Math 3272: Linear Programming[1] | Mikhail Lavrov

Lecture 13: Duality

September 29, 2022 | *Kennesaw State University*

1 An example of duality

Problem 1. *You visit a chocolate factory and want to buy as much chocolate as you can. The factory sells plain chocolate chips for $1 per pint and deluxe chocolate chips for $2 per pint. You can only carry one pint of chocolate in your hands; if you want more, you'll have to buy a bag. Empty bags of all sizes are available; an empty 3-pint bag costs $4, and all other sizes cost a proportional amount.*

If you have $7, what is the largest amount of chocolate you can buy and take home with you?

Let x_1 be the amount of plain chocolate chips and x_2 the amount of deluxe chocolate chips, in pints (so that we want to maximize $x_1 + x_2$). Let x_3 be the number of 3-pint bags you buy (if it is a fraction, we assume that you bought some other size of bag.) Then the amount of money you brought limits these variables to $x_1 + 2x_2 + 4x_3 \le 7$. Also, you can carry at most $1 + 3x_3$ pints of chocolate, so $x_1 + x_2 \le 1 + 3x_3$, or $x_1 + x_2 - 3x_3 \le 1$.

In summary, we get the linear program below:

$$\begin{array}{rl} \underset{x_1,x_2,x_3\in\mathbb{R}}{\text{maximize}} & x_1 + x_2 \\ \text{subject to} & x_1 + 2x_2 + 4x_3 \le 7 \\ & x_1 + x_2 - 3x_3 \le 1 \\ & \quad\; x_1, x_2, x_3 \;\; \ge 0 \end{array}$$

Today, we're going to be too lazy to try to solve this linear program. Instead, we want to prove some lower and upper bounds on the objective value of the solution.

Lower bounds for a maximization problem are easy to find.

- Setting $x_1 = x_2 = x_3 = 0$ satisfies both constraints, so clearly we can't do worse than an objective value of 0.
- We could try tweaking that: say $x_1 = 1$ and $x_2 = x_3 = 0$, then we get an objective value of 1.
- In general, any feasible solution gives us a lower bound on the objective value. If we wanted to get good lower bounds this way, we'd start trying to solve the linear program, which we said we didn't want to do.

What about upper bounds? Well, here are some ideas:

- $x_1 + x_2$ is always less than or equal to $x_1 + 2x_2 + 4x_3$. So if $x_1 + 2x_2 + 4x_3 \le 7$, we can immediately conclude $x_1 + x_2 \le 7$.

[1] This document comes from the Math 3272 course webpage: https://facultyweb.kennesaw.edu/mlavrov/courses/3272-fall-2022.php

- Note that we *can't* conclude from $x_1 + x_2 - 3x_3 \le 1$ that $x_1 + x_2 \le 1$, because the $-3x_3$ term could potentially make $x_1 + x_2 - 3x_3$ a lot smaller than $x_1 + x_2$.
- However, if we average the two constraints, we get an improvement:

$$\frac{1}{2}(x_1 + 2x_2 + 4x_3) + \frac{1}{2}(x_1 + x_2 - 3x_3) \le \frac{1}{2}(7+1) \implies x_1 + \frac{3}{2}x_2 + \frac{1}{2}x_3 \le 4$$

  and we always have $x_1 + x_2 \le x_1 + \frac{3}{2}x_2 + \frac{1}{2}x_3$, so we conclude that $x_1 + x_2 \le 4$.

More generally, we could try to combine the two constraints with any coefficients. As long as $u_1 \ge 0$ and $u_2 \ge 0$, we can try to combine the inequalities with weights u_1 and u_2 to get

$$u_1(x_1 + 2x_2 + 4x_3) + u_2(x_1 + x_2 - 3x_3) \le 7u_1 + u_2.$$

Rearranging the inequality to group the x_1, x_2, and x_3 terms together, we get

$$(u_1 + u_2)x_1 + (2u_1 + u_2)x_2 + (4u_1 - 3u_2)x_3 \le 7u_1 + u_2.$$

This is a valid inequality, but not necessarily a useful one. We want the left-hand side to be an upper bound on $x_1 + x_2$ if we want to apply the same logic that we did earlier. For this to happen, the coefficients of x_1 and x_2 must be at least 1, and the coefficient of x_3 must be nonnegative. This gives us three constraints on u_1 and u_2 in order for $7u_1 + u_2$ to be an upper bound.

