Math 482: Linear Programming[1] Mikhail Lavrov

Lecture 2: Constraints in Linear Programs

January 24, 2020 *University of Illinois at Urbana-Champaign*

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.

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:

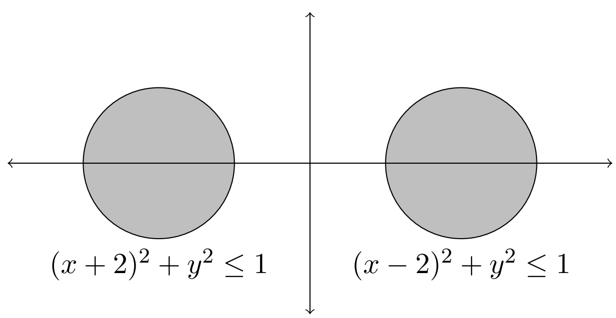

No matter how you try, you can never draw a linear inequality that includes both of these disks, but *excludes* the origin, $(0, 0)$. Here's a formal proof. Suppose you have any system of inequalities $A\mathbf{x} \le \mathbf{b}$ that includes both disks. Then in particular it includes the points $(-2,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

$$\begin{aligned} A\begin{bmatrix}-2\\0\end{bmatrix} \le \mathbf{b} \text{ and } A\begin{bmatrix}2\\0\end{bmatrix} \le \mathbf{b} &\implies A\begin{bmatrix}-2\\0\end{bmatrix} + A\begin{bmatrix}2\\0\end{bmatrix} \le 2\mathbf{b} \\ &\implies A\begin{bmatrix}0\\0\end{bmatrix} \le 2\mathbf{b} && \text{(in other words, } 2\mathbf{b} \ge \mathbf{0}\text{)} \\ &\implies A\begin{bmatrix}0\\0\end{bmatrix} \le \mathbf{b} && \text{(in other words, } \mathbf{b} \ge \mathbf{0}\text{).} \end{aligned}$$

[1] This document comes from the Math 482 course webpage: https://faculty.math.illinois.edu/~mlavrov/courses/482-spring-2020.html

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* $\mathbf{x} \ge \mathbf{0}$. 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.

Some of the textbooks we use say that a linear program in the form

$$\begin{array}{ll} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} & \mathbf{c}^\mathsf{T}\mathbf{x} \\ \text{subject to} & A\mathbf{x} \le \mathbf{b}, \\ & \mathbf{x} \ge 0 \end{array}$$

is in "standard" or "canonical" form. I won't ask you to learn the terminology for various forms, but it will sometimes be convenient to assume that a program has this structure.

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:

$$\begin{array}{ll} \underset{x,y\in\mathbb{R}}{\text{maximize}} & x - y \\ \text{subject to} & y \le 3, \\ & y \ge 2x - 5, \\ & x + y \ge 1 \end{array}$$

can be put in standard form as

$$\begin{array}{lrcr} \underset{x^+,x^-,y^+,y^-\in\mathbb{R}}{\text{maximize}} & \multicolumn{3}{l}{x^+ - x^- - y^+ + y^-} \\ \text{subject to} & y^+ - y^- & \le & 3, \\ & 2x^+ - 2x^- - y^+ + y^- & \le & 5, \\ & -x^+ + x^- - y^+ + y^- & \le & -1, \\ & \multicolumn{3}{l}{x^+, x^-, y^+, y^- \ge 0.} \end{array}$$

This tends to create infinitely many solutions: if the optimal solution was $(x, y) = (2, -1)$ before, then the simplest way to extend it to four variables is $(x^+, x^-, y^+, y^-) = (2, 0, 0, 1)$, but equally valid is $(x^+, x^-, y^+, y^-) = (4, 2, 2, 3)$. That's okay: as long as there's one solution, we're happy.

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".

The idea is this: if we have an inequality $\mathbf{a}^\mathsf{T}\mathbf{x} \le b$, we can rewrite it as an equation: $\mathbf{a}^\mathsf{T}\mathbf{x} + s = \mathbf{b}$, for some $s \ge 0$. This s is called a *slack variable*, because it measures how much "slack" or flexibility there was in satisfying the inequality constraint. Doing this for every single constraint in a linear program turns every inequality into an equation, except for some nonnegativity constraints. Picking up where we left off:

$$\begin{array}{ll} \underset{x^+,x^-,y^+,y^-,s_1,s_2,s_3\in\mathbb{R}}{\text{maximize}} & x^+ - x^- - y^+ + y^- \\ \text{subject to} & \begin{aligned} y^+ - y^- + s_1 &= 3, \\ 2x^+ - 2x^- - y^+ + y^- + s_2 &= 5, \\ -x^+ + x^- - y^+ + y^- + s_3 &= -1, \end{aligned} \\ & x^+, x^-, y^+, y^-, s_1, s_2, s_3 \ge 0. \end{array}$$

A general linear program in equational form looks like

$$\begin{array}{ll} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} & \mathbf{c}^\mathsf{T}\mathbf{x} \\ \text{subject to} & A\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \ge \mathbf{0}. \end{array}$$

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 $\mathbf{x} \ge \mathbf{0}$.

(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 $A\mathbf{x} = \mathbf{b}$.)

If linear inequalities are better than linear equations, we can always go the other way. The equation $\mathbf{a}^\mathsf{T}\mathbf{x} = b$ is a combination of two inequalities: $\mathbf{a}^\mathsf{T}\mathbf{x} \le b$, and $\mathbf{a}^\mathsf{T}\mathbf{x} \ge b$. (In particular, we can always express a linear program using only inequalities, and no equations at all.)