\tilde{B} = (B + B^\top)/2$.

Next, the Hessian of $Q_A = x^\top A x$ equals $2A$. To see this, observe that

$$Q_A(x) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i x_j,$$

so, using the symmetry $a_{ij} = a_{ji}$, the k th component of the gradient and the (k, l) second partial are

$$\frac{\partial Q_A}{\partial x_k}(x) = 2 \sum_{j=1}^n a_{kj} x_j, \quad \frac{\partial}{\partial x_l} \left(\frac{\partial Q_A}{\partial x_k} \right) (x) = 2a_{kl}.$$

In matrix form, $\nabla Q_A(x) = 2Ax$ and

$$H_{Q_A}(x) = 2A,$$

which indeed does not depend on x .

The sign of $x^\top Ax$ (which is a scalar) describes the possible shapes of the curvature of the quadratic form in $\mathbb{R}^n$.

Definition 1.7 (Definiteness). Let A be a symmetric $n \times n$ matrix.

A is <i>positive definite</i>	$\iff x^\top Ax > 0$	for every $x \neq 0$,
A is <i>positive semidefinite</i>	$\iff x^\top Ax \geq 0$	for every x ,
A is <i>negative definite</i>	$\iff x^\top Ax < 0$	for every $x \neq 0$,
A is <i>negative semidefinite</i>	$\iff x^\top Ax \leq 0$	for every x .

The matrix is *indefinite* if its quadratic form takes both positive and negative values.




Figure 1.1: Four representative quadratic forms in two variables. The light blue surface is the graph of Q over a square around the origin; the two thick curves show the values of Q along the x_1 -axis and along the x_2 -axis.

To read the pictures, fix a direction $v \neq 0$ and move through the origin along the line tv . Along this line, the quadratic form becomes the one-variable function $\phi_v(t) = Q_A(tv) = t^2 v^\top Av$, a parabola through the origin with second derivative

$$\phi_v''(t) = 2v^\top Av.$$

The sign of $v^\top Av$ therefore tells us whether this slice bends upward, bends downward, or is flat. Thus definiteness describes every one-dimensional slice through the origin. Positive definite means that every nonzero-direction slice bends upward, so the origin is the unique global minimum; negative definite means that every such slice bends downward, so the origin is the unique global maximum. A positive semidefinite form may be flat in some directions, while an indefinite form bends upward in some directions and downward in others, so the origin is neither a maximum nor a minimum. These are exactly the shapes in Figure 1.1.

1.3 Computational tests for definiteness

Checking definiteness of a matrix A using the definition in Definition 1.7 requires evaluating $x^\top Ax$ for every nonzero $x \in \mathbb{R}^n$; that is not very practical.

Instead, this section introduces two types of tests that are easier to apply. The first uses eigenvalues, and the second uses determinants of submatrices.

For a diagonal matrix

$$D = \begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{pmatrix},$$

the check is immediate:

$$x^\top Dx = \lambda_1 x_1^2 + \lambda_2 x_2^2 + \cdots + \lambda_n x_n^2,$$

a sum of squares weighted by the diagonal entries. If every $\lambda_i > 0$, the sum is positive for every $x \neq 0$, so D is positive definite; if every $\lambda_i \geq 0$, the sum is nonnegative, so D is positive semidefinite; the negative cases flip the inequalities. If some $\lambda_i > 0 > \lambda_j$, then $e_i^\top De_i = \lambda_i > 0$ and $e_j^\top De_j = \lambda_j < 0$, so D is indefinite. The definiteness of a diagonal matrix is thus decided by the signs of its diagonal entries.

Recall the diagonalization theorem of Lecture 3: a matrix A with n independent eigenvectors can be written as $A = PDP^{-1}$, where the columns of P are the eigenvectors and D is the diagonal matrix of the corresponding eigenvalues; we also noted, without proof, that symmetric matrices always have n independent eigenvectors. For symmetric matrices slightly more is true: the eigenvectors can be chosen so that $P^{-1} = P^\top$ — such a P is called *orthogonal* — and the factorization becomes $A = PDP^\top$. This refinement is the *spectral theorem*, which we again use without proof. It reduces any symmetric matrix to the diagonal case: changing variables to $z = P^\top x$ turns $x^\top Ax$ into $z^\top Dz$, and the diagonal computation above applies. This gives the first test.

Theorem 1.8 (Eigenvalue test). *Let A be a symmetric $n \times n$ matrix. Then*

$$\begin{aligned} A \text{ is positive definite} &\iff \text{every eigenvalue of } A \text{ is positive,} \\ A \text{ is positive semidefinite} &\iff \text{every eigenvalue of } A \text{ is nonnegative,} \\ A \text{ is negative definite} &\iff \text{every eigenvalue of } A \text{ is negative,} \\ A \text{ is negative semidefinite} &\iff \text{every eigenvalue of } A \text{ is nonpositive.} \end{aligned}$$

The matrix is indefinite if and only if it has eigenvalues of both signs.

