COMP 761: Lecture 20 – Linear Programs I

David Rolnick

October 21, 2020

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$ (remember that $||v||_1 = |v_1| + |v_2| + \ldots + |v_n|$). Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

(Please don't post your ideas in the chat just yet, we'll discuss the problem soon in class.)

Course Announcements

Course Announcements

- Office hours: Vincent (10:30 am Thu), David (10:00 am Fri).

Course Announcements

- Office hours: Vincent (10:30 am Thu), David (10:00 am Fri).
- If you use theorems on the homework, you should prove them (unless they are “common sense”), if we have not stated or proven them in class.

Hyperplanes

Hyperplanes

- The equation of any line in $\mathbb{R}^2$ can be expressed as $ax + by = c$.

Hyperplanes

- The equation of any line in $\mathbb{R}^2$ can be expressed as $ax + by = c$.
- Similarly, a plane in $\mathbb{R}^3$ can be expressed as $ax + by + cz = d$.

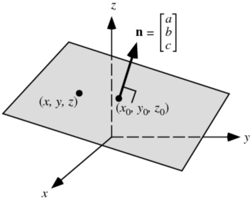

Hyperplanes

- The equation of any line in $\mathbb{R}^2$ can be expressed as $ax + by = c$.
- Similarly, a plane in $\mathbb{R}^3$ can be expressed as $ax + by + cz = d$.

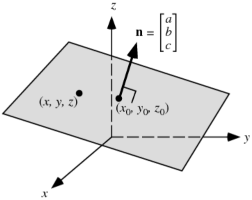

- A *hyperplane* in $\mathbb{R}^n$ is the higher-dimensional analogue:

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = b.$$

Hyperplanes

- The equation of any line in $\mathbb{R}^2$ can be expressed as $ax + by = c$.
- Similarly, a plane in $\mathbb{R}^3$ can be expressed as $ax + by + cz = d$.

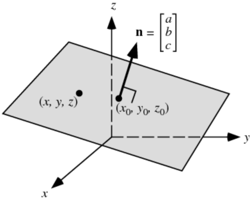

- A *hyperplane* in $\mathbb{R}^n$ is the higher-dimensional analogue:

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = b.$$

- Why is true that the vector $a = [a_1, a_2, \ldots, a_n]$ is the perpendicular vector to the hyperplane?

Hyperplanes

- The equation of any line in $\mathbb{R}^2$ can be expressed as $ax + by = c$.
- Similarly, a plane in $\mathbb{R}^3$ can be expressed as $ax + by + cz = d$.
- A *hyperplane* in $\mathbb{R}^n$ is the higher-dimensional analogue:

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = b.$$

- Why is true that the vector $a = [a_1, a_2, \ldots, a_n]$ is the perpendicular vector to the hyperplane?
- If we take any x, x' on the hyperplane, then the vector between them is $x - x'$, and we have:

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = b,$$
$$a_1 x_1' + a_2 x_2' + \cdots + a_n x_n' = b.$$

Hyperplanes

- The equation of any line in $\mathbb{R}^2$ can be expressed as $ax + by = c$.
- Similarly, a plane in $\mathbb{R}^3$ can be expressed as $ax + by + cz = d$.
- A *hyperplane* in $\mathbb{R}^n$ is the higher-dimensional analogue:

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b.$$

- Why is true that the vector $a = [a_1, a_2, \ldots, a_n]$ is the perpendicular vector to the hyperplane?
- If we take any x, x' on the hyperplane, then the vector between them is $x - x'$, and we have:

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b,$$
$$a_1x_1' + a_2x_2' + \cdots + a_nx_n' = b.$$

- Subtracting these gives:

$$0 = a_1(x_1 - x_1') + a_2(x_2 - x_2') + \cdots + a_n(x_n - x_n') = a \cdot (x - x'),$$

so a and $x - x'$ are perpendicular.

Convex sets

Convex sets

- A *convex set* S of points in $\mathbb{R}^n$ is a set such that if $x, y \in S$ then:

$$wx + (1 - w)y \in S,$$

for every constant w with $0 \leq w \leq 1$.

Convex sets

- A *convex set* S of points in $\mathbb{R}^n$ is a set such that if $x, y \in S$ then:

$$wx + (1 - w)y \in S,$$

for every constant w with $0 \le w \le 1$.

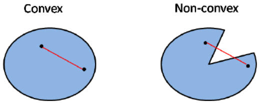

Convex sets

- A *convex set* S of points in $\mathbb{R}^n$ is a set such that if $x, y \in S$ then:

$$wx + (1 - w)y \in S,$$

for every constant w with $0 \leq w \leq 1$.

