Advanced Algorithms / Fall 2026

Linear Programming and Optimization

LP Foundations

Modeling, geometry, and integral polytopes

Lecture 1: 100 minutes

These notes cover only the linear-programming strand taught by Shi Li. This module comprises nine 100-minute lectures, numbered consistently with the slides. Parts I and II are reading divisions within this lecture; optional appendices are for further study.

Instructor: Shi Li

Nanjing University

Prepared from the instructor's course slides, with AI-assisted expansion of explanations, proofs, and independent-study exercises. The English and Chinese editions have matching mathematical content.

Learning goals and reading route

This unit introduces the language used throughout the optimization strand. Part I turns discrete choices into linear constraints. Part II explains why the geometry of the feasible region determines whether an LP gives an exact combinatorial algorithm or only a relaxation.

1. Translate a weighted set-selection problem into an integer program and a linear relaxation.
2. Distinguish a feasible solution, an optimum solution, and an optimum value.
3. Use active constraints to characterize a vertex, including lower-dimensional examples.
4. Explain how integrality turns a relaxation into an exact formulation.

Read Part I before Part II. The algorithm survey supplies context; this unit does not teach simplex pivot rules or interior-point analysis. Exercises belong to independent study. Throughout, input coefficients are rational when an algorithmic running-time claim is made.

Contents

1 Part I: Modeling and formulations **1**

1.1 Discrete objects as vectors . . . 1
1.2 The LP convention used in these notes . . . 1
1.3 A complete two-variable example . . . 2
1.4 Feasible regions and the role of LP . . . 2

2 Part II: Polyhedral geometry and integrality **3**

2.1 Convexity, dimension, and faces . . . 3
2.2 Vertices and active constraints . . . 3
2.3 Integral polytopes and exact formulations . . . 4
2.4 Description size and algorithms . . . 4

3 Exercises for independent study **5**

1 Part I: Modeling and formulations

1.1 Discrete objects as vectors

A typical combinatorial optimization problem has a ground set $[n] = \{1, \dots, n\}$, a family of feasible sets $\mathcal{S} \subseteq 2^{[n]}$, and weights w_i. We seek a feasible set maximizing or minimizing $\sum_{i \in S} w_i$. The family need not be listed explicitly: a graph, for example, specifies all its matchings implicitly.

Definition 1.1 (Indicator vector)**.** For $S \subseteq [n]$, its indicator vector $\chi^S \in \{0,1\}^n$ has coordinate $\chi^S_i = 1$ exactly when $i \in S$.

The objective becomes the linear function $w^\top x$. The modeling task is to find constraints whose 0–1 solutions are exactly the vectors of feasible objects. Shortest paths, spanning trees, matchings, independent sets, and knapsack solutions all admit this viewpoint, although the usefulness of their relaxations is very different.

Example 1.2 (A first matching formulation)**.** For a loopless graph $G = (V, E)$, let x_e indicate whether edge e is chosen. No two chosen edges may share a vertex, so a maximum-weight matching is described by

$$\max \sum_{e \in E} w_e x_e \quad \text{s.t.} \quad \sum_{e \in \delta(v)} x_e \le 1 \ (v \in V), \qquad x_e \in \{0,1\}.$$

Replacing $x_e \in \{0,1\}$ by $x_e \ge 0$ gives an LP relaxation; the degree constraints already imply $x_e \le 1$. Whether this LP has an integral optimum is a structural question about the graph. Unit 2 answers it for bipartite graphs and explains the obstruction in general graphs.

What a relaxation guarantees

A relaxation enlarges the feasible set. Therefore its optimum is a lower bound for a minimization problem and an upper bound for a maximization problem. A fractional feasible vector is generally not a valid discrete answer. One must either prove integrality or design a rounding algorithm and analyze its loss.

