COMP36111: Advanced Algorithms I
Lecture 4: Linear Programming

Ian Pratt-Hartmann

Room KB2.38: email: ipratt@cs.man.ac.uk

2017–18

Outline

The Linear Programming Problem

Geometrical analysis

The Simplex Method

- The originators of linear programming:

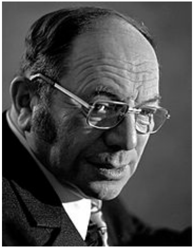

Leonid Kantorovich
1912–1986

George Dantzig
1914–2005

- The military situation around Leningrad, 1942

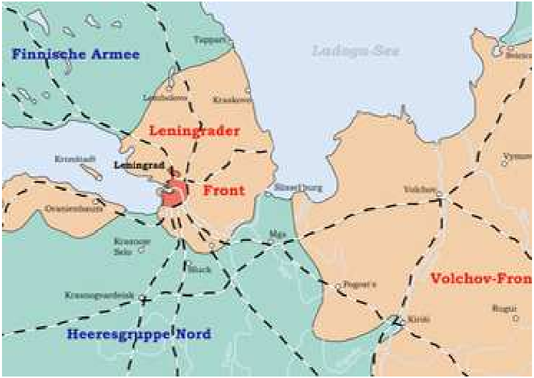

- Linear programming is an optimization problem in which we seek values for non-negative real variables which will maximize a linear objective function, subject to a set of linear constraints.
- Expressed in matrix form, we must maximize $\mathbf{c}^T\mathbf{x}$ subject to $A\mathbf{x} \leq \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$, where $\mathbf{x}$ is the variable vector $(x_1, \ldots, x_n)^T$
- The feasible region is the area of n-dimensional space $\mathbb{R}^n$ satisfying the constraints. This may be empty. Any point within this region is a feasible solution. The function $\mathbf{x} \mapsto \mathbf{c}^T\mathbf{x}$ is called the objective function.

- A Simple Example

Suppose Scatalogical Devices PLC (a chip company) makes profits of £10 on its α chip and £12 on its β chip.
The α chip costs 25 pence to produce, and takes 5 hours. The β chip costs 40 pence to produce and takes only 4 hours to produce.
The company has just £20 to spend and 320 hours of fabrication time.

- We express this problem as a set of inequalities
  - Let the number of α chips produced be x
  - Let the number of β chips produced be y
- Then the profit is $f = 10x + 12y$.
- And the constraints are given by

$$\begin{aligned} 0.25x + 0.4y &\leq 20 \\ 5x + 4y &\leq 320 \end{aligned}$$

- Or, if you like:

$$\begin{aligned} 5x + 8y &\leq 400 \\ 5x + 4y &\leq 320 \end{aligned}$$

- Pictorially:

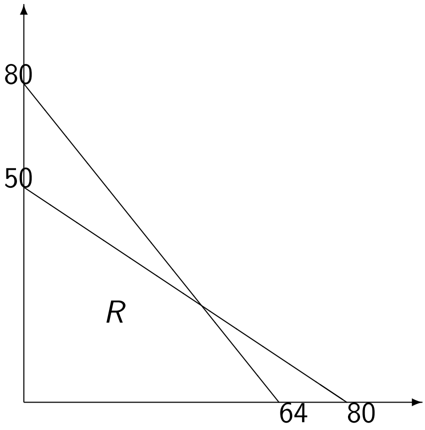

- We shall always take linear programming problems to have the form

$$\begin{aligned} \text{maximize:} \quad & \mathbf{c}^T\mathbf{x} \\ \text{subject to:} \quad & A\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq 0 \end{aligned}$$

- This is sometimes called standard form.
- Standard form is less restrictive than it looks:
  - An equality is the same as two inequalities
  - The inequality $\mathbf{a}^T\mathbf{b} \geq b$ can be re-written $-\mathbf{a}^T\mathbf{b} \leq -b$
  - Any real variable can be written as the difference of two non-negative real variables $x = x' - x''$.
- Strict inequalities cannot be eliminated in this way, but they are not really interesting in the context of maximization.

Outline

The Linear Programming Problem

Geometrical analysis

The Simplex Method

- A half-plane is the region consisting of a plane and one of its residual domains

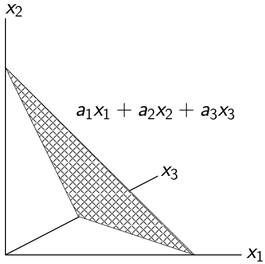

- A polyhedron is a non-empty intersection of half-planes (i.e. $A\mathbf{x} \leq \mathbf{b}$).

