CS675: Convex and Combinatorial Optimization
Fall 2016
Introduction to Linear Programming

Instructor: Shaddin Dughmi

Outline

1. Linear Programming Basics
2. Duality and Its Interpretations
3. Properties of Duals
4. Weak and Strong Duality
5. Formal Proof of Strong Duality of LP
6. Consequences of Duality
7. More Examples of Duality

Outline

A Brief History

- The forefather of convex optimization problems, and the most ubiquitous.
- Developed by Kantorovich during World War II (1939) for planning the Soviet army's expenditures and returns. Kept secret.
- Discovered a few years later by George Dantzig, who in 1947 developed the simplex method for solving linear programs
- John von Neumann developed LP duality in 1947, and applied it to game theory
- Polynomial-time algorithms: Ellipsoid method (Khachiyan 1979), interior point methods (Karmarkar 1984).

LP General Form

$$
\begin{array}{lll}
\text{minimize (or maximize)} & c^\intercal x & \\
\text{subject to} & a_i^\intercal x \leq b_i, & \text{for } i \in \mathcal{C}^1. \\
& a_i^\intercal x \geq b_i, & \text{for } i \in \mathcal{C}^2. \\
& a_i^\intercal x = b_i, & \text{for } i \in \mathcal{C}^3.
\end{array}
$$

- Decision variables: $x \in \mathbb{R}^n$
- Parameters:
  - $c \in \mathbb{R}^n$ defines the linear objective function
  - $a_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ define the i'th constraint.

Standard Form

$$\begin{array}{lll} \text{maximize} & c^\intercal x & \\ \text{subject to} & a_i^\intercal x \leq b_i, & \text{for } i = 1, \ldots, m. \\ & x_j \geq 0, & \text{for } j = 1, \ldots, n. \end{array}$$

Every LP can be transformed to this form

- minimizing $c^\intercal x$ is equivalent to maximizing $-c^\intercal x$
- $\geq$ constraints can be flipped by multiplying by -1
- Each equality constraint can be replaced by two inequalities
- Uconstrained variable x_j can be replaced by $x_j^+ - x_j^-$, where both x_j^+ and x_j^- are constrained to be nonnegative.

Geometric View

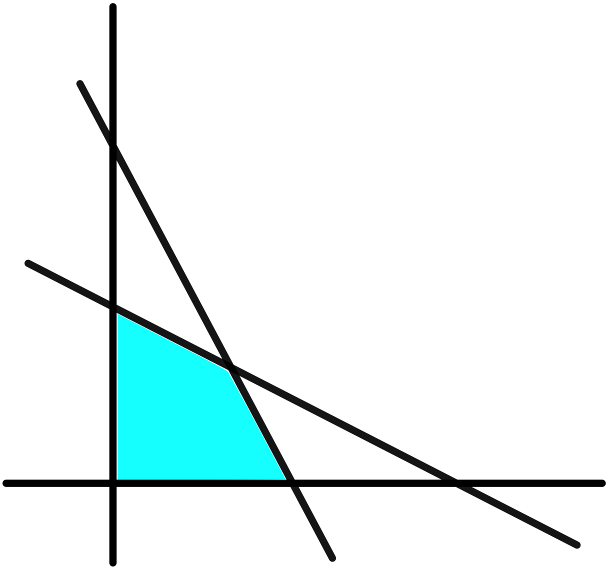

Geometric View

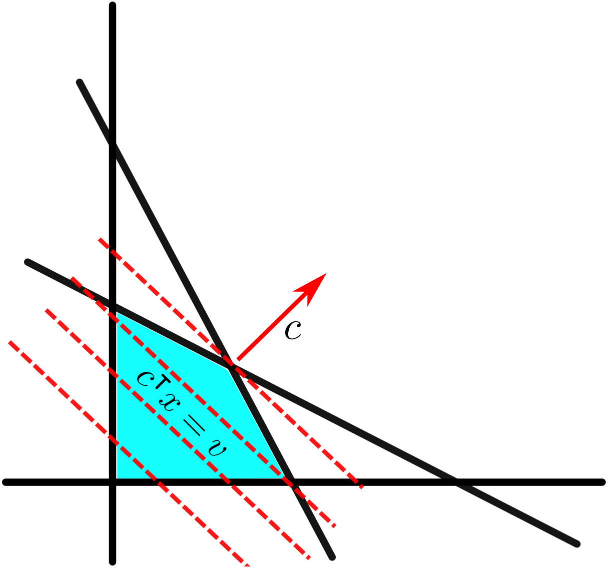

A 2-D example

$$\begin{array}{ll} \text{maximize} & x_1 + x_2 \\ \text{subject to} & x_1 + 2x_2 \leq 2 \\ & 2x_1 + x_2 \leq 2 \\ & x_1, x_2 \geq 0 \end{array}$$

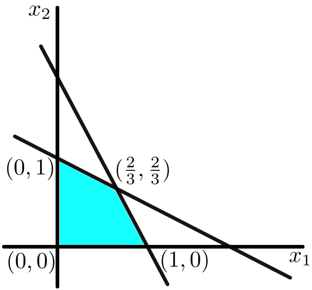

Application: Optimal Production

- n products, m raw materials
- Every unit of product j uses a_{ij} units of raw material i
- There are b_i units of material i available
- Product j yields profit c_j per unit
- Facility wants to maximize profit subject to available raw materials

$$
\begin{array}{lll}
\text{maximize} & c^\intercal x & \\
\text{subject to} & a_i^\intercal x \leq b_i, & \text{for } i = 1, \ldots, m. \\
 & x_j \geq 0, & \text{for } j = 1, \ldots, n.
\end{array}
$$

Terminology

- Hyperplane: The region defined by a linear equality
- Halfspace: The region defined by a linear inequality $a_i^\intercal x \leq b_i$.
- Polyhedron: The intersection of a set of linear inequalities
  - Feasible region of an LP is a polyhedron
- Polytope: Bounded polyhedron
  - Equivalently: convex hull of a finite set of points
- Vertex: A point x is a vertex of polyhedron P if $\nexists y \neq 0$ with $x + y \in P$ and $x - y \in P$
- Face of P: The intersection with P of a hyperplane H disjoint from the interior of P

