

## Supplementary Notes on Linear Programming

A linear program is an optimization problem. In an  $n$  dimensional space, whose points are described by variables  $x_1, \dots, x_n$ , we have a “feasible region” which is a “polytope” by which we mean a region whose boundaries are defined by linear constraints.

In two dimensions, a polytope is a polygon, (which can be unbounded and sometimes is, and can even conceivably be empty) In three dimensions a polytope can be envisioned as a polished gem with planar facets.

The standard LP consists of the problem of finding the maximum value of some linear function of the  $n$  variables, among points of the feasible region. This amounts to maximizing the dot product of your  $\mathbf{r}$  variable with some vector  $\mathbf{c}$ , which means finding a point in the polytope with maximum component in the direction of  $\mathbf{c}$ .

We use the following terminology. The points of our  $n$  dimensional space that obey a single one of our linear constraints as equalities, define a **hyperplane**. In three dimensions a hyperplane is just a plane, and in fact a hyperplane is the generalization of a plane in 3D to higher dimensions. It is a plane-like region of  $n-1$  dimensions in an  $n$  dimensional space. A hyperplane that actual forms part of the boundary of the feasible region is called an  $n-1$  face of that region.

The points that obey two constraints lie on the intersection of the hyperplanes defined by them, the  $n-2$  dimensional region of this intersection form part of the boundary of the feasible region they form an  $n-2$  face of it.

And we have similar definitions for the intersection of  $k$  of the constraints that form part of the boundary of the feasible region: these are called  $n-k$  faces of it.

A **zero face** of the feasible region, also called a vertex of it, is a point which lies on the intersection of  $n$  (linearly independent) constraints and is on the boundary of the feasible region. Two vertices are adjacent if  $n-1$  of their defining constraints are the same, and the intersection of the  $n-1$  shared constraints is a 1-face that is called the edge joining them.

The linear function, called the **objective function**, that we seek to maximize in our feasible region defines a direction on our space, and we seek a point that maximizes the component of the position vector  $\mathbf{r}$  in that direction, In one or two dimensions it is easy to see that there is always such a point that is a vertex of the feasible polytope (assuming that the feasible region is bounded) The same statement is actually true in any dimension: there is always a vertex of the feasible region at which a maximum is achieved.

There are many approaches to “solving” an LP, {which means finding a vertex at which the maximum of the objective function is achieved, by determining all the components of the position vector at that vertex, and the value of the maximum as well.)

The Simplex Algorithm of Dantzig (and perhaps von Neumann) is the classical approach to it. It has been remarkably successful in practice, and its successes contributed greatly to the development of the field of operations research, in which it has been a major tool.

Geometrically, it can be described very simply. You begin at a vertex of the feasible region. You then move to an adjacent vertex across a 1-face (edge) at which the objective function is larger than what it was. This is the basic step. You repeat this step until you are at the maximum point (or, if there is no maximum point you find a direction in which the feasible region is unbounded and the objective function increases as you move in that direction)

In three dimensions you can imagine yourself climbing along the edges of a diamond, always upward, until you reach the top.

So how do we actually do this?

We suppose first that all our constraints are inequalities, which each requires that feasible points lie either on or on one side of a hyper-plane that it defines. The feasible region is then the intersection of the feasible “half-spaces” so defined by each constraint. We also assume, for convenience of our discussion, that all of our variables are constrained to be positive, and that the origin of our coordinates is a vertex of the feasible region. We will relax these assumptions later.

A given point in our  $n$ -dimensional space is characterized by its  $n$  coordinates, but we can define additional coordinates as well. At each point  $\mathbf{r}$ ,  $\mathbf{r}=(x_1, \dots, x_n)$ , for each additional constraint beyond those that require our coordinates to be positive, we can calculate the distance from  $\mathbf{r}$  to that constraint (with a positive sign if it is on the feasible side of the constraint.) If there are  $m$  of these additional constraints, we can define  $n+m$  variables associated with each point  $\mathbf{r}$  in our space, these being the  $n$  original coordinates and the (signed) distances to each of the additional constraints.

