• The vectors $V = \{v_1, v_2, \dots, v_n\}$ are *linearly independent* if $\sum_i \alpha_i v_i = 0$ and $\forall_i \alpha_i \in \mathbb{R}$ imply that $\forall_i \alpha_i = 0$.

A vector space or linear space \mathcal{V} in \mathbb{R}^n is a nonempty subset of \mathbb{R}^n closed under vector addition and scalar multiplication.

A set of vectors in V are said to span V whenever every vector in V can be expressed as a linear combination of vectors in V and any vector not in V cannot be expressed as a linear combination of vectors in V.

The vectors in V form a *basis* of V iff vectors in V are linearly independent and span(V) = V iff vectors in V are linearly independent and including any vector not in V into V makes vectors in V linearly dependent. Though it is not necessary for a vector space to have a unique basis, the number of vectors in any basis of V is the same. The *dimension* of V is the number of vectors in any basis of V.

For a matrix A, the vector space spanned by the columns of A is called the *column space*, a.k.a., range space, C(A) of A. The rank of A is the dimension of C(A).

For a matrix $A_{m \times n}$ with $n \le m$, the *nullspace a.k.a.*, *kernel* of A is the set comprising vectors in $\{x : Ax = 0\}$. And, the rank of nullspace of A is equal to n - rank(A).

• Given vectors $x_1, \ldots, x_k \in \mathbb{R}^n$ and $\lambda_1, \ldots, \lambda_k \ge 0$ with $\sum_{i=1}^k \lambda_i = 1$, we call $x = \sum_{i=1}^k \lambda_i x_i$ a convex combination of x_1, \ldots, x_k .

For a point p belonging to line segment p'p'', there exists a $\alpha \in [0, 1]$ such that $p = \alpha p' + (1 - \alpha)p''$. That is, p is a convex combination of p' and p''.

A set S is called a *convex set* if for every two points x, y ∈ S, every point on the line segment joining x and y is contained in S.



Equivalently, a set S is convex if for every set S' of points located in S, the convex combination of points in S' also belongs to S.

The intersection of any number of convex sets in \mathbb{R}^n is convex.

- Let $S \subseteq \mathbb{R}^n$ be a convex set. The function $f : S \to \mathbb{R}$ is a *convex function* in S if for any two points $x, y \in S$, $f(\lambda x + (1 \lambda)y) \le \lambda f(x) + (1 \lambda)f(y)$ for $\lambda \in \mathbb{R}$ and $0 \le \lambda \le 1$. Ex., $x^2, e^x, |x|$.
- The *convex hull* of a set S, denoted by CH(S), is the smallest convex set that contains S. That is, the convex hull of S is the intersection of all convex sets that contain S.

Let CC(S) be the set that precisely comprises of all the connvex combinations of points in S. Then, CC(S) = CH(S).

For any set S and $S' \subseteq S$, $CH(S') \subseteq CH(S)$.

- *Caratheodory's theorem*: For a point $x \in \mathbb{R}^d$, suppose x is expressed as a convex combination of k > d+1 number of points in \mathbb{R}^d . Then, x can be expressed as a convex combination of at most d+1 points.
- An *extreme point* of a convex set S is a point p in S which cannot be expressed as a convex combination of points in S − {p}. In other words, a point p ∈ S is an extreme point of S if and only if p is a vertex of CH(S).
- The vectors v_1, \ldots, v_r are affinely independent if $\sum_i \alpha_i v_i = 0$, $\forall_i \alpha_i \in \mathbb{R}$, and $\sum_i \alpha_i = 0$ together imply that $\forall_i \alpha_i = 0$.

An *affine space* A in \mathbb{R}^n is the space resulting from adding a fixed vector t to all the elements of a linear space S, i.e., $A = \{t + y | y \in S\}$. (For example, a plane in \mathbb{R}^2 not necessarily passing through the origin is an affine space.) The dimension of A is the dimension of S. If B is an arbitrary subset of \mathbb{R}^n , then the dimension of B is the smallest dimension of any affine space containing B.

