MATH 4211/6211 – Optimization

Duality

Xiaojing Ye
Department of Mathematics & Statistics
Georgia State University

Every LP has a corresponding **dual** LP.

Consider the primal LP given by

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

Then the dual LP is given by

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{b}^\top \boldsymbol{\lambda} \\ \text{subject to} \quad & \boldsymbol{A}^\top \boldsymbol{\lambda} \leq \boldsymbol{c} \\ & \boldsymbol{\lambda} \geq \boldsymbol{0} \end{aligned}$$

where $\lambda \in \mathbb{R}^m$.

Note the exchanges $\boldsymbol{c} \leftrightarrow \boldsymbol{b}$, $\boldsymbol{x} \leftrightarrow \boldsymbol{\lambda}$, $\boldsymbol{A} \leftrightarrow \boldsymbol{A}^\top$, and $\geq \leftrightarrow \leq$.

The dual of the dual LP is the primal LP: first we write the dual LP as

$$\begin{aligned} \text{minimize} \quad & (-\boldsymbol{b})^\top \boldsymbol{\lambda} \\ \text{subject to} \quad & -\boldsymbol{A}^\top \boldsymbol{\lambda} \geq -\boldsymbol{c} \\ & \boldsymbol{\lambda} \geq \boldsymbol{0} \end{aligned}$$

Then the dual of the dual LP above is

$$\begin{aligned} \text{maximize} \quad & (-\boldsymbol{c})^\top \boldsymbol{x} \\ \text{subject to} \quad & -\boldsymbol{A}\boldsymbol{x} \leq -\boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

which is equivalent to the primal LP

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

We can also consider the standard form of LP and its dual.

The standard form of LP is

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

where the equality constraint $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ can be written as two inequality constraints: $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}$ (or equivalently $-\boldsymbol{A}\boldsymbol{x} \geq -\boldsymbol{b}$) and $\boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}$, i.e.,

$$\begin{bmatrix} \boldsymbol{A} \\ -\boldsymbol{A} \end{bmatrix} \boldsymbol{x} \geq \begin{bmatrix} \boldsymbol{b} \\ -\boldsymbol{b} \end{bmatrix}$$

Therefore the dual LP is

$$\begin{aligned} \text{maximize} \quad & [\boldsymbol{b}^\top, -\boldsymbol{b}^\top] \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix} \\ \text{subject to} \quad & [\boldsymbol{A}^\top, -\boldsymbol{A}^\top] \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix} \le \boldsymbol{c} \\ & \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix} \ge \boldsymbol{0} \end{aligned}$$

where $\begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix}$ is the variable.

Letting $\boldsymbol{\lambda} = \boldsymbol{u} - \boldsymbol{v}$, we obtain

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{b}^\top \boldsymbol{\lambda} \\ \text{subject to} \quad & \boldsymbol{A}^\top \boldsymbol{\lambda} \le \boldsymbol{c} \end{aligned}$$

Note that we do not have $\boldsymbol{\lambda} \ge \boldsymbol{0}$ anymore.

Symmetric form of duality

Primal

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

Dual

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{b}^\top \boldsymbol{\lambda} \\ \text{subject to} \quad & \boldsymbol{A}^\top \boldsymbol{\lambda} \leq \boldsymbol{c} \\ & \boldsymbol{\lambda} \geq \boldsymbol{0} \end{aligned}$$

Asymmetric form of duality

Primal

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

Dual

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{b}^\top \boldsymbol{\lambda} \\ \text{subject to} \quad & \boldsymbol{A}^\top \boldsymbol{\lambda} \leq \boldsymbol{c} \end{aligned}$$

Lemma (Weak Duality Lemma). Suppose $\boldsymbol{x}$ and $\boldsymbol{\lambda}$ are feasible solutions to the primal and dual LP problems, respectively, then

$$\boldsymbol{c}^\top \boldsymbol{x} \ge \boldsymbol{b}^\top \boldsymbol{\lambda}$$

That is, the primal objective value $\ge$ dual objective value.

