

Figure 1: Uniqueness of minima in the convex and non-convex case

2 Simplest problems and methods

The difference between Min-Cut and Max-Cut comes from the fact that the former can be reduced to a convex problem with polynomial size in the data, while the latter cannot.

Definition 2.1. A set $X \subset \mathbb{R}^n$ is called *convex* if for every $x, y \in X$ and $\lambda \in (0, 1)$ we have $\lambda x + (1 - \lambda)y \in X$.

A function $f : X \rightarrow \mathbb{R}$ on a convex set X is called *convex* if for every $x, y \in X$ and $\lambda \in (0, 1)$ we have $f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$, and *concave* if the opposite inequality holds.

If the feasible set X and the cost function $f(x)$ in the minimization problem $\min_{x \in X} f(x)$ are convex, then the problem is called convex. Convexity of the problem guarantees that every local minimum is also a global minimum.

2.1 Examples of simple convex problems

Convex quadratic function: We wish to minimize the function $f(x) = \frac{1}{2}x^T A x + b^T x$ with A real symmetric and with positive eigenvalues. The minimizer can be found analytically via the gradient condition $f'(x) = Ax + b = 0$. This problem hence reduces to solving a linear system of equations $Ax = -b$.

Remark 2.2. It is less costly to solve a linear system directly rather than to compute the inverse A^{-1} and to proceed by the analytic formula $x = -A^{-1}b$.

Linear function over ellipsoid: We wish to solve the problem

$$\min_x c^T x : \quad (x - x_0)^T A^{-1}(x - x_0) \leq 1,$$

with A real symmetric and with positive eigenvalues, and with $c \neq 0$. The feasible set of the problem is an ellipsoid centered on x_0 .

The minimizer x^* has to lie on the boundary of the ellipsoid, with the level surface through the minimizer being tangent to the ellipsoid and the gradient pointing inside the ellipsoid. Therefore there exists $\lambda < 0$ such that $\lambda c = 2A^{-1}(x^* - x_0)$, yielding $x^* = x_0 + \frac{\lambda}{2}Ac$. We also get

$$1 = (x^* - x_0)^T A^{-1}(x^* - x_0) = \frac{\lambda^2}{4} c^T A A^{-1} A c = \frac{\lambda^2}{4} c^T A c,$$

yielding $\lambda = -\frac{2}{\sqrt{c^T A c}}$ and finally

$$x^* = x_0 - \frac{Ac}{\sqrt{c^T A c}}.$$

Quadratic function over Euclidean ball: Consider the problem

$$\min_{x \in \mathbb{R}^n} x^T A x : \quad \|x\|_2 \leq 1.$$

Here A is real symmetric, but not necessarily positive semi-definite. If $A \succeq 0$, then the objective function is convex, and $x = 0$ is a local and hence global minimum. If $\lambda_{\min}(A) < 0$, then by homogeneity the minimum is achieved on the boundary of the ball. The gradient of the objective must then be proportional to the gradient $2x$ of the constraint, and there exists $\lambda \in \mathbb{R}$ such that $2Ax = 2\lambda x$. It follows that λ is an eigenvalue of A and x is the corresponding unit norm eigenvector. We then get for the value of the objective that $x^T A x = \lambda x^T x = \lambda$. Hence the minimum of the objective is achieved at the eigenvector corresponding to the minimum eigenvalue $\lambda_{\min}(A)$.

The problem is formally non-convex, but can be rewritten equivalently in convex form. With $X = xx^T$ we have

$$x^T A x = \langle A, X \rangle, \quad \|x\|_2^2 = \text{tr } X,$$

and hence the problem is equivalent to

$$\min_{X \in \mathcal{S}^n} \langle A, X \rangle : \quad \text{tr } X \leq 1, \quad X \succeq 0, \quad \text{rk } X \leq 1.$$

Here $\langle A, B \rangle = \sum_{i,j=1}^n A_{ij} B_{ij}$ is the Frobenius scalar product between symmetric matrices. Removing the rank constraint, we obtain a semi-definite program (SDP). However, this SDP has always a solution of rank not exceeding 1, and in generic position this solution is unique.

Quadratic function over sphere: Here we replace the inequality constraint by an equality constraint,

$$\min_{x \in \mathbb{R}^n} x^T A x : \quad \|x\|_2 = 1.$$

As in the previous problem the solution is the eigenvector to the minimal eigenvalue of A .

An equivalent convex problem is the SDP

$$\min_{X \in \mathcal{S}^n} \langle A, X \rangle : \quad \text{tr } X = 1, \quad X \succeq 0.$$

The Euclidean ball and the sphere in the previous examples can be replaced by an ellipsoid or the boundary of an ellipsoid, because these can be reduced to a ball or a sphere by a coordinate transformation.

Line search with bisection: We wish to minimize a convex function $f : \mathbb{R} \supset I \rightarrow \mathbb{R}$, or more generally a function which is strictly monotonically decreasing for $x \leq x^*$ and strictly monotonically increasing for $x \geq x^*$, where x^* is the minimizer of the function, and I is a (not necessarily finite) interval. We assume that a minimizer x^* exists and that given a point x , there is a means to determine whether $x \leq x^*$ or $x \geq x^*$, e.g., by computing the derivative $f'(x)$.

We commence with an initial guess x_0 and check whether x_0 is smaller or larger than x^* . Assume without loss of generality that $x_0 < x^*$. Then we choose $d > 0$ and $\lambda > 1$ and compute recursively a sequence of points $x_{k+1} = x_k + d \cdot \lambda^k$. After $n \approx \frac{\log(1 + \frac{(x^* - x_0)(\lambda - 1)}{d})}{\log \lambda}$ steps we arrive at a point x_n such that $x_{n-1} \leq x^* \leq x_n$, and $x_n - x_{n-1} < d \cdot \lambda^n$.

We now set $I_0 = [x_{n-1}, x_n]$ and construct recursively a sequence of nested intervals $I_k \subset I_{k-1}$ with mid-points y_k . Here I_k is either the left half or the right half of I_{k-1} , depending on whether $y_{k-1} \geq x^*$ or $y_{k-1} \leq x^*$. After $\log_2 \frac{x_n - x_{n-1}}{\epsilon} < \log_2 \frac{d \cdot \lambda^n}{\epsilon}$ steps we arrive at an interval of length ϵ which contains a minimizer.

Line search with golden ratio: We wish to solve the same problem as before, but this time given a point x we only have information on the value $f(x)$.

We commence with two initial points $x_0 < x_1$ and determine the corresponding function values. Suppose without loss of generality that $f(x_0) > f(x_1)$. Then there exists a minimizer $x^* > x_0$. We compute again a sequence of points $x_{k+1} = x_k + d \cdot \lambda^k$, until after $n = O(\log \frac{x^* - x_0}{d})$ steps we arrive at a point x_n such that $f(x_n) \geq f(x_{n-1})$. Thus we have three points $x_{n-2} < x_{n-1} < x_n$ such that $f(x_{n-1}) \leq \min(f(x_{n-2}), f(x_n))$. It follows that the minimizer lies in the interval $I_0 = [x_{n-2}, x_n]$.

We now construct recursively a sequence of nested intervals $I_{k+1} \subset I_k$ which contain the minimizer. At each step we have already computed the end-points x_{k-}, x_{k+} of the interval I_k and a point x_{km} in its interior. Assume now that $\frac{x_{k+} - x_{km}}{x_{k+} - x_{k-}} = \alpha = \frac{\sqrt{5}-1}{2} \approx 0.618$, i.e., x_{km} divides I_k at a golden ratio. To this end we have to choose $d = \frac{\alpha}{1-\alpha} = 1 + \alpha = \frac{\sqrt{5}+1}{2}$. We choose the next iterate $x_{k+1} = x_{k-} + x_{k+} - x_{km}$ such that it divides I_k also at a golden ratio, but lies in the other half of I_k . If now $f(x_{k+1}) > f(x_{km})$, then we discard the end-point x_{k+} of I_k which is closer to x_{k+1} , if $f(x_{k+1}) < f(x_{km})$, then we discard the other end-point x_{k-} which is closer to x_{km} . We again end up with three points $x_{k\pm}, x_{km}, x_{k+1}$, two of which are the end-points of the new interval I_{k+1} and one of which divides this interval at a golden ratio. At each step the length of the interval decreases by a factor α . After $\log_{\alpha^{-1}} \frac{x_n - x_{n-2}}{\epsilon}$ steps we arrive at an interval of length ϵ which contains a minimizer.

In the line search methods considered above, the number of known digits of the minimizer increases linearly with the index k of the iteration. Such a convergence behaviour is called *linear*.

2.2 Ellipsoid method

If a *separating oracle* is available, then there exists a simple polynomial-time iterative method to solve general convex optimization problems, the *ellipsoid method*. Here a separating oracle, given a point $x \in \mathbb{R}^n$, outputs a non-zero vector g such that a minimizer x^* of the problem is guaranteed to satisfy the condition $\langle g, x^* - x \rangle \leq 0$. This may be, e.g., the gradient $f'(x)$ of the cost function at a feasible point x .