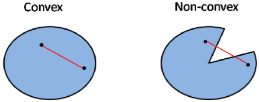

Convex sets

- A *convex set* S of points in $\mathbb{R}^n$ is a set such that if $x, y \in S$ then:

$$wx + (1 - w)y \in S,$$

for every constant w with $0 \leq w \leq 1$.

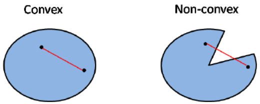

- Some more examples:

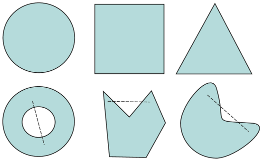

Polytopes

Polytopes

- A half-space is "one side" of a hyperplane:

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n \leq b,$$

  or with $\geq$ instead of $\leq$.

Polytopes

- A half-space is "one side" of a hyperplane:

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n \leq b,$$

  or with $\geq$ instead of $\leq$.
- A *polytope* is the intersection of a finite number of half-spaces. This gives us a connected, convex set.

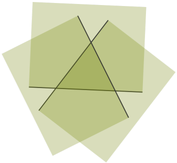

Polytopes

- A half-space is "one side" of a hyperplane:

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n \leq b,$$

  or with $\geq$ instead of $\leq$.
- A *polytope* is the intersection of a finite number of half-spaces. This gives us a connected, convex set.

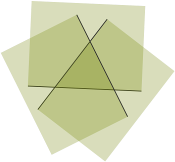

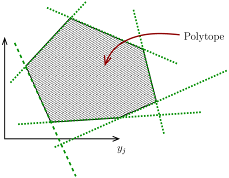

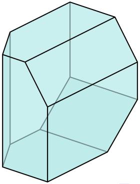

Polytopes

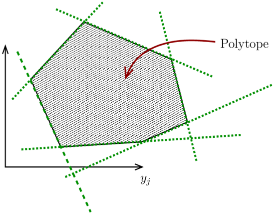

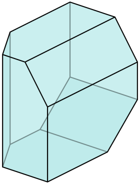

Polytopes

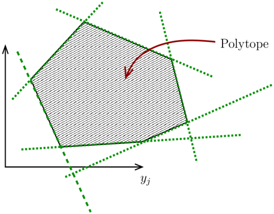

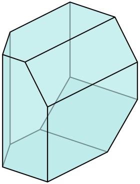

- A *vertex* of a polytope in $\mathbb{R}^n$ is a point that lies simultaneously on n of the bounding hyperplanes.

Polytopes

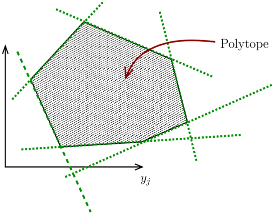

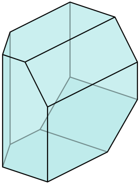

- A *vertex* of a polytope in $\mathbb{R}^n$ is a point that lies simultaneously on n of the bounding hyperplanes.
- An *edge* (1-dim face) is the set of points in the polytope that lie simultaneously on some $n - 1$ of the bounding hyperplanes.

Polytopes

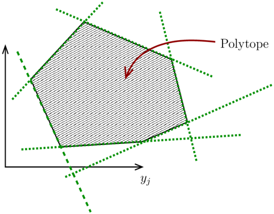

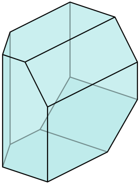

- A *vertex* of a polytope in $\mathbb{R}^n$ is a point that lies simultaneously on n of the bounding hyperplanes.
- An *edge* (1-dim face) is the set of points in the polytope that lie simultaneously on some $n - 1$ of the bounding hyperplanes.
- A *d-dimensional face* is the set of points in the polytope that lie simultaneously on some $n - d$ of the bounding hyperplanes.

Polytopes

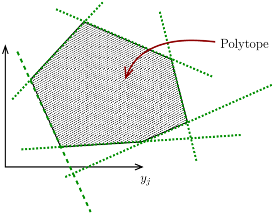

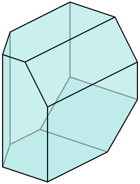

- A *vertex* of a polytope in $\mathbb{R}^n$ is a point that lies simultaneously on n of the bounding hyperplanes.
- An *edge* (1-dim face) is the set of points in the polytope that lie simultaneously on some $n - 1$ of the bounding hyperplanes.
- A *d-dimensional face* is the set of points in the polytope that lie simultaneously on some $n - d$ of the bounding hyperplanes.
- A *facet* or $(n-1)$*-dimensional face* is the set of points in the polytope that lie on one of the bounding hyperplanes.