Proof (optional). By the spectral theorem, write $A = PDP^\top$ with $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ and P orthogonal, and put $z = P^\top x$. Then

$$x^\top Ax = x^\top PDP^\top x = z^\top Dz = \sum_{i=1}^n \lambda_i z_i^2.$$

Since P is orthogonal, $x = Pz$, so $x = 0$ if and only if $z = 0$, and as x ranges over $\mathbb{R}^n$ so does z . The quantity $x^\top Ax$ therefore has the same range of signs as $\sum_i \lambda_i z_i^2$ over all z . That sum is positive for every $z \neq 0$ exactly when every $\lambda_i > 0$: the “if” direction is immediate, and taking $z = e_i$ gives λ_i itself, which forces the “only if” direction. The other three cases are the same argument with the inequality changed, and the indefinite case follows by taking $z = e_i$ for eigenvalues of each sign. $\square$

The second test requires no eigenvalues. Instead, it looks at the signs of the determinants of certain submatrices of A .

Definition 1.9 (Principal minors). Let A be an $n \times n$ matrix. For a nonempty index set $I \subseteq \{1, \dots, n\}$, let A_I be the square matrix obtained by retaining the rows and columns whose indices belong to I . The determinant $\det A_I$ is a *principal minor* of order k when I contains k indices. The minors

$$\Delta_k = \det A_{\{1, \dots, k\}}, \quad k = 1, \dots, n,$$

are the *leading principal minors* of A .

In general, a *minor* of A is the determinant of a submatrix obtained by retaining some rows and some columns; a minor is *principal* when the retained row and column index sets are the same set I . For example, for a 3×3 matrix $A = (a_{ij})$, the index set $I = \{1, 3\}$ retains rows 1, 3 and columns 1, 3:

$$A_{\{1,3\}} = \begin{pmatrix} a_{11} & a_{13} \\ a_{31} & a_{33} \end{pmatrix},$$

and $\det A_{\{1,3\}}$ is a principal minor of order 2. Counting nonempty index sets, A has $2^3 - 1 = 7$ principal minors. The three leading ones come from the sets $\{1\}$, $\{1, 2\}$, and $\{1, 2, 3\}$, so they grow from the top-left corner:

$$\Delta_1 = a_{11}, \quad \Delta_2 = \det \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, \quad \Delta_3 = \det A.$$

Definiteness can be read off the signs of these minors.

Theorem 1.10 (Principal-minor tests). Let A be a symmetric $n \times n$ matrix, and let $\Delta_1, \dots, \Delta_n$ be its leading principal minors.

1. A is positive definite if and only if $\Delta_k > 0$ for every $k = 1, \dots, n$.
2. A is negative definite if and only if $(-1)^k \Delta_k > 0$ for every $k = 1, \dots, n$.
3. A is positive semidefinite if and only if every principal minor of A is nonnegative.
4. A is negative semidefinite if and only if every principal minor of order k has sign satisfying

$$(-1)^k \det A_I \geq 0.$$

We use this result without proof. The first two parts are commonly called *Sylvester's criterion*. Notice the difference between the definite and semidefinite cases: definite matrices require only the n leading principal minors, while semidefinite matrices require all $2^n - 1$ nonempty principal minors.

Example 1.11 (The 2×2 tests). Let

$$A = \begin{pmatrix} a & b \\ b & c \end{pmatrix}.$$

Its leading principal minors are a and $ac - b^2$, while all of its principal minors are a , c , and $ac - b^2$. Therefore

$$\begin{aligned} A \text{ is positive definite} &\iff a > 0 \text{ and } ac - b^2 > 0, \\ A \text{ is negative definite} &\iff a < 0 \text{ and } ac - b^2 > 0, \\ A \text{ is positive semidefinite} &\iff a \geq 0, \ c \geq 0, \text{ and } ac - b^2 \geq 0, \\ A \text{ is negative semidefinite} &\iff a \leq 0, \ c \leq 0, \text{ and } ac - b^2 \geq 0. \end{aligned}$$

For the positive-definite case, the first line can also be seen directly by completing the square (possible whenever $a \neq 0$):

$$x^\top Ax = a \left(x_1 + \frac{b}{a}x_2 \right)^2 + \frac{ac - b^2}{a}x_2^2.$$

If $a > 0$ and $ac - b^2 > 0$, both coefficients are positive. Conversely, if A is positive definite, evaluating at $x = e_1$ gives $a > 0$, and evaluating at $x = (-b/a, 1)^\top$ gives $(ac - b^2)/a > 0$.

The next example shows why the semidefinite tests need more than just the leading minors.

Example 1.12 (Weak leading-minor inequalities are insufficient). Let $A = \text{diag}(0, -1)$. Its leading minors satisfy $\Delta_1 = \Delta_2 = 0$, but the principal minor indexed by $I = \{2\}$ is -1 . Thus weak inequalities on the leading minors alone do not establish positive semidefiniteness. The missing principal minor detects the negative direction e_2 , along which $e_2^\top Ae_2 = -1$.

