Linear Programming Class Notes

1 A Quick Example

Let $f(x,y) = x + y$. Find the maximum of f given that

$$\begin{aligned} 5x + 3y &\leq 15 \\ x &\geq 0 \\ y &\geq 0 \end{aligned}$$

To solve this problem, we consider a graph of the set of points that satisfy all the constraints. In Figure 1, we see that the set of points that satisfy all the constraints (found by plotting $5x + 3y = 15$, $x = 0$, $y = 0$) is a triangle. Now

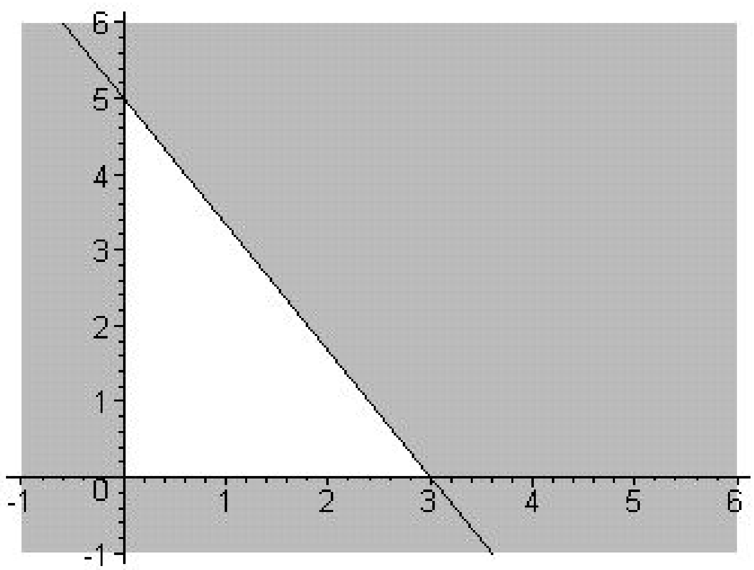

Figure 1: The set of all points that satisfy all of the constraints in our first example. This set is called the feasible set.

consider a plot of $x + y = k$ for different values of k. Then:

1. Any point on this line gives the same value for $f(x,y)$ (which is k).
2. Increasing k means that we shift the line upwards.
3. To find the maximum, we shift up as high as we can with the condition that the line still intersects the triangle. In this case, it is easy to see that the maximum occurs at $(0,5)$ (which also gives $k = 5$).

These are the types of problems we will consider in this chapter.

2 Theoretical Background

A *linear program* (LP) is a problem of the form: Maximize (or Minimize) $f(\mathbf{x}) = c_1x_1 + c_2x_2 + \ldots + c_nx_n$ (or $f(\mathbf{x}) = \mathbf{c}^T\mathbf{x}$) subject to the constraints that:

$$\begin{array}{lllll} a_{11}x_1 & +a_{12}x_2 & +\ldots & +a_{1n}x_n & \leq b_1 \\ a_{21}x_1 & +a_{22}x_2 & +\ldots & +a_{2n}x_n & \leq b_2 \\ \vdots & & & \vdots & \vdots \\ a_{m1}x_1 & +a_{m2}x_2 & +\ldots & +a_{mn}x_n & \leq b_n \end{array}$$

and $x_i \geq 0$ for all i. More compactly, we can write the constraints as $A\mathbf{x} \leq \mathbf{b}$. It is normal to simply assume all variables are positive, and not to include them in the system of equations, although this could be done.

Some definitions:

1. The function $f : \mathbb{R}^n \to \mathbb{R}$ is called *the objective function*.
2. A point that satisfies all of the constraints is called a *feasible solution*.
3. The set of all points that satisfy all constraints is called the feasible set, or polytope.
4. Suppose we have $\mathbf{x} \in \mathbb{R}^n$. A vertex of the polytope is the intersection of n constraints. A vertex is a single point, and can be found by converting the n inequality constraints to n equalities (when this happens, the constraints are said to be tight).
5. An $\epsilon-$neighborhood about a point $\mathbf{a} \in \mathbb{R}^n$ is the set of points that have distance less than ϵ to $\mathbf{a}$:
$$N_\epsilon(\mathbf{a}) = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{a}\|_2 < \epsilon\}$$
6. A point $\mathbf{x}$ is called an *interior point* of a set S if there exists an $N_\epsilon(\mathbf{x})$ that is totally contained in S.
7. A point $\mathbf{x}$ is called a *boundary point* of a set S if any ball about $\mathbf{x}$ contains both points of S and points not in S.
8. A set S is $\mathbb{R}^n$ is said to be *bounded* if S can be completely contained within a ball with finite radius about the origin.
9. A set S is *closed* if it contains all of its boundary points.

