CS 435 : Linear Optimization — Fall 2008

Lecture 12: The Simplex Algorithm: Proof of Correctness

Lecturer: *Sundar Vishwanathan*
Computer Science & Engineering — Indian Institute of Technology, Bombay

In this lecture we will analyse the stopping condition of the Simplex algorithm and prove that the algorithm is correct.

To recap, consider an given extreme point x_0; given by $A^{'}x_0 = \mathbf{b}^{'}$ and $A^{''}x_0 \leq \mathbf{b}^{''}$. The directions of the neighbouring extreme points are the columns of the matrix $-A^{'^{-1}}$.

Stopping Condition: *The algorithm stops at an extreme point x_0 and returns it as optimal when the cost at x_0 is greater than or equal to the cost at the neighbouring extreme points.*

1 Proof of Correctness

We now prove that the simplex algorithm is correct, *i.e.*, when it terminates, we indeed have found the globally optimal point.

Note that the stopping condition does not say that x_0 is a local maximum. This is because we only know that the cost at x_0 is maximum as compared to the values at its *neighbours*–not compared to the values at *all points* in a small enough *neighbourhood* around it.

It is natural to hope that we can express each point in the neighbourhood of x_0 as a convex combination of x_0 and its neighbours, in which case we are done. Why?

We will finally achieve something like this but will take a detour.

1.1 Intuition

To understand this phenomenon better, we look around the point x_0. Consider the following figure:

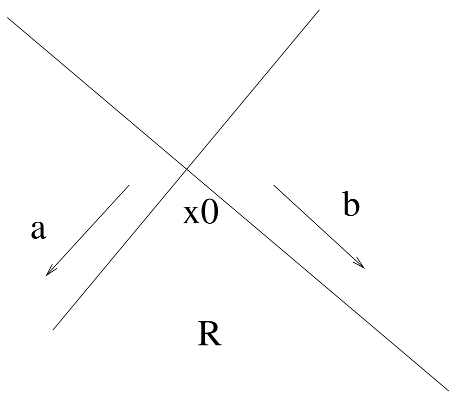

Let us say that the feasible region is R and x_0 is an extreme point. The vectors $\mathbf{a}$ and $\mathbf{b}$ are directions towards x_0's neighbours. The lines through x_0 partition the plane into four parts. Suppose that the cost decreases along the directions $\mathbf{a}$ and $\mathbf{b}$.

Can you identify other directions where the cost decreases?

Notice that in any direction in the region R the cost decreases. In the opposite region, the cost increases, and in the other two regions it varies from point to point.

Similarly in 3D, around a point, we see that there are 8 regions. The feasible points seem to lie in one region which is bounded by the vectors that determine the directions of the neighbours.

The main idea seems to be that every point in the feasible region can be given a non-negative linear combination of the direction vectors of the neighbours. Since the cost along the direction vectors of the neighbours does not decrease, no feasible point will have larger cost.

Why can we even write every feasible point as a *linear* combination of these direction vectors of the neighbours? Because they are linearly independent. Why?

1.2 The Proof

Assume that the algorithm terminates at x_0. Let x_{opt} be an optimal point. Therefore, $c^T x_{opt} \geq c^T x_0$.

As x_0 is an extreme point, $A^{'}$ has full rank. So, $-A^{'^{-1}}$ has full rank. (Why?)

In other words, the columns form a basis. Hence, the vector $x_{opt} - x_0$ can be written as a linear combination of the columns of $-A^{'^{-1}}$, *i.e,.*:

$$x_{opt} - x_0 = \sum_i \beta_i (-A^{'^{-1}})^{(i)} \tag{1}$$

We claim that β_i cannot be negative for any i. The geometric intuition is clear. We go outside the feasible region if we move in the direction of $(A'^{-1})^{(i)}$. But why cannot we come back using the other vectors?

Let us get back to the proof. Premultiplying the above equation with $A^{'}$, we get

$$A^{'} x_{opt} - A^{'} x_0 = \sum_i \beta_i A^{'} (-A^{'^{-1}})^{(i)} \tag{2}$$

What can we say about the resulting vector on the left hand side?

As x_{opt} is a feasible point, $A^{'} x_{opt} \leqslant \mathbf{b}^{'}$ whereas $A^{'} x_0 = \mathbf{b}^{'}$. Hence $A^{'} x_{opt} - A^{'} x_0$ will be a vector with each component at most 0.

How about the right hand side?

The product $A^{'}(-A^{'^{-1}})^{(i)} = -e^i$ has a zero at all positions except at the i^{th} row where it is -1. Hence $\sum_i \beta_i A^{'} (-A^{'^{-1}})^{(i)}$ is a vector β with the ith component $-\beta_i$.

These two observations imply that $\forall j, \beta_j \geqslant 0$. We are almost done, we have just concluded that the coefficients of the linear combination are non-negative.

Now multipling Equation 1 with c^T we get:

$$c^T x_{opt} - c^T x_0 = \sum_i \beta_i c^T (-A^{'^{-1}})^{(i)} \tag{3}$$

We know that as $(-A^{'^{-1}})^{(i)}$ are directions of the neighbours, and that the cost does not increase in these directions, $c^T(-A^{'^{-1}})^{(i)} \leq 0$. Coupled with the fact that $\beta_j \geq 0$, the right hand side of the above equation is at most 0. Hence, we get

$$c^T x_{opt} - c^T x_0 \leq 0 \tag{4}$$

And we conclude that x_0 is indeed an optimal point.

Point to ponder over: What would happen to this proof if we did not use the assumption that only n hyperplanes meet at a point? Where does it fail and how will you fix it.