Example 1.13 (A three-variable quadratic form). Let $Q : \mathbb{R}^3 \rightarrow \mathbb{R}$ be given by

$$Q(x_1, x_2, x_3) = -x_1^2 - x_2^2 - x_3^2 + x_1x_2 + x_2x_3.$$

Its Hessian is

$$H_Q = \begin{pmatrix} -2 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -2 \end{pmatrix},$$

whose leading principal minors are

$$\Delta_1 = -2, \quad \Delta_2 = 3, \quad \Delta_3 = -4.$$

Thus $(-1)^k \Delta_k$ equals 2, 3, and 4, respectively, so by Theorem 1.10 H_Q is negative definite. This establishes negative definiteness without computing any eigenvalues.

The tests apply equally when the Hessian entries are not constant.

Example 1.14 (A varying Hessian). Let $f : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}$ be given by $f(x_1, x_2) = \log x_1 + \log x_2$. Differentiating twice,

$$H_f(x) = \begin{pmatrix} -1/x_1^2 & 0 \\ 0 & -1/x_2^2 \end{pmatrix},$$

which is diagonal with negative entries and hence negative definite at every x , by Theorem 1.8. The Hessian varies with x , but its definiteness does not.

In the last example of this section, definiteness depends on a parameter.

Example 1.15 (An interaction term). Fix $b \in \mathbb{R}^2$ and $\gamma \in \mathbb{R}$, and define $q : \mathbb{R}^2 \rightarrow \mathbb{R}$ by

$$q(x) = b \cdot x - (x_1^2 + \gamma x_1 x_2 + x_2^2).$$

Here H_q is constant. Direct differentiation gives

$$H_q = \begin{pmatrix} -2 & -\gamma \\ -\gamma & -2 \end{pmatrix}.$$

The leading principal minors are -2 and $4 - \gamma^2$. Therefore Theorem 1.10 shows that H_q is negative definite when $|\gamma| < 2$. When $|\gamma| = 2$, its two order-one principal minors are negative and its determinant is zero, so it is negative semidefinite. When $|\gamma| > 2$, its determinant is negative, so it is indefinite: symmetry makes both eigenvalues real, and their negative product forces them to have opposite signs. The interaction term can therefore preserve, flatten, or overturn the downward curvature of the two squared terms.

At this point two exact tests are available: definiteness can be read off the signs of the eigenvalues (Theorem 1.8) or off the signs of the principal minors (Theorem 1.10). Often a full classification is unnecessary, and a cheaper observation settles the question. Table 1.1 collects these quick checks for the negative case; reversing every sign gives the positive case. A necessary condition can rule the property out but never establishes it.

Check	Verdict
$a_{ii} \leq 0$ for every i	necessary for negative semidefiniteness, not sufficient
$a_{ii} < 0$ for every i	necessary for negative definiteness, not sufficient
$(-1)^k \Delta_k \geq 0$ for every k	necessary for negative semidefiniteness, not sufficient
$(-1)^k \Delta_k > 0$ for every k	equivalent to negative definiteness
some x has $x^\top A x > 0$	counterexample to negative semidefiniteness
some $x \neq 0$ has $x^\top A x = 0$	counterexample to negative definiteness
negative semidefinite and $\det A \neq 0$	sufficient for negative definiteness

Table 1.1: Quick checks for a symmetric matrix A with entries a_{ij} and leading principal minors Δ_k .

The diagonal checks follow by evaluating the quadratic form at the coordinate vectors, since $e_i^\top A e_i = a_{ii}$. The reverse is not true in the first three rows: for example, $\begin{pmatrix} -1 & 2 \\ 2 & -1 \end{pmatrix}$ passes both diagonal checks yet is indefinite. The last row follows from the eigenvalue test: the eigenvalues are nonpositive and their product $\det A$ is nonzero, so all of them are negative.

2 Convex sets and concave functions

2.1 Convex sets and preservation

The Hessian and Taylor expansion describe a function near one point. Turning their curvature information into a statement about the function everywhere requires moving between arbitrary points along the segment joining them. The domain must contain that segment.

Definition 2.1 (Convex combination and convex set). Let $x, y \in \mathbb{R}^n$.

- A *convex combination* of x and y is a point

$$(1 - t)x + ty, \quad t \in [0, 1],$$

and the set of all of them is the *line segment* joining x and y .

- A set $C \subseteq \mathbb{R}^n$ is *convex* if

$$(1 - t)x + ty \in C \quad \text{for every } x, y \in C \text{ and every } t \in [0, 1].$$

At $t = 0$ the combination is x and at $t = 1$ it is y , so a set is convex exactly when the segment joining any two of its points stays inside it. For example, $\mathbb{R}^n$ itself, every affine subspace, every halfspace $\{x : a \cdot x \leq c\}$, and the orthants $\mathbb{R}_+^n$ and $\mathbb{R}_{++}^n$ are convex. The unit circle $\{x \in \mathbb{R}^2 : \|x\| = 1\}$ is not: it contains e_1 and $-e_1$ but not their midpoint 0. Figure 2.1 shows the definition in $\mathbb{R}^2$.




