

CS257 Linear and Convex Optimization

Lecture 2

Bo Jiang

John Hopcroft Center for Computer Science
Shanghai Jiao Tong University

September 14, 2020

Recap: Mathematical Optimization Problems

$$\min_{\mathbf{x} \in X} f(\mathbf{x})$$

- $f : \mathbb{R}^n \rightarrow \mathbb{R}$: **objective function**
- $\mathbf{x} = (x_1, x_2, \dots, x_n)^T \in \mathbb{R}^n$: **optimization/decision variables**
- $X \subset \mathbb{R}^n$: **feasible set or constraint set**

$$\begin{aligned} \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{s. t.} \quad & g_i(\mathbf{x}) \leq 0, \quad i = 1, 2, \dots, m \end{aligned}$$

- $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$: **constraint function**
- $X = \{\mathbf{x} : g_i(\mathbf{x}) \leq 0, i = 1, \dots, m\}$

Examples: linear regression, linear programming, SVM

Recap: Global Minimum and Local Minimum

$\mathbf{x}^* \in X$ is a **global minimum** of f if

$$f(\mathbf{x}^*) \leq f(\mathbf{x}), \quad \forall \mathbf{x} \in X$$

$\mathbf{x}^* \in X$ is a **local minimum** of f if there exists $\epsilon > 0$ s.t.

$$f(\mathbf{x}^*) \leq f(\mathbf{x}), \quad \forall \mathbf{x} \in X \cap B(\mathbf{x}^*, \epsilon)$$

Sufficient conditions for existence of global min

- f is continuous and X is compact (closed and bounded)
- f is continuous and **coercive** ($f(\mathbf{x}) \rightarrow \infty$ as $\|\mathbf{x}\| \rightarrow \infty$)

Contents

1. First-order Optimality Conditions for Unconstrained Optimization
2. Second-order Optimality Conditions for Unconstrained Optimization

Review: Derivative

\mathbf{x} is an **interior point** of $X \subset \mathbb{R}^n$ if there exists $\epsilon > 0$ s.t. $B(\mathbf{x}, \epsilon) \subset X$.

The **interior** of X , denoted by $\text{int } X$, is the set of interior points of X .

A function $f : X \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is **differentiable** at $\mathbf{x}_0 \in \text{int } X$, if there exists a matrix¹ $A \in \mathbb{R}^{m \times n}$ s.t.

$$\lim_{X \ni \mathbf{x} \rightarrow \mathbf{x}_0} \frac{\|f(\mathbf{x}) - f(\mathbf{x}_0) - A(\mathbf{x} - \mathbf{x}_0)\|}{\|\mathbf{x} - \mathbf{x}_0\|} = 0$$

The matrix A is called the **derivative** of f at \mathbf{x}_0 , and we write

$$f'(\mathbf{x}_0) = Df(\mathbf{x}_0) = A$$

The affine function $f(\mathbf{x}_0) + A(\mathbf{x} - \mathbf{x}_0)$ is the first-order approximation of f at \mathbf{x}_0 ,

$$f(\mathbf{x}) = f(\mathbf{x}_0) + f'(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0) + o(\|\mathbf{x} - \mathbf{x}_0\|)$$

¹More precisely, a linear transformation represented by matrix A

Review: Derivative

The derivative is given by the **Jacobian matrix** of $\mathbf{f} = (f_1, \dots, f_m)$

$$[\mathbf{f}'(\mathbf{x}_0)]_{ij} = \frac{\partial f_i(\mathbf{x}_0)}{\partial x_j}, \quad i = 1, \dots, m; j = 1, \dots, n$$

Example. An affine function $\mathbf{f}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$ from \mathbb{R}^n to \mathbb{R}^m has derivative $\mathbf{f}'(\mathbf{x}) = \mathbf{A}$ at all \mathbf{x} . In particular, when $m = 1$, $f(\mathbf{x}) = \mathbf{a}^T \mathbf{x} + b$ has derivative $f'(\mathbf{x}) = \mathbf{a}^T$, which is a $1 \times n$ matrix, i.e. a row vector.