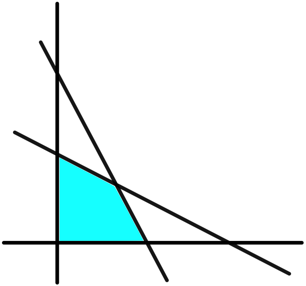

Basic Facts about LPs and Polyhedrons

Fact

Feasible regions of LPs (i.e. polyhedrons) are convex

Basic Facts about LPs and Polyhedrons

Fact

Feasible regions of LPs (i.e. polyhedrons) are convex

Fact

Set of optimal solutions of an LP is convex

- In fact, a face of the polyhedron
- intersection of P with hyperplane $c^\intercal x = OPT$

Basic Facts about LPs and Polyhedrons

Fact

Feasible regions of LPs (i.e. polyhedrons) are convex

Fact

Set of optimal solutions of an LP is convex

- In fact, a face of the polyhedron
- intersection of P with hyperplane $c^\intercal x = OPT$

Fact

At a vertex, n linearly independent constraints are satisfied with equality (a.k.a. tight)

Basic Facts about LPs and Polyhedrons

Fact

An LP either has an optimal solution, or is unbounded or infeasible

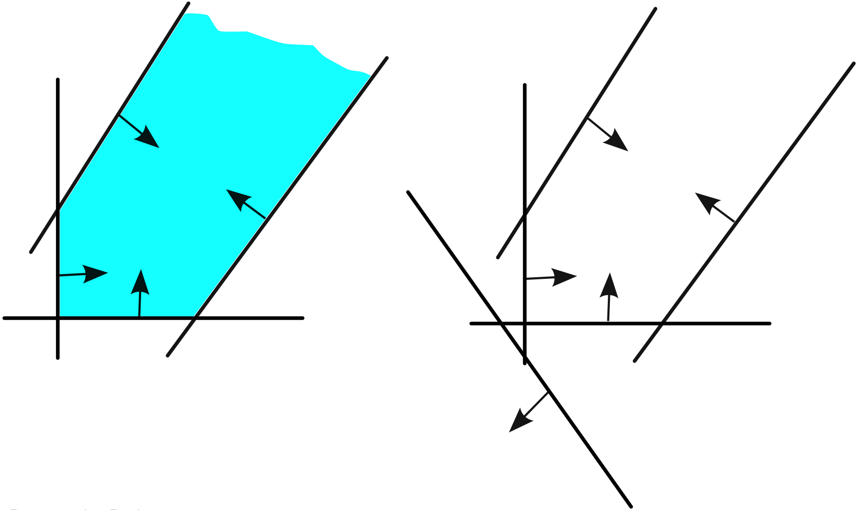

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Proof

- Assume not, and take a non-vertex optimal solution x with the maximum number of tight constraints

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Proof

- Assume not, and take a non-vertex optimal solution x with the maximum number of tight constraints
- There is $y \neq 0$ s.t. $x \pm y$ are feasible

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Proof

- Assume not, and take a non-vertex optimal solution x with the maximum number of tight constraints
- There is $y \neq 0$ s.t. $x \pm y$ are feasible
- y is perpendicular to the objective function and the tight constraints at x.
  - i.e. $c^\intercal y = 0$, and $a_i^\intercal y = 0$ whenever the i'th constraint is tight for x.

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Proof

- Assume not, and take a non-vertex optimal solution x with the maximum number of tight constraints
- There is $y \neq 0$ s.t. $x \pm y$ are feasible
- y is perpendicular to the objective function and the tight constraints at x.
  - i.e. $c^\intercal y = 0$, and $a_i^\intercal y = 0$ whenever the i'th constraint is tight for x.
- Can choose y s.t. $y_j < 0$ for some j

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Proof

- Assume not, and take a non-vertex optimal solution x with the maximum number of tight constraints
- There is $y \neq 0$ s.t. $x \pm y$ are feasible
- y is perpendicular to the objective function and the tight constraints at x.
  - i.e. $c^\intercal y = 0$, and $a_i^\intercal y = 0$ whenever the i'th constraint is tight for x.
- Can choose y s.t. $y_j < 0$ for some j
- Let α be the largest constant such that $x + \alpha y$ is feasible
  - Such an α exists

Fundamental Theorem of LP

If an LP in standard form has an optimal solution, then it has a vertex optimal solution.

Proof

- Assume not, and take a non-vertex optimal solution x with the maximum number of tight constraints
- There is $y \neq 0$ s.t. $x \pm y$ are feasible
- y is perpendicular to the objective function and the tight constraints at x.
  - i.e. $c^\intercal y = 0$, and $a_i^\intercal y = 0$ whenever the i'th constraint is tight for x.
- Can choose y s.t. $y_j < 0$ for some j
- Let α be the largest constant such that $x + \alpha y$ is feasible
  - Such an α exists
- An additional constraint becomes tight at $x + \alpha y$, a contradiction.

Counting non-zero Variables

Corollary

If an LP in standard form has an optimal solution, then there is an optimal solution with at most m non-zero variables.

$$\begin{array}{lll} \text{maximize} & c^\intercal x & \\ \text{subject to} & a_i^\intercal x \le b_i, & \text{for } i = 1, \ldots, m. \\ & x_j \ge 0, & \text{for } j = 1, \ldots, n. \end{array}$$

- e.g. for optimal production with n products and m raw materials, there is an optimal plan with at most m products.