Figure 2.1: The set on the left is convex: for any two points in the set, the segment joining them is inside the set. The right is not convex: the segment joining the marked points x and y is not entirely in the set.

Rather than verify Definition 2.1 for each new set, it is usually quicker to build the set out of convex pieces. An *affine map* $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ has the form $T(x) = Ax + b$ for an $m \times n$ matrix A and a vector $b \in \mathbb{R}^m$.

Theorem 2.2 (Preservation rules). *The following constructions preserve convexity.*

1. *The intersection of any collection of convex sets is convex.*
2. *The Cartesian product of convex sets is convex.*
3. *The image of a convex set under an affine map is convex.*
4. *The inverse image of a convex set under an affine map is convex.*

Proof (optional). Intersections. Let $C = \bigcap_{\alpha} C_{\alpha}$ with each C_{α} convex, and let $x, y \in C$ and $t \in [0, 1]$. For each α both x and y lie in C_{α} , so Definition 2.1 puts $(1 - t)x + ty$ in C_{α} . Being in every C_{α} , it lies in C .

Products. Convex combinations in a product are taken coordinate block by coordinate block, so the claim follows by applying Definition 2.1 in each factor.

Affine maps. Let $T(x) = Ax + b$. Expanding, and recombining the two copies of b as $(1-t)b + tb$,

$$T((1-t)x + ty) = A((1-t)x + ty) + b = (1-t)T(x) + tT(y), \quad (2)$$

so T carries the segment joining x and y onto the segment joining $T(x)$ and $T(y)$. If C is convex, (2) exhibits every point between $T(x)$ and $T(y)$ as the image of a point of C , so $T(C)$ is convex. If D is convex and $x, y \in T^{-1}(D)$, then (2) puts $T((1-t)x + ty)$ between $T(x)$ and $T(y)$, hence in D . $\square$

A useful consequence of the preservation rules: any finite list of linear equalities and inequalities defines a convex set. The example below shows why.

Example 2.3 (Linear restrictions define a convex set). Let A be $m \times n$, let B be $p \times n$, and let $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^p$. The set

$$C = \{x \in \mathbb{R}^n : Ax \leq b, Bx = c\}$$

is convex, where the inequality is coordinatewise. Each row of these restrictions has the form $a \cdot x = d$ or $a \cdot x \leq d$ for some $a \in \mathbb{R}^n$ and $d \in \mathbb{R}$. The set each restriction defines is the inverse image of the convex set $\{d\}$ or $(-\infty, d]$ under the affine map $x \mapsto a \cdot x$. Part (4) of Theorem 2.2 makes each restriction convex, and part (1) makes their intersection C convex.

2.2 Concavity and superlevel sets

Convexity of a set says that its line segments stay in the domain. Concavity then compares the height of a function along those same segments.

Definition 2.4 (Concave and convex functions). Let $C \subseteq \mathbb{R}^n$ be convex and let $f : C \rightarrow \mathbb{R}$.

- f is *concave* if

$$f((1-t)x + ty) \geq (1-t)f(x) + tf(y) \quad \text{for every } x, y \in C \text{ and every } t \in [0, 1].$$

- f is *strictly concave* if that inequality is strict whenever $x \neq y$ and $t \in (0, 1)$.
- f is *convex* if $-f$ is concave, and *strictly convex* if $-f$ is strictly concave.

The right-hand side is the height at $(1-t)x + ty$ of the chord joining the two points $(x, f(x))$ and $(y, f(y))$ on the graph, so a concave function lies on or above each of its chords. This is the global version of bending downward: it compares any two points, not only nearby ones. Strict concavity additionally forbids the graph from containing a line segment. If $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$, the affine function $f(x) = a \cdot x + b$ satisfies Definition 2.4 with equality, by the computation (2). Thus it is both concave and convex, and neither strictly.




Figure 2.2: The chord inequality of Definition 2.4. At each point $(1-t)x + ty$ of the segment joining x and y , the graph of a concave function lies on or above the chord joining $(x, f(x))$ and $(y, f(y))$. Here f is strictly concave, but establishing that requires checking the inequality for every pair of points $x \neq y$.

Concavity is a statement about the graph. It also imposes structure on the sets of points at which the function is at least as large as a given level.

