Math 3272: Linear Programming[1] Mikhail Lavrov

Lecture 4: The simplex method

August 25, 2022 *Kennesaw State University*

0 The plan

The simplex method can be roughly summarized as "go from one solution to another, improving every time, until you reach the best solution". We'll get there in two steps.

Today, we will talk about how we go from one solution to another. We will only think about the constraints of our linear program, and not even consider the objective function.

In the next lecture, we'll go back and think about which steps bring us closer to our goal, and which steps take us further away from it.

1 From linear algebra back to linear programming

The simplex method works on linear programs in equational form: the constraints are $A\mathbf{x} = \mathbf{b}$ with $\mathbf{x} \geq \mathbf{0}$. Written out in full:

$$\begin{cases} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n = b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n = b_2 \\ \quad \vdots \qquad\quad \vdots \qquad \ddots \qquad \vdots \qquad\quad \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n = b_m \\ \qquad\qquad\qquad x_1, x_2, \dots, x_n \geq 0 \end{cases}$$

That is, we have a perfectly ordinary system of linear equations, together with the added constraint that all variables must be nonnegative.

There are infinitely many feasible solutions, but on the first day, we saw a rule that cuts their number down to a manageable amount:

Rule #1: At least one optimal solution is a corner point of the feasible region.[2]

We understand what a corner point is geometrically, in two dimensions: it's a point where two of the boundary lines meet. Visualizing the same thing in higher dimensions is tricky, but let's try it anyway.

Suppose we have n variables $x_1, \dots, x_n$ and the system $A\mathbf{x} = \mathbf{b}$ consists of m linear equations, none of which are redundant. The solutions to this system live in $\mathbb{R}^n$. However, each linear equation reduces the dimension of the solution set by 1, so the solution set is an *affine subspace* of dimension

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

[2] Terms and conditions apply. Void if the linear program doesn't have an optimal solution. Also void if the feasible region doesn't have any vertices.

$n-m$. (An "affine subspace" is a subspace that has been shifted so it doesn't necessarily pass through the origin. "Dimension $n-m$" means it looks like $\mathbb{R}^{n-m}$. For example, when $n=3$ and $m=2$, the points live in $\mathbb{R}^3$, but the solutions to $A\mathbf{x}=\mathbf{b}$ look like $\mathbb{R}^1$: they are a line in 3-dimensional space.)

In two dimensions, a corner point is where two boundaries meet. In three dimensions, a corner is where three boundaries meet (imagine the corner of a cube). In $n-m$ dimensions, a corner is where $n-m$ boundaries meet. What are the boundaries of our feasible region? They come from the inequalities $x_1, x_2, \dots, x_n \geq 0$. When $n-m$ boundaries meet, it is because $n-m$ of our variables have been set to 0.

If that was intimidating—well, we have another way to think about the same thing. When solving a system of m linear equations in n variables, we pick m basic variables: one for each equation. Then, we solve for them in terms of the $n-m$ nonbasic variables. A basic solution is what we get if we set all $n-m$ nonbasic variables to 0: exactly the number that we wanted for a corner point!

In other words, we can deduce the following rule:

Rule #2: All corner points of the feasible region are basic solutions of the system of linear equations.

This gives a motivation to find as many basic solutions as possible.

2 An example of pivoting in the simplex method

In keeping with our intention to think about constraints only, let's pose half a problem: a set of constraints without an objective.

Problem 1. *You are trying to plan out a diet consisting entirely of french fries and ketchup. Your research says that the following conditions are required for a healthy diet:*[3]

1. *You need to eat at least* 10 *servings of food to avoid being hungry.*
2. *With 210 calories per serving of fries and 20 calories per serving of ketchup, you want to limit your intake to 2000 calories.*
3. *With* 0.1 *grams of sodium per serving of fries and* 0.2 *grams per serving of ketchup, you want to consume at most* 3 *grams of sodium.*

With x servings of fries and y servings of ketchup, the constraints are shown below on the left:

$$\begin{cases} x + y \geq 10 \\ 210x + 20y \leq 2000 \\ 0.1x + 0.2y \leq 3 \\ x, y \geq 0 \end{cases} \quad \rightsquigarrow \quad \begin{cases} x + y - w_1 = 10 \\ 210x + 20y + w_2 = 2000 \\ 0.1x + 0.2y + w_3 = 3 \\ x, y, w_1, w_2, w_3 \geq 0 \end{cases}$$

We can begin by practicing turning these into equations. Add a slack variable to each inequality, and we get the equations above on the right.