Outline

Linear Programming Duality

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & b^\intercal y \\ \text{subject to} & A^\intercal y = c \\ & y \geq 0 \end{array}$$

- $A \in \mathbb{R}^{m \times n}$, $c \in \mathbb{R}^n$, $b \in \mathbb{R}^m$
- y_i is the dual variable corresponding to primal constraint $A_i x \leq b_i$
- $A_j^T y \geq c_j$ is the dual constraint corresponding to primal variable x_j

Linear Programming Duality: Standard Form, and Visualization

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

Linear Programming Duality: Standard Form, and Visualization

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

Linear Programming Duality: Standard Form, and Visualization

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

- y_i is the dual variable corresponding to primal constraint $A_i x \leq b_i$
- $A_j^T y \geq c_j$ is the dual constraint corresponding to primal variable x_j

Interpretation 1: Economic Interpretation

Recall the Optimal Production problem from last lecture

- n products, m raw materials
- Every unit of product j uses a_{ij} units of raw material i
- There are b_i units of material i available
- Product j yields profit c_j per unit
- Facility wants to maximize profit subject to available raw materials

Interpretation 1: Economic Interpretation

Primal LP

$$
\begin{array}{lll}
\text{max} & \sum_{j=1}^{n} c_j x_j & \\
\text{s.t.} & \sum_{j=1}^{n} a_{ij} x_j \leq b_i, & \text{for } i \in [m]. \\
 & x_j \geq 0, & \text{for } j \in [n].
\end{array}
$$

Interpretation 1: Economic Interpretation

Primal LP

$$\begin{array}{lll} \text{max} & \sum_{j=1}^{n} c_j x_j & \\ \text{s.t.} & \sum_{j=1}^{n} a_{ij} x_j \leq b_i, & \text{for } i \in [m]. \\ & x_j \geq 0, & \text{for } j \in [n]. \end{array}$$

Dual LP

$$\begin{array}{lll} \text{min} & \sum_{i=1}^{m} b_i y_i & \\ \text{s.t.} & \sum_{i=1}^{m} a_{ij} y_i \geq c_j, & \text{for } j \in [n]. \\ & y_i \geq 0, & \text{for } i \in [m]. \end{array}$$

Interpretation 1: Economic Interpretation

Primal LP

$$\begin{aligned} \text{max} \quad & \sum_{j=1}^{n} c_j x_j \\ \text{s.t.} \quad & \sum_{j=1}^{n} a_{ij} x_j \le b_i, && \text{for } i \in [m]. \\ & x_j \ge 0, && \text{for } j \in [n]. \end{aligned}$$

Dual LP

$$\begin{aligned} \text{min} \quad & \sum_{i=1}^{m} b_i y_i \\ \text{s.t.} \quad & \sum_{i=1}^{m} a_{ij} y_i \ge c_j, && \text{for } j \in [n]. \\ & y_i \ge 0, && \text{for } i \in [m]. \end{aligned}$$

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

Interpretation 1: Economic Interpretation

Primal LP

$$\begin{array}{lll} \max & \sum_{j=1}^{n} c_j x_j & \\ \text{s.t.} & \sum_{j=1}^{n} a_{ij} x_j \le b_i, & \text{for } i \in [m]. \\ & x_j \ge 0, & \text{for } j \in [n]. \end{array}$$

Dual LP

$$\begin{array}{lll} \min & \sum_{i=1}^{m} b_i y_i & \\ \text{s.t.} & \sum_{i=1}^{m} a_{ij} y_i \ge c_j, & \text{for } j \in [n]. \\ & y_i \ge 0, & \text{for } i \in [m]. \end{array}$$

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

- Dual variable y_i is a proposed price per unit of raw material i
- Dual price vector is feasible if facility has incentive to sell materials
- Buyer wants to spend as little as possible to buy materials

Interpretation 2: Finding the Best Upperbound

Consider the simple LP from last lecture

$$\begin{array}{ll} \text{maximize} & x_1 + x_2 \\ \text{subject to} & x_1 + 2x_2 \leq 2 \\ & 2x_1 + x_2 \leq 2 \\ & x_1, x_2 \geq 0 \end{array}$$

- We found that the optimal solution was at $(\frac{2}{3}, \frac{2}{3})$, with an optimal value of $4/3$.

Interpretation 2: Finding the Best Upperbound

Consider the simple LP from last lecture

$$\begin{aligned}\text{maximize} \quad & x_1 + x_2 \\ \text{subject to} \quad & x_1 + 2x_2 \le 2 \\ & 2x_1 + x_2 \le 2 \\ & x_1, x_2 \ge 0\end{aligned}$$

- We found that the optimal solution was at $(\frac{2}{3}, \frac{2}{3})$, with an optimal value of $4/3$.
- What if, instead of finding the optimal solution, we saught to find an upperbound on its value by combining inequalities?
  - Each inequality implies an upper bound of 2
  - Multiplying each by $\frac{1}{3}$ and summing gives $x_1 + x_2 \le 4/3$.

Interpretation 2: Finding the Best Upperbound

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

- Multiplying each row i by y_i and summing gives the inequality

$$y^T Ax \leq y^T b$$

Interpretation 2: Finding the Best Upperbound

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

- Multiplying each row i by y_i and summing gives the inequality

$$y^T Ax \leq y^T b$$

- When $y^T A \geq c^T$, the right hand side of the inequality is an upper bound on $c^T x$ for every feasible x.

$$c^T x \leq y^T Ax \leq y^T b$$

Interpretation 2: Finding the Best Upperbound

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

- Multiplying each row i by y_i and summing gives the inequality

$$y^T Ax \leq y^T b$$

- When $y^T A \geq c^T$, the right hand side of the inequality is an upper bound on $c^T x$ for every feasible x.

$$c^T x \leq y^T Ax \leq y^T b$$

- The dual LP can be thought of as trying to find the best upperbound on the primal that can be achieved this way.

Interpretation 3: Physical Forces

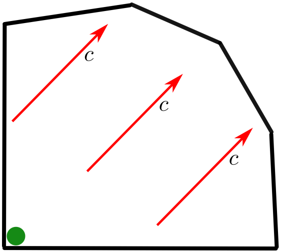

- Apply force field c to a ball inside bounded polytope $Ax \leq b$.

Interpretation 3: Physical Forces

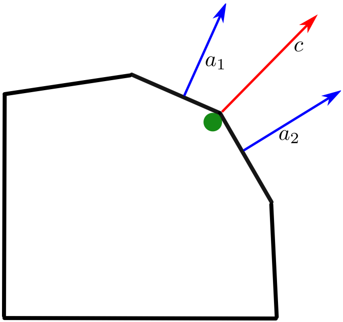

- Apply force field c to a ball inside bounded polytope $Ax \leq b$.
- Eventually, ball will come to rest against the walls of the polytope.

