Advanced Algorithms / Fall 2026

Linear Programming and Optimization

Separation and Applications

Cut polytopes, network flow, and zero-sum games

Lecture 4: 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

An LP can be useful even when its constraints cannot all be listed. Part I replaces explicit constraints by an algorithm that finds a violated one. Part II uses the same path inequalities to connect flow and cut, then uses LP duality to prove the minimax theorem for finite zero-sum games.

1. Implement a separation oracle for path-covering constraints.
2. Prove the cut-polytope theorem, including the step from coordinatewise domination to an exact convex combination.
3. Derive max-flow/min-cut and minimax from LP formulations and rounding.

The optional appendices expand the spanning-tree separation exercise and the ellipsoid volume argument from the slides. All capacities and edge lengths below are nonnegative unless explicitly stated otherwise.

Contents

1 Part I: Separation oracles and cut polytopes **1**
1.1 Asking for a violated inequality . . . 1
1.2 Cuts defined as edge-deletion sets . . . 1
1.3 Threshold rounding and exact marginals . . . 1

2 Part II: Flow and zero-sum games **2**
2.1 Two flow formulations . . . 2
2.2 Max-flow/min-cut from duality . . . 2
2.3 Finite zero-sum games . . . 3

3 Exercises for independent study **3**

A Optional: spanning-tree and matching separation **4**
A.1 Spanning trees via a closure network . . . 4
A.2 Odd cuts and redundant inequalities . . . 4

B Optional: the ellipsoid volume calculation **5**

1 Part I: Separation oracles and cut polytopes

1.1 Asking for a violated inequality

Definition 1.1. A separation oracle for $P \subseteq \mathbb{R}^n$ receives a candidate z and either certifies $z \in P$ or returns a, β such that $a^\top x \le \beta$ for every $x \in P$, but $a^\top z > \beta$. It must return a valid separating inequality, not merely the answer "outside".

For an explicitly listed polynomial-size LP, test every row. For an exponential family, the problem is to optimize the amount of violation without enumerating every row. Oracle running time includes the encoding length of the rational query point.

Ellipsoid feasibility: the geometric outline

1. Start with an ellipsoid containing the target feasible set.
2. Query its center. If the center is feasible, return it.
3. Otherwise keep the side of the separating hyperplane containing the feasible set, and enclose the corresponding central half-ellipsoid in a smaller-volume ellipsoid.
4. Repeat with quantitative radius, precision, and stopping bounds.

Optimization reduces to feasibility by adding a trial objective constraint, such as $c^\top x \le \gamma$. This extra row can itself be checked directly. With the usual rational encoding, size bounds, and appropriate treatment of lower-dimensional sets, separation yields polynomial-time optimization. The geometric outline alone is not a complete bit-complexity proof; the appendix explains the volume bound and its assumptions.

1.2 Cuts defined as edge-deletion sets

Let $G = (V, E)$ be a directed graph with distinct terminals s, t. Here an s–t cut means any edge set C whose deletion destroys every directed s–t path. Thus a superset of a cut is also a cut. This convention is broader than requiring C to equal the outgoing boundary $\delta^+(S)$ of a vertex set.

Theorem 1.2 (Path description of the cut polytope)**.** *Let $\mathcal{P}$ be all simple directed s–t paths. Then*

$$\operatorname{conv}\{\chi^C : C \text{ is an } s\text{–}t \text{ cut}\} = Q := \{x \in [0,1]^E : x(P) \ge 1 \text{ for every } P \in \mathcal{P}\}.$$

A cut hits every path, so its indicator satisfies these inequalities. For separation, first check $0 \le x_e \le 1$, then compute a shortest s–t path with edge lengths x_e. A path of length below 1 is a violated row; otherwise all path rows hold. If no path exists, there are no path rows to violate.

