Linear Programming Notes II: Graphical Solutions

1 Graphing Linear Inequalities in the Plane

You can solve linear programming problems involving just two variables by drawing a picture. The method works for problems with more than two variables, but it is hard to visualize the higher dimensional problems.

There are essentially two things you need to know in order to find graphical solutions to linear programming problems. First, you need to be able to graph the solution to linear inequalities in the plane. Second, you need to be able to see how the relationship between these points and the value of the objective function. I will discuss the first topic in this section. The next section will discuss the second topic.

As a warm up, remember how to graph a line. You find two points that are on the line and then connect them. For example, if the line is described by the equation: $2x_1 + x_2 = 2$, then you can observe that the points $(x_1, x_2) = (1, 0)$ and $(x_1, x_2) = (0, 2)$ are on the line. Connecting them leads to a straight line.

The inequality $2x_1 + x_2 \geq 2$, consists of all of the points above and to the right of the straight line. (In general, inequalities are satisfied by points on one side of the line. In order to determine which set consists of the point that satisfies the inequality, I test by checking an arbitrary point not on the line. For example, $(x_1, x_2) = (0, 0)$ does not satisfy the inequality $2x_1 + x_2 \geq 2$. Consequently the set of points that satisfies the inequality consists of the points on the side of the line $2x_1 + x_2 = 2$ that does not contain $(0, 0)$. Now you can figure out how to graph one inequality. If your linear programming problem had only one constraint, you would be in business. LPs can have many constraints (and typically do). In order to complete the process, you graph the constraints one at a time. The feasible region is the intersection. Consider, for example, the set determined by the five inequalities

$$\begin{array}{rcrcr} 2x_1 & + & x_2 & \geq & 2 \\ -2x_1 & + & x_2 & \leq & 2 \\ 4x_1 & + & x_2 & \leq & 8 \\ & & x & \geq & 0. \end{array}$$

This is region bounded by the quadrilateral pictured. (The four corners are $(0, 2)$, $(1, 0)$, $(2, 0)$, and $(1, 4)$.) (I only have so much patience for doing the work needed to include graphs. Please accept the humble offerings at the end of the notes.)

There are several things to note. Why did I say that there were five inequalities? The first three lines describe one inequality each. The fourth line describes two: $x_1 \geq 0$ and $x_2 \geq 0$.

If you have five inequalities, you would expect the feasible set of have five sides. This set has only four. The reason is that the constraint that $x_1 \geq 0$

is redundant. If you satisfy the other four constraints, then you automatically satisfy $x_1 \geq 0$. (In fact, if the first two constraints hold, then x_1 must be non-negative.)

Most of you have an intuition from high school algebra (or college linear algebra), that you should have as many variables as equations to have a system that makes sense. In this example there are two variables. There are five constraints. What is wrong with your intuition? One problem is that it is not clear what it means to make sense. In linear algebra courses, you want to have as many equations as unknowns because then (and only then) should you expect the system of equations to have a unique solution. I am dealing with inequalities, not equations. You can convince yourself that when you have two (or more) variables, it is possible to have solutions to an arbitrarily large number of distinct linear inequalities. (You describe a polygon having n sides as the set of points that satisfies n linear inequalities.) Furthermore, I do not want a unique solution to the constraints. This would mean that the feasible set had only one point. So the optimization problem would be simple. (In the diet problem, if the feasible set had only one point, that would mean that there was only one possible way in which you could meet the nutritional constraints. Of course this is impossible - you could meet the constraints by eating more of everything - but the point is that you should expect feasible sets to be large.)

