

Duality

This lecture covers weak and strong duality, and also explains the rules for finding the dual of a linear program, with an example. Before we move on to duality, we shall first see some general facts about the location of the optima of a linear program.

1.1 Structure of LP solutions

1.1.1 Some intuition in two dimensions

Consider a linear program -

Maximize $y^T b$
subject to $y^T A \leq c$

The feasible region of this LP is in general, a convex polyhedron. Visualize it as a polygon in 2 dimensions, for simplicity. Now, maximizing $y^T b$ is the same as maximizing the projection of the vector y in the direction represented by vector b . For whichever direction b we choose, the point y that maximizes $y^T b$ cannot lie strictly in the interior of the feasible region. The reason is that, from an interior point, we can move further in any direction, and still be feasible. In particular, by moving along b , we can get to a point with a larger projection along b . This intuition suggests that the optimal solution of an LP will never lie in the interior of the feasible region, but only on the boundaries. In fact, we can say more. We can show that for any LP, the optimal solutions are always at the “corners” of the feasible region polyhedron. This notion is formalized in the next subsection.

1.1.2 Some definitions

Definition 1 (Vertex of a Polyhedron) *A point in the polyhedron which is uniquely optimal for some linear objective, is called a vertex of the polyhedron.*

Definition 2 (Extreme Point of a Polyhedron) *A point in the polyhedron which is not*

a convex combination of two other points in the polyhedron is called an extreme point of the polyhedron.

Definition 3 (Tightness) A constraint of the form $a^T x \leq b$, $a^T x = b$ or $a^T x \geq b$ in a linear program is said to be tight for a certain point y , if $a^T y = b$.

Definition 4 (Basic Solution) For an n -dimensional linear program, a point is called a basic solution, if n linearly independent constraints are tight for that point.

Definition 5 (Basic Feasible Solution) A point is a basic feasible solution, iff it is a basic solution that is also feasible.

Note: If x is a basic feasible solution, then it is in fact, the unique point that is tight for all its tight constraints. This is because, there can be only one solution for a set of n linearly independent equalities, in n -dimensional space.

Theorem 1 For a polyhedron P and a point $x \in P$, the following are equivalent:

1. x is a basic feasible solution
2. x is a vertex of P
3. x is an extreme point of P

Proof: Assume the LP is in the canonical form.

1. Vertex \Rightarrow Extreme Point

Let v be a vertex. Then for some objective function c , $c^T x$ is uniquely minimized at v . Assume v is not an extreme point. Then, v can be written as $v = \lambda y + (1 - \lambda)z$ for some y, z neither of which is v , and some λ satisfying $0 \leq \lambda \leq 1$.

$$\text{Now, } c^T v = c^T [\lambda y + (1 - \lambda)z] = \lambda c^T y + (1 - \lambda)c^T z$$

This means $c^T y \leq c^T v \leq c^T z$. But, since v is a minimum point, $c^T v \leq c^T y$ and $c^T v \leq c^T z$. Thus, $c^T y = c^T v = c^T z$. This is a contradiction, since v is the unique point at which $c^T x$ is minimized.

2. Extreme Point \Rightarrow Basic Feasible Solution

Let x be an extreme point. By definition, it lies in the polyhedron and is therefore feasible. Assume x is not a basic solution. Let T be the set of rows of the constraint matrix A for which the constraints are tight at x . Let a_i (a $1 \times n$ vector) denote the

i^{th} row of A . For $a_i \notin T$, $a_i \cdot x > b_i$. Since x is not a basic solution, T does not span \mathcal{R}^n . So, there is a vector $d \neq 0$ such that $a_i \cdot d = 0 \forall a_i \in T$.

Consider $y = x + \epsilon d$ and $z = x - \epsilon d$. If $a_i \in T$, then $a_i \cdot y = a_i \cdot z = b_i$. If $a_i \notin T$, then, by choosing a sufficiently small ϵ : $0 < \epsilon \leq \min_{i \notin T} \frac{a_i \cdot x - b_i}{|a_i \cdot d|}$, we can ensure that $a_i \cdot y \geq b_i$ and $a_i \cdot z \geq b_i$. Thus y and z are feasible. Since $x = y/2 + z/2$, x cannot be an extreme point – a contradiction.

3. Basic Feasible Solution \Rightarrow Vertex

Let x be a basic feasible solution. Let $T = \{i \mid a_i \cdot x = b_i\}$. Consider the objective as minimizing $c \cdot y$ for $c = \sum_{i \in T} a_i$. Then, $c \cdot x = \sum_{i \in T} (a_i \cdot x) = \sum_{i \in T} b_i$.

