
Algorithmic Game Theory and Applications

Lecture 7: The LP Duality Theorem

Kousha Etessami



recall LP's in "Primal Form"

Maximize $c_1 x_1 + c_2 x_2 + \dots + c_n x_n$

Subject to:

$$a_{1,1} x_1 + a_{1,2} x_2 + \dots + a_{1,n} x_n \leq b_1$$

$$a_{2,1} x_1 + a_{2,2} x_2 + \dots + a_{2,n} x_n \leq b_2$$

...

...

$$a_{m,1} x_1 + a_{i,2} x_2 + \dots + a_{m,n} x_n \leq b_m$$

$$x_1, \dots, x_n \geq 0$$

An LP can concisely be represented in matrix notation:

Define the $(m \times n)$ -matrix A by: $(A)_{i,j} = a_{i,j}$.

Define (column) vectors $x = [x_1, \dots, x_n]^T$,

$b = [b_1, \dots, b_m]^T$ & $c = [c_1, \dots, c_n]^T$.

The LP is:

Maximize $c^T x$

Subject to:

$$Ax \leq b$$

$$x_1, \dots, x_n \geq 0$$

Solving LPs against an Adversary

Suppose you want to optimize this Primal LP:

Maximize $c^T x$

Subject to:

$$Ax \leq b$$

$$x_1, \dots, x_n \geq 0$$

Suppose an “Adversary” comes along with a m -vector $y \in \mathbb{R}^m$, $y \geq 0$, such that:

$$c^T \leq y^T A$$

For any solution, x , you may have for the Primal:

$$c^T x \leq (y^T A)x \quad (\text{because } x \geq 0 \ \& \ c^T \leq y^T A)$$

$$= y^T (Ax)$$

$$\leq y^T b \quad (\text{because } y^T \geq 0 \ \& \ Ax \leq b)$$

So, the adversary’s goal is to minimize $y^T b$, subject to $c^T \leq y^T A$, i.e., to optimize the following Dual LP:

Minimize $b^T y$

Subject to:

$$A^T y \geq c$$

$$y_1, \dots, y_m \geq 0$$

The LP Duality Theorem

As we saw, a feasible value of the primal LP is at most the optimal value of the dual, i.e.:

Proposition (“Weak Duality”) If $x^* \in R^n$ and $y^* \in R^m$ are optimal feasible solutions to the primal and dual LPs, then:

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

Amazingly, equality holds!

Theorem (“Strong Duality”, von Neumann’47) One of the following four situations holds:

1. Both the primal and dual LPs are feasible, and for any optimal solutions x^* of the primal and y^* of the dual:

$$c^T x^* = b^T y^*$$

2. The primal is infeasible and the dual is unbounded.
3. The primal is unbounded and the dual is infeasible.
4. Both LPs are infeasible.

Simplex and the Duality Theorem

- It turns out, LP Duality follows from the “proof of correctness” of the Simplex algorithm (which we didn’t provide).
- In fact, suppose Simplex on the Primal LP halts and produces an optimal solution vector x^* .

Let $y_j^* = -\hat{c}_{n+j}$, where \hat{c}_{n+j} is the coefficient of the “slack variable” x_{n+j} in the last objective function $f(x)$ associated with the final dictionary for the primal LP. This defines $y^* \in \mathbb{R}^m$, and it turns out the following holds:

Fact $y^* \geq 0$, $A^T y^* \geq c$, and $c^T x^* = b^T y^*$.

- So, using Simplex, we can actually compute an optimal solution for the dual while we compute an optimal solution for the primal (or find out neither exists).

Complementary Slackness

Useful Corollary of LP Duality: solutions x^* and y^* to the primal and dual LPs, respectively, are both optimal if and only if both of the following hold:

- For each primal constraint, $(Ax)_i \leq b_i$, $i = 1, \dots, m$, either

$$(Ax^*)_i = b_i \quad \text{or} \quad y_i^* = 0 \quad (\text{or both}).$$
- For each dual constraint, $(A^T y)_j \geq c_j$, $j = 1, \dots, n$, either

$$(A^T y^*)_j = c_j \quad \text{or} \quad x_j^* = 0 \quad (\text{or both}).$$

Proof By weak duality

$$c^T x^* \leq ((y^*)^T A)x^* = (y^*)^T (Ax^*) \leq (y^*)^T b$$

By strong duality each inequality holds with equality precisely when x^* and y^* are optimal.

So, when optimal, $((y^*)^T A) - c^T)x^* = 0$.

Since both $((y^*)^T A) - c^T \geq 0$ and $x^* \geq 0$, it must be that for each $j = 1, \dots, n$,

$$(A^T y^*)_j = c_j \quad \text{or} \quad x_j^* = 0$$

Likewise, when optimal, $(y^*)^T (b - Ax^*) = 0$, and thus, for each $i = 1, \dots, m$, $(Ax^*)_i = b_i$ or $y_i^* = 0$.

■

general recipe for LP duals

If the “primal” is:

Maximize $c^T x$

Subject to:

$$(Ax)_i \leq b_i, \quad i = 1, \dots, d,$$

$$(Ax)_i = b_i, \quad i = d + 1, \dots, m,$$

$$x_1, \dots, x_r \geq 0$$

Then the “dual” is:

Minimize $b^T y$

Subject to:

$$(A^T y)_j \geq c_j, \quad j = 1, \dots, r,$$

$$(A^T y)_j = c_j, \quad j = r + 1, \dots, n,$$

$$y_1, \dots, y_d \geq 0$$

Strong Duality holds also in this more general setting.

Fact: The “dual” of the “dual” is the primal.

