Advanced Algorithms / Fall 2026

Linear Programming and Optimization

Linear Programming Duality

Certificates, alternatives, and complementary slackness

Lecture 3: 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

A feasible solution shows what can be achieved. A dual solution proves what no feasible solution can improve upon. This unit develops that certificate viewpoint, proves why optimal certificates exist for finite LP optima, and identifies exactly when the weak-duality inequalities become equalities.

1. Derive a dual by taking weighted sums of constraints, including equality and free-variable cases.
2. Prove weak duality and the Farkas alternatives.
3. Follow the ε-certificate proof of strong duality and justify attainment.
4. Use complementary slackness to verify and recover optimal solutions.

Part I constructs the dual and its certificate interpretation. Part II proves strong duality and complementary slackness. The short geometric appendix fills in the closed-cone and attainment facts used by the proof.

Contents

1 Part I: Constructing dual programs **1**
1.1 From lower bounds to an optimization problem . . . 1
1.2 Weak duality and optimality certificates . . . 1
1.3 Signs, equalities, and free variables . . . 1

2 Part II: Strong duality and complementary slackness **2**
2.1 A theorem of alternatives . . . 2
2.2 Strong duality via arbitrarily accurate certificates . . . 2
2.3 Feasibility and unboundedness . . . 3
2.4 Complementary slackness . . . 3

3 Exercises for independent study **4**

A Geometric details used by the proof **4**
A.1 Why a finitely generated cone is closed . . . 4
A.2 Constructing the separating vector . . . 4
A.3 Why the finite optimum is attained . . . 4

1 Part I: Constructing dual programs

1.1 From lower bounds to an optimization problem

Return to the LP from Unit 1. A feasible point $(1, 4)$ has value 23, but we need a lower bound valid for every feasible point. Nonnegative combinations of its constraints provide such bounds. For example, the coefficient vectors $(2, 1, 0)$, $(1, 1, 1)$, and $(4, 0, 0)$ certify lower bounds 16, 19, and 20, respectively. The combination $(3, 0, 1)$ certifies 23. Why are these multipliers legal?

With $y_1, y_2, y_3 \geq 0$, summing the three constraints gives

$$(y_1 + y_2 + 4y_3)x_1 + (y_1 + 2y_2 + y_3)x_2 \geq 5y_1 + 6y_2 + 8y_3.$$

If the two left coefficients do not exceed 7 and 4, nonnegativity of x makes the left side at most $7x_1 + 4x_2$. Thus the strongest bound obtainable this way is itself the LP

$$\begin{aligned} \max \quad & 5y_1 + 6y_2 + 8y_3 \\ \text{s.t.} \quad & y_1 + y_2 + 4y_3 \leq 7, \\ & y_1 + 2y_2 + y_3 \leq 4, \\ & y_1, y_2, y_3 \geq 0. \end{aligned} \tag{1.1}$$

The original LP is the primal; this bound-maximization LP is its dual. A primal variable becomes a dual constraint, and a primal constraint becomes a dual variable. The right-hand side b becomes the dual objective, and the primal objective c becomes the dual right-hand side.

1.2 Weak duality and optimality certificates

$$\begin{array}{ll|ll} (P) & \min c^\top x & (D) & \max b^\top y \\ & Ax \geq b,\ x \geq 0 & & A^\top y \leq c,\ y \geq 0. \end{array} \tag{1.2}$$

Theorem 1.1 (Weak duality)**.** *For every primal-feasible x and dual-feasible y, $b^\top y \leq c^\top x$.*

Proof.

$$b^\top y \leq (Ax)^\top y = x^\top A^\top y \leq x^\top c.$$

The first inequality uses $Ax \geq b$ and $y \geq 0$; the second uses $A^\top y \leq c$ and $x \geq 0$. Neither step may discard a sign condition. □

Corollary 1.2. *If feasible x, y satisfy $c^\top x = b^\top y$, then both are optimal. More generally, their gap bounds the suboptimality of each feasible solution.*

When the common optimum is V, the inequality $b^\top y \leq V \leq c^\top x$ gives $c^\top x - V \leq c^\top x - b^\top y$. This is useful even before an optimization algorithm finishes: it turns intermediate feasible iterates into a numerical certificate. In approximation algorithms, a feasible dual objective is a lower bound on the integral minimization optimum as well.

Example 1.3. For the running example, $x = (1, 4)$ and $y = (3, 0, 1)$ are feasible and both have objective 23. Checking these two vectors is a complete optimality proof; no knowledge of the method that produced them is needed.

1.3 Signs, equalities, and free variables

For a minimization primal, the rule is always to construct a lower bound. A $\geq$ row receives a nonnegative multiplier, a $\leq$ row a nonpositive multiplier, and an equality an unrestricted multiplier. On the variable side, $c - A^\top y$ must have the right sign to ensure $x^\top(c - A^\top y) \geq 0$.