1.2 The LP convention used in these notes

A linear program optimizes a linear objective over finitely many linear equalities and inequalities. We use the following minimization convention:

$$\min c^\top x \quad \text{s.t.} \quad Ax \ge b, \quad x \ge 0, \tag{1.1}$$

where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, and $c \in \mathbb{R}^n$. Vector inequalities are coordinatewise: the ith constraint says $\sum_j a_{ij} x_j \ge b_i$. Here n is the number of variables and m is the number of displayed rows, before counting nonnegativity. Other books use other meanings of "standard form"; the actual signs matter more than the name.

Original feature	Conversion
$a^\top x \le \beta$	$-a^\top x \ge -\beta$
$a^\top x = \beta$	$a^\top x \ge \beta$ and $-a^\top x \ge -\beta$
$z \in \mathbb{R}$	$z = z^+ - z^-$, $z^+, z^- \ge 0$
$\max c^\top x$	$-\min(-c^\top x)$ at the level of optimum values

These transformations preserve the information needed to recover an original feasible solution and objective value. A free variable requires two nonnegative variables; simply declaring it nonnegative would change the problem. Likewise, changing a maximization to a minimization reverses the sign of the reported optimum.

1.3 A complete two-variable example

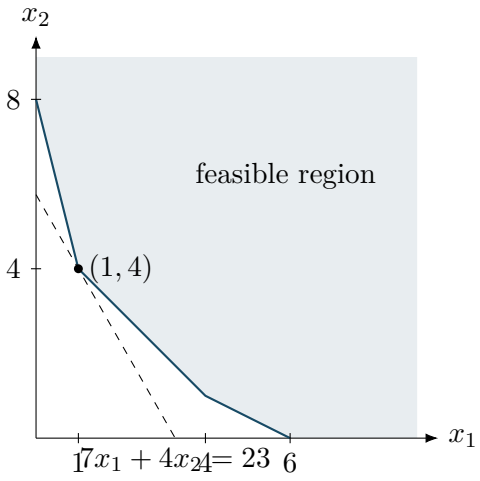
Consider the running example in the slides:

$$\begin{aligned} \min \quad & 7x_1 + 4x_2 \\ \text{s.t.} \quad & x_1 + x_2 \geq 5, \\ & x_1 + 2x_2 \geq 6, \\ & 4x_1 + x_2 \geq 8, \\ & x_1, x_2 \geq 0. \end{aligned} \tag{1.2}$$

The point $(1, 4)$ is feasible and has value 23. Feasibility alone does not show optimality. To prove a matching lower bound, multiply the first constraint by 3 and add the third:

$$7x_1 + 4x_2 = 3(x_1 + x_2) + (4x_1 + x_2) \geq 3 \cdot 5 + 8 = 23.$$

Every feasible solution costs at least 23, and $(1, 4)$ attains this bound. It is therefore optimal. The multipliers $(3, 0, 1)$ are an early example of a dual certificate, developed systematically in Unit 3. Notice that the second constraint is slack at this optimum.



The shaded region continues upward and to the right. It is unbounded, but the minimum is finite. This distinction is fundamental: unboundedness of the feasible region does not imply unboundedness of the objective in the optimization direction.

1.4 Feasible regions and the role of LP

Definition 1.3. A polyhedron is an intersection of finitely many closed half-spaces. A polytope is a bounded polyhedron. The feasible region of an LP is a polyhedron; it is a polytope only when it is bounded.

For a minimization LP, three possibilities must be distinguished: infeasibility, a finite attained optimum, and an objective unbounded below. The convention $\inf \varnothing = +\infty$ is useful, while an objective unbounded below has infimum $-\infty$. A feasible LP with finite infimum attains it; this is a property of polyhedra and does not hold for arbitrary nonclosed feasible sets. LPs arise in transportation, routing, scheduling, resource allocation, and facility location. They can be exact algorithms, lower-bound subroutines inside integer-programming methods, or starting points for approximation algorithms. Their value is not just that a solver returns numbers: the formulation can reveal structure and provide certificates.