For any $x' \in \mathcal{P}$, $c \cdot x' = \sum_{i \in T} (a_i \cdot x') \geq \sum_{i \in T} b_i$ with equality only if $a_i \cdot x' = b_i \forall i \in T$.

This implies that $x' = x$ and that x uniquely minimizes the objective $c \cdot y$.

This proves that vertex, extreme point and basic feasible solution are equivalent terms. ■

Theorem 2 *Any bounded LP in standard form has an optimum at a basic feasible solution.*

Proof: Consider an optimal x which is not a basic feasible solution. Being optimal, it is feasible, hence it is not basic. As in the previous proof, let T be the set of rows of the constraint matrix A for which the constraints are tight at x . Since x is not a basic solution, T does not span \mathcal{R}^n . So, there is a vector $d \neq 0$ such that $a_i \cdot d = 0 \forall a_i \in T$. For a scalar ϵ with sufficiently small absolute value, $y = x + \epsilon d$ is feasible, and represents a line containing x in the direction d . The objective function at y is $c^T x + \epsilon c^T d$. Since x is optimal, $c^T d = 0$, as otherwise, an ϵ of the opposite sign can reduce the objective. This means, all feasible points on this line are optimal. One of the directions of motion on this line will reduce some x_i . Keep going till some x_i reduces to 0. This results in one more tight constraint than before.

This technique can be repeated, till the solution becomes basic. ■

Thus, we can convert any feasible solution to a basic feasible solution of no worse value. In fact, this proof gives an algorithm for solving a linear program: evaluate the objective at all basic feasible solutions, and take the best one. Suppose there are m constraints and n variables. Since a set of n constraints defines a basic feasible solution, there can be upto $\binom{m}{n}$ basic feasible solutions. For each set of n constraints, a linear system of inequalities has to be solved, which by Gaussian elimination, takes $O(n^3)$ time. This is in general an exponential complexity algorithm in n . Note that the output size is polynomial in n , since the optimal solution is just the solution of a system of linear equalities.

1.2 The dual of a linear program

Given an LP in the standard form:

$$\begin{array}{l} \text{Minimize } c \cdot x \\ \text{subject to: } Ax = b; x \geq 0 \end{array}$$

We call the above LP the primal LP. The decision version of the problem is: Is the optimum $c \cdot x \leq \delta$? This problem is in *NP*, because, if we find a feasible solution with optimum value $\leq \delta$, we can verify that it satisfies these requirements, in polynomial time. A more interesting question is whether this problem is in *co-NP*. We need to find an easily verifiable proof for the fact that there is no x which satisfies $c \cdot x < \delta$. To do this, we require the concept of duality.

1.2.1 Weak Duality

We seek a lower bound on the optimum. Consider a vector y (treat it as a row vector here). For any feasible x , $yAx = yb$ holds. If we require that $yA \leq c$, then $yb = yAx \leq cx$. Thus, yb is a lower bound on cx , and in particular on the optimum cx . To get the best lower bound, we need to maximize yb . This new linear program:

$$\begin{array}{l} \text{Maximize } yb \\ \text{subject to: } yA \leq c \end{array}$$

is called the *dual linear program*. (Note: The dual of a dual program is the primal). Thus primal optimum is lower bounded by the dual optimum. This is called *weak duality*.

Theorem 3 (Weak Duality) Consider the LP $z = \text{Min}\{c \cdot x \mid Ax = b, x \geq 0\}$ and its dual $w = \text{max}\{y \cdot b \mid yA \leq c\}$. Then $z \geq w$.

Corollary 1 If the primal is feasible and unbounded, then the dual is infeasible.

1.3 Strong Duality

In fact, if either the primal or the dual is feasible, then the two optima are equal to each other. This is known as *strong duality*. In this section, we first present an intuitive explanation of the theorem, using a gravitational model. The formal proof follows that.

1.3.1 A gravitational model

Consider the LP $\min\{y \cdot b \mid yA \geq c\}$. We represent this feasible region as a hollow polytope, with the vector b pointing “upwards”. If a ball is dropped into the polytope, it will settle down at the lowest point, which is the optimum of the above LP. Note that any minimum is a global minimum, since the feasible region of an LP is a convex polyhedron. At the equilibrium point, there is a balance of forces – the gravitational force and the normal reaction of the floors (constraints). Let x_i represent the amount of force exerted by the i^{th} constraint. The direction of this force is given by the i^{th} column of A . Then the total force exerted by all the constraints Ax balances the gravity b : $Ax = b$.