In the example, the feasible set has four corners. These corners are determined by the intersection of pairs of constraints, solved as equations. That is, $(0,2)$ is the solution to

$$\begin{array}{rcrcl} 2x_1 & + & x_2 & = & 2 \\ -2x_1 & + & x_2 & = & 2, \end{array}$$

$(1,4)$ is the solution to

$$\begin{array}{rcrcl} -2x_1 & + & x_2 & = & 2 \\ 4x_1 & + & x_2 & = & 8, \end{array}$$

$(1,0)$ is the solution to $2x_1 + x_2 = 2$ and $x_2 = 0$, and $(2,0)$ is the solution to $4x_1 + x_2 = 8$ and $x_2 = 0$. This is what typically happens. That is, the feasible region of a linear programming problem has corners determined by solving subsets of the constraints as equations (here you do want to use as many constraints as you have variables). Once you have these corners, you get the feasible set by connecting the dots and identifying the region that satisfies all of the constraints.

Warnings: The feasible set may be empty. (Imagine that you replaced the constraint that $x_1 \geq 0$ with one that said that $x_1 \leq -1$.) There is nothing mathematically mysterious about this. It means that you need to be careful about which side of a constraint line is in the feasible set. The feasible set may be unbounded. That is, it may go out forever in one or more directions. (After all, having no constraints is perfectly ok.) The only way to have a problem that has an unbounded solution is to have an unbounded feasible set.

2 Graphical Solutions

Now you know how to graph the feasible set. To solve a linear programming problem graphically, that is the first thing you do. If the feasible set is empty, then stop. It does not matter what the objective function is, the linear programming problem is not feasible.

If the feasible set is non-empty, then you must decide whether the linear programming problem has a solution or is unbounded. If the problem has a solution, then you would like to find it (and, if the problem has more than one solution, you would like to find all of them). If the problem is unbounded, then you would like to be able to explain why it is unbounded. It should be clear that the linear program cannot be unbounded if the feasible set is not unbounded. It is possible for the LP to have a solution even if the feasible set is unbounded.

I will now discuss graphical solutions using the feasible region described in the first section:

$$\begin{array}{rcrcl} 2x_1 & + & x_2 & \geq & 2 \\ -2x_1 & + & x_2 & \leq & 2 \\ 4x_1 & + & x_2 & \leq & 8 \\ & & x & \geq & 0. \end{array}$$

For a start, assume that the objective function (I call it x_0) is $x_1 + 2x_2$. (I will discuss other possible objective functions later.)

The next step is to graph a **level set** of the objective function. A level set of a function is the set of points at which the function takes on the same value. For the example, a level set is a set of points for which $x_1 + 2x_2 = c$ for some constant c. Level sets of linear functions of two variables are straight lines (level sets of linear functions of three variables are planes; level sets of linear functions of more than two variables are flat things called hyperplanes). Concretely, the points at which $x_1 + 2x_2 = 2$ consist of a straight line through $(2, 0)$ and $(0, 1)$.

Superimpose a level set of the objective function (which is a line) on the graph of the feasible region. When you do so, either the line intersects with the graph or it does not. For example, the line $x_1 + 2x_2 = 2$ does intersect the feasible region (at the point $(2, 0)$ for example). The line $x_1 + 2x_2 = -2$ does not intersect the feasible region (it lies below the region). Neither does the line $x_1 + 2x_2 = 20$ (it lies above the region).

If the level set $x_1 + 2x_2 = c$ intersects the feasible set, that means that there exists feasible points that make the objective function's value equal to c. If the level set does not intersect the feasible set, then it is not possible to make the objective function's value equal to c. So it is possible to make the objective function's value equal to 2, while it is not possible to make the value equal -2 or 20. These observations are important. They tell us that the value of the optimization problem is at least 2 but no more than 20. Our job is to find the highest value of the objective function subject to staying within the feasible set. Geometrically, we want to find a level set of x_0 that intersects the feasible set, with the additional property that no higher level set intersects the feasible set. The level set $x_1 + 2x_2 = 2$ intersects interior of the feasible set. (Interior is a technical term, but I hope that its intuitive meaning is clear.) This means that level sets for higher values still intersect the feasible region. Now consider the level set $x_1 + 2x_2 = 9$. This level set describes points that make $x_0 = 9$. This level set is parallel to the others. It intersects the feasible region. It intersects the region in just one place, however (the point $(1, 4)$). We can conclude that it is possible to satisfy constraints and make the objective function's value be (at least) 9. We can further conclude that it is not possible to make the objective function any higher and still satisfy the constraints. It follows that the solution to the problem is $(x_1, x_2) = (1, 4)$. It leads to the optimal value 9. This is the solution to the linear programming problem.

