

Lecture 3: Duality Theory

Rajat Mittal

IIT Kanpur

One of the most beautiful and important concept in linear programming is called *duality*. According to Dantzig (inventor of simplex method), duality theory was first conjectured by John von Neumann.

The main intuition is that a linear program (more generally a convex program) can be viewed from two perspectives, a primal problem and a dual problem. When the optimal value of these two perspectives match, the program is said to have strong duality. We know that strong duality holds for linear programs.

The origin of such a strong theorem starts from hyperplane separation theorems, which itself can be viewed as a form of duality. The objective of this lecture note will be to learn duality theory of linear programming. We are planning to answer following questions.

- What are hyperplane separation theorems?
- What is Farkas lemma?
- How do we define the dual of a linear program?
- How to show strong duality for linear programs?

For this lecture, it will be helpful to know the definitions of open, closed and compact sets. We will only require the geometric intuition behind these definitions.

To begin, we will state another useful property of convex sets: there exists a closest point in a convex set for any point outside the convex set. The concept of orthogonal projection is known to us in the context of linear subspaces. Given a point x outside a linear subspace L , the projection of x on L is the closest point inside the subspace L to x . This concept can be extended to have the closest point to x with respect to a convex set C .

Lemma 1. *Suppose x is a point outside a non-empty closed convex set C . Then there exist a unique point x_c , on the boundary of C , which is closer to x than any other point in C . Say $d(x, y)$ denote the Euclidean distance between x and y , then*

$$d(x, x_c) < d(x, y) \quad \forall y \neq x \in C.$$

Note 1. We will omit the proof of this lemma, but the uniqueness part can be shown easily.

For the relevance of this theorem, note that uniqueness is not true for closed concave sets. For example, take all points in \mathbb{R}^n not inside a circle. The uniqueness is lost if x is the center of the circle.

Exercise 1. Come up with an example to prove that C is required to be closed in the lemma.

1 Separating hyperplanes

Given a point x and a set of points C , how can we possibly prove that x is not part of C ?

Suppose, there exist a vector a and a scalar b , s.t., $a^T x > b$. If it can be shown that $a^T y \leq b$ for every $y \in C$, this gives a proof that x is not part of C . We can say that a, b (which define the hyperplane $a^T x = b$) separate point x and set C .

It is natural that we ask the question, can we find such separating a, b for all pairs of x, C where $x \notin C$?

Exercise 2. Show that there exist C and $x \notin C$, s.t., no a, b separate them. What happens if C is a convex set?

It turns out, it is always possible to separate C and x by a hyperplane (if C is a closed convex set). Generalizing it further, even for two closed convex sets non-intersecting C_1 and C_2 , similar separation can be found.

Exercise 3. Give an example of two non-intersecting sets which cannot be separated.

Let us define this notion of separation more formally. A hyperplane $a^T x - b = 0$ *separates* two sets C and D , if

$$a^T x - b \geq 0, \forall x \in C \text{ and } a^T x - b \leq 0, \forall x \in D.$$

We will be interested in the case when both the convex sets are not contained in the hyperplane (*proper separation*). A stronger separation can be defined by taking strict inequalities. The separation is called a *strict separation* if strict inequalities hold on both sides.

$$a^T x - b > 0, \forall x \in C \text{ and } a^T x - b < 0, \forall x \in D.$$

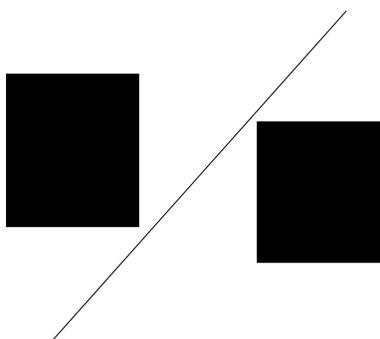


Fig. 1. Convex sets and hyperplanes separating them

There is a general theorem that two disjoint convex sets can be *separated* by a hyperplane. Depending on the convex sets the kind of separation can be different. Below we look at various cases.

1.1 Two convex sets

Theorem 1. *Given two convex sets C and D , which are mutually disjoint and non-empty, there always exists a hyperplane $a^T x - b = 0$ separating them. Hence,*

$$a^T x - b \geq 0, \forall x \in C \text{ and } a^T x - b \leq 0, \forall x \in D.$$

Note 2. The separation here is not strict.

