Introduction to Linear Programming

Linear Programming is the study of optimization problems in which the objective function and all constraints are linear.

Linear Programming is the study of optimization problems in which the objective function and all constraints are linear.

A **linear function** in n variables is one of the form

$$f(x_1, x_2, \ldots, x_n) = c_1x_1 + c_2x_2 + \cdots + c_nx_n$$

for some constants $c_1, c_2, \ldots, c_n$.

Linear Programming is the study of optimization problems in which the objective function and all constraints are linear.

A **linear function** in n variables is one of the form

$$f(x_1, x_2, \dots, x_n) = c_1x_1 + c_2x_2 + \cdots + c_nx_n$$

for some constants $c_1, c_2, \dots, c_n$.

A **linear inequality** in n variables if one of the form

$$f(x_1, x_2, \dots, x_n) \leq b \text{ or } f(x_1, x_2, \dots, x_n) \geq b$$

where f is a linear function and b is a constant.

Linear Programming is the study of optimization problems in which the objective function and all constraints are linear.

A **linear function** in n variables is one of the form

$$f(x_1, x_2, \ldots, x_n) = c_1 x_1 + c_2 x_2 + \cdots + c_n x_n$$

for some constants $c_1, c_2, \ldots, c_n$.

A **linear inequality** in n variables if one of the form

$$f(x_1, x_2, \ldots, x_n) \leq b \text{ or } f(x_1, x_2, \ldots, x_n) \geq b$$

where f is a linear function and b is a constant.

Linear Programming is concerned with optimizing a linear function subject to a set of constraints given by linear inequalities.

A **linear program** (an **LP**) is a linear optimization problem taking the following form:

Maximize (or minimize)
$f(x_1, x_2, \ldots, x_n) = c_1x_1 + c_2x_2 + \cdots + c_nx_n$ subject to

$$
\begin{aligned}
a_{1,1}x_1 + a_{1,2}x_2 + \cdots a_{1,n}x_n &\ \substack{\leq \\ \geq}\ b_1 \\
a_{2,1}x_1 + a_{2,2}x_2 + \cdots a_{2,n}x_n &\ \substack{\leq \\ \geq}\ b_2 \\
&\vdots \\
&\vdots \\
a_{m,1}x_1 + a_{m,2}x_2 + \cdots a_{m,n}x_n &\ \substack{\leq \\ \geq}\ b_m \\
x_1, x_2, \ldots x_n &\geq 0
\end{aligned}
$$

The inequalities, except for the last one, can be greater than or equal or less than or equal.

This general form is often expressed with vectors and matrices:

This general form is often expressed with vectors and matrices:

$$\text{Maximize (or minimize) } c^T\vec{x}$$

$$\text{subject to } A\vec{x} \underset{\geq}{\overset{\leq}{=}} \vec{b}$$

$$\text{and } \vec{x} \geq 0$$

This general form is often expressed with vectors and matrices:

$$\begin{aligned} \text{Maximize (or minimize) } & c^T\vec{x} \\ \text{subject to } & A\vec{x} \begin{matrix} \le \\ = \\ \ge \end{matrix} \vec{b} \\ \text{and } & \vec{x} \ge 0 \end{aligned}$$

This looks very concise but it obscures a lot of things we will want to talk about, so I will not use this form at all. You will run across it in some papers and books on the subject.

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

One unit of salsa requires 5 tomatoes and 1 clove of garlic.

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

One unit of salsa requires 5 tomatoes and 1 clove of garlic.

One unit of guacamole requires 1/2 tomato and 4 cloves of garlic.

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

One unit of salsa requires 5 tomatoes and 1 clove of garlic.

One unit of guacamole requires $1/2$ tomato and 4 cloves of garlic.

We only have 30 tomatoes and 20 cloves of garlic.

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

One unit of salsa requires 5 tomatoes and 1 clove of garlic.

One unit of guacamole requires 1/2 tomato and 4 cloves of garlic.

We only have 30 tomatoes and 20 cloves of garlic.

