LP. Lecture 1. Chapter 1 and 2: example and the simplex algorithm

This course gives an introduction to linear optimization and related areas.

- what is LP (lin.opt.=lin.programmering)
- more generally: mathematical optimization
- theory, methods, applications
- these notes are based on the textbook we use: R. Vanderbei: "Linear programming: fundations and extensions". Third edition, Springer (2008). (You may also use Second Ed., Kluwer (2001)).

- a practical example: production planning
- simplex algorithm, some concepts
- algorithm

What is linear optimization?

- Linear optimization is to maximize (or minimize) a linear function in several variables subject to constraints that are linear equations and linear inequalities.

Example: production planning

Products:

- door (with glass)
- window

Production facility:

- Factory 1: produces metal frame
- Factory 2: produces wooden frame
- Factory 3: produces glass and mounts the parts

Production of each product is made in series of 200 items.
Data:

	Hours/series		Hours at disposal
	door	window	
Factory 1	1	0	4
Factory 2	0	2	12
Factory 3	3	2	18
Revenue/series	3000	5000	

Problem: How much should be produced of each product in order to maximize the revenue?

The production plan as an LP problem:

$$
\begin{array}{llll}
\text{maximize} & 3x_1 + & 5x_2 & \\
\text{subject to} & & & \\
 & x_1 & & \leq 4 \\
 & & 2x_2 & \leq 12 \\
 & 3x_1 + & 2x_2 & \leq 18 \\
 & & x_1 \geq 0,\ x_2 \geq 0. &
\end{array}
$$

- We want to find an optimal solution, i.e., a vector (x_1, x_2) which satisfies all the constraints and has a maximum value of the function $f(x_1, x_2) = 3x_1 + 5x_2$.
- A vector satisfying all the constraints is called a feasible solution.
- The function we want to maximize is called the objective function.

Later we will work with LP problems in matrix form; then the problem is as follows

$$\begin{aligned} \max \quad & c^T x \\ \text{f.a.} \quad & \\ & Ax \leq b \\ & x \geq O \end{aligned}$$

Here c and x are column vectors with n components, A is a $m \times n$ matrix and b is a column vector with m components. O denotes the zero vector (of suitable length). The inequality $Ax \leq b$ is a vector inequality and means that $\leq$ holds componentwise (for every component).

Analysis of this problem and methods for solving it are based on linear algebra.

LP is closely tied to theory/methods for solving systems of linear inequalities. Such systems have the form

$$Hx \leq f$$

where H is a $m \times n$ matrix and f is a column vector (length m).

Example.

$$\begin{array}{rcrcr} 3x_1 & + & x_2 & \leq & 4 \\ x_1 & - & 2x_2 & \leq & 17 \\ -x_1 & & & \leq & 0 \end{array}$$

Central questions;

- existence of solution,
- how to find a solution, possibly *all* solutions.

Such problems may be written as LP problems: let the objective function have all its coefficients equal to 0.

More about linear inequalities later.

The simplex method

The simplex method is a general method for solving LP problems.

- Later we distinguish between the simplex method and the simplex algorithm, but this is not important now.
- The method was developed by George B. Dantzig around 1947 in connection with the investigation of transportation problems for the U.S. Air Force.
- The work was published in 1951.
- An interesting article in Washington Post i 2005 may be found on the course web page; Dantzig passed away in 2005.

The simplex algorithm, an example

Some other scientists who were early contributors to the development of linear programming were T.J.Koopmans and L.V.Kantorovich, and they were both awarded the Nobel prize in economics for this work in 1975.

Example: We want to solve

$$\begin{array}{lllll} \max & 5x_1+ & 4x_2+ & 3x_3 & \\ \text{subject to} & & & & \\ \text{(i)} & 2x_1+ & 3x_2+ & x_3 & \leq 5 \\ \text{(ii)} & 4x_1+ & x_2+ & 2x_3 & \leq 11 \\ \text{(iii)} & 3x_1+ & 4x_2+ & 2x_3 & \leq 8 \\ & & x_1, x_2, x_3 \geq 0. & & \end{array}$$

First, we convert to equations by introducing slack variables for every $\leq$-inequality: for instance (i) is replaced by

$$w_1 = 5 - 2x_1 - 3x_2 - x_3 \quad \text{and } w_1 \geq 0.$$

Then the problem may be written in the following form which we call a dictionary:

$$
\begin{array}{llllllll}
\max & \eta & = & & 5x_1 & + & 4x_2 & + & 3x_3 \\
\text{subj. to} & & & & & & & & \\
\text{(i)} & w_1 & = 5 & - & 2x_1 & - & 3x_2 & - & x_3 \\
\text{(ii)} & w_2 & = 11 & - & 4x_1 & - & x_2 & - & 2x_3 \\
\text{(iii)} & w_3 & = 8 & - & 3x_1 & - & 4x_2 & - & 2x_3 \\
& & & & & x_1, x_2, x_3, w_1, w_2, w_3 \geq 0. & & &
\end{array}
$$

- Left-hand side: dependent variables = basic variables.
- Right-hand side: independent variables = nonbasic variables.