Proof. We prove the asymmetric form only. Since $\boldsymbol{x}$ and $\boldsymbol{\lambda}$ are both feasible in their corresponding problems, we know $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ and $\boldsymbol{x} \ge \boldsymbol{0}$, as well as $\boldsymbol{A}^\top \boldsymbol{\lambda} \le \boldsymbol{c}$. Hence

$$\boldsymbol{b}^\top \boldsymbol{\lambda} = (\boldsymbol{A}\boldsymbol{x})^\top \boldsymbol{\lambda} = \boldsymbol{x}^\top \boldsymbol{A}^\top \boldsymbol{\lambda} \le \boldsymbol{x}^\top \boldsymbol{c}$$

Theorem. Suppose $\boldsymbol{x}_0$ and $\boldsymbol{\lambda}_0$ are feasible points to the primal and dual LP problems, respectively. If $\boldsymbol{c}^\top \boldsymbol{x}_0 = \boldsymbol{b}^\top \boldsymbol{\lambda}_0$, then $\boldsymbol{x}_0$ and $\boldsymbol{\lambda}_0$ are optimal solutions to their respective problems.

Proof. By the weak duality lemma above, we know for every primal feasible point $\boldsymbol{x}$ there is

$$\boldsymbol{c}^\top \boldsymbol{x} \geq \boldsymbol{b}^\top \boldsymbol{\lambda}_0 = \boldsymbol{c}^\top \boldsymbol{x}_0$$

Therefore $\boldsymbol{x}_0$ is optimal. Similarly we can show the optimality of $\boldsymbol{\lambda}_0$.

Theorem. If the primal problem has an optimal solution, then so does the dual, and the optimal primal objective value is equal to the optimal dual objective value.

Proof. Consider the asymmetric form first. If the primal problem has an optimal solution, then by the fundamental theorem of LP, it has an optimal basic feasible solution. That is $\boldsymbol{A} = [\boldsymbol{B}, \boldsymbol{D}]$ (WLOG we assume the first m columns of $\boldsymbol{A}$ are basic columns) and $\boldsymbol{x} = [\boldsymbol{x}_B; \boldsymbol{0}]$ is optimal.

As $\boldsymbol{x}$ is optimal, we set $\boldsymbol{\lambda}$ such that $\boldsymbol{\lambda}^\top = \boldsymbol{c}_B^\top \boldsymbol{B}^{-1}$, and have

$$\boldsymbol{0} \le \boldsymbol{r}_D^\top = \boldsymbol{c}_D^\top - (\boldsymbol{c}_B^\top \boldsymbol{B}^{-1})\boldsymbol{D} = \boldsymbol{c}_D^\top - \boldsymbol{\lambda}^\top \boldsymbol{D}$$

i.e., $\boldsymbol{\lambda}^\top \boldsymbol{D} \le \boldsymbol{c}_D^\top$. So $\boldsymbol{\lambda}$ is feasible and optimal:

$$\begin{aligned}\boldsymbol{\lambda}^\top \boldsymbol{A} = \boldsymbol{\lambda}^\top [\boldsymbol{B}, \boldsymbol{D}] = \boldsymbol{c}_B^\top \boldsymbol{B}^{-1} [\boldsymbol{B}, \boldsymbol{D}] \le [\boldsymbol{c}_B^\top, \boldsymbol{c}_D^\top] = \boldsymbol{c}^\top \\ \boldsymbol{\lambda}^\top \boldsymbol{b} = \boldsymbol{c}_B^\top \boldsymbol{B}^{-1} \boldsymbol{b} = \boldsymbol{c}_B^\top \boldsymbol{x}_B = \boldsymbol{c}^\top \boldsymbol{x}.\end{aligned}$$

For symmetric form, convert it to the asymmetric form and apply the proof.

Theorem (Complementary Slackness Condition). The feasible solution $\boldsymbol{x}$ and $\boldsymbol{\lambda}$ are optimal to a dual pair of LPs iff $(\boldsymbol{c} - \boldsymbol{A}^\top \boldsymbol{\lambda})^\top \boldsymbol{x} = 0$ and $\boldsymbol{\lambda}^\top (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) = 0$.

Proof. We consider the asymmetric form.

($\Rightarrow$) Since $\boldsymbol{x}$ is feasible, there is $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ and hence $\boldsymbol{\lambda}^\top (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) = 0$. In addition, since both $\boldsymbol{x}$ and $\boldsymbol{\lambda}$ are optimal, we have $\boldsymbol{c}^\top \boldsymbol{x} = \boldsymbol{b}^\top \boldsymbol{\lambda} = (\boldsymbol{A}\boldsymbol{x})^\top \boldsymbol{\lambda}$, which implies the first equality.

($\Leftarrow$) The two equalities imply $\boldsymbol{c}^\top \boldsymbol{x} = \boldsymbol{\lambda}^\top \boldsymbol{A}\boldsymbol{x} = \boldsymbol{\lambda}^\top \boldsymbol{b}$. Hence $\boldsymbol{x}$ and $\boldsymbol{\lambda}$ are optimal.