The physical world also gives the constraints that $x \geq 0$, since the floors’ force is always outwards. Only those floors which the ball touches exert a force. This means that for the constraints which are not tight, the corresponding x_i ’s are zero: $x_i = 0$ if $yA_i > c_i$. This can be summarized as

$$(c_i - yA_i)x_i = 0$$

. This means x and y satisfy:

$$y \cdot b = \sum yA_ix_i = \sum c_ix_i = c \cdot x$$

But weak duality says that $yb \leq cx$, for every x and y . Hence the x and y are the optimal solutions of their respective LP’s. This implies strong duality – the optima of the primal and dual are equal.

1.3.2 A formal proof

Theorem 4 (Strong Duality) Consider $w = \min\{y \cdot b \mid yA \geq c\}$ and $z = \min\{c \cdot x \mid Ax = b, x \geq 0\}$. Then $z = w$.

Proof: Consider the LP $\min\{y \cdot b \mid yA \geq c\}$. Consider the optimal solution y^* . Without loss of generality, ignore all the constraints that are loose for y^* . If there are any redundant constraints, drop them. Clearly, these changes cannot alter the optimal solution. Dropping these constraints leads to a new A with fewer columns and a new shorter c . We will prove that the dual of the new LP has an optimum equal in value to the primal. This dual optimal solution can be extended to an optimal solution of the dual of the original LP, by filling in zeros at places corresponding to the dropped constraints. The point is that we do not need those constraints to come up with the dual optimal solution.

After dropping those constraints, at most n tight constraints remain (where n is the length of the vector y). Since we have removed all redundancy, these constraints are linearly independent. In terms of the new A and c , we have new constraints $yA = c$. y^* is still the optimum.

Claim: There exists an x , such that $Ax = b$.

Proof: Assume such an x does not exist, i.e. $Ax = b$ is infeasible. Then “duality” for linear equalities implies that there exists a z such that $zA = 0$, but $zb \neq 0$. Without loss of generality, assume $z.b < 0$ (otherwise, just negate the z). Now consider $(y^* + z)$. $A(y^* + z) = Ay^* + Az = Ay^*$. Hence, it is feasible. $(y^* + z).b = y^*.b + z.b < y^*.b$, which is better than the assumed optimum – a contradiction. So, there is an x such that $Ax = b$. Let this be called x^* .

Claim: $y^*.b = c.x^*$.

Proof: $y^*.b = y^*. (Ax^*) = (y^*A).x^* = c.x^*$ (since $Ax^* = b$ and $y^*A = c$)

Claim: $x^* \geq 0$

Proof: Assume the contrary. Then, for some i , $x_i^* < 0$. Let $c' = c + e_i$, where e_i is all 0's except at the i^{th} position, where it has a 1. Since A has full rank, $yA \geq c'$ has a solution, say y' . Besides, since $c' \geq c$, y' is feasible for the original constraints $yA \geq c$. But, $y'.b = y'Ax^* = c'.x^* < c.x^* = y^*.b$ (since c'_i is now higher and $x_i < 0$). This means y' gives a better objective value than the optimal solution – a contradiction. Hence, $x^* \geq 0$.

Thus, there is an x^* which is feasible in the dual, and whose objective is equal to the primal optimum. Hence, x^* must be the dual optimal solution, using weak duality. Thus, the optima of primal and dual are equal.

■

Corollary 2 *Checking for feasibility of a linear system of inequalities and optimizing an LP are equally hard.*

Proof: Optimizer \rightarrow Feasibility checker

Use the optimizer to optimize any arbitrary function with the linear system of inequalities as the constraints. This will automatically check for feasibility, since every optimal solution is feasible.

Feasibility checker \rightarrow Optimizer

We construct a reduction from the problem of finding an optimal solution of LP_1 to the problem of finding a feasible solution of LP_2 . LP_1 is $\min\{c.x \mid Ax = b, x \geq 0\}$. Consider $LP_2 = \min\{0.x \mid Ax = b, x \geq 0, yA \leq c, c.x = b.y\}$. Any feasible solution of LP_2 gives an optimal solution of LP_1 due to the strong duality theorem. Finding an optimal solution is thus no harder than finding a feasible solution. ■

1.4 Rules for duals

Usually the primal is constructed as a minimization problem and hence the dual becomes a maximization problem. For the standard form, the primal is given by:

$$\begin{aligned} z &= \min (c^T x) \\ Ax &\geq b \\ x &\geq 0 \end{aligned}$$

while the dual is given by:

$$\begin{aligned} w &= \max (b^T y) \\ A^T y &\leq c \\ y &\geq 0 \end{aligned}$$

