

II. Unconstrained convex optimization

Unconstrained optimization

We will start our discussion about solving convex optimization programs by considering the unconstrained case. Our template problem is

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad f(\mathbf{x}), \quad (1)$$

where f is convex. While we state this problem as a search over all of \mathbb{R}^N , almost everything we say here can be applied to minimized a convex function over an *open* set.¹

Before we go too deep into optimization, however, we need to provide a bit more mathematical rigor to how we think about convexity.

Convex sets

Before we discuss convex *functions*, it will be useful to start with the mathematical fundamentals of convex *sets*.

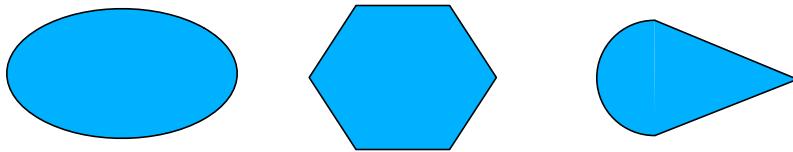
A set $\mathcal{C} \subset \mathbb{R}^N$ is **convex** if

$$\mathbf{x}, \mathbf{y} \in \mathcal{C} \Rightarrow (1 - \theta)\mathbf{x} + \theta\mathbf{y} \in \mathcal{C} \quad \text{for all } \theta \in [0, 1].$$

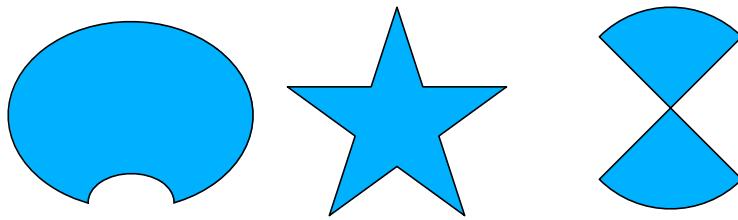
In English, this means that if we travel on a straight line between any two points in \mathcal{C} , then we never leave \mathcal{C} .

¹A formal definition of what it means for a set to be open is provided in the technical details at the end of these notes. Informally, an open set is one that doesn't have a boundary. The standard example of such a set is an open interval on the real line, e.g., $(0, 1)$. In the context of constrained optimization where the constraint set has a boundary, we must consider the fact that the solution can (and probably is) on this boundary, which complicates the picture considerably.

These sets in \mathbb{R}^2 are convex:



These sets are not:



Examples of convex (and nonconvex) sets

- Subspaces. Recall that if \mathcal{S} is a subspace of \mathbb{R}^N , then $\mathbf{x}, \mathbf{y} \in \mathcal{S} \Rightarrow a\mathbf{x} + b\mathbf{y} \in \mathcal{S}$ for all $a, b \in \mathbb{R}$. So \mathcal{S} is clearly convex.
- Bound constraints. Rectangular sets of the form

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^N : \ell_1 \leq x_1 \leq u_1, \ell_2 \leq x_2 \leq u_2, \dots, \ell_N \leq x_N \leq u_N\}$$

for some $\ell_1, \dots, \ell_N, u_1, \dots, u_N \in \mathbb{R}$ are convex.

- The *simplex* in \mathbb{R}^N

$$\{\mathbf{x} \in \mathbb{R}^N : \sum_{n=1}^N x_n = 1, x_n \geq 0\}$$

is convex.

- The “*filled-in*” *simplex* in \mathbb{R}^N

$$\{\mathbf{x} \in \mathbb{R}^N : \sum_{n=1}^N x_n \leq 1, x_n \geq 0\}$$

is convex.

- Any subset of \mathbb{R}^N that can be expressed as a set of linear inequality constraints

$$\{\mathbf{x} \in \mathbb{R}^N : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$$

is convex. Notice that both rectangular sets and the simplex fall into this category — for the previous example, take

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ -1 & 0 & 0 & \cdots & 0 \\ 0 & -1 & 0 & \cdots & 0 \\ \vdots & & & \ddots & \\ 0 & & \cdots & & -1 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

In general, when sets like these are bounded, the result is a polyhedron.