Example. $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{Q} \mathbf{x} = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n Q_{ij} x_i x_j$ has derivative

$$f'(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T (\mathbf{Q} + \mathbf{Q}^T)$$

Proof.

$$\frac{\partial f}{\partial x_k} = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n Q_{ij} \left(x_j \frac{\partial x_i}{\partial x_k} + x_i \frac{\partial x_j}{\partial x_k} \right) = \frac{1}{2} \sum_{j=1}^n Q_{kj} x_j + \frac{1}{2} \sum_{i=1}^n Q_{ik} x_i$$

Review: Gradient

For a real-valued function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the **gradient** of f at \mathbf{x} , denoted by $\nabla f(\mathbf{x})$, is the transpose of $f'(\mathbf{x})$,

$$\nabla f(\mathbf{x}) = [f'(\mathbf{x})]^T, \quad [\nabla f(\mathbf{x})]_i = \frac{\partial f(\mathbf{x})}{\partial x_i}, \quad i = 1, \dots, n$$

$\nabla f(\mathbf{x})$ is a column vector and satisfies

$$f'(\mathbf{x})\Delta\mathbf{x} = \langle \nabla f(\mathbf{x}), \Delta\mathbf{x} \rangle = \nabla f(\mathbf{x})^T \Delta\mathbf{x}$$

The first-order approximation of f at \mathbf{x}_0 is

$$f(\mathbf{x}_0) + \nabla f(\mathbf{x}_0)^T (\mathbf{x} - \mathbf{x}_0)$$

Example. For symmetric \mathbf{Q} , the gradient of $f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{Q}\mathbf{x} + \mathbf{b}^T \mathbf{x} + c$ is

$$\nabla f(\mathbf{x}) = \mathbf{Q}\mathbf{x} + \mathbf{b}$$

Review: Chain Rule

If $f : X \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at $\mathbf{x}_0 \in X$, $g : Y \subset \mathbb{R}^m \rightarrow \mathbb{R}^p$ is differentiable at $\mathbf{y}_0 = f(\mathbf{x}_0)$, then the composition of f and g defined by $h(\mathbf{x}) = g(f(\mathbf{x}))$ is differentiable at \mathbf{x}_0 , and

$$h'(\mathbf{x}_0) = g'(\mathbf{y}_0)f'(\mathbf{x}_0) = g'(f(\mathbf{x}_0))f'(\mathbf{x}_0)$$

Note. The order is important since $g'(\mathbf{y}_0) \in \mathbb{R}^{p \times m}$ and $f'(\mathbf{x}_0) \in \mathbb{R}^{m \times n}$ are matrices. In general $f'(\mathbf{x}_0)g'(\mathbf{y}_0)$ is **undefined**.

$$\begin{array}{ccccc} \mathbb{R}^n & \xrightarrow{f} & \mathbb{R}^m & \xrightarrow{g} & \mathbb{R}^p \\ \mathbf{x}_0 & \mapsto & \mathbf{y}_0 = f(\mathbf{x}_0) & \mapsto & h(\mathbf{x}_0) = g(\mathbf{y}_0) \\ \Delta \mathbf{x} & \xrightarrow{f'} & f'(\mathbf{x}_0)\Delta \mathbf{x} & \xrightarrow{g'} & g'(\mathbf{y}_0)f'(\mathbf{x}_0)\Delta \mathbf{x} \end{array}$$

Review: Chain Rule

Example. $h(\mathbf{x}) = f(\mathbf{Ax} + \mathbf{b})$ has derivative $h'(\mathbf{x}_0) = f'(\mathbf{Ax}_0 + \mathbf{b})\mathbf{A}$. If f is real-valued,

$$\nabla h(\mathbf{x}_0) = \mathbf{A}^T [f'(\mathbf{Ax}_0 + \mathbf{b})]^T = \mathbf{A}^T \nabla f(\mathbf{Ax}_0 + \mathbf{b})$$

Example. Given $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\mathbf{x}, \mathbf{v} \in \mathbb{R}^n$, define

$$\tilde{f}(t) = f(\mathbf{x} + t\mathbf{v})$$

Then

$$\tilde{f}'(t) = f'(\mathbf{x} + t\mathbf{v})\mathbf{v} = \nabla f(\mathbf{x} + t\mathbf{v})^T \mathbf{v} = \mathbf{v}^T \nabla f(\mathbf{x} + t\mathbf{v})$$

Note. \tilde{f} is the restriction of f to the straight line through \mathbf{x} with direction \mathbf{v} . We can often get useful information about f by looking at \tilde{f} , which is usually easier to deal with.