The initial data are an ellipsoid $E_0 = \{x \in \mathbb{R}^n \mid (x - x_0)^T P_0^{-1} (x - x_0) \leq 1\}$ centered at the initial iterate x_0 and containing a minimizer. At iteration step k we call the separation oracle at the center x_k of the ellipsoid E_k to obtain a vector g_{k+1} such that $\langle g_{k+1}, x_k - x^* \rangle \geq 0$. Then we compute the new ellipsoid E_{k+1} centered on the new iterate x_{k+1} by

$$x_{k+1} = x_k - \frac{1}{n+1} \frac{P_k g_{k+1}}{\sqrt{g_{k+1}^T P_k g_{k+1}}},$$

$$P_{k+1} = \frac{n^2}{n^2 - 1} \left(P_k - \frac{2}{n+1} \frac{P_k g_{k+1} (P_k g_{k+1})^T}{g_{k+1}^T P_k g_{k+1}} \right).$$

The volume of the ellipsoids E_k decreases exponentially and each ellipsoid is guaranteed to contain a minimizer of the problem.

The method is impractical in most situations due to its slow convergence and numerical instabilities. It can be used for small-dimensional convex problems if little information other than the separation oracle is available.

2.3 Non-convex problems

Non-convex problems are in general much more difficult to solve than convex problems. When using local descent methods one can hope only for convergence to a local minimum.

Below we present a number of methods for unconstrained minimization.

Line search: At each iteration, we choose a *descent direction* d_k at the current iterate x_k . Then we solve the one-dimensional problem

$$\min_{\alpha > 0} f(x_k + \alpha d_k)$$

as a sub-problem. Note that we do not need to solve the sub-problem to high accuracy. A coarse approximation of the minimum will suffice for serving as the next iterate x_{k+1} .

If the function is of class C^1 , then we may choose $d_k = -f'(x_k)$ and obtain a gradient descent algorithm. Note that we implicitly assume the presence of a Euclidean structure on the underlying space, since we identify the co-vector $-f'(x_k)$ with the vector d_k .

If the function is of class C^2 and with positive definite Hessian, then we may choose $d_k = -(f''(x_k))^{-1}f'(x_k)$. This will be a Newton-type method.

There exist schemes which combine directions from several previous steps and achieve an acceleration of convergence. Quasi-Newton schemes build an approximation of the Hessian based on previously computed values of the gradient.

Trust region methods: At each iteration, we construct an *approximation* of the cost function around the current iterate x_k . Then we construct a sub-problem from the original problem with the approximation instead of f and the additional constraint $\|x - x_k\| \leq c$. The feasible set of this constraint is the *trust region* where we are confident in our approximation of f . If the minimizer of the sub-problem gives a lower as expected decrease of the original cost function, then we decrease the constant c .

Regularization: At each iteration we add a *regularizing term* to the objective function, usually proportional to $\|x - x_k\|^p$ for some $p \geq 1$, and approximate the objective function by, e.g., a second order Taylor polynomial. The regularizing term has the effect of holding the next iterate close to the previous one, where the approximation is valid.

Constrained optimization problems can be converted to unconstrained ones by adding barriers or penalty functions. While the former are defined on the feasible set and tend to $+\infty$ at its boundary, the latter are zero on the feasible set and penalize constraint violations by an increasing value outside of the feasible set.

Another strategy is to approximate the original problem globally by a convex one with polynomial complexity (relaxation). Relaxations may consist in

- dropping non-convex constraints (rank constraints, integer or binary constraints);
- replacing non-convex constraints by stronger or weaker ones (convexification of objective, replace positivity conditions on polynomials by sums of squares conditions).

Problems with integer or binary constraints can be solved by dedicated branch-and-bound methods. Here the feasible set is split into several smaller parts (branching), and the problem restricted to a given part is relaxed and solved (bounding), resulting in a search tree.

These methods will be considered in more detail in the second part of the course.

3 Mathematical background

3.1 Norms

Optimization problems can rarely be solved exactly. Most often a solution algorithm delivers a sequence of iterates that converge to an optimal solution of the problem. In order to define convergence and to measure the quality of approximation, we shall need the notions of topology and norm.

Definition 3.1. Let V be a real vector space. A function $\cdot : V \times V \rightarrow \mathbb{R}$ is called a *scalar product* if

- $u \cdot v = v \cdot u$ for all $u, v \in V$,
- $(au + bv) \cdot w = a(u \cdot w) + b(v \cdot w)$ for all $a, b \in \mathbb{R}, u, v, w \in V$,
- $u \cdot u \geq 0$ for all $u \in V$ and $u \cdot u = 0$ if and only if $u = 0$.

Equivalently, a scalar product is a symmetric positive definite bilinear form on V .

Definition 3.2. Let V be a real vector space. A function $\|\cdot\| : V \rightarrow \mathbb{R}_+$ is called a *norm* if

- $\|u\| = 0$ if and only if $u = 0$,
- $\|au\| = |a| \cdot \|u\|$ for all $a \in \mathbb{R}$ and $u \in V$,
- $\|u + v\| \leq \|u\| + \|v\|$ for all $u, v \in V$.

Every scalar product on V defines a norm by $\|u\| = \sqrt{u \cdot u}$, but not every norm can be represented in such a way. Norms defined by scalar products are called *Euclidean*.

Every norm on V defines a distance function, or *metric*, by $d(u, v) := \|u - v\|$.

Two norms $\|\cdot\|, \|\cdot\|'$ are called *strongly equivalent* if there exist constants $\alpha, \beta > 0$ such that

$$\alpha\|u\| \leq \|u\|' \leq \beta\|u\|$$

for all $u \in V$. In finite-dimensional vector spaces every two norms are strongly equivalent, because the continuous positive function $\frac{\|u\|}{\|u\|'}$ attains its minimum and its maximum on the unit sphere.

The *unit ball* and the *open unit ball* of a norm $\|\cdot\|$ are the sets

$$B_1 = \{u \in V \mid \|u\| \leq 1\}, \quad B_1^o = \{u \in V \mid \|u\| < 1\},$$

respectively.

Examples:

- p -norms on \mathbb{R}^n : $\|x\|_p = (\sum_{k=1}^n |x_k|^p)^{1/p}$, $p \in [1, +\infty)$; $\|x\|_\infty = \max_k |x_k|$;
- Euclidean norms on \mathbb{R}^n : $\|x\| = \sqrt{x^T A x}$, A real symmetric with all eigenvalues positive;
- matrix (Schatten) p -norms on $\mathbb{R}^{n \times m}$: $\|A\|_p = (\sum_{k=1}^{\min(n,m)} \sigma_k(A)^p)^{1/p}$, $p \in [1, +\infty)$; $\|A\|_\infty = \max_k \sigma_k(A)$;
- matrix (Schatten) p -norms on \mathcal{S}^n : $\|A\|_p = (\sum_{k=1}^n |\lambda_k(A)|^p)^{1/p}$, $p \in [1, +\infty)$; $\|A\|_\infty = \max_k |\lambda_k(A)|$.

Here \mathcal{S}^n is the $\frac{n(n+1)}{2}$ -dimensional real vector space of real symmetric $n \times n$ matrices, $\sigma_k(A)$ are the singular values, and $\lambda_k(A)$ the eigenvalues of the matrix A .

The matrix 1-norms are also called *nuclear norms*, the matrix 2-norms *Frobenius norms*, and the matrix ∞ -norms *spectral norms*. The Frobenius norm can be computed as $\|A\|_2 = \sum_{i,j} A_{ij}^2$ without resorting to the eigenvalues. On positive semi-definite matrices the nuclear norm equals the trace $\sum_i A_{ii}$.

Attention: There exist also other matrix p -norms defined by $\|A\|_p = \sup_{\|x\|_p=1} \|Ax\|_p$ which are derived from the p -norms on the vector spaces the matrix is acting on as a linear operator. In this notion the spectral norm is, e.g., given by $\|A\|_2$, not $\|A\|_\infty$.




Figure 2: Unit balls of some vector p -norms in the plane

0-”norms” and sparsity: The expressions for the p -norms above are also defined for $p \in (0, 1)$, but these are not anymore norms, because the unit balls cease to be convex, which is equivalent to a violation of the triangle inequality. Nevertheless, the expressions

$$\lim_{p \rightarrow 0} \sum_{k=1}^n |x_k|^p, \quad \lim_{p \rightarrow 0} \sum_{k=1}^{\min(n,m)} \sigma_k(A)^p, \quad \lim_{p \rightarrow 0} \sum_{k=1}^n |\lambda_k(A)|^p$$

are of interest for optimization, because the measure the number of non-zero components of a vector $x \in \mathbb{R}^n$ and the rank of a matrix $A \in \mathbb{R}^{n \times m}$ or $A \in \mathcal{S}^n$, respectively.

It is sometimes a desirable feature of the solution of an optimization problem that it be sparse (low number of non-zero components) or low rank in case of a matrix decision variable. In order to enforce such behaviour, one can add a penalty term $\mu\|x\|_1$ or $\mu\|A\|_1$ to the objective function of the problem, where $\mu > 0$ is a weight which emphasizes the importance of the sparsity / low rank property over the original objective function value of the solution.

Dual norm: Let $\|\cdot\|$ be a norm on a vector space V . The *dual norm* is defined on the dual vector space V^* by

$$\|y\|_* = \sup_{\|x\| \leq 1} \langle x, y \rangle.$$

Exercise: Show that the dual to the p -norm $\|\cdot\|_p$ is the q -norm $\|\cdot\|_q$ with q defined by $\frac{1}{p} + \frac{1}{q} = 1$.

Scalar products and duality: A scalar product on a vector space V can be interpreted as a bijection between V and its dual:

$$x \mapsto \langle x, \cdot \rangle = (y \mapsto \langle x, y \rangle).$$

If in some algorithm dual vectors are identified with primal vectors, then implicitly the presence of a scalar product is assumed. Such algorithms are *not* invariant with respect to general affine transformations of the coordinate system, they are invariant only with respect to orthogonal transformations preserving the scalar product.