Primal feature (minimization)	Dual requirement
$a_i^\top x \geq b_i$	$y_i \geq 0$
$a_i^\top x \leq b_i$	$y_i \leq 0$
$a_i^\top x = b_i$	$y_i \in \mathbb{R}$
$x_j \geq 0$	$(A^\top y)_j \leq c_j$
$x_j \leq 0$	$(A^\top y)_j \geq c_j$
$x_j \in \mathbb{R}$	$(A^\top y)_j = c_j$

One can always convert to the standard form first, but the sign table avoids doubling variables and constraints unnecessarily. For a maximization primal, either derive an upper bound directly or negate the objective and apply the minimization convention. Mixing rows from different conventions is a common source of incorrect duals.

Example 1.4 (An equality and a free variable)**.** For $\min 2x + z$ subject to $x + z = 3$, $x \geq 0$, $z \in \mathbb{R}$, the multiplier y is free, the x column requires $y \leq 2$, and the z column requires $y = 1$. The dual is therefore $\max 3y$ subject to these two constraints. It has value 3, attained in the primal at $x = 0, z = 3$.

2 Part II: Strong duality and complementary slackness

2.1 A theorem of alternatives

Theorem 2.1 (Farkas lemma)**.** *Exactly one of the following systems is feasible:*

$$Ax = b,\ x \geq 0; \qquad A^\top y \geq 0,\ b^\top y < 0,\ y \in \mathbb{R}^m.$$

Proof. They cannot both be feasible because $b^\top y = x^\top A^\top y \geq 0$ would contradict $b^\top y < 0$. For existence of the certificate when b is infeasible, consider the finitely generated closed convex cone $K = \{Ax : x \geq 0\}$. Separate $b \notin K$ from K. Because $0 \in K$ and K is closed under nonnegative scaling, the separating functional can be chosen with $y^\top u \geq 0$ for every $u \in K$ and $y^\top b < 0$. Applying it to each column of A yields $A^\top y \geq 0$. The appendix supplies the geometric details. □

Corollary 2.2 (Inequality form)**.** *The system $Ax \leq b,\ x \geq 0$ is infeasible if and only if some $y \geq 0$ satisfies $A^\top y \geq 0$ and $b^\top y < 0$.*

Proof. Introduce slack $s \geq 0$ and write $Ax + s = b$. Apply Theorem 2.1 to the matrix $[A\ I]$. The columns of I force $y \geq 0$. □

A Farkas certificate is a concise reason for infeasibility. It does not merely say that a search failed; it gives a linear combination that every alleged feasible solution would have to contradict.

2.2 Strong duality via arbitrarily accurate certificates

Theorem 2.3 (Strong duality)**.** *If the primal in (1.2) is feasible with finite optimum P, its dual attains an optimum D and $D = P$.*

Proof. Fix any $\varepsilon > 0$. No nonnegative x satisfies both $Ax \geq b$ and $c^\top x \leq P - \varepsilon$. Apply Corollary 2.2 to

$$\begin{pmatrix} -A \\ c^\top \end{pmatrix} x \leq \begin{pmatrix} -b \\ P - \varepsilon \end{pmatrix}, \qquad x \geq 0.$$

We obtain $y \geq 0$ and $\alpha \geq 0$ with

$$-A^\top y + \alpha c \geq 0, \qquad -b^\top y + \alpha(P - \varepsilon) < 0. \tag{2.1}$$

We must justify dividing by α. If $\alpha = 0$, then $A^\top y \leq 0$ and $b^\top y > 0$. For any primal-feasible x_0, however, $b^\top y \leq x_0^\top A^\top y \leq 0$, a contradiction. Thus $\alpha > 0$. Set $\widehat{y} = y/\alpha$; then (2.1) gives

$$\widehat{y} \geq 0, \qquad A^\top \widehat{y} \leq c, \qquad b^\top \widehat{y} > P - \varepsilon.$$

For every $\varepsilon > 0$ there is a dual-feasible value above $P - \varepsilon$, while weak duality bounds all such values by P. The dual supremum is therefore P. Finally, a linear objective with finite supremum on a nonempty polyhedron attains it (appendix). This last step is needed: equality of a supremum with P alone would not yet supply an optimal dual vector. □

2.3 Feasibility and unboundedness

Dualizing twice returns the primal after harmless sign conversions. Weak and strong duality give the following complete implications for a primal–dual pair:

- If one side has a finite optimum, both sides do, with equal values.
- If the primal is unbounded below, the dual is infeasible. If the dual is unbounded above, the primal is infeasible.
- If one side is infeasible, the other can be infeasible or unbounded; infeasibility does not force unboundedness.