We have an unlimited supply of all other ingredients (salt, cilantro, lime juice, etc.)

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

One unit of salsa requires 5 tomatoes and 1 clove of garlic.

One unit of guacamole requires $1/2$ tomato and 4 cloves of garlic.

We only have 30 tomatoes and 20 cloves of garlic.

We have an unlimited supply of all other ingredients (salt, cilantro, lime juice, etc.)

How many units of salsa and guacamole should we make to maximize the total number of units (salsa+guacamole) we make?

LP example

Suppose we are making salsa and guacamole (dips or sauces originally from Mexico).

One unit of salsa requires 5 tomatoes and 1 clove of garlic.

One unit of guacamole requires $1/2$ tomato and 4 cloves of garlic.

We only have 30 tomatoes and 20 cloves of garlic.

We have an unlimited supply of all other ingredients (salt, cilantro, lime juice, etc.)

How many units of salsa and guacamole should we make to maximize the total number of units (salsa+guacamole) we make?

Fractional units are okay.

Let x_1 be the number of units of salsa we make.

Let x_2 be the number of units of guacamole we make.

Let x_1 be the number of units of salsa we make.

Let x_2 be the number of units of guacamole we make.

Then the quantity we want to **maximize** is $x_1 + x_2$.

This is our **objective function**.

Let x_1 be the number of units of salsa we make.

Let x_2 be the number of units of guacamole we make.

Then the quantity we want to **maximize** is $x_1 + x_2$.

This is our **objective function**.

We see that we cannot simply make x_1 and x_2 huge due to our limited amount of garlic and tomatoes.

If we make x_1 units of salsa and x_2 units of guacamole, then the amount of tomatoes we'll need is

$$5x_1 + \frac{1}{2}x_2.$$

If we make x_1 units of salsa and x_2 units of guacamole, then the amount of tomatoes we'll need is

$$5x_1 + \frac{1}{2}x_2.$$

Since we only have 30 tomatoes, we have the following **constraint**

$$5x_1 + \frac{1}{2}x_2 \leq 30.$$

If we make x_1 units of salsa and x_2 units of guacamole, then the amount of tomatoes we'll need is

$$5x_1 + \frac{1}{2}x_2.$$

Since we only have 30 tomatoes, we have the following **constraint**

$$5x_1 + \frac{1}{2}x_2 \leq 30.$$

Similarly, since we have a limited amount of garlic, we have another **constraint**:

$$x_1 + 4x_2 \leq 20.$$

If we make x_1 units of salsa and x_2 units of guacamole, then the amount of tomatoes we'll need is

$$5x_1 + \frac{1}{2}x_2.$$

Since we only have 30 tomatoes, we have the following **constraint**

$$5x_1 + \frac{1}{2}x_2 \leq 30.$$

Similarly, since we have a limited amount of garlic, we have another **constraint**:

$$x_1 + 4x_2 \leq 20.$$

We also require $x_1, x_2 \geq 0$ since we cannot make a negative amount of salsa or guacamole.

Thus, the LP we wish to solve is:

$$\begin{aligned} &\text{Maximize } x_1 + x_2 \\ &\text{subject to:} \\ &\qquad 5x_1 + \frac{1}{2}x_2 \leq 30 \\ &\qquad x_1 + 4x_2 \leq 20 \\ &\qquad x_1, x_2 \geq 0 \end{aligned}$$

Since the non-negative constraints are always with us, we will often refer to such an LP as having two variables and two constraints.

In this course, we will not be concerned with the methods used or solving LPs.

In this course, we will not be concerned with the methods used or solving LPs.

Instead, we will focus on lots of different applications of the LP idea, and we will use software to solve the LPs for us.

In Math 407, you will learn methods for solving LPs.

Just once, though, let's look at how we might solve our salsa and guacamole LP "by hand".

Just once, though, let's look at how we might solve our salsa and guacamole LP "by hand".

A classic approach (for two-variable LPs) is to consider the LP graphically.

Just once, though, let's look at how we might solve our salsa and guacamole LP "by hand".