You should experiment with this procedure using alternative objective functions. You should be able to construct objective functions that yield solutions at any of the corners of the feasible set.

In brief, here is the process.

1. Graph feasible set. If feasible set is empty, then stop. The problem is infeasible. Otherwise continue.
2. Graph a level set of the objective function.
3. Shift the level set (parallel movement) until it intersects the feasible region.
4. Continue to shift the level set until it reaches the maximum value the intersects the feasible region.

You follow the same steps for a minimization problem, taking care to move the objective function in the opposite direction. You know which direction increases the objective function value by drawing two level sets and comparing (the direction of increase never changes). In the example, the level set $x_1 + 2x_2 = 9$ lies above and to the right of the level set $x_1 + 2x_2 = 2$; you always increase the objective function (in this example) by moving up and to the right.

If your feasible set is unbounded, then it may be that the linear programming problem is unbounded. You will be able to see this graphically if level sets for arbitrarily large values of the objective function continue to intersect the feasible region.

The graphical method illustrates an important, and general, property of solutions to linear programming problems. *If a linear programming problem has a solution, then it has a solution at a "corner" of the feasible set.* Graphically, this fact follows from the observation that if the level set of the objective function does not intersect a corner of the feasible set, then you can move the level set a little bit in either direction and still intersect the feasible region. You cannot be at either a minimum or a maximum. It is a really important fact. It means that if you want to solve a LP and you know the corners of the feasible set, then all you need to do is plug the corners into the objective function and pick the best one. (In the example, the feasible set has only four corners: $(0, 2)$, $(1, 0)$, $(2, 0)$, and $(1, 4)$.) This realization turns out be the idea behind the simplex algorithm that can be used to solve any linear programming problem.

Linear programming problems may have solutions that are not at corners. For example, if $x_0 = 4x_1 + x_2$, then all points on the segment connecting $(1, 4)$

to $(2,0)$ solve the LP. Here the level sets of the objective function are parallel to the upper right boundary of the feasible set. This example does not contradict the fact. Points $(1,4)$ and $(2,0)$ are corners. (It is generally true that if a linear programming problem has two solutions, then everything on the segment connecting the two solutions is also a solution.)

If the feasible set is unbounded, then the "trick" of testing only at the corners of the feasible region will tell you the solution to the problem *if* the LP has a solution. The problem may be unbounded.

Experiment with unbounded problems. For example, take the feasible set that I have used throughout these notes and omit the constraint that $4x_1 + x_2 \leq 8$. You obtain an unbounded feasible set. Even though the feasible set is unbounded, certain problems still have a solution. For example, the solution to $\min x_1 + x_2$ subject to:

$$\begin{array}{rcrcl} 2x_1 & + & x_2 & \geq & 2 \\ -2x_1 & + & x_2 & \leq & 2 \\ & & x & \geq & 0 \end{array}$$

is $(1,0)$. It should be clear, however, that if you tried to solve $\max x_1 + x_2$ subject to these constraints you cannot find a solution. For example, since the point $(z,0)$ is in the feasible set for all $z \geq 1$, it is possible to make the objective function larger than anything (just let $z > M$). You can see this graphically by noting that the level sets of $x_1 + x_2$ intersect the feasible region at the point $(z,0)$.

If I ask you to solve a linear programming problem graphically, then you should give one of three answers. If the problem is not feasible, then you demonstrate that the graph of the constraint set is empty (and state that the problem is infeasible). If the problem has a solution, then you should state the solution (or solutions, if there is more than one); show a level set of the objective function that intersects the feasible region at the solution; and show that if you increase the objective function the corresponding level set would no longer intersect the feasible region. If the problem is unbounded, then you should show how you can find a feasible point that makes the objective function larger (or, in the case of a minimization problem, smaller) than any arbitrarily chosen value M.