Math 3272: Linear Programming[1] Mikhail Lavrov

Lecture 2: Constraints in linear programs

August 18, 2022 *Kennesaw State University*

1 What can a linear program model?

Linear programs are almost always a simplification: real life is nonlinear a lot of the time. Sometimes we're lucky and our constraints do end up being linear. Sometimes we're slightly less lucky, and can still approximate real life by a linear program. Sometimes we're not lucky at all.

For example, suppose we're optimizing over the disk $x^2 + y^2 \le 1$. That's not a linear constraint. But we can replace the circle by a polygon with many sides. Each side is a straight line, so we can describe the polygon by a bunch of linear inequalities. Probably, optimizing over the polygon will not be too different from optimizing over the circle—and if not, we can give the polygon more sides to improve the approximation.

Similarly, *strict* inequalities like $x + y < 1$ are not okay in our linear programs, but also not a huge problem. We can always replace such an inequality by either $x + y \le 1$ (including slightly more points) or $x + y \le 0.999$ (including slightly fewer points). The second approximation can be made arbitrarily good.

On the other hand, suppose our region is the union of two disks:

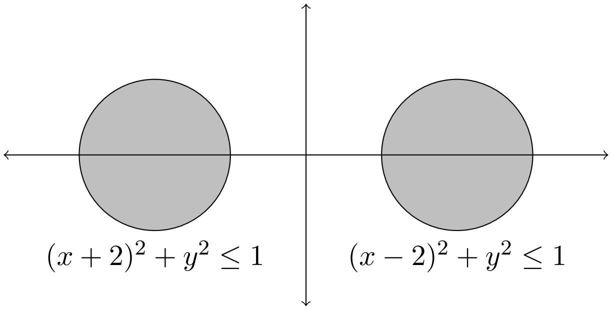

No matter how you try, you can never draw a linear inequality that includes both of these disks, but *excludes* the origin, $(0, 0)$. Here's a formal proof. Suppose you have any system of inequalities $A\mathbf{x} \le \mathbf{b}$ that includes both disks. Then in particular it includes the points $(-2, 0)$ and $(2, 0)$ at their centers. So

$$\begin{aligned} A\begin{bmatrix}-2\\0\end{bmatrix} \le \mathbf{b} \text{ and } A\begin{bmatrix}2\\0\end{bmatrix} \le \mathbf{b} &\implies A\begin{bmatrix}-2\\0\end{bmatrix} + A\begin{bmatrix}2\\0\end{bmatrix} \le 2\mathbf{b} \\ &\implies A\begin{bmatrix}0\\0\end{bmatrix} \le 2\mathbf{b} \\ &\implies A\begin{bmatrix}0\\0\end{bmatrix} \le \mathbf{b}. \end{aligned}$$

[1] This document comes from an archive of the Math 3272 course webpage: http://misha.fish/archive/3272-fall-2022

Therefore $(0,0)$ satisfies the system of inequalities as well.

In other words, this region can't be approximated by a linear program. No matter what, you will never exclude the point $(0,0)$, which is pretty far from any point that's actually in your region.

There's a generalization of this idea. We call a subset S of $\mathbb{R}^n$ a **convex set** if, whenever $\mathbf{x}, \mathbf{y} \in S$, the entire line segment joining $\mathbf{x}$ and $\mathbf{y}$ is also contained in S. Algebraically, the line segment joining $\mathbf{x}$ and $\mathbf{y}$ can be described as

$$[\mathbf{x}, \mathbf{y}] = \{t\mathbf{x} + (1-t)\mathbf{y} : 0 \le t \le 1\}$$

and so we can also state this definition as "whenever $\mathbf{x}, \mathbf{y} \in S$ and $0 \le t \le 1$, $t\mathbf{x}+(1-t)\mathbf{y} \in S$."

The feasible region of a linear program is always convex. We can check this by an algebraic proof: if $A\mathbf{x} \le \mathbf{b}$ and $A\mathbf{y} \le \mathbf{b}$, then

$$A(t\mathbf{x} + (1-t)\mathbf{y}) = t(A\mathbf{x}) + (1-t)(A\mathbf{y}) \le t\mathbf{b} + (1-t)\mathbf{b} = \mathbf{b}.$$

There is also an argument from geometric intuition. If $\mathbf{x}$ and $\mathbf{y}$ satisfy an linear inequality, this means that they both fall on one side of a straight line. Then the entire line segment $[\mathbf{x}, \mathbf{y}]$ must be on the same side of that line, so it also satisfies that linear inequality. The same is true for a system of inequalities: we just consider the inequalities one at a time.