Exercise 4. Construct two convex sets where separation cannot be strict. Can you construct two convex sets such that none of the inequalities can be strict?

We will show the proof for the case when at least one of the sets is compact and the other one is at least closed. We will mostly be dealing with such situations in future.

Theorem 2. *Given two convex sets C and D , mutually disjoint and non-empty, such that C is compact and D is closed. There always exists a hyperplane $a^T x - b = 0$ strictly separating them. In other words,*

$$a^T x - b > 0, \forall x \in C \text{ and } a^T x - b < 0, \forall x \in D.$$

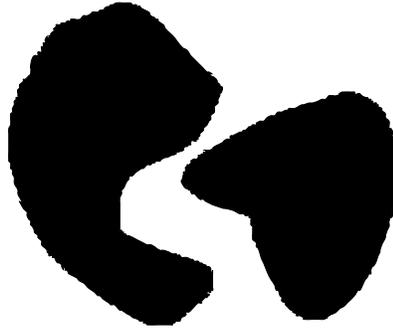


Fig. 2. Non-convex sets

Proof. In this case actually strict separation can be achieved. Since one set is compact and other is closed, there exist two points $c \in C$ and $d \in D$ whose distance is minimum among any pair of points, one from C and one from D . This distance can't be zero since C and D are mutually disjoint.

This statement seems quite intuitive but requires a proof. For people with a background of analysis (others can assume this statement), this follows from the fact that $d(x, D)$ (the minimum distance of point x from set D , which exists by Lemma 1) is a continuous function on a compact set C and hence its minimum should exist.

Let us look at the perpendicular bisector of the line between c and d . Our claim is, this will be the separating hyperplane.

Exercise 5. What is the equation of this hyperplane?

Notice that $(c - d)^T c \geq (c - d)^T d$. We will first show that if,

$$\exists x \in C \text{ such that } (c - d)^T x < (c - d)^T c, \quad (1)$$

then c and d are not the closest points.

So, suppose there exists a point $x \in C$, s.t.,

$$(c - d)^T x < (c - d)^T c \quad (2)$$

Since C is convex, for every $\theta \in [0, 1]$, point $\theta x + (1 - \theta)c$ lies in C . What is the distance between this point and d ? Consider,

$$\begin{aligned} & \|d - (\theta x + (1 - \theta)c)\|^2 \\ = & \|(d - c) - \theta(x - c)\|^2 \\ = & \|d - c\|^2 + \theta^2 \|x - c\|^2 - 2\theta(d - c)^T(x - c) \\ = & \|d - c\|^2 + \theta^2 \|x - c\|^2 - 2\theta(c - d)^T(c - x) \end{aligned}$$

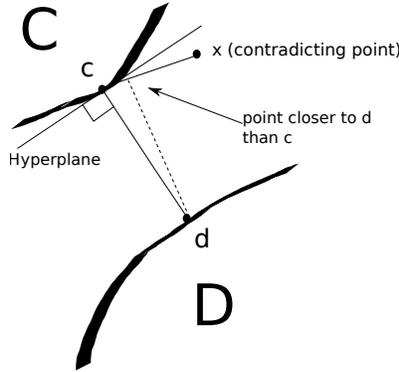


Fig. 3. Closest points

The last term $2\theta(c-d)^T(c-x)$ is positive because of our assumption, Eq. 2. Hence, by choosing θ small enough, we can make sure that the distance between point $\theta x + (1-\theta)c$ and d is less than the distance between c and d . This gives us the contradiction.

Exercise 6. What is the exact value of θ at which this will happen?

So,

$$\forall x \in C : (c-d)^T x \geq (c-d)^T c.$$

Eq. 1 shows that set C lies on opposite side of the hyperplane $(c-d)^T x = (c-d)^T c$ as compared with d . Similar calculations will show,

$$\forall x \in D : (c-d)^T x \leq (c-d)^T d.$$

Exercise 7. Show that $(c-d)^T x = (c-d)^T (c+d)/2$ is the separating hyperplane for sets C and D .

□

A point and a convex set Our next example will be a point and a convex set. In this case we get a strict separation by the hyperplane, i.e., point lies on one side of the hyperplane and the set on the other side. Here, strict means both the point and the set are disjoint with the hyperplane.

