LECTURE 7: SIMPLEX ALGORITHM (III)

Today: Why the simplex method works and what can go wrong.

1. Why the simplex method works

How come the simplex method gives us *the* optimal vertex?

Main Issue: The simplex method walks from vertex to vertex, and checks if a vertex is better than its neighbors.

Example: In the example from last time, $(0,3)$ is better than its neighbor $(0,0)$, and $(1,4)$ is better than its neighbors $(0,3)$ and $(3.4,2.8)$

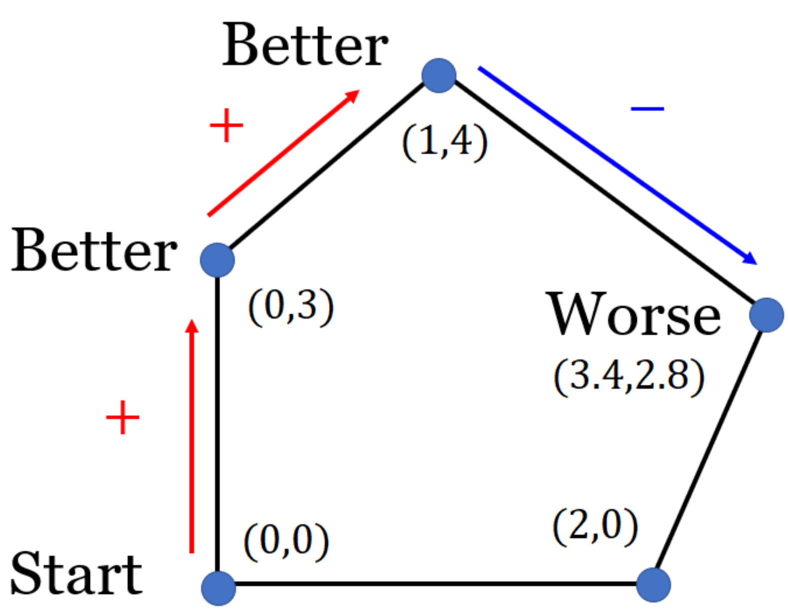

Date: Thursday, September 29, 2022.

So the simplex method gives us a **local maximum**, a vertex than is better than its neighbors

But does it give you a **global maximum**, meaning the biggest value of z in the feasible region?

In general, the answer is **NO**:

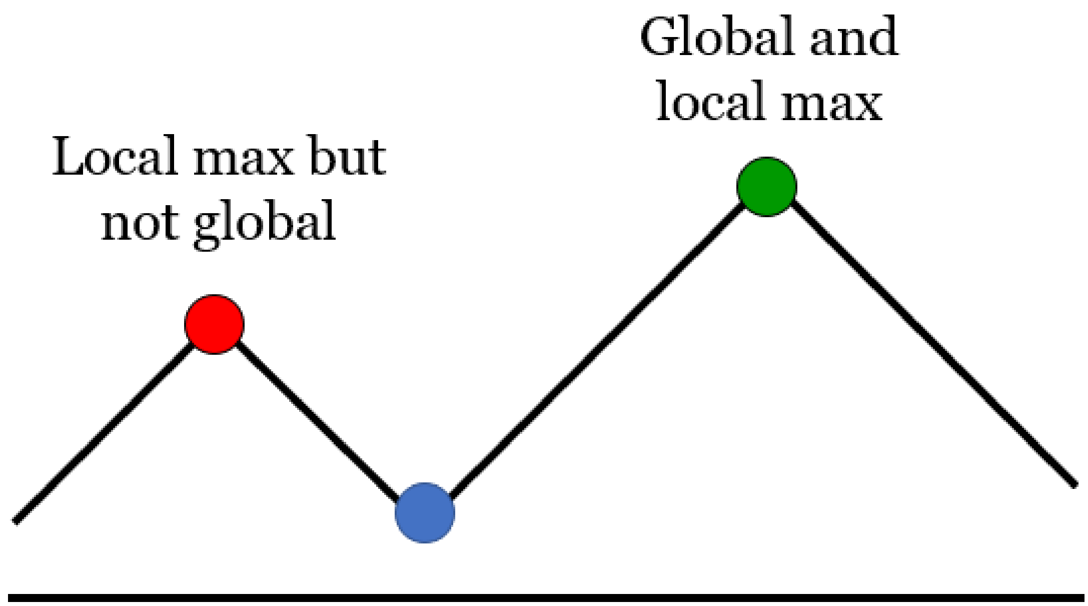

Here the vertex on the left is a local max (better than its neighbors), but not a global max, it does not give us the biggest value of z.