A hyperplane in \mathbb{R}^n is an affine subspace of dimension n-1; in other words, it is the set of all solutions of a linear equation of the form $a_1x_1 + \ldots + a_nx_n = b$, where a_1, \ldots, a_n are not all 0.

The hyperplane with equation $a_1x_1 + \ldots + a_nx_n = b$ induces two *closed half-spaces*: $\{x \in \mathbb{R}^n : x_1 + \ldots + a_nx_n \leq b\}$ and $\{x \in \mathbb{R}^n : x_1 + \ldots + a_nx_n \geq b\}$.

• For some matrix $A \in \mathbb{R}^{m \times n}$ and some vector $b \in \mathbb{R}^m$, the set *P* comprising points in $\{x \in \mathbb{R}^n : Ax \le b\}$ is a *polyhedron* (a.k.a., *H-polyhedron* or *polyhedral set*). That is, the convex set obtained by the intersection of a finite number of affine halfspaces is a polyhedron. If *A* and *b* are rational, then *P* is a *rational polyhedron*. A bounded H-polyhedron is called a *H-polytope*.

A subset P of \mathbb{R}^n is called a *polytope* (a.k.a., V-*polytope*) if P is the convex hull of a finite number of vectors.

Minkowski-Weyl's theorem: A subset P of \mathbb{R}^n is a polytope if and only if it is a bounded polyhedron.

A polyhedron P is *pointed* if it has at least one vertex.

- The dimension of a convex polyhedron P ⊆ ℝⁿ is the smallest dimension of an affine subspace containing P. Equivalently, it is the largest d ≤ n for which P contains points x₀, x₁,..., x_d such that the d-tuple of vectors (x₁ x₀,..., x_d x₀) is linearly independent. That is, the dimension is the maximum number of affinely independent points in P minus 1. If the dimension of P is n then P is full-dimensional.
- For a hyperplane H, if H ∩ P ≠ φ and P is entirely contained in one of the closed half-spaces defined by H then H is called a *supporting hyperplane of* P. Formally, let P = {x : Ax ≤ b} be a nonempty polyhedron. If c is a nonzero vector for which δ = max{cx : x ∈ P} is finite, then {x : cx = δ} is called a supporting hyperplane of P.
- A *face* of *P* is *P* itself or the intersection of *P* with a supporting hyperplane of *P*.

F is a face of P if and only if $F = \{x \in P : A'x = b'\} \neq \phi$ for some subsystem $A'x \leq b'$ of $Ax \leq b$.

A point x for which $\{x\}$ is a face is called a *vertex* of P. That is, a vertex is a 0-dimensional face. An edge is a 1-dimensional face.

A *facet* of P is any maximal face distinct from P. An inequality $cx \le \delta$ is *facet-defining* for P if $cx \le \delta$ for all $x \in P$ and $\{x \in P : cx = \delta\}$ is a facet of P.

• Asymptotic upper bound theorem: A d-dimensional polytope with n vertices has $O(n^{\lfloor d/2 \rfloor})$ facets and $O(n^{\lfloor d/2 \rfloor})$ faces.

Since any polytope formed with f half-spaces can have at most f facets, in d-dimensions, the maximum number of vertices a polytope has is $\Theta(f^{\lfloor d/2 \rfloor})$.

- Separation theorem: Let $S \subseteq \mathbb{R}^n$ be a finite set and let $v \in \mathbb{R}^n \setminus CH(S)$. Then there exists an inequality $w^T x \leq t$ that separates v from CH(S), that is, $w^T s \leq t$ for all $s \in CH(S)$ but $w^T v > t$.
- Theorem of the alternatives: The system $Ax \leq b$ has no solution $x \in \mathbb{R}^n$ if and only if there exists $y \in \mathbb{R}^m$ such that $y \geq 0$, $A^T y = 0$ and $b^T y < 0$.

Farkas' lemma: The system Ax = b has a nonnegative solution if and only if there is no vector y satisfying $y^T A \ge 0$ and $y^T b < 0$.

• The linear constraints together form a convex feasible region, which is a convex polyhedra.