In particular, in gradient descent methods gradients (dual vectors) are interpreted as directions (primal vectors). Hence, e.g., a change of the unit of measurement in one of the coordinates modifies the sequence of iterates produced by the method.

3.2 Affine space

We have seen that the notion of convexity relies on the ability to define the segment $[x, y] = \{\lambda x + (1 - \lambda)y \mid \lambda \in [0, 1]\}$ between points of the set under consideration. Obviously the notion of segment is invariant not only under automorphisms of the underlying vector space, but also under translations which do not preserve the zero vector and are hence not automorphisms. We therefore do not need the full structure of a vector space in order to work with convex sets. It is sufficient to keep the structure of an affine space, which can be obtained from a vector space by forgetting the location of the zero vector.

Definition 3.3. Let V be a real vector space. An *affine space* with associated vector space V is a set A together with a map $+$: $A \times V \rightarrow A$ such that

- $x + 0 = x$ for all $x \in A$,
- $(x + u) + v = x + (u + v)$ for all $x \in A, u, v \in V$,
- $v \mapsto x + v$ is a bijection between V and A for all $x \in A$.

Example: An affine subspace A of a vector space W is an affine space in the sense above. The associated vector space V is the linear subspace of W spanned by the differences of the elements of A . In particular, the vector space W itself is also an affine space, with the associated vector space being again W .

The third property of the definition allows to consider the elements of V as differences between points in A : for $x, y \in A$ we define $x - y$ to be the unique vector $v \in V$ such that $x = y + v$. The points of the segment $[x, y]$, where $x, y \in A$, can then be written as

$$\lambda x + (1 - \lambda)y = y + \lambda(x - y).$$

Since $x - y$ is a vector, its multiple $\lambda(x - y)$ is also a vector, and hence the right-hand side of the relation is again a point of the affine space. Generally, every combination of points of affine space with coefficients summing to 0 can be interpreted as a vector in V , since it can be written as a linear combination of differences of elements of A . Therefore every combination with coefficients summing to 1 can be seen again as a point in affine space, because it can be written as an element of affine space plus a combination of points with coefficients summing to 0.




Figure 3: Affine hulls of sets of points

Definition 3.4. Let x_1, \dots, x_k be points in an affine space A . Then $\sum_{i=1}^k \lambda_i x_i$ is called an *affine combination* of the points x_1, \dots, x_k if $\sum_{i=1}^k \lambda_i = 1$.

A finite set $\{x_1, \dots, x_k\}$ of points in an affine space A is called *affinely independent* if the relations $\sum_{i=1}^k \lambda_i x_i = 0$, $\sum_{i=1}^k \lambda_i = 0$ imply $\lambda_i = 0$ for all $i = 1, \dots, k$.

An *affine basis* of an affine space A is an affinely independent set $\{x_1, \dots, x_n\}$ such that every element of A is an affine combination of x_1, \dots, x_n .

If $\{x_1, \dots, x_n\}$ is an affine basis of A , then $\{x_2 - x_1, \dots, x_n - x_1\}$ is a basis of the vector space associated to A . Hence the dimension of A is equal to $n - 1$, i.e., the number of elements in any of its affine bases, minus 1.

Definition 3.5. The *affine hull* of a subset $X \subset A$ of an affine space is the set of all affine combinations of elements of X .

An *affine subspace* of A is a subset of A which equals its affine hull.

The affine hull of an arbitrary subset X is the smallest affine subspace of A which contains X , or equivalently, the intersection of all affine subspaces containing X .

Given an affine basis $\{x_1, \dots, x_n\}$ of A , we can represent every point $x \in A$ uniquely as an affine combination $x = \sum_{i=1}^n \lambda_i x_i$. The coefficients $\lambda_1, \dots, \lambda_n$ are called the *barycentric coordinates* of x with respect to the basis $\{x_1, \dots, x_n\}$.

Definition 3.6. A map $f : A \rightarrow B$ between affine spaces A, B with associated vector spaces U, V is called *affine* if it preserves affine combinations, i.e., $f(\sum_k \lambda_k x_k) = \sum_k \lambda_k f(x_k)$ for all $x_k \in A$ and all λ_k such that $\sum_k \lambda_k = 1$.

An *affine isomorphism* is a bijective affine map.

If we turn the affine spaces into vector spaces by arbitrarily designating a point x_0 as the origin and identifying an arbitrary point x of the affine space with the vector $x - x_0$, then the affine maps turn out to be linear maps plus translations by a constant vector.

3.3 Topology

Defining a topology on a set allows one to define the notion of convergence of sequences of points to points of the set. This is necessary since usually the output of an optimization algorithm will not consist of an optimal solution of the optimization problem, but merely of a sequence of iterates which converge to an optimal solution.

Definition 3.7. Let X be a set. A collection \mathcal{T} of subsets $U \subset X$ is called a *topology* on X if

- $\emptyset, X \in \mathcal{T}$,
- finite intersections of sets in \mathcal{T} are again in \mathcal{T} ,
- arbitrary unions of sets in \mathcal{T} are again in \mathcal{T} .

The set X equipped with a topology \mathcal{T} is called a *topological space*.

The topology *defines* which subsets of X are *open* and which are *closed*: $U \subset X$ is open if $U \in \mathcal{T}$ and it is closed if $X \setminus U \in \mathcal{T}$. If $x \in X$ is a point, then any open set containing x is called a *neighbourhood* of x .

In this way, finite unions of closed sets and arbitrary intersections of closed sets are again closed.

Infinite unions of closed sets do not need to be closed, however: in \mathbb{R}^n we have $\bigcup_{r < 1} B_r = B_1^o$, where B_r is the closed ball of radius r , and B_1^o the open unit ball.

Definition 3.8. Let $U \subset X$ be a subset of a topological space. The *interior* of U , denoted $\text{int}U$ or U^o , is the largest open set contained in U , or equivalently, the union of all open sets contained in U .

The *closure* of U , denoted $\text{cl}U$ or \bar{U} , is the smallest closed set containing U , or equivalently, the intersection of all closed sets containing U .

The *boundary* of U , denoted $\text{bd}U$ or ∂U , is the difference $\bar{U} \setminus U^o$.

Let U be a subset of an affine space. The *relative interior* of U , denoted $\text{ri}U$, is the interior of U in the topology of the affine hull of U .

The *relative boundary* of U , denoted $\text{rbd}U$, is the difference $\bar{U} \setminus \text{ri}U$.

Example: Let $X = \mathbb{R}$ and $U = \mathbb{Q}$ the subset of rational numbers. Then $U^o = \emptyset$ and $\bar{U} = \mathbb{R}$.

Definition 3.9. Let $U, V \subset X$ be subsets of a topological space. The U is called *dense* in V if $U \subset V \subset \bar{U}$.

The topology allows to define the notion of convergence.

Definition 3.10. Let X be a topological space. A sequence x_1, x_2, \dots of points *converges* to a point x^* if for every open set $U \subset X$ containing x^* , there exists a number N_U such that $x_k \in U$ for all $k \geq N_U$.

Thus $\{x_k\}$ converges to x^* if the sequence eventually enters and no more leaves arbitrarily small neighbourhoods of x^* .

Often it is convenient to define or describe topologies by the following simpler notion.

Definition 3.11. Let X be a set. A collection \mathcal{B} of subsets $U \subset X$ is called a *base* if

- $\bigcup_{U \in \mathcal{B}} U = X$,
- for every $U_1, U_2 \in \mathcal{B}$ and every $x \in U_1 \cap U_2$ there exists $U_3 \in \mathcal{B}$ such that $x \in U_3 \subset U_1 \cap U_2$.

Every base defines a topology by taking the open subsets of X to be arbitrary unions of elements of \mathcal{B} . On the other hand, every topology can be defined by a base, the largest such base being the topology itself.

A norm on a vector space V induces a *topology* on V by the base

$$\mathcal{B} = \{u + \varepsilon B_1^o \mid u \in V, \varepsilon > 0\}.$$

Thus a set $U \subset V$ is *open* if for every $u \in U$ there exists a constant $\varepsilon > 0$ such that $u + \varepsilon B_1 \subset U$.

The notion of convergence can then be reformulated as follows: a sequence $\{x_k\}$ converges to x^* if and only if $\lim_{k \rightarrow \infty} \|x_k - x^*\| = 0$.

Strongly equivalent norms define the same topology on V . In finite dimension all norms are strongly equivalent, hence every norm defines the same topology on V .

Definition 3.12. Let S be a set. A *hull operator* on S is a map $\text{cl} : 2^S \rightarrow 2^S$ assigning subsets of S to subsets of S , and satisfying the following properties:

- extensivity: $X \subset \text{cl} X$ for all $X \subset S$
- monotonicity: $\text{cl} X \subset \text{cl} Y$ for all $X \subset Y \subset S$
- idempotence: $\text{cl}(\text{cl} X) = \text{cl} X$ for all $X \subset S$.

The closure operator in a topological space is a hull operator, but not every hull operator defines a topology on S . For this also the property

$$\text{cl}(X \cup Y) = \text{cl} X \cup \text{cl} Y \quad \forall X, Y \subset S$$

is needed.

The affine hull is an example of a hull operator violating this condition. We shall see more examples of hull operators below.




Figure 4: Convex hulls of sets of points

4 Convex sets

Convex sets are defined via affine combinations of two elements with nonnegative coefficients.

Definition 4.1. A subset $X \subset A$ of a real vector space or a real affine space is called *convex* if for all $x, y \in X$ and all $\lambda \in [0, 1]$ we have

$$\lambda x + (1 - \lambda)y \in X.$$

Examples:

- the empty set \emptyset ,
- the whole space A ,
- singletons $\{x\}$,
- affine subspaces,
- open or closed affine half-spaces,
- open or closed norm balls $x + rB_1^o$, $x + rB_1$ around arbitrary points.