First-order Necessary Condition

Consider unconstrained optimization problem, i.e. $X = \mathbb{R}^n$.

Theorem. If \mathbf{x}^* is a local minimum of f and f is differentiable at \mathbf{x}^* , then its gradient at \mathbf{x}^* vanishes, i.e.

$$\nabla f(\mathbf{x}^*) = \left(\frac{\partial f(\mathbf{x}^*)}{\partial x_1}, \dots, \frac{\partial f(\mathbf{x}^*)}{\partial x_n} \right)^T = \mathbf{0}.$$


Proof. Let $\mathbf{d} \in \mathbb{R}^n$. Define $g(\alpha) = f(\mathbf{x}^* + \alpha\mathbf{d})$.

- Since \mathbf{x}^* is local minimum, $g(\alpha) \geq g(0)$
- For $\alpha > 0$,

$$\frac{g(\alpha) - g(0)}{\alpha} \geq 0 \implies g'(0) = \lim_{\alpha \downarrow 0} \frac{g(\alpha) - g(0)}{\alpha} \geq 0$$


- By chain rule, $g'(0) = \sum_{i=1}^n d_i \frac{\partial f(\mathbf{x}^*)}{\partial x_i} = \mathbf{d}^T \nabla f(\mathbf{x}^*) \geq 0$
- Replacing \mathbf{d} by $-\mathbf{d} \implies -\mathbf{d}^T \nabla f(\mathbf{x}^*) \geq 0 \implies \mathbf{d}^T \nabla f(\mathbf{x}^*) = 0$
- Setting $\mathbf{d} = \nabla f(\mathbf{x}^*) \implies \|\nabla f(\mathbf{x}^*)\|^2 = 0 \implies \nabla f(\mathbf{x}^*) = \mathbf{0}$

First-order Necessary Condition (cont'd)




First-order Necessary Condition (cont'd)

A point \mathbf{x}^* with $\nabla f(\mathbf{x}^*) = \mathbf{0}$ is called a **stationary point** of f .




$$x^* = 0$$

minimum




$$x^* = 0$$

inflection point




$$x^* = 0$$

maximum




$$f(\mathbf{x}) = x_1^2 + x_2^2$$

minimum




$$f(\mathbf{x}) = x_1^2 - x_2^2$$

saddle point



$$f(\mathbf{x}) = -x_1^2 - x_2^2$$

maximum



$$f(\mathbf{x}) = -x_1|x_1| + x_2^2$$

Note. Will see stationarity is sufficient for convex optimization.


First-order Necessary Condition (cont'd)

For constrained optimization problem, i.e. $X \neq \mathbb{R}^n$,

- if \mathbf{x}^* is in the interior of X , i.e. $B(\mathbf{x}^*, \epsilon) \subset X$ for some $\epsilon > 0$, then the proof still works, so $\nabla f(\mathbf{x}^*) = \mathbf{0}$
- otherwise, the proof shows $\mathbf{d}^T \nabla f(\mathbf{x}^*) \geq 0$ for any feasible direction \mathbf{d} at \mathbf{x}^*
 - ▶ \mathbf{d} is a feasible direction at $\mathbf{x} \in X$ if $\mathbf{x} + \alpha \mathbf{d} \in X$ for all sufficiently small $\alpha > 0$
- will revisit later

Example. $X = [a, b]$

- $f'(x_1) = 0$
- $d_1 f'(a) \geq 0 \implies f'(a) \geq 0$
- $d_2 f'(b) \geq 0 \implies f'(b) \leq 0$



Contents

1. First-order Optimality Conditions for Unconstrained Optimization
2. Second-order Optimality Conditions for Unconstrained Optimization

Review: Second Derivative

The second-order partial derivatives of $f : X \subset \mathbb{R}^n \rightarrow \mathbb{R}$ at $\mathbf{x}_0 \in \text{int } X$ are

$$\frac{\partial^2 f(\mathbf{x}_0)}{\partial x_i \partial x_j}, \quad i, j = 1, 2, \dots, n$$

The **Hessian (matrix)** of f at \mathbf{x}_0 , denoted by $\nabla^2 f(\mathbf{x}_0)$, is given by

$$[\nabla^2 f(\mathbf{x}_0)]_{ij} = \frac{\partial^2 f(\mathbf{x}_0)}{\partial x_i \partial x_j}, \quad i, j = 1, 2, \dots, n$$