[3] Not medical advice.

2.1 Step 1: a basic solution

If all we want is a basic solution, that's easy to find—and there's a generally useful strategy for how to do it. Just solve each equation for its slack variable: the first equation for w_1, the second equation for w_2, and the third equation for w_3. Then our system can be rewritten as

$$\begin{cases} w_1 = -10 + x + y \\ w_2 = 2000 - 210x - 20y \\ w_3 = 3 - 0.1x - 0.2y \end{cases}$$

where the nonnegativity conditions $x, y, w_1, w_2, w_3 \geq 0$ still hold, but I'll stop writing them every time. To find a basic solution, set the nonbasic variables x, y to 0, and read off the values of the basic variables w_1, w_2, w_3.

Is this one of the corner points? No! When $x = y = 0$, we get $w_1 = -10$, $w_2 = 2000$, and $w_3 = 0$. These are *not* all nonnegative. We should have expected this: setting $x = y = 0$ means you're not eating anything, so you're violating the constraint "eat at least 10 servings".

A corner point must be a basic solution, but a corner point must also be feasible: all the variables must be nonnegative. We are looking for a **basic feasible solution**: you will hear these words a lot this semester. This term (sometimes cryptically abbreviated **bfs**) is just the sum of its parts: a feasible solution which is also a basic solution.

We won't get anywhere with an infeasible solution, so let's start from scratch.

2.2 Step 1, again: a basic feasible solution

In general, finding any starting basic feasible solution can be tricky, and we'll return to the hard cases of the problem later. Today, I will just give a set of basic variables that works: y, w_2, w_3. This basic feasible solution will correspond to the strategy "Eat enough ketchup to satisfy your hunger".

Starting from our first set of equations, we can do the row reduction to solve for y, w_2, w_3 in terms of x, w_1. If you want more practice with this, you can try this yourself and check your work; you should get

$$\begin{cases} y = 10 - x + w_1 \\ w_2 = 1800 - 190x - 20w_1 \\ w_3 = 1 + 0.1x - 0.2w_1 \end{cases}$$

Setting $x = w_1 = 0$ gives us $y = 10$, $w_2 = 1800$, and $w_3 = 1$: no arithmetic is required, you can just read those off from a column in the system of equations above. These are all positive, so $(x, y, w_1, w_2, w_3) = (0, 10, 0, 1800, 1)$ is our first basic feasible solution!

2.3 Step 2: pivoting (intuitively)

Right now, we don't have an objective, so we don't have a reason to get more basic feasible solutions. But let's see how we do it anyway.

The simplex method, which we'll finish learning in the next lecture, works by a strategy called **pivoting**. The idea is that:

1. We start with a basic feasible solution.
2. We modify it slightly to make one nonbasic variable become basic (**enter the basis**). One of the basic variables will have to make room and become nonbasic (**leave the basis**).

   If x_i is the entering variable, we call this **pivoting around** x_i.
3. We choose the leaving variable to avoid negative signs, so that we arrive at a new basic feasible solution.

Any nonbasic variable can be chosen to enter the basis; as an example, we'll make w_1 our entering variable, starting from our previous basic feasible solution. The intuition is this: keeping our other nonbasic variables at 0, we try to increase w_1 as much as we can without breaking anything!

What can break? Well, let's look at the equations in the previous step, one at a time.

- We have $y = 10 - x + w_1$, so when $x = w_1 = 0$, we get $y = 10$. Increasing w_1 from this point will increase y at the same rate. When $w_1 = 1$, we get $y = 11$; when $w_2 = 2$, we get $y = 12$; when $w_2 = 100$, we get $y = 110$. We can keep going forever, and *this* equation will be just fine.
- We have $w_2 = 1800 - 190x - 20w_1$, so when $x = w_1 = 0$, we get $w_2 = 1800$. Increasing w_1 from here will *decrease* w_2 by 20 units per increase in w_1. This could cause a problem: we don't want to make w_2 negative. Since w_2 drops to 0 when $w_1 = \frac{1800}{20} = 90$, we want to keep $w_1 \le 90$.
- We have $w_3 = 1 + 0.1x - 0.2w_1$, so when $x = w_1 = 0$, we get $w_3 = 1$. Increasing w_1 from here will *decrease* w_3 by 0.2 units per increase in w_1. Again, we want to keep w_3 nonnegative. How far can we go? w_3 drops to 0 when $w_1 = \frac{1}{0.2} = 5$, so we want to keep $w_1 \le 5$.