The actual LP problem can then be defined by  $m$  equality constraints among these  $n+m$  variables, as we shall see, and by the objective function..

The initial form of the equation has the original position variables as “basis variables” among the  $n+m$  variables, and the  $m$  variables that define distances to the other constraints are expressed in terms of these, again as we shall see.

The basic step of the algorithm consists of switching one of these basis variables with one of the other (distance to constraint) variables, so that we have a new set of basis variables that differs by one substitution from the original ones.

What does this basic step have to do with moving from one vertex of the polytope to another?

We associate with each set of basis variables the point at which all of them are set equal to 0. This is initially the origin of our coordinate system, which, by our assumption here is a vertex of our feasible region. When we make a switch to our basis, we switch from the initial origin to the origin of the new basis, The operation of removing one variable from the basis and replacing it by another is called a pivot, and the new origin will be a vertex of the feasible region if the old one was one.

So the simplex algorithm consists of a sequence of pivot operations, and we move from our initial origin to the origin in a sequence of new bases, until we reach a vertex, if any exist, at which the objective function is maximized.

Each pivot moves us across a 1-face of the feasible polytope to a new origin-vertex. We choose ourselves where to move. And we choose always to move to a new vertex at each stage that has a larger value for the objective function than the previous origin-vertex had. Since the objective function keeps increasing at each step, the algorithm cannot cycle, and, since there are only a finite number of possible vertices, it must eventually reach the top value in the polytope, if there is one.

Is this really doable?

Yes. The geometric picture can be translated into an algebraic one with remarkable ease. It is easy and straightforward to find a pair of variables to pivot on, and to perform a pivot, as we shall now see.

So let us now define the problem algebraically.

Our  $n$  original variables are  $x_1, \dots, x_n$ . We refer to a typical variable as  $x_j$ . We refer to a typical constraint as  $A_i$ , which is a vector whose  $j$ -th component is  $A_{ij}$ .

Our  $m$  constraints can be defined by the matrix  $\{A_{ij}\}$  and take the form

$$\sum_j A_{ij} x_j \leq b_j, \text{ for each } i$$

We also have the constraints that our variables are all positive,  $x_j \geq 0$  for each  $j$

Our objective function can be written as the dot product  $(\mathbf{c}, \mathbf{r})$  or  $\sum_j c_j x_j$ .

For each constraint  $i$  the distance variable to it is called its “slack” variable, and is usually denoted by  $s_i$ . The  $i$ -th constraint then becomes  $s_i \geq 0$ , where  $s_i$  is defined by the equation

$$s_i + \sum_j A_{ij} x_j = b_i.$$

So we have  $n+m$  variables ( $x$ 's and  $s$ 's) the constraints are all that these variables are non-negative, and there are  $m$  linear equations among them.

Notice that the initial origin will only be feasible when setting the  $x$ 's to 0 makes the  $s$ 's all non-negative. This means here that all the  $b$ 's must be non-negative. In fact, we really want all the  $b$ 's to be positive, so if any are 0 we add a tiny amount to each to make them all positive.

A pivot consists of picking an  $x$  variable,  $x_{j0}$ , and an  $s$  variable,  $s_{i0}$ , and interchanging their roles here.

And what does that mean?

Well you use the  $i_0$  equation to eliminate  $x_{j_0}$  from all other constraints and the objective function, substituting for it according to the  $i_0$  equation everywhere. You also divide the  $i_0$  equation by  $A_{i_0 j_0}$  so the coefficient of  $x_{j_0}$  in it will be 1 as was the coefficient of  $s_i$  before the pivot.

And that is all there is to a pivot. Which leaves two questions.

What pivot should we perform?  
How actually do we do it?

What we are doing with an  $(s_{i_0}, x_{j_0})$  pivot is to move ourselves from the origin along the  $x_{j_0}$  axis in the positive direction until we meet the first constraint which will be the  $i_0$ -th constraint. .