It turns out (though it's harder to prove) that any convex set can be approximated as well as you like by enough linear inequalities. If the set is bounded by straight lines (or higher-dimensional surfaces), you can even describe it exactly. On the other hand, if a set is *not* convex, there's no hope to even get close.

2 Different formulations of linear programs

We've talked already about expressing the constraints of a linear program as a system of inequalities $A\mathbf{x} \le \mathbf{b}$. There are several variations, and we can convert linear programs from one form to the other.

2.1 Nonnegativity constraints

It's common to automatically include the *nonnegativity constraints* $x_1, x_2, \dots, x_n \ge 0$. There are several reasons for this:

- Lots of real-world problems already include them. (Many actual quantities can't be negative.)
- Mathematically they are fairly nice. (We'll see some ways this comes up later.)
- In the previous lecture, we saw that if a linear program has any optimal solutions, we can always find one at a vertex. There's one exception to this: some linear programs don't have *any* vertices (for example, if there's only one inequality, or if the feasible region looks like an infinite prism in three or more dimensions).

  When nonnegativity constraints are present, this case is guaranteed not to happen.

Our textbook says that a linear program in the form

$$\begin{array}{ll} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} & \mathbf{c}^{\mathsf{T}}\mathbf{x} \\ \text{subject to} & A\mathbf{x} \le \mathbf{b} \\ & \mathbf{x} \ge 0 \end{array} \qquad \text{or} \qquad \begin{array}{rrcrcrcl} \underset{x_1,\dots,x_n\in\mathbb{R}}{\text{maximize}} & c_1x_1 & + & c_2x_2 & + \cdots + & c_nx_n & & \\ \text{subject to} & a_{11}x_1 & + & a_{12}x_2 & + \cdots + & a_{1n}x_n & \le & b_1 \\ & a_{21}x_1 & + & a_{22}x_2 & + \cdots + & a_{2n}x_n & \le & b_2 \\ & \vdots & & \vdots & \ddots & \vdots & & \vdots \\ & a_{m1}x_1 & + & a_{m2}x_2 & + \cdots + & a_{mn}x_n & \le & b_m \\ & & & & & x_1, x_2, \dots, x_n & \ge & 0 \end{array}$$

is in **standard form**. I won't ask you to learn this terminology, because different sources have different ideas of standard form, and it's not a very self-explanatory term, but sometimes it will be useful to put linear programs in this form.

What if there are no nonnegativity constraints? We can introduce them by a standard trick: whenever a variable x can be positive or negative, replace it (everywhere it occurs) by the difference $x^+ - x^-$, where x^+ and x^- are two variables with $x^+, x^- \ge 0$. Any real number can be written as the difference of two nonnegative numbers.

For instance, the example yesterday can be rewritten as a linear program in four nonnegative variables instead of two unconstrained variables:

$$\begin{array}{lrrl} \underset{x,y\in\mathbb{R}}{\text{maximize}} & \multicolumn{3}{l}{1000x - 300y} \\ \text{subject to} & x + & y & \le 50 \\ & & -y & \le 20 \\ & 2x - & y & \le 40 \end{array} \quad \rightsquigarrow \quad \begin{array}{lrrrrl} \underset{x^+,x^-,y^+,y^-\in\mathbb{R}}{\text{maximize}} & \multicolumn{5}{l}{1000x^+ - 1000x^- - 300y^+ + 300y^-} \\ \text{subject to} & x^+ - & x^- + & y^+ - & y^- & \le 50 \\ & & & -y^+ + & y^- & \le 20 \\ & 2x^+ - & 2x^- - & y^+ + & y^- & \le 40 \\ & & & \multicolumn{2}{r}{x^+, x^-, y^+, y^-} & \ge \ 0 \end{array}$$

This tends to create infinitely many solutions, because each point (x, y) has infinitely many representations (x^+, x^-, y^+, y^-). For example, the point $(x, y) = (20, -10)$ can be represented by $(x^+, x^-, y^+, y^-) = (20, 0, 0, -10)$, but equally valid is $(x^+, x^-, y^+, y^-) = (20, 0, 100, -110)$. That's okay: as long as there's at least one solution, we're happy.