A classic approach (for two-variable LPs) is to consider the LP graphically.

We begin by making a sketch of our inequalities.

On a set of x_1, x_2 axes, we draw the lines that define our constraints, and indicate with shading which side of the line satisfies the constraints.

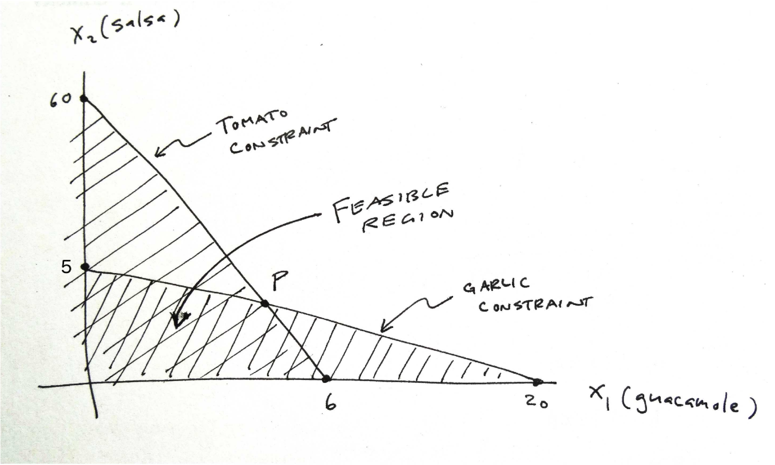

The **feasible region** is the set of points (x_1, x_2) that satisfy both constraints. This is the doubly-shaded portion in the picture.

The **feasible region** is the set of points (x_1, x_2) that satisfy both constraints. This is the doubly-shaded portion in the picture.

The points in the feasible region represent combinations of amounts of salsa and quacamole that we can make.

The **feasible region** is the set of points (x_1, x_2) that satisfy both constraints. This is the doubly-shaded portion in the picture.

The points in the feasible region represent combinations of amounts of salsa and quacamole that we can make.

Other points are not feasible.

We note that our objective function $x_1 + x_2$ is an increasing function of x_1 and x_2.

We note that our objective function $x_1 + x_2$ is an increasing function of x_1 and x_2.

This means that, if (x_1, x_2) is a point in the feasible region, it will not yield the maximum value of $x_1 + x_2$ unless the point is "pushed up" against one of the constraint lines.

We note that our objective function $x_1 + x_2$ is an increasing function of x_1 and x_2.

This means that, if (x_1, x_2) is a point in the feasible region, it will not yield the maximum value of $x_1 + x_2$ unless the point is "pushed up" against one of the constraint lines.

That is, if (x_1, x_2) is in the feasible region, and not on one of the constraint lines, then we can increase the value of $x_1 + x_2$ by increasing x_1 or x_2. Hence, that point does not yield the maximum.

We note that our objective function $x_1 + x_2$ is an increasing function of x_1 and x_2.

This means that, if (x_1, x_2) is a point in the feasible region, it will not yield the maximum value of $x_1 + x_2$ unless the point is "pushed up" against one of the constraint lines.

That is, if (x_1, x_2) is in the feasible region, and not on one of the constraint lines, then we can increase the value of $x_1 + x_2$ by increasing x_1 or x_2. Hence, that point does not yield the maximum.

So, the maximum must occur **on** one of the line segments bounding the feasible region.

Where does the maximum occur?

Where does the maximum occur?

Let's consider the function $f(x_1, x_2) = x_1 + x_2$.

Where does the maximum occur?

Let's consider the function $f(x_1, x_2) = x_1 + x_2$.

If we draw level curves of this function, i.e., curves given by $f(x_1, x_2) = k$ for various values of k, we see that these are all lines with slope -1.

Let's draw one of these level curves, the one given by $f(x_1, x_2) = 0$. Here it is in red.