Here open and closed affine half-spaces are sets of the form $\{x \in A \mid a(x) < b\}$ and $\{x \in A \mid a(x) \leq b\}$, respectively, where a is a non-constant linear functional on A and $b \in \mathbb{R}$.

4.1 Convex hull

Definition 4.2. Let x_1, \dots, x_k be points in an affine space A . Then $\sum_{i=1}^k \lambda_i x_i$ is called a *convex combination* of the points x_1, \dots, x_k if $\sum_{i=1}^k \lambda_i = 1$ and $\lambda_i \geq 0$, $i = 1, \dots, k$.

The *convex hull* of a subset $X \subset A$ of an affine space is the set of all convex combinations of elements of X . It is denoted by $\text{conv}X$.

Lemma 4.3. A set X is convex if and only if it equals its convex hull.

Proof. Let $X = \text{conv}X$. Then, in particular, convex combinations of any two elements of X belong to X . Hence X is convex.

Let X be convex. We show by induction on k that a convex combination of k elements of X is in X . The definition of convexity yields the base of the induction for $k = 2$. Suppose we have proven that any convex combination of $k - 1$ elements of X is in X . Let $x_1, \dots, x_k \in X$ and let $x = \sum_{i=1}^k \lambda_i x_i$ be a convex combination. If any of the coefficients λ_i vanishes, then x is actually a convex combination of strictly less than k elements and is in X by the induction hypothesis. Assume $\lambda_i > 0$ for all $i = 1, \dots, k$. Then we have

$$x = \sum_{i=1}^{k-1} \lambda_i x_i + \lambda_k x_k = \left(\sum_{i=1}^{k-1} \lambda_i \right) \sum_{i=1}^{k-1} \frac{\lambda_i}{\sum_{j=1}^{k-1} \lambda_j} x_i + \lambda_k x_k = (1 - \lambda_k) y + \lambda_k x_k.$$

Here $y = \sum_{i=1}^{k-1} \frac{\lambda_i}{\sum_{j=1}^{k-1} \lambda_j} x_i$ is a convex combination of $k-1$ elements of X and is hence in X . The point x has then been represented as convex combination of two elements of X and is hence also in X . \square

The following assertion follows immediately from Definition 4.1.

Lemma 4.4. *Arbitrary intersections of convex sets are convex.*

Corollary 4.5. *The convex hull of a set X is the smallest convex set which contains X , namely the intersection of all convex sets containing X .*

Proof. Since convex combinations of convex combinations are again convex combinations of the original points, the convex hull of X is equal to its own convex hull. By Lemma 4.3 it is hence convex. On the other hand, any convex set Y containing X must contain at least the convex hull of X , because $Y \supset X$ implies $Y = \text{conv}Y \supset \text{conv}X$. \square

Further examples of convex sets:

- polytopes (convex hulls of a finite set of points),
- polyhedra (finite intersections of closed affine half-spaces),
- simplices (convex hull of an affinely independent set of points).

4.2 Operations preserving convexity

We now consider more operations which preserve convexity.

Definition 4.6. Let X, Y be subsets of a vector space. The set

$$X + Y := \{x + y \mid x \in X, y \in Y\}$$

is called *Minkowski sum* of X, Y .

This definition can be extended to the case where one of the sets X, Y is a subset of an affine space and the other a subset of the underlying vector space.

The following assertions follow easily from the definition of convexity.


- the Minkowski sum of convex sets is convex,
- images of convex sets under affine maps are convex,
- pre-images of convex sets under affine maps are convex,
- the interior X° of a convex set X is convex,
- the relative interior $ri X$ of a convex set X is convex,
- the closure $cl X$ of a convex set X is convex.

We now come to the interplay between convexity and topology.

Lemma 4.7. *Let $X \neq \emptyset$ be convex. Then $ri X \neq \emptyset$.*

For non-convex sets this is in general not the case (consider $X = \mathbb{Q} \subset \mathbb{R}$, then $ri X = \emptyset$).

Proof. The affine hull $\text{aff} X$ possesses an affine basis of points in X . To construct such a basis, pick an arbitrary point $x_1 \in X$. If $\text{aff}\{x_1\} = \text{aff} X$, then $\{x_1\}$ is an affine basis of $\text{aff} X$. If $\text{aff}\{x_1\} \neq \text{aff} X$, then there exists a point $x_2 \in X \setminus \text{aff}\{x_1\}$. This point x_2 is affinely independent of x_1 . We now repeat the process by comparing $\text{aff}\{x_1, x_2\}$ with $\text{aff} X$ and adjoin another affinely independent point $x_3 \in X$ if these affine hulls are not equal. Obviously the affine hulls become equal after $\dim \text{aff} X + 1$ steps.


 Figure 5: Proof of Lemma 4.8. Radii are shown in *italic*.

Let hence $x_1, \dots, x_k \in X$ form an affine basis of the affine hull of X . Then the simplex $\Sigma = \text{conv}\{x_1, \dots, x_k\}$ is a subset of X , and the relative interior of Σ is given by the set

$$\text{ri } \Sigma = \left\{ \sum_{i=1}^k \lambda_i x_i \mid \lambda_i > 0, \sum_{i=1}^k \lambda_i = 1 \right\}.$$

Since $\text{aff } \Sigma = \text{aff } X$, any point in $\text{ri } \Sigma$ is also in $\text{ri } X$. □

We now need an auxiliary lemma.

Lemma 4.8. *Let X be a convex set, let $x \in \text{ri } X$ and $y \in \text{cl } X$. Then the half-open segment $[x, y) = \{\lambda x + (1 - \lambda)y \mid \lambda \in (0, 1]\}$ is a subset of $\text{ri } X$.*

Proof. By definition there exists $r > 0$ such that $(x + rB_1) \cap \text{aff } X \subset X$. Let $\lambda \in (0, 1]$ and $z = \lambda x + (1 - \lambda)y$. Set $\rho = \frac{\lambda r}{1 + \lambda}$. Since $y \in \text{cl } X$, there exists $w \in X$ such that $\|y - w\| < \rho$.

Set $u = x + w - y$. Then $u \in \text{aff } X$ as an affine combination of points in $\text{aff } X$. Moreover, $\|u - x\| = \|w - y\| < r$. Hence $(u + (r - \|u - x\|)B_1) \cap \text{aff } X \subset (x + rB_1) \cap \text{aff } X \subset X$. We then get

$$\lambda[(u + (r - \|u - x\|)B_1) \cap \text{aff } X] + (1 - \lambda)w = [z + w - y + \lambda(r - \|y - w\|)B_1] \cap \text{aff } X \subset X$$

by the convexity of X . But

$$z + w - y + \lambda(r - \|y - w\|)B_1 \supset z + (\lambda(r - \|y - w\|) - \|y - w\|)B_1$$

and $\lambda(r - \|y - w\|) - \|y - w\| = (1 + \lambda)(\rho - \|y - w\|) > 0$. Therefore $(z + (1 + \lambda)(\rho - \|y - w\|)B_1) \cap \text{aff } X \subset X$, and $z \in \text{ri } X$. □

This will allow us to show that for convex sets the relative interior and the closure can be obtained from each other.

Lemma 4.9. *Let X be a convex set. Then $\text{cl } \text{ri } X = \text{cl } X$ and $\text{ri } \text{cl } X = \text{ri } X$.*

Proof. Clearly $\text{cl } \text{ri } X \subset \text{cl } X$ and $\text{ri } \text{cl } X \supset \text{ri } X$.

Let $y \in \text{cl } X$. Then $X \neq \emptyset$ and there exists a point $x \in \text{ri } X$. It follows that $[x, y) \subset \text{ri } X$, and hence $y \in \text{cl } \text{ri } X$.

Let now $z \in \text{ri } \text{cl } X$. Then $X \neq \emptyset$ and there exists $x \in \text{ri } X$. Further there exists $\varepsilon > 0$ such that $(z + \varepsilon B_1) \cap \text{aff } X \subset \text{cl } X$. We have $[x, z] \subset \text{aff } X$, and there exists $y \in (z + \varepsilon B_1) \cap \text{aff } X$ such that y lies on the line through x and z and such that $z \in [x, y)$. But then $z \in \text{ri } X$ by Lemma 4.8. □




Figure 6: Were $z \in X$, then y' , which is closer to x than y , would also be in X .

4.3 Distance from a convex set

Let $X \subset \mathbb{R}^n$ be a non-empty closed convex set, and let $d(x, y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$ be the Euclidean distance function on \mathbb{R}^n . Then the distance function

$$D(x) = \inf_{y \in X} d(x, y)$$

is everywhere defined. Moreover, the infimum is attained for every $x \in \mathbb{R}^n$.

Indeed, let $z \in X$ be a point and set $r = \|x - z\|$. Then

$$\inf_{y \in X} d(x, y) = \inf_{y \in X \cap (x+rB_1)} d(x, y).$$

But $X \cap (x+rB_1)$ is compact, and hence the infimum on this intersection is attained. But then also the infimum over X is attained at the same point.

Moreover, the infimum is unique. Assume there exist two points $y, y' \in X$ such that $d(x, y) = d(x, y') = D(x)$. Then $D(x) \leq d(x, \frac{y+y'}{2}) = \sqrt{D^2(x) - \frac{\|y'-y\|^2}{4}}$, which gives $\|y' - y\|^2 \leq 0$ and hence $y' = y$.

