EECS 127
Convex Optimization
Notes

Instructor: Laurent El Ghaoui

Kelvin Lee

UC Berkeley

Contents

1 Introduction **4**
- 1.1 Standard Form of Optimization . . . 4
- 1.2 Convex Problems . . . 9
- 1.3 Non-convex Problems . . . 9

2 Vectors and Functions **10**
- 2.1 Basics . . . 10
- 2.2 Vector Spaces . . . 10
- 2.3 Inner Product . . . 13
- 2.4 Orthogonality and Orthonormality . . . 15
- 2.5 Functions and Maps . . . 19
- 2.6 Hyperplanes and Halfspaces . . . 21
- 2.7 Gradients . . . 21

3 Matrices and Linear Maps **24**
- 3.1 Matrix Basics . . . 24
- 3.2 Matrices as linear maps . . . 26
- 3.3 Determinants . . . 27
- 3.4 Matrix factorizations . . . 30
- 3.5 Matrix Norms . . . 30

4 Matrices II **32**
- 4.1 Orthogonalization . . . 32

5 Symmetric Matrices **33**
- 5.1 The Spectral Theorem . . . 33

6 Singular Value Decomposition **34**
- 6.1 Dyads . . . 34

7 Linear Equations **35**
- 7.1 Set of solutions of linear equations . . . 35
- 7.2 Use case in optimization . . . 36
- 7.3 Solving via SVD . . . 36

8 Convexity **37**
8.1 Affine Sets . . . 37
8.2 Convex Sets . . . 37
8.3 Cones . . . 39
8.4 Hyperplanes and Halfspaces . . . 40
8.5 Operations Preserving Convexity . . . 41
8.6 Convex Functions . . . 42
8.7 Operations Preserving Convexity (continued) . . . 44

9 Convex Optimization Problems **47**
9.1 Trouble with nonlinear models . . . 47
9.2 Convex Problem . . . 47
9.3 Problem Transformations . . . 49

10 Weak Duality **53**
10.1 Lagrangian . . . 53
10.2 Minimax Inequality . . . 54
10.3 Weak Duality . . . 54
10.4 . . . 54

11 Strong Duality **55**
11.1 Strong Duality for Convex Problems . . . 55

1

Introduction

1.1 Standard Form of Optimization

$$\begin{array}{rl} p^* = \min_{\boldsymbol{x}} & f_0(\boldsymbol{x}) \\ \text{subject to:} & f_i(\boldsymbol{x}) \leq 0, \quad i = 1, \ldots, m \end{array}$$

where

- vector $\boldsymbol{x} \in \mathbb{R}^n$ is the **decision variable**;
- $f_0 : \mathbb{R}^n \to \mathbb{R}$ is the **objective function**, or **cost**;
- $f_i : \mathbb{R}^n \to \mathbb{R}, i = 1, \ldots, m,$ represent the **constraints**;
- p^* is the **optimal value**.

1.1.1 Least-squares Regression

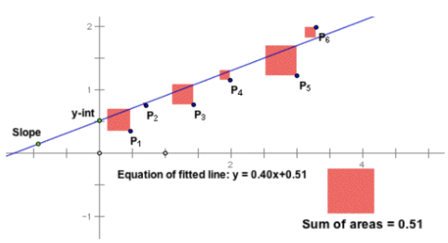

Figure 1.1: Least-squares regression.

$$\min_{\boldsymbol{x}} \sum_{i=1}^{m} \left(\boldsymbol{y}_i - \boldsymbol{x}^\top \boldsymbol{z}^{(i)} \right)^2$$

where

- $\boldsymbol{z}^{(i)} \in \mathbb{R}^n, i = 1, \ldots, n$ are **data points**;
- $\boldsymbol{y} \in \mathbb{R}^m$ is a **response vector**;

- $\boldsymbol{x}^\top \boldsymbol{z}$ is the **scalar product** $z_1 x_1 + \ldots + z_n x_n$ between the two vectors $\boldsymbol{x}, \boldsymbol{z} \in \mathbb{R}^n$.
- Many variants exist.
- Once $\boldsymbol{x}$ is found, allows to predict the output $\hat{\boldsymbol{y}}$ corresponding to a new data point $z : \hat{\boldsymbol{y}} = \boldsymbol{x}^\top \boldsymbol{z}$.

1.1.2 Linear Classification

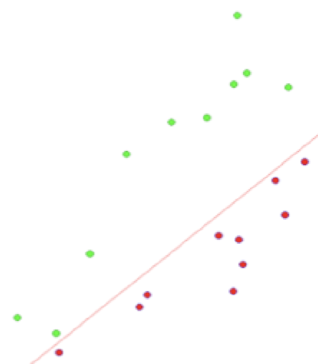

Figure 1.2: Linear classification.

Support Vector Machine (SVM):

$$\min_{\boldsymbol{x},b} \sum_{i=1}^{m} \max\left(0, 1 - y_i \left(\boldsymbol{x}^\top \boldsymbol{z}^{(i)} + b\right)\right)$$

where

- $\mathbf{0} z^{(i)} \in \mathbb{R}^n, i = 1, \ldots, n$ are data points;
- $\boldsymbol{y} \in \{-1, 1\}^m$ is a **binary response vector**;
- $x^\top z + b = 0$ defines a **separating hyperplane** in data space.
- Once $\boldsymbol{x}, n$ are found, we can predict the binary output $\hat{y}$ corresponding to a new data point $\boldsymbol{z}$: $\hat{y} = \textbf{sign}\left(\boldsymbol{x}^T \boldsymbol{z} + b\right)$.
- Very useful for classifying data.

1.1.3 Nomenclature

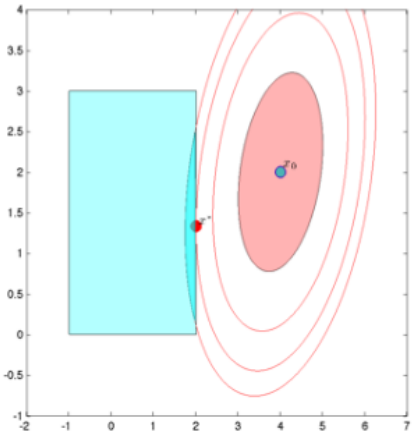

Figure 1.3: A toy optimization problem.

$$\begin{array}{ll} \min_{\boldsymbol{x}} & 0.9x_1^2 - 0.4x_1x_2 - 0.6x_2^2 - 6.4x_1 - 0.8x_2 \\ \text{s.t.} & -1 \le x_1 \le 2, 0 \le x_2 \le 3. \end{array}$$

- **Feasible set**: a set of possible values that satisfy the constraints. (light blue region)
- **Unconstrained minimizer**: x_0.
- **Optimal Point**: $\boldsymbol{x}^*$.
- **Level sets** of objective functions: $\{\boldsymbol{x} \mid g(\boldsymbol{x}) = c\}$ for some c. (red lines)
- **Sub-level sets**: $\{\boldsymbol{x} \mid g(\boldsymbol{x}) \le c\}$ for some c. (red region)

1.1.4 Problems with equality constraints

Sometimes the problem may have equality constraints, along with inequality ones:

$$\begin{array}{rll} p^* = \min_{\boldsymbol{x}} & f_0(\boldsymbol{x}) & \\ \text{s.t.} & f_i(\boldsymbol{x}) \le 0, & i = 1, \dots, m \\ & h_i(\boldsymbol{x}) = 0, & i = 1, \dots, p \end{array}$$

where h_i's are given functions.
However, we can always reduce it to a **standard form** with inequality constraints only, using the following method:

$$h_i(\boldsymbol{x}) = 0 \quad \Longrightarrow \quad h_i(\boldsymbol{x}) \le 0, \quad h_i(\boldsymbol{x}) \ge 0.$$

1.1.5 Problems with set constraints

Sometimes, the constraints of the problem are described abstractly via a set-membership condition of the form $\boldsymbol{x} \in \mathcal{X}$, for some subset $\mathcal{X}$ of $\mathbb{R}^n$.

The corresponding notation is

$$p^* = \min_{\boldsymbol{x} \in \mathcal{X}} f_0(\boldsymbol{x}),$$

or, equivalently,

$$\begin{aligned} p^* = \min_{\boldsymbol{x}} \quad & f_0(x) \\ \text{s.t.} \quad & x \in \mathcal{X}. \end{aligned}$$

1.1.6 Problems in maximization form

Some optimization problems come in the form of maximization (instead of minimization) of an objective function, i.e.,

$$p^* = \max_{\boldsymbol{x} \in \mathcal{X}} g_0(\boldsymbol{x}).$$

We can recast it as a **standard minimization form** using the following fact:

$$\max_{\boldsymbol{x} \in \mathcal{X}} g_0(\boldsymbol{x}) = -\min_{\boldsymbol{x} \in \mathcal{X}} -g_0(\boldsymbol{x}).$$

Thus we can reformulate the problem as the following:

$$-p^* = \min_{\boldsymbol{x} \in \mathcal{X}} f_0(\boldsymbol{x}),$$

where $f_0 = -g_0$.

1.1.7 Feasible Set

The **feasible set** of a problem is defined as

$$\mathcal{X} = \{\boldsymbol{x} \in \mathbb{R}^n \mid f_i(\boldsymbol{x}) \leq 0, \quad i = 1, \ldots, m\}.$$

Definition 1.1.1 (Infeasible)**.** A problem is **infeasible** if the feasible set is empty, i.e., the constraints cannot be satisfied simultaneously.

Remark. We take the convention that the optimal value is $p^* = +\infty$ for infeasible minimization problems, while $p^* = -\infty$ for infeasible maximization problems.

1.1.8 Feasibility Problems

Sometimes an objective function is not provided. This means that we are just interested in finding a feasible point, or determine that the problem is infeasible.

In this case, we set f_0 to be a constant to reflect the fact that we are indifferent to the choice of a point $\boldsymbol{x}$, as long as it is feasible.

1.1.9 Solution to an optimization problem

The optimal value p^* is *attained* if there exists a feasible $\boldsymbol{x}^*$ such that

$$f_0(\boldsymbol{x}^*) = p^*.$$

1.1.9.1 Optimal Set

Definition 1.1.2 (Optimal Set)**.** The **optimal set** is defined as

$$\mathcal{X}_{\mathsf{OPT}} = \{\boldsymbol{x} \in \mathbb{R}^n \mid f_0(\boldsymbol{x}) = p^*, f(\boldsymbol{x}) \leq 0, i = 1, \dots, m\},$$

or equivalently,

$$\mathcal{X}_{\mathsf{OPT}} = \arg\min_{\boldsymbol{x}\in\mathcal{X}} f_0(\boldsymbol{x}).$$

A point $\boldsymbol{x}$ is **optimal** if $\boldsymbol{x} \in \mathcal{X}_{\mathsf{OPT}}$.

1.1.9.2 Empty Optimal Set

The optimal set can be empty for two reasons:

1. The problem is infeasible.
2. The optimal value is only reached in the limit.
   - For example, the problem
     $$p^* = \min_x e^{-x}$$
     has no optimal points because $p^* = 0$ is only reached in the limit for $x \to +\infty$.
   - Another example is when constraints include strict inequalities:
     $$p^* = \min_x x \quad \text{s.t } 0 < x \leq 1.$$
     In this case, $p^* = 0$ but cannot be attained by any $\boldsymbol{x}$ that satisfies the constraints.

1.1.9.3 Sub-optimality

Definition 1.1.3 (Suboptimal)**.** We say that a point $\boldsymbol{x}$ is **ϵ-suboptimal** for a problem if it is feasible, and satisfies

$$p^* \leq f_0(\boldsymbol{x}) \leq p^* + \epsilon.$$

In other words, $\boldsymbol{x}$ is ϵ-close to p^*.

1.1.9.4 Local vs. global optimal points

Definition 1.1.4 (Locally optimal)**.** A point $\boldsymbol{z}$ is **locally optimal** if there exist a value $R > 0$ such that $\boldsymbol{z}$ is optimal for problem

$$\min_{\boldsymbol{x}} f_0(\boldsymbol{x}) \text{ s.t. } f_i(\boldsymbol{x}) \leq 0, \quad i = 1, \dots, m, \quad |x_i - z_i| \leq R, \quad i = 1, \dots, n.$$

1.1.10 Problem Transformations

Sometimes we can cast a problem in a tractable formulation. For example, consider

$$\max_{\boldsymbol{x}} x_1 x_2^3 x_3 \quad \text{s.t.} \quad x_i \geq 0, \quad i = 1, 2, 3, \quad x_1 x_2 \leq 2, \quad x_2^2 x_3 \leq 1$$

can be transformed into the following by taking the log, in terms of $z_i = \log x_i$:

$$\max_{\boldsymbol{z}} z_1 + 3z_2 + z_3 \quad \text{s.t.} \quad z_1 + z_2 \leq \log 2, \quad 2z_2 + z_3 \leq 0.$$

Now the objective function and the constraints are all linear.

1.2 Convex Problems

Convex optimization problems are problems where the objective and constraint functions have the special property of **convexity**.

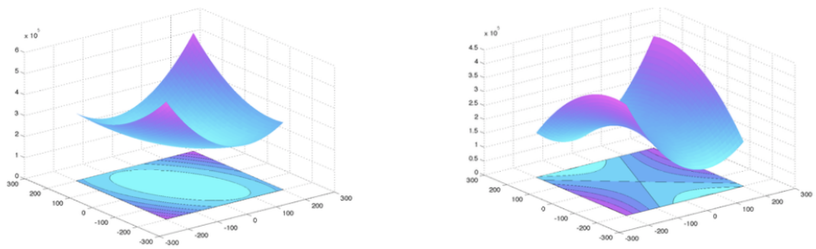

Figure 1.4: *Left.* Convex function. *Right.* Non-convex function.

For a convex function, any local minimum is global.

1.2.1 Special convex models

Convex optimization problems with special structure:

- Least-Squares (LS)
- Linear Programs (LP)
- Convex Quadratic Programs (QP)
- Geometric Programs (GP)
- Second-order Cone Programs (SOCP)
- Semi-definite Programs (SDP).

1.3 Non-convex Problems

- **Boolean/integer optimization**: some variables are constrained to be Boolean or integers. Convex optimization can be used for getting good approximations.
- **Cardinality-constrained problems**: we seek to bound the number of non-zero elements in a vector variable. Convex optimization can be used for getting good approximations.
- **Non-linear programming**: usually non-convex problems with differentiable objective and functions. Algorithms provide only local minima.

Remark. Most non-convex problems are hard.

2

Vectors and Functions

2.1 Basics

Definition 2.1.1 (Vector)**.** A **vector** is a collection of numbers, arranged in a column or a row, representing the **coordinates** of a point in n-dimensional space. We write vectors in column format:

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix},$$

where each element x_i is the $\boldsymbol{i}$**-th component** of vector $\boldsymbol{x}$ and n is the **dimension** of $\boldsymbol{x}$. If $\boldsymbol{x}$ is a *real* vector, then we write $\boldsymbol{x} \in \mathbb{R}^n$. If $\boldsymbol{x}$ is a *complex* vector, then we write $\boldsymbol{x} \in \mathbb{C}^n$.

Definition 2.1.2 (Transpose)**.** The **transpose** of a vector $\boldsymbol{x}$ is defined as

$$\boldsymbol{x}^\top = \begin{bmatrix} x_1 & x_2 & \cdots & x_n \end{bmatrix}$$

and the transpose of the transpose of $\boldsymbol{x}$ is itself, i.e., $\boldsymbol{x}^{\top\top} = \boldsymbol{x}$.