So we want to move along an axis in which moving positively increases the objective function, This means that  $c_{j_0}$  should be positive. You can pick any  $j_0$  for which this condition holds to pivot with.

Once you choose  $j_0$  then  $i_0$  is defined by the first constraint that you reach along the 1-face defined by the  $x_{j_0}$  axis.

And how?

Well as you increase  $x_{j_0}$  for those constraints  $i$  for which  $A_{ij_0}$  is negative, the left hand side of the constraint decreases as  $x_{j_0}$  increases. This means you will never hit the constraint. (you are moving away from it.),

On the other hand you move toward constraints for which  $A_{ij_0}$  is positive when you increase  $x_{j_0}$ . And  $A_{ij_0}$  represents the derivative of that constraint with respect to that variable, And  $b_i$  represents the distance from the origin to that constraint. So the amount you can increase  $x_{j_0}$  until you hit the constraint is  $b_i/A_{ij_0}$ ;

You must stop increasing  $x_{j_0}$  at the first constraint you encounter, which is therefore the  $i$  that minimizes  $b_i/A_{ij_0}$  among those  $i$  for which  $A_{ij_0}$  is positive, That  $i$  will be  $i_0$ .

And the objective function at the point you reach, which will be the new origin after this pivot, will increase from its previous value by this increase times the derivative of the objective function with respect to  $x_{j_0}$ , (which is  $c_{j_0}$ ) hence by  $b_{i_0} * c_{j_0} / A_{i_0 j_0}$ .

Doing pivots by hand is tedious and errors are hard to avoid. Fortunately it easy to do them on a spreadsheet, You need to enter two instructions for each pivot, and do some copying.

Consider the following example:

	Constraints
1:	$x_1 + 3x_2 - x_3 + x_4 \leq 2$
1	$-x_1 + x_2 + 2x_3 + x_4 \leq 1$
2	$x_1 - 2x_2 + x_3 - x_4 \leq 1$

Maximize  $x_1+x_2+x_3+x_4$

s

We set up a “tableau” with a column for each variable including the three slack variables, and also for the  $-b$ 's which we get by moving the right hand sides here to the left side. The tableau for this LP then looks like.

s1	s2	s3	x1	x2	x3	x4	-b	
1	0	0	1	3	-1	1	-2	first equation
0	1	0	-1	1	2	1	-1	second
0	0	1	1	-2	1	-1	-1	third
0	0	0	1	1	1	1	0	objective function

At this point you can pivot on any  $x_j$  since all coefficients in the objective function are positive. If you choose  $x_1$  you must pivot on the third equation, and the objective function will increase by 1. If you choose  $x_2$  you must pivot on the first equation, and the objective function will increase by  $2/3$ . If you choose  $x_3$  you must pivot on the second, and the objective function will increase by  $1/2$ . If you choose  $x_4$  you must pivot on the second, and the objective function will increase by 1.

We have here a 4 by 8 matrix and suppose its upper left entry is in box A4 of the spreadsheet.

Suppose we decide to pivot on row 2 and column  $x_4$  which would be the entry in g5.

You could then enter in say a11:  $=a4 - a5 * g4 / g5$  and copy it into the 4 by 8 area of which it is the upper left corner.

This will eliminate  $x_4$  from all the equations, but the second, and it will entirely obliterate the second equation. To bring it back you enter into a12:  $=a5 / g5$  and copy that across row 12.

That completes the pivot.

You now have a new tableau and can pivot again. To do so you can copy the entry in a11 into a18 and change the indices in it that have dollar signs to reflect the new pivot and then copy it.

When are you done?

When every coefficient in the objective function (except the constant term in the  $-b$  column) is non-positive, you are at the top and have the solution.

What else can happen?

You could end with a positive coefficient  $c_j$  in the objective function and every  $A_{ij}$  in its column is negative. This means increasing  $x_j$  moves away from all constraints. You can therefore move on forever, always increasing the objective function, without violating any constraint. This is unbounded.

And what happens when you are at the top?

