CS 170 Section 7
Linear Programming

Owen Jow | owenjow@berkeley.edu

Agenda

- Linear programming
  - Chocolate milk factory
  - Job assignment
  - Understanding convex polytopes

Linear Programming

Linear Programming

- A linear optimization problem (i.e. a **linear program**) is one where every function is affine.
- This means that the feasible set is a polytope.
- With linear programming, our goal is to assign real values to a set of variables so as to (1) maximize or minimize an objective function, and (2) meet all of the constraints.

The objective function and constraints **must** be affine.
A linear program is described by its objective function and constraints!

$$
\begin{array}{ll}
\textbf{Objective function} & \max x_1 + 6x_2 \\
\textbf{Constraints} & x_1 \leq 200 \\
& x_2 \leq 300 \\
& x_1 + x_2 \leq 400 \\
& x_1, x_2 \geq 0
\end{array}
$$

The Feasible Region

- The feasible region consists of the points which meet all constraints.
- An optimum will always be found at a vertex of the feasible region, unless perhaps the linear program is **infeasible** or the feasible region is **unbounded**.

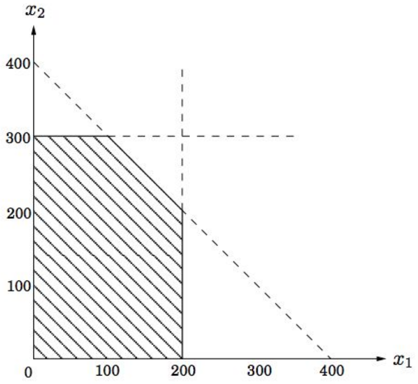

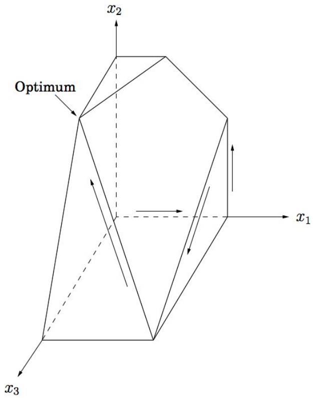

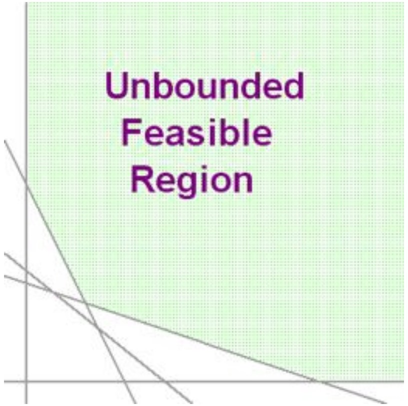

Solving a Linear Program

- One method for solving a linear program is the **simplex method**.
- This method starts at a vertex, and repeatedly moves to a better adjacent vertex until none exists.

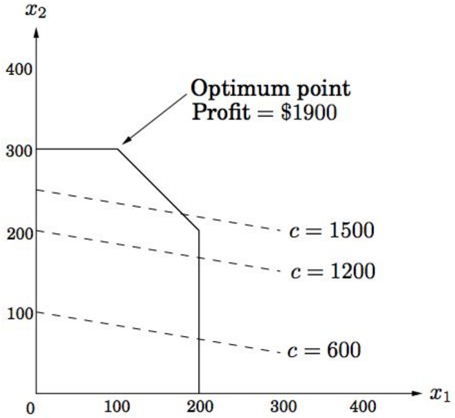

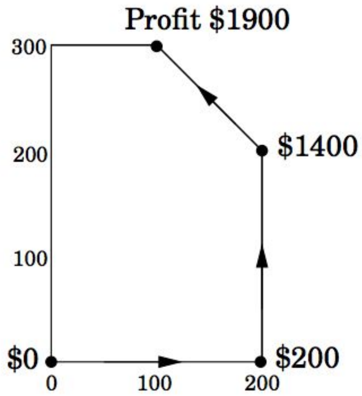

This works because our polyhedra are convex. Consider the hyperplane for the objective function that passes through the ending vertex. If all neighbors lie on one side of this hyperplane, then so must the rest of the polytope!

Rewriting Linear Programs

- We can reduce a linear program to an equivalent (but perhaps more manageable) linear program through simple transformations.

To go from maximization to minimization (or vice-versa), we multiply the objective function by -1.
To go from inequalities to equalities, we introduce a slack variable s:

$$\sum_{i=1}^{n} a_i x_i \leq b \longrightarrow \begin{aligned} \sum_{i=1}^{n} a_i x_i + s &= b \\ s &\geq 0 \end{aligned}$$

[a vector (x_1, …, x_n) satisfies the original inequality iff there is some s ≥ 0 for which it satisfies the new equality]

To go from equalities to inequalities, we rewrite ax = b as ax ≤ b and ax ≥ b.
To go from a signed x to nonnegative variables, replace x with $x^+ - x^-$.