2.2 Vector Spaces

Definition 2.2.1 (Vector Space)**.** A **vector space** $\mathcal{V}$ is a set of vectors on which two operations: **vector addition** and **scalar multiplication**, are defined.

2.2.1 Subspaces and Span

Definition 2.2.2 (Subspace)**.** A nonempty subset $\mathcal{S}$ of a vector space $\mathcal{V}$ is a **subspace** of $\mathcal{V}$ if, for $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{S}$ and any scalars $\alpha, \beta \in \mathbb{R}$,

$$\alpha\boldsymbol{x} + \beta\boldsymbol{y} \in \mathcal{S}.$$

In other words, $\mathcal{S}$ is *closed* under addition and scalar multiplication.

Definition 2.2.3 (Linear Combination)**.** A **linear combination** of a set of vectors $S = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}\}$ in a vector space $\mathcal{X}$ is a vector

$$\boldsymbol{x} = \sum_{i=1}^{m} \alpha\boldsymbol{x}^{(i)},$$

where each α_i is a given scalar.

Definition 2.2.4 (Span)**.** The **span** of a set of vectors $S = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}\}$ in a vector space $\mathcal{X}$ is the set of all vectors that is a linear combination of that set of vectors

$$\text{span}(S) = \left\{\boldsymbol{x} \mid \exists\alpha_1, \ldots, \alpha_m \text{ s.t. } \boldsymbol{x} = \sum_{i=1}^{m} \alpha_i\boldsymbol{x}^{(i)}\right\}.$$

Definition 2.2.5 (Direct Sum)**.** Given two subspaces $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$, the **direct sum** of $\mathcal{X}, \mathcal{Y}$, denoted by $\mathcal{X} \oplus \mathcal{Y}$, is the set of vectors of the form $\boldsymbol{x} + \boldsymbol{y}$, where $\boldsymbol{x} \in \mathcal{X}, \boldsymbol{y} \in \mathcal{Y}$. The direct sum is itself a subspace.

2.2.2 Bases and Dimensions

Definition 2.2.6 (Linearly Independent)**.** A set of vectors $\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}$ in a vector space $\mathcal{X}$ is **linearly independent** if

$$\sum_{i=1}^{m} \alpha_i\boldsymbol{x}^{(i)} = 0 \implies \alpha_1 = \ldots = \alpha_n = 0.$$

Definition 2.2.7 (Basis)**.** Given a subspace of $\mathcal{S}$ of a vector space $\mathcal{X}$, a **basis** of $\mathcal{S}$ is a set $\mathcal{B}$ of vectors of minimal cardinality, such that $\text{span}(\mathcal{B}) = \mathcal{S}$.

Definition 2.2.8 (Dimension)**.** The **dimension** of a subspace is the cardinality of a basis of that subspace.

If we have a basis $\{\boldsymbol{x}^{(1)}, \dots, \boldsymbol{x}^{(d)}\}$ for a subspace $\mathcal{S}$, then any element in the subspace can be expressed as a linear combination of the elements in the basis. That is, any $\boldsymbol{x} \in \mathcal{S}$ can be written as

$$\boldsymbol{x} = \sum_{i=1}^{d} \alpha_i \boldsymbol{x}^{(i)},$$

for some scalars α_i.

2.2.3 Affine Sets

Definition 2.2.9 (Affine Set)**.** An **affine set** is a set of the form

$$\mathcal{A} = \left\{\boldsymbol{x} \in \mathcal{X} \mid \boldsymbol{x} = \boldsymbol{v} + \boldsymbol{x}^{(0)}, \boldsymbol{v} \in \mathcal{V}\right\}$$

where $\boldsymbol{x}^{(0)}$ is a given point and $\mathcal{V}$ is a given subspace of $\mathcal{X}$. Subspaces are just affine spaces containing the origin.

Geometric interpretation: An affine set is a flat plane passing through $\boldsymbol{x}^{(0)}$.
The dimension of an affine set $\mathcal{A}$ is defined as the dimension of its generating subspace $\mathcal{V}$.

2.2.4 Euclidean Length

Definition 2.2.10 (Euclidean Length)**.** The **Euclidean length** of a vector $\boldsymbol{x} \in \mathbb{R}^n$ is defined as

$$\|\boldsymbol{x}\|_2 \doteq \sqrt{\sum_{i=1}^{n} x_i^2}.$$

2.2.5 Norms

Definition 2.2.11 (Norm)**.** A **norm** on a vector space $\mathcal{X}$ is a real-valued function with special properties that maps any element $\boldsymbol{x} \in \mathcal{X}$ into a real number $\|\boldsymbol{x}\|$.

Definition 2.2.12. A function from $\mathcal{X}$ to $\mathbb{R}$ is a **norm**, if

- $\forall \boldsymbol{x} \in \mathcal{X}, \|\boldsymbol{x}\| \geq 0$ and $\|\boldsymbol{x}\| = 0$ if and only if $\boldsymbol{x} = \mathbf{0}$;
- $\forall \boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}, \|\boldsymbol{x} + \boldsymbol{y}\| \leq \|\boldsymbol{x}\| + \|\boldsymbol{y}\|$ (**triangle inequality**);
- $\forall \boldsymbol{x} \in \mathcal{X}, \|\alpha \boldsymbol{x}\| = |\alpha| \|\boldsymbol{x}\|$ for any scalar α.

Definition 2.2.13 (ℓ_p norms)**.** **ℓ_p norms** are defined as

$$\|\boldsymbol{x}\|_p \doteq \left(\sum_{i=1}^n |x_i|^p\right)^{1/p} \qquad 1 \leq p < \infty.$$

For $p = 2$, we have the **Euclidean length**

$$\|\boldsymbol{x}\|_2 \doteq \sqrt{\sum_{i=1}^n x_i^2},$$

or $p = 1$ we get the **sum-of-absolute-values length**

$$\|\boldsymbol{x}\|_1 \doteq \sum_{i=1}^n |x_i|.$$

The limit case $p = \infty$ defines the **ℓ_∞ norm** (**max absolute value norm**, or **Chebyshev norm**)

$$\|\boldsymbol{x}\|_\infty \doteq \max_{i=1,\ldots,n} |x_i|.$$

The cardinality of a vector $\boldsymbol{x}$ is called the **ℓ_0** (**pseudo**) **norm** and denoted by $\|\boldsymbol{x}\|_0$.

2.3 Inner Product

Definition 2.3.1 (Inner Product)**.** An **inner product** on a real vector space $\mathcal{X}$ is a real-valued function which maps any pair of elements $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$ into a scalar denoted as $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$. It satisfies the following axioms: for any $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathcal{X}$ and scalar α

(i) $\langle \boldsymbol{x}, \boldsymbol{x} \rangle \geq 0$;

(ii) $\langle \boldsymbol{x}, \boldsymbol{x} \rangle = 0$ if and only if $\boldsymbol{x} = \mathbf{0}$;

(iii) $\langle \boldsymbol{x} + \boldsymbol{y}, \boldsymbol{z} \rangle = \langle \boldsymbol{x}, \boldsymbol{z} \rangle + \langle \boldsymbol{y}, \boldsymbol{z} \rangle$;

(iv) $\langle \alpha \boldsymbol{x}, \boldsymbol{y} \rangle = \alpha \langle \boldsymbol{x}, \boldsymbol{y} \rangle$;

(v) $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \langle \boldsymbol{y}, \boldsymbol{z} \rangle$.

Definition 2.3.2 (Standard Inner Product)**.** The **standard inner product**, also called the **dot product** is defined as

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \boldsymbol{x}^\top \boldsymbol{y} = \sum_{i=1}^n x_i y_i.$$

An inner product naturally induces an associated **norm**: $\|\boldsymbol{x}\| = \sqrt{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}$.

2.3.1 Angle between vectors

The angle between $\boldsymbol{x}$ and $\boldsymbol{y}$ is defined via the relation

$$\cos\theta = \frac{\boldsymbol{x}^\top \boldsymbol{y}}{\|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2}.$$

There is a right angle between $\boldsymbol{x}$ and $\boldsymbol{y}$ when $\boldsymbol{x}^\top\boldsymbol{y} = 0$, i.e., $\boldsymbol{x}$ and $\boldsymbol{y}$ are **orthogonal**.
When $\theta = 0°$, or $\pm 180°$, then $\boldsymbol{y} = \alpha\boldsymbol{x}$ for some scalar α, i.e. $\boldsymbol{x}$ and $\boldsymbol{y}$ are **parallel**. Then $|\boldsymbol{x}^\top\boldsymbol{y}|$ achieves its **maximum value** $|\alpha|\|\boldsymbol{x}\|_2^2$.

2.3.2 Cauchy-Schwartz and Hölder Inequality

Theorem 2.3.3 (Cauchy-Schwartz's Inequality)**.** For any vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, it holds that

$$|\langle \boldsymbol{x}, \boldsymbol{y}\rangle| = |\boldsymbol{x}^\top\boldsymbol{y}| \leq \|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2,$$

Proof. Note that $|\cos\theta| \leq 1$, then using the angle equation, we have

$$|\cos\theta| = \frac{|\boldsymbol{x}^\top\boldsymbol{y}|}{\|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2} \leq 1 \implies |\boldsymbol{x}^\top\boldsymbol{y}| \leq \|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2.$$

$\square$

Theorem 2.3.4 (Hölder's Inequality)**.** For any vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ and for any $p, q \geq 1$ such that $1/p + 1/q = 1$, it holds that

$$|\langle \boldsymbol{x}, \boldsymbol{y}\rangle| = |\boldsymbol{x}^\top\boldsymbol{y}| \leq \sum_{i=1}^n |x_i y_i| \leq \|\boldsymbol{x}\|_p\|\boldsymbol{y}\|_q.$$

2.3.3 Maximization of inner product over norm balls

Given a nonzero vector $\boldsymbol{y} \in \mathbb{R}^n$, we want to find some vector $\boldsymbol{x} \in \mathcal{B}_p$ (the unit ball in ℓ_p norm) that maximizes the inner product $\boldsymbol{x}^\top\boldsymbol{y}$, i.e., we want to solve the following:

$$\max_{\|\boldsymbol{x}\|_p \leq 1} \boldsymbol{x}^\top\boldsymbol{y}.$$

If the level set $\alpha = 0$, then we are solving for

$$\boldsymbol{x}^\top\boldsymbol{y} = 0,$$

which are the set of vectors that are on a line that is orthogonal to $\boldsymbol{y}$ and passes through the origin. However, if we have $\neq 0$, then we have

$$\boldsymbol{x}^\top\boldsymbol{y} = \alpha \implies \boldsymbol{x}_0 = \alpha\frac{\boldsymbol{y}}{\|\boldsymbol{y}\|_p}.$$

Note that $\boldsymbol{x}_0$ is parallel to $\boldsymbol{y}$. Then we can rewrite the equation as

$$\boldsymbol{y}^\top(\boldsymbol{x} - \boldsymbol{x}_0) = 0.$$

Geometrically, $\boldsymbol{x} - \boldsymbol{x}_0$ represents the vectors on the that are shifted α away (towards $\boldsymbol{y}$ if $\alpha > 0$ and away from $\boldsymbol{y}$ otherwise) and are orthogonal to $\boldsymbol{y}$.

Question. What is the distance (margin) between the two separating hyperplanes $\boldsymbol{w}^\top \boldsymbol{x} + b = 1$ and $\boldsymbol{w}^\top \boldsymbol{x} + b = -1$?

Answer. $\frac{2}{\|\boldsymbol{w}\|_2}$. (why?)

2.4 Orthogonality and Orthonormality

2.4.1 Orthogonal Vectors

Definition 2.4.1 (Orthogonal)**.** Two vectors $\boldsymbol{x}, \boldsymbol{y}$ in an inner product space $\mathcal{X}$ are **orthogonal** if $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = 0$, i.e., $\boldsymbol{x} \perp \boldsymbol{y}$.

Definition 2.4.2 (Mutually Orthogonal)**.** Nonzero vectors $\boldsymbol{x}^{(1)}, \dots, \boldsymbol{x}^{(d)}$ are said to be **mutually orthogonal** if $\left\langle \boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)} \right\rangle = 0$ whenever $i \neq j$. In other words, each vector is orthogonal to all other vectors in the collection.

Proposition 3. Mutually orthogonal vectors are linearly independent.

Proof. Suppose for the sake of contradiction that $\boldsymbol{x}^{(1)}, \dots, \boldsymbol{x}^{(d)}$ are orthogonal but linearly dependent vectors. Then this implies that there exist scalars $\alpha_1 \dots, \alpha_d$ that are not all identically zero, such that

$$\sum_{i=1}^{d} \alpha_i \boldsymbol{x}^{(i)} = 0.$$

□

Taking the linear product of both sides of this equation with $\boldsymbol{x}^{(j)}$ for $j = 1, \dots, d$, we have

$$\left\langle \sum_{i=1}^{d} \alpha_i \boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)} \right\rangle = 0.$$

Since

$$\left\langle \sum_{i=1}^{d} \alpha_i \boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)} \right\rangle = 0,$$

this means that $\alpha_i = 0$ for all $i = 1, \dots, d$, hence a contradiction.

Definition 2.4.4 (Orthonormal)**.** A collection of vectors $S = \left\{ \boldsymbol{x}^{(1)}, \dots, \boldsymbol{x}^{(d)} \right\}$ is **orthonormal** if, for $i, j = 1, \dots, d$

$$\left\langle \boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)} \right\rangle = \begin{cases} 0 & \text{if } i \neq j; \\ 1 & \text{if } i = j, \end{cases}$$

i.e., S is orthonormal if every element has **unit norm**, and all elements are **orthogonal** to each other. A collection of orthonormal vectors S forms an **orthonormal basis** for the span of S.

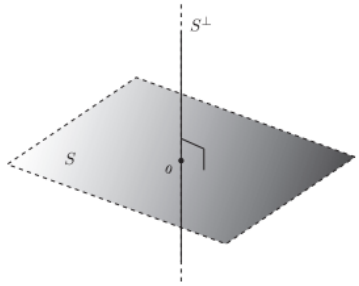

Figure 2.1: Orthogonal complement of $\mathcal{S}$.

2.4.2 Orthogoanl Complement

Definition 2.4.5 (Orthogonal Complement)**.** The set of vectors in $\mathcal{X}$ that are orthogonal to $\mathcal{S}$ is called the **orthogonal complement** of $\mathcal{S}$, denoted by $\mathcal{S}^\perp$.

Theorem 2.4.6 (Orthogonal Decomposition)**.** If $\mathcal{S}$ is a subspace of an inner product space $\mathcal{X}$, then any vector $\boldsymbol{x} \in \mathcal{X}$ can be written in an **unique** way as the sum of an element in $\mathcal{S}$ and one in the orthogonal complement $\mathcal{S}^\perp$:

$$\mathcal{X} = \mathcal{S} \oplus \mathcal{S}^\perp$$

for any subspace $\mathcal{S} \subseteq \mathcal{X}$.

Proof. □

2.4.3 Projections

Definition 2.4.7 (Projection)**.** Given a vector $\boldsymbol{x}$ in an inner product space $\mathcal{X}$ and a closed set $\mathcal{S} \subseteq \mathcal{X}$, the projection of $\boldsymbol{x}$ onto $\mathcal{S}$, denoted as $\Pi_{\mathcal{S}}(\boldsymbol{x})$, is defined as the point in $\mathcal{S}$ at minimal distance from $\boldsymbol{x}$:

$$\Pi_{\mathcal{S}}(\boldsymbol{x}) = \arg\min_{\boldsymbol{y} \in \mathcal{S}} \|\boldsymbol{y} - \boldsymbol{x}\|,$$

called **Euclidean projection**.