Definition 2.5 (Superlevel set). Let $C \subseteq \mathbb{R}^n$, let $f : C \rightarrow \mathbb{R}$, and let $\alpha \in \mathbb{R}$. The *superlevel set* of f at level α is

$$S_\alpha(f) = \{x \in C : f(x) \geq \alpha\}.$$

Theorem 2.6 (Superlevel sets of a concave function). *If C is convex and $f : C \rightarrow \mathbb{R}$ is concave, then $S_\alpha(f)$ is convex for every $\alpha \in \mathbb{R}$.*

Proof. Let $x, y \in S_\alpha(f)$ and $t \in [0, 1]$, and put $z = (1-t)x + ty$. Convexity of C puts z in C by Definition 2.1, and Definition 2.4 gives

$$f(z) \geq (1-t)f(x) + tf(y) \geq (1-t)\alpha + t\alpha = \alpha,$$

the second inequality because $f(x) \geq \alpha$ and $f(y) \geq \alpha$ by Definition 2.5. Hence $z \in S_\alpha(f)$. $\square$

2.3 Quasiconcavity

The proof of Theorem 2.6 used concavity only to keep $f(z)$ above α , never to keep it above the weighted average. Demanding just the conclusion gives a weaker property, and it is the one that many results actually need.

Definition 2.7 (Quasiconcavity). Let $C \subseteq \mathbb{R}^n$ be convex. A function $f : C \rightarrow \mathbb{R}$ is *quasiconcave* if $S_\alpha(f)$ is convex for every $\alpha \in \mathbb{R}$.

In one variable, the convex sets are the intervals, so quasiconcavity requires each superlevel set to be an interval. Figure 2.3 shows a single-peaked function that passes this test and a two-peaked function that fails it.




Figure 2.3: Superlevel sets in one variable. The single-peaked function on the left is quasiconcave: at each level α , the superlevel set $S_\alpha(f)$ is an interval. The two-peaked function on the right is not: at the level shown, its superlevel set is a union of two disjoint intervals, which is not convex.

Quasiconcavity can also be stated as an inequality along segments, parallel to Definition 2.4 but with the minimum in place of the weighted average.

Theorem 2.8 (Minimum characterization). *Let $C \subseteq \mathbb{R}^n$ be convex. A function $f : C \rightarrow \mathbb{R}$ is quasiconcave if and only if*

$$f((1-t)x + ty) \geq \min\{f(x), f(y)\} \quad \text{for every } x, y \in C \text{ and every } t \in [0, 1].$$

Proof. Convex superlevel sets imply the inequality. Let $x, y \in C$ and $t \in [0, 1]$, and put $\alpha = \min\{f(x), f(y)\}$. Then $x, y \in S_\alpha(f)$ by Definition 2.5, and that set is convex by Definition 2.7, so $(1-t)x + ty \in S_\alpha(f)$, which is the inequality.

The inequality implies convex superlevel sets. Fix α and let $x, y \in S_\alpha(f)$ and $t \in [0, 1]$. Then $(1-t)x + ty \in C$ by Definition 2.1, and

$$f((1-t)x + ty) \geq \min\{f(x), f(y)\} \geq \alpha,$$

so the point lies in $S_\alpha(f)$. □

Concavity implies quasiconcavity, since a weighted average of $f(x)$ and $f(y)$ is at least their minimum, so Definition 2.4 gives the inequality of Theorem 2.8. The converse fails.

Example 2.9 (Quasiconcave but not concave). Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be given by $f(x) = x^3$. Since f is increasing, each superlevel set is the interval $[\alpha^{1/3}, \infty)$, which is convex, so f is quasiconcave by Definition 2.7. It is not concave: taking the midpoint of 0 and 2,

$$f(1) = 1 < 4 = \frac{1}{2}f(0) + \frac{1}{2}f(2),$$

which violates Definition 2.4.

What made Example 2.9 work is that x^3 is an increasing relabelling of the values of x , and quasiconcavity, unlike concavity, survives any such relabelling.

Theorem 2.10 (Increasing transformations). *Let $C \subseteq \mathbb{R}^n$ be convex, let $f : C \rightarrow \mathbb{R}$ be quasiconcave, and let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be nondecreasing. Then $\varphi \circ f$ is quasiconcave.*

Proof. Let $x, y \in C$ and $t \in [0, 1]$. Applying φ to the inequality of Theorem 2.8 preserves it, because φ is nondecreasing, so

$$\varphi\left(f((1-t)x + ty)\right) \geq \varphi(\min\{f(x), f(y)\}) = \min\{\varphi(f(x)), \varphi(f(y))\},$$

the equality because a nondecreasing φ sends the smaller of two numbers to the smaller of their images. By Theorem 2.8 again, $\varphi \circ f$ is quasiconcave. $\square$

3 Concavity criteria, global optimality, and separation

3.1 Restriction to line segments

To decide whether a function of n variables is concave, we reduce the question to one variable. The reduction rests on three results: concavity of f is equivalent to concavity of its restriction to every line segment of the domain (Lemma 3.1); such a restriction has second derivative $v^\top H_f(x + tv)v$ (Theorem 3.2); and a one-variable function is concave exactly when its second derivative is nonpositive (Lemma 3.3).