2 Part II: Polyhedral geometry and integrality

2.1 Convexity, dimension, and faces

Definition 2.1. A convex combination of points $v^1, \dots, v^k$ is a vector $\sum_i \lambda_i v^i$ with $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$. Their convex hull, denoted $\operatorname{conv}\{v^1, \dots, v^k\}$, is the set of all such combinations.

Every polyhedron is convex: if two vectors satisfy $a^\top x \leq b$, so does every point on the segment joining them. Its affine dimension is the dimension of its smallest containing affine subspace. If the maximal independent system of equations true throughout a nonempty polyhedron in $\mathbb{R}^n$ has rank r, its dimension is $n - r$. An equation is counted by the rank of its normal vectors, not by how often it is written.

If $a^\top x \leq \beta$ is valid on P and equality is attained, the set $F = \{x \in P : a^\top x = \beta\}$ is a face. A proper face of dimension $\dim(P) - 1$ is a facet; a zero-dimensional face is a vertex. For a triangle embedded in $\mathbb{R}^3$, edges are facets because the triangle has dimension 2, not 3.

2.2 Vertices and active constraints

Definition 2.2. A point $x \in P$ is an extreme point, or vertex, if it cannot be written as $x = \lambda u + (1 - \lambda)v$ for distinct $u, v \in P$ and $0 < \lambda < 1$. A constraint is active, or tight, at x when it holds with equality.

Theorem 2.3 (Vertex criterion)**.** *Let $P = \{x \in \mathbb{R}^n : Ax \leq b\}$, with equalities represented by pairs of inequalities. A feasible x is a vertex if and only if the normals of its active constraints span $\mathbb{R}^n$. Equivalently, some n independent active constraints determine x uniquely.*

Proof. If the active normals have rank less than n, choose a nonzero vector d orthogonal to all of them. Every inactive constraint has positive slack. Because there are finitely many constraints, a sufficiently small $\varepsilon > 0$ makes both $x + \varepsilon d$ and $x - \varepsilon d$ feasible. Their midpoint is x, so x is not a vertex.

Conversely, suppose n independent tight rows form $Bx = b'$, and $x = \lambda u + (1-\lambda)v$ with u, v feasible. A tight inequality at a strict convex combination must be tight at both endpoints; otherwise the combination would have slack. Hence $Bu = Bv = b'$. Nonsingularity gives $u = v = x$. □

The criterion also covers degenerate vertices: more than n constraints may be tight, but only n independent ones are required. In a lower-dimensional polytope, equations defining the affine hull contribute to the ambient rank.

Theorem 2.4 (A polytope is the convex hull of its vertices)**.** *A nonempty polytope has finitely many vertices and equals their convex hull. Consequently, every linear objective has an optimum at a vertex.*

Proof. There are finitely many choices of n independent active rows, and each choice determines at most one vertex. To express an arbitrary nonvertex x, choose a direction orthogonal to its tight rows as above. Boundedness makes the feasible interval along this line closed and bounded. Its two endpoints u, v express x as a convex combination, and each endpoint has an additional tight row independent of the old tight rows: a row already in their span cannot bound motion in the chosen direction. Repeat at nonvertex endpoints. The active rank increases along every branch, so the process ends after at most n rank increases, giving a convex combination of vertices. Finally, a linear function at a convex combination is the same convex combination of its values; at least one vertex does at least as well as the combination. □

Remark 2.5. Boundedness is essential for this convex-hull statement. An unbounded polyhedron can contain rays, and an affine line has no vertex at all. Also, an LP solver may return a nonvertex optimum when the entire optimal face is nontrivial. Integrality of vertices does not say every optimum returned by every solver is integral.