• The LP $\min\{c^T x : Ax \ge b, x \ge 0\}$ is *unbounded* from below if there exists a feasible x such that $c^T x < \lambda$ for every real number λ . (Analogous definition applies to maximization LP as well.) When the feasible region is empty, then the LP is said to be *infeasible*. Any linear program either (i) has an optimal solution with a finite objective value, (ii) is infeasible, or (iii) is unbounded.



Consider an LP max{c^Tx : Ax = b, x ≥ 0}. Let B be a basis of A. The basic solution corresponding to basis B is a vector x ∈ ℝⁿ with (i) for every j ∉ B, x_j = 0, and (ii) for every j ∈ B, the kth element of column vector x is equal to the kth element of B⁻¹b. If x is both basic and feasible, then x is a basic feasible solution. Every basic feasible solution is uniquely determined by the set B. The variables x_j with j ∈ B the basic variables and the remaining variables are the nonbasic variables.

• Let A be a $m \times n$ matrix with full row rank. Then, every feasible x to $P = \{x : Ax = b, x \ge 0\}$ is a basic feasible solution if and only if x is an extreme point solution. (Note that after introducing slack variables for each constraint, any LP can be transformed to a linear program in *standard form*: $\{x : Ax = b, x \ge 0\}$.)

The rank lemma: Let $P = \{x \mid Ax \leq b\}$ be a polyhedron in \mathbb{R}^n and let $z \in P$. Also, let A_z be the submatrix of A consisting of those rows a_i of A for which $a_i z = b_i$. Then z is an extreme point of P if and only if $rank(A_z) = n$.

Every basic feasible solution corresponds to one and only one vertex. And, there exists a basic feasible solution corresponding to every vertex of F. However, the correspondance between bases and basic feasible solutions is not one to one: there can be many bases, each of which corresponds to the same basic feasible solution. When more than two different bases correspond to the same basic feasible solution x, then x contains more than n - m zeros, and such an x is said to be a *degenerate* solution.

The objective function is a family of parallel hyperplanes; the objective function value is fixed along any one such hyperplane. Significantly, for the set S of points satisfying the given constraints, max{w^Tx : x ∈ S} = max{w^Tx : x ∈ bd(CH(S))}.



An optimal solution corresponding to a basic feasible solution is called a *basic optimal solution*. As every supporting hyperplane contains at least one vertex, there exists a vertex at which optimum is guaranteed to occur.



Let $P = \{x : Ax \ge b, x \ge 0\}$ and assume that $\min\{c^T x : x \in P\}$ is finite. Then, for every $x \in P$, there is an extreme point solution $x' \in P$ such that $c^T x' \le c^T x$, i.e., there is always an extreme point optimal solution.

Hence, LP is a combinatorial optimization problem even though it seem to be a continuous optimization problem! (Some say it lies on the boundary of continous and discrete optimization problems.)

[•] Local optima is the global optima for any LP: Consider an optimization problem with feasibility constraints F and objective function f, where $F \subseteq R^n$ is a convex set and f is a convex function defined over

the domain F. Then, the global optimality equals the local optimality whenever the locality is defined with respect to an ϵ -radius Euclidean closed ball B, that is, if f(x) is locally optimum (say, minimum), then $\forall_{y \in B} f(x) \leq f(y)$.

• Dantzig's *simplex algorithm*:

A variable whose value is currently set to 0 is called a nonbasic variable; and, a variable whose value is currently not equal to 0 is called a basic variable. At any time of execution of the simplex algorithm, every variable is either a basic or a nonbasic variable.

In every iteration of simplex algorithm, one nonbasic variable becomes a basic variable and one basic variable becomes a nonbasic variable; for two adjacent vertices of a LP polytope, the basic and nonbasic sets are identical except for these. Hence, in every iteration, choose the nonbasic variable that gives the fastest rate of increase in the objective function value as the entering basic variable, and choose the basic variable that most limits the increase in the value of the entering basic variable as the leaving basic variable. Significantly, the updated solution vector in any iteration is guaranteed to improve the objective function value; besides, the new solution vector corresponds to a vertex of the LP polytope.

If the objective function value cannot be increased by increasing the values of any of the variables on which it depends, then the simplex must have arrived to optimality.