What saves us is… **convexity** !!

Definition:

A function $f : S \to \mathbb{R}$ is **convex** if for all $x, y \in S$ and all $0 \leq \lambda \leq 1$, we have

$$f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda) f(y)$$

Interpretation: f lies below the segment connecting $f(x)$ and $f(y)$

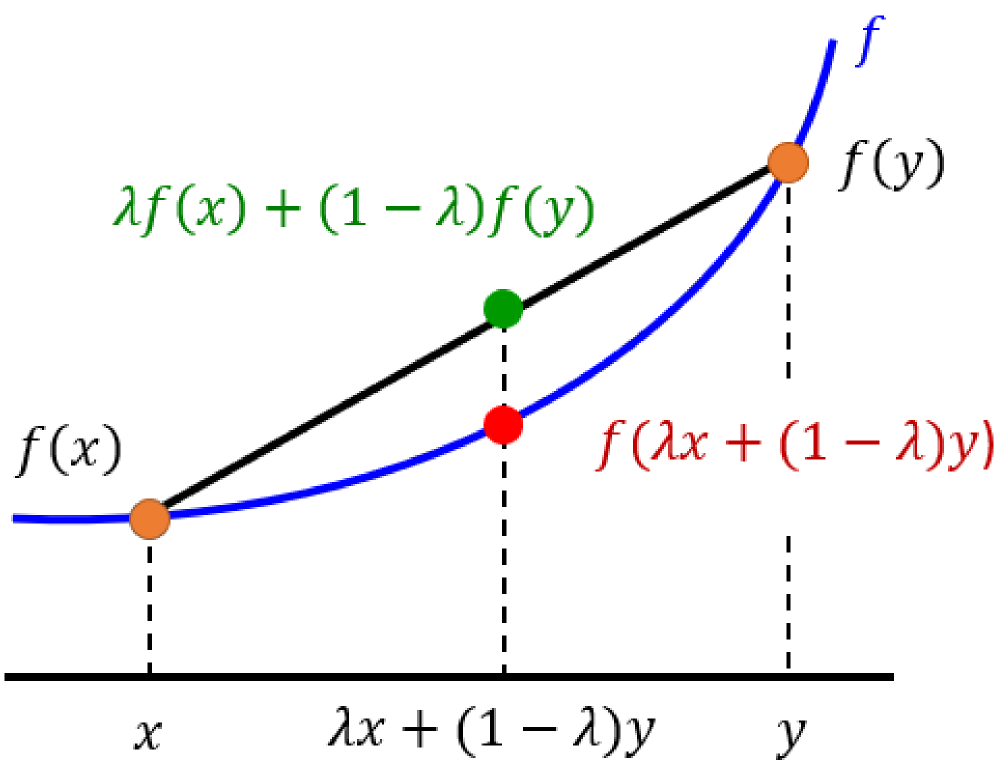

Typical examples include $f(x) = |x|$ or $f(x) = x^2$ or $f(x) = e^x$. In calculus, convex is the same as *concave up*.

What makes the simplex method work is the following

Fact:

Any local min of a **convex** function is automatically a global min

You will prove this on the homework, but intuitively it's because if you're at a local min (for example) and f "bends upwards," this forces the local minimum to become global.