What is the best upper bound we can find by combining the inequalities in this way? The answer can be found by solving a different linear program in terms of u_1 and u_2:

$$\begin{array}{rl} \underset{u_1,u_2 \in \mathbb{R}}{\text{minimize}} & 7u_1 + u_2 \\ \text{subject to} & u_1 + u_2 \ge 1 \\ & 2u_1 + u_2 \ge 1 \\ & 4u_1 - 3u_2 \ge 0 \\ & u_1, u_2 \ge 0 \end{array}$$

2 Weak duality

In matrix form, we can write our constraints in the original problem as $A\mathbf{x} \le \mathbf{b}$, where

$$A = \begin{bmatrix} 1 & 2 & 4 \\ 1 & 1 & -3 \end{bmatrix} \qquad \mathbf{b} = \begin{bmatrix} 7 \\ 1 \end{bmatrix}.$$

The combined inequality $u_1(x_1+2x_2+4x_3)+u_2(x_1+x_2-3x_3) \le 7u_1+u_2$ can be written in matrix form as

$$\begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} 1 & 2 & 4 \\ 1 & 1 & -3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \le \begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} 7 \\ 1 \end{bmatrix}.$$

In matrix notation: starting from $A\mathbf{x} \le \mathbf{b}$, we deduced that $\mathbf{u}^\mathsf{T} A\mathbf{x} \le \mathbf{u}^\mathsf{T}\mathbf{b}$. (As a reminder, combining these inequalities in this way is only valid provided that $\mathbf{u} \ge \mathbf{0}$.)

We don't just want to deduce valid inequalities: we want to deduce *useful* inequalities. We want the coefficients of x_1, x_2, x_3 on the left-hand side to be at least as big as the coefficients of x_1, x_2, x_3

in the objective function $x_1 + x_2$, so that we get an upper bound on $x_1 + x_2$. The three constraints $u_1 + u_2 \geq 1$, $2u_1 + u_2 \geq 1$, and $4u_1 - 3u_2 \geq 0$ can be written in matrix form as

$$\begin{bmatrix} u_1 & u_2 \end{bmatrix} \begin{bmatrix} 1 & 2 & 4 \\ 1 & 1 & -3 \end{bmatrix} \geq \begin{bmatrix} 1 & 1 & 0 \end{bmatrix}.$$

If $\mathbf{c}^\mathsf{T} = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix}$ is the cost vector in our objective function, then these constraints can be written as $\mathbf{u}^\mathsf{T} A \geq \mathbf{c}^\mathsf{T}$. (We can also take the transpose of both sides, and write $A^\mathsf{T}\mathbf{u} \geq \mathbf{c}$.) You can see that in the two linear programs we wrote down in the previous section, the matrix of coefficients of $\mathbf{u}$ is the transpose of the matrix of coefficients of $\mathbf{x}$.

We can do this for any linear program. Write down a general linear program in the form

$$\textbf{(P)} \begin{cases} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} & \mathbf{c}^\mathsf{T}\mathbf{x} \\ \text{subject to} & A\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0} \end{cases}$$

where A is an $m \times n$ matrix, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{c} \in \mathbb{R}^n$. The linear program of "what is the best upper bound we can deduce on **(P)** by taking a linear combination of its inequalities?" is called the **dual linear program**, and has the form below:

$$\textbf{(D)} \begin{cases} \underset{\mathbf{u}\in\mathbb{R}^m}{\text{minimize}} & \mathbf{u}^\mathsf{T}\mathbf{b} \\ \text{subject to} & \mathbf{u}^\mathsf{T}A \geq \mathbf{c}^\mathsf{T} \\ & \mathbf{u} \geq \mathbf{0} \end{cases} \iff \begin{matrix} \underset{\mathbf{u}\in\mathbb{R}^m}{\text{minimize}} & \mathbf{b}^\mathsf{T}\mathbf{u} \\ \text{subject to} & A^\mathsf{T}\mathbf{u} \geq \mathbf{c} \\ & \mathbf{u} \geq \mathbf{0} \end{matrix}$$

(When **(D)** is the dual linear program of **(P)**, we call **(P)** the **primal linear program**.)

The dual linear program above is written in two forms. On the right, we took the transpose of both sides, putting it into a form more usual for linear programs. But when we think about the dual relationship between **(P)** and **(D)**, it's more convenient to use the formulation on the left, because then the dual program is distinguished by being in terms of a row vector $\mathbf{u}^\mathsf{T}$ instead of a column vector $\mathbf{u}$.

The reasoning by which **(D)** gives upper bounds for **(P)** holds in general. Formally, this relationship is called **weak duality**, and is summarized in the theorem below:

Theorem 1 (Weak duality of linear programs)**.** *For any $\mathbf{x} \in \mathbb{R}^n$ which is feasible for the primal linear program* (**P**) *(or **primal feasible**) and for any $\mathbf{u} \in \mathbb{R}^m$ which is feasible for the dual linear program* (**D**) *(or **dual feasible**), we have $\mathbf{c}^\mathsf{T}\mathbf{x} \leq \mathbf{u}^\mathsf{T}\mathbf{b}$.*

In particular, the objective value of the dual optimal solution is an upper bound for the objective value of the primal optimal solution (assuming both optimal solutions exist).