Definition 4.10. Let X be a closed convex set and x be a point. The unique point $y \in X$ which minimizes the distance to x is called the *projection* of x on X .

Fix $x \in \mathbb{R}^n \setminus X$ and let $y \in X$ be such $D(x) = d(x, y) > 0$. Set $u = x - y$ and $e = \frac{u}{\|u\|}$ and define the closed half-space

$$H = \{z \in \mathbb{R}^n \mid e^T z \leq e^T y\}.$$

The boundary to this half-space is the hyperplane through y which is perpendicular to e , and H is the half-space opposite to x .

Claim: $X \subset H$. Indeed, let $z \in X \setminus H$. Then on the line segment $[z, y]$ there exists a point y' which is strictly closer to x than y . But $z, y \in X$, hence $y' \in X$, contradicting the optimality of y (see Fig. 6).

Therefore we have the inclusion $\{y\} \subset X \subset H$. Now consider $D(x')$ at an arbitrary point $x' \notin H$. We have the bounds

$$d(x', y) \geq \min_{y' \in X} d(x', y') = D(x') \geq \min_{y' \in H} d(x', y').$$

Set $u' = x' - x$ and decompose this vector into a component v which is orthogonal to e and a component αe which is collinear with e , $u' = \alpha e + v$ (see Fig. 7). We get

$$d(x', y) = \|(\alpha + \|u\|)e + v\| = \sqrt{(\alpha + \|u\|)^2 + \|v\|^2}, \quad \min_{y' \in H} d(x', y') = \alpha + \|u\|,$$

in the second case the minimizer $y' \in H$ of the distance to x' being given by $y' = y + v$. Therefore we obtain the explicit bounds

$$\alpha + \|u\| \leq D(x') \leq (\alpha + \|u\|) \sqrt{1 + \frac{\|v\|^2}{(\alpha + \|u\|)^2}} = \alpha + \|u\| + O(\|v\|^2).$$




Figure 7: Upper and lower bound on $D(x')$.

Since $D(x) = \|u\|$ and $\alpha = e^T u' = e^T(x' - x)$, we get

$$D(x') = D(x) + e^T(x' - x) + O(\|x' - x\|^2).$$

This decomposition can be seen as the first order Taylor approximation of D around x . It follows that $D(x)$ is differentiable at x and its gradient is given by e .

We get the following result.

Theorem 4.11. *Let $X \subset \mathbb{R}^n$ be a closed convex set, and let $D(x)$ be the Euclidean distance from x to X . Let further $u(x) \in \mathbb{R}^n$ be the difference between x and the projection of x on X . Then at any point $x \notin X$ the function $D(x)$ is differentiable, and its gradient is given by $\frac{u(x)}{\|u(x)\|}$.*

4.4 Separation

In this subsection we consider a fundamental property of convex sets, namely that convex sets can be separated or supported by hyperplanes. The latter property is closely linked to duality which we shall consider later.

Definition 4.12. Let X, Y be non-empty convex sets. The distinct parallel affine hyperplanes

$$H_1 = \{x \mid a^T x = b_1\}, \quad H_2 = \{x \mid a^T x = b_2\}$$

separate X and Y strongly if

$$\sup_{x \in X} a^T x \leq b_1 < b_2 \leq \inf_{y \in Y} a^T y.$$

We say that X and Y can be separated strongly if there exist affine hyperplanes which separate them strongly.

Here $a(x) := a^T x$ is a non-zero linear functional, represented by a non-zero vector in any given coordinate system.

We may also say that the linear functional a separates X, Y strongly.

Definition 4.13. Let X, Y be non-empty convex sets. The affine hyperplane

$$H = \{x \mid a^T x = b\}$$

separates X and Y properly if

$$\sup_{x \in X} a^T x \leq b \leq \inf_{y \in Y} a^T y, \quad \inf_{x \in X} a^T x < \sup_{y \in Y} a^T y.$$

We say that X and Y can be separated properly if there exists an affine hyperplane which separates them properly.

Thus H separates X and Y properly if X and Y lie in opposite closed half-spaces with respect to H and at least one of the sets X, Y is not contained in H .

We may also say that the linear functional a separates X, Y properly.




Figure 8: (a) H_1, H_2 separate X, Y strongly; (b) H separates X, Y properly; (c) H does not separate X, Y .

Clearly strong separation implies proper separation, but there exist non-intersecting closed convex sets which can be separated properly but not strongly.

Example: Consider the sets $X = \{(x_1, x_2) \mid x_1 > 0, x_1 x_2 \geq 1\}$, $Y = \{(y_1, y_2) \mid y_1 < 0, y_1 y_2 \leq -1\}$. Then the hyperplane $H = \{x \mid e_1^T x = 0\}$ separates X and Y properly, where $e_1 = (1, 0)^T$. However, X and Y cannot be separated strongly, because the distance between these sets is zero.

We shall now consider the special case when one of the sets is a singleton.

Lemma 4.14. *Let V be a real vector space equipped with a Euclidean distance function. Let X be a non-empty closed convex set and y a point. Then the singleton $\{y\}$ can be separated strongly from X if and only if $d(y, X) = \min_{x \in X} d(x, y) > 0$.*

Proof. If $y \in X$, then $\{y\}$ cannot be separated strongly from X . Hence assume that $y \notin X$. Then $d(y, X) > 0$. Let $y^* \in X$ be the projection of y on X . Set $a = y - y^* \neq 0$. Then the hyperplanes $H_1 = \{x \mid a^T x = a^T y\}$ and $H_2 = \{x \mid a^T x = a^T y^*\}$ separate X and $\{y\}$ strongly. \square

Lemma 4.15. *Let $X \subset V$ be a convex set and let $y \notin \text{ri } X$. Then $\{y\}$ and X can be properly separated.*

Proof. Introduce a Euclidean distance in the ambient vector space V .

If $y \notin \text{cl } X$, then $d(y, \text{cl } X) > 0$ and $\{y\}$ can be strongly separated from $\text{cl } X$. The same hyperplanes then separate $\{y\}$ from X strongly, and thus $\{y\}$ and X can be separated properly.

Hence assume that $y \in \text{cl } X \setminus \text{ri } X = \text{rbd } X$, where $\text{rbd } X$ denotes the relative boundary of X . Let A be the affine hull of X . Then there exists a sequence of points $\{y_k\}_{k \geq 1}$ such that $y_k \in A \setminus \text{cl } X$ and $\lim_{k \rightarrow \infty} y_k = y$. Let y_k^* be the projection of y_k on $\text{cl } X$, and let $a_k = y_k - y_k^*$, $e_k = \frac{a_k}{\|a_k\|}$. Let e^* be an accumulation point of the sequence $\{e_k\}$ which must exist because all e_k are elements of the unit sphere. Without loss of generality we may assume that $\{e_k\}$ converges to e .

For every fixed $z \in \text{cl } X$ we have $a_k^T y_k > a_k^T y_k^* \geq a_k^T z$ and hence $e_k^T y_k > e_k^T z$. Passing to the limit $k \rightarrow \infty$, we get that $e^T y \geq e^T z$. However, the linear functional $x \mapsto e^T x$ is not constant on A and thus not constant on X by construction. Therefore the hyperplane $H = \{x \mid e^T x = e^T y\}$ separates $\{y\}$ and $\text{cl } X$ properly. Thus it separates $\{y\}$ from X properly. \square

The main result of this subsection is the following *Separation Theorem*.

Theorem 4.16. *The non-empty convex sets X and Y can be separated properly if and only if their relative interiors do not intersect.*

Proof. For the sake of contradiction, suppose there exists $z \in \text{ri } X \cap \text{ri } Y$ and a hyperplane $H = \{x \mid a^T x = b\}$ which separates X and Y properly. Then

$$a^T z \leq \sup_{x \in X} a^T x \leq b \leq \inf_{y \in Y} a^T y \leq a^T z,$$

and hence all inequalities are satisfied with equality.

Let A be the affine hull of X . Then there exists $\varepsilon > 0$ such that $(z + \varepsilon B_1) \cap A \subset X$, and hence

$$\sup_{x \in (z + \varepsilon B_1) \cap A} a^T x \leq a^T z.$$

It follows that the linear functional $x \mapsto a^T x$ is constant on A and hence also on X . Likewise, this functional is constant on the affine hull of Y and hence also on Y . But this contradicts the second condition $\inf_{x \in X} a^T x < \sup_{y \in Y} a^T y$ implied by proper separation of X and Y by H .

Therefore, if X and Y can be separated properly, then $ri X \cap ri Y = \emptyset$. Let us prove the converse implication.

Assume that $ri X \cap ri Y = \emptyset$. Then $0 \notin ri X - ri Y = \{x - y \mid x \in ri X, y \in ri Y\}$. By Lemma 4.15 the singleton $\{0\}$ can then be properly separated from $ri X - ri Y$. This implies that there exists a non-zero a such that

$$0 \leq \inf_{x \in ri X, y \in ri Y} a^T(x - y), \quad 0 < \sup_{x \in ri X, y \in ri Y} a^T(x - y).$$

Dropping the relative interior will not change the infimum or the supremum, which yields

$$\sup_{x \in X} a^T x \leq \inf_{y \in Y} a^T y, \quad \inf_{x \in X} a^T x < \sup_{y \in Y} a^T y.$$

Therefore X and Y can be separated properly. □

A similar reasoning allows to prove the following necessary and sufficient condition for strong separation.