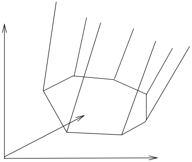

- All polyhedra have a simple characterization.
- A polytope is a bounded polyhedron.

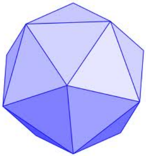

- A cone is a polyhedron defined by planes passing through the origin (i.e. $A\mathbf{x} \leq \mathbf{0}$).

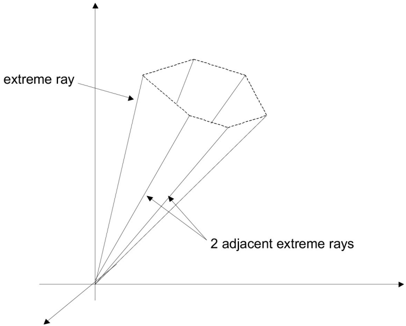

- The following facts are not hard to prove:
- The polytpes are exactly the convex combinations of finite non-empty sets of vectors:

$$\lambda_1 \mathbf{a}_1 + \cdots \lambda_p \mathbf{a}_p \qquad (\lambda_k \text{ all non-negative, } \lambda_1 + \cdots \lambda_p = 1)$$

- The cones are exactly the non-negative combinations of finite non-empty sets of vectors.

$$\lambda_1 \mathbf{a}_1 + \cdots \lambda_p \mathbf{a}_p \qquad (\lambda_k \text{ all non-negative}).$$

- Every polyhedron P defined by $A\mathbf{x} \leq \mathbf{b}$ is the Minkowski sum of a polytope U and a cone V (where we allow $V = \{\mathbf{0}\}$):

$$P = U \oplus V = \{\mathbf{u} + \mathbf{v} \mid \mathbf{u} \in U, \ \mathbf{v} \in V\}.$$

- Let's just think about polytopes for the moment. (This is the usual case in practice.)
- The objective function always has a maximum in this case. (Why?)
- The objective function is a convex combination of its value at the vertices: if

$$\mathbf{x} = \lambda_1 \mathbf{a}_1 + \cdots \lambda_p \mathbf{a}_p,$$

  then

$$f(\mathbf{x}) = \lambda_1 f(\mathbf{a}_1) + \cdots \lambda_p f(\mathbf{a}_p).$$

- It follows that the maximum must be achieved at a vertex. (Why?)

- Not only must the maximum be achieved at a vertex—any vertex **a** at which f is greater than its value at any neighbour is a global maximum.

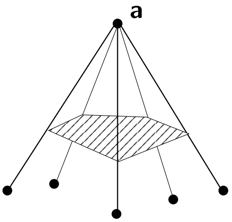

- This suggests an algorithm based on moving from one vertex to another in such a way as to increase the objective function.

- The simplex algorithm is:
  1. select a vertex of the feasible region as starting point;
  2. choose an edge on the surface of the feasible region through this point such that the objective function increases;
  3. proceed along this edge to the next vertex;
  4. repeat steps 2 and 3 until the objective function cannot be increased.

- Consider a system of m inequalities

$$A\mathbf{x} \leq \mathbf{b}$$

  in n variables, and assume all m rows of A to be linearly independent (generally, $m > n$).
- We can take $k \leq n$ of these inequalities and replace them with the corresponding equations (keep the remaining inequalities).
- This will define a face of the feasible region of dimension $n - k$.
- If $k = n$, we have a vertex.

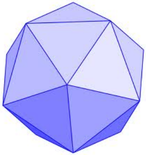

- Note that any inequality

$$a_1 x_1 + \cdots + a_n x_n \leq b$$

can be re-written as an equation

$$a_1 x_1 + \cdots + a_n x_n + w = b$$

where w is a new variable subject to the constraint $w \geq 0$.
- Such a variable is called a slack variable.
- Obviously, we can rewrite a whole system of inequalities in this way: we need a new slack variable for every inequality.

- Now think about faces of the feasible region in terms of the slack variables.
- A vertex is a point lying on n bounding planes—in particular, it is a solution of the equational system in which n variables are 0.
- We call such a solution in which n variables are chosen to be zero a basic feasible solution (bfs).
- These are called the basic variables.
- The remainder are called the non-basic variables.

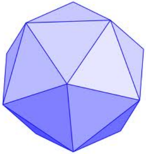

- Now imagine taking a bfs in which n slack variables are 0, and allowing one of the variables to become positive. (But stay in the polytope: don't violate any of the remaining inequalities!)
- This corresponds to moving away from the vertex corresponding to the bfs, and moving to a point lying on just $n-1$ of the n bounding planes involved—in other words, on an edge of the polytope adjacent to that vertex.