Lemma 3.1 (Line-segment restriction). *Let $U \subseteq \mathbb{R}^n$ be convex and let $f : U \rightarrow \mathbb{R}$. For $x, y \in U$ let $v = y - x$ and*

$$\phi(t) = f(x + tv), \quad t \in [0, 1].$$

Then f is concave if and only if ϕ is concave on $[0, 1]$ for every choice of $x, y \in U$, and f is strictly concave if and only if ϕ is strictly concave for every choice of $x, y \in U$ with $x \neq y$.

Proof (optional). Concavity of f implies concavity of every ϕ . Fix $x, y \in U$ and let $s, s' \in [0, 1]$ and $t \in [0, 1]$. The point $x + ((1-t)s + ts')v$ is the convex combination $(1-t)(x + sv) + t(x + s'v)$ of two points of U , so Definition 2.4 gives


$$\phi((1-t)s + ts') = f((1-t)(x + sv) + t(x + s'v)) \geq (1-t)\phi(s) + t\phi(s').$$

Concavity of every ϕ implies concavity of f . Given $x, y \in U$ and $t \in [0, 1]$, take $s = 0$ and $s' = 1$ above. Since $\phi(0) = f(x)$ and $\phi(1) = f(y)$,


$$f((1-t)x + ty) = \phi(t) \geq (1-t)f(x) + tf(y).$$

For the strict statements, note that $x \neq y$ makes $v \neq 0$, so distinct values of s give distinct points $x + sv$; both displays above then hold with strict inequalities under the corresponding strict hypothesis. $\square$

For a function of one variable, the second derivative tells us whether the graph bends upward or downward. For a function of several variables, the input can move in many directions, and the curvature may depend on the direction. We therefore restrict the function to the line $x + tv$ and study it as a function of the single variable t . Figure 3.1 shows the construction.



(a) the segment
 $x + tv$ in the domain



(b) the restriction
 $\phi(t) = f(x + tv)$

Figure 3.1: Restricting f to a line segment. The segment joining x and y in the n -dimensional domain $U \subseteq \mathbb{R}^n$ (left) is parameterized as $x + tv$ with $v = y - x$; along it, the n -variable function f collapses to the function $\phi(t) = f(x + tv)$ of the single real variable t (right), with $\phi(0) = f(x)$ and $\phi(1) = f(y)$.

Theorem 3.2 (Directional curvature). *Let $U \subseteq \mathbb{R}^n$ be open, let $f : U \rightarrow \mathbb{R}$ be C^2 , let $x \in U$, and let $v \in \mathbb{R}^n$. Put $\phi(t) = f(x + tv)$ for those t with $x + tv \in U$. Then*

$$\phi'(t) = \nabla f(x + tv) \cdot v, \quad \phi''(t) = v^\top H_f(x + tv) v.$$

Proof (optional). Define the curve $\gamma : \mathbb{R} \rightarrow \mathbb{R}^n$ by

$$\gamma(t) = x + tv.$$

Its i th coordinate is the affine function $\gamma_i(t) = x_i + tv_i$, whose derivative is $\gamma'_i(t) = v_i$. Therefore $\gamma'(t) = v$.

Fix any t such that $\gamma(t) \in U$. Because U is open, there is $\varepsilon > 0$ such that $B(\gamma(t), \varepsilon) \subseteq U$. If $v \neq 0$, then

$$\|\gamma(s) - \gamma(t)\| = |s - t|\|v\| < \varepsilon$$

whenever $|s - t| < \varepsilon/\|v\|$. If $v = 0$, then $\gamma(s) = \gamma(t)$ for every s . In either case, $f \circ \gamma$ is defined on an open interval containing t . Since $\phi = f \circ \gamma$, the derivative-along-a-curve formula from Lecture 4 gives

$$\phi'(t) = \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\gamma(t)) \gamma'_i(t) = \sum_{i=1}^n \frac{\partial f}{\partial x_i}(x + tv) v_i = \nabla f(x + tv) \cdot v.$$

For each i , define $g_i : U \rightarrow \mathbb{R}$ by

$$g_i(y) = \frac{\partial f}{\partial x_i}(y).$$

Because f is C^2 , each g_i is C^1 and hence differentiable. Applying the same derivative-along-a-curve formula to $g_i \circ \gamma$ gives

$$\begin{aligned} (g_i \circ \gamma)'(t) &= \sum_{j=1}^n \frac{\partial g_i}{\partial x_j}(\gamma(t)) \gamma'_j(t) \\ &= \sum_{j=1}^n (H_f(x + tv))_{ij} v_j, \end{aligned}$$

where the second equality follows from the definition of the Hessian. For every s in the open interval above, the formula for the first derivative can be written as

$$\phi'(s) = \sum_{i=1}^n v_i (g_i \circ \gamma)(s).$$

Since each v_i is constant, differentiating this equality at $s = t$ gives

$$\phi''(t) = \sum_{i=1}^n v_i (g_i \circ \gamma)'(t) = \sum_{i=1}^n \sum_{j=1}^n v_i (H_f(x + tv))_{ij} v_j,$$

which equals $v^\top H_f(x + tv)v$ by the definition of matrix multiplication. Since t was arbitrary, both formulas hold for every t in the domain of ϕ . $\square$

Thus $v^\top H_f(x)v$ is the second derivative of f along the line through x in direction v : the Hessian measures curvature one direction at a time, just as the gradient measures slope one direction at a time. We use this result below to relate the signs of $v^\top H_f(x)v$ across directions v to concavity, and in Lecture 6 to derive necessary second-order conditions for a local optimum.