Polytopes

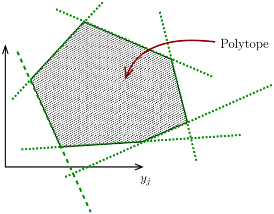

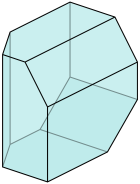

- A *vertex* of a polytope in $\mathbb{R}^n$ is a point that lies simultaneously on n of the bounding hyperplanes.
- An *edge* (1-dim face) is the set of points in the polytope that lie simultaneously on some $n-1$ of the bounding hyperplanes.
- A *d-dimensional face* is the set of points in the polytope that lie simultaneously on some $n-d$ of the bounding hyperplanes.
- A *facet* or $(n-1)$*-dimensional face* is the set of points in the polytope that lie on one of the bounding hyperplanes.
- A polytope can be thought of as the *convex hull* of its vertices – the smallest convex set containing them all.

The hypercube

The hypercube

- The polytope with vertices $[0, 0], [1, 0], [0, 1], [1, 1]$ is a square.

The hypercube

- The polytope with vertices $[0, 0], [1, 0], [0, 1], [1, 1]$ is a square.
- The polytope with vertices $[0, 0, 0]$, $[1, 0, 0]$, $[0, 1, 0]$, $[0, 0, 1]$, $[0, 1, 1]$, $[1, 0, 1]$, $[1, 1, 0]$, $[1, 1, 1]$ is a cube:

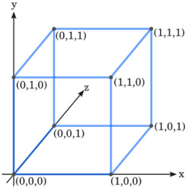

The hypercube

- The polytope with vertices $[0,0], [1,0], [0,1], [1,1]$ is a square.
- The polytope with vertices $[0,0,0]$, $[1,0,0]$, $[0,1,0]$, $[0,0,1]$, $[0,1,1]$, $[1,0,1]$, $[1,1,0]$, $[1,1,1]$ is a cube:

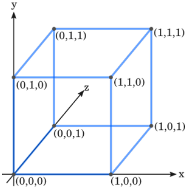

- The *hypercube* is the generalization to higher dimensions>

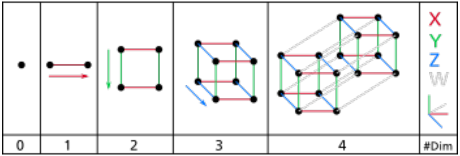

The hypercube

- The polytope with vertices $[0, 0], [1, 0], [0, 1], [1, 1]$ is a square.
- The polytope with vertices $[0, 0, 0]$, $[1, 0, 0]$, $[0, 1, 0]$, $[0, 0, 1]$, $[0, 1, 1]$, $[1, 0, 1]$, $[1, 1, 0]$, $[1, 1, 1]$ is a cube:

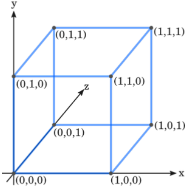

- The *hypercube* is the generalization to higher dimensions>

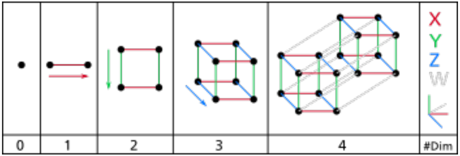

- How many vertices does the n-dimensional hypercube have?

The hypercube

- The polytope with vertices $[0,0], [1,0], [0,1], [1,1]$ is a square.
- The polytope with vertices $[0,0,0]$, $[1,0,0]$, $[0,1,0]$, $[0,0,1]$, $[0,1,1]$, $[1,0,1]$, $[1,1,0]$, $[1,1,1]$ is a cube:

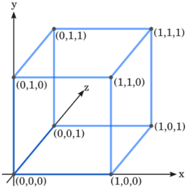

- The *hypercube* is the generalization to higher dimensions>

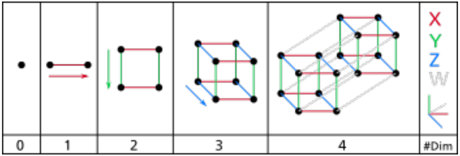

- The n-dimensional hypercube has 2^n vertices.

The hypercube

- The polytope with vertices $[0, 0], [1, 0], [0, 1], [1, 1]$ is a square.
- The polytope with vertices $[0, 0, 0]$, $[1, 0, 0]$, $[0, 1, 0]$, $[0, 0, 1]$, $[0, 1, 1]$, $[1, 0, 1]$, $[1, 1, 0]$, $[1, 1, 1]$ is a cube:

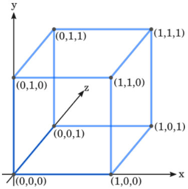

- The *hypercube* is the generalization to higher dimensions>

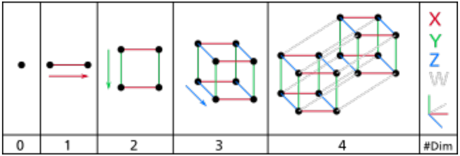

- The n-dimensional hypercube has 2^n vertices.
- How many facets does it have?