1.3 Threshold rounding and exact marginals

Proof of Theorem 1.2. For feasible x, let $d(v)$ be the shortest-path distance from s, capped at 1; unreachable vertices also receive distance 1. Then $d(s) = 0$, $d(t) = 1$, and $d(v) - d(u) \le x_{uv}$ for every arc uv. Choose θ uniformly in $(0, 1)$, set $S_\theta = \{v : d(v) \le \theta\}$, and let $C_\theta = \delta^+(S_\theta)$. This is an s–t cut.

$$p_{uv} := \mathbb{P}[uv \in C_\theta] = \max\{d(v) - d(u), 0\} \le x_{uv}.$$

This only proves $\mathbb{E}\chi^{C_\theta} \le x$, not equality. To complete the convex-hull proof, whenever an edge e is absent, independently add it with probability

$$q_e = \frac{x_e - p_e}{1 - p_e} \quad \text{if } p_e < 1.$$

If $p_e = 1$, necessarily $x_e = 1$, and no change is needed. The enlarged set C' remains a cut and has $\mathbb{P}[e \in C'] = p_e + (1 - p_e)q_e = x_e$. There are only finitely many edge subsets, so the distribution expresses x as a convex combination of cut indicators. □

Two different uses of threshold rounding

For a nonnegative-cost objective, coordinatewise domination already proves that some threshold cut costs at most the fractional solution. To prove equality with the convex hull of all deletion cuts, we additionally need exact marginals and therefore the enlargement step.

2 Part II: Flow and zero-sum games

2.1 Two flow formulations

An arc-flow variable f_e satisfies $0 \le f_e \le c_e$ and conservation at every vertex except s, t. Maximize net inflow at t. The conservation matrix formulation is compact, with one variable per arc. For duality, another formulation is especially transparent: put nonnegative flow f_P on each simple s–t path.

$$\begin{aligned} \max \quad & \sum_{P \in \mathcal{P}} f_P && \qquad \min \quad \sum_{e \in E} c_e y_e \\ \text{s.t.} \quad & \sum_{P \ni e} f_P \le c_e \quad (e \in E) && \qquad \text{s.t.} \quad \sum_{e \in P} y_e \ge 1 \quad (P \in \mathcal{P}) \\ & f_P \ge 0 && \qquad y_e \ge 0. \end{aligned} \tag{2.1}$$

A path solution immediately induces an arc flow. Conversely, a nonnegative arc flow of nonnegative value can be decomposed into s–t paths and directed cycles. Subtract the minimum residual flow on a discovered path or cycle; at least one positive arc disappears each time. Discarding cycles never increases arc usage and preserves net flow value. Hence the two formulations have the same optimum.

The left program is a packing LP: path flows compete for capacities. Its dual assigns a nonnegative length to each edge and makes every s–t path have length at least 1. It has exponentially many rows, but the shortest-path separation oracle applies. The compact arc LP is often preferable computationally; the path LP makes the theorem easier to see.

2.2 Max-flow/min-cut from duality

Theorem 2.1. *For a finite directed network with nonnegative finite capacities, the maximum s–t flow value equals the minimum capacity of an s–t cut.*

Proof. The primal in (2.1) is feasible at zero and bounded, so its optimum equals the finite dual optimum by strong duality. Given any dual-feasible y, define capped distances using lengths y and take a random threshold cut as before. For each edge, its crossing probability is at most y_e, so

$$\mathbb{E}[c(C_\theta)] \le \sum_e c_e y_e.$$

Some threshold cut costs at most the dual optimum. Conversely, every cut indicator is dual-feasible, so the dual optimum is at most the minimum cut. The two inequalities force equality. If y is optimal, every threshold cut occurring with positive probability must itself be optimal: none can be cheaper than the minimum cut, while their average is no larger. □

This proof separates two integrality facts. Integral capacities imply an integral optimum flow by TU. In contrast, existence of an integral optimal cut vector holds for arbitrary nonnegative real capacities; it comes from threshold rounding, not from integrality of capacities.