The value of the objective function at the solution will be the constant term (in the  $-b$  column) in the objective function.

The value of the current basis variables will all be 0,  
The value of each of the other variables will be the corresponding  $b$  in the equation in which it occurs,

Try doing this and see what happens.

This is essentially the simplex algorithm, and this is all there is to it.

However there are several technical points as follows.

1. What do you do when you have a variable that is not constrained to be positive?

You can find a constraint to pivot on and take it out of the basis. Then leave it alone and do not consider pivoting on that row. This is because setting this variable to 0 does not represent a constraint so it should never be a basis variable.

2. What do you do when you have an equality constraint?

Then there is no point in introducing a slack variable for it. You can change one of your variables into a slack variable by using the equation to remove that variable from all the other equations and the objective function. (This could however make some of the other  $b$ 's negative)

3. What do you do if some  $b$ 's are 0?

Just change them slightly to be very small numbers. This will prevent pivoting in circles.

4. What do you do if your origin is not initially feasible? (that is, one or more  $b$ 's are negative,)

This is kind of fun. We create a new origin that is feasible.

How?

We introduce a new dimension with coordinate  $x_{new}$ . The old world represents the hyperplane  $x_{new}=1$ .

The new origin occurs where  $x_{new}=0$ .

And how do we do this?

We can add terms  $d_i(x_{new}-1)$  to the left hand side of any  $i$  equation for which  $b_i$  is negative. This will increase  $b_i$  by  $d_i$  and if we choose  $d_i$  such that  $b_i+d_i=1$  (or anything else positive) for each such  $i$  all the  $b$ 's will become positive and our new origin will be feasible.

Notice that when  $x_{new}=1$  these added terms all disappear and the equations have the same content as before.

We also add a constraint:  $x_{new} \leq 1$  and we add a new objective function,  $x_{new}$ .

So now we have added a new variable and a new constraint and a new objective function. This may seem strange but there is no harm in having two objective functions.

One strategy is to maximize the new objective function,  $x_{new}$ , first. If the maximum occurs at  $x_{new} = 1$ , then obviously  $x_{new}$  is no longer in the basis, so you can ignore its equation, and forget it and proceed to maximize the original objective function.

It can happen that the maximum of this new objective function is less than 1. In that case there is no real world at  $x_{new} = 1$ , which means that there are no feasible points at all.

There is another strategy that involves pivoting on linear combinations of the two objective functions, that is kind of fun, but we will not go into it here.

The short answer to the last question is: finding a feasible origin involves solving a modified LP starting from a feasible origin.

We first give an example for making the origin feasible when it is not.

So here is an LP for which the origin in the  $x$  variables is not feasible, because it violates two constraints.

$$\begin{aligned} 3x_1 + 2x_2 - x_3 &\leq 1. \\ -x_1 + 2x_2 + x_3 &\leq -1. \\ x_1 - x_2 + x_3 &\leq -2 \\ \text{maximize } x_1 + x_2 + x_3 \\ \text{all } x\text{'s must be positive} \end{aligned}$$

We add a new variable,  $x_4$  and a new constraint,  $x_4 \leq 1$ , and add  $2(x_4 - 1)$  to the left side of constraint 2 and  $3(x_4 - 1)$  to the left side of constraint 3. We also require  $x_4$  to be non-negative, and make  $x_4$  our new objective function.

Notice that when  $x_4 = 1$  happens, if it does,  $x_4$  will be out of the basis, and so will appear in only one equation. When  $x_4 = 1$  happens, the added terms are all 0s so that we are back to our old problem.

So our new constraints are

$$\begin{aligned} 3x_1 + 2x_2 - x_3 &\leq 1. \\ -x_1 + 2x_2 + x_3 + 2x_4 &\leq 1. \\ x_1 - x_2 + x_3 + 3x_4 &\leq 1 \\ x_4 &\leq 1 \end{aligned}$$

and we have two objective functions,  $x_4$  and  $x_1 + x_2 + x_3$ . We use the  $x_4$  objective function first. We first pivot on the  $x_4$  column and row 3 (which increases that objective function to  $1/3$ .) After a second pivot, we find that the maximum of  $x_4$  occurs at  $.375$ , which means there are no feasible points with  $x_4 = 1$  and hence the feasible region is empty.