- Norm balls. If $\|\cdot\|$ is a valid norm on \mathbb{R}^N , then

$$\mathcal{B}_r = \{\mathbf{x} : \|\mathbf{x}\| \leq r\},$$

is a convex set.

- Ellipsoids. An ellipsoid is a set of the form

$$\mathcal{E} = \{\mathbf{x} : (\mathbf{x} - \mathbf{x}_0)^T \mathbf{P}^{-1} (\mathbf{x} - \mathbf{x}_0) \leq r\},$$

for a symmetric positive-definite matrix \mathbf{P} . Geometrically, the ellipsoid is centered at \mathbf{x}_0 , its axes are oriented with the eigenvectors of \mathbf{P} , and the relative widths along these axes are proportional to the eigenvalues of \mathbf{P} .

- A single point $\{\mathbf{x}\}$ is convex.

- The empty set is convex.

- The set

$$\{\mathbf{x} \in \mathbb{R}^2 : x_1^2 - 2x_1 - x_2 + 1 \leq 0\}$$

is convex. (Sketch it!)

- The set

$$\{\mathbf{x} \in \mathbb{R}^2 : x_1^2 - 2x_1 - x_2 + 1 \geq 0\}$$

is **not** convex.

- The set

$$\{\mathbf{x} \in \mathbb{R}^2 : x_1^2 - 2x_1 - x_2 + 1 = 0\}$$

is certainly not convex.

- Sets defined by linear equality constraints where only some of the constraints have to hold are in general not convex. For example

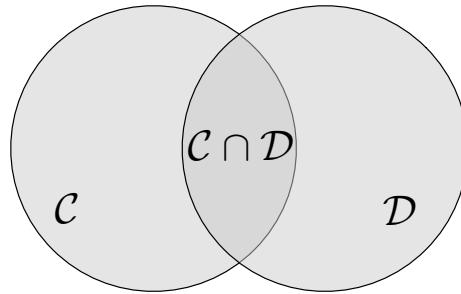
$$\{\mathbf{x} \in \mathbb{R}^2 : x_1 - x_2 \leq -1 \text{ and } x_1 + x_2 \leq -1\}$$

is convex, while

$$\{\mathbf{x} \in \mathbb{R}^2 : x_1 - x_2 \leq -1 \text{ or } x_1 + x_2 \leq -1\}$$

is not convex.

- It also follows easily from the definition of convexity that the **intersection** of any two convex sets is convex. However, the **union** of two convex sets is not necessarily convex. Example:



Cones

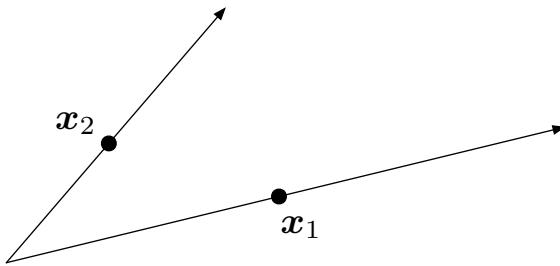
A **cone** is a set \mathcal{C} such that

$$\mathbf{x} \in \mathcal{C} \Rightarrow \theta \mathbf{x} \in \mathcal{C} \text{ for all } \theta \geq 0.$$

Convex cones are sets which are both convex and a cone. \mathcal{C} is a convex cone if

$$\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C} \Rightarrow \theta_1 \mathbf{x}_1 + \theta_2 \mathbf{x}_2 \in \mathcal{C} \text{ for all } \theta_1, \theta_2 \geq 0.$$

Given an $\mathbf{x}_1, \mathbf{x}_2$, the set of all linear combinations with positive weights makes a wedge. For practice, sketch the region below that consists of all such combinations of \mathbf{x}_1 and \mathbf{x}_2 :



We will mostly be interested in **proper cones**, which in addition to being convex, are closed, have a non-empty interior² (“solid”), and do not contain entire lines (“pointed”).

