

Linear Programming Problems

Linear programming problems come up in many applications. In a linear programming problem, we have a function, called the *objective function*, which depends linearly on a number of independent variables, and which we want to optimize in the sense of either finding its minimum value or maximum value.

Depending linearly means that the objective function is in the form of a linear polynomial, a polynomial in which each of the variables occurs to the first power and none of the variables are multiplied together.

Constraints

The variables are also subject to conditions, in the form of linear inequalities. These are called constraints.

The variables must also satisfy the *non-negativity condition*: they can't be negative.

The set of points, or values of the variables, which satisfy the constraints and the non-negativity condition is called the *feasible set*.

Fundamental Theorem of Linear Programming

Theorem 1 (Fundamental Theorem of Linear Programming). *The optimal value of the objective functions must occur at one of the vertices of the feasible set.*

We can understand this fairly easily for a case where there are two independent variables. The notation gets more complicated when there are more variables, but the basic ideas carry through.

Consider the values of an objective function $p = \alpha x + \beta y$ along a line $y = mx + b$. If we plug $mx + b$ in for y in the formula for p , we get $p = \alpha x + \beta(mx + b) = \alpha x + \beta mx + \beta b$ or

$$p = (\alpha + \beta m)x + \beta b.$$

If $\alpha + \beta m > 0$, then p increases when x increases and p decreases when x decreases.

If $\alpha + \beta m < 0$, then p decreases when x increases and p increases when x decreases.

In either case, p increases when we go along the line in one direction and decreases when we go along the line in the opposite direction.

A third possibility is that $\alpha + \beta m = 0$, in which case p never changes.

In any of those cases, along any segment of the line, any extreme value for p must occur at an endpoint.

If an extreme value occurs in a closed polygon, we easily see the extreme value must occur along the boundary, and then if it occurs along the boundary it must occur at one of the vertices, which is what the Fundamental Theorem of Linear Programming says.

In a linear programming problem with just two variables and a handful of constraints, it's easy to sketch the feasible set and find its vertices. This is the essence of solving linear programming problems geometrically.

- Find the feasible set.
- Find the vertices.
- Evaluate the objective function at each vertex.
- Pick out the optimal value for the objective function.

We'll do some examples to help understand linear programming problems, but most linear programming problems that come up in real life involve numerous variables and constraints and effectively require a more efficient approach. The most common approach is called the *Simplex Method*.

The Elephant in the Room

Linear programming problems don't come out of thin air; there are real problems that, when translated, turn into linear programming problems.

The process of turning a real problem into a linear programming problem is the same involved in any other word problem: Read the problem!!!

In particular,

- Look for variables and unknowns.
- Find connections among the variables and unknowns. In this case, the connections translate into constraints along with the objective function.

The Simplex Method

In most word problems studied before, the connections translated to equations, which were then solved.

In linear programming problems, we generally use the *Simplex Method*.

An Example

Consider the following maximum problem:

$$\begin{aligned}
 2x + y &\leq 3 \\
 3x + y &\leq 4 \\
 x \geq 0, y &\geq 0 \\
 p &= 17x + 5y
 \end{aligned}$$

Vertices

The vertices of the feasible set and the values of the objective function $p = 17x + 5y$ at those points are the following:

$$\begin{array}{ll}
 (0, 0) & 0 \\
 (4/3, 0) & 22\frac{2}{3} \\
 (1, 1) & 22 \\
 (0, 3) & 15
 \end{array}$$

Thus, the maximum value for p is $22\frac{2}{3}$, which occurs at $(4/3, 0)$.

The Simplex Method

With the Simplex Method, we do not need to evaluate the objective function at each vertex.

We will illustrate the Simplex Method using this example and then generalize the method so we can perform it mechanically, as an algorithm.

The first step is to change the constraints from inequalities into equations by introducing what are called *slack variables*. The slack variables essentially take up the slack in the inequalities.

Looking at the inequality $2x + y \leq 3$, we introduce the slack variable $u = 3 - (2x + y)$, so $u \geq 0$ while $2x + y + u = 3$.

Similarly, looking at the inequality $3x + y \leq 4$, we introduce the slack variable $v = 4 - (3x + y)$, so $v \geq 0$ while $3x + y + v = 4$.

Restating the Problem

We can thus restate the problem as

$$\begin{aligned}
 2x + y + u &= 3 \\
 3x + y + v &= 4 \\
 x, y, u, v &\geq 0 \\
 p &= 17x + 5y
 \end{aligned}$$

where we find the maximum value for the objective function p which satisfies the constraints, now given as equations, and the non-negativity condition.