Theorem 2.4.8 (Projection Theorem)**.** Let $\mathcal{X}$ be an inner product space, let $\boldsymbol{x}$ be a given element in $\mathcal{X}$, and let $\mathcal{S}$ be a subspace of $\mathcal{X}$. Then, there exists a unique vector $\boldsymbol{x}^* \in \mathcal{S}$ which is solution to the problem

$$\min_{\boldsymbol{y} \in \mathcal{S}} \|\boldsymbol{y} - \boldsymbol{x}\|$$

Moreover, a necessary and sufficient condition for $\boldsymbol{x}^*$ being the optimal solution for this problem is that

$$\boldsymbol{x}^* \in \mathcal{S}, \quad (\boldsymbol{x} - \boldsymbol{x}^*) \perp \mathcal{S}.$$

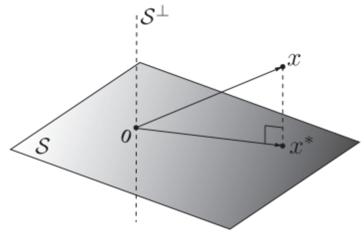

Figure 2.2: Projection onto a subspace.

Proof. □

Theorem 2.4.9 (Projection on affine set)**.** Let $\mathcal{X}$ be an inner product space, let x be a given element in $\mathcal{X}$, and let $\mathcal{A} = x^{(0)} + \mathcal{S}$ be the affine set obtained by translating a given subspace $\mathcal{S}$ by a given vector $x^{(0)}$. Then, there exists a unique vector $x^* \in \mathcal{A}$ which is solution to the problem

$$\min_{y \in \mathcal{A}} \|y - x\|$$

Moreover, a necessary and sufficient condition for x^* to be the optimal solution for this problem is that

$$x^* \in \mathcal{A}, \quad (x - x^*) \perp \mathcal{S}.$$

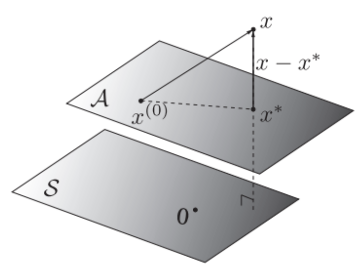

Figure 2.3: Projection on affine set.

Proof. □

2.4.3.1 Euclidean projection of a point onto a line

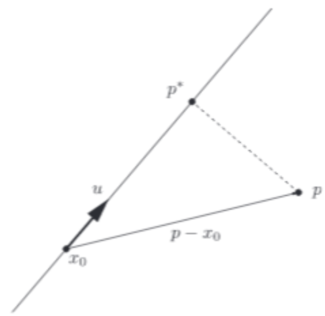

Figure 2.4: Euclidean projection of a point onto a line.

Let $\boldsymbol{p} \in \mathbb{R}^n$ be a given point. We want to compute the Euclidean projection $\boldsymbol{p}^*$ of $\boldsymbol{p}$ onto a line $L = \{\boldsymbol{x}_0 + \text{span}(\boldsymbol{u})\}$, where $\|\boldsymbol{u}\|_2 = 1$:

$$\boldsymbol{p}^* = \arg\min_{\boldsymbol{x} \in L} \|\boldsymbol{x} - \boldsymbol{p}\|_2.$$

Since any point $\boldsymbol{x} \in L$ can be written as $\boldsymbol{x} = \boldsymbol{x}_0 + \boldsymbol{v}$, for some $\boldsymbol{v} \in \text{span}(\boldsymbol{u})$, the problem is equivalent to finding a value $\boldsymbol{v}^*$ such that

$$\boldsymbol{v}^* = \arg\min_{\boldsymbol{v} \in \text{span}(\boldsymbol{u})} \|\boldsymbol{v} - (\boldsymbol{p} - \boldsymbol{x}_0)\|_2.$$

2.4.3.2 Euclidean projection of a point onto an hyperplane

A **hyperplane** is an affine set defined as

$$H = \left\{ \boldsymbol{z} \in \mathbb{R}^n \mid \boldsymbol{a}^\top \boldsymbol{z} = b \right\}$$

where $\boldsymbol{a} \neq 0$ is called a **normal direction** of the hyperplane, since for any two vectors $\boldsymbol{z}_1, \boldsymbol{z}_2 \in H$ it holds that $(\boldsymbol{z}_1 - \boldsymbol{z}_2) \perp \boldsymbol{a}$.

Our goal is that given $\boldsymbol{p} \in \mathbb{R}^n$ we want to determine the Euclidean projection $\boldsymbol{p}^*$ of $\boldsymbol{p}$ onto H.

The projection theorem requires $\boldsymbol{p} - \boldsymbol{p}^*$ to be orthogonal to H. Since $\boldsymbol{a}$ is a direction orthogonal to H, the condition $(\boldsymbol{p} - \boldsymbol{p}^*) \perp H$ is equivalent to saying that $\boldsymbol{p} - \boldsymbol{p}^* = \alpha \boldsymbol{a}$, for some $\alpha \in \mathbb{R}$.

To find α, consider that $\boldsymbol{p}^* \in H$, thus $\boldsymbol{a}^\top \boldsymbol{p}^* = b$, then consider the optimality condition

$$\boldsymbol{p} - \boldsymbol{p}^* = \alpha a$$

and multiply it on the left by $\boldsymbol{a}^\top$, obtaining

$$\boldsymbol{a}^\top \boldsymbol{p} - b = \alpha \|\boldsymbol{a}\|^2$$

whereby

$$\alpha = \frac{\boldsymbol{a}^\top \boldsymbol{p} - b}{\|\boldsymbol{a}\|_2^2}$$

and

$$\boldsymbol{p}^* = \boldsymbol{p} - \frac{\boldsymbol{a}^\top \boldsymbol{p} - b}{\|\boldsymbol{a}\|_2^2}\boldsymbol{a}.$$

The distance from $\boldsymbol{p}$ to H is

$$\|\boldsymbol{p} - \boldsymbol{p}^*\|_2 = |\alpha| \cdot \|\boldsymbol{a}\|_2 = \frac{|\boldsymbol{a}^\top \boldsymbol{p} - b|}{\|\boldsymbol{a}\|_2}.$$

2.4.3.3 Projection on a vector span

Suppose we have a basis for a subspace $\mathcal{S} \subseteq \mathcal{X}$, that is

$$\mathcal{S} = \operatorname{span}\left(\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(d)}\right).$$

Given $\boldsymbol{x} \in \mathcal{X}$, the Projection Theorem states that the unique projection $\boldsymbol{x}^*$ of $\boldsymbol{x}$ onto $\mathcal{S}$ is characterized by $(\boldsymbol{x} - \boldsymbol{x}^*) \perp \mathcal{S}$.

Since $\boldsymbol{x}^* \in \mathcal{S}$, we can write $\boldsymbol{x}^*$ as some (unknown) linear combination of the elements in the basis of $\mathcal{S}$, that is

$$\boldsymbol{x}^* = \sum_{i=1}^{d} \alpha_i \boldsymbol{x}^{(i)}$$

Then $(\boldsymbol{x} - \boldsymbol{x}^*) \perp \mathcal{S} \Leftrightarrow \langle \boldsymbol{x} - \boldsymbol{x}^*, \boldsymbol{x}^{(k)} \rangle = 0, k = 1, \ldots, d:$

$$\sum_{i=1}^{d} \alpha_i \left\langle \boldsymbol{x}^{(k)}, \boldsymbol{x}^{(i)} \right\rangle = \left\langle \boldsymbol{x}^{(k)}, \boldsymbol{x} \right\rangle, \quad k = 1, \ldots, d$$

Solving this system of linear equations (Gram equations) provides the coefficients α, and hence the desired $\boldsymbol{x}^*$.

2.5 Functions and Maps

Definition 2.5.1 (Function)**.** A **function** takes a vector argument in $\mathbb{R}^n$, and returns a unique value in $\mathbb{R}$. We write

$$f : \mathbb{R}^n \to \mathbb{R}.$$

Definition 2.5.2 (Domain)**.** The **domain** of a function f, denoted $\mathbf{dom} f$, is defined as the set of points where the function is finite.

Definition 2.5.3 (Map)**.** **Maps** are functions that return a vector of values. We write

$$f : \mathbb{R}^n \to \mathbb{R}^m.$$

2.5.1 Sets related to functions

Definition 2.5.4 (Graph)**.** The **graph** of f is the set of input-output pairs that f can attain, that is:

$$f = \left\{(\boldsymbol{x}, f(\boldsymbol{x})) \in \mathbb{R}^{n+1} \mid \boldsymbol{x} \in \mathbb{R}^n\right\}$$

Definition 2.5.5 (Epigraph)**.** The **epigraph**, denoted f, describes the set of input-output pairs that f can achieve, as well as *anything above*:

$$f = \left\{(\boldsymbol{x}, t) \in \mathbb{R}^{n+1} \mid \boldsymbol{x} \in \mathbb{R}^n, t \geq f(\boldsymbol{x})\right\}.$$

Definition 2.5.6 (Level Set)**.** A **level set** (or **contour line**) is the set of points that achieve exactly some value for the function f. For $t \in \mathbb{R}$, the t-level set of the function f is defined as

$$C_f(t) = \{\boldsymbol{x} \in \mathbb{R}^n \mid f(\boldsymbol{x}) = t\}.$$

Definition 2.5.7 (t-sublevel set)**.** The ***t*-sublevel set** of f is the set of points that achieve at most a certain value for f :

$$L_f(t) = \{\boldsymbol{x} \in \mathbb{R}^n \mid f(\boldsymbol{x}) \leq t\}.$$

2.5.2 Linear and Affine Functions

Definition 2.5.8 (Linear)**.** A function $f : \mathbb{R}^n \to \mathbb{R}$ is **linear** if and only if

- $\forall \boldsymbol{x} \in \mathbb{R}^n, \alpha \in \mathbb{R},\ f(\alpha \boldsymbol{x}) = \alpha f(\boldsymbol{x})$;
- $\forall \boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n,\ f(\boldsymbol{x}_1 + \boldsymbol{x}_2) = f(\boldsymbol{x}_1) + f(\boldsymbol{x}_2)$.

Definition 2.5.9 (Affine)**.** A function f is **affine** if and only if the function $\tilde{f}(\boldsymbol{x}) = f(\boldsymbol{x}) - f(\mathbf{0})$ is linear (affine = linear + constant). In addition, f is affine if and only if it can be expressed as

$$f(\boldsymbol{x}) = \boldsymbol{a}^\top \boldsymbol{x} + b,$$

for some unique pair $(\boldsymbol{a}, b)$ where $\boldsymbol{a} \in \mathbb{R}^n$ and $b \in \mathbb{R}$.

For any affine function f, we can obtain $\boldsymbol{a}$ and b as follows:

$$b = f(\mathbf{0}),$$

$$a_i = f(\boldsymbol{e}_i) - b, \quad \text{for } i = 1, \ldots, n.$$

2.6 Hyperplanes and Halfspaces

Definition 2.6.1 (Hyperplane)**.** A **hyperplane** in $\mathbb{R}^n$ is a set of the form

$$\mathcal{H} = \left\{ \boldsymbol{x} \in \mathbb{R}^n \mid \boldsymbol{a}^\top \boldsymbol{x} = b \right\},$$

where $\boldsymbol{a} \in \mathbb{R}^n, \boldsymbol{a} \neq \boldsymbol{0}$, and $b \in \mathbb{R}$ are given.

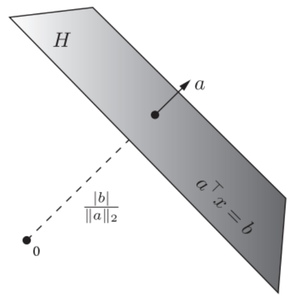

Figure 2.5: Hyperplane.

Definition 2.6.2 (Halfspace)**.** A hyperplane $\mathcal{H}$ separates the whole space in two regions called **halfspaces** ($\mathcal{H}_-$ is a **closed halfspace**, $\mathcal{H}$ is an **open halfspace**).

$$\mathcal{H}_- = \left\{ \boldsymbol{x} \mid \boldsymbol{a}^\top \boldsymbol{x} \leq b \right\}, \quad \mathcal{H}_{++} = \left\{ \boldsymbol{x} \mid \boldsymbol{a}^\top \boldsymbol{x} > b \right\}.$$

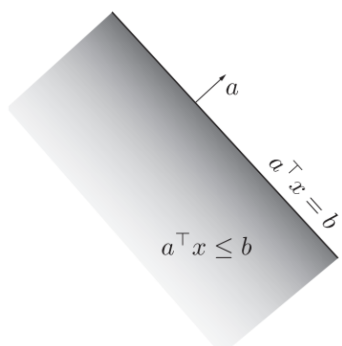

Figure 2.6: Halfspace.

2.7 Gradients

Definition 2.7.1 (Gradient)**.** The **gradient** of a function $f : \mathbb{R}^n \to \mathbb{R}$ at a point $\boldsymbol{x}$ where f is differentiable, denoted with $\nabla f(\boldsymbol{x})$, is a column vector of first derivatives of f with respect to $x_1, \ldots, x_n$

$$\nabla f(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial f(\boldsymbol{x})}{\partial x_1} & \cdots & \frac{\partial f(\boldsymbol{x})}{\partial x_n} \end{bmatrix}^\top$$

An affine function $f : \mathbb{R}^n \to \mathbb{R}$, represented as $f(\boldsymbol{x}) = \boldsymbol{a}^\top \boldsymbol{x} + b$, has a very simple gradient: $\nabla f(\boldsymbol{x}) = \boldsymbol{a}$.

Example 2.7.2. The distance function $\rho(x) = \|\boldsymbol{x} - \boldsymbol{p}\|_2 = \sqrt{\sum_{i=1}^n (x_i - p_i)^2}$ has gradient

$$\nabla \rho(\boldsymbol{x}) = \frac{1}{\|\boldsymbol{x} - \boldsymbol{p}\|_2}(\boldsymbol{x} - \boldsymbol{p}).$$

2.7.1 Affine approximation of non-linear functions

A non-linear function $f : \mathbb{R}^n \to \mathbb{R}$ can be approximated locally via an affine function, using a **first-order Taylor series expansion**:

Theorem 2.7.3 (First-order Taylor Series Expansion)**.** If f is differentiable at point $\boldsymbol{x}_0$, then for all points $\boldsymbol{x}$ in a neighborhood of $\boldsymbol{x}_0$, we have that

$$f(\boldsymbol{x}) = f(\boldsymbol{x}_0) + \nabla f(\boldsymbol{x}_0)^\top (\boldsymbol{x} - \boldsymbol{x}_0) + \epsilon(\boldsymbol{x})$$

where the error term $\epsilon(\boldsymbol{x})$ goes to zero faster than first order, as $\boldsymbol{x} \to \boldsymbol{x}_0$, that is

$$\lim_{\boldsymbol{x} \to \boldsymbol{x}_0} \frac{\epsilon(\boldsymbol{x})}{\|\boldsymbol{x} - \boldsymbol{x}_0\|_2} = 0$$

In practice, this means that for $\boldsymbol{x}$ sufficiently close to $\boldsymbol{x}_0$, we can write the approximation

$$\boxed{f(\boldsymbol{x}) \simeq f(\boldsymbol{x}_0) + \nabla f(\boldsymbol{x}_0)^\top (\boldsymbol{x} - \boldsymbol{x}_0)}.$$

2.7.2 Geometric interpretation of the gradient

Geometrically, the gradient of f at a point $\boldsymbol{x}_0$ is a vector $\nabla f(\boldsymbol{x}_0)$ perpendicular to the contour line of f at level $\alpha = f(\boldsymbol{x}_0)$, pointing from $\boldsymbol{x}_0$ outwards the α-sublevel set (i.e., it points towards higher values of the function).