The hypercube

- The polytope with vertices $[0, 0], [1, 0], [0, 1], [1, 1]$ is a square.
- The polytope with vertices $[0, 0, 0]$, $[1, 0, 0]$, $[0, 1, 0]$, $[0, 0, 1]$, $[0, 1, 1]$, $[1, 0, 1]$, $[1, 1, 0]$, $[1, 1, 1]$ is a cube:

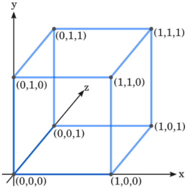

- The *hypercube* is the generalization to higher dimensions>

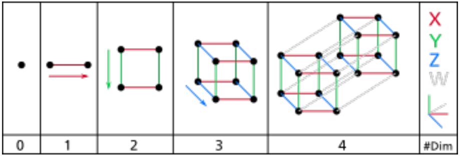

- The n-dimensional hypercube has 2^n vertices.
- It has $2n$ facets ($x_k \geq 0$ and $x_k \leq 1$).

The simplex

The simplex

- The regular *simplex* is the generalization of the equilateral triangle and regular tetrahedron into higher dimensions.

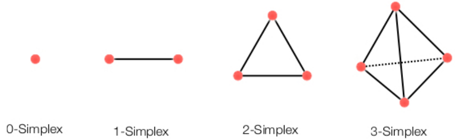

The simplex

- The regular *simplex* is the generalization of the equilateral triangle and regular tetrahedron into higher dimensions.

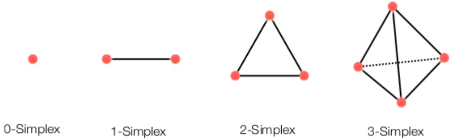

- One can think of the n-dimensional simplex within $\mathbb{R}^{n+1}$ as the convex hull of $[1, 0, \ldots, 0]$, $[0, 1, 0, \ldots, 0]$, …, $[0, 0, \ldots, 0, 1]$:

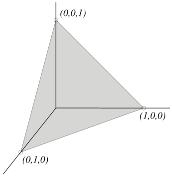

The simplex

- The regular *simplex* is the generalization of the equilateral triangle and regular tetrahedron into higher dimensions.

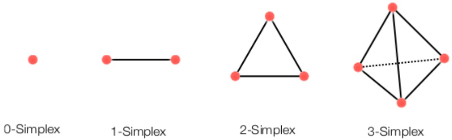

- One can think of the n-dimensional simplex within $\mathbb{R}^{n+1}$ as the convex hull of $[1, 0, \ldots, 0]$, $[0, 1, 0, \ldots, 0]$, ..., $[0, 0, \ldots, 0, 1]$:

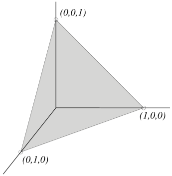

- How many vertices does the n-dimensional simplex have?

The simplex

- The regular *simplex* is the generalization of the equilateral triangle and regular tetrahedron into higher dimensions.

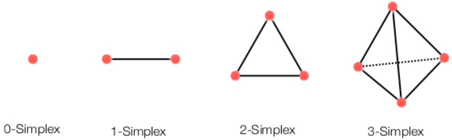

- One can think of the n-dimensional simplex within $\mathbb{R}^{n+1}$ as the convex hull of $[1, 0, \ldots, 0]$, $[0, 1, 0, \ldots, 0]$, ..., $[0, 0, \ldots, 0, 1]$:

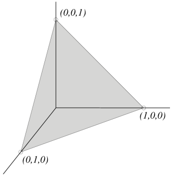

- The n-dimensional simplex has $n + 1$ vertices.

The simplex

- The regular *simplex* is the generalization of the equilateral triangle and regular tetrahedron into higher dimensions.

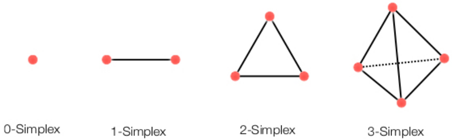

- One can think of the n-dimensional simplex within $\mathbb{R}^{n+1}$ as the convex hull of $[1, 0, \ldots, 0]$, $[0, 1, 0, \ldots, 0]$, …, $[0, 0, \ldots, 0, 1]$:

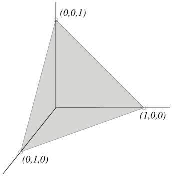

- The n-dimensional simplex has $n + 1$ vertices.
- How many facets does it have?