If $\frac{\partial^2 f(\mathbf{x})}{\partial x_i \partial x_j}$ and $\frac{\partial^2 f(\mathbf{x})}{\partial x_j \partial x_i}$ exist in a neighborhood of \mathbf{x}_0 and are continuous at \mathbf{x}_0 , then

$$\frac{\partial^2 f(\mathbf{x}_0)}{\partial x_i \partial x_j} = \frac{\partial^2 f(\mathbf{x}_0)}{\partial x_j \partial x_i}$$

so $\nabla^2 f(\mathbf{x}_0)$ is symmetric.

Will assume twice continuous differentiability when considering $\nabla^2 f$.

Review: Second-order Taylor Expansion

The second-order Taylor expansion for $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is

$$f(\mathbf{x} + \mathbf{d}) = f(\mathbf{x}) + \sum_{i=1}^n \frac{\partial f(\mathbf{x})}{\partial x_i} d_i + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 f(\mathbf{x})}{\partial x_i \partial x_j} d_i d_j + o(\|\mathbf{d}\|^2)$$

or in vector notation,

$$f(\mathbf{x} + \mathbf{d}) = f(\mathbf{x}) + \nabla f(\mathbf{x})^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}) \mathbf{d} + o(\|\mathbf{d}\|^2)$$

Note. This can be used to find the expressions for ∇f and $\nabla^2 f$.

Example. For affine function $f(\mathbf{x}) = \mathbf{b}^T \mathbf{x} + \mathbf{c}$

$$\nabla f(\mathbf{x}) = \mathbf{b}, \quad \nabla^2 f(\mathbf{x}) = \mathbf{O}$$

Proof. Compare the following with Taylor expansion.

$$f(\mathbf{x} + \mathbf{d}) - f(\mathbf{x}) = [\mathbf{b}^T (\mathbf{x} + \mathbf{d}) + \mathbf{c}] - [\mathbf{b}^T \mathbf{x} + \mathbf{c}] = \mathbf{b}^T \mathbf{d}$$

Review: Second-order Taylor Expansion

Example. For quadratic function $f(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x}$ with general \mathbf{A} ,

$$\nabla f(\mathbf{x}) = (\mathbf{A} + \mathbf{A}^T)\mathbf{x}, \quad \nabla^2 f(\mathbf{x}) = \mathbf{A} + \mathbf{A}^T$$

If \mathbf{A} is symmetric, then $\nabla f(\mathbf{x}) = 2\mathbf{A}\mathbf{x}$, $\nabla^2 f(\mathbf{x}) = 2\mathbf{A}$.

Proof.

$$\begin{aligned} f(\mathbf{x} + \mathbf{d}) &= (\mathbf{x} + \mathbf{d})^T \mathbf{A} (\mathbf{x} + \mathbf{d}) = \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{d}^T \mathbf{A} \mathbf{x} + \mathbf{x}^T \mathbf{A} \mathbf{d} + \mathbf{d}^T \mathbf{A} \mathbf{d} \\ &= f(\mathbf{x}) + (\mathbf{A}\mathbf{x} + \mathbf{A}^T \mathbf{x})^T \mathbf{d} + \mathbf{d}^T \mathbf{A} \mathbf{d} \end{aligned}$$

Since quadratic functions are twice continuously differentiable, $\nabla^2 f$ is symmetric. Need to rewrite the above as

$$f(\mathbf{x} + \mathbf{d}) = f(\mathbf{x}) + (\mathbf{A}\mathbf{x} + \mathbf{A}^T \mathbf{x})^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T (\mathbf{A} + \mathbf{A}^T) \mathbf{d}$$

Since a quadratic function is **exactly** equal to its second-order Taylor expansion, we must have $\nabla f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{A}^T \mathbf{x}$ and $\nabla^2 f(\mathbf{x}) = \mathbf{A} + \mathbf{A}^T$.