Interpretation 3: Physical Forces

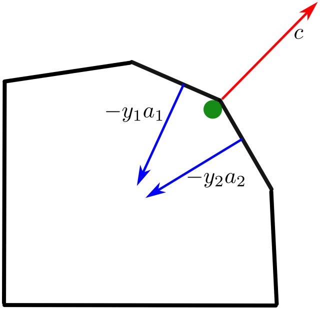

- Apply force field c to a ball inside bounded polytope $Ax \leq b$.
- Eventually, ball will come to rest against the walls of the polytope.
- Wall $a_ix \leq b_i$ applies some force $-y_ia_i$ to the ball

Interpretation 3: Physical Forces

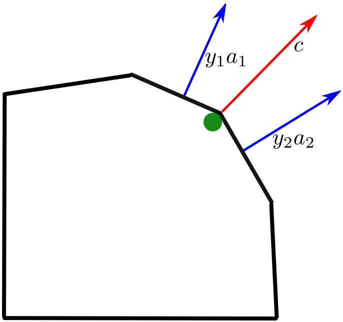

- Apply force field c to a ball inside bounded polytope $Ax \leq b$.
- Eventually, ball will come to rest against the walls of the polytope.
- Wall $a_i x \leq b_i$ applies some force $-y_i a_i$ to the ball
- Since the ball is still, $c^T = \sum_i y_i a_i = y^T A$.

Interpretation 3: Physical Forces

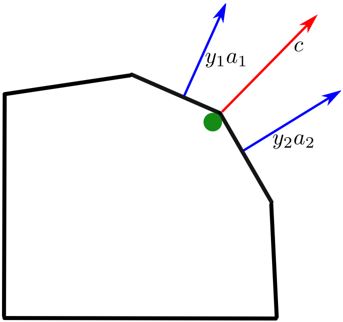

- Apply force field c to a ball inside bounded polytope $Ax \leq b$.
- Eventually, ball will come to rest against the walls of the polytope.
- Wall $a_i x \leq b_i$ applies some force $-y_i a_i$ to the ball
- Since the ball is still, $c^T = \sum_i y_i a_i = y^T A$.
- Dual can be thought of as trying to minimize "work" $\sum_i y_i b_i$ to bring ball back to origin by moving polytope
- We will see that, at optimality, only the walls adjacent to the ball push (Complementary Slackness)

Outline

Duality is an Inversion

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & b^\intercal y \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Duality is an Inversion

Given a primal LP in standard form, the dual of its dual is itself.

Correspondance Between Variables and Constraints

Primal LP

$$
\begin{array}{lll}
\max & \sum_{j=1}^{n} c_j x_j & \\
\text{s.t.} & & \\
& \sum_{j=1}^{n} a_{ij} x_j \leq b_i, & \text{for } i \in [m]. \\
& x_j \geq 0, & \text{for } j \in [n].
\end{array}
$$

Dual LP

$$
\begin{array}{lll}
\min & \sum_{i=1}^{m} b_i y_i & \\
\text{s.t.} & & \\
& \sum_{i=1}^{m} a_{ij} y_i \geq c_j, & \text{for } j \in [n]. \\
& y_i \geq 0, & \text{for } i \in [m].
\end{array}
$$

Correspondance Between Variables and Constraints

Primal LP

$$\begin{array}{llr} \max & \sum_{j=1}^{n} c_j x_j & \\ \text{s.t.} & & \\ y_i: & \sum_{j=1}^{n} a_{ij} x_j \le b_i, & \text{for } i \in [m]. \\ & x_j \ge 0, & \text{for } j \in [n]. \end{array}$$

Dual LP

$$\begin{array}{llr} \min & \sum_{i=1}^{m} b_i y_i & \\ \text{s.t.} & & \\ & \sum_{i=1}^{m} a_{ij} y_i \ge c_j, & \text{for } j \in [n]. \\ & y_i \ge 0, & \text{for } i \in [m]. \end{array}$$

- The i'th primal constraint gives rise to the i'th dual variable y_i

Correspondance Between Variables and Constraints

Primal LP

$$\begin{array}{lll} \text{max} & \sum_{j=1}^{n} c_j x_j & \\ \text{s.t.} & & \\ y_i: & \sum_{j=1}^{n} a_{ij} x_j \leq b_i, & \text{for } i \in [m]. \\ & x_j \geq 0, & \text{for } j \in [n]. \end{array}$$

Dual LP

$$\begin{array}{lll} \text{min} & \sum_{i=1}^{m} b_i y_i & \\ \text{s.t.} & & \\ x_j: & \sum_{i=1}^{m} a_{ij} y_i \geq c_j, & \text{for } j \in [n]. \\ & y_i \geq 0, & \text{for } i \in [m]. \end{array}$$

- The i'th primal constraint gives rise to the i'th dual variable y_i
- The j'th primal variable x_j gives rise to the j'th dual constraint

Syntactic Rules

Primal LP

$$
\begin{array}{lll}
\text{max} & c^\intercal x & \\
\text{s.t.} & & \\
y_i : & a_i x \leq b_i, & \text{for } i \in \mathcal{C}_1. \\
y_i : & a_i x = b_i, & \text{for } i \in \mathcal{C}_2. \\
 & x_j \geq 0, & \text{for } j \in \mathcal{D}_1. \\
 & x_j \in \mathbb{R}, & \text{for } j \in \mathcal{D}_2.
\end{array}
$$

Dual LP

$$
\begin{array}{lll}
\text{min} & b^\intercal y & \\
\text{s.t.} & & \\
x_j : & \overline{a}_j^\intercal y \geq c_j, & \text{for } j \in \mathcal{D}_1. \\
x_j : & \overline{a}_j^\intercal y = c_j, & \text{for } j \in \mathcal{D}_2. \\
 & y_i \geq 0, & \text{for } i \in \mathcal{C}_1. \\
 & y_i \in \mathbb{R}, & \text{for } i \in \mathcal{C}_2.
\end{array}
$$

Rules of Thumb

- Loose constraint (i.e. inequality) $\Rightarrow$ tight dual variable (i.e. nonnegative)
- Tight constraint (i.e. equality) $\Rightarrow$ loose dual variable (i.e. unconstrained)

Outline

Weak Duality

Primal LP

maximize $c^\intercal x$
subject to $Ax \leq b$
$x \geq 0$

Dual LP

minimize $b^\intercal y$
subject to $A^\intercal y \geq c$
$y \geq 0$

Theorem (Weak Duality)

For every primal feasible x and dual feasible y, we have $c^\intercal x \leq b^\intercal y$.