2.3 Finite zero-sum games

Let $M \in \mathbb{R}^{m\times n}$ be a payoff matrix. The row player minimizes the payoff; the column player maximizes it. A mixed strategy is a probability distribution: $x \in \Delta_m = \{x \geq 0 : \mathbf{1}^\top x = 1\}$ and $y \in \Delta_n$. Independent mixed play has expected payoff $x^\top My$. Against a fixed x, the column player can choose a pure best response, because a weighted average cannot exceed the largest column payoff. Similarly, against a fixed y, a pure row attains the smallest row payoff. Therefore the two order-of-play values are

$$\min_{x\in\Delta_m} \max_j x^\top M_{\cdot j}, \qquad \max_{y\in\Delta_n} \min_i M_{i\cdot} y.$$

Theorem 2.2 (Minimax theorem)**.**

$$\min_{x\in\Delta_m} \max_{y\in\Delta_n} x^\top My = \max_{y\in\Delta_n} \min_{x\in\Delta_m} x^\top My =: V. \tag{2.2}$$

Proof. The row player's optimization and its dual are

$$\begin{array}{lll} \min & R & \max \quad C \\ \text{s.t.} & M^\top x \leq R\mathbf{1} & \text{s.t.} \quad My \geq C\mathbf{1} \\ & \mathbf{1}^\top x = 1,\ x \geq 0 & \quad\ \mathbf{1}^\top y = 1,\ y \geq 0, \\ & R \in \mathbb{R} & \quad\ C \in \mathbb{R}. \end{array}$$

The unrestricted variable R forces the dual normalization $\mathbf{1}^\top y = 1$; the primal normalization gives the free variable C. Both programs are feasible and have values bounded by the smallest and largest matrix entries. Strong duality proves (2.2). □

Optimal strategies x^*, y^* form a saddle point:

$$x^{*\top} My \leq x^{*\top} My^* = V \leq x^\top My^* \qquad (x \in \Delta_m,\ y \in \Delta_n).$$

Each player guarantees the value without knowing the other's chosen strategy. This does not mean every strategy is optimal, nor that pure strategies suffice. In rock–paper–scissors, the uniform mixed strategies give value zero; a revealed pure strategy can be exploited.

Complementary slackness sharpens the equilibrium description:

$$x^*_i > 0 \Longrightarrow M_{i\cdot} y^* = V, \qquad y^*_j > 0 \Longrightarrow x^{*\top} M_{\cdot j} = V.$$

Thus every pure action used with positive probability has the same payoff against the opponent's optimal strategy. Unused rows may have larger payoff, and unused columns may have smaller payoff, consistently with the players' optimization directions. Unit 7 will approximate this equilibrium using multiplicative weights.

3 Exercises for independent study

1. Implement the cut separation oracle on a graph of your choice and exhibit a violated path inequality for a fractional candidate.
2. Give an example where $\mathbb{E}\chi^{C_\theta} < x$ in some coordinate. Compute the enlargement probabilities q_e.

3. Show how sorting the distinct distances derandomizes threshold-cut selection: evaluate one threshold between each consecutive pair.
4. Solve the zero-sum game $M = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$ and verify the saddle-point inequalities.
5. Derive both game LPs directly using the sign rules of Unit 3. Explain why the variables R, C cannot be constrained nonnegative for an arbitrary payoff matrix.
6. Optional: verify the minimum-cut reduction for spanning-tree separation below, including why the forced vertex is necessary.