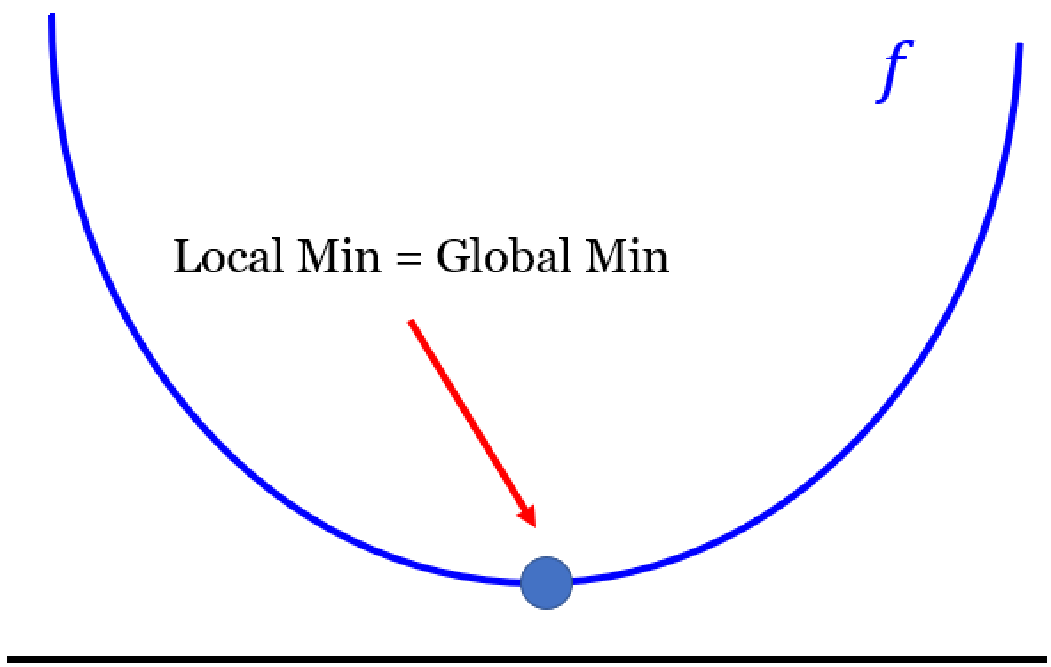

Luckily for LP problems, our objective function is always convex:

Fact:

$$z = c^T \mathbf{x} \text{ is convex}$$

(To prove this, just use the definition of convex function)

This tells you that any local min is a global min, but then use

$$\max z = -\min(-z)$$

And $-z = -c^T\mathbf{x}$ is convex to get the statement about global max.

2. Starting Vertex

So far in our simplex method, we got lucky because we were always easily able to find a starting vertex, namely the origin. But what if it's not obvious where to start?

Luckily the following trick helps us figure out a starting vertex:

Example:

$$\begin{aligned} \max\ z =& 2x_1 - 3x_2 + 5x_3 \\ \text{subject to } & x_1 - x_2 + 3x_3 = 5 \\ & x_2 + 4x_3 \leq 4 \\ & x_1, x_2, x_3 \geq 0 \end{aligned}$$

Note: Here $m = 2$, which is the number of constraints other than $x_i \geq 0$

Trick: Introduce $m = 2$ new 'artificial' variables $z_1, z_2, \cdots, z_m$ with and solve the **new** LP:

$$\begin{aligned} \min\ & z_1 + \cdots + z_m \\ \text{subject to } & x_1 - x_2 + 3x_3 + z_1 = 5 \\ & x_2 + 4x_3 + z_2 \leq 4 \\ & x_i \geq 0 \\ & z_i \geq 0 \end{aligned}$$

Then simply set

$$\begin{aligned} x_1 &= 0 \\ x_2 &= 0 \\ x_3 &= 0 \\ z_1 &= 5 \\ z_2 &= 4 \end{aligned}$$

Then $(0, 0, 0, 5, 4)$ is a **basic feasible solution**, that is:

(1) $(0, 0, 0, 5, 4)$ is in the feasible region

(2) $m+n = 2+3 = 5$ (lin ind) constraints are satisfied with equality at $(0, 0, 0, 5, 4)$

So *by definition*, $(0, 0, 0, 5, 4)$ is a vertex ✓

In other words, our new LP problem *has* a vertex.

Then two things can happen

Case 1: The optimal solution for the new LP is $z = 0$.

This means there is a vertex $(x_1, x_2, x_3, z_1, z_2)$ such that

$$z = z_1 + z_2 = 0$$

But since $z_1, z_2 \geq 0$ this implies $z_1 = 0$ and $z_2 = 0$, so the vertex is $(x_1, x_2, x_3, z_1, z_2)$, and so (x_1, x_2, x_3) is a vertex of the original problem (since it's basic feasible) and we can use that vertex as our starting point to solve that original LP problem. ✓

For example, if you find that vertex is $(1, 7, 4, 0, 0)$ then your starting vertex is $(1, 7, 4)$

Case 2: The optimal solution is $z > 0$. Then $z = z_1 + z_2 > 0$, so either $z_1 > 0$ or $z_2 > 0$. But this means that the original problem is infeasible!

Why? Suppose for instance that $(6, 1, 0)$ is in the feasible region of the original problem, then $(6, 1, 0, 0, 0, 0)$ would be in the feasible region of the new problem, and in that case at that point, $z = z_1 + z_2 = 0 + 0 = 0$, so there *is* a point where $z = 0$ which contradicts the fact that $\min z$ is positive $\Rightarrow\Leftarrow$

To summarize: To find a starting vertex, first write the new LP and solve it. If you find $z = 0$, then you get a starting vertex for the original LP. If you find $z > 0$, then the original LP is unsolvable and you stop.