Class Exercise: From the last example, identify the feasible set, the set of boundary points, the set of interior points. Identify the vertices, and which constraints were tight to form each vertex.
Class Exercise: In $\mathbb{R}$, $\mathbb{R}^2$ and $\mathbb{R}^3$, give examples of sets that are closed and not closed.
Class Exercise: If $\mathbf{x} \in \mathbb{R}^3$, then each tight constraint is a plane which cuts 3-d into two regions. Given any three planes, do they necessarily define a vertex?

2.1 A Reminder of some Calculus

Before continuing farther, let us put this material into context. Linear programming is an optimization problem, therefore the Extreme Value Theorem (EVT) of Calculus will come into play. In one dimension, $y = f(x)$, the EVT says that:

> Let f be continuous on $[a, b]$. Then f attains a global maximum and a global minimum on $[a, b]$. This occurs at either (i) a critical point[1] of f or (ii) at an endpoint.

If f is a function of more than one variable, then we make the appropriate conversions:

1. The domain changes from an interval to a closed, bounded region.
2. Change $f'(x)$ to the gradient of f
3. Change the endpoints to the boundary points.

We are able to go even farther, since our function f is especially simple: If $f(x_1, \ldots, x_n) = c_1x_1 + c_2x_2 + \ldots c_nx_n$, then the gradient is just the vector $\mathbf{c}$. Therefore, we won't be considering the critical points of f- only the boundary points. We summarize this with the following theorem:

Theorem: Let $D = \{\mathbf{x} \in \mathbb{R}^n \,|\, A\mathbf{x} \le \mathbf{b}\}$ be not empty, closed, and bounded. Then the LP problem has a solution, and that solution lies along the boundary of D.

The boundary, ∂D, is contained within the set made up when each single inequality is changed to an equality. Then:

$$\partial D \subset \{\mathbf{x} \in R^n \,|\, (A\mathbf{x})_i = \mathbf{b}_i, \quad i = 1, 2, \ldots, n\}$$

In our first example, the boundary was made up of a subset of the set of points where $x = 0$, $y = 0$, and $5x + 3y = 15$. Notice that we cannot say that the boundary is made up of points from $A\mathbf{x} = \mathbf{b}$, although that may be possible - recall that the solution set to $A\mathbf{x} = \mathbf{b}$ has three possibilities.

2.2 The (Non-empty) Feasible Set is Convex

Before continuing, you might recall the definition of a *convex set*: Let $\mathbf{x}_1$ and $\mathbf{x}_2$ be any two points of the set (in $\mathbb{R}^n$). Then any point on the line between, denoted:

$$\mathbf{x} = t\mathbf{x}_1 + (1 - t)\mathbf{x}_2, \quad 0 \le t \le 1$$

is also in the set.

Let us show that D is convex, which follows from the linearity of $A\mathbf{x}$: We show that, if $\mathbf{x}_1$ and $\mathbf{x}_2$ are in D, then so is any point in between. Here we go: let $\mathbf{x}_1$ and $\mathbf{x}_2$ be in D. Then it follows that

$$A\mathbf{x}_1 \le \mathbf{b}, \qquad A\mathbf{x}_2 \le \mathbf{b}$$

[1] Recall that a critical point of f is where $f'(x) = 0$ or where $f'(x)$ does not exist

Given that $\mathbf{x} = t\mathbf{x}_1 + (1-t)\mathbf{x}_2, \quad 0 \le t \le 1$, we show that $A\mathbf{x} \le \mathbf{b}$:

$$A\mathbf{x} = A(t\mathbf{x}_1 + (1-t)\mathbf{x}_2) = tA\mathbf{x}_1 + (1-t)A\mathbf{x}_2 \le t\mathbf{b} + (1-t)\mathbf{b} = \mathbf{b}$$

Therefore, for any two points in D, the line segment connecting them is also in D.

Is it always possible to write a given point $\mathbf{x} \in D$ in terms of two other points; that is, as $t\mathbf{x}_1 + (1-t)\mathbf{x}_2$, where $0 < t < 1$? If we restrict ourselves to the plane, you can convince yourself that just about every point can be written as point in between two others. But can a corner point (vertex) be written that way? No- See Exercises (1) and (2).