- (1) corresponding to a vertex of the LP polytope, initialize the basic variables.
- (2) while (the current solution is not optimal)
 - (i) choose entering and leaving basic variables
 - (ii) update the current solution vector

Using the asymptotic upper bound theorem, the worst-case time was proven to be exponential in the number of variables. Klee & Minty provided instances (known as Minty's cube) on which the algorithm indeed visits all the vertices of the polytope. However, it does well in practice. Specifically, under various probability distributions over the input, the simplex algorithm is proven to be taking polynomial time in the average-case.

• Kachiyan's Ellipsoid algorithm:



- (1) start with an ellpsoid E that contains all the basic feasible solutions of the given LP
- (2) invoke a *separation oracle O*, which is given as input, with center c of the ellipsoid E:
 if c is feasible, O returns a constraint C: ∑_j d_jx_j ≤ ∑_j d_jc_j for the objective function ∑_j d_jx_j
 else O returns a *violated constraint C*: ∑_j a_jx_j ≤ b for the center c, i.e., the current set of basic feasible solutions and c are on different sides of C

- in both the cases, one of the hyperspaces H_C defined by C contains a basic optimal solution (if one exists for the given LP)
- (3) compute a minimum volume elliposid E' that contains the region common to both E and H_C
- (4) set E := E'; and go to step (2) unless the ellipsoid is sufficiently small so that it can contain at most one basic feasible solution, if any (if one exists, then this must be a basic optimal solution)

This algorithm takes $O(n^6Lt)$ time, where n is the number of variables, L is the number of bits in encoding the input instance, and t is the time complexity of the separation oracle; the time does not depend on the number of constraints! In practice, this algorithm does not perform as good as the simplex algorithm. And, due to square-root operations involved, rounding errors need to be handled carefully.

However, the user of this algorithm requires to provide a separation oracle: given a point p either determine that p is in the polytope P corresponding to the given LP or find a constraint $a^T x \leq b$ that contains P but $a^T p > b$. For some LPs, oracles that do better than explicitly checking p versus each constraint are possible, i.e., a polynomial-time oracle exists even when the number of constraints is exponential! Hence, subjected to the efficiency of separation oracle provided, the Ellipsoid method take only polynomial time even when the number of constraints is exponential.

Ex. relaxed LP for arborescence rooted at r:

 $\begin{array}{l} \min \sum_{(i,j)\in E} c_{ij} x_{ij} \\ \sum_{(i,j)\in E, i\in S, j\in V-S} x_{ij} \geq 1 \\ 0 \leq x_{ij} \leq 1 \end{array} \begin{array}{l} \text{for every } S \subseteq V, r \in S \end{array}$ (1)

separation oracle: detecting a violated constraint among inequalities in (1) is equivalent to finding a r-t min-cut for some $t \in V - \{r\}$ such that the cut capacity is less than one when the weight of arc (i, j) is set to x_{ij} for every arc (i, j); other constraints can be checked explicitly

• Karmakar's *interior point algorithm*: takes $O(n^{3.5}L^2(\lg L)(\lg \lg L))$ time; in practice, competes with (or, does better than) the simplex algorithm. Here, n is the number of variables, and L is the number of bits in encoding the input instance. In practice, this algorithm does not perform as good as the simplex algorithm. And, like in Ellipsoid algorithm, there are rounding errors due to square-root operations.

The hunt is on for a strongly polynomial time algorithm.

• Primal-dual LP pairs -

primal LP:		dual LP:
$\min \sum_{i=1}^{n} c_i x_i$		$\max \sum_{j=1}^{m} b_j y_j$
$a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n = b_i$	$i \in M$	$y_i \ll 0$ $i \in M$
$a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n \ge b_i$	$i\in\overline{M}$	$y_i \ge 0$ $i \in \overline{M}$
$x_j \ge 0 \qquad j \in N$		$a_{1j}y_1 + a_{2j}y_2 + \ldots + a_{mj}y_m \le c_j j \in N$
$x_j <> 0 \qquad j \in \overline{N}$		$a_{1j}y_1 + a_{2j}y_2 + \ldots + a_{mj}y_m = c_j j \in \overline{N}$

• Weak duality theorem: Let x^* , y^* be feasible solutions of primal and dual LPs respectively. Assuming the primal LP is a minimization LP,

 $\sum_{i=1}^{m} b_i y_i^* \le \sum_{j=1}^{n} c_j x_j^*$

A consequence of weak duality theorem: If $\sum_{i=1}^{m} b_i y_i^* = \sum_{j=1}^{n} c_j x_j^*$ then x^* and y^* are optimal solutions to the primal and dual LPs respectively.