3. Degeneracy

Now let's talk about potential problems that might arise with the simplex method. One problem is not so much an issue of the simplex method, but more of the computer implementation.

Consider the following scenario, where our constraints intersect at a single point:

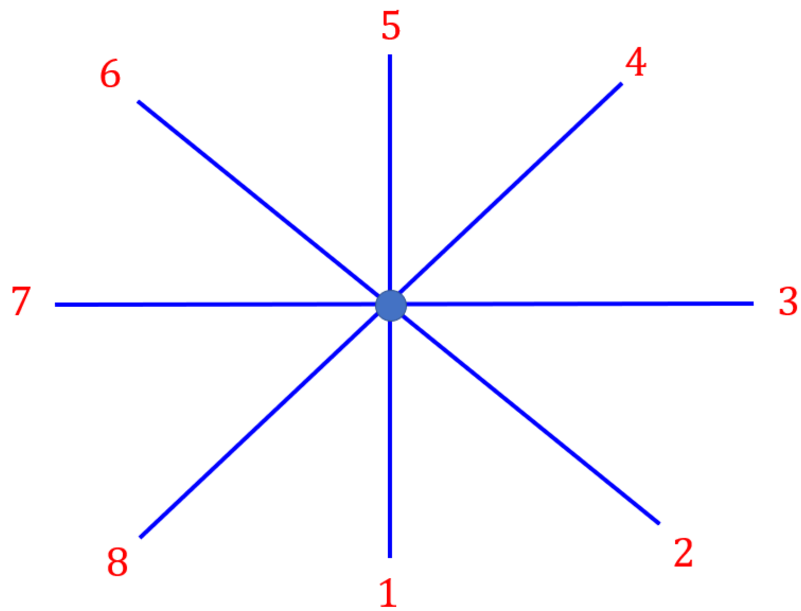

What could happen is that the algorithm starts with constraints ①+② being tight, then releases ① to make ② + ③ tight, then releases ② to make ③ + ④ tight ... then releases ⑧ to make ① + ② tight, which results in an infinite loop!!

To get around this, one possible solution is to move one constraint just a little bit, as in this figure

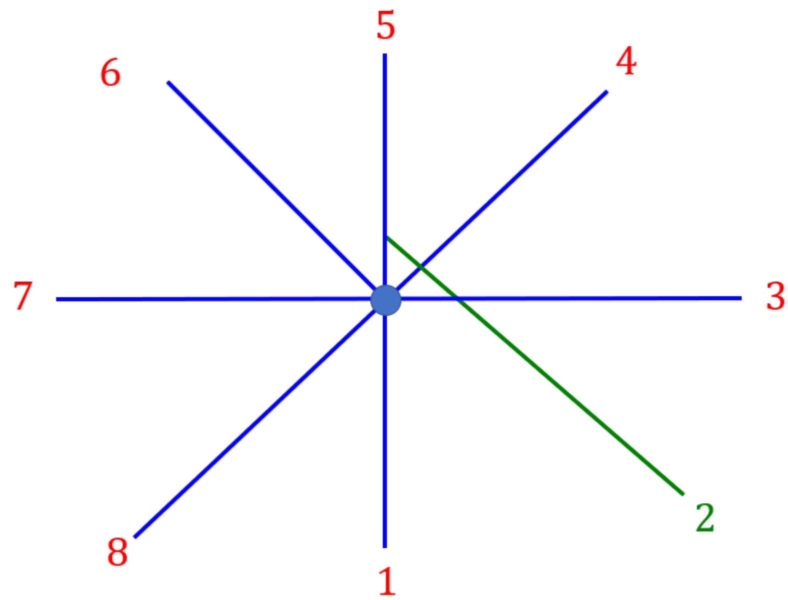

This would give us an approximate solution, but at least it resolves the cycling issue

4. Efficiency

How fast is the simplex algorithm? This is a question that has only been partially resolved until quite recently!

Because since the simplex method walks through a small subset of a polyhedron, so the question essentially boils down to: What is the smallest distance between vertices x and y?