Let l be any line passing through the convex set D. What kinds of line segments will be produced by intersecting l with D? We'll have one of three choices:

- A closed line segment
- A ray.
- All the points on l.

The second two choices might occur if D is unbounded. If D is bounded, we will only have the first choice (see Exercise 5).

Consider the maximization/minimization of $f(\mathbf{x}) = \mathbf{c}^T\mathbf{x}$ restricted to the line l. To be more precise, let us take l as the line defined through two arbitrary (but fixed) points $\mathbf{u}$ and $\mathbf{v}$,

$$l = \{\mathbf{x} \mid \mathbf{x} = t\mathbf{u} + (1-t)\mathbf{v}, \text{ for some } t \in R\}$$

If we restrict the domain of f to the line l, we get (by substitution):

$$f(\mathbf{x}) = \mathbf{c}^T\left(t\mathbf{u} + (1-t)\mathbf{v}\right) = t(\mathbf{c}^T\mathbf{u} - \mathbf{c}^T\mathbf{v}) + \mathbf{c}^T\mathbf{v} \tag{1}$$

so that we can now write f (restricted to the line l) very simply as $f(t) = mt + b$ for scalars m and b.

Now it is a simple matter to describe where f has a maximum or minimum- its graph is a line, so it is either strictly increasing ($m > 0$), strictly decreasing ($m < 0$), or it is a horizontal ($m = 0$). In all cases *the maximum and minimum of f occur at the endpoints* (if there are endpoints). In particular, this shows that if D is a bounded set made up of vertices and edges, the max/min of f must occur at a vertex. This is summarized by the following theorem:

The Fundamental Theorem of Linear Programming. Given an LP problem with a nonempty, bounded feasible set D, the maximum and minimum of f occurs at a vertex in D.

This gives us a new method for finding the optimal value of f- we could examine all of the edge intersections. In our first example, we can list them:

$$(0,0), (3,0), (0,5)$$

The maximum value of $x+y$ occurs at $(0, 5)$. In the plane, these are combinations of two equations at a time. If we were in three dimensions, this would mean taking three equations at a time, and so on.

To summarize, we have shown that the solution to the LP problem exists if D is not empty, closed and bounded. We have also shown that D is a convex set, and that the solution will be found at an intersection of two or more of the edges.

Is it possible that the LP problem has no solution? Yes- it might be that all the constraints cannot be satisfied simultaneously, in which case D is the empty set. It might also happen that D is not bounded.

Is it possible that the solution to the LP problem is not unique? Yes- Consider our very first example. Suppose now that the isclines of f have the same slope as the hypotenuse of the triangle. Then the function has a maximum at every point of the hypotenuse that lies in the feasible region.

For two dimensional problems, we will focus on our geometric technique for finding the max/min- Draw the feasible set, then draw isoclines of f.

3 Drawing the Feasible Set

Maple has a nice feature that allows us to draw the feasible set for a two dimensional problem. We'll use `inequal` from the `plots` library (that is, type `with(plots):` before using `inequal`). Here's an example:

```
>> with(plots):
>> inequal( { x+y>=1/2, x-y<=1,y<=2,x>=0,y>=0}, x=-3..3,
   y=-3..3, optionsfeasible=(color=yellow),
   optionsexcluded=(color=white) );
```

4 Matlab's Optimization Toolbox

We can solve higher dimensional problems using Matlab's optimization toolbox. The relevant command is `linprog`.

The function `linprog` solves the following general problem:

$$\min_{\mathbf{x}} \quad \mathbf{c}^T\mathbf{x} \tag{2}$$

such that:

$$\begin{aligned} A\mathbf{x} &\le \mathbf{b} \\ A_{\text{eq}}\mathbf{x} = \mathbf{b}_{\text{eq}} & \qquad \text{Optional Argument} \\ lb \le \mathbf{x} \le ub & \qquad \text{Optional Argument} \end{aligned}$$

We will call `linprog` one of several ways, depending on what we want to input, and what we want to output. Here's a basic example:

4.1 Worked Example

You have \$60,000.00 to invest in Certificates of Deposit (CDs) and mutual funds. The CDs pay an annual yield of 5%. The mutual funds pay an annual yield of 9%, with a minimum investment of \$10,000.00. As a safety net, you decide that you need to put in at least twice as much into CDs as mutual funds. How much should you invest in each in order to maximize your annual yield?
SOLUTION:

First, our variables: Let x_1 be the amount we invest in CDs, and let x_2 be the amount in mutual funds. We want to maximize our annual yield, which is $0.05x_1 + 0.09x_2$.