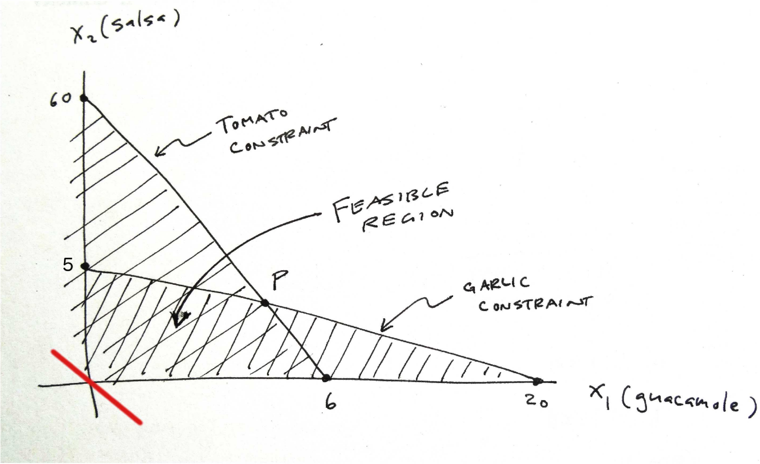

Let's draw one of these level curves, the one given by $f(x_1, x_2) = 0$. Here it is in red.

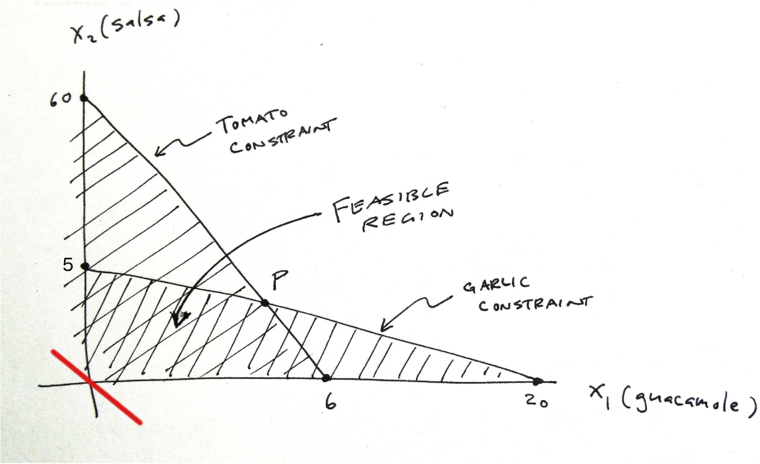

I haven't been careful to make my picture to scale, so we'll need to be careful what conclusions we make here.

Now, if we increase k, the line $f(x_1, x_2) = k$ moves away from the origin into the first quadrant.

Now, if we increase k, the line $f(x_1, x_2) = k$ moves away from the origin into the first quadrant.

For a while, the line intersects the feasible region: there are combinations of salsa and guacamole that we can make to achieve a total output of k units.

Now, if we increase k, the line $f(x_1, x_2) = k$ moves away from the origin into the first quadrant.

For a while, the line intersects the feasible region: there are combinations of salsa and guacamole that we can make to achieve a total output of k units.

But eventually, the line $f(x_1, x_2) = k$ will intersect the feasible region for the last time, and then for any larger k will not intersect the feasible region at all.

Now, if we increase k, the line $f(x_1, x_2) = k$ moves away from the origin into the first quadrant.

For a while, the line intersects the feasible region: there are combinations of salsa and guacamole that we can make to achieve a total output of k units.

But eventually, the line $f(x_1, x_2) = k$ will intersect the feasible region for the last time, and then for any larger k will not intersect the feasible region at all.

We want to figure out what this largest value of k is. This is the maximum we are looking for.

Looking at the drawing again, we can convince ourselves that the level curves (the moving red line) will intersect the feasible region last in one of a few ways: either at the point P, at $(6, 0)$, at $(0, 5)$, or by meeting up with one of the two constraint lines.