There's a more compact but less intuitive way to do the same thing: replace n unconstrained variables $(x_1, x_2, \dots, x_n)$ by just $n+1$ nonnegative variables $(x_0', x_1', \dots, x_n')$, and replace every occurrence of x_i by the difference $x_i' - x_0'$. See if you can convince yourself that this works!

2.2 Equations and inequalities

Nonnegativity constraints are the simplest kind of inequality, and so you might wish: what if those were the only kinds of inequalities we had to deal with? This is possible, and the resulting form of the linear program is sometimes called **equational form**. (I will use this term more freely, because it's at least self-explanatory.)

The idea is this: if we have an inequality

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n \le b,$$

we can rewrite it as an equation with one new nonnegative variable: for some $w \ge 0$,

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n + w = b.$$

This new nonnegative variable w is called a **slack variable**, because it measures how much "slack" or flexibility there was in satisfying the constraint. When $w = 0$, the constraint is tight: we are right on the edge of the linear inequality. When $w > 0$, we have some wiggle room to change $x_1, \dots, x_n$ without running up against this constraint.

Doing this for every single constraint in a linear program turns every inequality into an equation, except for some nonnegativity constraints. If we pick up where we left off and do this to our linear program, we get:

$$\begin{array}{lrrrrrrrl}
\underset{x^+,x^-,y^+,y^-,w_1,w_2,w_3\in\mathbb{R}}{\text{maximize}} & \multicolumn{8}{l}{1000x^+ - 1000x^- - 300y^+ + 300y^-} \\
\text{subject to} & x^+ - & x^- + & y^+ - & y^- + w_1 & & & = & 50 \\
 & & & -y^+ + & y^- & + w_2 & & = & 20 \\
 & 2x^+ - & 2x^- - & y^+ + & y^- & & + w_3 & = & 40 \\
 & \multicolumn{6}{r}{x^+, x^-, y^+, y^-, w_1, w_2, w_3} & \geq & 0
\end{array}$$

A general linear program in equational form looks like

$$\begin{array}{ll}
\underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} & \mathbf{c}^\mathsf{T}\mathbf{x} \\
\text{subject to} & A\mathbf{x} = \mathbf{b} \\
 & \mathbf{x} \geq \mathbf{0}
\end{array}
\qquad \text{or} \qquad
\begin{array}{lrrrrrl}
\underset{x_1,\dots,x_n\in\mathbb{R}}{\text{maximize}} & c_1x_1 + & c_2x_2 + & \cdots + & c_nx_n & & \\
\text{subject to} & a_{11}x_1 + & a_{12}x_2 + & \cdots + & a_{1n}x_n & = & b_1 \\
 & a_{21}x_1 + & a_{22}x_2 + & \cdots + & a_{2n}x_n & = & b_2 \\
 & \vdots & \vdots & \ddots & \vdots & & \vdots \\
 & a_{m1}x_1 + & a_{m2}x_2 + & \cdots + & a_{mn}x_n & = & b_m \\
 & \multicolumn{4}{r}{x_1, x_2, \dots, x_n} & \geq & 0
\end{array}$$

This is convenient to deal with, because linear algebra gives us a lot of tools for understanding the system of equations $A\mathbf{x} = \mathbf{b}$. We just need to figure out what happens when we also require $\mathbf{x} \geq \mathbf{0}$.

(In particular, this is the form of linear program that the simplex method will use: this method is built on top of Gaussian elimination for solving the system of equations $A\mathbf{x} = \mathbf{b}$.)

In some cases, we want to go the other way: we want to turn equations into inequalities. This is also possible. To write down an equation

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b$$

simply write down two inequalities, one in each direction:

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n \leq b \qquad \text{and} \qquad a_1x_1 + a_2x_2 + \cdots + a_nx_n \geq b$$

This means we can express any linear program using only inequalities, and no equations at all.