Review: Chain Rule for Second Derivative

The composition with affine function $g(\mathbf{x}) = f(\mathbf{A}\mathbf{x} + \mathbf{b})$ has Hessian

$$\nabla^2 g(\mathbf{x}) = \mathbf{A}^T \nabla^2 f(\mathbf{A}\mathbf{x} + \mathbf{b}) \mathbf{A}$$

Proof. Let $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{b}$, i.e. $y_k = \sum_i A_{ki} x_i$. Recall $\nabla g(\mathbf{x}) = \mathbf{A}^T \nabla f(\mathbf{y})$, i.e.

$$\frac{\partial g(\mathbf{x})}{\partial x_j} = \sum_k \frac{\partial f(\mathbf{y})}{\partial y_k} \frac{\partial y_k}{\partial x_j} = \sum_k \frac{\partial f(\mathbf{y})}{\partial y_k} A_{kj}$$

$$\frac{\partial^2 g(\mathbf{x})}{\partial x_i \partial x_j} = \sum_k \frac{\partial}{\partial x_i} \frac{\partial f(\mathbf{y})}{\partial y_k} A_{kj} = \sum_k \sum_\ell \frac{\partial^2 f(\mathbf{y})}{\partial y_\ell \partial y_k} A_{\ell i} A_{kj} = [\mathbf{A}^T \nabla^2 f(\mathbf{y}) \mathbf{A}]_{ij}$$

Special case. For $\tilde{f}(t) = f(\mathbf{x} + t\mathbf{v})$,

$$\tilde{f}''(t) = \mathbf{v}^T \nabla^2 f(\mathbf{x} + t\mathbf{v}) \mathbf{v}$$

Proof. Set $\mathbf{A} \leftarrow \mathbf{v}$, $\mathbf{x} \leftarrow t$, $\mathbf{b} \leftarrow \mathbf{x}$ in the general formula above.

Review: Definite Matrices

Matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is **positive semidefinite**, denoted by $\mathbf{A} \succeq \mathbf{O}$, if

1. it is symmetric, i.e. $\mathbf{A} = \mathbf{A}^T$
2. $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0, \forall \mathbf{x} \in \mathbb{R}^n$

It is **positive definite**, denoted by $\mathbf{A} \succ \mathbf{O}$, if condition 2 is replaced by

- 2'. $\mathbf{x}^T \mathbf{A} \mathbf{x} > 0, \forall \mathbf{x} \in \mathbb{R}^n$ and $\mathbf{x} \neq \mathbf{0}$.

\mathbf{A} is **negative (semi)definite** if $-\mathbf{A}$ is positive (semi)definite.

\mathbf{A} is **indefinite** if it is neither positive semidefinite nor negative semidefinite, i.e. there exists $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$ s.t.

$$\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 > 0 > \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$$

Note. For quadratic forms $\mathbf{x}^T \mathbf{A} \mathbf{x}$, can always assume \mathbf{A} is symmetric, since

$$\mathbf{x}^T \mathbf{A} \mathbf{x} = \mathbf{x}^T \mathbf{A}^T \mathbf{x} = \mathbf{x}^T \left(\frac{\mathbf{A} + \mathbf{A}^T}{2} \right) \mathbf{x}$$

Review: Test for Positive Definiteness

Vector \mathbf{x} is **eigenvector** of matrix \mathbf{A} with associated **eigenvalue** λ if

$$\mathbf{Ax} = \lambda\mathbf{x}$$

Find eigenvalues by solving $\det(\lambda\mathbf{I} - \mathbf{A}) = 0$.

Theorem. Let \mathbf{A} be a symmetric matrix.

- $\mathbf{A} \succ \mathbf{O}$ iff all its eigenvalues $\lambda > 0$.
- $\mathbf{A} \succeq \mathbf{O}$ iff all its eigenvalues $\lambda \geq 0$.

Example. $\mathbf{A} = \begin{pmatrix} 1 & 2 \\ 2 & 5 \end{pmatrix}$ is positive definite.

$$\det(\lambda\mathbf{I} - \mathbf{A}) = (\lambda - 1)(\lambda - 5) - 4 = 0 \implies \lambda = 3 \pm 2\sqrt{2} > 0$$

Example. $\mathbf{A} = \begin{pmatrix} 1 & 2 \\ 2 & 4 \end{pmatrix}$ is positive semidefinite.

$$\det(\lambda\mathbf{I} - \mathbf{A}) = (\lambda - 1)(\lambda - 4) - 4 = 0 \implies \lambda_1 = 0, \lambda_2 = 5$$

Review: Test for Positive Definiteness

Given matrix $A = (a_{ij}) \in \mathbb{R}^{n \times n}$, a $k \times k$ principal submatrix of A consists of k rows and k columns with the same indices $I = \{i_1 < i_2 < \cdots < i_k\}$,

$$A_I = \begin{pmatrix} a_{i_1 i_1} & \cdots & a_{i_1 i_k} \\ \vdots & \ddots & \vdots \\ a_{i_k i_1} & \cdots & a_{i_k i_k} \end{pmatrix}$$

A **principal minor of order k** of A is $\det A_I$ for some I with $|I| = k$.