Chocolate Milk Factory

- We make \$3 per gallon of dark chocolate.
- We make \$5 per gallon of milk chocolate.
- We cannot make negative amounts of anything.
- We can make at most 400 gallons of chocolate combined.

We want to maximize our profit. *What is the linear program for this problem?*

Chocolate Milk Factory

- We make \$3 per gallon of dark chocolate.
- We make \$5 per gallon of milk chocolate.
- We cannot make negative amounts of anything.
- We can make at most 400 gallons of chocolate combined.

We want to maximize our profit. *What is the linear program for this problem?*

Let d = # gallons of dark chocolate. Let m = # gallons of milk chocolate. Then our LP is

$$\max 3d + 5m$$
$$d, m \geq 0$$
$$d + m \leq 400$$

Chocolate Milk Factory

What does the feasible region look like? Draw the contour lines of the objective.

Chocolate Milk Factory

What does the feasible region look like? Draw the contour lines of the objective.

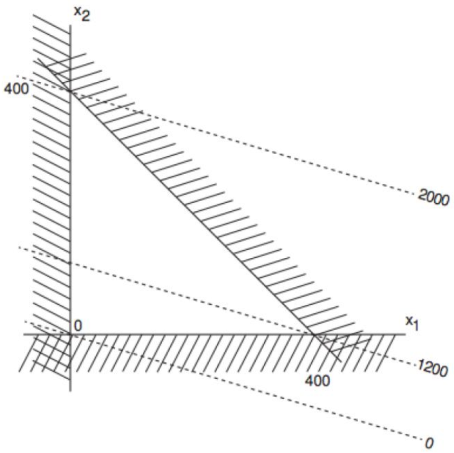

The contours are {(d, m) | 3d + 5m = c}, i.e. points with a profit of c dollars lie on the line 3d + 5m = c.

In the graph to the left, x_1 = d and x_2 = m.

Chocolate Milk Factory

Solve again with the additional constraint that you can't make more than 300 gallons of dark chocolate, and 200 gallons of milk chocolate.

Chocolate Milk Factory

Solve again with the additional constraint that you can't make more than 300 gallons of dark chocolate, and 200 gallons of milk chocolate.

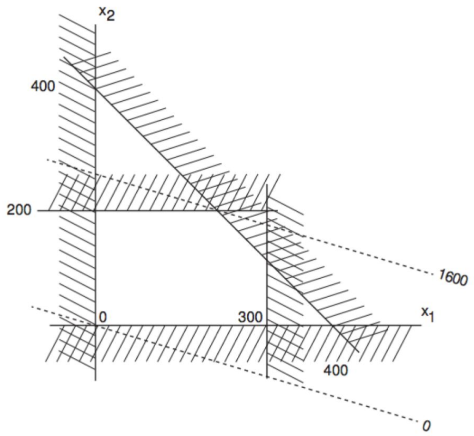

We should make
200 gallons of each.

Job Assignment

- There are I people and J jobs.
- The value of person i working 1 day at job j is a_{ij} for i = 1, …, I and j = 1, …, J.
- Each job is completed after it has been worked on for a total of 1 day (over all workers).

We want to optimally assign jobs to each person for one day. Note: a person doesn't have to spend all of his day on one job; he can spend ½ of his day on job 1, ⅓ on job 2… and so on.

In our linear program, what are the variables going to be?

Job Assignment

- There are I people and J jobs.
- The value of person i working 1 day at job j is a_{ij} for i = 1, …, I and j = 1, …, J.
- Each job is completed after it has been worked on for a total of 1 day (over all workers).

We want to optimally assign jobs to each person for one day. Note: a person doesn't have to spend all of his day on one job; he can spend ½ of his day on job 1, ⅓ on job 2… and so on.

In our linear program, what are the variables going to be?

x_{ij} for all i and j, which represents the portion of person i's day spent on job j.

Job Assignment

- There are I people and J jobs.
- The value of person i working 1 day at job j is a_{ij} for $i = 1, \ldots, I$ and $j = 1, \ldots, J$.
- Each job is completed after it has been worked on for a total of 1 day (over all workers).

We want to optimally assign jobs to each person for one day. Note: a person doesn't have to spend all of his day on one job; he can spend ½ of his day on job 1, ⅓ on job 2… and so on.

What are the constraints?

Job Assignment

- There are I people and J jobs.
- The value of person i working 1 day at job j is a_{ij} for i = 1, …, I and j = 1, …, J.
- Each job is completed after it has been worked on for a total of 1 day (over all workers).

We want to optimally assign jobs to each person for one day. Note: a person doesn't have to spend all of his day on one job; he can spend ½ of his day on job 1, ⅓ on job 2… and so on.

What are the constraints?