Theorem 3. *Given a closed convex set C and a point p . There always exist a hyperplane $a^T x - b$ which strictly separates them. So,*

$$a^T p - b > 0 \text{ and } a^T x - b < 0, \forall x \in C.$$

The proof follows by constructing a small enough ball around point p which is compact and disjoint with C . Then, the result follows from Thm. 2.

An interesting corollary follows.

Corollary 1. *Any closed convex set is the intersection of all the half spaces which contain it.*

Proof is left as an exercise. The set of half spaces can be restricted to all the supporting hyperplanes.

1.2 Farkas lemma: A point and a cone

A special case of separating hyperplane theorems is the separation between a point and a finitely generated cone. Remember that a finitely generated cone is convex. The proof of the lemma below follows from Thm. 3, but it will be instructive to see another proof of the same statement.

Lemma 2 (Farkas). *Given a set of vectors $x_1, \dots, x_k \in \mathbb{R}^n$ (equivalently $C \in \mathbb{R}^{n \times k}$ with x_1, \dots, x_k as columns) and a point $b \in \mathbb{R}^n$. Exactly one of the following two conditions are satisfied.*

1. $\exists \alpha \in \mathbb{R}_+^k$ (α in positive orthant), such that, $C\alpha = b$.
2. $\exists a \in \mathbb{R}^n$, such that, $a^T b > 0$ and $a^T C \leq 0$ (entry-wise).

Exercise 8. How does Farkas lemma follow from Thm. 3?

Note 3. The inequality is strict (strict separation) and the hyperplane is of special kind.

Proof. It is clear that both conditions cannot be satisfied simultaneously. Look at the following series of implications.

$$\begin{aligned} a^T C &\leq 0 \\ \Rightarrow a^T C \alpha &\leq 0 \\ \Rightarrow a^T b &\leq 0. \end{aligned}$$

But, we know that $a^T b > 0$, giving a contradiction.

Let us prove the exclusivity. Suppose, b can't be written as the conic combination of columns of C , i.e., first condition does not hold. We will show that there exists $a \in \mathbb{R}^n$ satisfying the second condition.

Look at cone generated by columns of C . It is finitely generated cone and hence by Weyl's theorem, it can be expressed as a bunch of inequalities.

$$\text{Cone}(x_1, \dots, x_k) = \{x : Ax \leq 0\}$$

Since b is not a member of this cone, there exists a row of A whose inner product with b is strictly positive (call it a_i). This implies,

$$a_i^T b > 0 \text{ and } \forall i, a_i^T x_i \leq 0 \Rightarrow a_i^T C \leq 0.$$

□

Exercise 9. Interpret Farkas lemma as a hyperplane separation theorem. What do you know about this hyperplane?

This theorem converts the membership question in a cone to finding a separating hyperplane question. The question of membership in the cone is of real importance in convex optimization. It can be shown that optimization over a cone can be done using polynomially many calls to membership/separation algorithm for the cone.

We will now look at taking dual of a linear program. In the beginning, you will not see much connection of the newly defined duals with hyperplane separating theorems. Later, in Sec. 4 you will see that hyperplane separation theorems are preliminary versions of duality.

2 Dual of a linear program

One of the strongest aspect of linear optimization is the presence of duality theory. The dual of a linear program (called *primal*) is another linear optimization program (called *dual*), which provides *tight* upper/lower bounds on the original linear program.

We will motivate duality theory by looking at ways to give bounds on the value of a linear program. Let us consider a maximization problem, a feasible solution of this program gives a lower bound on the optimal value. Is there a way to give an upper bound on the optimal value?

We will start with a very simple linear program.

$$\begin{aligned} \max \quad & 2x_1 + 3x_2 + x_3 \\ \text{s.t.} \quad & x_1 + x_2 + x_3 = 5 \\ & x_1 + 2x_2 = 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

On a close inspection, it will be clear that the objective function is a linear combination of constraints. Hence whatever be the feasible solution, the objective value will be 15.

Notice that it is required to have a feasible solution. In case there is no feasible solution then the objective value of the program will be $-\infty$. Remember that if a max optimization program does not have a feasible solution then its optimal value is $-\infty$ and if the min optimization program is infeasible then its optimal value is $+\infty$.

Taking another example,

$$\begin{aligned} \max \quad & 2x_1 + 3x_2 + x_3 \\ \text{s.t.} \quad & x_1 + x_2 + x_3 = 5 \\ & x_1 + 2x_2 + x_3 = 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

In this case it is not clear if 15 is the optimal value. But since $x_3 \geq 0$, the objective value is $\leq 15 - x_3$ and hence lesser than or equal to 15. So 15 is an upper bound on the objective value. Let us make this argument more precise.