If $I = \{1, 2, \dots, k\}$, $D_k(A) \triangleq \det A_I$ is called the **leading principal minor of order k** .

Theorem (Sylvester). Let A be a symmetric matrix.

- $A \succ \mathbf{O}$ iff $D_k(A) > 0$ for $k = 1, 2, \dots, n$.
- $A \succeq \mathbf{O}$ iff $\det A_I \geq 0$ for all $I \subset \{1, 2, \dots, n\}$

Note. For positive semidefiniteness, we need to check **all** principal minors, not just the leading principal minors.

Review: Test for Positive Definiteness

Exmample. $A = \begin{pmatrix} 1 & 2 \\ 2 & 5 \end{pmatrix}$ is positive definite.

$$D_1(A) = \det(1) = 1 > 0, \quad D_2(A) = \det A = 1 > 0$$

Example. $A = \begin{pmatrix} 1 & 2 \\ 2 & 4 \end{pmatrix}$ is positive semidefinite.

$$D_1(A) = \det(1) = 1, \quad \det A_{\{2\}} = \det(4) = 4, \quad D_2(A) = \det A = 0$$

Note. It is **not** enough to check $D_k(A) \geq 0$ for all k !

Example. $A = \begin{pmatrix} 0 & 0 \\ 0 & -2 \end{pmatrix}$ is negative semidefinite,

$$D_1(A) = \det(0) = 0, \quad D_2(A) = \det A = 0,$$

but

$$\det A_{\{2\}} = \det(-2) = -2 < 0$$

Review: Test for Positive Definiteness

Example. $A = \begin{pmatrix} 1 & 2 \\ 2 & 5 \end{pmatrix}$ is positive definite.

- Use definition,

$$\mathbf{x}^T \mathbf{A} \mathbf{x} = x_1^2 + 4x_1x_2 + 5x_2^2 = (x_1 + 2x_2)^2 + x_2^2 \geq 0, \quad \forall \mathbf{x} \in \mathbb{R}^2$$

$$\text{with equality} \iff \begin{cases} x_1 + 2x_2 = 0 \\ x_2 = 0 \end{cases} \iff \mathbf{x} = \mathbf{0}$$

- Find eigenvalues by solving $\det(\lambda \mathbf{I} - \mathbf{A}) = 0$

$$\det \begin{pmatrix} \lambda - 1 & -2 \\ -2 & \lambda - 5 \end{pmatrix} = (\lambda - 1)(\lambda - 5) - 4 = 0 \implies \lambda = 3 \pm 2\sqrt{2} > 0$$

- Check leading principal minors

$$D_1(\mathbf{A}) = \det(1) = 1 > 0, \quad D_2(\mathbf{A}) = \det \mathbf{A} = 1 > 0$$

Review: Test for Positive Definiteness

Exmample. $A = \begin{pmatrix} 1 & 2 & 1 \\ 2 & 5 & 8 \\ 1 & 8 & 1 \end{pmatrix}$ is not positive definite.

Check leading principal minors

$$D_1(\mathbf{A}) = \det(1) = 1 > 0, \quad D_2(\mathbf{A}) = \det \begin{pmatrix} 1 & 2 \\ 2 & 5 \end{pmatrix} = 1 > 0$$

$$D_3(\mathbf{A}) = \det \mathbf{A} = 1 \times \begin{vmatrix} 5 & 8 \\ 8 & 1 \end{vmatrix} - 2 \times \begin{vmatrix} 2 & 8 \\ 1 & 1 \end{vmatrix} + 1 \times \begin{vmatrix} 2 & 5 \\ 1 & 8 \end{vmatrix} = -36 < 0$$