So to increase w_1 as much as possible, we set it to 5, driving w_3 down to 0. This tells us which variable should leave the basis: w_3 will become a nonbasic variable, since the nonbasic variables are the ones that are set to 0 in a basic solution.

This means we want to solve for y, w_2, w_1 on terms of x, w_3. We've already seen that this can be done from our previous set of equations, saving some effort.

First, divide the last equation by 0.2, so that the coefficient of w_1 is -1. Row-reduce: add the third equation to the first, and subtract 20 times the third equation from the second. Finally, move w_3 to the right and w_1 to the left.

$$\begin{cases} y = 10 - x + w_1 \\ w_2 = 1800 - 190x - 20w_1 \\ 5w_3 = 5 + 0.5x - w_1 \end{cases} \quad \rightsquigarrow \quad \begin{cases} y + 5w_3 = 15 - 0.5x \\ w_2 - 100w_3 = 1700 - 200x \\ 5w_3 = 5 + 0.5x - w_1 \end{cases}$$

$$\rightsquigarrow \quad \begin{cases} y = 15 - 0.5x - 5w_3 \\ w_2 = 1700 - 200x + 100w_3 \\ w_1 = 5 + 0.5x - 5w_3 \end{cases}$$

We can read off our new basic feasible solution from here: $(x, y, w_1, w_2, w_3) = (0, 15, 5, 1700, 0)$.(This is the "eat as much ketchup as you can without having too much sodium" strategy.)

2.4 Step 3: pivoting (algebraically)

Let's try to add some french fries to our diet and pivot around x, making it a basic variable. Which variable should leave the basis? This time, let's try to take our experience for the previous pivot, and come up with rules to follow to make this decision.

1. The leaving variable is the first one that will be driven to 0 as x increases. For this to happen at all, it should increase as x increases. Therefore:

   In the leaving variable's equation, the coefficient of x should be negative.

   In this example, we are choosing between y and w_2.

2. The leaving variable is the *first* one that will be driven to 0 as x increases. At which value of x will it get to 0? Solving $15 - 0.5x = 0$, we divide 15 (the current value of y) by 0.5 (the negative coefficient of x: the rate at which y decreases as x increases). So the rule is:

   From these options, pick the variable with the least value of $\frac{\textbf{current value}}{\textbf{rate of decrease}}$ to be the leaving variable.

   Here, y's ratio is $\frac{15}{0.5} = 30$ and w_2's ratio is $\frac{1700}{200} = 8.5$, so we pick w_2.

These are the rules the simplex method always follows! (With x replaced by whatever the entering variable is, of course.)

Pivoting as before, we get our new set of equations:

$$\begin{cases} y = 15 - 0.5x - 5w_3 \\ \frac{1}{200}w_2 = 8.5 - x + 0.5w_3 \\ w_1 = 5 + 0.5x - 5w_3 \end{cases} \quad \rightsquigarrow \quad \begin{cases} y - \frac{1}{400}w_2 = 10.75 - 5.25w_3 \\ \frac{1}{200}w_2 = 8.5 - x + 0.5w_3 \\ w_1 + \frac{1}{400}w_2 = 9.25 - 4.75w_3 \end{cases}$$

$$\rightsquigarrow \quad \begin{cases} y = 10.75 + \frac{1}{400}w_2 - 5.25w_3 \\ x = 8.5 - \frac{1}{200}w_2 + 0.5w_3 \\ w_1 = 9.25 - \frac{1}{400}w_2 - 4.75w_3 \end{cases}$$

Our new basic feasible solution is $(x, y, w_1, w_2, w_3) = (8.5, 10.75, 9.25, 0, 0)$.

This was just aimless wandering around; in the next lecture, we'll reintroduce the objective function, and think about pivoting with purpose. Think of what we've done today as driving around the parking lot; next, we'll get on the highway.

2.5 Troubleshooting

The only goal of the pivoting algorithm we learned today is to go from a basic feasible solution to another basic feasible solution. You know that you've picked the correct leaving variable if your new basic solution is still feasible—if it's not, then go back and rethink your choice of leaving variable.

Aside from that, remember the cardinal rule: always do the same thing to both sides of an equation. Finally, watch out for mistakes with lost negative signs, as those are very easy to make here.