A New System of Equations

If we rewrite the formula for the objective function in the form $-17x - 5y + p = 0$, we get a system of equations:

$$\begin{aligned} 2x + y + u &= 3 \\ 3x + y + v &= 4 \\ -17x - 5y + p &= 0. \end{aligned}$$

We will write down the augmented matrix for this system, use it to easily pick out one of the vertices and the value of the objective function at that vertex, and then pivot in a way that will enable us to find another vertex at which the value for the objective function is larger.

The Initial Simplex Tableau

The augmented matrix for the system

$$\begin{aligned} 2x + y + u &= 3 \\ 3x + y + v &= 4 \\ -17x - 5y + p &= 0. \end{aligned}$$

is

$$\begin{pmatrix} x & y & u & v & p \\ 2 & 1 & 1 & 0 & 0 & 3 \\ 3 & 1 & 0 & 1 & 0 & 4 \\ -17 & -5 & 0 & 0 & 1 & 0 \end{pmatrix}$$

The first row isn't actually part of the augmented matrix, but we'll write it down as a reminder of which variable each column is associated with.

The Initial Tableau

$$\begin{pmatrix} x & y & u & v & p \\ 2 & 1 & 1 & 0 & 0 & 3 \\ 3 & 1 & 0 & 1 & 0 & 4 \\ -17 & -5 & 0 & 0 & 1 & 0 \end{pmatrix}$$

Looking at the matrix, it's fairly obvious that the columns corresponding to x and y are the most complicated. We can easily read off a solution if we set $x = 0$ and $y = 0$. If we do that, the first row, which corresponds to the equation $2x + y + u = 3$, gives a solution $u = 3$. Similarly, the second row, which corresponds to the equation $3x + y + v = 4$, gives a solution $v = 4$. The third row, which corresponds to the equation $-17x - 5y + p = 0$, gives a solution $p = 0$.

A Solution

So we have a solution of the equations giving $x = 0$, $y = 0$, $u = 3$, $v = 4$ and $p = 0$. This corresponds to the vertex $(0, 0)$, which gives a value of 0 for the objective function.

We will call x and y Group I variables and we will call u and v Group II variables.

Each time we pivot, we will exchange one pair of variables, one from Group I and one from Group II, to get another set of Group I variables which will be set to 0. This will give another fairly obvious solution, which will correspond to another vertex and another value of the objective function.

The Objective Function

The last row gives a formula $-17x - 5y + p = 0$ involving the objective function. We can solve this for the objective function, getting $p = 17x + 5y$. *It may seem as if we're going around in circles, but right now we're trying to understand a process; later we will synthesize the ideas into an efficient algorithm.*

Since at the solution we found, $x = y = 0$, p takes on the value 0. On the other hand, it's clear that if we found another solution where either x or y was positive, we'd get a larger value for p . It should also be clear that increasing x will increase p faster than increasing y would, so we'll pivot in a way that will make x a Group II variable.

We'll do this by pivoting about the column corresponding to x , in other words, we'll pivot about the first column.

Finding the Pivot Row

We now have to decide which entry in the first column to pivot around. Since the last row corresponds to the objective function, which is different in character from the others, we won't even consider pivoting about the last row.

If we pivoted about the first row, we'd start by dividing the first row by 2, getting $3/2$ in the last column. We'd next subtract 3 times the first row from the second, which would give $-\frac{1}{2}$ in the last column. Since negative numbers are evil, we'll avoid that by pivoting about the second row. So we'll pivot about the entry in the second row, first column.

The Pivot

$$\begin{pmatrix} 2 & 1 & 1 & 0 & 0 & 3 \\ 3 & 1 & 0 & 1 & 0 & 4 \\ -17 & -5 & 0 & 0 & 1 & 0 \end{pmatrix}$$

First, we divide the second row by 3 to get a 1 at the pivot point:

$$\begin{pmatrix} 2 & 1 & 1 & 0 & 0 & 3 \\ 1 & 1/3 & 0 & 1/3 & 0 & 4/3 \\ -17 & -5 & 0 & 0 & 1 & 0 \end{pmatrix}$$

To get 0's elsewhere in the first column, we'll subtract $2\times$ the pivot row from the first row and add $17\times$ the pivot row to the third row:

$$R_1 \leftarrow R_1 - 2R_2, R_3 = R_3 + 17R_2:$$

$$\begin{pmatrix} 0 & 1/3 & 1 & -2/3 & 0 & 1/3 \\ 1 & 1/3 & 0 & 1/3 & 0 & 4/3 \\ 0 & 2/3 & 0 & 17/3 & 1 & 68/3 \end{pmatrix}$$

Interpreting the New Matrix

We'll rewrite the matrix with the columns again labeled:

$$\begin{pmatrix} x & y & u & v & p \\ 0 & 1/3 & 1 & -2/3 & 0 & 1/3 \\ 1 & 1/3 & 0 & 1/3 & 0 & 4/3 \\ 0 & 2/3 & 0 & 17/3 & 1 & 68/3 \end{pmatrix}$$

This time, the Group II variables are y and v and we'll set them to 0.