Finally, we state the one-variable second-derivative test.

Lemma 3.3 (One-variable second-derivative test). *Let $I \subseteq \mathbb{R}$ be an interval, and suppose ϕ is twice continuously differentiable on an open interval containing I . Then ϕ is concave on I if and only if $\phi''(t) \leq 0$ for every $t \in I$. If $\phi''(t) < 0$ for every $t \in I$, then ϕ is strictly concave.*

We use this result without proof. The second half of the statement does not reverse: a strictly concave ϕ may have $\phi'' = 0$ somewhere, as $\phi(t) = -t^4$ does at $t = 0$.

3.2 Tangent and Hessian criteria

The first criterion replaces the chord inequality by a comparison with the tangent plane at one point.

Theorem 3.4 (Tangent characterization). *Let $U \subseteq \mathbb{R}^n$ be open and convex, and let $f : U \rightarrow \mathbb{R}$ be differentiable. Then f is concave if and only if*

$$f(y) \leq f(x) + \nabla f(x) \cdot (y - x) \quad \text{for every } x, y \in U. \quad (3)$$

Proof (optional). Concavity implies (3). Let $x, y \in U$ and $t \in (0, 1)$. By Definition 2.4 applied to x and y ,

$$f(x + t(y - x)) \geq (1 - t)f(x) + tf(y),$$

and subtracting $f(x)$ and dividing by $t > 0$ gives

$$\frac{f(x + t(y - x)) - f(x)}{t} \geq f(y) - f(x).$$

As t decreases to zero the left-hand side tends to $\phi'(0)$, where $\phi(t) = f(x + t(y - x))$, and the chain rule gives $\phi'(0) = \nabla f(x) \cdot (y - x)$. The inequality passes to the limit, which is (3).

(3) *implies concavity*. Let $x, y \in U$, let $t \in [0, 1]$, and put $z = (1 - t)x + ty$, which lies in U because U is convex. Applying (3) at the point z to each of x and y ,

$$\begin{aligned} f(x) &\leq f(z) + \nabla f(z) \cdot (x - z), \\ f(y) &\leq f(z) + \nabla f(z) \cdot (y - z). \end{aligned}$$

Multiply the first by $1 - t$ and the second by t and add. The gradient terms combine into $\nabla f(z) \cdot ((1 - t)x + ty - z)$, which is zero by the definition of z , leaving $(1 - t)f(x) + tf(y) \leq f(z)$. $\square$

A differentiable concave function thus lies on or below each of its tangent planes. Note that the right-hand side of (3) uses the value and gradient of f at the point x only, yet it bounds f everywhere. Figure 3.2 shows the inequality in one variable.




Figure 3.2: The tangent characterization of Theorem 3.4. The graph of a differentiable concave function lies on or below the tangent line at x , so the value $f(x) + \nabla f(x) \cdot (y - x)$ computed from data at x alone bounds $f(y)$ at every other point y .

The second criterion uses the Hessian. Negative semidefiniteness means that every line restriction bends downward or remains flat.

Theorem 3.5 (Hessian characterization). *Let $U \subseteq \mathbb{R}^n$ be open and convex, and let $f : U \rightarrow \mathbb{R}$ be C^2 .*

1. *f is concave if and only if $H_f(x)$ is negative semidefinite for every $x \in U$.*
2. *If $H_f(x)$ is negative definite for every $x \in U$, then f is strictly concave.*

The corresponding statements for convexity reverse every sign.

Proof (optional). Negative semidefinite implies concave. Fix $x, y \in U$, put $v = y - x$, and let $\phi(t) = f(x + tv)$ as in Lemma 3.1. By openness of U , this restriction is C^2 on an open interval containing $[0, 1]$. By Theorem 3.2, $\phi''(t) = v^\top H_f(x + tv)v$, which is at most 0 by Definition 1.7. So ϕ is concave by Lemma 3.3, and since x and y were arbitrary, f is concave by Lemma 3.1.

Concave implies negative semidefinite. Fix $x \in U$ and $v \in \mathbb{R}^n$. Since U is open, $x + tv \in U$ for every t in some open interval containing 0. The restriction $\phi(t) = f(x + tv)$ is concave by Lemma 3.1, so Lemma 3.3 gives $\phi''(0) \leq 0$. By Theorem 3.2,

$$\phi''(0) = v^\top H_f(x)v \leq 0.$$

Since v was arbitrary, $H_f(x)$ is negative semidefinite.