We have the following constraints, taken as they appear in the description:

- $x_1 + x_2 \leq 60000$
- $x_2 \geq 10000$
- $x_1 \geq 2x_2$

We of course need x_1, x_2 to be non-negative. This gives the following LP problem:

$$\max_{x_1,x_2} 0.05x_1 + 0.09x_2 \text{ such that } A\mathbf{x} \leq \mathbf{b}, \text{ where}$$

$$A = \begin{bmatrix} 1 & 1 \\ 0 & -1 \\ -1 & 2 \\ -1 & 0 \\ 0 & -1 \end{bmatrix} \qquad \mathbf{b} = \begin{bmatrix} 60000 \\ -10000 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

And in Matlab:

```
>> A=[1 1;0 -1;-1 2;-1 0;0 -1];
>> b=[60000;-10000;0;0;0];
>> c=[-0.05 -0.09];
>> [f,fopt]=linprog(c,A,b)
```

The output argument `f` will return the minimizer (the value of the domain where the minimum occurs). The argument `fopt` will return the value of the minimum.

The solution is \$40,000 for CDs, \$20,000 for mutual funds, and we get a maximum annual yield of \$3800.00

We can also determine which of the inequalities were "tight" and which were not. If f is the minimizer, then Af gives the values of the constraints at the minimizer- Compute Af and see which values occurred at the boundary, and which were in the interior. In this case, we should see that the first and third constraints were tight, but the second was not.

4.2 Other Matlab Arguments

The full `linprog` command, with most of the options is:

```
[f,fopt,exitflag,output]=linprog(c,A,b,Aeq,beq,LB,UB)
```

The inputs were described in our general equation 2. The outputs are:

- f is the minimizer.
- $fopt$ is the value of the minimum.
- `exitflag` is an integer describing why Matlab terminated the program (See `help linprog` for a description).
- `output` is a Matlab structure with the details for how Matlab arrived at the solution.

We don't need to specify all the input arguments. For example, another way to set up our example problem would be to separate the non-negative constraints on the decision variables so that:

$$A = \begin{bmatrix} 1 & 1 \\ 0 & -1 \\ -1 & 2 \end{bmatrix} \qquad \mathbf{b} = \begin{bmatrix} 60000 \\ -10000 \\ 0 \end{bmatrix}$$

and set `lb=[0;0];`. Call Matlab:

```
[f,fopt,exitflag,options]=linprog(c,A,b,[],[],lb);
```

In this case, `exitflag`$= 1$, which means that the program terminated normally at a solution. The `options` structure says that it took 5 iterations to come to the solution.

5 Exercises

1. Let D be the set of points satisfying some constraints, $A\mathbf{x} \leq \mathbf{b}$. Let R, S be two points interior to D so that they satisfy:

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

   Show directly that, any point in between R and S must also satisfy the strict inequalities. This shows that any point in between two interior points is also interior.

2. Let D be as defined previously. Let R, S be two points in D that satisfy the equations:

   $$(A\mathbf{x})_j = b_j, \text{ for some j}$$
   $$(A\mathbf{x})_i < b_i, \quad i \neq j$$

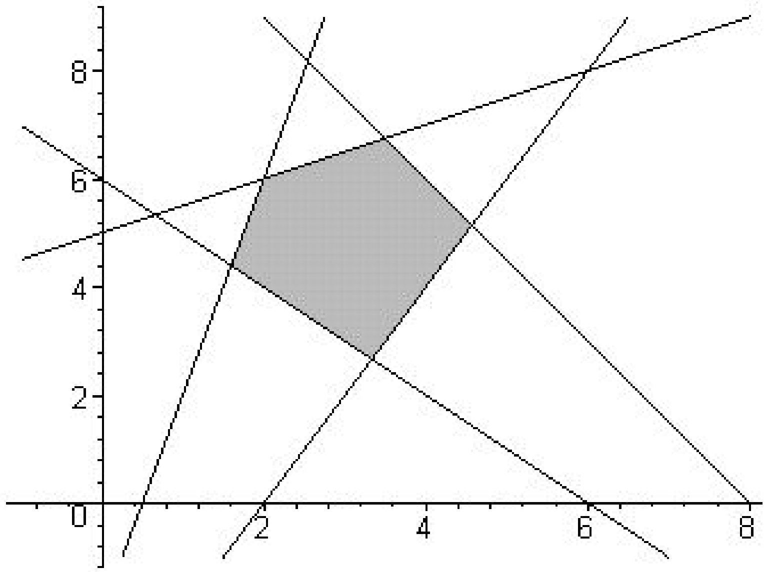

Figure 2: The shaded region is the feasible region. This is for Question 6.