Theorem 4.17. *Let $X, Y \subset V$ be non-empty convex sets and let V be equipped with a Euclidean distance. Then X and Y can be separated strongly if and only if $d(X, Y) = \inf_{x \in X, y \in Y} d(x, y) > 0$.*

Proof. If the distinct parallel affine hyperplanes H_1, H_2 separate X and Y strongly, then $d(X, Y) \geq d(H_1, H_2) > 0$.

Let on the other hand $d(X, Y) > 0$. Then also $d(\{0\}, X - Y) = d(\{0\}, cl(X - Y)) > 0$. Let $a \neq 0$ be the projection of $\{0\}$ on $cl(X - Y)$. Then

$$\inf_{x \in X} a^T x - \sup_{y \in Y} a^T y = \inf_{z \in X - Y} a^T z = \inf_{z \in cl(X - Y)} a^T z = \|a\|^2$$

and hence $\inf_{x \in X} a^T x > \sup_{y \in Y} a^T y$. It is now easily seen that X and Y can be strongly separated. □

In particular, a compact convex set can always be strongly separated from a closed convex set if the sets do not intersect. This follows from the fact that the distance function reaches its minimum on the compact set.

Definition 4.18. Let X be a non-empty convex set and $x \in rbd X$. A hyperplane H which separates $\{x\}$ from X properly is called a *supporting hyperplane* to X at x .

By Lemma 4.15 supporting hyperplanes exist at every relative boundary point. Moreover, the intersection of X with a supporting hyperplane has strictly lower dimension than X , because X is not contained in the hyperplane by definition of proper separation.

Lemma 4.19. *Let $X \neq V$ be a closed convex set. Then X equals the intersection of all closed half-spaces which contain X .*

Proof. Let C be the intersection of all closed half-spaces which contain X . By assumption there exists a point $x \notin X$. Then $\{x\}$ can be strongly separated from X , and there exists a closed half-space containing X . Thus C is non-empty and $X \subset C$.

Let now $z \in C \setminus X$. Then $\{z\}$ can be strongly separated from X , and there exists a closed half-space which contains X but not z . This leads to a contradiction, and such z cannot exist. □

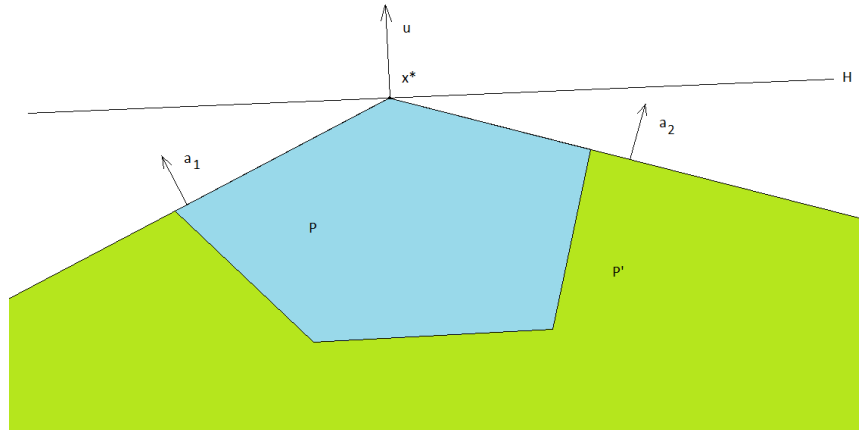


Figure 9: Lemma 4.20: The supporting linear functional u can be represented as a combination of the functionals a_1, a_2 corresponding to the active constraints with nonnegative coefficients.

4.5 Theorem on the alternative

Polyhedral sets are particularly simple convex sets. If a hyperplane is supporting to a polyhedral set P , then this relation can be certified in the form of a convex combination of the inequalities defining the polyhedral set.

Lemma 4.20. *Let $P = \{x \mid Ax \leq b\}$ be a non-empty polyhedral set, and let $H = \{x \mid u^T x = b_0\}$ be a hyperplane containing a point $x^* \in P$ and such that the open half-space $C = \{x \mid u^T x > b_0\}$ has an empty intersection with P , i.e., $u^T x^* = b_0$, $u^T x \leq b_0$ for all P . Then there exists a nonnegative vector $\mu \geq 0$ such that $u = A^T \mu$, $b_0 = b^T \mu$.*

Proof. Let I be the set of indices of rows for which the inequality $Ax^* \leq b$ is an equality, i.e., the index set of active constraints at x^* . Then there exists $\epsilon > 0$ such that for all $x \in x^* + B_\epsilon$ we have $(Ax)_j < b_j$ for all $j \notin I$. Define another polyhedral set by $P' = \{x \mid (Ax)_i \leq b_i \ \forall i \in I\}$. Then $P \subset P'$. We claim that C has also an empty intersection with P' , i.e., $u^T x \leq b_0$ for all $x \in P'$.

Indeed, suppose there exists $z \in P'$ such that $u^T z > u^T x^*$. Then for all $\lambda \in (0, 1]$ we have $z_\lambda = \lambda z + (1 - \lambda)x^* \in P'$ and $u^T z_\lambda > u^T x^*$. But for λ small enough we have $z_\lambda \in x^* + B_\epsilon$, and hence $z_\lambda \in P$, a contradiction.

We now define the polyhedral cone $K = \{A^T \mu \mid \mu \geq 0, \mu_j = 0 \ \forall j \notin I\}$. We claim that $u \in K$.

Indeed, suppose $u \notin K$. Then u can be separated from K , and there exists δ such that $u^T \delta > 0$, $v^T \delta \leq 0$ for all $v \in K$. In particular, $(A\delta)_i \leq 0$ for all $i \in I$. Hence $x^* + \delta \in P'$. But then $u^T(x^* + \delta) \leq u^T x^*$ and $u^T \delta \leq 0$, a contradiction.

Hence there exists $\mu \geq 0$, $\mu_j = 0$ for all $j \notin I$, such that $u = A^T \mu$. It follows that $\mu^T(Ax^* - b) = 0$ and $b_0 = \mu^T Ax^* = \mu^T b$, which yields the desired assertion. \square

As a consequence, we obtain the following Theorem on the Alternative.

Theorem 4.21. (Farkas) *Let $P = \{x \mid Ax \leq b\}$ be a polyhedral set. Then either $P \neq \emptyset$, or there exists $\mu \geq 0$ such that $\mu^T A = 0$, $\mu^T b = -1$.*

Proof. Clearly if $P \neq \emptyset$, then such a μ cannot exist.

Let $P = \emptyset$. Then the non-empty polyhedral set

$$P' = \{(x, t) \mid Ax - bt = (A, -b)(x^T, t)^T \leq 0\}$$

has an empty intersection with the open half-space $C = \{(x, t) \mid t = (0, 1)(x^T, t)^T > 0\}$. By the Lemma 4.20 there exists $\mu \geq 0$ such that $(0, 1) = \mu^T(A, -b)$. \square

We provide also a version of the theorem with equalities.

Corollary 4.22. Let $P = \{x \mid Ax \leq b, Cx = d\}$ be a polyhedral set. Then either $P \neq \emptyset$, or there exists $\mu \geq 0$, ν such that $\mu^T A + \nu^T C = 0$, $\mu^T b + \nu^T d = -1$.

Proof. Represent $P = \{x \mid Ax \leq b, Cx \leq d, -Cx \leq -d\}$. By the previous theorem, either $P \neq \emptyset$, or there exist $\mu, \nu_+, \nu_- \geq 0$ such that $\mu^T A + \nu_+^T C - \nu_-^T C = 0$, $\mu^T b + \nu_+^T d - \nu_-^T d = -1$. In the latter case, set $\nu = \nu_+ - \nu_-$. \square

4.6 Faces and extremal points

We shall now investigate the structure of the boundary of a convex set.

Definition 4.23. A convex subset F of a convex set X is called a *face* of X if for every line segment $l \subset X$ such that $F \cap \text{ri } l \neq \emptyset$ we have $l \subset F$.

A face F of X is called *proper* if $F \neq \emptyset$ and $F \neq X$.

A face F of a convex set X is called *exposed* if there exists a hyperplane H such that $F = X \cap H$ and $X \not\subset H$.

The hyperplane H in the last definition necessarily separates F and X properly.

Definition 4.24. Let X be a convex set. A point $x \in X$ is called *extremal* if the singleton $\{x\}$ is a face of X . It is called *exposed* if $\{x\}$ is an exposed face of X .

From the definition it follows that $x \in X$ is an extremal point of X if and only if $X \setminus \{x\}$ is convex.

Theorem 4.25. (Straszewicz, [1, Theorem 18.6]) Let X be closed convex. Then the set of exposed points of X is dense in the set of its extreme points.

Example: Let $\mathbb{R}^2 \supset X = ([-1, 0] \times [-1, 1]) \cup B_1$ be the union of a rectangle and a half-disc. The faces of X are

- X ,
- $\{-1\} \times [-1, 1]$, $[-1, 0] \times \{-1\}$, $[-1, 0] \times \{1\}$,
- $\{(x, y)\}$, where $x \geq 0$ and $x^2 + y^2 = 1$,
- \emptyset .

Of these, the zero-dimensional faces $\{(0, -1)\}$ and $\{(0, 1)\}$ are not exposed. They can be represented as limits of exposed points, however.

Lemma 4.26. Intersections of faces and faces of faces are again faces.

Proof. Let $F \subset X$ be a face of X and $F' \subset F$ a face of F . Let $l \subset X$ be a line segment such that $\text{ri } l \cap F' \neq \emptyset$. Then $\text{ri } l \cap F \neq \emptyset$, because $F' \subset F$. Hence $l \subset F$, because F is a face of X . But then $l \subset F'$, because F' is a face of F . Hence F' is a face of X .