Proof. Since $\mathbf{x}$ is primal feasible, we have $A\mathbf{x} \leq \mathbf{b}$. Since $\mathbf{u}$ is dual feasible, we have $\mathbf{u} \geq \mathbf{0}$. Therefore the inequality $\mathbf{u}^\mathsf{T}A\mathbf{x} \leq \mathbf{u}^\mathsf{T}\mathbf{b}$ is valid: we have multiplied the inequalities in $A\mathbf{x} \leq \mathbf{b}$ by nonnegative coefficients $u_1, \dots, u_m$ and added them together.

Since $\mathbf{u}$ is dual feasible, we have $\mathbf{u}^\mathsf{T}\mathbf{A} \geq \mathbf{c}^\mathsf{T}$. Since $\mathbf{x}$ is primal feasible, we have $\mathbf{x} \geq \mathbf{0}$. By the same logic as above, we can deduce that $\mathbf{u}^\mathsf{T}\mathbf{A}\mathbf{x} \geq \mathbf{c}^\mathsf{T}\mathbf{x}$: again, we have multiplied the inequalities in $\mathbf{u}^\mathsf{T}\mathbf{A} \geq \mathbf{c}^\mathsf{T}$ by nonnegative coefficients $x_1, \dots, x_n$, then added them together.

Putting these together, we get $\mathbf{c}^\mathsf{T}\mathbf{x} \le \mathbf{u}^\mathsf{T} A\mathbf{x} \le \mathbf{u}^\mathsf{T}\mathbf{b}$, so $\mathbf{c}^\mathsf{T}\mathbf{x} \le \mathbf{u}^\mathsf{T}\mathbf{b}$. □

3 Duals of other kinds of programs

So far we've discussed starting with a primal program that's a maximization problem with nonnegative variables and an $A\mathbf{x} \le \mathbf{b}$ constraint. Duality is more general than this: it can handle any kind of linear program. The only thing that never changes is that

Variables in one program correspond to constraints in the other.

Just to give a few examples of how things change:

- Suppose that we drop the requirement that $x_1 \ge 0$ in our original linear program. (We can buy negative chocolate chips, which have negative weight and cost negative money.)

  As before, an expression such as $\frac{1}{2}x_1 + 2x_2$ is not an upper bound on $x_1 + x_2$, because $\frac{1}{2}x_1$ might be less than x_1: if x_1 is large and x_2 is small, then $\frac{1}{2}x_1 + 2x_2 < x_1 + x_2$. However, this time, an expression such as $2x_1 + 2x_2$ is also not an upper bound on $x_1 + x_2$: if x_1 is a large negative number, then $2x_1 < x_1$.

  So we see that in any inequality which gives an upper bound on $x_1 + x_2$, the coefficient of x_1 has to be *exactly* 1. Our inequality $u_1 + u_2 \ge 1$ would become $u_1 + u_2 = 1$.

  In general, an unconstrained variable gives us $=$ constraints in the dual linear program.

- Suppose that we reverse the first constraint to say $x_1 + 2x_2 + 4x_3 \ge 7$. (We have unlimited money and must spend *at least* \$7 of it.)

  In this case, if we still want upper bounds on some expression in terms of x_1 and x_2, we have to multiply this constraint by a *negative* coefficient to reverse the inequality. Instead of wanting $u_1 \ge 0$, we'd want $u_1 \le 0$.

  In general, a $\ge$ constraint gives us a nonpositive variable in the dual linear program.

- Suppose that the primal program asks to *minimize* $x_1 + x_2$ instead of maximizing it.

  This changes everything, because now we are trying to get lower bounds instead of upper bounds. In particular, the relationship between **(P)** and **(D)** is reversed: a feasible solution for **(P)** will always have a greater or equal objective value compared to a feasible solution for **(D)**.

  Now $\le$ constraints in **(P)** correspond to nonnegative variables in **(D)** (they are the "natural" kind of constraint when we're minimizing) and $\ge$ constraints in **(P)** correspond to nonpositive variables in **(D)**.

Actually, the relationship between **(P)** and **(D)** is symmetric: if **(D)** is the dual of **(P)**, then **(P)** is the dual of **(D)**. It's easiest to describe the duality relationship as a relationship between a maximization problem and a minimization problem, never mind which one of them was the primal and which was the dual.