We will consider a linear program in the standard form,

$$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0. \end{aligned} \tag{3}$$

If there exist a linear combination of constraints, say y , s.t., $A^T y = c$. In this case, every feasible solution will have objective value $b^T y$. So, if the above program is feasible, the optimal value of the program is $b^T y$.

Though, getting a y such that $A^T y = c$ is not easy. We saw instead that if $A^T y \geq c$, even then $b^T y$ is an upper bound on the value of the linear program.

Since we want to consider the best upper bound, it is natural to consider the following problem.

$$\begin{aligned} \min \quad & b^T y \\ \text{s.t.} \quad & A^T y \geq c. \end{aligned} \tag{4}$$

We first observe that this new optimization problem is also linear. From the way this program is constructed, the optimal value of Eqn. 4 is always higher than optimal value of Eqn. 3.

Eqn. 3 is known as the *primal linear program*. Eqn. 4 is known as the *dual linear program*.

Exercise 10. What is the dual if the constraint was $Ax \leq b$ instead of $Ax = b$ in the primal linear program?

The argument above showed that the dual objective value is higher than the primal objective value. This phenomenon is known as the *weak duality*. Importantly, in case of linear programs, it is known that the dual value is same as the primal value, given that primal is feasible and bounded. This is known as *strong duality* and will be proved in Sec. 4.

The terms primal and dual are interchangeable. The dual of the dual is the primal program. So, if Eqn. 4 is considered as the primal program then the dual will be Eqn. 3.

Exercise 11. Verify it by the lower bound argument.

Another important thing to notice is that for every primal constraint there is a dual variable and for every primal variable there is a dual constraint. This relationship is much more stronger than it seems now and will be explored more in Sec. 3.

Exercise 12. Find the dual of the following linear program.

$$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & Ax \geq b \\ & x \leq 0, \end{aligned}$$

Exercise 13. What is the relationship between dual and separating hyperplanes?

For easier access to taking dual, this table might be helpful.

	Primal variable to dual constraint	Primal constraint to dual variable
Max	$\geq \rightarrow \geq$, $\leq \rightarrow \leq$, unrestricted $\rightarrow =$	$\geq \rightarrow \leq$, $\leq \rightarrow \geq$, $= \rightarrow$ unrestricted
Min	$\geq \rightarrow \leq$, $\leq \rightarrow \geq$, unrestricted $\rightarrow =$	$\geq \rightarrow \geq$, $\leq \rightarrow \leq$, $= \rightarrow$ unrestricted

Exercise 14. Verify that the table is correct.

Unbounded/infeasible linear programs It was mentioned that if primal is feasible and bounded then the optimal value of primal and dual coincide (shown in Sec. 4).

What happens when primal is unbounded or infeasible?

Assume that the primal problem is unbounded. In this case, we can't have a feasible solution of the dual problem. Otherwise, the dual feasible solution will give a concrete upper bound on the value of the primal program contradicting unboundedness.

Similarly, if the dual is unbounded, primal will be infeasible.

Can both programs be infeasible. Take a look at the following linear program,

$$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & -x_1 \geq 1, x_2 \geq 1 \\ & x \geq 0, \end{aligned}$$

Clearly this program is infeasible.

Exercise 15. Find the dual of this linear program and show that it is infeasible.

The discussion above shows,

- Both primal and dual could be infeasible.
- If primal is unbounded then dual is infeasible.
- If dual is unbounded then primal is infeasible.

3 Complementary slackness

In the last section we saw how to formulate a dual program from a linear program. From the formulation, it turned out that the dual program gave an upper/lower bound depending on whether the problem was maximization/minimization respectively.

The strength of duality theory lies in the fact that these bounds are in fact tight. This implies that the dual value always agrees with the primal value. This fact has been useful in numerous applications in computer science. Let us study this relationship in detail.

Note 4. The fact that these bounds are tight only holds true for linear programming. For other convex programming instances, we need to be careful.

Let's look at the primal-dual pair for a linear program.

Primal	Dual
$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$	$\begin{aligned} \min \quad & y^T b \\ \text{s.t.} \quad & A^T y - c \geq 0 \\ & y \geq 0 \end{aligned}$

Note 5. We have changed the form slightly. Convince yourself that it does not matter.