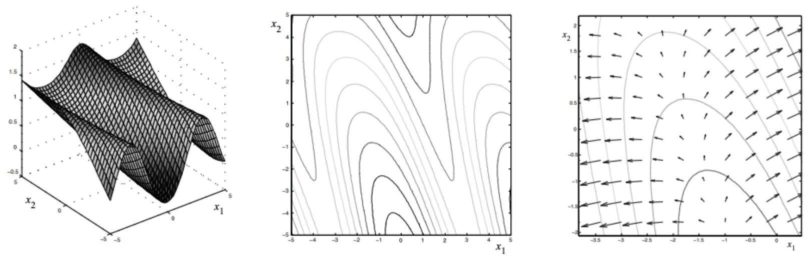

Figure 2.7: *Left.* Graph of a function. *Center.* Its contour lines. *Right.* Gradient vectors (arrows) at some grid points.

The gradient $\nabla f\left(\boldsymbol{x}_0\right)$ also represents the direction along which the function has the **maximum rate of increase** (**steepest ascent direction**).

Let $\boldsymbol{v}$ be a unit direction vector (i.e., $\|\boldsymbol{v}\|_2=1$), let $\epsilon \geq 0$, and consider moving away at distance ϵ from $\boldsymbol{x}_0$ along direction $\boldsymbol{v}$, that is, consider a point $\boldsymbol{x}=\boldsymbol{x}_0+\epsilon \boldsymbol{v}$. We have that

$$f\left(\boldsymbol{x}_0+\epsilon \boldsymbol{v}\right) \simeq f\left(\boldsymbol{x}_0\right)+\epsilon \nabla f\left(\boldsymbol{x}_0\right)^{\top} \boldsymbol{v}, \text{ for } \epsilon \rightarrow 0,$$

equivalently,

$$\lim_{\epsilon \rightarrow 0} \frac{f\left(\boldsymbol{x}_0+\epsilon \boldsymbol{v}\right)-f\left(\boldsymbol{x}_0\right)}{\epsilon}=\nabla f\left(\boldsymbol{x}_0\right)^{\top} \boldsymbol{v}.$$

Whenever $\epsilon>0$ and v is such that $\nabla f\left(\boldsymbol{x}_0\right)^{\top} \boldsymbol{v}>0$, then f is increasing along the direction $\boldsymbol{v}$, for small ϵ.

Remark. The inner product $\nabla f\left(\boldsymbol{x}_0\right)^{\top} \boldsymbol{v}$ measures the rate of variation of f at $\boldsymbol{x}_0$, along direction $\boldsymbol{v}$, and it is called the **directional derivative** of f along $\boldsymbol{v}$.

If $\boldsymbol{v}$ is orthogonal to $\nabla f(\boldsymbol{x}_0)$, the rate of variation is zero: along such a direction the function value remains constant. Contrary, the rate of variation is maximal when $\boldsymbol{v}$ is parallel to $\nabla f\left(\boldsymbol{x}_0\right)$, hence along the normal direction to the contour line at $\boldsymbol{x}_0$.

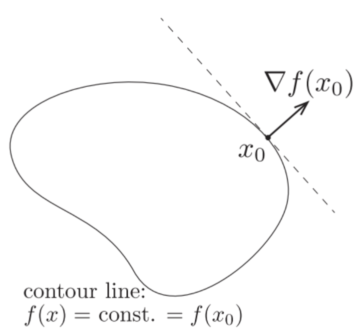

Figure 2.8: The gradient $\nabla f(\boldsymbol{x}_0)$ is normal to the contour line of f at $\boldsymbol{x}_0$, and defines the direction of maximum increase rate.

3

Matrices and Linear Maps

3.1 Matrix Basics

Definition 3.1.1 (Matrix)**.** A **matrix** is a collection of numbers, arranged in columns and rows in a tabular format:

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix},$$

where m is the number of rows and n is the number of columns. If A contains only real elements, we write $A \in \mathbb{R}^{m,n}$ and $A \in \mathbb{C}^{m,n}$ if A contains complex elements.

Definition 3.1.2 (Transpose)**.** The **transposition** operation is defined as

$$A^\top_{ij} = A_{ji},$$

where A_{ij} is the element of A positioned in row i and column j.

3.1.1 Matrix Products

Definition 3.1.3 (Matrix Multiplication)**.** Two matrices can be multiplied if conformably sized, i.e., if $A \in \mathbb{R}^{m,n}$ and $B \in \mathbb{R}^{n,p}$, then the matrix product $AB \in \mathbb{R}^{m,p}$ is defined as a matrix whose (i, j)-th entry is

$$(AB)_{ij} = \sum_{k=1}^{n} A_{ik} B_{kj}.$$

Remark. The matrix product is **non-commutative**, i.e., $AB \neq BA$.

Definition 3.1.4 (Identity Matrix)**.** The $n \times n$ **identity matrix** (denoted I_n, or I), is a matrix with all zero elements, except for the elements on the diagonal, which are equal to one. This matrix satisfies $AI_n = A$ for every matrix A with n columns, and $I_n B = B$ for every matrix B with n rows.

3.1.2 Matrix-vector Product

Definition 3.1.5 (Matrix-vector Product)**.** Let $A \in \mathbb{R}^{m,n}$ be a matrix with columns $\boldsymbol{a}_1, \dots, \boldsymbol{a}_n \in \mathbb{R}^m$ and $\boldsymbol{b} \in \mathbb{R}^n$ a vector. The **matrix-vector product** is defined as

$$A\boldsymbol{b} = \sum_{k=1}^{n} \boldsymbol{a}_k b_k, \quad A \in \mathbb{R}^{m,n}, \boldsymbol{b} \in \mathbb{R}^n$$

which is a linear combination of the columns of A, using the elements in $\boldsymbol{b}$ as coefficients.

Similarly, we can multiply matrix $A \in \mathbb{R}^{m,n}$ on the left by (the transpose of) vector $\boldsymbol{c} \in \mathbb{R}^m$ as follows:

$$\boldsymbol{c}^\top A = \sum_{k=1}^{m} c_k \alpha_k^\top, \quad A \in \mathbb{R}^{m,n}, \boldsymbol{c} \in \mathbb{R}^m$$

forming a linear combination of the rows α_k of A, using the elements in $\boldsymbol{c}$ as coefficients.

3.1.3 Matrix Representations

A matrix $A \in \mathbb{R}^{m,n}$ can be expressed in the following two forms:

$$A = \begin{bmatrix} \boldsymbol{a}_1 & \boldsymbol{a}_2 & \cdots & \boldsymbol{a}_n \end{bmatrix}, \text{ or } A = \begin{bmatrix} \boldsymbol{\alpha_1}^\top \\ \boldsymbol{\alpha_2}^\top \\ \vdots \\ \boldsymbol{\alpha_m}^\top \end{bmatrix},$$

where $\boldsymbol{a}_1, \dots, \boldsymbol{a}_n \in \mathbb{R}^m$ denote the columns of A, and $\boldsymbol{\alpha_1}^\top, \dots, \boldsymbol{\alpha_m}^\top \in \mathbb{R}^n$ denote the rows of A.

AB can be written as

$$AB = A \begin{bmatrix} \boldsymbol{b}_1 & \dots & \boldsymbol{b}_p \end{bmatrix} = \begin{bmatrix} A\boldsymbol{b}_1 & \dots & A\boldsymbol{b}_p \end{bmatrix}.$$

In other words, AB results from transforming each column $\boldsymbol{b}_i$ of B into $A\boldsymbol{b}_i$.
Similarly, we can also write

$$AB = \begin{bmatrix} \boldsymbol{\alpha_1}^\top \\ \vdots \\ \boldsymbol{\alpha_m}^\top \end{bmatrix} B = \begin{bmatrix} \boldsymbol{\alpha_1}^\top B \\ \vdots \\ \boldsymbol{\alpha_m}^\top B \end{bmatrix}.$$

Finally, the product AB can be given the interpretation as the sum of so-called **dyadic** matrices (matrices of rank one, of the form $\boldsymbol{a}_i \boldsymbol{\beta_i}^\top$, where $\boldsymbol{\beta_i}^\top$ denote the rows of B :

$$AB = \sum_{i=1}^{n} \boldsymbol{a}_i \boldsymbol{\beta_i}^\top, \quad A \in \mathbb{R}^{m,n}, B \in \mathbb{R}^{n,p}.$$

For any two conformably sized matrices A, B, it holds that

$$(AB)^\top = B^\top A^\top.$$

Then for a generic chain of n products, we have

$$(A_1 A_2 \cdots A_p)^\top = A_p^\top \cdots A_2^\top A_1^\top.$$

3.2 Matrices as linear maps

We can interpret matrices as linear maps (vector-valued functions), or **operators**, acting from an **input** space to an **output** space.

Recall that a map $f : \mathcal{X} \to \mathcal{Y}$ is **linear** if any points $\boldsymbol{x}$ and $\boldsymbol{z}$ in $\mathcal{X}$ and any scalars λ, μ satisfy $f(\lambda \boldsymbol{x} + \mu \boldsymbol{z}) = \lambda f(x) + \mu f(\boldsymbol{z})$.
Any linear map $f : \mathbb{R}^n \to \mathbb{R}^m$ can be represented by a matrix $A \in \mathbb{R}^{m,n}$, mapping input vectors $\boldsymbol{x} \in \mathbb{R}^n$ to output vectors $\boldsymbol{y} \in \mathbb{R}^m$:

$$y = Ax.$$

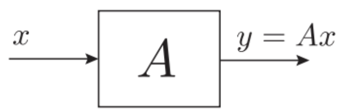

Figure 3.1: Linear map defined by a matrix A.

Affine maps are simply linear functions plus a constant term, thus any affine map $f : \mathbb{R}^n \to \mathbb{R}^m$ can be represented as

$$f(\boldsymbol{x}) = A\boldsymbol{x} + \boldsymbol{b},$$

for some $A \in \mathbb{R}^{m,n}, \boldsymbol{b} \in \mathbb{R}^m$.

3.2.1 Range, rank, and nullspace

Definition 3.2.1 (Range)**.** The **range** of a matrix A is defined as

$$\mathcal{R}(A) = \{A\boldsymbol{x} \mid \boldsymbol{x} \in \mathbb{R}^n\},$$

which is a subspace.

Definition 3.2.2 (Rank)**.** The **rank** of $\mathcal{R}(A)$, denoted by $\operatorname{rank} A$, is the **dimension** of A, which is the number of linearly independent columns of A.

Remark. The rank is also equal to the number of linearly independent rows of A; that is,

$$\operatorname{rank} A = \operatorname{rank} A^\top.$$

Thus,

$$1 \leq \operatorname{rank} A \leq \min(m, n).$$

Definition 3.2.3 (Nullspace)**.** The **nullspace** of a matrix A, denoted $\mathcal{N}(A)$ is defined as:

$$\mathcal{N}(A) = \{\boldsymbol{x} \in \mathbb{R}^n \mid A\boldsymbol{x} = \boldsymbol{0}\},$$

which is also a subspace.

Corollary 3.2.4. $\mathcal{R}(A^\top)$ and $\mathcal{N}(A)$ are mutually orthogonal subspaces, i.e., $\mathcal{N}(A) \perp \mathcal{R}(A^\top)$.

Corollary 3.2.5.

$$\mathbb{R}^n = \mathcal{N}(A) \oplus \mathcal{N}(A)^\perp = \mathcal{N}(A) \oplus \mathcal{R}(A^\top).$$

Theorem 3.2.6 (Fundamental Theorem of Linear Algebra)**.** For any given matrix $A \in \mathbb{R}^{m,n}$, it holds that $\mathcal{N}(A) \perp \mathcal{R}\left(A^\top\right)$ and $\mathcal{R}(A) \perp \mathcal{N}\left(A^\top\right)$, hence

$$\begin{aligned} \mathcal{N}(A) \oplus \mathcal{R}\left(A^\top\right) &= \mathbb{R}^n \\ \mathcal{R}(A) \oplus \mathcal{N}\left(A^\top\right) &= \mathbb{R}^m. \end{aligned}$$

Consequently, we can decompose any vector $\boldsymbol{x} \in \mathbb{R}^n$ as the sum of two vectors orthogonal to each other, one in the range of $A^\top$, and the other in the nullspace of A :

$$\boldsymbol{x} = A^\top \xi + \boldsymbol{z}, \quad \boldsymbol{z} \in \mathcal{N}(A)$$

Similarly, we can decompose any vector $\boldsymbol{w} \in \mathbb{R}^m$ as the sum of two vectors orthogonal to each other, one in the range of A, and the other in the nullspace of $A^\top$:

$$\boldsymbol{w} = A\boldsymbol{\varphi} + \boldsymbol{\zeta}, \quad \boldsymbol{\zeta} \in \mathcal{N}\left(A^\top\right).$$

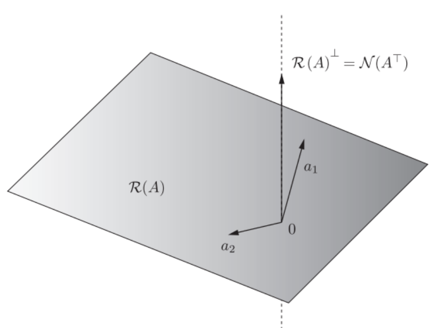

Figure 3.2: Illustration of the fundamental theorem of linear algebra in $\mathbb{R}^3$.

3.3 Determinants

Definition 3.3.1 (Determinants)**.** The **determinant** of a generic (square) matrix $A \in \mathbb{R}^{n,n}$ can be computed by defining $\det\{a\} = a$ for a scalar a, and then applying the following inductive formula (**Laplace's determinant expansion**):

$$\det\{A\} = \sum_{j=1}^{n} (-1)^{i+j} a_{ij} \det\{A_{(i,j)}\},$$

where i is any row, chosen at will, and $A_{(i,j)}$ denotes a $(n-1) \times (n-1)$ submatrix of A obtained by eliminating row i and column j from A.

$$A \in \mathbb{R}^{n,n} \text{ is singular } \iff \det\{A\} = 0 \iff \mathcal{N}(A) \text{ is not equal to } \{0\}.$$

For any square matrices $A, B \in \mathbb{R}^{n,n}$ and scalar α:

$$\begin{aligned} \det\{A\} &= \det\left\{A^\top\right\} \\ \det\{AB\} &= \det\{BA\} = \det\{A\}\det\{B\} \\ \det\{\alpha A\} &= \alpha^n \det\{A\}. \end{aligned}$$

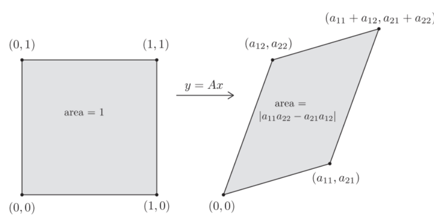

Figure 3.3: Linear mapping of the unit square. The absolute value of the determinant equals the area of the transformed unit square.