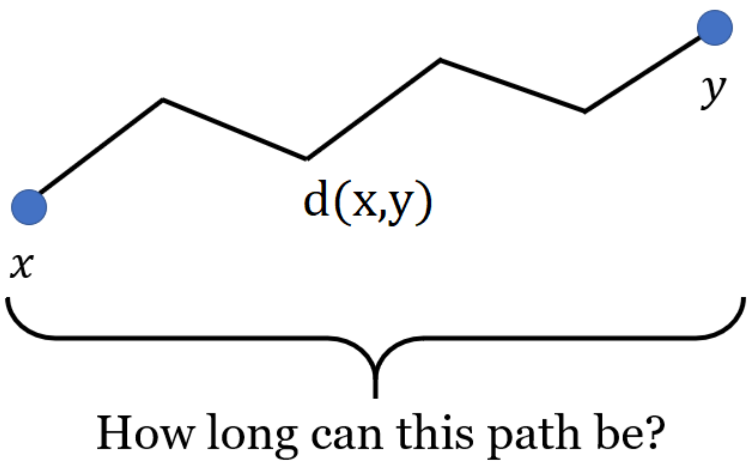

Definition: If x and y are vertices of a polyhedron, then

$$d(x,y) = \text{ Minimum length between } x \text{ and } y$$

Think of it as the most efficient way of going from x to y

Now given a polyhedron P, we just need to do this over all the vertices

Definition: If P is a bounded polyhedron, then the **diameter of** P:

$$D(P) = \max_{\text{vertices } x,y \in P} d(x,y)$$

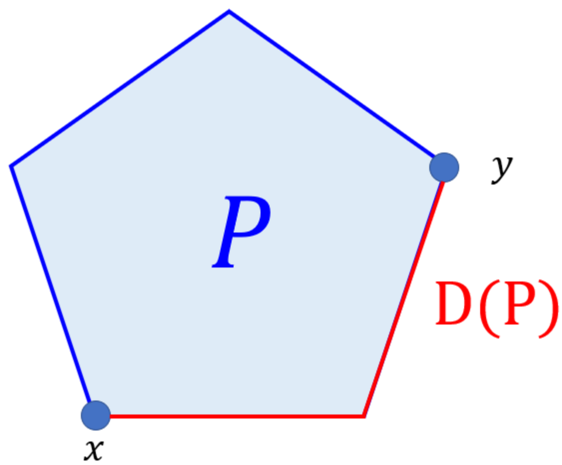

Think of it as the biggest path in P.

Finally, we just need to maximize this over all the vertices.

Definition:

$$\Delta(n,m) = \max_{P} D(P)$$

Where the max is taken over all bounded polyhedra P in $\mathbb{R}^n$ defined by m inequalities.

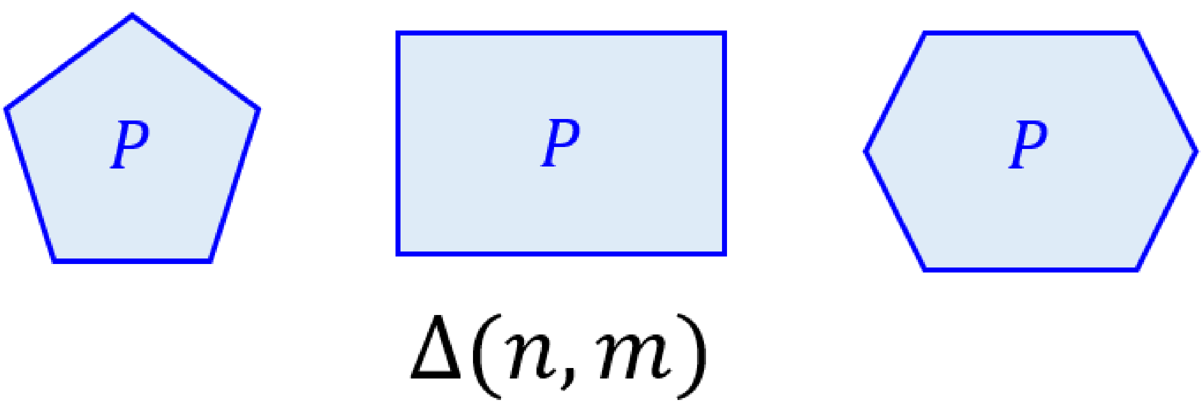

Hirsch Conjecture:

Do we have $\Delta(n, m) \leq m - n$?

Surprisingly the answer is **NO**! There is a counterexample found only recently by F. Santos in 2010 where we have

$$\text{Dimension: } n = 43$$
$$\text{Constraints: } m = 86$$
$$\text{But } \Delta(n, m) > 44$$