The simplex

- The regular *simplex* is the generalization of the equilateral triangle and regular tetrahedron into higher dimensions.

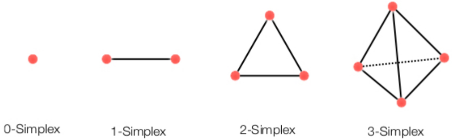

- One can think of the n-dimensional simplex within $\mathbb{R}^{n+1}$ as the convex hull of $[1, 0, \ldots, 0]$, $[0, 1, 0, \ldots, 0]$, …, $[0, 0, \ldots, 0, 1]$:

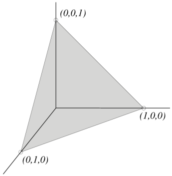

- The n-dimensional simplex has $n + 1$ vertices.
- It has $n + 1$ facets (one opposite each vertex).

Linear programs

Linear programs

- A *linear program (LP)* is an optimization problem minimizing or maximizing a linear function of variables, given linear equalities and inequalities as constraints.

Linear programs

- A *linear program (LP)* is an optimization problem minimizing or maximizing a linear function of variables, given linear equalities and inequalities as constraints.
- Examples:

$$\begin{aligned}\min_{x\in\mathbb{R}^4} \quad & x_1 + 100x_4 \\ \text{such that} \quad & 2x_1 + x_3 \geq 1 \\ & x_2 - x_4 \leq -3 \\ & x_1 + x_2 + x_3 + x_4 = 0 \\ & 5x_3 = -2\end{aligned}$$

$$\begin{aligned}\max_{x\in\mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0.\end{aligned}$$

Feasibility

Feasibility

- A *feasible point* of an LP is any $x \in \mathbb{R}^n$ that satisfies all the constraints (but might not be at the optimum).

Feasibility

- A *feasible point* of an LP is any $x \in \mathbb{R}^n$ that satisfies all the constraints (but might not be at the optimum).
- Since we have a bunch of linear inequalities and equalities, the set of feasible points is a polytope.

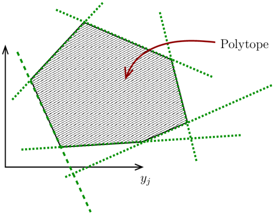

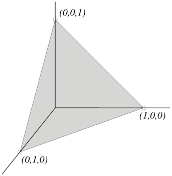

Feasibility

- A *feasible point* of an LP is any $x \in \mathbb{R}^n$ that satisfies all the constraints (but might not be at the optimum).
- Since we have a bunch of linear inequalities and equalities, the set of feasible points is a polytope.

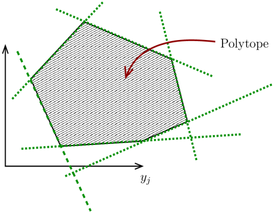

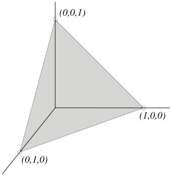

- (The dimension of the polytope might be lower than the dimension of x, e.g. if there are equality constraints.)

Feasibility

- A *feasible point* of an LP is any $x \in \mathbb{R}^n$ that satisfies all the constraints (but might not be at the optimum).
- Since we have a bunch of linear inequalities and equalities, the set of feasible points is a polytope.

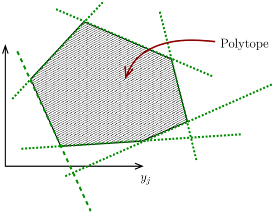

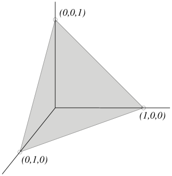

- (The dimension of the polytope might be lower than the dimension of x, e.g. if there are equality constraints.)
- If there are *no* feasible points (the *empty polytope*) then the LP is said to be *infeasible*.

Optima

If an LP has an optimum, it occurs at one of the vertices of the polytope of feasible points.

Optima

If an LP has an optimum, it occurs at one of the vertices of the polytope of feasible points.

- Why is this true?

Optima

If an LP has an optimum, it occurs at one of the vertices of the polytope of feasible points.

- Suppose we must maximize $c_1 x_1 + \ldots + c_n x_n = c \cdot x$ for $x \in \mathbb{R}^n$.

Optima

If an LP has an optimum, it occurs at one of the vertices of the polytope of feasible points.

- Suppose we must maximize $c_1 x_1 + \ldots + c_n x_n = c \cdot x$ for $x \in \mathbb{R}^n$.
- Equivalent to maximizing the projection of x in the direction of c:

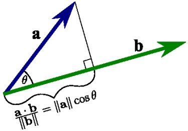

Optima

If an LP has an optimum, it occurs at one of the vertices of the polytope of feasible points.