Show that, for any point in between R and S, it also must satisfy the same constraints- only one equality, the rest strict inequalities. This shows that a vertex (which satisfies at least two equality constraints) cannot be written as a point interior to two others.

3. True or False, and explain (for example, if false, provide a counterexample):

   (a) The following two sets are equal:

      - The set of points where at least one constraint is an equality.
      - The set of boundary points.

   (b) The following two sets are equal: Suppose we have n constraints

      - The set of all vertices.
      - The set of all solutions to n rows being equal (that is, all of the possible solutions to all of the possible n equality constraints).

   (c) It is possible that there are an infinite number of solutions to an LP problem.

   (d) The intersection of a line l with a set D constructed by the points $A\mathbf{x} \leq \mathbf{b}$ could be two distinct line segments.

4. In Figure 2, how many vertices are there? How many are actually feasible? How many variables would be in the objective function? How many constraints?

5. First, solve the problem below by graphing the feasible region by hand or using Maple, then draw in several isoclines of the objective function. Verify your solution using Matlab:

   A dog breeder feeds her dogs a combination of two types of food, A and B. She is concerned that her dogs require a certain combination of 4 nutritional factors, a, b, c, d each month. The chart below shows the amount of each factor per bag of food, together with the minimum nutritional requirements:

	a	b	c	d
A	3	2	1	2
B	2	4	3	1
Min Req	28	30	20	25

   The costs per bag are \$50.00 for A, \$40.00 for B. How many bags of each type of dog food should she blend to get the minimum monthly nutritional requirements, and minimize her cost? At the minimum, which constraints are tight, and which are not?

6. First, solve the problem below by graphing the feasible region by hand or using Maple, then draw in several isoclines of the objective function. Verify your solution using Matlab:

   Its time to reinvest some maturing funds. You have \$55000.00 to invest in Doug's fund and Mary's fund.

   Doug's fund is cheaper at \$200.00 per note with an expected return of 9%.

   Mary's fund has a buy in of \$300.00 per note with an expected return of 15%. However, Mary's fund is riskier than Doug's fund- For diversification of your portfolio, you want to buy at least twice as much into Doug's fund as Mary's.

   There are currently 200 notes available for Doug's fund and 100 notes for Mary's fund. How much of each should you buy?

7. First, solve the problem below by graphing the feasible region by hand or using Maple, then draw in several isoclines of the objective function. Verify your solution using Matlab:

   Shady Lane Grass Seed Company sells two brands, Evergreen and Quick-Green in its grass seed business. Each bag contains certain combinations of Fescue, Rye and Bluegrass seed, as summarized in the table (in pounds per bag), as well as the total number of pounds of each type of grass seed that the company normally stocks:

	Fescue	Rye	Bluegrass
Evergreen	3	1	1
QuickGreen	2	2	1
Total Avail	1200	800	450

If the company can sell Evergreen for \$2.00 of profit per bag, and \$3.00 of profit on QuickGreen, how many bags of each should the company produce to maximize its profit?

8. The following is not strictly a linear programming problem, but we will use the following when we use a linear program to find a line of best fit.

   Definition: Let $\mathbf{x}$ be a vector in $\mathbb{R}^n$. The $p-$norm for $\mathbf{x}$ is defined as:

   $$\|\mathbf{x}\|_p = \left(|x_1|^p + |x_2|^p + \ldots |x_n|^p\right)^{1/p}$$

   (a) Find the $2-$, $3-$, and $10-$norms of $\mathbf{x} = [3, -1]^T$. Guess at what the norm is as $p \to \infty$.

   (b) If $\mathbf{x} \in \mathbb{R}^2$, use Maple to plot the set of points where

   $$\|\mathbf{x}\|_p = 1$$

   for $p = 1, 2, 3$ and 10. What do you notice about the shapes of these (closed) neighborhoods about the origin? What will the shape be as $p \to \infty$?

   Compare your answer with the following definition of the ∞ norm:

   $$\|\mathbf{x}\|_\infty = \max_i \ |x_i|$$