Examples of convex cones

Non-negative orthant. The set of non-negative vectors,

$$\mathbb{R}_+^N = \{\mathbf{x} \in \mathbb{R}^N : x_n \geq 0, \text{ for } n = 1, \dots, N\},$$

is a proper cone.

²See Technical Details for precise definition.

Positive semi-definite cone. The set of $N \times N$ symmetric matrices with non-negative eigenvalues, \mathcal{S}_+^N , is a proper cone.

Non-negative polynomials. Vectors of coefficients of non-negative polynomials on $[0, 1]$,

$$\{\mathbf{x} \in \mathbb{R}^N : x_1 + x_2 t + x_3 t^2 + \cdots + x_N t^{N-1} \geq 0 \text{ for all } 0 \leq t \leq 1\},$$

form a proper cone. Notice that it is not necessary that all the $x_n \geq 0$; for example $t - t^2$ ($x_1 = 0, x_2 = 1, x_3 = -1$) is non-negative on $[0, 1]$.

Norm cones. The subset of \mathbb{R}^{N+1} defined by

$$\{(\mathbf{x}, t), \mathbf{x} \in \mathbb{R}^N, t \in \mathbb{R} : \|\mathbf{x}\| \leq t\}$$

is a proper cone for any valid norm $\|\cdot\|$ and $t > 0$. We have seen this already for the Euclidean norm with $N = 2$, but this holds for arbitrary norms and dimensions.

Every proper cone \mathcal{K} defines a **partial ordering** or **generalized inequality**. We write

$$\mathbf{x} \preceq_{\mathcal{K}} \mathbf{y} \quad \text{when} \quad \mathbf{y} - \mathbf{x} \in \mathcal{K}.$$

For example, for vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$, we say

$$\mathbf{x} \preceq_{\mathbb{R}_+^N} \mathbf{y} \quad \text{when} \quad x_n \leq y_n \quad \text{for all } n = 1, \dots, N.$$

For symmetric matrices \mathbf{X}, \mathbf{Y} , we say

$$\mathbf{X} \preceq_{\mathcal{S}_+^N} \mathbf{Y} \quad \text{when} \quad \mathbf{Y} - \mathbf{X} \text{ has non-negative eigenvalues.}$$

We will typically just use \preceq when the context makes it clear. In fact, for \mathbb{R}_+^N we will just write $\mathbf{x} \leq \mathbf{y}$ to denote $\mathbf{x} \preceq_{\mathbb{R}_+^N} \mathbf{y}$, as we have already done several times.

Partial orderings obey share of the properties of the standard \leq on the real line. For example:

$$\mathbf{x} \preceq \mathbf{y}, \quad \mathbf{u} \preceq \mathbf{v} \quad \Rightarrow \quad \mathbf{x} + \mathbf{u} \preceq \mathbf{y} + \mathbf{v}.$$

But other properties do not hold; for example, it is not necessary that either $\mathbf{x} \preceq \mathbf{y}$ or $\mathbf{y} \preceq \mathbf{x}$. For an extensive list of properties of partial orderings (most of which will make perfect sense on sight) can be found in [BV04, Chapter 2.4].