- Suppose we must maximize $c_1 x_1 + \ldots + c_n x_n = c \cdot x$ for $x \in \mathbb{R}^n$.
- Equivalent to maximizing the projection of x in the direction of c:

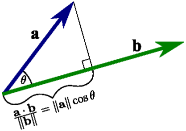

- Since x lives in a polytope of feasible solutions, we are trying to find points that are furthest in the direction of c:

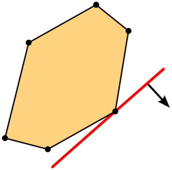

Optima

If an LP has an optimum, it occurs at one of the vertices of the polytope of feasible points.

- Suppose we must maximize $c_1x_1 + \ldots + c_nx_n = c \cdot x$ for $x \in \mathbb{R}^n$.
- Equivalent to maximizing the projection of x in the direction of c:

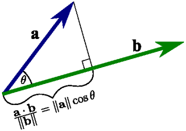

- Since x lives in a polytope of feasible solutions, we are trying to find points that are furthest in the direction of c:

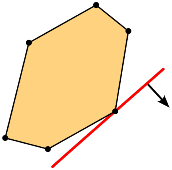

- This has to occur at a vertex (though it might not be *only* that vertex - it could be a whole edge, facet, etc.)

Unboundedness

Unboundedness

- Sometimes there isn't an optimum if the polytope is unbounded, e.g.:

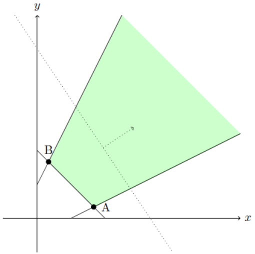

Unboundedness

- Sometimes there isn't an optimum if the polytope is unbounded, e.g.:

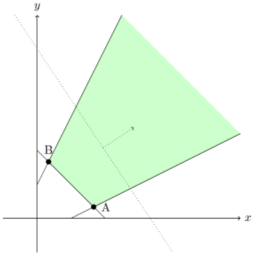

- If the polytope is unbounded in one direction, though, there can still be optima for LPs optimizing *other* directions.

Unboundedness

- Sometimes there isn't an optimum if the polytope is unbounded, e.g.:

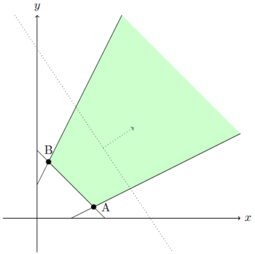

- If the polytope is unbounded in one direction, though, there can still be optima for LPs optimizing *other* directions.
- If the polytope is bounded, there exist optima in every direction.

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

- What is the set of v in $\mathbb{R}^n$ such that $||v||_1 \leq 1$?

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

- What is the set of v in $\mathbb{R}^n$ such that $||v||_1 \leq 1$?
- The constraint $||v||_1 \leq 1$ is

$$|v_1| + |v_2| + \ldots + |v_n| \leq 1.$$

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

- What is the set of v in $\mathbb{R}^n$ such that $||v||_1 \leq 1$?
- This can be rewritten as the 2^n inequalities:

$$\pm v_1 \pm v_2 \pm \ldots \pm v_n \leq 1.$$

  defining a polytope (called the *cross-polytope*).

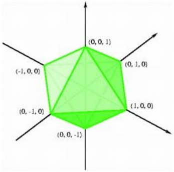

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

- What is the set of v in $\mathbb{R}^n$ such that $||v||_1 \leq 1$?
- This can be rewritten as the 2^n inequalities:

$$\pm v_1 \pm v_2 \pm \ldots \pm v_n \leq 1.$$

defining a polytope (called the *cross-polytope*).

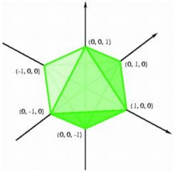

- What are the vertices of this polytope?

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

- What is the set of v in $\mathbb{R}^n$ such that $||v||_1 \leq 1$?
- This can be rewritten as the 2^n inequalities:

$$\pm v_1 \pm v_2 \pm \ldots \pm v_n \leq 1.$$

  defining a polytope (called the *cross-polytope*).

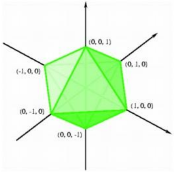

- The vertices are $[\pm 1, 0, \ldots, 0]$, $[0, \pm 1, 0, \ldots, 0]$, ..., $[0, 0, \ldots, \pm 1]$.

Problem

Suppose that c, v are vectors in $\mathbb{R}^n$, with c fixed and v variable under the constraint $||v||_1 \leq 1$. Prove that the maximum value of $c \cdot v$ is attained at some v where all entries are 0, except for one ± 1.