Corollary

- *If primal and dual both feasible and bounded,* $OPT(Primal) \leq OPT(Dual)$
- *If primal is unbounded, dual is infeasible*
- *If dual is unbounded, primal is infeasible*

Weak Duality

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & b^\intercal y \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Theorem (Weak Duality)

For every primal feasible x and dual feasible y, we have $c^\intercal x \leq b^\intercal y$.

Corollary

If x is primal feasible, and y is dual feasible, and $c^\intercal x = b^\intercal y$, then both are optimal.

Interpretation of Weak Duality

Economic Interpretation

If selling the raw materials is more profitable than making any individual product, then total money collected from sale of raw materials would exceed profit from production.

Interpretation of Weak Duality

Economic Interpretation

If selling the raw materials is more profitable than making any individual product, then total money collected from sale of raw materials would exceed profit from production.

Upperbound Interpretation

Self explanatory

Interpretation of Weak Duality

Economic Interpretation

If selling the raw materials is more profitable than making any individual product, then total money collected from sale of raw materials would exceed profit from production.

Upperbound Interpretation

Self explanatory

Physical Interpretation

Work required to bring ball back to origin by pulling polytope is at least potential energy difference between origin and primal optimum.

Proof of Weak Duality

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \le b \\ & x \ge 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & b^\intercal y \\ \text{subject to} \quad & A^\intercal y \ge c \\ & y \ge 0 \end{aligned}$$

$$c^\intercal x \le y^\intercal A x \le y^\intercal b$$

Strong Duality

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & b^\intercal y \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

Theorem (Strong Duality)

If either the primal or dual is feasible and bounded, then so is the other and $OPT(Primal) = OPT(Dual)$*.*

Interpretation of Strong Duality

Economic Interpretation

Buyer can offer prices for raw materials that would make facility indifferent between production and sale.

Interpretation of Strong Duality

Economic Interpretation

Buyer can offer prices for raw materials that would make facility indifferent between production and sale.

Upperbound Interpretation

The method of scaling and summing inequalities yields a tight upperbound on the primal optimal value.

Interpretation of Strong Duality

Economic Interpretation

Buyer can offer prices for raw materials that would make facility indifferent between production and sale.

Upperbound Interpretation

The method of scaling and summing inequalities yields a tight upperbound on the primal optimal value.

Physical Interpretation

There is an assignment of forces to the walls of the polytope that brings ball back to the origin without wasting energy.

Informal Proof of Strong Duality

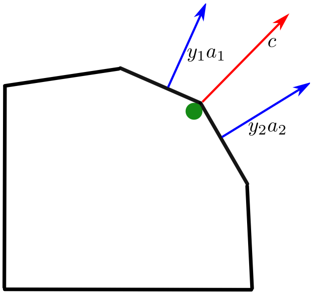

- Recall the physical interpretation of duality

Informal Proof of Strong Duality

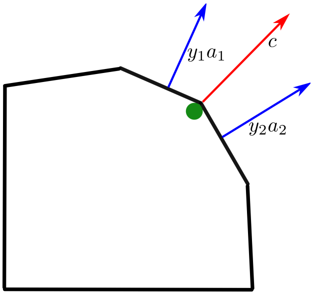

- Recall the physical interpretation of duality
- When ball is stationary at x, we expect force c to be neutralized only by constraints that are tight. i.e. force multipliers $y \geq 0$ s.t.
  - $y^\intercal A = c$
  - $y_i(b_i - a_i x) = 0$

Informal Proof of Strong Duality

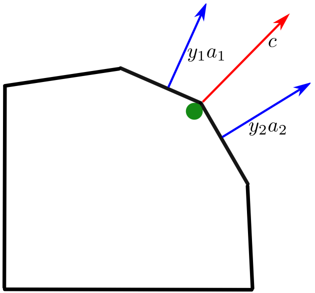

- Recall the physical interpretation of duality
- When ball is stationary at x, we expect force c to be neutralized only by constraints that are tight. i.e. force multipliers $y \geq 0$ s.t.
  - $y^\intercal A = c$
  - $y_i(b_i - a_i x) = 0$

$$y^\intercal b - c^\intercal x = y^\intercal b - y^\intercal A x = \sum_i y_i(b_i - a_i x) = 0$$

We found a primal and dual solution that are equal in value!

Outline

Separating Hyperplane Theorem

If $A, B \subseteq \mathbb{R}^n$ are disjoint convex sets, then there is a hyperplane separating them. That is, there is $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$ such that $a^\intercal x \leq b$ for every $x \in A$ and $a^\intercal y \geq b$ for every $y \in B$. Moreover, if both A and B are closed and at least one of them is compact, then there is a hyperplane strictly separating them (i.e. $a^T x < b$ for $x \in A$ and $a^T y > b$ for $y \in B$).

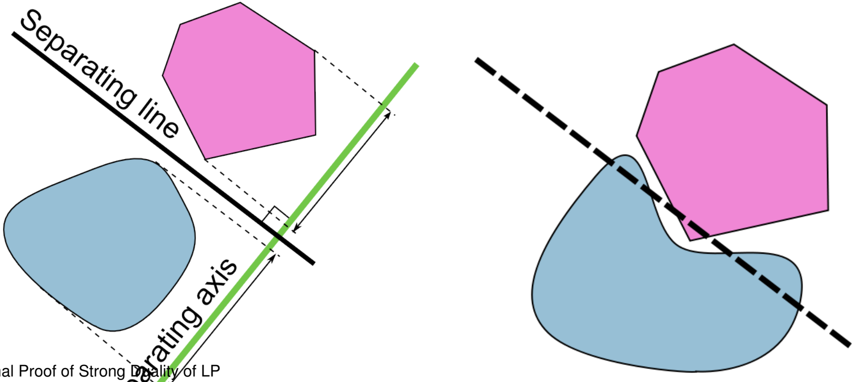

Definition

A convex cone is a convex subset of $\mathbb{R}^n$ which is closed under nonnegative scaling and convex combinations.