A Optional: spanning-tree and matching separation

A.1 Spanning trees via a closure network

First check $x \geq 0$ and $x(E) = |V| - 1$. A violated spanning-tree inequality has $x(E[S]) > |S| - 1$. Guess a vertex $r \in S$, assign vertex costs $w_r = 0$ and $w_v = 1$ for $v \neq r$, and maximize

$$\sum_{e \in E[S]} x_e - \sum_{v \in S} w_v, \qquad r \in S.$$

Build a network with source a, sink b, one node for every original edge and vertex, an arc $a \to e$ of capacity x_e, arcs $e \to u, e \to v$ of capacity H for $e = uv$, and an arc $v \to b$ of capacity w_v. Add $a \to r$ of capacity H. Choose $H > \sum_e x_e + \sum_v w_v$. A minimum cut never crosses an H-arc, so an edge-node on the source side forces both endpoints there and r is present. For a chosen vertex set S, the cheapest choice puts every internal edge-node on the source side. The cut capacity becomes

$$\sum_e x_e - \sum_{e \in E[S]} x_e + \sum_{v \in S} w_v.$$

It is below $\sum_e x_e$ exactly when S violates a spanning-tree row. Repeat for all r. A singleton cannot have positive profit in a loopless graph, so the test automatically finds a set of size at least two. This expands the densest-subgraph/max-flow hint in the slides.

A.2 Odd cuts and redundant inequalities

For perfect matching, first check nonnegativity and degree equations, then find a minimum-capacity odd cut under capacities x_e. A Gomory–Hu cut tree of the undirected capacitated graph allows this: among tree edges whose removal has an odd-cardinality shore, choose the minimum weight; its fundamental graph cut is a minimum odd cut. Here the number of vertices is even. The parity reason is that these odd-shore tree edges form a set with odd degree at every tree vertex, so every odd vertex set crosses at least one of them. For such a crossing tree edge uv, its weight is a minimum u–v cut value and cannot exceed the capacity of the candidate odd cut. An inequality description need not be irredundant. In a spanning-tree system, a disconnected induced subgraph's inequality follows by summing inequalities for its components. If an induced connected subgraph has a bridge and at least three vertices, sum the two side inequalities and the bridge bound $x_e \leq 1$. Within the affine hull, only an irredundant description corresponds one-for-one to facets.

B Optional: the ellipsoid volume calculation

An invertible affine map preserves volume ratios and reduces a central half-ellipsoid to the half unit ball $\{x : \|x\|_2 \le 1,\ x_1 \ge 0\}$. For dimension $n \ge 2$, put

$$c = \frac{1}{n+1}, \qquad a = \frac{n}{n+1}, \qquad b = \frac{n}{\sqrt{n^2-1}}.$$

The ellipsoid centered at $(c, 0, \dots, 0)$ with semiaxes $(a, b, \dots, b)$ contains the half-ball. Indeed, for $t = x_1 \in [0,1]$ and $\sum_{i\ge 2} x_i^2 \le 1 - t^2$,

$$\frac{(t-c)^2}{a^2} + \frac{1-t^2}{b^2} = 1 - \frac{2(n+1)}{n^2} t(1-t) \le 1.$$

Its volume relative to the unit ball is ab^{n-1}. Using $\ln(1+u) \le u$,

$$\begin{aligned} \ln(ab^{n-1}) &= \ln \frac{n}{n+1} + \frac{n-1}{2} \ln \frac{n^2}{n^2-1} \\ &\le -\frac{1}{n+1} + \frac{n-1}{2(n^2-1)} = -\frac{1}{2(n+1)}. \end{aligned}$$

After N unsuccessful iterations the containing volume shrinks by a factor at most $\exp(-N/(2(n+1)))$. If the initial region lies in a radius-R ball and every nonempty target contains a radius-r ball, then

$$N > 2n(n+1)\ln(R/r)$$

would make a remaining nonempty target impossible. This is $O(n^2 \log(R/r))$ oracle calls. Dimension one uses interval bisection.

What the volume argument does not assume for free

A lower-dimensional nonempty polytope has zero ambient volume and contains no positive-radius ambient ball. The inner-radius promise therefore needs an affine-hull reduction, controlled thickening, or the full rational LP treatment. Finite-precision updates also require analysis. The displayed geometric count is not by itself a proof for arbitrary oracle-presented convex sets.

Source and reading notes

Primary source: the instructor's slide folder

`4.SeparationAndApplications`

`Topics.tex` and `Optional.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.