3.3.1 Matrix Inverses

If $A \in \mathbb{R}^{n,n}$ is **nonsingular** (i.e., $\det\{A\} \neq 0$), then the inverse matrix A^{-1} is defined as the unique $n \times n$ matrix such that

$$AA^{-1} = A^{-1}A = I_n.$$

If A, B are **square** and **nonsingular**, then

$$(AB)^{-1} = B^{-1}A^{-1}.$$

If A is **square** and **nonsingular**, then

$$\begin{aligned} (A^\top)^{-1} &= (A^{-1})^\top \\ \det\{A\} &= \det\{A\}^\top = \frac{1}{\det\{A^{-1}\}}. \end{aligned}$$

3.3.2 Eigenvalues and Eigenvectors

Definition 3.3.2 (Eigenvalue/Eigenvector). $\lambda \in \mathbb{C}$ is an **eigenvalue** of matrix $A \in \mathbb{R}^{n,n}$, and $\boldsymbol{u} \in \mathbb{C}^n$ is a corresponding **eigenvector**, if it holds that

$$A\boldsymbol{u} = \lambda \boldsymbol{u}, \quad \boldsymbol{u} \neq 0,$$

or equivalently, $(\lambda I_n - A)\boldsymbol{u} = 0, \boldsymbol{u} \neq 0$.

Definition 3.3.3 (Characteristic Polynomial). Eigenvalues can be characterized as those real or complex numbers that satisfy the equation

$$p(\lambda) \doteq \det(\lambda I_n - A) = 0$$

where $p(\lambda)$ is a polynomial of degree n in λ, known as the **characteristic polynomial** of A

Any matrix $A \in \mathbb{R}^{n,n}$ has n eigenvalues $\lambda_i, i = 1, \ldots, n$, counting multiplicities. To each distinct eigenvalue $\lambda_i, i = 1, \ldots, k$, there corresponds a whole subspace $\phi_i \doteq \mathcal{N}(\lambda_i I_n - A)$ of eigenvectors associated to this eigenvalue, called the **eigenspace**.

3.3.3 Diagonalizable Matrices

Theorem 3.3.4. Let $\lambda_i, i = 1, \ldots, k \leq n$ be the distinct eigenvalues of $A \in \mathbb{R}^{n,n}$. Let μ_i, $i = 1, \ldots, k$ denote the coreesponding algebraic multiplicites. Let $\phi_i = (\lambda_i I_n - A)$, and $U^{(i)} = \begin{bmatrix} u_1^(i) & \cdots & u_{\nu_i}^{(i)} \end{bmatrix}$ be a matrix containing by columns a basis of ϕ_i, being $\nu_i \doteq \dim \phi_i$. It holds that $\nu_i \leq \mu_i$ and, if $\nu_i = \mu_i, i = 1, \ldots, k$, then

$$U = \begin{bmatrix} U^{(1)} \cdots U^{(k)} \end{bmatrix}$$

is invertible, and

$$A = U \Lambda U^{-1}$$

where

$$\Lambda = \begin{bmatrix} \lambda_1 l_{\mu_1} & 0 & \cdots & 0 \\ 0 & \lambda_2 I_{\mu_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \lambda_k I_{\mu_k} \end{bmatrix}.$$

3.3.4 Matrices with special structure

- Square, diagonal triangular (upper or lower)
- Symmetric: a square matrix A such that $A = A^\top$.
- Orthogonal: a square matrix A such that $AA^\top = A^\top A = I$.
- **Dyad**: a rank-one matrix $A = \boldsymbol{u}\boldsymbol{v}^\top$, where $\boldsymbol{u}, \boldsymbol{v}$ are vectors.

- Block-structured matrices: block diagonal, block triangular, etc.

3.4 Matrix factorizations

A factorization can be interpreted as a decomposition of the map into a series of successive stages.

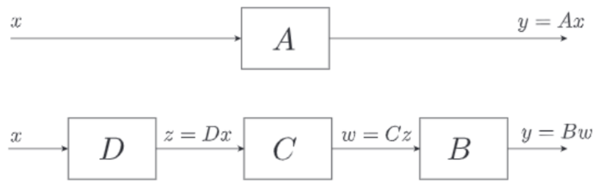

Figure 3.4: Matrix decomposition.

3.4.1 QR decomposition

Definition 3.4.1 (Orthogonal-triangular decomposition (QR)). Any square $A \in \mathbb{R}^{n,n}$ can be decomposed as

$$A = QR,$$

where Q is an orthogonal matrix, and R is an upper triangular matrix. If A is nonsingular, then the factors Q, R are uniquely defined, if the diagonal elements in R are imposed to be positive.

3.4.2 SVD

Definition 3.4.2 (Singular Value Decomposition). Any non-zero $A \in \mathbb{R}^{m,n}$ can be decomposed as

$$A = U\tilde{\Sigma}V^\top,$$

where $V \in \mathbb{R}^{n,n}$ and $U \in \mathbb{R}^{m,m}$ are orthogonal matrices, and

$$\tilde{\Sigma} = \begin{bmatrix} \Sigma & 0_{r,n-r} \\ 0_{m-r,r} & 0_{m-r,n-r} \end{bmatrix}, \quad \Sigma = \text{diag}(\sigma_1, \ldots, \sigma_r),$$

where r is the rank of A, and the scalars $\sigma_i > 0, i = 1, \ldots, r$ are called **singular values** of A. The first r columns $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_r$ of U ($\boldsymbol{v}_1, \ldots, \boldsymbol{v}_r$ of V) are called the **left/right singular vectors**, and satisfy

$$A\boldsymbol{v}_i = \sigma_i \boldsymbol{u}_i \qquad A^\top \boldsymbol{u}_i = \sigma_i \boldsymbol{v}_i, i = 1, \ldots, r.$$

3.5 Matrix Norms

A function $f : \mathbb{R}^{m,n} \to \mathbb{R}$ is a **matrix norm** if, analogously to the vector case, it satisfies three standard axioms. Namely, for all $A, B \in \mathbb{R}^{m,n}$ and all $\alpha \in \mathbb{R}$:

(i) $f(A) \geq 0$, and $f(A) = 0$ if and only if $A = 0$;

(ii) $f(\alpha A)=|\alpha| f(A)$;

(iii) $f(A+B) \leq f(A)+f(B)$.

Many of the popular matrix norms also satisfy a fourth condition called **sub-multplicativity**: for any conformably sized matrices A, B

$$f(AB) \leq f(A) f(B).$$

3.5.1 Frobenius Norm

Definition 3.5.1 (Frobenius norm)**.** The **Frobenius norm** $\|A\|_F$ is nothing but the standard Euclidean (ℓ_2) vector norm applied to the vector formed by all elements of $A \in \mathbb{R}^{m,n}$:

$$\|A\|_F=\sqrt{\operatorname{trace} A A^{\top}}=\sqrt{\sum_{i=1}^m \sum_{j=1}^n\left|a_{ij}\right|^2}.$$

The Frobenius norm also has an interpretation in terms of the eigenvalues of the symmetric matrix $A A^{\top}$:

$$\|A\|_F=\sqrt{\operatorname{trace} A A^{\top}}=\sqrt{\sum_{i=1}^m \lambda_i\left(A A^{\top}\right)}.$$

For any $x \in \mathbb{R}^n$, it holds that $\|Ax\|_2 \leq \|A\|_F\|x\|_2$. (consequence of the Cauchy-Schwartz inequality applied to $\left|a_i^{\top} x\right|$). The Frobenius norm is sub-multiplicative: for any $B \in \mathbb{R}^{n,p}$, it holds that

$$\|AB\|_F \leq \|A\|_F\|B\|_F.$$

3.5.2 Operator norms

The so-called operator norms give a characterization of the maximum input-output gain of the linear map $u \rightarrow y=Au$. Choosing to measure both inputs and outputs in terms of a given ℓ_p norm, with typical values $p=1,2, \infty$, leads to the definition

$$\|A\|_p \doteq \max_{u \neq 0} \frac{\|Au\|_p}{\|u\|_p}=\max_{\|u\|=1}\|Au\|_p$$

By definition, for every $u, \|Au\|_p \leq \|A\|_p\|u\|_p$. From this property follows that any operator norm is sub-multiplicative, that is, for any two conformably sized matrices A, B, it holds that

$$\|AB\|_p \leq \|A\|_p\|B\|_p$$

This fact is easily seen by considering the product AB as the series connection of the two operators B, A:

$$\|Bu\|_p \leq\|B\|_p\|u\|_p, \quad\|ABu\|_p \leq\|A\|_p\|Bu\|_p \leq\|A\|_p\|B\|_p\|u\|_p,$$

4

Matrices II

4.1 Orthogonalization

Definition 4.1.1 (Orthonormal basis)**.** A basis $(\boldsymbol{u}_i)_{i=1}^n$ is said to be **orthogonal** if $\boldsymbol{u}_i^\top \boldsymbol{u}_j = 0$ if $i \neq j$. If in addition, $\|\boldsymbol{u}_i\|_2 = 1$, we say that the basis is **orthonormal**.

Definition 4.1.2 (Orthogonalization)**.** **Orthogonalization** refers to a procedure that finds an orthonormal basis of the span of given vectors.

5

Symmetric Matrices

Definition 5.0.1 (Symmetric)**.** A square matrix $A \in \mathbb{R}^{n,n}$ is **symmetric** if $A = A^\top$.

5.1 The Spectral Theorem

Theorem 5.1.1 (Spectral Theorem)**.** Let $A \in \mathbb{R}^{n,n}$ be symmetric, let $\lambda_i \in \mathbb{R}, i = 1, \ldots, n$, be the eigenvalues of A (counting multiplicities). Then, there exist a set of orthonormal vectors $\boldsymbol{u}_i \in \mathbb{R}^n, i = 1, \ldots, n$, such that $A\boldsymbol{u}_i = \lambda_i \boldsymbol{u}_i$. Equivalently, there exist an orthogonal matrix $U = [\boldsymbol{u}_1 \cdots \boldsymbol{u}_n]$ (i.e., $UU^\top = U^\top U = I_n$) such that

$$A = U\Lambda U^\top = \sum_{i=1}^{n} \lambda_i \boldsymbol{u}_i \boldsymbol{u}_i^\top, \quad \Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_n).$$

6

Singular Value Decomposition

6.1 Dyads

Definition 6.1.1 (Dyad)**.** A matrix $A \in \mathbb{R}^{m,n}$ is called a **dyad** if it can be written as

$$A = \boldsymbol{p}^\top$$

for some vectors $\boldsymbol{p} \in \mathbb{R}^m, \in \mathbb{R}^n$. Element-wise we have

$$A_{ij} = p_i q_i, \qquad 1 \leq i \leq m, 1 \leq j \leq n.$$

6.1.1 Sums of dyads

The SVD theorem, discussed below, states that any matrix can be written as a sum of dyads:

$$A = \sum_{i=1}^{r} \boldsymbol{p}_i{}_i^\top$$

for vectors $\boldsymbol{p}_{i,i}$ that are mutually orthogonal. This allows us to intepret data matrices as sums of simpler matrices (dyads).

Theorem 6.1.2 (Singular Value Decomposition)**.** Any matrix $A \in \mathbb{R}^{m,n}$ can be factored as

$$A = U\tilde{\Sigma}V^\top$$

where $V \in \mathbb{R}^{n,n}$ and $U \in \mathbb{R}^{m,m}$ are orthogonal matrices and $\tilde{\Sigma} \in \mathbb{R}^{m,n}$ is a matrix having the first $r \doteq \operatorname{rank}(A)$ digaonal entries $(\sigma_1, \ldots, \sigma_r)$ positive and decreasing in magnitude, and all other entries zero:

$$\tilde{\Sigma} \doteq [$$

7

Linear Equations

7.1 Set of solutions of linear equations

Generic linear equations can be expressed in vecctor format as

$$A\boldsymbol{x} = \boldsymbol{y},$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is the vector of unknowns, $y \in \mathbb{R}^m$ is a given vector, and $A \in \mathbb{R}^{m,n}$ is a matrix containing the coefficients of the linear equations.

Key issues are: existence, uniqueness of solutions; characterization of the **solution set**:

$$S \doteq \{\boldsymbol{x} \in \mathbb{R}^n \mid A\boldsymbol{x} = \boldsymbol{y}\}.$$

Let $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \mathbb{R}^m$ denote the columns of A, i.e. $A = \begin{bmatrix}\boldsymbol{a}_1 & \cdots & \boldsymbol{a}_n\end{bmatrix}$. $A\boldsymbol{x}$ is simply a linear combination of the columns of A, with coefficients given by $\boldsymbol{x}$:

$$A\boldsymbol{x} = x_1\boldsymbol{a}_1 + \cdots + x_n\boldsymbol{a}_n.$$

Remark. $A\boldsymbol{x} \in \mathcal{R}(A)$. Thus, $S \neq \emptyset \iff \boldsymbol{y} \in \mathcal{R}(A)$.

The linear equation

$$A\boldsymbol{x} = \boldsymbol{y}, \qquad A \in \mathbb{R}^{m,n}$$

admits a solution if and only if $\operatorname{rank}\left(\begin{bmatrix}A & \boldsymbol{y}\end{bmatrix}\right) = \operatorname{rank}(A)$ and $\mathcal{N}(A) = \{0\}$.

When this existence condition is satisfied, the set of all solutions is the affine set

$$S = \{\boldsymbol{x} = \bar{\boldsymbol{x}} + N\boldsymbol{z}\},$$

where $\hat{\boldsymbol{x}}$ is any vector such that $A\hat{\boldsymbol{x}} = \boldsymbol{y}$ and $N \in \mathbb{R}^{n,n-r}$ is a matrix whose columns span the nullspace of A ($AN = 0$).

Overall, the system has a unique solution if $\operatorname{rank}(\begin{bmatrix}A & \boldsymbol{y}\end{bmatrix}) = \operatorname{rank}(A)$ and $\mathcal{N}(A) = \{0\}$.

7.2 Use case in optimization

Consider an optimization problem with linear equality constraints:

$$\min_{\boldsymbol{x}} f_0(\boldsymbol{x}) \quad : \quad A\boldsymbol{x} = \boldsymbol{b},$$

with $f_0 : \mathbb{R}^n \to \mathbb{R}$, and $A \in \mathbb{R}^{m,n}, \boldsymbol{b} \in \mathbb{R}^m$, and $\boldsymbol{x} \in \mathbb{R}^n$ the variable. We assume that the problem is feasible, i.e., the solution set of $A\boldsymbol{x} = \boldsymbol{y}$ is not empty.

Since the solution set is affine, any solution is of the form $\boldsymbol{x}_0 + N\boldsymbol{z}$, with $\boldsymbol{x}_0$ a particular solution and N a matrix whose columns span the nullspace of A.
We can formulate the above problem as an unconstrained one:

$$\min_{\boldsymbol{z}} f_0(\boldsymbol{x}_0 + Nz).$$

7.3 Solving via SVD

The linear equation $A\boldsymbol{x} = \boldsymbol{y}$ can be fully analyzed via SVD. If $A = U\tilde{\Sigma}V^\top$ is the SVD of A, then $A\boldsymbol{x} = \boldsymbol{y}$ is equivalent to

$$\tilde{\Sigma}\tilde{\boldsymbol{x}} = \tilde{\boldsymbol{y}},$$

where $\tilde{\boldsymbol{x}} \doteq V^\top\boldsymbol{x}, \tilde{\boldsymbol{y}} \doteq U^\top\boldsymbol{y}$.