Definition

The convex cone generated by vectors $u_1, \dots, u_m \in \mathbb{R}^n$ is the set of all nonnegative-weighted sums of these vectors (also known as conic combinations).

$$Cone(u_1, \dots, u_m) = \left\{ \sum_{i=1}^{m} \alpha_i u_i : \alpha_i \geq 0 \ \forall i \right\}$$

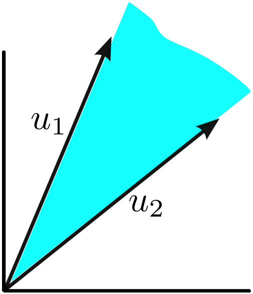

The following follows from the separating hyperplane Theorem (try to prove it).

Farkas' Lemma

Let $\mathcal{C}$ be the convex cone generated by vectors $u_1, \ldots, u_m \in \mathbb{R}^n$, and let $w \in \mathbb{R}^n$. Exactly one of the following is true:

- $w \in \mathcal{C}$
- There is $z \in \mathbb{R}^n$ such that $z \cdot u_i \leq 0$ for all i, and $z \cdot w > 0$.

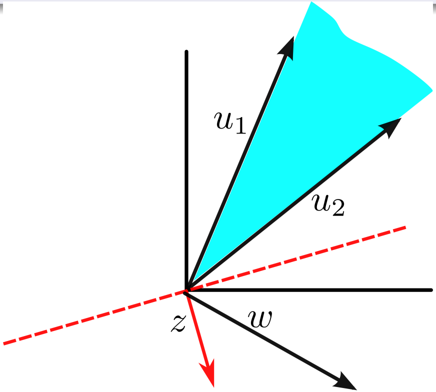

Equivalently: Theorem of the Alternative

Exactly one of the following is true for $U = [u_1, \ldots, u_m]$

- The system $Uy = w$, $y \geq 0$ has a solution
- The system $U^\intercal z \leq 0$, $z^\intercal w > 0$ has a solution.

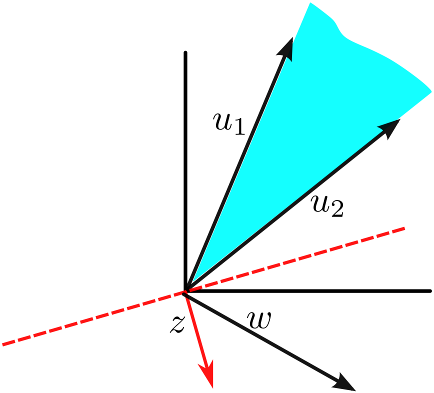

Formal Proof of Strong Duality

Primal LP

maximize $c^\intercal x$
subject to $Ax \leq b$

Dual LP

minimize $b^\intercal y$
subject to $A^\intercal y = c$
$y \geq 0$

Given $v \in \mathbb{R}$, by Farkas' Lemma exactly one of the following is true

1. The system $\begin{pmatrix} A^\intercal & 0 \\ b^\intercal & 1 \end{pmatrix} w = \begin{pmatrix} c \\ v \end{pmatrix}$, $w \geq 0$ has a solution.
   - Let $y \in \mathbb{R}^m_+$ and $\delta \in \mathbb{R}_+$ be such that $w = \begin{pmatrix} y \\ \delta \end{pmatrix}$
   - Implies dual is feasible and $OPT(dual) \leq v$
2. The system $\begin{pmatrix} A & b \\ 0 & 1 \end{pmatrix} z \leq 0$, $z^\intercal \begin{pmatrix} c \\ v \end{pmatrix} > 0$ has a solution.
   - Let $z = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}$, where $z_1 \in \mathbb{R}^n$ and $z_2 \in \mathbb{R}$ with $z_2 \leq 0$
   - When $z_2 \neq 0$, $x = -z_1/z_2$ is feasible and $c^T x > v$
   - When $z_2 = 0$, primal is either infeasible or unbounded, and dual is infeasible (prove it)

Outline

Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \le b \\ & x \ge 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \ge c \\ & y \ge 0 \end{aligned}$$

Complementary Slackness

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \le b \\ & x \ge 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \ge c \\ & y \ge 0 \end{array}$$

- Let $s_i = (b - Ax)_i$ be the i'th primal slack variable
- Let $t_j = (A^\intercal y - c)_j$ be the j'th dual slack variable

Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let $s_i = (b - Ax)_i$ be the i'th primal slack variable
- Let $t_j = (A^\intercal y - c)_j$ be the j'th dual slack variable

Complementary Slackness

x and y are optimal if and only if

- $x_j t_j = 0$ for all $j = 1, \ldots, n$
- $y_i s_i = 0$ for all $i = 1, \ldots, m$

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

Interpretation of Complementary Slackness

Economic Interpretation

Given an optimal primal production vector x and optimal dual offer prices y,

- Facility produces only products for which it is indifferent between sale and production.
- Only raw materials that are binding constraints on production are priced greater than 0

Interpretation of Complementary Slackness

Physical Interpretation

Only walls adjacent to the balls equilibrium position push back on it.

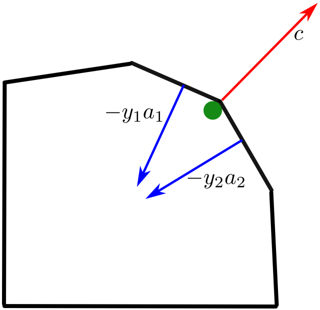

Proof of Complementary Slackness

Primal LP

$$
\begin{array}{ll}
\text{maximize} & c^\intercal x \\
\text{subject to} & Ax \leq b \\
 & x \geq 0
\end{array}
$$

Dual LP

$$
\begin{array}{ll}
\text{minimize} & y^\intercal b \\
\text{subject to} & A^\intercal y \geq c \\
 & y \geq 0
\end{array}
$$

Proof of Complementary Slackness

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax + s = b \\ & x \geq 0 \\ & s \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y - t = c \\ & y \geq 0 \\ & t \geq 0 \end{array}$$

- Can equivalently rewrite LP using slack variables

Proof of Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax + s = b \\ & x \geq 0 \\ & s \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y - t = c \\ & y \geq 0 \\ & t \geq 0 \end{aligned}$$

- Can equivalently rewrite LP using slack variables

$$y^\intercal b - c^\intercal x = y^\intercal (Ax + s) - (y^\intercal A - t^\intercal)x = y^\intercal s + t^\intercal x$$

Proof of Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax + s = b \\ & x \geq 0 \\ & s \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y - t = c \\ & y \geq 0 \\ & t \geq 0 \end{aligned}$$

- Can equivalently rewrite LP using slack variables

$$y^\intercal b - c^\intercal x = y^\intercal (Ax + s) - (y^\intercal A - t^\intercal) x = y^\intercal s + t^\intercal x$$

Gap between primal and dual objectives is 0 if and only if complementary slackness holds.