Can also check eigenvalues, e.g. using `numpy.linalg.eig`,

$$\lambda_1 = 11.69585173, \quad \lambda_2 = 0.58307572, \quad \lambda_3 = -5.27892745$$

Review: Eigendecomposition

A symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ has the following eigendecomposition

$$\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T = \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^T$$


where $\mathbf{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_n\}$, $\mathbf{Q} = (\mathbf{v}_1, \dots, \mathbf{v}_n)$ is an orthogonal matrix, i.e. $\mathbf{Q}^T \mathbf{Q} = \mathbf{I}$, and $\mathbf{A}\mathbf{v}_i = \lambda_i \mathbf{v}_i$.

Example. $\mathbf{A} = \frac{1}{4} \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix}$ has eigenvalues $\lambda_1 = \frac{1}{2}$ and $\lambda_2 = 1$, with corresponding eigenvectors $\mathbf{v}_1 = \frac{1}{\sqrt{2}}(1, 1)^T$ and $\mathbf{v}_2 = \frac{1}{\sqrt{2}}(-1, 1)^T$. The eigendecomposition is

$$\mathbf{A} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} = \frac{1}{2} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}^T + \begin{pmatrix} \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}^T$$


Review: Geometry of Quadratic Forms

Quadratic form $f(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x}$ in \mathbb{R}^2




$$\mathbf{A} = \text{diag}\{1, 1\}$$

positive definite




$$\mathbf{A} = \text{diag}\{0, 1\}$$

positive semidefinite




$$\mathbf{A} = \text{diag}\{1, -1\}$$

indefinite



$$\mathbf{A} = \text{diag}\{-1, -1\}$$

negative definite




$$\mathbf{A} = \text{diag}\{-1, 0\}$$


negative semidefinite

Review: Geometry of Quadratic Forms


Quadratic form $f(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x}$ in \mathbb{R}^2






$$\mathbf{A} = \text{diag}\{1, 1\}$$



$$\mathbf{A} = \text{diag}\left\{\frac{1}{2}, 1\right\}$$



$$\mathbf{A} = \frac{1}{4} \begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix}$$



Second-order Necessary Condition

Theorem. If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is twice continuously differentiable and \mathbf{x}^* is a local minimum of f , then its **Hessian matrix** $\nabla^2 f(\mathbf{x}^*)$ is positive semidefinite, i.e.

$$\mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} \geq 0, \quad \forall \mathbf{d} \in \mathbb{R}^n$$

Proof. Fix $\mathbf{d} \in \mathbb{R}^n$. By the first-order necessary condition, $\nabla f(\mathbf{x}^*) = \mathbf{0}$. By the second-order Taylor expansion, for any $t > 0$,

$$f(\mathbf{x}^* + t\mathbf{d}) = f(\mathbf{x}^*) + \frac{t^2}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(t^2 \|\mathbf{d}\|^2) \geq f(\mathbf{x}^*)$$

So

$$\frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(\|\mathbf{d}\|^2) \geq 0$$

Taking the limit $t \rightarrow 0$ yields $\mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} \geq 0$.

Note. Can apply the same argument to $g(\alpha) = f(\mathbf{x}^* + \alpha\mathbf{d})$ with local minimum $\alpha^* = 0$ and use chain rule to obtain $g''(0) = \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} \geq 0$.

Second-order Sufficient Condition

Theorem. Suppose f is twice continuously differentiable. If




1. $\nabla f(\mathbf{x}^*) = 0$
2. $\nabla^2 f(\mathbf{x}^*)$ is positive definite, i.e.

$$\mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} > 0, \quad \forall \mathbf{d} \neq \mathbf{0}$$

then \mathbf{x}^* is a local minimum.


Proof. Use second-order Taylor expansion.

Note. In condition 2, positive definiteness **cannot** be replaced by positive semidefiniteness.




Second-order Sufficient Condition (cont'd)


$\nabla f(\mathbf{0}) = \mathbf{0}$ and $\nabla^2 f(\mathbf{0}) = \mathbf{O}$ for all examples below.




$$f(x) = |x_1|^3 + |x_2|^3$$



$$f(x) = |x_1|^3 - |x_2|^3$$



$$f(x) = -|x_1|^3 - |x_2|^3$$



$$f(x) = -x_1^3 + |x_2|^3$$