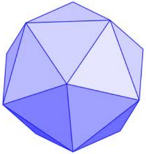

- So that is what it means to move from one vertex from another along an edge: take a bfs, and allow one of the basic variables to become positive, until you hit some other constraint, at which point another slack variable will become zero, and you have another bfs.

Outline

- Consider the following example

$$\text{maximize:} \ \ f = 4x - 2y - z$$

$$\text{subject to:} \quad \begin{array}{rcl} x + y + z & \leq & 3 \\ 2x + 2y + z & \leq & 4 \\ x - y & \leq & 1 \end{array}$$

$$x, y, z \geq 0.$$

- First, we introduce slack variables to turn the inequalities into equalities.
- Thus,

$$\begin{aligned} x + y + z &\leq 3 \\ 2x + 2y + z &\leq 4 \\ x - y &\leq 1 \end{aligned}$$

becomes

$$\begin{array}{ccccl} 3 & -x & -y & -z & = u \\ 4 & -2x & -2y & -z & = v \\ 1 & -x & +y & & = w \end{array}$$

with $u, v, w \geq 0$.

- So now we work with the equations

$$\begin{array}{lllll} 3 & -x & -y & -z & = u \\ 4 & -2x & -2y & -z & = v \\ 1 & -x & +y & & = w \end{array}$$

and the objective function

$$f = 4x - 2y - z,$$

and we take all variables to be non-negative.

- Equations:

$$\begin{array}{ccccc} 3 & -x & -y & -z & = u \\ 4 & -2x & -2y & -z & = v \\ 1 & -x & +y & & = w \end{array}$$

objective function:

$$f = 4x - 2y - z,$$

- We note that $x = y = z = 0$ is a solution. (This is not always so, but in real examples it quite often is.)
- This solution is a bfs, and the variables x, y, z are its basic variables.
- The value of f at this bfs is 0.

- Equations:

$$\begin{array}{rrrrcl} 3 & -x & -y & -z & = & u \\ 4 & -2x & -2y & -z & = & v \\ 1 & -x & +y & & = & w \end{array}$$

objective function:

$$f = 4x - 2y - z.$$

- We start off with the bfs $x = 0, y = 0, z = 0$.
- We can increase f by increasing x.
- By how much? Well, until one of the equations becomes insoluble. (Remember: all variables are non-negative).
- We can increase x by 1, which zeroes w.

- Equations:

$$\begin{array}{ccccc} 3 & -x & -y & -z & = u \\ 4 & -2x & -2y & -z & = v \\ 1 & -x & +y & & = w \end{array}$$

objective function:

$$f = 4x - 2y - z.$$

- We start off with the bfs $x = 0, y = 0, z = 0$.
- We can increase f by increasing x.
- By how much? Well, until one of the equations becomes insoluble. (Remember: all variables are non-negative).
- We can increase x by 1, which zeroes w.

- Now that w is 0, it becomes a basic variable, and x ceases to be one.
- So we pivot on x, rewriting

$$1 - x + y = w$$

as

$$1 - w + y = x$$

- Substituting the left-hand-side for x in the the above problem, we have the equations:

$$\begin{array}{lllll} 2 & +w & -2y & -z & = u \\ 2 & +2w & -4y & -z & = v \\ 1 & -w & +y & & = x \end{array}$$

and objective function:

$$f = 4 - 4w + 2y - z.$$

- Now we have the bfs $w = 0, y = 0, z = 0$.
- We now increase y by $\frac{1}{2}$, which zeroes v.

- Now that v is 0, it becomes a basic variable, and y ceases to be one.
- So we pivot on y, rewriting

$$2 + 2w - 4y - z = v$$

as

$$\frac{1}{2} + \frac{1}{2}w - \frac{1}{4}v - \frac{1}{4}z = y$$

- Substituting the left-hand-side for y in the the above problem, we have the objective function:

$$f = 5 - 3w - \frac{1}{2}v - \frac{3}{2}z.$$

and the equations ...

- ... Just a minute, we can't increase the objective function by increasing any variables!

- We had the equations

$$\begin{array}{lllll} 2 & +w & -2y & -z & = u \\ 2 & +2w & -4y & -z & = v \\ 1 & -w & +y & & = x \end{array}$$

and objective function

$$f = 5 - 3w - \frac{1}{2}v - 3z;$$

and we agreed to set $y = \frac{1}{2}$ with $w = z = 0$.

- We have arrived at the basic solution

$$(x, y, z, u, v, w) = (\frac{3}{2}, \frac{1}{2}, 0, 1, 0, 0)$$

and

$$f = 5.$$

- Reading
  - G+T: Ch 26.1–26.2, pp. 731–743.