Negative definite implies strictly concave. Fix distinct $x, y \in U$, put $v = y - x$, and let $\phi(t) = f(x + tv)$. The direction v is nonzero, so $\phi''(t) < 0$ throughout by Theorem 3.2 and Definition 1.7. Hence Lemma 3.3 makes ϕ strictly concave. Since x and y were arbitrary, Lemma 3.1 makes f strictly concave. $\square$

Part (2) does not reverse, for the same reason Lemma 3.3 does not: $f(x) = -x^4$ is strictly concave on $\mathbb{R}$ while $f''(0) = 0$. Negative definiteness is sufficient for strict concavity, never necessary.

3.3 Global optimality and uniqueness

For concave functions, a first-order condition at a single point is sufficient for a global maximum.

Corollary 3.6 (First-order sufficiency). *Let $U \subseteq \mathbb{R}^n$ be open and convex, let $C \subseteq U$ be convex, and let $f : U \rightarrow \mathbb{R}$ be differentiable and concave. If $x^* \in C$ satisfies*

$$\nabla f(x^*) \cdot (x - x^*) \leq 0 \quad \text{for every } x \in C, \quad (4)$$

then x^ maximizes f on C . If f is strictly concave, then x^* is the only maximizer.*

Proof. For $x \in C$, Theorem 3.4 at the point x^* and then (4) give

$$f(x) \leq f(x^*) + \nabla f(x^*) \cdot (x - x^*) \leq f(x^*).$$

Uniqueness. Suppose f is strictly concave and $x \neq y$ both maximize f on C , with common value m . Their midpoint lies in C by Definition 2.1, and Definition 2.4 gives $f(\frac{1}{2}x + \frac{1}{2}y) > \frac{1}{2}m + \frac{1}{2}m = m$, contradicting that m is the maximum. $\square$

In particular, any $x^* \in C$ with $\nabla f(x^*) = 0$ satisfies (4). It is therefore a global maximizer of a concave function, and under strict concavity it is the only one. The more general inequality (4) also allows x^* to lie on the boundary of C , where the gradient need not vanish.

Example 3.7 (Maximizing a strictly concave quadratic). Fix $b \in \mathbb{R}^2$ and $\gamma \in \mathbb{R}$ with $|\gamma| < 2$, and define $q : \mathbb{R}^2 \rightarrow \mathbb{R}$ by

$$q(x) = b \cdot x - (x_1^2 + \gamma x_1 x_2 + x_2^2).$$

Its Hessian has leading principal minors -2 and $4 - \gamma^2 > 0$, so it is negative definite by Theorem 1.10. Hence q is strictly concave on $\mathbb{R}^2$ by Theorem 3.5. Setting $\nabla q(x) = 0$ gives

$$\begin{pmatrix} 2 & \gamma \\ \gamma & 2 \end{pmatrix} x = b, \quad \text{so} \quad x^* = \frac{1}{4 - \gamma^2} \begin{pmatrix} 2 & -\gamma \\ -\gamma & 2 \end{pmatrix} b.$$

By Corollary 3.6, x^* is the unique global maximizer of q over $\mathbb{R}^2$. The first-order equations produce the candidate; strict concavity is what makes it the answer.

3.4 Separation

We end with a geometric consequence of convexity. Separating hyperplanes place two convex sets in opposite halfspaces.

Definition 3.8 (Hyperplane). Let $a \in \mathbb{R}^n \setminus \{0\}$ and $c \in \mathbb{R}$. The *hyperplane* with normal a at level c is

$$H(a, c) = \{x \in \mathbb{R}^n : a \cdot x = c\},$$

and its two *closed halfspaces* are $\{x : a \cdot x \leq c\}$ and $\{x : a \cdot x \geq c\}$.

Theorem 3.9 (Separating hyperplane). Let $A, B \subseteq \mathbb{R}^n$ be nonempty, disjoint, and convex. There exist $a \neq 0$ and $c \in \mathbb{R}$ such that

$$a \cdot x \leq c \leq a \cdot y \quad \text{for every } x \in A \text{ and every } y \in B.$$

We use this result without proof. The two sets lie in opposite closed halfspaces of $H(a, c)$. The inequalities are weak: the common level c may be approached, or even attained, by points of the two sets. Figure 3.3 shows the configuration in $\mathbb{R}^2$.




Figure 3.3: A hyperplane separating two disjoint convex sets. Every point of A satisfies $a \cdot x \leq c$ and every point of B satisfies $a \cdot y \geq c$.

Example 3.10 (No gap between the sets). Let $A = \{x \in \mathbb{R}^2 : x_2 < 0\}$ and $B = \{x \in \mathbb{R}^2 : x_2 \geq 0\}$, which are nonempty, disjoint, and convex. Taking $a = e_2$ and $c = 0$ gives

$$a \cdot x = x_2 < 0 \leq y_2 = a \cdot y \quad (x \in A, y \in B).$$

Thus the sets are separated by the level 0, but there is no positive gap between their separating levels: points of A have second coordinate arbitrarily close to 0.

A positive gap therefore requires hypotheses beyond disjointness.