Since $\tilde{\Sigma}$ is a diagonal matrix

$$\tilde{\Sigma} = \begin{bmatrix} \Sigma & \mathbf{0}_{r,n-r} \\ \mathbf{0}_{m-r,r} & \mathbf{0}_{m-r,n-r} \end{bmatrix}, \qquad \Sigma = \mathsf{diag}(\sigma_1, \ldots, \sigma_r) \succ 0,$$

the system above is very easy to solve.

8

Convexity

8.1 Affine Sets

Definition 8.1.1 (Affine Set). A set $\mathcal{C} \subseteq \mathbb{R}^n$ is **affine** if the line through any two distinct points in $\mathcal{C}$ lies in $\mathcal{C}$, i.e. if for any $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{C}$ and $\lambda \in \mathbb{R}$, we have $\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2 \in \mathcal{C}$.

Definition 8.1.2 (Affine Combination). A point of the form

$$\lambda_1 \boldsymbol{x}^{(1)} + \cdots + \lambda_m \boldsymbol{x}^{(m)}$$

where $\sum_{i=1}^m \lambda_i = 1$ is an **affine combination** of the points $\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}$.

Definition 8.1.3 (Affine Hull). The set of all affine combinations of points in some set $\mathcal{C} \subseteq \mathbb{R}^n$ is called the **affine hull** of $\mathcal{C}$, denoted **aff** $\mathcal{C}$:

$$\textbf{aff } \mathcal{C} = \left\{ \sum_{i=1}^m \lambda_i \boldsymbol{x}^{(i)} \middle| \boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)} \in \mathcal{C}, \sum_{i=1}^m \lambda_i = 1 \right\}.$$

8.2 Convex Sets

Definition 8.2.1 (Linear Hull). Given a set of points in $\mathbb{R}^n$:

$$\mathcal{P} = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}\},$$

the **linear hull** (subspace), or **span** generated by these points is the set of all possible *linear combinations* of the points:

$$\text{span}\mathcal{P} = \left\{ \sum_{i=1}^m \lambda_i \boldsymbol{x}^{(i)} \; \middle| \boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)} \in \mathcal{P}, \lambda_1, \ldots, \lambda_m \in \mathbb{R} \right\}.$$

Definition 8.2.2 (Convex Set)**.** A set $\mathcal{C}$ is **convex** if for any $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{C}$ and any $\lambda \in [0, 1]$:

$$\lambda \boldsymbol{x}_1 + (1 - \lambda)\boldsymbol{x}_2 \in \mathcal{C},$$

i.e., the **line segment** between any two points in $\mathcal{C}$ lies in $\mathcal{C}$.

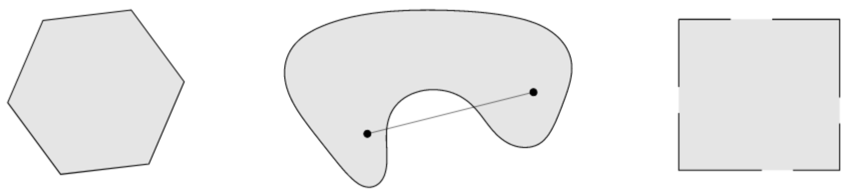

Figure 8.1: Convex and nonconvex sets. *Left.* Convex. *Middle.* Not convex as the line segment between the two points in the set is not contained. *Right.* Not convex as it contains some boundary points but not other.

Remark. Subspaces, halfspaces, and affine sets, such as lines and hyperplanes are convex, as they contain the entire line between any two distinct points in it, and thus also the line segment in $\mathbb{R}^2$. Hence, **every affine set is convex. However, not every convex set is affine.**

Definition 8.2.3 (Convex Combination)**.** A point of the form

$$\lambda_1 \boldsymbol{x}^{(1)} + \cdots + \lambda_m \boldsymbol{x}^{(m)}$$

where $\sum_{i=1}^m \lambda_i = 1$ and $\lambda_i \geq 0$ is a **convex combination** of the points $\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}$.

Definition 8.2.4 (Convex Hull)**.** The **convex hull** of a set $\mathcal{C}$, denoted **conv** $\mathcal{C}$, is the set of all convex combinations of points in $\mathcal{C}$:

$$\textbf{conv } \mathcal{C} = \left\{ \sum_{i=1}^m \lambda_i \boldsymbol{x}^{(i)} \mid \boldsymbol{x}^{(i)} \in \mathcal{C}, \lambda_i \geq 0, i = 1, \ldots, m, \sum_{i=1}^m \lambda_i = 1 \right\}.$$

Remark. conv $\mathcal{P}$ is always convex and it is the smallest convex set that contains $\mathcal{P}$, i.e., if $\mathcal{C}$ is any convex set that contains $\mathcal{P}$, then **conv** $\mathcal{P} \subseteq \mathcal{C}$.

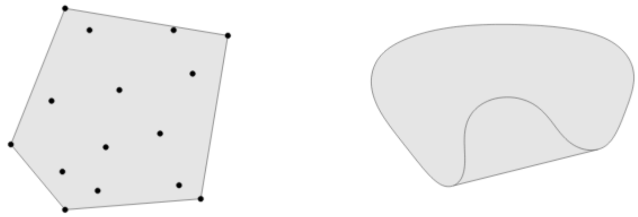

Figure 8.2: *Left.* The convex hull of a set of 15 points is the pentagon. *Right.* The convex hull of the kidney shped set is the shaded set.

8.3 Cones

Definition 8.3.1 (Cone)**.** A set $\mathcal{C}$ is a **cone** if $\lambda \boldsymbol{x} \in \mathcal{C}$ for every $x \in \mathcal{C}$ and $\lambda \geq 0$.

Definition 8.3.2 (Convex Cone)**.** A set $\mathcal{C}$ is a **convex cone** if it is convex and it is a cone. Equivalently, it means that for any $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{C}$ and $\lambda_1, \lambda_2 \geq 0$, we have

$$\lambda_1 \boldsymbol{x}_1 + \lambda_2 \boldsymbol{x}_2 \in \mathcal{C}.$$

Definition 8.3.3 (Conic Combination)**.** A point of the form

$$\lambda_1 \boldsymbol{x}^{(1)} + +\lambda_m \boldsymbol{x}^{(m)}$$

with $\lambda_1, \ldots, \lambda_m \geq 0$ is a **conic combination** of $\boldsymbol{x}^{(1)} \ldots, \boldsymbol{x}^{(m)}$.

Definition 8.3.4 (Conic Hull)**.** The **conic hull** of a set $\mathcal{C}$ is the set of all conic combinations of points in $\mathcal{C}$, i.e.,

$$\left\{ \sum_{i=1}^{m} \lambda_i \boldsymbol{x}_i \mid \boldsymbol{x}_i \in C, \lambda_i \geq 0, i = 1, \ldots, m \right\},$$

which is also the smallest convex cone that contains C.

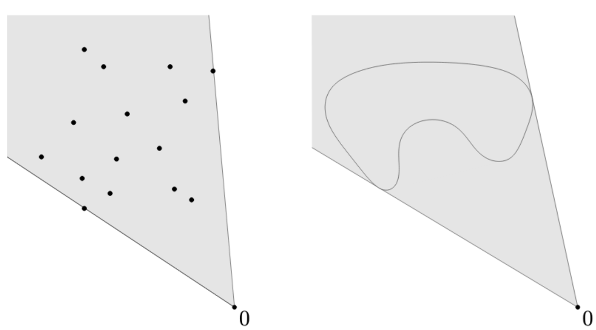

Figure 8.3: The conic hulls of the two sets of figure 1.2.

8.4 Hyperplanes and Halfspaces

Definition 8.4.1 (Hyperplane)**.** A **hyperplane** is a set of the form

$$\{\boldsymbol{x} \mid \boldsymbol{a}^\top \boldsymbol{x} = \boldsymbol{b}\},$$

where $\boldsymbol{a} \in \mathbb{R}^n, \boldsymbol{a} \neq \mathbf{0}$, and $\boldsymbol{b} \in \mathbb{R}$, i.e., the solutions set of a nontrivial linear equation among the components of $\boldsymbol{x}$ (and hence an affine set).

Geometric interpretation: The hyperplane is a set of points with a constant inner product to a given vector $\boldsymbol{a}$, which can also be viewed as a **normal vector**; the constant b determines the offset from the origin.

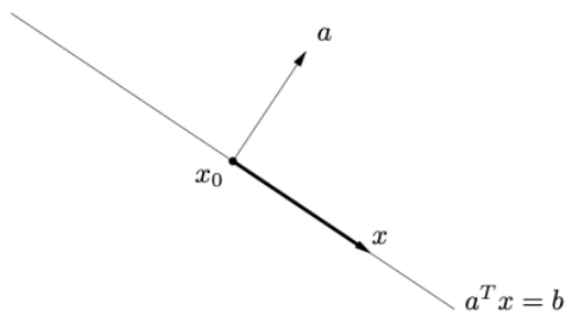

Figure 8.4: Hyperplane in $\mathbb{R}^2$, with normal vector $\boldsymbol{a}$. $\boldsymbol{x} - \boldsymbol{x}_0$ (arrow) is orthogonal to $\boldsymbol{a}$ for any $\boldsymbol{x}$ in the hyperplane.

A hyperplane divides $\mathbb{R}^n$ into two **halfspaces**, defined as follows:

Definition 8.4.2 (Halfspace)**.** A **halfspace** is a set of the form

$$\{\boldsymbol{x} \mid \boldsymbol{a}^\top \boldsymbol{x} \leq \boldsymbol{b}\},$$

where $\boldsymbol{a} \neq \mathbf{0}$, i.e., the solution set of a nontrivial linear inequality.

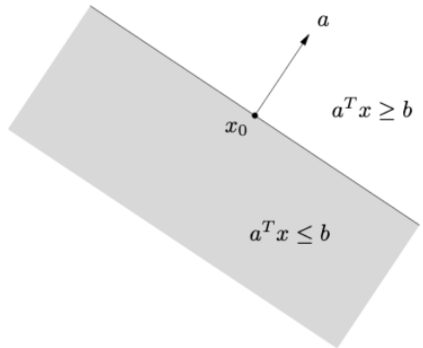

Figure 8.5: The halfspace determined by $\boldsymbol{a}^\top \boldsymbol{x} \leq \boldsymbol{b}$ (shaded) extends in the direction $-\boldsymbol{a}$.

Remark. Halfspaces are **convex**, but not affine.

Proof. Let $\boldsymbol{x}_1, \boldsymbol{x}_2$ be two points in a halfspace. Then for any $\lambda \in [0, 1]$, we have

$$\begin{aligned}\boldsymbol{a}^\top(\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2) &= \lambda a^\top \boldsymbol{x}_1 + (1-\lambda) a^\top \boldsymbol{x}_2 \\ &\leq \lambda b + (1-\lambda) b \\ &= b.\end{aligned}$$

Thus, halfspaces are convex. □

8.5 Operations Preserving Convexity

Proving convexity using the definition is very difficult and so we need some tools to help us with that. In particular, operations that preserve convexity will come in handy.

8.5.1 Intersection

Theorem 8.5.1. If $\mathcal{C}_1, \ldots, \mathcal{C}_m$ are convex sets, then their intersection

$$\mathcal{C} = \bigcap_{i=1}^{m} \mathcal{C}_i$$

is also a convex set.

Remark. This also holds for possibly infinite families of convex sets.
If $\mathcal{C}(\alpha), \alpha \in \mathcal{A} \subseteq \mathbb{R}^q$, is a family of convex sets, parameterized by α, then $\mathcal{C} = \bigcap_{\alpha \in \mathcal{A}} \mathcal{C}_\alpha$ is convex.

Proof. Let $\{\mathcal{C}_i\}_{i=1}^m$ be convex sets. For any $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \bigcap_{i=1}^m \mathcal{C}_i, \lambda \in [0, 1]$, $\boldsymbol{x}_1 \in \mathcal{C}_i$ and $\boldsymbol{x}_2 \in \mathcal{C}_i$ implies

$$\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2 \in \mathcal{C}_i$$

for $i = 1, 2, \ldots, m$ by convexity of $\mathcal{C}_i$. Hence,

$$\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2 \in \bigcap_{i=1}^{m} \mathcal{C}_i.$$

Thus, $\bigcap_{i=1}^m \mathcal{C}_i$ is convex. □

Example 8.5.2 (Second-order cone)**.** The second-order cone in $\mathbb{R}^{n+1}$:

$$\mathcal{K}_n = \{(\boldsymbol{x}, t) \mid \boldsymbol{x} \in \mathbb{R}^n, t \in \mathbb{R}, \|\boldsymbol{x}\|_2 \leq t\}$$

is convex since it is the intersection of half-spaces (which are convex):

$$\mathcal{K}_n = \bigcap_{\boldsymbol{u}:\|\boldsymbol{u}\|_2 \leq 1} \left\{(\boldsymbol{x}, t) \mid \boldsymbol{x} \in \mathbb{R}^n, t \in \mathbb{R}, \boldsymbol{u}^\top \boldsymbol{x} \leq t\right\}.$$

8.5.2 Affine Transformation

Theorem 8.5.3. If a map $f : \mathbb{R}^n \to \mathbb{R}^m$ is affine, and $\mathcal{C} \subset \mathbb{R}^n$ is convex, then the image set

$$f(\mathcal{C}) = \{f(\boldsymbol{x}) \mid \boldsymbol{x} \in \mathcal{C}\}$$

is convex.

Proof. Any affine map has a matrix representation

$$f(\boldsymbol{x}) = A\boldsymbol{x} + \boldsymbol{b}.$$

Then for any $\boldsymbol{y}_1, \boldsymbol{y}_2 \in f(\mathcal{C})$, there exists $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{C}$ such that $\boldsymbol{y}_1 A\boldsymbol{x}_1 + \boldsymbol{b}$, $\boldsymbol{y}_2 = A\boldsymbol{x}_2 + \boldsymbol{b}$. Hence, for $\lambda \in [0, 1]$, we have

$$\lambda\boldsymbol{y}_1 + (1 - \lambda)\boldsymbol{y}_2 = A(\lambda\boldsymbol{x}_1 + (1 - \lambda)\boldsymbol{x}_2) + \boldsymbol{b} = f(\boldsymbol{x}),$$

where $\boldsymbol{x} = \lambda\boldsymbol{x}_1 + (1 - \lambda)\boldsymbol{x}_2 \in \mathcal{C}$. $\square$

Remark. The projection of a convex set $\mathcal{C}$ onto a subspace is representable by means of a linear map, hence the projected set is convex.

8.6 Convex Functions

Definition 8.6.1 (Domain)**.** The **domain** of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$\mathbf{dom} f = \{\boldsymbol{x} \in \mathbb{R}^n \mid -\infty < f(\boldsymbol{x}) < \infty\}.$$

Definition 8.6.2 (Convex Function)**.** A function $f : \mathbb{R}^n \to \mathbb{R}$ is **convex** if $\mathbf{dom} f$ is a convex set, and for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbf{dom} f$ and all $\lambda \in [0, 1]$ it holds that

$$f(\lambda\boldsymbol{x} + (1 - \lambda)\boldsymbol{y}) \leq \lambda f(\boldsymbol{x}) + (1 - \lambda)f(\boldsymbol{y}).$$

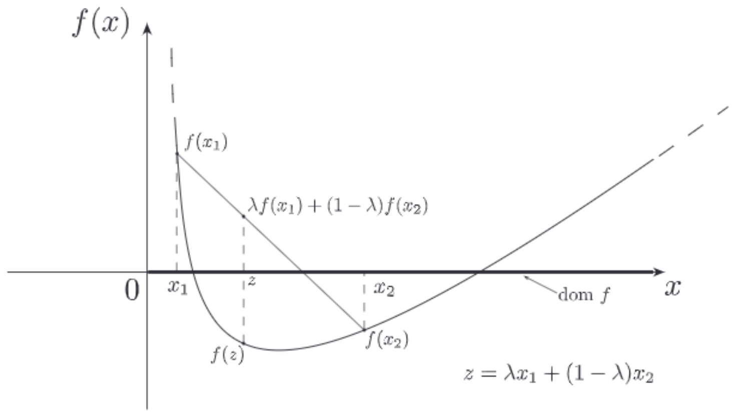

Figure 8.6: Convex function.