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: At every vertex of primal [dual] there are exactly n [m] tight constraints which are linearly independent.

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: At every vertex of primal [dual] there are exactly n [m] tight constraints which are linearly independent.

Primal LP
(n variables, $m+n$ constraints)

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP
(m variables, $m+n$ constraints)

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: At every vertex of primal [dual] there are exactly n [m] tight constraints which are linearly independent.

Primal LP
(n variables, $m + n$ constraints)

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP
(m variables, $m + n$ constraints)

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let y be dual optimal. By non-degeneracy:
  - Exactly m of the $m + n$ dual constraints are tight at y
  - Exactly n dual constraints are loose

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: At every vertex of primal [dual] there are exactly n [m] tight constraints which are linearly independent.

Primal LP
(n variables, $m+n$ constraints)

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP
(m variables, $m+n$ constraints)

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let y be dual optimal. By non-degeneracy:
  - Exactly m of the $m+n$ dual constraints are tight at y
  - Exactly n dual constraints are loose
- n loose dual constraints impose n tight primal constraints

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: At every vertex of primal [dual] there are exactly n [m] tight constraints which are linearly independent.

Primal LP
(n variables, $m + n$ constraints)

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP
(m variables, $m + n$ constraints)

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let y be dual optimal. By non-degeneracy:
  - Exactly m of the $m + n$ dual constraints are tight at y
  - Exactly n dual constraints are loose
- n loose dual constraints impose n tight primal constraints
  - Assuming non-degeneracy, solving the linear equation yields a unique primal optimum solution x.

Sensitivity Analysis

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Sometimes, we want to examine how the optimal value of our LP changes with its parameters c and b

Sensitivity Analysis

Primal LP

maximize $c^\intercal x$
subject to $Ax \leq b$
$x \geq 0$

Dual LP

minimize $y^\intercal b$
subject to $A^\intercal y \geq c$
$y \geq 0$

Sometimes, we want to examine how the optimal value of our LP changes with its parameters c and b

Sensitivity Analysis

Let $OPT = OPT(A, c, b)$ be the optimal value of the above LP. Let x and y be the primal and dual optima.

- $\frac{\partial OPT}{\partial c_j} = x_j$ when x is the unique primal optimum.
- $\frac{\partial OPT}{\partial b_i} = y_i$ when y is the unique dual optimum.

Sensitivity Analysis

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

Sometimes, we want to examine how the optimal value of our LP changes with its parameters c and b

Economic Interpretation of Sensitivity Analysis

- A small increase δ in c_j increases profit by $\delta \cdot x_j$
- A small increase δ in b_i increases profit by $\delta \cdot y_i$
  - y_i measures the "marginal value" of resource i for production

Outline

Shortest Path

Given a directed network $G = (V, E)$ where edge e has length $\ell_e \in \mathbb{R}_+$, find the minimum cost path from s to t.

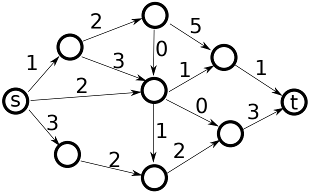

Shortest Path

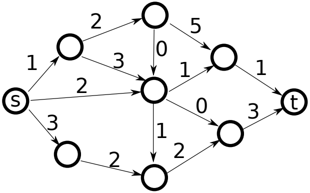

Primal LP

$$\begin{array}{lll} \text{min} & \sum_{e \in E} \ell_e x_e & \\ \text{s.t.} & \sum\limits_{e \to v} x_e - \sum\limits_{v \to e} x_e = \delta_v, & \forall v \in V. \\ & x_e \geq 0, & \forall e \in E. \end{array}$$

Dual LP

$$\begin{array}{lll} \text{max} & y_t - y_s & \\ \text{s.t.} & y_v - y_u \leq \ell_e, & \forall (u,v) \in E. \end{array}$$

Where $\delta_v = -1$ if $v = s$, 1 if $v = t$, and 0 otherwise.

Shortest Path

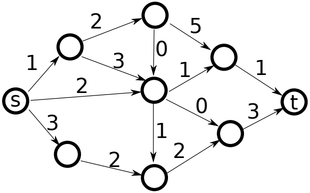

Primal LP

$$\begin{array}{lll} \text{min} & \sum_{e \in E} \ell_e x_e & \\ \text{s.t.} & \sum\limits_{e \to v} x_e - \sum\limits_{v \to e} x_e = \delta_v, & \forall v \in V. \\ & x_e \geq 0, & \forall e \in E. \end{array}$$

Dual LP

$$\begin{array}{lll} \text{max} & y_t - y_s & \\ \text{s.t.} & y_v - y_u \leq \ell_e, & \forall (u,v) \in E. \end{array}$$

Where $\delta_v = -1$ if $v = s$, 1 if $v = t$, and 0 otherwise.

Interpretation of Dual

Stretch s and t as far apart as possible, subject to edge lengths.

Maximum Weighted Bipartite Matching

Set B of buyers, and set G of goods. Buyer i has value w_{ij} for good j, and interested in at most one good. Find maximum value assignment of goods to buyers.