Let F_1, F_2 be faces of X , and let $F = F_1 \cap F_2$. Let $l \subset X$ be a line segment such that $\text{ri } l \cap F \neq \emptyset$. Then $\text{ri } l \cap F_i \neq \emptyset$ for $i = 1, 2$, because $F \subset F_i$. Hence $l \subset F_i$, because F_i are faces of X . But then $l \subset F$, and F is also a face of X . \square

Definition 4.27. Let X be a closed convex set and $x \in X$. The *minimal face* of x in X is the smallest face of X containing x .

It is given by the intersection of all faces containing x . A face F is the minimal face of x if and only if $x \in \text{ri } F$. The set X then decomposes into a disjoint union of the relative interiors of its faces.

We shall now investigate when a closed convex set possesses extremal points. We need the following auxiliary result.

Lemma 4.28. Let X be a closed convex set. Suppose there exists a point $z \in X$ and a non-zero vector v such that the ray $\{z + \alpha v \mid \alpha \geq 0\}$ is contained in X . Then for every $x \in X$, the ray $\{x + \alpha v \mid \alpha \geq 0\}$ is contained in X .

Proof. Let $\mu > 0$ be arbitrary, and set $\alpha = \frac{\beta}{1-\lambda}$ for $\lambda \in (0, 1)$. Then we have

$$\lambda x + (1 - \lambda)(z + \alpha v) = \lambda x + (1 - \lambda)z + \beta v \in X$$

for all $\lambda \in (0, 1)$. Let $\lambda \rightarrow 1$, then by closedness of X we also get $x + \beta v \in X$. \square

Lemma 4.29. *Let X be a non-empty closed convex set. Then X has extremal points if and only if X does not contain a line (a 1-dimensional affine subspace).*

Proof. Suppose X contains a line l . Let $z, w \in l$ be two distinct points and set $v = z - w$. By the previous lemma, for every $x \in X$ we have $\{x + \alpha v \mid \alpha \in \mathbb{R}\} \subset X$. Hence through every point of X there runs a line which belongs to X . Therefore X cannot have extreme points.

Suppose now that X does not contain a line. We proceed by induction over the dimension of X . If $\dim X = 0$, then X is a singleton and contains an extremal point. Suppose we have shown the assertion of the lemma for all convex sets X with $\dim X < d$. Let $\dim X = d$. Choose a point $x \in X$ and consider a line l through x which lies in the affine hull of X . Since $l \not\subset X$, there exists a point $y \in l$ which lies on the relative boundary of X . Let H be an affine hyperplane separating y from X properly. Then $X' = H \cap X$ is a non-empty closed convex set of dimension $\dim X' < d$, and it does not contain a line. By the induction hypothesis, X' has extremal points. But these points have then also to be extremal points of X . \square

Lemma 4.30. *A compact convex set equals the convex hull of its extreme points.*

Proof. Clearly the convex hull of the extremal points is a subset of the set itself. We then have to show that every point is a convex combination of extremal points.

We also proceed by induction on the dimension. If $X = \emptyset$ or $\dim X = 0$, then the assertion is evident. Assume the assertion is proven for $\dim X < d$, and let X be a compact convex set of dimension d . Let $x \in X$ be a point, and let l be a line through x which lies in the affine hull of X . Then $l \cap X$ must be a closed line segment, because X is compact. The end-points y, z of this segment lie on the relative boundary of X and x is a convex combination of y, z . Let H_y, H_z be affine hyperplanes which separate y, z from X properly, and let $X_y = H_y \cap X, X_z = H_z \cap X$. The sets X_y, X_z are convex, compact, and have dimension strictly smaller than d . Hence y, z are convex combinations of extremal points of X_y, X_z , respectively, by the induction hypothesis. But these extremal points are also extremal points of X . Hence y, z are convex combinations of extremal points of X , and x is also such a combination. \square

Now we shall consider the extremal points of polyhedral sets.

Lemma 4.31. *A polyhedral set has a finite number of extremal points. In particular, a bounded polyhedral set is a polytope.*

Proof. Let $X = \{x \in \mathbb{R}^n \mid Ax \leq b\}$ be a polyhedral set, where $b \in \mathbb{R}^m$, and let $x \in X$ be an extremal point of X . Let a_1, \dots, a_m be the rows of A . Let $I = \{i \in \{1, \dots, m\} \mid a_i x = b_i\}$ be the index set of the active inequality constraints. In particular, we have $a_i x < b_i$ for all $i \notin I$. Then among the row vectors $a_i, i \in I$, there are n linearly independent vectors. Indeed, suppose there exists a non-zero vector $v \in \mathbb{R}^n$ such that $a_i v = 0$ for all $i \in I$. Then there exists $\varepsilon > 0$ such that $x \pm \varepsilon v \in X$. This contradicts the extremality of x .

Thus there exist n linearly independent row vectors a_{i_1}, \dots, a_{i_n} of A such that $a_{i_j} x = b_{i_j}$ for all $j = 1, \dots, n$. On the other hand, for any linearly independent set of n row vectors a_{i_1}, \dots, a_{i_n} of A the system $a_{i_j} x = b_{i_j}, j = 1, \dots, n$, uniquely determines a solution x . If this point x is in X , then it is an extremal point of X .

Since there exists only a finite number of sets of n linearly independent row vectors of A , the number of extremal points of X must also be finite. \square

In principle one can therefore compute all extremal points of a polyhedral set by checking for each set of n linearly independent row vectors a_{i_1}, \dots, a_{i_n} of A whether the solution of the system $a_{i_j} x = b_{i_j}, j = 1, \dots, n$ is feasible.

Finally, we consider extremal points of intersections of convex sets with affine subspaces.

Lemma 4.32. *Let $X \subset \mathbb{R}^n$ be a closed convex set and let A be an affine subspace of dimension m . Then every extreme point of $C = X \cap A$ is located in a face of X of dimension at most $n - m$.*

Proof. Let $x^* \in C$ be an extreme point, let $F \subset X$ be the minimal face of X containing x^* , and let $k = \dim F$. Then $x^* \in \text{ri } F$. We have $\dim F \cap A \geq m + k - n$. Since $x^* \in \text{ri}(F \cap A) \subset C$ and x^* is extremal in C we get $\dim F \cap A = 0$, and therefore $k \leq n - m$. \square

By a similar reasoning, if the minimal face of x^* in C has dimension m' , then $\dim F \cap A = m'$ and $\dim F = k \leq n + m' - m$.

4.7 Polar

Definition 4.33. Let $X \subset \mathbb{R}^n$ be a closed convex set containing the origin. The *polar* X° of X is defined as the set

$$\{p \in (\mathbb{R}^n)^* \mid \langle x, p \rangle \leq 1 \ \forall x \in X\}.$$

It follows that $(X^\circ)^\circ$.

Clearly if $X \subset X'$ are closed convex sets containing the origin, then $(X')^\circ \subset X^\circ$.

Lemma 4.34. Let $X \subset \mathbb{R}^n$ be a closed convex set containing the origin. Then X° is bounded if and only if the origin is an interior point of X .

Proof. We have $0 \in \text{int } X$ if and only if there exists $r > 0$ such that the ball of radius r around the origin is contained in X . But this is equivalent to the condition that X° is contained in the ball of radius r^{-1} around the origin. \square

Definition 4.35. Let X be a closed convex set containing the origin, and let F be a face of X . The *complementary face* $F^* \subset X^\circ$ is defined to be the set of points

$$\{p \in X^\circ \mid \langle x, p \rangle = 1 \ \forall x \in F\}.$$

Let $X \subset \mathbb{R}^n$ be a compact convex set containing the origin in its interior. Then the sets ∂X , ∂X° , $\{(x, p) \in X \times X^\circ \mid \langle x, p \rangle = 1\}$ are homeomorphic to the sphere S^{n-1} . While this is evident for the first two sets due to their compactness, the last one is in a continuous bijection with the 1-level set of the distance function to X .

4.8 Cones

In this section we consider properties of convex cones and present the most important cones for conic optimization.

4.8.1 Basic properties

Definition 4.36. A subset $C \subset V$ of a real vector space is called *conic set* if for every $x \in C$ and every $\lambda \geq 0$ we have $\lambda x \in C$. A convex conic set is called a *cone*.

A cone is called *regular* (also *proper*) if it is closed, has non-empty interior, and contains no lines.

Clearly intersections of conic sets are again conic sets. Therefore intersections of cones are again cones.

Examples:

- the full space V ,
- linear subspaces of V ,
- closed linear half-spaces,
- polyhedral cones (finite intersections of closed linear half-spaces),
- simplicial cones (linear bijective images of an orthant).

Here closed linear half-spaces are sets of the form $\{x \in V \mid a(x) \leq 0\}$, where a is a non-constant linear functional on V .

Definition 4.37. Let x_1, \dots, x_k be points in a vector space V . Then $\sum_{i=1}^k \lambda_i x_i$ is called a *conic combination* of the points x_1, \dots, x_k if $\lambda_i \geq 0$, $i = 1, \dots, k$.

The *convex conic hull* of a set X is the set of all conic combinations of elements of X .

The convex conic hull of X is the smallest convex cone which contains X , namely the intersection of all convex cones which contain X .

Definition 4.38. Let K be a closed convex cone containing no line. A point $x \in K$ is said to lie on an *extreme ray* if for every $x_1, x_2 \in K$ such that $x_1 + x_2 = x$ there exist nonnegative scalars λ_1, λ_2 such that $x_1 = \lambda_1 x$, $x_2 = \lambda_2 x$.