Initial solution: Let $x_1 = x_2 = x_3 = 0$ and this gives $w_1 = 5, w_2 = 11, w_3 = 8$.

We always let the nonbasic variables be equal to zero. The basic variables are then *uniquely* determined; they become equal to the constants on the right-hand side.

Is this an optimal solution? No !!

For instance, we can increase x_1 while keeping $x_2 = x_3 = 0$. Then

- η (the value of the objective function) will increase
- we obtain new values for the basic variables; these new values are determined by x_1
- the more we increase x_1, the more η increases!
- watch out: the w_j's approach 0!

Maximum increase of x_1: want to avoid that the basic variables, one or more, become negative. From $w_1 = 5 - 2x_1$, $w_2 = 11 - 4x_1$ and $w_3 = 8 - 3x_1$ we get $x_1 \leq 5/2$, $x_1 \leq 11/4$, $x_1 \leq 8/3$ so we can increase x_1 to the smallest value, namely 5/2.

This gives the new solution $x_1 = 5/2$, $x_2 = x_3 = 0$ and therefore $w_1 = 0, w_2 = 1, w_3 = 1/2$. And now $\eta = 25/2$. Thus: a (much) better solution!!

How to proceed? The dictonary is well suited for testing optimality, swo we want to transform to a new dictionary.

- We want x_1 and w_1 to "switch sides". So: x_1 should go *into* the basis, while w_1 goes *out* of the basis. This can be done by using the w_1-equation in order to eliminate x_1 from all other equations.
- Equivalent: we may use elementary row operations on the system in order to eliminate x_1: (i) solve for x_1: $x_1 = 5/2 - (1/2)w_1 - (3/2)x_2 - (1/2)x_3$, and (ii) add a suitable multiple of this equation to the other equations so that x_1 disappears and is replaced by terms with w_1.

Remember: elementary row operations do not change the solution set of the linear system of equations.

Result:

$$
\begin{array}{lllllllll}
\eta & = 12.5 & - & 2.5w_1 & - & 3.5x_2 & + & 0.5x_3 \\
\hline
x_1 & = 2.5 & - & 0.5w_1 & - & 1.5x_2 & - & 0.5x_3 \\
w_2 & = 1 & + & 2w_1 & + & 5x_2 & & \\
w_3 & = 0.5 & + & 1.5w_1 & + & 0.5x_2 & - & 0.5x_3
\end{array}
$$

We have now performed a pivot: the use of elementary row operations (or elimination) to switch two variables (one into and one out of the basis).

Repeat the process!

Optimal? No: we can increase η by increasing x_3 from zero! May increase to $x_3 = 1$ for then $w_3 = 0$ (while the other basic variables are nonnegative).

Then we do another pivot: x_3 goes into the basis, and w_3 leaves the basis. This gives the new dictionary:

$$
\begin{array}{lllllll}
\eta & = 13 & - & w_1 & - & 3x_2 & - & w_3 \\
\hline
x_1 & = 2 & - & 2w_1 & - & 2x_2 & + & w_3 \\
w_2 & = 1 & + & 2w_1 & + & 5x_2 & & \\
x_3 & = 1 & + & 3w_1 & + & x_2 & - & 2w_3
\end{array}
$$

Here we see that all coefficients of the nonbasic variables are nonpositive (in fact negative) in the equation for η. Then every increase of one or more nonbasic variables will result in a soluton where $\eta \leq 13$.

But: *any* feasible solution is obtained by a suitable choice of the nonbasic variables! Why?

Conclusion: we have found an optimal solution! It is $w_1 = x_2 = w_3 = 0$ and $x_1 = 2, w_2 = 1, x_3 = 1$.
The corresponding value of η is 13, and this is called the optimal value.

The simplex method, in general

Consider a general LP problem

$$
\begin{array}{lrll}
\max & \sum_{j=1}^{n} c_j x_j & & \\
\text{subj. to} & & & \\
& \sum_{j=1}^{n} a_{ij} x_j & \leq b_i & \text{for } i = 1, \ldots, m \\
& x_j & \geq 0 & \text{for } j = 1, \ldots, n.
\end{array}
$$

where we (now) assume that $b_i \geq 0$ for all $i \leq m$.

Introduce slack variables

$$\begin{array}{rcl} \eta & = & \sum_{j=1}^{n} c_j x_j \\ w_i & = & b_i - \sum_{j=1}^{n} a_{ij} x_j \quad \text{for } i = 1, \ldots, m \end{array}$$

We do not need to distinguish between slack variables and the original variabels so we get the following dictionary:

$$\begin{array}{rcl} \eta & = & \sum_{j=1}^{n} c_j x_j \\ x_{n+i} & = & b_i - \sum_{j=1}^{n} a_{ij} x_j \quad \text{for } i = 1, \ldots, m \end{array}$$

where we have replaced w_i by x_{n+i}. So $x \in \mathbb{R}^{n+m}$.

The algorithm starts with this dictionary where $x_{n+1}, \dots, x_{n+m}$ are basic variables and $x_1, \dots, x_n$ are nonbasic variables. Let

- B be the index set of the basic variables.
- N be the index set of the nonbasic variables.