With that in mind, here is the complete list of possible correspondences betwene a constraint in one problem and a variable in the other:

Maximization problem	Minimization problem
$\leq$ constraint	variable ≥ 0
$=$ constraint	unconstrained variable
$\geq$ constraint	variable ≤ 0
variable ≥ 0	$\geq$ constraint
unconstrained variable	$=$ constraint
variable ≤ 0	$\leq$ constraint

Memorizing the rules in the table is possible, but it probably isn't very satisfying. It is healthier to practice figuring out the correspondence for yourself, by asking the questions in the examples above: how can we combine the constraints of the primal problem to get bounds on its optimal value, of whichever kind makes sense?

4 Strong duality

4.1 A stronger theorem

In fact, a stronger relationship between **(P)** and **(D)** holds, which is appropriately enough called **strong duality**. It says that:

Theorem 2 (Strong duality of linear programs). *If either one of **(P)** or **(D)** has an optimal solution, then so does the other one. The objective values of the optimal solutions are equal.*

In other words, the dual program is *good* at finding bounds on the primal program: the best bound it finds is exactly correct.

We have not yet proved strong duality. (We will see a proof later.)

However, keep in mind the word "if" at the beginning of this theorem. We are not guaranteed that a linear program has an optimal solution: it could be unbounded, or infeasible!

In fact, just from weak duality, we can already deduce a relationship between unbounded and infeasible linear programs.

- Suppose that **(P)** has a feasible solution $\mathbf{x}$. Then we know that for every dual feasible $\mathbf{u}$, we have $\mathbf{c}^\mathsf{T}\mathbf{x} \leq \mathbf{u}^\mathsf{T}\mathbf{b}$. Therefore $\mathbf{u}^\mathsf{T}\mathbf{b}$ cannot be arbitrarily low: it is bounded below by whatever $\mathbf{c}^\mathsf{T}\mathbf{x}$ is. So **(D)** cannot be unbounded.

  Conversely, if **(D)** is unbounded, it tells us that **(P)** is infeasible.

- By similar reasoning, any dual feasible $\mathbf{u}$ proves that **(P)** cannot be unbounded. Therefore if **(P)** *is* unbounded, **(D)** must be infeasible.

(It is also possible for both **(P)** and **(D)** to be infeasible in exceptionally unfortunate cases.)

4.2 Examples with infeasible primal and dual

Here are some very simple examples that illustrate all three possibilities where **(P)** or **(D)** is infeasible. (I have labeled each constraint in one program with the corresponding variable in the other program, which we're going to keep doing in the future for all our primal-dual pairs.)

In the pair

$$
\textbf{(P)}\ \begin{cases} \underset{x\in\mathbb{R}}{\text{maximize}} & x \\ \text{subject to} & x \le -1 \quad (u) \\ & x \ge 0 \end{cases}
\qquad
\textbf{(D)}\ \begin{cases} \underset{u\in\mathbb{R}}{\text{minimize}} & -u \\ \text{subject to} & u \ge 1 \quad (x) \\ & u \ge 0 \end{cases}
$$

the primal program is infeasible (we can't have $x \le -1$ and $x \ge 0$ at the same time) and the dual program is unbounded (by setting u to be very large, we make $-u$ very small).

We can get an example where the primal program is infeasible and the dual program is unbounded simply by reversing the roles of the two programs. Or, if we want to keep **(P)** a maximization problem and **(D)** a minimization problem, we could do a slight variant of the example above:

$$
\textbf{(P)}\ \begin{cases} \underset{y\in\mathbb{R}}{\text{maximize}} & y \\ \text{subject to} & -y \le 1 \quad (v) \\ & y \ge 0 \end{cases}
\qquad
\textbf{(D)}\ \begin{cases} \underset{v\in\mathbb{R}}{\text{minimize}} & v \\ \text{subject to} & -v \ge 1 \quad (y) \\ & v \ge 0 \end{cases}
$$

Here, any nonnegative y is primal feasible, but no v is dual feasible.

To get an example where both linear programs are infeasible, just combine these two examples:

$$
\textbf{(P)}\ \begin{cases} \underset{x,y\in\mathbb{R}}{\text{maximize}} & x+y \\ \text{subject to} & x \le -1 \quad (u) \\ & -y \le 1 \quad (v) \\ & x, y \ge 0 \end{cases}
\qquad
\textbf{(D)}\ \begin{cases} \underset{u,v\in\mathbb{R}}{\text{minimize}} & -u+v \\ \text{subject to} & u \ge 1 \quad (x) \\ & -v \ge 1 \quad (y) \\ & u, v \ge 0 \end{cases}
$$

Here, the primal is infeasible because we can't choose a value of x, and the dual is infeasible because we can't choose a value of v.