- What is the set of v in $\mathbb{R}^n$ such that $||v||_1 \leq 1$?
- This can be rewritten as the 2^n inequalities:

$$\pm v_1 \pm v_2 \pm \ldots \pm v_n \leq 1.$$

defining a polytope (called the *cross-polytope*).

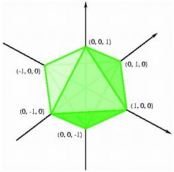

- The vertices are $[\pm 1, 0, \ldots, 0]$, $[0, \pm 1, 0, \ldots, 0]$, …, $[0, 0, \ldots, \pm 1]$.
- So the maximum of $c \cdot v = c_1 v_1 + \ldots + c_n v_n$ must be attained at one of these points.

Standard form

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- How can we deal with (1)?

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.
- How can we deal with (3)?

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.
- To satisfy (3), we can similarly flip any inequality constraints: e.g. $2x_1 + 3x_2 \geq -1$ can be $-2x_1 - 3x_2 \leq 1$.

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.
- To satisfy (3), we can similarly flip any inequality constraints: e.g. $2x_1 + 3x_2 \geq -1$ can be $-2x_1 - 3x_2 \leq 1$.
- How can we deal with (4)?

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.
- To satisfy (3), we can similarly flip any inequality constraints: e.g. $2x_1 + 3x_2 \geq -1$ can be $-2x_1 - 3x_2 \leq 1$.
- To satisfy (4), we can replace an equality with two inequalities: e.g. $2x_1 + 3x_2 = -1$ can be $2x_1 + 3x_2 \leq -1$ and $2x_1 + 3x_2 \geq -1$ (where the second one becomes $-2x_1 - 3x_2 \leq 1$).

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.
- To satisfy (3), we can similarly flip any inequality constraints: e.g. $2x_1 + 3x_2 \geq -1$ can be $-2x_1 - 3x_2 \leq 1$.
- To satisfy (4), we can replace an equality with two inequalities: e.g. $2x_1 + 3x_2 = -1$ can be $2x_1 + 3x_2 \leq -1$ and $2x_1 + 3x_2 \geq -1$ (where the second one becomes $-2x_1 - 3x_2 \leq 1$).
- How can we deal with (2)?

Standard form

- The *standard form* of a LP has:
  1. The optimization as a max not a min
  2. All variables x_k constrained by $x_k \geq 0$
  3. All other inequality constraints as $\leq$ not $\geq$
  4. No equality constraints
- We can convert any LP into an equivalent LP in standard form.
- To satisfy (1), we can flip the signs of coefficients in the objective: e.g. $\min(2x_1 + 3x_2)$ can be $\max(-2x_1 - 3x_2)$.
- To satisfy (3), we can similarly flip any inequality constraints: e.g. $2x_1 + 3x_2 \geq -1$ can be $-2x_1 - 3x_2 \leq 1$.
- To satisfy (4), we can replace an equality with two inequalities: e.g. $2x_1 + 3x_2 = -1$ can be $2x_1 + 3x_2 \leq -1$ and $2x_1 + 3x_2 \geq -1$ (where the second one becomes $-2x_1 - 3x_2 \leq 1$).
- To satisfy (2), we can e.g. replace x_1 by $(x_3 - x_4)$ with $x_3, x_4 \geq 0$, so $2x_1 + 3x_2 \leq -1$ is $2x_3 - 2x_4 + 3x_2 \leq -1$ and $x_3 \geq 0$ and $x_4 \geq 0$.

The dual

The dual

- Suppose we have an LP in standard form, e.g.

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

The dual

- Suppose we have an LP in standard form, e.g.

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- We can work out some things about the optimum.

The dual

- Suppose we have an LP in standard form, e.g.

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- We can work out some things about the optimum.
- Let's add the first inequality constraint and 2 times the second:

$$24 = 4 + 2(10) \geq (2x_1 + x_3) + 2(x_2 - x_3) = 2x_1 + 2x_2 - x_3.$$

The dual

- Suppose we have an LP in standard form, e.g.

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- We can work out some things about the optimum.
- Let's add the first inequality constraint and 2 times the second:

$$24 = 4 + 2(10) \geq (2x_1 + x_3) + 2(x_2 - x_3) = 2x_1 + 2x_2 - x_3.$$

- Why does this help?

The dual

- Suppose we have an LP in standard form, e.g.

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- We can work out some things about the optimum.
- Let's add the first inequality constraint and 2 times the second:

$$24 = 4 + 2(10) \geq (2x_1 + x_3) + 2(x_2 - x_3) = 2x_1 + 2x_2 - x_3.$$

- Why does this help?
- We know $x_2 \geq 0$, so $2x_1 + 2x_2 - x_3 \geq 2x_1 + x_2 - x_3$.