Definition 4.39. Let X be a convex set. The *recession cone* of X is the set of vectors v such that for all $x \in X$ and all $\alpha \geq 0$ we have $x + \alpha v \in X$.

If the set X is defined in an affine space, then the recession cone is a subset of the underlying vector space.

Lemma 4.40. Let $P = \{x \mid Ax \leq b\}$ be a non-empty polyhedral set. The recession cone of P is the polyhedral cone $K = \{v \mid Av \leq 0\}$.

Proof. Let $v \in K$ and $x \in P$. Then $A(x + \alpha v) = Ax + \alpha Av \leq Ax \leq b$ for all $\alpha \geq 0$. Hence v is in the recession cone of P .

Let v be in the recession cone of P , and let $x \in P$. Then $A(x + \alpha v) \leq b$ for all $\alpha \geq 0$. Hence we must have $Av \leq 0$, and $v \in K$. \square

Projective transformations: Cones allow us to define another class of transformations which preserve convexity.

Definition 4.41. Let $X \subset A$ be a convex subset of an affine space with underlying vector space V , and let $x_0 \in A$ be an arbitrary point. The set $\{(\alpha, \alpha u) \in \mathbb{R} \times V \mid \alpha \geq 0, x_0 + u \in X\}$ is called the *homogenization* of the set X .

The homogenization is a cone, and different choices of the point x_0 lead to linearly isomorphic homogenizations, so we may loosely speak of *the* homogenization of X .

Definition 4.42. Two convex sets $X, X' \subset A$ of an affine space A are called *projectively isomorphic* if their homogenizations are linearly isomorphic.

One set is then obtained by a projective transformation from the other.

Example: The open unit disc $B_1^o = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$ is projectively isomorphic to the unbounded set $\{(x, y) \in \mathbb{R}^2 \mid y > x^2\}$. Indeed, choosing x_0 as the origin of \mathbb{R}^2 , the homogenizations of these two sets are given by the cones

$$K_1 = \{(0, 0, 0)\} \cup \{(x_1, x_2, x_3) \mid x_1 > \sqrt{x_2^2 + x_3^2}\}, \quad K_2 = \{(0, 0, 0)\} \cup \{(x_1, x_2, x_3) \mid x_1 > 0, x_1 x_3 > x_2^2\}.$$

However, the linear map given by $\begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$ takes K_1 to K_2 .

4.8.2 Dual cones

Definition 4.43. Let $K \subset V$ be a cone. The set $K^* = \{y \in V^* \mid \langle y, x \rangle \geq 0 \forall x \in K\}$ is called the *dual cone* to K .

Here V^* is the vector space dual to V , i.e., the space of linear functionals on V , and $\langle y, x \rangle$ denotes the value of the functional y on the vector x .

Alternatively, if $\langle \cdot, \cdot \rangle$ is a scalar product on V , then K^* can be seen as a subset of V , with the same formula for the definition.

For finite-dimensional vector spaces, we may canonically identify V with the dual $(V^*)^*$ to the dual space. Namely, $x \in V$ is identified with the linear functional $p \mapsto \langle p, x \rangle$ on V^* .

Remark: For infinite-dimensional complete normed vector spaces (Banach spaces), "linear functional" is replaced by "continuous linear functional", and V is in general only a subset of $(V^*)^*$.

The dual cone hence consists of all linear functionals on V which take on only nonnegative values on K . It is not hard to see that K^* is a closed convex cone.

Lemma 4.44. *Let K be a closed convex cone. Then $(K^*)^* = K$.*

Proof. For every $x \in K$ and $p \in K^*$ we have $\langle p, x \rangle \geq 0$. Hence $K \subset (K^*)^*$.

Suppose for the sake of contradiction that there exists $y \in (K^*)^*$ such that $y \notin K$. Then y has a positive distance from K , since K is closed, and can be strongly separated from K . Therefore there exists a linear functional $p \in V^*$ such that $\langle p, y \rangle < 0$ and $\langle p, x \rangle \geq 0$ for all $x \in K$. From the second condition it follows that $p \in K^*$. But this is in contradiction with $y \in (K^*)^*$. Thus $(K^*)^* \subset K$. \square

Lemma 4.45. *Let K be a closed convex cone. Then K has non-empty interior if and only if K^* does not contain a line, and K does not contain a line if and only if K^* has non-empty interior.*

Proof. Let K have empty interior. Then K is contained in a proper linear subspace, and there exists a non-zero linear functional $p \in V^*$ such that $\langle p, x \rangle = 0$ for all $x \in K$. Then K^* contains the line $\{\alpha p \mid \alpha \in \mathbb{R}\}$.

On the other hand, let $p \in V^*$ be a non-zero functional such that the line $\{\alpha p \mid \alpha \in \mathbb{R}\}$ is contained in K^* . Then $\langle p, x \rangle \geq 0$ and $\langle -p, x \rangle \geq 0$, and hence $\langle p, x \rangle = 0$ for all $x \in K$. Hence K is contained in a proper linear subspace and has empty interior.

The second assertion is proved in a similar way. \square

Therefore the dual cone K^* of a closed convex cone K is regular if and only if K itself is regular.

We now consider how K^* behaves with respect to linear isomorphisms. We introduce the notion of the adjoint map. If $A : V \rightarrow W$ is a linear map between vector spaces, then its *adjoint* $A^* : W^* \rightarrow V^*$ is the unique linear map satisfying $\langle y, Ax \rangle = \langle A^*y, x \rangle$ for all $x \in V$, $y \in W^*$. Here $\langle \cdot, \cdot \rangle$ is the dual pairing between W^* and W on the left-hand side and between V^* and V on the right-hand side of the equation.

Note that if a basis $\{e_1, \dots, e_n\}$ is given in V , V^* is equipped with the dual basis $\{e^1, \dots, e^n\}$ (i.e., $\langle e_i, e^j \rangle = \delta_{ij}$ with $\delta_{ij} = 0$ for $i \neq j$ and $\delta_{ij} = 1$ for $i = j$ being the Kronecker symbol), and the linear automorphism A is expressed as a real $n \times n$ matrix in this basis, then A^* expressed in the dual basis will have coordinate matrix A^T .

Lemma 4.46. *Let $K \subset V$ be a convex cone, and let $K' = A[K]$ be its image under a linear isomorphism A of V . Then $(K')^* = (A^*)^{-1}[K^*]$.*

Proof. We have

$$\begin{aligned} (K')^* &= \{y \in V^* \mid \langle x, y \rangle \geq 0 \ \forall x \in K'\} = \{y \in V^* \mid \langle Ax, y \rangle \geq 0 \ \forall x \in K\} \\ &= \{y \in V^* \mid \langle x, A^*y \rangle \geq 0 \ \forall x \in K\} = \{y \in V^* \mid A^*y \in K^*\} = (A^*)^{-1}[K^*]. \end{aligned} \quad \square$$

Lemma 4.46 allows to compute the dual cone of a cone K if the dual of an isomorphic cone is already known. The following result is helpful for computing the dual cone if such information is not available.

Lemma 4.47. *Let $K \subset V$ be a regular convex cone. A non-zero point $y \in V^*$ lies on the boundary of K^* if and only if the kernel H of y is a supporting hyperplane to K at some non-zero boundary point $x \in \partial K$ such that K lies in the closed half-space with respect to H where y assumes nonnegative values.*

Proof. First note that y is by definition a linear functional on V . If $y \neq 0$, then its kernel $H = \{x \in V \mid \langle x, y \rangle = 0\}$ is indeed a linear subspace of V of codimension 1.

Let $y \in \partial K^* \setminus \{0\}$. The linear functional y is then non-constant on V , and in particular on K . Then $\langle x, y \rangle > 0$ for all x in the interior of K , otherwise there would exist $x \in K$ such that $\langle x, y \rangle < 0$. Hence the interior of K lies in the open half-space delineated by the hyperplane H where y assumes positive values.

Since $y \in \partial K^*$, there exists a sequence of points $y_k \in V \setminus K^*$ which converges to y . To every such $y_k \notin K^*$ we find a point $x_k \in K$ such that $\langle x_k, y_k \rangle < 0$. Since $x_k \neq 0$, we may normalize it such that $\|x_k\| = 1$ in some fixed Euclidean norm on V . Then there exists a subsequence of $\{x_k\}$ which converges to some point $x \in K$ of unit norm, because the intersection of the unit sphere with K is compact. By continuity we have $\langle x, y \rangle \leq 0$ and hence $\langle x, y \rangle = 0$, because $x \in K$ and $y \in K^*$. Therefore $x \in H$, and H separates $\{x\}$ properly from K . This proves one direction.

Let now $x \in \partial K$ be non-zero, and let H be a hyperplane separating $\{x\}$ from K properly. Then K lies in one of the closed half-spaces delineated by H . Let $y \in V^*$ be non-zero such that H is the kernel of y and such that y assumes nonnegative values on K . Then $y \in K^*$ by definition. However, $\langle x, y \rangle = 0$, because $x \in H$, and x is not constant as a linear functional on V^* . Therefore $y \notin \text{int}K^*$, which shows the other direction. \square

In order to find the boundary of K^* , we hence have to compute all hyperplanes which support K at non-zero boundary points. If the boundary $\partial K \setminus \{0\}$ is smooth, then at every non-zero boundary point of K there exists exactly one such hyperplane, namely the tangent plane to ∂K . In this case $\partial K^* \setminus \{0\}$ is comprised of the normals to all tangent planes to $\partial K \setminus \{0\}$.

References

- [1] R. Tyrrell Rockafellar. *Convex Analysis*, volume 28 of *Princeton Mathematical Series*. Princeton University Press, Princeton, 1970.