Maximum Weighted Bipartite Matching

Primal LP

$$
\begin{aligned}
\text{max} \quad & \sum_{i,j} w_{ij} x_{ij} \\
\text{s.t.} \quad & \sum_{j \in G} x_{ij} \leq 1, && \forall i \in B. \\
& \sum_{i \in B} x_{ij} \leq 1, && \forall j \in G. \\
& x_{ij} \geq 0, && \forall i \in B, j \in
\end{aligned}
$$

Dual LP

$$
\begin{aligned}
\text{min} \quad & \sum_{i \in B} u_i + \sum_{j \in G} p_j \\
\text{s.t.} \quad & u_i + p_j \geq w_{ij}, && \forall i \in B, j \in G. \\
& u_i \geq 0, && \forall i \in B. \\
& p_j \geq 0, && \forall j \in G.
\end{aligned}
$$

Maximum Weighted Bipartite Matching

Primal LP

$$\begin{aligned}
\max \quad & \sum_{i,j} w_{ij} x_{ij} \\
\text{s.t.} \quad & \sum_{j \in G} x_{ij} \le 1, && \forall i \in B. \\
& \sum_{i \in B} x_{ij} \le 1, && \forall j \in G. \\
& x_{ij} \ge 0, && \forall i \in B, j \in
\end{aligned}$$

Dual LP

$$\begin{aligned}
\min \quad & \sum_{i \in B} u_i + \sum_{j \in G} p_j \\
\text{s.t.} \quad & u_i + p_j \ge w_{ij}, && \forall i \in B, j \in G. \\
& u_i \ge 0, && \forall i \in B. \\
& p_j \ge 0, && \forall j \in G.
\end{aligned}$$

Interpretation of Dual

- p_j is price of good j
- u_i is utility of buyer i
- Complementary Slackness:
  - A buyer i only grabs goods j maximizing $w_{ij} - p_j$
  - Only fully assigned goods have non-zero price
  - A buyer witn nonzero utility must receive an item

2-Player Zero-Sum Games

Rock-Paper-Scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- Two players, row and column
- Game described by matrix A
- When row player plays pure strategy i and column player plays pure strategy j, row player pays column player A_{ij}

2-Player Zero-Sum Games

Rock-Paper-Scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- Two players, row and column
- Game described by matrix A
- When row player plays pure strategy i and column player plays pure strategy j, row player pays column player A_{ij}
- Mixed Strategy: distribution over pure strategies

2-Player Zero-Sum Games

Rock-Paper-Scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- Two players, row and column
- Game described by matrix A
- When row player plays pure strategy i and column player plays pure strategy j, row player pays column player A_{ij}
- Mixed Strategy: distribution over pure strategies
- Assume players know each other's mixed strategies but not coin flips

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$

	x_1	x_2	x_3	x_4
y_1	a_{11}	a_{12}	a_{13}	a_{14}
y_2	a_{21}	a_{22}	a_{23}	a_{24}
y_3	a_{31}	a_{32}	a_{33}	a_{34}

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$

Row Moves First

min $\quad \max_j (y^\intercal A)_j$
s.t. $\sum_{i=1}^m y_i = 1$
$y \geq \vec{0}$

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$

Row Moves First

$$\begin{aligned} \text{min} \quad & u \\ \text{s.t.} \quad & u\vec{1} - y^\intercal A \geq \vec{0} \\ & \textstyle\sum_{i=1}^m y_i = 1 \\ & y \geq \vec{0} \end{aligned}$$

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$
  - Similarly when column moves first

Row Moves First

$$\begin{aligned} \text{min} \quad & u \\ \text{s.t.} \quad & u\vec{1} - y^\intercal A \geq \vec{0} \\ & \textstyle\sum_{i=1}^{m} y_i = 1 \\ & y \geq \vec{0} \end{aligned}$$

Column Moves First

$$\begin{aligned} \text{max} \quad & v \\ \text{s.t.} \quad & v\vec{1} - Ax \leq \vec{0} \\ & \textstyle\sum_{j=1}^{n} x_j = 1 \\ & x \geq \vec{0} \end{aligned}$$

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$
  - Similarly when column moves first

Row Moves First

$$\begin{aligned} \min \quad & u \\ \text{s.t.} \quad & u\vec{1} - y^\intercal A \geq \vec{0} \\ & \textstyle\sum_{i=1}^m y_i = 1 \\ & y \geq \vec{0} \end{aligned}$$

Column Moves First

$$\begin{aligned} \max \quad & v \\ \text{s.t.} \quad & v\vec{1} - Ax \leq \vec{0} \\ & \textstyle\sum_{j=1}^n x_j = 1 \\ & x \geq \vec{0} \end{aligned}$$

These two optimization problems are LP Duals!

Duality and Zero Sum Games

Weak Duality

- $u^* \geq v^*$
- Zero sum games have a second mover advantage (weakly)

Duality and Zero Sum Games

Weak Duality

- $u^* \geq v^*$
- Zero sum games have a second mover advantage (weakly)

Strong Duality (Minimax Theorem)

- $u^* = v^*$
- There is no second or first mover advantage in zero sum games with mixed strategies
- Each player can guarantee $u^* = v^*$ regardless of other's strategy.
- y^*, x^* are simultaneously best responses to each other (Nash Equilibrium)

Duality and Zero Sum Games

Weak Duality

- $u^* \geq v^*$
- Zero sum games have a second mover advantage (weakly)

Strong Duality (Minimax Theorem)

- $u^* = v^*$
- There is no second or first mover advantage in zero sum games with mixed strategies
- Each player can guarantee $u^* = v^*$ regardless of other's strategy.
- y^*, x^* are simultaneously best responses to each other (Nash Equilibrium)

Complementary Slackness

x^* randomizes over pure best responses to y^*, and vice versa.

Saddle Point Interpretation

Consider the matching pennies game

	H	T
H	-1	1
T	1	-1

- Unique equilibrium: each player randomizes uniformly
- If row player deviates, he pays out more
- If column player deviates, he gets paid less

Saddle Point Interpretation

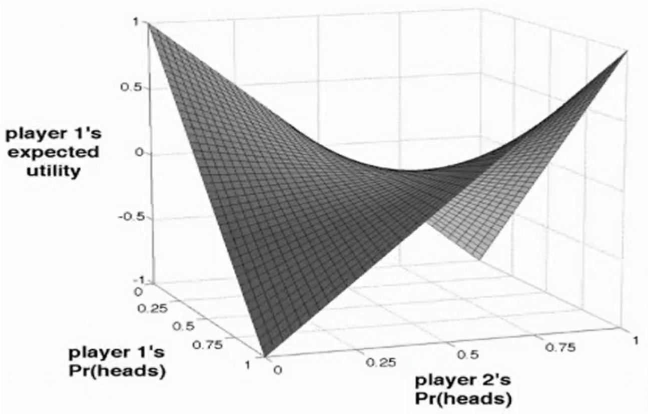

- Unique equilibrium: each player randomizes uniformly
- If row player deviates, he pays out more
- If column player deviates, he gets paid less