So, initially $B = \{n+1, \dots, n+m\}$, $N = \{1, \dots, n\}$. The initial solution is

$$\begin{aligned} x_j &= 0 \quad \text{for } j = 1, \dots, n \\ x_{n+i} &= b_i \quad \text{for } i = 1, \dots, m \end{aligned}$$

and the corresponding value of η is $\eta = 0$. Such a solution is called a basic solution.

Every iteration is a pivot (or change of basis) where

- one index k is moved from N to B (x_k is entering (ingoing) variable; it is a new basic variable because it results in an increase of η),
- another index l is moved from B to N (x_l is leaving (outgoing) variable; this variable leaves the basis because it becomes 0 as the first one, and
- we find the new dictionary from the old by performing row operations (or elimination)
- the basic solution that corresponds to the new dictionary is feasible.

At the start of every pivot we have the dictionary (with $\bar{b}_i \geq 0$):

$$\begin{aligned} \eta &= \bar{\eta} + \sum_{j \in N} \bar{c}_j x_j \\ x_i &= \bar{b}_i - \sum_{j \in N} \bar{a}_{ij} x_j \quad \text{for } i \in B. \end{aligned}$$

Selection of entering variable: choose a $k \in N$ with $\bar{c}_k > 0$. If no such index exists, the current solution is optimal and we terminate. Often several $\bar{c}_j$'s are positive. There are several principles for selection of the entering variable, but a simple, and often used, prnciple is to choose $k = j$ with $\bar{c}_j$ largest possible. Why?

Selection of leaving variable: Also here we may have several choices. First we have to determine the maximum increase of the entering variable x_k. From

$$x_i = b_i - \bar{a}_{ik} x_k \quad \text{for } i \in B$$

we see that

- if $\bar{a}_{ik} \leq 0$, x_i will increase when x_k is increased. Such basic variables will not become zero when x_k is increased (we assume now that $\bar{b}_i > 0$)
- if, however, $\bar{a}_{ik} > 0$, then x_i will decrease and it becomes zero when

$$x_k = b_i / \bar{a}_{ik}.$$

Selection of leaving variable, cont.

So, we can increase x_k to the value

$$\theta := \min\{b_i/\bar{a}_{ik} : \bar{a}_{ik} > 0\}.$$

What happens when $x_k = \theta$?

Well, all variables are still nonnegative. Good! And at least one basic variable has become zero, in fact $x_i = 0$ for all $i \in B$ satisfying

$$b_i/\bar{a}_{ik} = \theta.$$

Conclusion: Leaving variable x_l is selected so that

$$b_l/\bar{a}_{l,k} = \min\{b_i/\bar{a}_{ik} : \bar{a}_{ik} > 0\}.$$

Pivot rule: a rule which tells us which entering variable to choose and which leaving variable to choose.

Several pivot rules are around, so we get several variants of the simplex algorithm.

The pivot is terminated by the row operations:
assume x_k is entering variable and x_l is leaving variable. Then x_l is on the left-hand side in "equation number l":

$$x_l = \bar{b}_l - \sum_{j \in N} \bar{a}_{lj} x_j$$

For every equation $i \neq l$ we add
$\bar{a}_{ik}/\bar{a}_{lk}$ times equation l to equation i.
Furthermore, we use equation l to solve for x_k and therefore express x_k as a function of the other variables.

Result: an equivalent system of equations is obtained, so the same solutions, and with coefficient 0 associated with each x_k in every equation $i \neq l$. Further, the new basis variable is on the left-hand side! This is the new dictionary.

Comments

Some questions remain:

- how to find an initial basic feasible solution if there are negative b_i's?
- in a dictionary: what happens if some $\bar{b}_i$'s are 0? Does this cause problems for the pivot rule?
- will the algorithm terminate?
- and, if not, can we repair this somehow?

We will start working on these questions in Lecture 2!

Finally, let us consdier an important application of LP in connection with linear l_1-approximasjon.

Application: linear approximation

Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$, and let a_i be the i'th row in A considered as a column vector. Recall that the l_1-norm of a vector $y \in \mathbb{R}^n$ is $\|y\|_1 = \sum_{i=1}^n |y_i|$.
The linear approximation problem

$$\min\{\|Ax - b\|_1 : x \in \mathbb{R}^n\}$$

may be solved as the following LP problem

$$\begin{array}{lll} \min & \sum_{i=1}^m z_i & \\ \text{subj.to} & & \\ & a_i^T x - b_i & \leq z_i \quad (i \leq m) \\ & -(a_i^T x - b_i) & \leq z_i \quad (i \leq m) \end{array}$$

Proof: Because every optimal solution in this LP satisfies $z_i = |a_i^T x - b_i|$ for every $i \leq m$. □

This means that one has an alternative method to the traditional least squares method based on solving $\min\{\|Ax - b\|_2 : x \in \mathbb{R}^n\}$ and this problem has lots of important application (see any linear algebra textbook).

A similar LP approach works for the linear approximation problem

$$\min\{\|Ax - b\|_\infty : x \in \mathbb{R}^n\}$$

in the ℓ_∞-norm given by $\|z\|_\infty = \max_i |z_i|$.