If we change the problem to

$$3x_1 - x_2 - 2x_3 \leq 2.$$

$$\begin{aligned}
 2x_1 + 2x_2 + x_3 &\leq 10. \\
 x_1 - x_2 + x_3/10 &\leq -2 \\
 (\text{or } x_1 - x_2 + x_3/10 + 3x_4 &\leq 1 \\
 x_4 &\leq 1
 \end{aligned}$$

The following sequence of pivots seems to give a feasible answer.

s1	s2	s3	s4	x1	x2	x3	x4	b		
1	0		0	0	1	-1	-2	0	-2	
0	1		0	0	2	1	1	0	-10	
0	0		1	0	1	-1	0.1	<b>3</b>	-1	
0	0		0	1	0	0	0	1	-1	
					1	1	1	0		of1
								1		of2
1	0		0	0	1	-1	-2	0	-2	
0	1		0	0	2	<b>1</b>	1	0	-10	
0	0	0.333333	0	0.333333	-0.333333	0.033333	1	-0.333333		
0	0	-0.333333	1	-0.333333	<b>0.333333</b>	-0.033333	0	-0.666667		
0	0		0	0	1	1	1	0	0	
0	0	-0.333333	0	-0.333333	0.333333	-0.033333	0	0.333333		
1	0		-1	3	0	0	-2.1	0	-4	
0	1		1	-3	3	0	1.1	0	-8	
0	0		0	1	0	0	0	1	-1	x
0	0		-1	3	-1	1	-0.1	0	-2	
0	0		1	-3	2	0	1.1	0	2	
0	0		0	-1	0	0	0	0	1	x

That is, if you omit the rows and columns marked with an x, you find that the origin is feasible and you can pivot to a solution. The columns to be omitted are those of  $x_4$  and the slack variable  $s_4$ , and the rows are those of the  $x_4 + s_4 = 1$  equation and that of the second objective function.

### Duality

If you look at the pivot operation, you see that the step of replacing  $A_{ij}$  by  $A_{ij} - A_{i0} * A_{ij0} / A_{i0j0}$  is symmetric on transposing the A matrix. (and  $-b$  is like another column

of it and  $c$  another row.) This change describes how the rows and columns other than the  $i_0$  row and  $j_0$  column change under the pivot operation.

The change that takes place to the  $i_0$  row and to the columns that are involved in the pivot is not exactly symmetric but is almost so. The  $i_0$  row gets divided by  $A_{i_0j_0}$ . The  $j_0$  column gets changed into all 0's except for 1 in the  $i_0$  row. But the column for the variable traded for it, which used to have this exact form, now has entries given by the old  $A_{j_0}$  column divided by minus  $A_{i_0j_0}$ . If we ignore the position of this column, we see that in content it gets divided by  $-A_{i_0j_0}$  while the row got divided by  $+A_{i_0j_0}$ .

All this means that there is a symmetry that involves transposing the  $A$  matrix, but it is not the mere transpose. You have to transpose the  $A$  matrix and also reverse its sign. If you reverse the sign and transpose the equation  $A_{ij}$  new =  $A_{ij}/A_{i_0j_0}$  you get  $A_{i_0j}$  new = -

$A_{i_0j}/A_{i_0j_0}$  (the sign is changed because there are two  $A$ 's on the right which product doesn't change sign, and only one on the left. Because there is an odd number of  $A$ 's in every term of the other transformation, the change in sign does not appear: everything changes sign.

So an  $i_0j_0$  pivot has the same effect as far as content of the matrix is concerned as transposing  $A$  (which switches  $b$  to  $c$ ) and reversing the sign of its elements (so  $b$  ends up as  $-c$ ) and making a  $j_0i_0$  pivot.