For example, $A = (0)$, $b = (1)$, $c = (-1)$ makes the primal constraint $0x \geq 1$ and the dual constraint $0y \leq -1$, so both sides are infeasible. This prevents a common overstatement of duality.

2.4 Complementary slackness

For a feasible pair (x, y), expand the objective gap into nonnegative terms:

$$c^\top x - b^\top y = \sum_j x_j(c_j - (A^\top y)_j) + \sum_i y_i((Ax)_i - b_i). \tag{2.2}$$

Theorem 2.4 (Complementary slackness)**.** *A feasible pair is optimal if and only if*

$$x_j(c_j - (A^\top y)_j) = 0 \quad \forall j, \qquad y_i((Ax)_i - b_i) = 0 \quad \forall i.$$

Proof. Optimality and strong duality make (2.2) zero. Since every summand is nonnegative, each must vanish. Conversely, vanishing products give equal objectives, so weak duality certifies optimality. □

Thus a positive dual variable forces its primal row tight, and a positive primal variable forces its dual row tight. The converses are false: a tight row may have multiplier zero. Complementary slackness always comes with feasibility; products equal to zero alone do not prove optimality.

Example 2.5 (Recovering the numerical certificate)**.** At $x = (1, 4)$ the second primal row has positive slack, so $y_2 = 0$. Both primal coordinates are positive, so both dual inequalities must be tight. Solving $y_1 + 4y_3 = 7$ and $y_1 + y_3 = 4$ gives $y_3 = 1$, $y_1 = 3$. Conversely, these positive multipliers require $x_1 + x_2 = 5$ and $4x_1 + x_2 = 8$, recovering $(1, 4)$.

Dual multipliers also quantify sensitivity. Holding A, c fixed, any optimal multiplier y^* at right-hand side b remains dual-feasible when b changes to b'. Whenever the new primal has finite optimum $V(b')$,

$$V(b') \geq b'^{\top} y^* = V(b) + (b' - b)^{\top} y^*.$$

This is a rigorous lower supporting bound. Equality requires additional stability of the optimal basis or face and must not be assumed automatically.

3 Exercises for independent study

1. Derive the dual of $\min(3x_1 + 2x_2)$ subject to $x_1 + x_2 \geq 4$, $x_1 - x_2 = 1$, and $x \geq 0$. Find a primal–dual optimal pair.
2. Give a Farkas certificate for infeasibility of $x_1 + x_2 \leq -1$, $x_1, x_2 \geq 0$.
3. Construct examples showing that a tight primal row can have zero optimal multiplier, and a zero primal variable can correspond to a tight dual row.
4. In the strong-duality proof, identify the precise point at which primal feasibility is used. What fails if it is omitted?
5. Verify the dual of the dual algebraically for (1.2), keeping every sign conversion explicit.
6. If a primal-feasible point has cost 10.2 and a dual-feasible point has value 10, what additive guarantee follows? Under what positivity assumption does a multiplicative guarantee follow?

A Geometric details used by the proof

A.1 Why a finitely generated cone is closed

Every vector in $K = \{Ax : x \geq 0\}$ has a representation using linearly independent columns of A. To see this, take a representation with dependent positive-support columns, choose a nonzero dependence $Ad = 0$ on that support, change its sign so some $d_i > 0$, and subtract θd from the coefficient vector until one positive coordinate reaches zero. Nonnegativity is preserved and the support shrinks. Repeat.

For a convergent sequence in K, pass to a subsequence using the same independent column set; there are only finitely many choices. If its column matrix is B, the coefficient vectors equal $(B^{\top}B)^{-1}B^{\top}z_k$ and converge to a nonnegative vector. The limit is therefore in K. This proves closedness.

A.2 Constructing the separating vector

Let z be a nearest point in the closed convex cone K to $b \notin K$; a nearest point exists by restricting to a sufficiently large closed ball. Put $y = z - b$. Minimality along the segment from z to any $u \in K$ gives $y^{\top}(u - z) \geq 0$. Taking $u = 0$ and $u = 2z$ yields $y^{\top}z = 0$. Thus $y^{\top}u \geq 0$ for every $u \in K$, while $y^{\top}b = -\|z - b\|^2 < 0$. This is exactly the Farkas certificate.

A.3 Why the finite optimum is attained

The image of a polyhedron under a linear map is a polyhedron. One way to prove this is to introduce equations for the image coordinates and eliminate original variables by Fourier–Motzkin elimination: upper and lower bounds on an eliminated variable generate finitely many pairwise compatibility inequalities. The image of a feasible polyhedron under $y \mapsto b^{\top}y$ is therefore a one-dimensional

polyhedron, hence a closed interval, ray, point, or the whole line. If it is nonempty with finite supremum, that endpoint belongs to the image. This proves attainment without assuming the dual feasible set is bounded.

Source and reading notes

Primary source: the instructor's slide folder

`3.Duality`

`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.