The dual

- Suppose we have an LP in standard form, e.g.

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- We can work out some things about the optimum.
- Let's add the first inequality constraint and 2 times the second:

$$24 = 4 + 2(10) \geq (2x_1 + x_3) + 2(x_2 - x_3) = 2x_1 + 2x_2 - x_3.$$

- Why does this help?
- We know $x_2 \geq 0$, so $2x_1 + 2x_2 - x_3 \geq 2x_1 + x_2 - x_3$.
- So the optimum is at most 24.

The dual

$$
\begin{aligned}
\max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\
\text{such that} \quad & 2x_1 + x_3 \leq 4 \\
& x_2 - x_3 \leq 10 \\
& x_1, x_2, x_3 \geq 0.
\end{aligned}
$$

The dual

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- This inspires us to take a linear combination of the constraints in the LP to see how tight a bound we can get.

The dual

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- This inspires us to take a linear combination of the constraints in the LP to see how tight a bound we can get.
- Let's multiply the first constraint by y_1, the second by y_2, etc.

The dual

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- This inspires us to take a linear combination of the constraints in the LP to see how tight a bound we can get.
- Let's multiply the first constraint by y_1, the second by y_2, etc.
- If the coefficient of x_i in constraint j is a_{ij}, and the coefficient of x_i in the objective is c_i, then we need

$$\sum_j a_{ij} y_j \geq c_i.$$

The dual

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- This inspires us to take a linear combination of the constraints in the LP to see how tight a bound we can get.
- Let's multiply the first constraint by y_1, the second by y_2, etc.
- If the coefficient of x_i in constraint j is a_{ij}, and the coefficient of x_i in the objective is c_i, then we need

$$\sum_j a_{ij} y_j \geq c_i.$$

- We also want all $y_j \geq 0$.

The dual

$$\begin{aligned} \max_{x \in \mathbb{R}^3} \quad & 2x_1 + x_2 - x_3 \\ \text{such that} \quad & 2x_1 + x_3 \leq 4 \\ & x_2 - x_3 \leq 10 \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

- This inspires us to take a linear combination of the constraints in the LP to see how tight a bound we can get.
- Let's multiply the first constraint by y_1, the second by y_2, etc.
- If the coefficient of x_i in constraint j is a_{ij}, and the coefficient of x_i in the objective is c_i, then we need

$$\sum_j a_{ij} y_j \geq c_i.$$

- We also want all $y_j \geq 0$.
- Then, our bound is $\sum_j b_j y_j$, where b_j are upper bounds in the inequalities on x (4 and 10 in this case).

The dual

$$\max_{x\in\mathbb{R}^3} \quad 2x_1 + x_2 - x_3$$
$$\text{such that} \quad 2x_1 + x_3 \le 4$$
$$x_2 - x_3 \le 10$$
$$x_1, x_2, x_3 \ge 0.$$

- This inspires us to take a linear combination of the constraints in the LP to see how tight a bound we can get.
- Let's multiply the first constraint by y_1, the second by y_2, etc.
- If the coefficient of x_i in constraint j is a_{ij}, and the coefficient of x_i in the objective is c_i, then we need

$$\sum_j a_{ij} y_j \ge c_i.$$

- We also want all $y_j \ge 0$.
- Then, our bound is $\sum_j b_j y_j$, where b_j are upper bounds in the inequalities on x (4 and 10 in this case).
- We want this as *small* as possible so it gives a good upper bound.

The dual

The dual

- A general LP in standard form can be written as:

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & c \cdot x \\ \text{such that} \quad & Ax \leq b \\ & x \geq 0. \end{aligned}$$

where $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$.

The dual

- A general LP in standard form can be written as:

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & c \cdot x \\ \text{such that} \quad & Ax \leq b \\ & x \geq 0. \end{aligned}$$

  where $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$.
- Then, the dual can be written as:

$$\begin{aligned} \min_{y \in \mathbb{R}^m} \quad & b \cdot y \\ \text{such that} \quad & A^T y \geq c \\ & y \geq 0. \end{aligned}$$

The dual

- A general LP in standard form can be written as:

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & c \cdot x \\ \text{such that} \quad & Ax \leq b \\ & x \geq 0. \end{aligned}$$

  where $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$.
- Then, the dual can be written as:

$$\begin{aligned} \min_{y \in \mathbb{R}^m} \quad & b \cdot y \\ \text{such that} \quad & A^T y \geq c \\ & y \geq 0. \end{aligned}$$

- (You can look up how to get the dual if the LP *isn't* in standard form – it's also pretty easy to work out yourself.)

Next time!

Linear Programs II