For a mixed form of the primal, the following describes the dual:

Primal:

$$\begin{aligned} z &= \min c_1 x_1 + c_2 x_2 + c_3 x_3 \\ A_{11}x_1 + A_{12}x_2 + A_{13}x_3 &= b_1 \\ A_{21}x_1 + A_{22}x_2 + A_{23}x_3 &\geq b_2 \\ A_{31}x_1 + A_{32}x_2 + A_{33}x_3 &\leq b_3 \\ x_1 &\geq 0 \\ x_2 &\leq 0 \\ x_3 &\text{UIS} \end{aligned}$$

(UIS = unrestricted in sign)

Dual:

$$\begin{aligned} w &= \max y_1 b_1 + y_2 b_2 + y_3 b_3 \\ y_1 A_{11} + y_2 A_{21} + y_3 A_{31} &\leq c_1 \\ y_1 A_{12} + y_2 A_{22} + y_3 A_{32} &\geq c_2 \\ y_1 A_{13} + y_2 A_{23} + y_3 A_{33} &= c_3 \end{aligned}$$

$$\begin{aligned}
 y_1 & \quad \text{UIS} \\
 y_2 & \geq 0 \\
 y_3 & \leq 0
 \end{aligned}$$

These rules are summarized in the following table.

PRIMAL	Minimize	Maximize	DUAL
Constraints	$\geq b_i$ $\leq b_i$ $= b_i$	≥ 0 ≤ 0 Free	Variables
Variables	≥ 0 ≤ 0 Free	$\leq c_j$ $\leq c_j$ $= c_j$	Constraints

Each variable in the primal corresponds to a constraint in the dual, and vice versa. For a maximization, an upper bound constraint is a “natural” constraint, while for a minimization, a lower bound constraint is natural. If the constraint is in the natural direction, then the corresponding dual variable is non-negative.

An interesting observation is that, the tighter the primal gets, the looser the dual gets. For instance, an equality constraint in the primal leads to an unrestricted variable in the dual. Adding more constraints in the primal leads to more variables in the dual, hence more flexibility.

1.5 Shortest Path – an example

Consider the problem of finding the shortest path in a graph. Given a graph G , we wish to find the shortest path from a specified source node, to all other nodes. This can be formulated as a linear program:

$$\begin{aligned}
 w &= \max (d_t - d_s) \\
 \text{s.t. } & d_j - d_i \leq c_{ij}, \quad \forall i, j
 \end{aligned}$$

In this formulation, d_i represents the distance of node i from the source node s . The c_{ij} constraints are essentially the triangle inequalities – the distance from the source to a node i should not be more than the distance to some node j plus the distance from j to

i. Intuitively, one can imagine stretching the network physically, to increase the source-destination distance. When we cannot pull any further without breaking an edge, we have found a shortest path.

The dual to this program is found thus. The constraint matrix in the primal has a row for every pair of nodes (i, j) , and a column for every node. The row corresponding to (i, j) has a +1 in the i^{th} column and a -1 in the j^{th} column, and zeros elsewhere.

1. Using this, we conclude that the dual has a variable for each pair (i, j) , say y_{ij} .
2. It has a constraint for each node i . The constraint has a coefficient of +1 for each edge entering node i and a -1 for each edge leaving i . The right side for the constraints are -1 for the node s constraint, 1 for the node t constraint, and 0 for others, based on the objective function in the primal. Moreover, all the constraints are equality constraints, since the d_i variables were unrestricted in sign in the primal.
3. The dual variables will have to have a non-negativity constraint as well, since the constraints in the primal were “natural” (upper bounds for a maximization).
4. The objective is to minimize $\sum_{i,j} c_{ij}y_{ij}$, since the right side of the primal constraints are c_{ij} .

Thus the dual is:

$$z = \min \sum_{i,j} c_{ij}y_{ij}$$

$$\sum_j (y_{js} - y_{sj}) = -1$$

$$\sum_j (y_{jt} - y_{tj}) = 1$$

$$\sum_j (y_{ji} - y_{ij}) = 0, \forall i \neq s, t$$

$$y_{ij} \geq 0, \forall i, j$$

This is precisely the linear program to solve the minimum cost unit flow, in a gross flow formulation. The constraints correspond to the flow conservation at all nodes except at the source and sink. The value of the flow is forced to be 1. Intuitively, this says that we can use minimum cost unit flow algorithms to find the shortest path in a network.

Duality is a very useful concept, especially because it helps to view the optimization problem on hand from a different perspective, which might be easier to handle.