1. No person can work more than 1 day's worth of time.
2. No job can be worked past completion.
3. $x_{ij} \geq 0$.

Job Assignment

- There are I people and J jobs.
- The value of person i working 1 day at job j is a_{ij} for i = 1, …, I and j = 1, …, J.
- Each job is completed after it has been worked on for a total of 1 day (over all workers).

We want to optimally assign jobs to each person for one day. Note: a person doesn't have to spend all of his day on one job; he can spend ½ of his day on job 1, ⅓ on job 2… and so on.

What is the maximization function?

Job Assignment

- There are I people and J jobs.
- The value of person i working 1 day at job j is a_{ij} for i = 1, …, I and j = 1, …, J.
- Each job is completed after it has been worked on for a total of 1 day (over all workers).

We want to optimally assign jobs to each person for one day. Note: a person doesn't have to spend all of his day on one job; he can spend ½ of his day on job 1, ⅓ on job 2… and so on.

What is the maximization function?

We want to maximize the sum of $a_{ij}x_{ij}$ over all i and j. (If person i works job j for x_{ij} of a day, he/she contributes $a_{ij}x_{ij}$ value.)

Understanding Convex Polytopes

- In the standard form of a linear program, we maximize $c^T x$ such that $Ax \leq b$.
- Let's examine the properties of the feasible set $\Omega = \{x : Ax \leq b\}$.

A set X is convex if $\lambda x + (1 - \lambda)y \in X$ for any $x, y \in X$ and $\lambda \in [0, 1]$.
In other words, if we take any two points x and y in X, the entire line segment xy must also be in X.

Argue that Ω is convex.

Understanding Convex Polytopes

- In the standard form of a linear program, we maximize $c^T x$ such that $Ax \leq b$.
- Let's examine the properties of the feasible set $\Omega = \{x : Ax \leq b\}$.

A set X is convex if $\lambda x + (1 - \lambda)y \in X$ for any $x, y \in X$ and $\lambda \in [0, 1]$.
In other words, if we take any two points x and y in X, the entire line segment xy must also be in X.

Argue that Ω is convex.
Let $x, y \in \Omega$. To show convexity, we must show that the point $\lambda x + (1 - \lambda)y \in \Omega$, i.e. that $A(\lambda x + (1 - \lambda)y) \leq b$.
Since $x, y \in \Omega$, we know that $Ax \leq b$ and $Ay \leq b$. Therefore $A(\lambda x + (1 - \lambda)y) = \lambda Ax + (1 - \lambda)Ay \leq \lambda b + (1 - \lambda)b = b$.
We conclude that Ω is convex.

Understanding Convex Polytopes

- Let's show that linear maximizations over convex polytopes achieve their maximums at the vertices.
- (Again, a **polytope** is a bounded intersection of half-spaces, i.e. the generalization of a polyhedron.)
- Define a vertex as a point $v \in \Omega$ s.t. v cannot be expressed as a point on the line yz for $v \neq y$, $v \neq z$, and $y, z \in \Omega$.

Argue in favor of the following assertion:
Let Ω be a convex set and f be a linear function $f(x) = c^T x$. Show that for a line yz (with $y, z \in \Omega$), f(x) is maximized at either y or z. In other words, show that $\max f(\lambda y + (1 - \lambda)z)$ achieves its maximum at either $\lambda = 0$ or $\lambda = 1$.

Understanding Convex Polytopes

Argue in favor of the following assertion:
Let Ω be a convex set and f be a linear function $f(x) = c^T x$. Show that for a line yz (with $y, z \in \Omega$), $f(x)$ is maximized at either y or z. In other words, show that $\max f(\lambda y + (1 - \lambda)z)$ achieves its maximum at either $\lambda = 0$ or $\lambda = 1$.

Assume WLOG that $f(y) \geq f(z)$. Then $c^T y \geq c^T z$. We now aim to show that the maximum is achieved for $\lambda = 1$.

$$f(\lambda y + (1 - \lambda)z) = c^T(\lambda y + (1 - \lambda)z)$$
$$= \lambda c^T y + (1 - \lambda)c^T z$$
$$\leq \lambda c^T y + (1 - \lambda)c^T y = c^T y = f(y)$$

Understanding Convex Polytopes

Finally, argue that global maximums will be achieved at vertices.
For simplicity, assume there is a unique global maximum.

Understanding Convex Polytopes

Finally, argue that global maximums will be achieved at vertices.
For simplicity, assume there is a unique global maximum.

Proof by contradiction. Assume the maximum were achieved at a non-vertex point x. Since x is not a vertex, there must exist a line yz containing x such that x ≠ y and x ≠ z (otherwise it *would* be a vertex). However, by the previous argument, the function must achieve a maximum at either y or z!... and hence not x!

But this is a contradiction. So we conclude that the global maximum *is* achieved at a vertex.