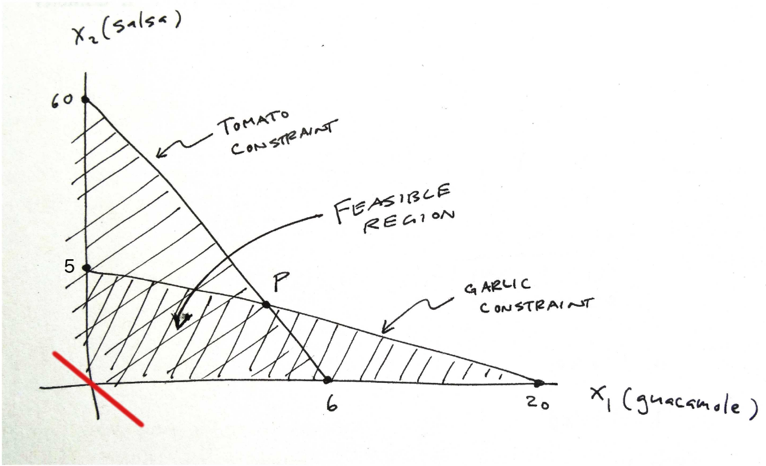

We can note that the two constraint lines have slope -10 and -0.25.

We can note that the two constraint lines have slope -10 and -0.25.

Since the level curves have slope -1, and $-10 < -1 < -0.25$, we can conclude that the level curves will last hit the feasible region at P, the intersection of the two constraint lines.

We can note that the two constraint lines have slope -10 and -0.25.

Since the level curves have slope -1, and $-10 < -1 < -0.25$, we can conclude that the level curves will last hit the feasible region at P, the intersection of the two constraint lines.

Calculating the intersection of the two constraint lines, we find $P = \left(\frac{220}{39}, \frac{140}{39}\right) \approx 5.64, 3.59$.

We can note that the two constraint lines have slope -10 and -0.25.

Since the level curves have slope -1, and $-10 < -1 < -0.25$, we can conclude that the level curves will last hit the feasible region at P, the intersection of the two constraint lines.

Calculating the intersection of the two constraint lines, we find $P = \left(\frac{220}{39}, \frac{140}{39}\right) \approx 5.64, 3.59$.

Thus, to maximize our production of salsa and guacamole, we should make 5.64 units of salsa and 3.59 units of guacamole, for a total of 9.23 units of stuff.

We see that it is the slope of the level curves of our objective function $f(x_1, x_2)$ the determine the nature of the solution.

We see that it is the slope of the level curves of our objective function $f(x_1, x_2)$ the determine the nature of the solution.

Suppose we change the situation slightly.

We see that it is the slope of the level curves of our objective function $f(x_1, x_2)$ the determine the nature of the solution.

Suppose we change the situation slightly.

Suppose now we want to **sell** our salsa and guacamole, say at \$p per unit of salsa and \$q per unit of guacamole. What should we produce to maximize the money we make from our sales (we'll assume we will sell all we make)?

We see that it is the slope of the level curves of our objective function $f(x_1, x_2)$ the determine the nature of the solution.

Suppose we change the situation slightly.

Suppose now we want to **sell** our salsa and guacamole, say at \$p per unit of salsa and \$q per unit of guacamole. What should we produce to maximize the money we make from our sales (we'll assume we will sell all we make)?

Now, the objective function is $f(x_1, x_2) = px_1 + qx_2$ and level curves have slope $-\frac{p}{q}$.

We see that it is the slope of the level curves of our objective function $f(x_1, x_2)$ the determine the nature of the solution.

Suppose we change the situation slightly.

Suppose now we want to **sell** our salsa and guacamole, say at \$p per unit of salsa and \$q per unit of guacamole. What should we produce to maximize the money we make from our sales (we'll assume we will sell all we make)?

Now, the objective function is $f(x_1, x_2) = px_1 + qx_2$ and level curves have slope $-\frac{p}{q}$.

So, depending on the relationship between the value of $-\frac{p}{q}$ and the slopes of the constraint lines, our solution might be the same as earlier, or it might be to create all salsa, or all guacamole.

We see then that it is a combination of the objective function and the constraints that determine the solution to an LP.

We see then that it is a combination of the objective function and the constraints that determine the solution to an LP.

Next time: more complex LPs that we will not try to solve by hand.