- *Strong duality theorem*: The primal LP has an optimal solution if and only if the dual LP has an optimal solution. At optimality, cost of primal LP equals to the cost of dual LP.
- Complementary slackness conditions: A pair x^*, y^* respectively feasible in a primal-dual LP pair is optimal if and only if

$$\forall_i y_i((a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n) - b_i) = 0, \text{ and}$$

 $\forall_j (c_j - (a_{1j}y_1 + a_{2j}y_2 + \ldots + a_{mj}y_m))x_j = 0.$

The former conditions are known as the pimal complementary slackness conditions whereas the latter are the dual complementary slackness conditions.

- For primal and dual LPs P and D, one of the following four statements must hold: (i) both P and D are feasible, (ii) P is infeasible and D is unbounded, (iii) P is unbounded and D is infeasible, and (iv) both P and D are infeasible.
- If the coefficients in the constraint matrix, objective function, and the RHS vector are all non-negative in the LP, then that LP is called a *covering LP*. The dual of a covering LP is called a *packing LP*.

- The integer linear programming (ILP) is a linear program in which the variables are restricted to be integers. The decision version of the ILP problem is NP-complete.
- A square, integer matrix M is called unimodular (UM) if det(M) = ±1. An integer matrix M is called totally unimodular (TUM) if every square, nonsingular submatrix of M is unimodular.
 Examples for TUM matrices include node-arc incidence matrix of a directed graph (occurs in ILPs corresponding to shortest path, maximum flow, Hitchcock problems) and the node-edge incidence matrix of an undirected bipartite graph (occurs in ILPs of maximum weighted bipartite perfect matching).

If A is TUM and b is integral, then the basic solutions corresponding to both LPs $\{x : Ax = b, x \ge 0, x \text{ integer}\}$ and $\{x : Ax = b, x \ge 0\}$ are same.

Analogously, if A is TUM and b is integral, then the basic solutions corresponding to both LPs $\{x : Ax \le b, x \ge 0, x \text{ integer}\}$ and $\{x : Ax \le b, x \ge 0\}$ are same.

• For a minimization (resp. maximization) ILP, the *integrality gap* of an ILP is defined as the supremum of $\frac{OPT(ILP)}{OPT(LP_{relaxed})}$ (resp. $\frac{OPT(ILP)}{OPT(LP_{relaxed})}$) over all problem instances.

[•] Let P be a polytope and let x be an extreme point solution of P, then x is *integral* if each coordinate of x is an integer. The polytope P is called *integral* if every extreme point of P is integral.

The integrality gap enforces a limit on the approximation power of our relaxation: It is not possible, at least for any algorithm that derives its performance guarantee by comapring its value to that of the LP relaxation, to get an approximation factor better than the integrality gap of that ILP.

Ex: Consider the following LP relaxation for the vertex cover ILP:

$$\min \sum_{v \in V} x_v \quad \text{s.t.}$$
$$x_v + x_w \ge 1 \quad \forall (v, w) \in E$$
$$x_v \ge 0 \quad \forall v \in V$$

For K_n , an integral solution need to choose n-1 vertices. Again, a feasible fractional solution assigns $x_v = \frac{1}{2}$ to all vertices, with the objective value $\frac{n}{2}$. Therefore, the integrality gap for K_n is at least $\frac{n-1}{n/2}$, which is equal to $2(1-\frac{1}{n})$, i.e., the integrality gap approaches 2 as n tends to infinity.

References:

- * Combinatorial Optimization by C. H. Papadimitriou and K. Steiglitz.
- * Lectures on Discrete Geometry by J. Matousek.