We usually describe this transposed problem by keeping the sign of  $A$  fixed. We can accomplish this by multiplying the matrix for the transposed problem by  $-1$ . This has the effect of switching less than or equal with greater than or equal, and switching maximizing to minimizing. (maximizing  $-b$  is minimizing  $b$ )

The transposed problem that pivots the same way as the original LP is called the dual to it, while the original is called the primal problem. (Though if you start with the dual, the primal is dual to it.)

After multiplying by  $-1$  we find that if the primal problem is to maximize  $c \cdot x$  subject to  $A_i \cdot x \leq b_i$ , subject to  $x_j \geq 0$ , the dual is to minimize  $b \cdot y$  subject to  $\sum y_i A_{ij} \geq c_j$  and  $y_i \geq 0$  for each  $i$ .

There are two simple but remarkable features about the dual. First, it not only changes under a pivot exactly as the primal changes, but the condition for it to be a solution is identical to that for the primal. And the value of its optimum at the solution is the same for both problems.

Recall that for the primal problem you are at a solution when the  $b$ 's are all non-negative, and the  $c$ 's are non-positive. Similarly, you are feasible for the dual when the  $c$ 's are all non-positive, and at the solution then when you also have the  $b$ 's all non-negative. (which implies that making any combination of your current  $y$  basis variables positive cannot decrease your dual objective function)

This means that when you have found a solution for either problem you can read off a solution for the other. (for the primal the values of the basis variables are all 0 and the others are the b's of the equations they appear in. for the dual, the values of the dual variables corresponding to the basis variables are given by the -c's in their columns, while those corresponding to the non-basis variables are all 0.

(Notice that if you are feasible for the primal but not at a maximizing vertex, the dual problem corresponding to where you are is not feasible. But so what?)

What dual variables correspond to what?

For each  $x_j$  there is a dual slack variable,  $t_j$  and for each  $s_i$  there is a dual variable  $y_i$ .

The second wonderful property is derivable by evaluating the sum over  $i$  and  $j$  of  $y_i A_{ij} x_j$

Recall that the primal problem, after the  $s_j$  slack variables were introduced obeyed the equations  $\sum A_{ij} x_j + s_i = b_i$  The dual variables with similar slack variables  $t_j$  (which go on the other side because the inequalities are reversed, obey  $\sum y_j A_{ij} = t_j + c_j$

Putting these together, we can evaluate  $\sum y_i A_{ij} x_j$  either as  $\sum b_i y_i - \sum s_i y_i$  or as  $\sum t_j x_j + \sum c_j x_j$ . All of this tells us

$\sum b_i y_i$  (which is the dual objective function) =  $\sum c_j x_j$  (the primal objective function) +  $\sum_i s_i y_i + \sum_j t_j x_j$

Notice that if the  $x_j$  define a feasible point for the primal then the  $s_i$  and  $x_j$  are non-negative for all arguments  $i$  or  $j$ . And if the  $y_i$  define a feasible point for the dual, the  $t_j$  and  $y_i$  are non-negative as well. This means the last two sums on the right here are non-negative When the  $x$  variables are feasible and the  $y$  variables are feasible.

This tells us that the value of the dual objective function at any feasible point is greater than the value of the primal objective function at any feasible point, And it is greater by the two last sums above. The fact that the primal and dual solution optimizing values are the same then tells us that at least one of each pair  $s_i$  and  $y_i$ , and each pair  $x_j$  and  $t_j$  must be 0 at the solutions optimizing points.

We have not really talked about variables that do not have positivity constraints, and about constraints that are equalities rather than inequalities. When a constraint is an equality, its slack variable is always 0. Thus the fact that its dual can be negative changes nothing in the discussion here. Their product is always 0.

There are many problems that can be described by linear programs and the equality of primal and dual objective functions at their solution turns out to tell us interesting things about the problem. For example, the Philip Hall marriage theorem is a duality theorem, as is the more general "maximum flow = minimum cut" statement about network flows.