Definition 8.6.3 (Concave Function)**.** A function f is **concave** if $-f$ is convex.

8.6.1 Domain of a convex function

Convex functions must be $+\infty$ outside their domain in order for the convexity inequality to hold even if $\boldsymbol{x}$ or $\boldsymbol{y} \notin \mathbf{dom} f$. The function

$$f(\boldsymbol{x}) = \begin{cases} -\sum_{i=1}^{n} \log x_i & \text{if } \boldsymbol{x} > 0, \\ +\infty & \text{otherwise,} \end{cases}$$

is convex, but the function

$$f(\boldsymbol{x}) = \begin{cases} -\sum_{i=1}^{n} \log x_i & \text{if } \boldsymbol{x} > 0, \\ -\infty & \text{otherwise,} \end{cases}$$

is not.

Remark. To check the convexity of a extended-value function f, $\mathbf{dom} f$ must be convex and anything outside the domain must be $+\infty$.

8.6.2 Epigraph

One method of proving convexity of a function is to prove the convexity of its epigraph.

Definition 8.6.4 (Epigraph)**.** Recall that the **epigraph** of a function $f : \mathbb{R}^n \to (-\infty, +\infty)$ is the set

$$\mathbf{epi}\, f = \{(\boldsymbol{x}, t) \mid \boldsymbol{x} \in \mathbf{dom} f, t \in \mathbb{R}, f(\boldsymbol{x}) \leq t\}.$$

Remark. f is a convex function if and only if $\mathbf{epi}\, f$ is a convex set.

Example 8.6.5 (Log-sum-exp)**.** Consider the *log-sum-exp* function arising in logistic regression:

$$f(\boldsymbol{x}) = \log\left(\sum_{i=1}^{n} e^{x_i}\right).$$

The epigraph is the set of pairs $(\boldsymbol{x}, t)$ characterized by the inequality $t \geq f(\boldsymbol{x})$, which can be rewritten as

$$\textbf{epi}\, f = \left\{(\boldsymbol{x}, t) \in \mathbb{R}^n \times \mathbb{R} \mid \sum_{i=1}^{n} e^{x_i - t} \leq 1\right\},$$

which is convex, due to the convexity of the exponential function.

8.6.3 Sublevel Sets

Definition 8.6.6 (α-sublevel set)**.** For $\alpha \in \mathbb{R}$, the **α-sublevel set** of f is defined as

$$\mathcal{S}_\alpha \doteq \{\boldsymbol{x} \in \mathbb{R}^n \mid f(\boldsymbol{x}) \leq \alpha\}.$$

8.7 Operations Preserving Convexity (continued)

8.7.1 Non-negative Linear Combinations

Theorem 8.7.1. If $f_i : \mathbb{R}^n \to \mathbb{R}, i = 1, \ldots, m$, are convex functions, then the function

$$f(\boldsymbol{x}) = \sum_{i=1}^{m} \alpha_i f_i(\boldsymbol{x}), \qquad \alpha_i \geq 0, i = 1, \ldots, m$$

is also convex over $\bigcap_i \textbf{dom}\, f_i$.

Proof. This follows from the definition of convexity, since for any $\boldsymbol{x}, \boldsymbol{y} \in \textbf{dom}$ and $\lambda \in [0, 1]$,

$$\begin{aligned} f(\lambda \boldsymbol{x} + (1-\lambda)\boldsymbol{y}) &= \sum_{i=1}^{m} \alpha_i f_i(\lambda \boldsymbol{x} + (1-\lambda)\boldsymbol{y}) \\ &\leq \sum_{i=1}^{m} \alpha_i (\lambda f_i(\boldsymbol{x}) + (1-\lambda) f_i(\boldsymbol{y})) \\ &= \lambda f(\boldsymbol{x}) + (1-\lambda) f(\boldsymbol{y}). \end{aligned}$$

□

8.7.2 Affine Variable Transformation

Theorem 8.7.2. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex, and define

$$g(\boldsymbol{x}) = f(A\boldsymbol{x} + \boldsymbol{b}), \qquad A \in \mathbb{R}^{n,m}, \boldsymbol{b} \in \mathbb{R}^n.$$

Then g is convex over $\textbf{dom}\, g = \{\boldsymbol{x} \mid A\boldsymbol{x} + \boldsymbol{b} \in \textbf{dom}\, f\}$.

8.7.3 Pointwise Maximum

Theorem 8.7.3. If $(f_\alpha)_{\alpha\in\mathcal{A}}$ is a family of convex functions indexed by parameter α, and $\mathcal{A}$ is a set, then the pointwise max function

$$f(\boldsymbol{x}) = \max_{\alpha\in\mathcal{A}} f_\alpha(\boldsymbol{x})$$

is convex over the domain $\{\bigcap_{\alpha\in\mathcal{A}} \textbf{dom}\, f_\alpha\} \cap \{\boldsymbol{x} \mid f(\boldsymbol{x}) < \infty\}$.

Proof. The epigraph of f is the set of pairs $(\boldsymbol{x}, t)$ such that

$$\forall\alpha \in \mathcal{A}, \quad f_\alpha(\boldsymbol{x}) \le t.$$

Hence, the epigraph of f is the intersection of the epigraphs of all the functions involved, therefore f is convex. □

Remark. A convex function can be thought of as the maximum of possibly infinite number of linear functions (tangent lines of all points on the function).

Example 8.7.4 (Sum of k largest elements)**.** Consider the function $f : \mathbb{R}^n \to \mathbb{R}$ with values

$$f(\boldsymbol{x}) = \sum_{i=1}^{k} x_{[i]}$$

where $\boldsymbol{x}_{[i]}$ denotes the i-the largest element in $\boldsymbol{x}$. Then

$$f(\boldsymbol{x}) = \max_{\boldsymbol{u}} \boldsymbol{u}^\top \boldsymbol{x} \quad : \quad \boldsymbol{u} \in \{0,1\}^n, \mathbf{1}^\top \boldsymbol{u} = k.$$

For every $\boldsymbol{u}, \boldsymbol{x} \to \boldsymbol{u}^\top \boldsymbol{x}$ is linear, hence f is convex.

Example 8.7.5 (Largest eigenvalue of a symmetric matrix)**.** Consider the function $f : \mathbb{S}^n \to \mathbb{R}$, with values for a given $X = X^\top \in \mathbb{S}^n$ given by

$$f(X) = \lambda_{\max}(X),$$

where $\lambda_{\max}$ denotes the largest eigenvalue.
The function is the pointwise maximum of *linear* functions of X:

$$F(X) = \max_{\boldsymbol{u}:\|\boldsymbol{u}\|_2=1} \boldsymbol{u}^\top X \boldsymbol{u} = \lambda_{\max}(X).$$

Hence, f is convex.

8.7.4 First-order Conditions

Theorem 8.7.6. If f is differentiable (that is, $\textbf{dom} f$ is open and the gradient exists everywhere on the domain), then f is convex if and only if

$$\forall \boldsymbol{x}, \boldsymbol{y} \in \textbf{dom} f, \quad f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}).$$

Proof. Assume that f is convex. Then, the definition implies that for any $\lambda \in (0, 1]$

$$\frac{f(\boldsymbol{x} + \lambda(\boldsymbol{y} - \boldsymbol{x})) - f(\boldsymbol{x})}{\lambda} \leq f(\boldsymbol{y}) - f(\boldsymbol{x})$$

which, for $\lambda \to 0$ yields $\nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) \leq f(\boldsymbol{y}) - f(\boldsymbol{x})$ - Conversely, take any $\boldsymbol{x}, \boldsymbol{y} \in \textbf{dom} f$ and $\lambda \in [0, 1]$, and let $\boldsymbol{z} = \lambda \boldsymbol{x} + (1 - \lambda \boldsymbol{y})$

$$f(\boldsymbol{x}) \geq f(\boldsymbol{z}) + \nabla f(\boldsymbol{z})^\top (\boldsymbol{x} - \boldsymbol{z}), \quad f(\boldsymbol{y}) \geq f(\boldsymbol{z}) + \nabla f(\boldsymbol{z})^\top (\boldsymbol{y} - \boldsymbol{z})$$

Taking a convex combination of these inequalities, we get

$$\lambda f(\boldsymbol{x}) + (1 - \lambda) f(\boldsymbol{y}) \geq f(\boldsymbol{z}) + \nabla f(\boldsymbol{z})^\top 0 = f(\boldsymbol{z})$$

which concludes the proof. □

8.7.5 Second-order Conditions

Theorem 8.7.7. If f is twice differentiable, then f is convex if and only if its Hessian matrix $\nabla^2 f$ is positive semi-definite everywhere on the (open) domain of f, that is if and only if $\nabla^2 f \succeq 0$ for all $\boldsymbol{x} \in \textbf{dom} f$.

Example 8.7.8. A generic quadratic function

$$f(\boldsymbol{x}) = \frac{1}{2} \boldsymbol{x}^\top H \boldsymbol{x} + c^\top \boldsymbol{x} + d$$

has Hessian $\nabla^2 f(\boldsymbol{x}) = H$. Hence f is convex if and only if H is positive semidefinite.

9

Convex Optimization Problems

9.1 Trouble with nonlinear models

General optimization model:

$$
\begin{aligned}
p^* = &\min_{\boldsymbol{x}} f_0(\boldsymbol{x}) \\
\text{subject to:} \quad & f_i(\boldsymbol{x}) \le 0, \quad i = 1, \ldots, m,
\end{aligned}
$$

with f_i's arbitrary nonlinear functions.

- Algorithms may deliver very suboptimal solutions for unconstrained problems.
- Can fail for constrained problems-find no feasible point even though one exists.

9.2 Convex Problem

9.2.1 Standard Form

$$
\begin{aligned}
p^* = &\min_{\boldsymbol{x}} f_0(\boldsymbol{x}) \\
\text{subject to:} \quad & f_i(\boldsymbol{x}) \le 0, \quad i = 1, \ldots, m, \\
& A\boldsymbol{x} = \boldsymbol{b},
\end{aligned}
$$

where

- $f_0, \ldots, f_m$ are convex functions;
- Th equality constraints are affine, and represented via the matrix $A \in \mathbb{R}^{q \times n}$ and vector $b \in \mathbb{R}^q$.

Remark. There is an implicit constraint on $\boldsymbol{x}$, that it belongs to the domain of the problem, which is the set

$$D \doteq \bigcap_{i=0}^{m} \mathbf{dom} f_i.$$

Since f_i's are convex, the domain is convex.

9.2.2 Convexity of Conic Optimization Problems

The class of convex problems includes the conic optimization problems we've seen so far: LP, QP, QCQP, and SOCP, where

$$\text{LP} \subseteq \text{QP} \subseteq \text{QCQP} \subseteq \text{SOCP}.$$

Since the SOCP class includes all the others, it suffices to show that SOCPs are convex problems. Recall the SOCP model:

$$\begin{aligned} p^* = \min_{\boldsymbol{x} \in \mathbb{R}^n} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{s.t.:} \quad & \|A_i \boldsymbol{x} + \boldsymbol{b}_i\|_2 \leq \boldsymbol{c}_i^\top \boldsymbol{x} + \boldsymbol{d}_i, \quad i = 1, \ldots, m, \end{aligned}$$

with $A_i, \boldsymbol{b}_i, \boldsymbol{c}_i, \boldsymbol{d}_i$ matrices of appropriate size. Convexity of the problem stems from the fact that for every i,

$$\boldsymbol{x} \to \|A_i \boldsymbol{x} + \boldsymbol{b}_i\|_2 - (\boldsymbol{c}_i^\top \boldsymbol{x} + \boldsymbol{d}_i)$$

is a convex function.

9.2.3 Feasibility and Boundedness

- If the feasible set $\mathcal{X}$ is empty, the problem is **infeasible**: no solution that satisfies the constraints exists. Then it is customary to set $p^* = +\infty$. When $\mathcal{X}$ is nonempty, the problem is **feasible**.
- We usually do not know in advance if the feasible set $\mathcal{X}$ is empty or not; the task of determining this is referred to as a **feasibility problem**.

Remark. If the problem is feasible and $p^* = -\infty$, the problem is **unbounded below**. Note that it can also happen that the problem is feasible but stil no optimal solution exists, in which case we say that the optimal value p^* is not *attained* at any finite point.

9.2.4 Active vs. Inactive Constraints

- If $\boldsymbol{x}^* \in \mathcal{X}_{\text{opt}}$ is such that $f_i(\boldsymbol{x}^*) < 0$, we say that the i-th inequality constraint is **inactive** (or **slack**) at the optimal solution $\boldsymbol{x}^*$.
- Conversely, if $f_i(\boldsymbol{x}^*) = 0$, we say that the i-th inequality constraint is **active** at $\boldsymbol{x}^*$.

9.2.5 Optimal Set

- The **optimal set** is defined as the set of feasible points for which the objective function attains the optimal value:

  $$\mathcal{X}_{\text{opt}} = \{x \in \mathcal{X} : f_0(x) = p^*\}$$

  Equivalently,

  $$\mathcal{X}_{\text{opt}} = \arg\min_{x \in \mathcal{X}} f_0(x)$$

- For convex problems, the optimal set is a convex set.

9.2.6 Local and Global Optima

Theorem 9.2.1. Consider the optimization problem: $\min_{\boldsymbol{x} \in \mathcal{X}} f_0(\boldsymbol{x})$. If f_0 is a convex function and $\mathcal{X}$ is a convex set, then any locally optimal solution is also globally optimal. Moreover, the set $\mathcal{X}_{\text{opt}}$ of optimal points is convex.

Proof. Proof. Let $\boldsymbol{x}^* \in \mathcal{X}$ be a local minimizer of f_0, let $p^* = f_0(\boldsymbol{x}^*)$, and consider any point $\boldsymbol{y} \in \mathcal{X}$. We need to prove that $f_0(\boldsymbol{y}) \geq f_0(\boldsymbol{x}^*) = p^*$. By convexity of f_0 and $\mathcal{X}$ we have that, for $\theta \in [0,1], \boldsymbol{x}_\theta = \theta \boldsymbol{y} + (1-\theta)\boldsymbol{x}^* \in \mathcal{X}$, and

$$f_0(\boldsymbol{x}_\theta) \leq \theta f_0(\boldsymbol{y}) + (1-\theta) f_0(\boldsymbol{x}^*)$$

Subtracting $f_0(\boldsymbol{x}^*)$ from both sides of this equation, we obtain

$$f_0(\boldsymbol{x}_\theta) - f_0(\boldsymbol{x}^*) \leq \theta\left(f_0(\boldsymbol{y}) - f_0(\boldsymbol{x}^*)\right)$$

Since $\boldsymbol{x}^*$ is a local minimizer, the left-hand side in this inequality is nonnegative for all small enough values of $\theta > 0$. We thus conclude that the right hand side is also nonnegative, i.e., $f_0(\boldsymbol{y}) \geq f_0(\boldsymbol{x}^*)$, as claimed. Also, the optimal set is convex, since it can be expressed as the p^*-sublevel set of a convex function:

$$\mathcal{X}_{\text{opt}} = \{\boldsymbol{x} \in \mathcal{X} : f_0(\boldsymbol{x}) \leq p^*\}$$

□

Example 9.2.2 (Feasible problem with empty feasible set)**.**
The problem

$$\begin{aligned} p^* = \min_{x \in \mathbb{R}} \quad & e^{-x} \\ \text{s.t.:} \quad & x \geq 0. \end{aligned}$$

is feasible with $p^* = 0$. However, the optimal set is empty since p^* is not attained at any finite point (it is only attained in the limit, as $x \to \infty$).