Duality and the Minimax Theorem

Recall the LP for solving Minimax for Player 1 in a zero-sum game given by matrix A :

Maximize v

Subject to:

$$v - (x^T A)_j \leq 0 \text{ for } j = 1, \dots, m_2,$$

$$x_1 + \dots + x_{m_1} = 1,$$

$$x_i \geq 0 \text{ for } i = 1, \dots, m_1.$$

What is the Dual to this LP? It can be shown to be:

Minimize u

Subject to constraints:

$$u - (Ay)_i \geq 0 \text{ for } i = 1, \dots, m_1,$$

$$y_1 + \dots + y_{m_2} = 1$$

$$y_j \geq 0 \text{ for } j = 1, \dots, m_2.$$

This is exactly the LP for Player 2's optimal strategy!

Thus, the minimax theorem follows from Duality: the optimal value for the two LPs is the same.

From Minimax to LP Duality

In fact, LP Duality can also be “derived” from the Minimax Theorem. Consider (A a $(m \times n)$ -matrix):

Maximize $c^T x$

Subject to:

$$Ax \leq b$$

$$x_1, \dots, x_n \geq 0$$

Proof (sketch) that Minimax \Rightarrow LP-duality:
Consider the zero-sum payoff matrix:

$$B = \begin{bmatrix} 0 & A & -b \\ -A^T & 0 & c \\ b^T & -c^T & 0 \end{bmatrix}$$

This is a symmetric zero-sum game: $B = -B^T$. By your homework (Hw1, problem 3(a)), the game has minimax value 0, and every minmaximizer for player 1 is also a maximizer for player 2. Consider a symmetric minimax profile

$$((Y^*, X^*, z), (Y^*, X^*, z))$$

for this game. Here Y^* is a (row) vector of length m , and X^* is a (row) vector of length n , and $z \in \mathbb{R}$.

more in Minimax \Rightarrow LP-duality

Suppose $z > 0$, in the minimax strategy (Y^*, X^*, z) .

Note that, by the Minimax Theorem, we have

$$AX^* - bz \leq 0 \quad \text{and} \quad cz - A^T Y^* \leq 0$$

Letting $y^* = (1/z)Y^*$, and $x^* = (1/z)X^*$, we would have $Ax^* \leq b$ and $A^T y^* \geq c$, i.e., feasible solutions for the LP and its Dual.

Moreover, by Useful Corollary to Minimax, player 1 can switch to ANY pure strategy j in its “support” (i.e., where $x_1(j) > 0$) and not change its profit. Let it switch to the last row (i.e., letting $z=1$). Then we also have: $bY^* - cX^* = 0$, and hence $by^* - cx^* = 0$. Thus,

$$b^T y^* = c^T x^*$$

The only thing left to prove is:

Claim If both the LP and its Dual are feasible, the game B has some minimax profile where $z > 0$.

remarks

- The claim can be proved using facts related to the “geometry” of LP and Minimax: the “Farkas lemmas”.
- Such “Farkas Lemmas” can actually be proved very nicely using Fourier-Motzkin elimination. Here is a Farkas lemma:
 $Ax \leq b$ has no solutions if and only if there exists $y \geq 0$ such that $y^T A = 0$ and $y^T b < 0$.
(Extra credit on HW1, problem 6: prove this.)
- This proof is somewhat unsatisfactory because the Farkas lemmas are essentially “equivalent” to LP-duality. Unfortunately, every (correct) proof I’ve seen that Minimax \Rightarrow LP-duality uses them in some form.
- Note: We can view this proof of LP-duality as a “polynomial time reduction” from the problem of computing an optimal feasible solution to an LP, to that of computing a minimax solution of a zero sum game with a particular pure strategy in its support.

Significance of LP Duality

- The Duality Theorem is an extremely important fact in subjects like mathematical economics and combinatorial optimization.
So much so that we can not possibly do it justice.

- Often, you can “learn something” about an LP by looking at its dual.

- In Economics, one often sets up an LP to optimize some economic goal. Often, the dual LP can be meaningfully interpreted as a problem of optimizing some real economic “counter” goal.

The fact that the optimal solutions of the two goals are the same is very powerful.

- Duality has many consequences in algorithms.

For example, duality implies that LPs can be solved in “ $NP \cap co-NP$ ”.

(Of course, we know from [Khachian’79] that the LP problem is in “P”.)

more on algorithmic significance

- approximation algorithms: hard combinatorial optimization problems can often be formulated as “Integer Linear Programs” (ILPs). One can often use the “LP-relaxation” of the ILP together with its Dual to find approximate solutions and to bound the proximity of the approximate solution to the optimal. (See, e.g., [Vazirani’2001].)
- Many important results in combinatorics can be viewed as particular manifestations of LP Duality.
 - Max-Flow Min-Cut Theorem;
 - Hall’s Theorem;
 - Dilworth’s Theorem;
 - Konig-Egervary Theorem.....

Each result says “the maximum value of one useful quantity associated with a combinatorial object is the same as the minimum value of another useful quantity associated with it.”

These follow from LP-Duality: dual LPs for these “complementary” quantities can be set up (with solutions that consist necessarily of integers, due to the LPs’ structure).

food for thought (in this case “thought for food”)

- Recall the “Diet Problem”. It has the form:

Minimize $c^T x$

Subject to:

$$Ax \geq b$$

$$x_1, \dots, x_n \geq 0$$

- Construct the dual to this LP.
- What do the dual variables “mean” in the context of the diet problem?

Try to come up with an interpretation.

(Hint: try to assign consistent “units of measure” to the primal variables, constants, and coefficients. These will determine the “units of measure” of the dual variables, and will guide you toward an interpretation.)

- What is the dual LP trying to optimize?