2.3 Integral polytopes and exact formulations

For a feasible family $\mathcal{S}$, define its combinatorial polytope

$$P_{\mathcal{S}} = \operatorname{conv}\{\chi^S : S \in \mathcal{S}\}.$$

Optimizing any linear function over this polytope has exactly the same value as optimizing over the discrete objects. The difficulty is describing the convex hull efficiently. Matching, spanning-tree, and travelling-salesperson polytopes are examples of this construction, with very different descriptions and algorithmic consequences.

Definition 2.6. A polytope is integral if all its vertices belong to $\mathbb{Z}^n$. For a polytope contained in $[0,1]^n$, this means all vertices are 0–1 vectors.

Theorem 2.7 (Exactness from integrality)**.** *Let $Q \subseteq [0,1]^n$ be a polytope satisfying $Q \cap \{0,1\}^n = \{\chi^S : S \in \mathcal{S}\}$. If Q is integral, then $Q = P_{\mathcal{S}}$.*

Proof. Every feasible indicator vector lies in Q, so convexity gives $P_{\mathcal{S}} \subseteq Q$. Conversely, every vertex of Q is a 0–1 point and therefore the indicator of a feasible set. Since Q is the convex hull of its vertices, $Q \subseteq P_{\mathcal{S}}$. □

Showing that the integer points are correct is usually the easy part of an LP formulation. Showing that there are no additional fractional vertices is often the central theorem. Next unit develops two methods: perturb a fractional point in opposite directions, and use determinants of totally unimodular matrices.

2.4 Description size and algorithms

A formal description using all supporting inequalities always exists:

$$w^\top x \le \max_{S \in \mathcal{S}} w^\top \chi^S \qquad \text{for every } w \in \mathbb{R}^n.$$

This does not yet give an algorithm: evaluating the right-hand side is the original optimization problem. We must distinguish a polynomial-size explicit formulation from an exponentially large formulation equipped with an efficient separation oracle. Unit 4 studies separation; Unit 8 asks whether extra variables can make a formulation smaller.

Method	Geometric idea	Complexity perspective
Simplex	Move between adjacent bases/vertices.	Common pivot rules have exponential worst cases; degeneracy needs care.
Ellipsoid	Shrink a containing ellipsoid using separation.	Polynomial bit complexity under the usual rational-encoding assumptions.
Interior point	Follow suitably controlled interior points.	Polynomial-time methods with strong practical performance.

"Polynomial time" includes dependence on the binary encoding length of coefficients. Strongly polynomial time asks for an arithmetic-operation bound polynomial in the combinatorial dimensions, together with polynomially bounded intermediate encodings. This distinction explains why the slides separately raise the strongly polynomial LP question. We will use LP solvability as a tool, while making size and precision assumptions explicit.

3 Exercises for independent study

1. Convert $\max(2u - v)$ subject to $u + v = 3$, $u - 2v \leq 4$, $u \in \mathbb{R}$, and $v \geq 0$ into form (1.1). Explain how to recover the original optimum value.
2. List every vertex of (1.2). Identify its tight constraints and objective value. Which intersections of boundary lines fail feasibility?
3. Give a bounded LP with infinitely many optimal solutions, and an unbounded feasible region with a unique finite optimum.
4. For the segment $\{(t, t, 0) : 0 \leq t \leq 1\} \subseteq \mathbb{R}^3$, determine its dimension, vertices, and independent tight constraints at each vertex.
5. On a triangle with unit edge weights, compare the integer optimum of the matching formulation with the value at $x_e = 1/2$ for every edge. What does this prove about integrality?
6. Suppose Q is integral and a solver returns a fractional optimal point. Explain why an integral optimum still exists and how repeated optimization over the optimal face can select a vertex.

Source and reading notes

Primary source: the instructor's slide folder

`1.LPFoundations`

`Topics.tex`.

Optional sections follow the optional slides where present. Explanatory examples and exercises are added to support independent reading; they do not change the lecture allocation.

Course website. The early-course notes linked there were used as examples of organization and level of detail, not as the source of this module's subject matter.