We can then interpret the equations. The first row corresponds to the equation $\frac{1}{3}y + u - \frac{2}{3}v = \frac{1}{3}$. Since $y = v = 0$, we get $u = \frac{1}{3}$.

The second row corresponds to the equation $x + \frac{1}{3}y + \frac{1}{3}v = \frac{4}{3}$. Again, since $y = v = 0$, we get $x = \frac{4}{3}$.

The third row corresponds to the equation $\frac{2}{3}y + \frac{17}{3}v + p = \frac{68}{3}$. Once again, since $y = v = 0$, we easily get $p = \frac{68}{3}$.

This corresponds to the vertex $(4/3, 0)$, giving the value $\frac{68}{3} = 22\frac{2}{3}$ for the objective function.

Another Look at the Last Row

Looking at the last row, corresponding to the equation $\frac{2}{3}y + \frac{17}{3}v + p = \frac{68}{3}$, more closely, we can solve for p to get:

$$p = \frac{68}{3} - \frac{2}{3}y - \frac{17}{3}v.$$

We just found a vertex corresponding to values of 0 for y and v . Any other solution would either involve still having 0 values for y and v , which would give the same value for p , or would involve positive values for one or both of y and v , which would give a smaller value for p .

Obviously, we have found the largest possible value for the objective function.

We'll analyze what we just did, synthesize it, and come up with the *Simplex Method*.

The Simplex Method

- (1) The first step is to take each inequality and introduce a slack variable to get an equation.
- (2) We also rewrite the formula for the objective function, bringing all the variables over to the left side.
- (3) We now write down the augmented matrix for the system of equations, with the following optional adjustment: Since we never pivot about the last row, the next to last column never changes. Effectively, the next to last column is unnecessary, so we don't have to write it down.

At this point, we have the *Initial Simplex Tableau*.

Pivoting

We then pivot, choosing where to pivot as follows:

- (1) We look at the bottom row, look at the negative entries, and choose the column for which the negative entry in the bottom row has the largest magnitude. That will be the pivot column. *If there are no negative entries, then we are finished.*
- (2) Look at the rows with positive entries in both the pivot column and the last column. In each such row, compute the quotient of the entry in the last column divided by the entry in the pivot column. The row for which this quotient is smallest is the pivot row.

Once we have chosen the pivot row and pivot column, we pivot.

The End Game

We continue pivoting until we can't pivot any more, which occurs when there are no negative entries in the bottom row.

At this point, we can set the Group I variables to 0, easily find the values for the Group II variables, and the entry in the bottom right of the matrix will be the value of the objective function.

Complications

This works well as long as all the constraints are in the form
 (Linear Polynomial) \leq (Positive Constant)
 and we are trying to maximize the objective function.

Things get more complicated if the inequality goes the other way, which generally occurs when we are trying to minimize the objective function and can also occur occasionally in maximum problems. Dealing with minimum problems is another complication.

Dealing with \geq

If the inequality goes the other way, we can deal with it by multiplying both sides by -1 . In other words, if we have an equality of the form

$$P \geq k,$$

where P is a linear polynomial and k is a constant, we can instead use the equivalent inequality

$$-P \leq -k.$$

The complication is that this is likely to lead to a negative number on the right hand side. We need to see how to deal with that.

Dealing with Negative Constants

The problem with having a negative constant is that when we pick out what we hope will be a feasible solution, one of the variables will fail to satisfy the non-negativity condition. We deal with that via some preliminary pivoting on the Initial Simplex Tableau.

Preliminary Pivoting

If there is a negative entry in the last column of a tableau (excluding the bottom row), we look in that row for a negative entry.

The column containing that entry becomes our pivot column.

Note:

- If there is more than one row with a negative entry in the last column, we can look at any of those rows.
- If there is more than one negative entry in that row, we may choose any of the columns with a negative entry as the pivot column.

Once we have the pivot column, we choose the pivot row in the usual way, with the following adjustment: We consider all rows where the quotient of the entry in the last column divided by the entry in the pivot column is positive.

We then pivot and repeat this preliminary process until there are no more negative entries in the last column. At that point, we proceed with the Simplex Method in the usual way.

Minimizing the Objective Function

We minimize the objective function by maximizing its additive inverse, that is, its negative.

In other words, if we need to minimize p , we simply maximize $M = -p$.

When we are finished pivoting, we then merely recognize that the entry in the lower right hand corner of the final simplex tableau is the negative of minimum value for the objective function.