First, it is instructive to see the direct proof of weak duality. Given any feasible solution X for primal and y for dual,

$$c^T x \leq y^T Ax \leq y^T b. \tag{5}$$

Note 6. Above equation implies, any feasible solution of the dual give an upper bound on the optimal value of the primal. Similarly, any feasible solution for the primal gives a lower bound on the optimal value of the dual.

Suppose, the optimal value of the primal is p^* , and is attained at point x^* . Similarly, the optimal value of the dual is d^* and is obtained for the point y^* .

Weak duality implies that $p^* \leq d^*$. Assume that this two values are equal, i.e., $p^* = d^*$ (strong duality, Thm. 4). Then, from Eqn. 5,

$$(A^T y^* - c)^T x^* = 0 \text{ and } (Ax^* - b)^T y^* = 0. \tag{6}$$

These two conditions are called the *complementary slackness* condition.

We get the following two observations.

1. Let x^*, y^* be optimal solutions of primal and dual respectively. Since primal is both bounded and feasible, strong duality (Thm. 4) as well as complementary slackness (Eqn. 6) holds.
2. Conversely, if x, y are feasible solutions of primal and dual respectively and satisfy the complementary slackness condition then strong duality holds and $p^* = d^*$.

The three conditions (primal feasibility, dual feasibility and complementary slackness for x, y) ensure that x, y are optimal. These three conditions combined are called *Karush-Kuhn-Tucker (KKT) conditions*. They can be generalized for convex programming.

Looking at complementary slackness more deeply,

- Given an optimal solution x^* for the primal, if $x_i^* \neq 0$ then $A_{*,i}^T y - c_i = 0$, where $A_{*,i}$ is the i^{th} column of A . In other words, if the primal optimal variable is non-zero, then the corresponding constraint in dual is tight.
- Given an optimal solution y^* for the dual, if $y_j^* \neq 0$ then $A_{j,*}^T x - b_j = 0$, where $A_{j,*}$ is the j^{th} column of A . In other words, if the dual optimal variable is non-zero, then the corresponding constraint in primal is tight.

There is another way to interpret the same thing. Assign a variable to each equation, i.e., we write $A_{*,i}^T y - c_i = w_i$ and $A_{j,*}^T x - b_j = z_j$. Then, both primal and dual problems have $m + n$ variables. Complementary slackness tells us that out of these $m + n$ variables, if a primal variable is non-zero then the corresponding dual variable is 0 and vice versa.

4 Strong duality

Strong duality states that the optimal value of primal is same as optimal value of dual. The next theorem states that strong duality holds for all linear programs. We again look at the standard format for primal and dual linear programs,

$$\begin{array}{ll} \max & c^T x \\ \text{s.t.} & Ax - b = 0 \\ & x \geq 0 \end{array} \qquad \begin{array}{ll} \min & b^T y \\ \text{s.t.} & A^T y - c \geq 0 \end{array}$$

We assume that A is an $m \times n$ matrix.

Theorem 4. *Given a linear program in above format with parameters c, A, b (A is $m \times n$ matrix), suppose the feasible set of primal is \mathcal{P} and feasible set of dual is \mathcal{D} . Then one of the following is true,*

- If \mathcal{P} is empty (primal infeasible), then dual is unbounded or infeasible.
- If \mathcal{D} is empty (dual infeasible), then primal is unbounded or infeasible.
- If both \mathcal{D}, \mathcal{P} are non-empty, then strong duality holds. So for optimal x^*, y^* , $p^* = c^T x^* = b^T y^* = d^*$.

In other words, if any one of primal or dual is feasible and bounded then strong duality holds.

The proof is based on Farkas Lemma.

Proof. Suppose \mathcal{P} is empty. By Farkas lemma, there exist y , s.t., $b^T y < 0, A^T y \geq 0$. If dual is infeasible, then there is nothing to prove, else there exist $y' : A^T y' \geq c$.

Exercise 16. Show that using y, y', z'_i s can be constructed, s.t., $A^T z'_i \geq c, b^T z'_i \rightarrow -\infty$ as $i \rightarrow \infty$.

From the previous exercise, dual is unbounded. Similarly, we can deal with the case when \mathcal{D} is empty.

Exercise 17. Show that if primal is unbounded, then dual is infeasible.