9.3 Problem Transformations

- An optimization problem can be transformed, or reformulated, into an equivalent one by means of several useful *tricks*, such as:
  - monotone transformation of the objective (e.g., scaling, logarithm, squaring) and constraint functions;
  - change of variables;
  - addition of slack variables;
  - epigraphic reformulation;
  - replacement of equality constraints with inequality ones;
  - elimination of inactive constraints;
  - discovering hidden convexity.

9.3.1 Monotone Objective Transformation

Theorem 9.3.1. Consider an optimization problem in standard form. Let $\varphi : \mathbb{R} \to \mathbb{R}$ be a continuous and strictly increasing function over $\mathcal{X}$, and consider the transformed problem

$$\begin{aligned} g^* = \min_{\boldsymbol{x} \in \mathbb{R}^n} \quad & \varphi(f_0(\boldsymbol{x})) \\ \text{s.t:} \quad & f_i(\boldsymbol{x}) \leq 0, \quad i = 1, \ldots, m, \\ & A\boldsymbol{x} = \boldsymbol{b}. \end{aligned}$$

The original and transformed problems have the same set of optimal solutions.

Proof. Suppose $\boldsymbol{x}^*$ is optimal i.e., $f_0(\boldsymbol{x}^*) = p^*$. Then, $\boldsymbol{x}^*$ is feasible, thus it holds that $\varphi(f_0(\boldsymbol{x}^*)) = \varphi(p^*) \geq g^*$. Assume next that $\tilde{\boldsymbol{x}}^*$ is optimal, i.e., $\varphi(f_0(\tilde{\boldsymbol{x}}^*)) = g^*$. Then, $\tilde{\boldsymbol{x}}^*$ is feasible, thus it holds that $f_0(\tilde{\boldsymbol{x}}^*) \geq p^*$ Now, since φ is continuous and strictly increasing over $\mathcal{X}$, it has a well-defined inverse φ^{-1}, thus we may write $\varphi(f_0(\tilde{x}^*)) = g^* \iff \varphi^{-1}(g^*) = f_0(\tilde{\boldsymbol{x}}^*)$, which yields

$$\varphi^{-1}(g^*) \geq p^*$$

Since φ is strictly increasing and $\varphi\left(\varphi^{-1}(g^*)\right) = g^*$, the latter relation also implies that $g^* \geq \varphi(p^*)$, which implies that it must be $\varphi(p^*) = g^*$. This means that for any optimal solution $\boldsymbol{x}^*$ it holds that

$$\varphi(f_0(\boldsymbol{x}^*)) = g^*$$

which implies that $\boldsymbol{x}^*$ is also optimal for the problem. Vice-versa, for any optimal solution $\tilde{\boldsymbol{x}}^*$, it holds that

$$f_0(\tilde{\boldsymbol{x}}^*) = \varphi^{-1}(g^*) = p^*$$

which implies that $\tilde{\boldsymbol{x}}^*$ is also optimal. ☐

Example 9.3.2 (Least-squares)**.** The least-squares problem, with objective $f_0(\boldsymbol{x}) = \|A\boldsymbol{x} - \boldsymbol{y}\|_2$, where $A \in \mathbb{R}^{m \times n}$, $\boldsymbol{b} \in \mathbb{R}^m$ are given. We apply the result with the function $\boldsymbol{z} \geq 0 \to \varphi(\boldsymbol{z}) = \boldsymbol{z}^2$, which is increasing on $\mathbb{R}_+$.

Example 9.3.3 (Logistic Regression)**.** A random variable $\dot{y} \in \{-1, 1\}$ has a distribution modelled as

$$p = \mathbb{P}(y = 1) = \frac{1}{1 + \exp\left(-(w^\top \boldsymbol{x} + b)\right)} = 1 - \mathbb{P}(y = -1)$$

where $\boldsymbol{w} \in \mathbb{R}^n, b \in \mathbb{R}$ are parameters, and $\boldsymbol{x} \in \mathbb{R}^n$ contains explanatory variables (features). The estimation problem is to estimate $\boldsymbol{w}, b$ from given observations $(x_i, y_i), i = 1, \ldots, m$.
In the **maximum likelihood** approach, we maximize the **likelihood function**

$$\mathcal{L}(\boldsymbol{w}, b) \doteq \prod_{i=1}^{m} \left(\frac{1}{1 + \exp\left(-y_i(w^\top \boldsymbol{x} + b)\right)} \right)$$

$\mathcal{L}$ is not concave in $(\boldsymbol{w}, b)$, but $\log \mathcal{L}$ is, since the log-sum-exp function is convex:

$$\log \mathcal{L}(\boldsymbol{w}, b) = -\sum_{i=1}^{m} \log\left(1 + \exp\left(-y_i\left(\boldsymbol{w}^\top x_i + b\right)\right)\right).$$

9.3.2 Addition of Slack Variables

Equivalent problem formulations are also obtained by introducing new **slack variables** into the problem. Here is a typical case that arises when a constraint or the objective involves the sum of terms, as in the following problem

$$\begin{aligned} p^* = \min_{\boldsymbol{x}} \quad & f_0(\boldsymbol{x}) + \sum_{i=1}^{r} \varphi_i(x) \\ \text{s.t.:} \quad & \boldsymbol{x} \in \mathcal{X}. \end{aligned}$$

Introducing slack variables $t_i, i = 1, \ldots, p$, we reformulate this problem as

$$\begin{aligned} g^* = \min_{\boldsymbol{x}, t} \quad & f_0(\boldsymbol{x}) + \sum_{i=1}^{r} t_i \\ \text{s.t.:} \quad & \boldsymbol{x} \in \mathcal{X} \\ & \varphi_i(x) \leq t_i \quad i = 1, \ldots, r, \end{aligned}$$

where this new problem has the original variable $\boldsymbol{x}$, plus the vector of slack variables æ $(t_1, \ldots, t_r)$.

Example 9.3.4 (LASSO)**.** The LASSO problem

$$p^* = \min_{\boldsymbol{x}} \|A\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \|\boldsymbol{x}\|_1$$

is equivalent to the QP

$$\begin{aligned} p^* = \min_{\boldsymbol{x}, t} \; & \|A\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \sum_{i=1}^{n} t_i \\ \text{s.t.:} \quad & -t_i \leq x_i \leq t_i, \quad i = 1, \ldots, n. \end{aligned}$$

At optimum, we have $t_i^* = |x_i^*|, i = 1, \ldots, n$.

9.3.3 Epigraphic Reformulation

A common use of the slack variable trick described above consists in transforming a convex optimization problem with generic convex objective f_0, into an equivalent convex problem having **linear** objective. Introducing a new slack variable $t \in \mathbb{R}$, we can reformulate the problem as

$$\begin{aligned} t^* = \min_{\boldsymbol{x} \in \mathbb{R}^n, t \in \mathbb{R}} \quad & t \\ \text{s.t.:} \quad & f_i(\boldsymbol{x}) \leq 0, \quad i = 1, \ldots, m \\ & h_i(x) = 0, \quad i = 1, \ldots, q \\ & f_0(x) \leq t. \end{aligned}$$

This is referred to as the **epigraphic reformulation** of the original problem.

Remark. Any convex optimization problem can thus be reformulated in the form of an equivalent convex problem with linear objective.

9.3.4 Replacement of Equality Constraints with Inequality Ones

Consider a (not necessarily convex) problem of the form

$$\begin{aligned} p^* = \min_{\boldsymbol{x} \in \mathcal{X}} \quad & f_0(\boldsymbol{x}) \\ \text{s.t.:} \quad & b(\boldsymbol{x}) = u, \end{aligned}$$

where u is a given scalar, together with the related problem in which the equality constraint is substituted by an inequality one:

$$\begin{aligned} g^* = \min_{\boldsymbol{x} \in \mathcal{X}} \quad & f_0(\boldsymbol{x}) \\ \text{s.t.:} \quad & b(\boldsymbol{x}) \leq u, \end{aligned}$$

- Clearly, since the feasible set of the first problem is included in the feasible set of the second problem, it always holds that $g^* \leq p^*$.
- It actually holds that $g^* = p^*$, under the following conditions:
  1. f_0 is nonincreasing over $\mathcal{X}$ (i.e., $f_0(\boldsymbol{x}) \leq f_0(y) \iff \boldsymbol{x} \geq \boldsymbol{y}$ elementwise)
  2. b is nondecreasing over $\mathcal{X}$, and
  3. the optimal value p^* is attained at some optimal point $\boldsymbol{x}^*$, and the optimal value g^* is attained at some optimal point $\tilde{\boldsymbol{x}}^*$.
- In certain cases, we can substitute an equality constraint of the form $b(\boldsymbol{x}) = \boldsymbol{u}$ with an inequality constraint $b(\boldsymbol{x}) \leq \boldsymbol{u}$.
- This can be useful, in some cases, for gaining convexity. Indeed, if $b(\boldsymbol{x})$ is a convex function, then the set described by the equality constraint $\{\boldsymbol{x} : b(\boldsymbol{x}) = \boldsymbol{u}\}$ is non-convex in general (unless b is affine); contrary, the set described by the inequality constraint $\{\boldsymbol{x} : b(\boldsymbol{x}) \leq \boldsymbol{u}\}$ is the sublevel set of a convex function, which is convex.

9.3.5 Hidden Convexity

- Sometimes a problem as given is not convex, but we can transform it into an equivalent problem that is.
- An approach is to *relax* the problem into a convex one, and then prove that the relaxation is *exact*.

Remark. None of these approaches is full-proof and a guaranteed path to finding a convex expression of a given problem.

10

Weak Duality

10.1 Lagrangian

Definition 10.1.1 (Lagrangian)**.** The **Lagrangian** is a function with values for $\boldsymbol{x} \in \mathbb{R}^n, \lambda \in \mathbb{R}^m$ and $\nu \in \mathbb{R}^q$:

$$\mathcal{L}(\boldsymbol{x}, \lambda, \nu) = f_0(\boldsymbol{x}) + \sum_{i=1}^{m} \lambda_i f_i(\boldsymbol{x}) + \sum_{i=1}^{q} \nu_i h_i(\boldsymbol{x})$$

Vectors λ and ν are referred to as **Lagrange multipliers,** or dual variables.

Consider an optimization problem in standard form:

$$\begin{aligned} p^* = \min_{\boldsymbol{x} \in \mathbb{R}^n} \quad & f_0(\boldsymbol{x}) \\ \text{s.t.:} \quad & f_i(\boldsymbol{x}) \leq 0, \quad i = 1, \dots, m, \\ & h_i(\boldsymbol{x}) = 0, \quad i = 1, \dots, q, \end{aligned}$$

where f_i, h_i are not convex or concave.
With Lagrangian, we can express a problem in **min-max** form:

$$p^* = \min_{\boldsymbol{x}} \max_{\lambda \geq 0, \nu} \mathcal{L}(\boldsymbol{x}, \lambda, \nu).$$

This follows from the fact that for any $\boldsymbol{x}$:

$$\max_{\lambda \geq 0, \nu} \mathcal{L}(\boldsymbol{x}, \lambda, \nu) = \begin{cases} f_0(\boldsymbol{x}) & \text{if } \boldsymbol{x} \text{ is feasible,} \\ +\infty & \text{otherwise.} \end{cases}$$

Remark. When $\boldsymbol{x}$ is feasible, we have $f_i(\boldsymbol{x}) \leq 0$, and thus we set $\lambda_i = 0$ to maximize. Similarly, we also have $h_i(\boldsymbol{x}) = 0$, which means that the last term of the Lagrangian disappears. If $\boldsymbol{x}$ is not feasible, it means that we either have $f_i(\boldsymbol{x}) > 0$ for some f_i or $h_i(\boldsymbol{x}) \neq 0$ for some h_i. In the first case, λ will be set to $+\infty$ to maximize the value. Likewise, ν_i will be set to $\pm\infty$ depending on the sign of $h_i(\boldsymbol{x})$. Hence, we have $+\infty$ when $\boldsymbol{x}$ is infeasible.

10.2 Minimax Inequality

Theorem 10.2.1 (Minimax Inequality)**.** For any sets X, Y and any function $F : X \times Y \to \mathbb{R}$:

$$\min_{x \in X} \max_{y \in Y} F(x, y) \geq \max_{y \in Y} \min_{x \in X} F(x, y).$$

Proof. For any $(x_0, y_0) \in X \times Y$:

$$h(y_0) \doteq \min_{x \in X} F(x, y_0) \leq F(x_0, y_0) \leq \max_{y \in Y} F(x_0, y) \doteq g(x_0).$$

Hence, $h(y_0) \leq g(x_0)$ for any (x_0, y_0), which implies that

$$\max_{y \in Y} \min_{x \in X} F(x, y) = \max_{y_0 \in Y} h(y_0) \leq \min_{x_0 \in X} g(x_0) = \min_{x \in X} \max_{y \in Y} F(x, y).$$

□

10.3 Weak Duality

Applying the minimax inequality to the Lagrangian, we have the **weak duality**.

Theorem 10.3.1 (Weak Duality)**.**

$$\min_{\boldsymbol{x}} \max_{\lambda \geq 0, \nu} \mathcal{L}(\boldsymbol{x}, \lambda, \nu) \geq \max_{\lambda \geq 0, \nu} \min_{\boldsymbol{x}} \mathcal{L}(\boldsymbol{x}, \lambda, \nu).$$

The problem on the right is called the **dual problem**; it involves maximizing over $\lambda \geq 0, \nu$ the **dual function**:

$$g(\lambda, \nu) \doteq \min_{\boldsymbol{x}} \mathcal{L}(\boldsymbol{x}, \lambda, \nu).$$

Remark. Since g is the pointwise minimum of affine (hence, concave) functions, g is concave. Hence, the dual problem, a concave maximization problem over a convex set is convex.

10.4

11

Strong Duality

11.1 Strong Duality for Convex Problems

11.1.1 Slater's Condition for Strong Duality

Definition 11.1.1 (Strictly feasible)**.** A problem is **strictly feasible** if

Theorem 11.1.2 (Slater's Conditions for Convex Programs)**.** If the convex problem is strictly feasible, then strong duality holds: $p^* = d^*$.

11.1.2 Geometry

11.1.3 Recovering a Primal Solution from the Dual

11.1.4 Duality in Unconstrained Problems

11.1.5 Sion's Minimax Theorem

Theorem 11.1.3 (Minimax Theorem)**.** Let $X \subseteq \mathbb{R}^n$ be convex and let $Y \subseteq \mathbb{R}^m$ be a compact set. Let $F : X \times Y \to \mathbb{R}$ be a function such that for every $y \in Y$, $F(\cdot, y)$ is convex and continuous over X, and for every $x \in X$, $F(x, \cdot)$ is concave and continuous over Y. Then

$$\max_{y \in Y} \min_{x \in X} F(x, y) = \min_{x \in X} \max_{y \in Y} F(x, y).$$