We now assume that both \mathcal{P}, \mathcal{D} are non-empty and the optimal value is bounded. Say, $p^*(d^*)$ are the optimal values of primal(dual) achieved at $x^*(y^*)$ respectively.

Since p^* is optimal, there does not exist x , s.t., $(Ax, c^T x) = (b, p^* + \epsilon)$ (for all $\epsilon > 0$). Considering $A' = (A, c)$ and $b' = (b, p^* + \epsilon)$, Farkas lemma implies,

$$\exists y' : A'^T y' \leq 0 \text{ and } b'^T y' > 0$$

We can scale y' to ensure that the last co-ordinate of y' (y'_{m+1}) is either $+1, -1$ or 0 . Let y denote y' except the last co-ordinate (which can be $1, -1$ or 0).

Suppose $y'_{m+1} = -1$. Then,

$$\begin{aligned} & b^T y > p^* + \epsilon \text{ and } A^T y \leq c \\ \Rightarrow & b^T y > p^* \text{ and } (x^*)^T A^T y \leq c^T x^* = p^* \\ \Rightarrow & b^T y > p^* \text{ and } (b^T) y = p^* \text{ (since } Ax^* = b) \end{aligned}$$

This is a contradiction.

Exercise 18. Show that y'_{m+1} is not equal to zero.

So, we can assume $y'_{m+1} = 1$. Putting back the definition of A', b' ,

$$\exists y : A^T y + c \leq 0 \text{ and } b^T y + p^* + \epsilon > 0$$

Replacing y by $-y$,

$$\exists y' : A^T y' \geq c \text{ and } b^T y' < p^* + \epsilon$$

So, we have a feasible dual solution with objective value less than $p^* + \epsilon$ for all $\epsilon \geq 0$. That means, the dual optimal value is less than $p^* + \epsilon$ for every $\epsilon > 0$. Hence, the dual optimal value is p^* (remember that it can't be less than p^*). □

Note 7. There can be cases when both primal and dual are infeasible.

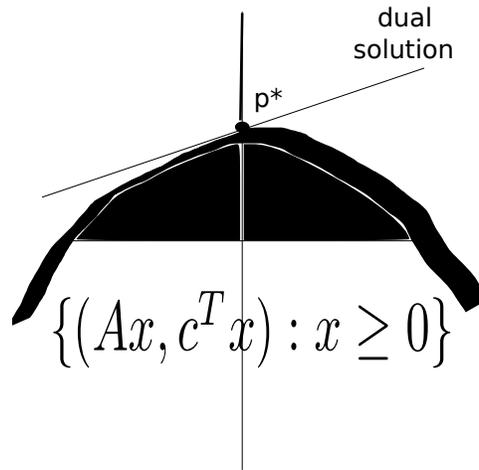


Fig. 4. Strong duality through picture

Extra reading: why do we NOT need infimum (supremum) in linear programming For people familiar with real analysis, it might be confusing that our optimization problems were phrased as max/min and not inf/sup. In other words, is it possible that the infimum/supremum is never attained for a linear program. The answer is NO, and that is why we can phrase problems as max/min. Below, we outline an argument to show that inf/sup is always attained in a linear program.

Consider weak duality,

$$\sup\{c^T x : Ax = b, x \geq 0\} \leq \inf\{b^T y : y^T A \geq c\}. \quad (7)$$

Define $\Gamma = \inf\{b^T y : y^T A \geq c\}$. It implies that for all y such that $y^T A \geq c$, we know $b^T y \geq \Gamma$.

Exercise 19. Using Farkas Lemma, show that the above condition is equivalent to,

$$\exists x : Ax = b, c^T x \geq \Gamma.$$

So, we get that the supremum in Eq. 7 is attained. So, inf/sup is attained in any linear program.

5 Assignment

Exercise 20. Show that the closest point to a convex set is unique (from Lemma 1).

Exercise 21. Modify the proof of Thm. 1 to show the Thm. 3 directly.

Exercise 22. Prove Cor. 1.

Exercise 23. Show that the dual of the dual is the primal program (for the standard linear program).

Hint: First convert the dual into primal form and then take the dual.

Exercise 24. Define dual cone of a cone C to be,

$$\bar{C} = \{y : y^T x \geq 0 \quad \forall x \in C\}.$$

What is the dual cone of $\{x : x \geq 0\}$ and \mathbb{R}^n ? Can you guess the relationship between dual cone and duality?