Convex functions

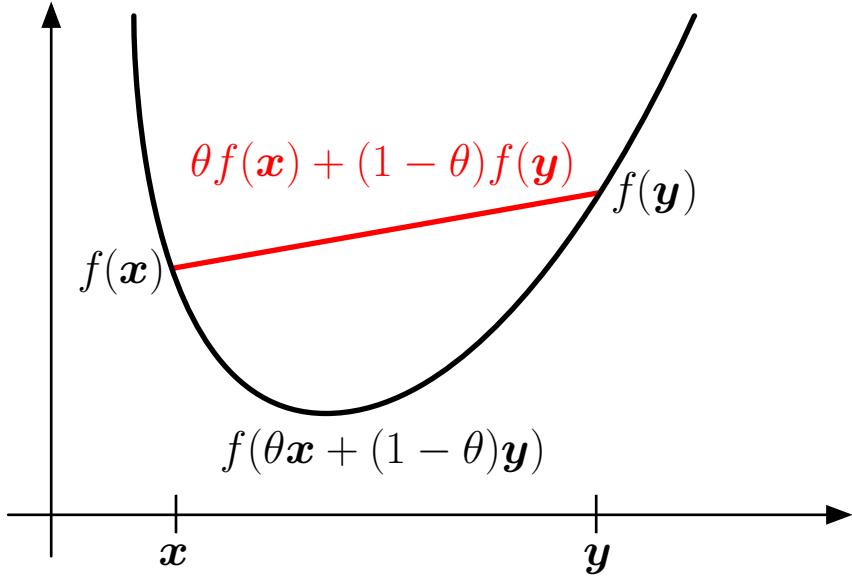
Convex *sets* are a fundamental concept in optimization. An equally important (and closely related) notion is that of convex *functions*.

To define this rigorously, we must sometimes be specific about the subset of \mathbb{R}^N where a function can be applied. Specifically, the **domain** $\text{dom } f$ of a function $f : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is the subset of \mathbb{R}^N where f is well-defined. We then say that a function f is **convex** if $\text{dom } f$ is a convex set, and

$$f(\theta \mathbf{x} + (1 - \theta) \mathbf{y}) \leq \theta f(\mathbf{x}) + (1 - \theta) f(\mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in \text{dom } f$ and $0 \leq \theta \leq 1$.

This inequality is easier to interpret with a picture. The left-hand side of the inequality above is simply the function f evaluated along a line segment between \mathbf{x} and \mathbf{y} . The right-hand side represents a straight line segment between $f(\mathbf{x})$ and $f(\mathbf{y})$ as we move along this line segment, which for a convex function must lie above f .



We say that f is **strictly convex** if $\text{dom } f$ is convex and

$$f(\theta\mathbf{x} + (1 - \theta)\mathbf{y}) < \theta f(\mathbf{x}) + (1 - \theta)f(\mathbf{y})$$

for all $\mathbf{x} \neq \mathbf{y} \in \text{dom } f$ and $0 < \theta < 1$.

Note also that we say that a function is f is **concave** if $-f$ is convex, and similarly for strictly concave functions. We are mostly interested in convex functions, but this is only because we are mostly restricting our attention to *minimization* problems. We justified this because any maximization problem can be converted to a minimization one by multiplying the objective function by -1 . Everything that we say about minimizing convex functions also applies maximizing concave ones.

We make a special note here that *affine* functions of the form

$$f(\mathbf{x}) = \langle \mathbf{x}, \mathbf{a} \rangle + b,$$

are both convex and concave (but neither strictly convex nor strictly concave). This is the only kind of function that has this property. (Why?)

Note that in the definition above, the domain matters. For example,

$$f(x) = x^3$$

is convex if $\text{dom } f = \mathbb{R}_+ = [0, \infty]$ but not if $\text{dom } f = \mathbb{R}$.

It will also sometimes be useful to consider the **extension** of f from $\text{dom } f$ to all of \mathbb{R}^N , defined as

$$\tilde{f}(\mathbf{x}) = f(\mathbf{x}), \quad \mathbf{x} \in \text{dom } f, \quad \tilde{f}(\mathbf{x}) = +\infty, \quad \mathbf{x} \notin \text{dom } f.$$

If f is convex on $\text{dom } f$, then its extension is also convex on \mathbb{R}^N .

It is an easily proven fact that if f is convex, then it is continuous on the interior of its domain. It can be the case, however that f is discontinuous at the boundary of its domain, as shown in our next example.

A special, but weird, convex function is the **indicator function** of a convex set \mathcal{C} , defined as

$$\iota_{\mathcal{C}}(\mathbf{x}) = \begin{cases} 0, & \mathbf{x} \in \mathcal{C} \\ \infty, & \mathbf{x} \notin \mathcal{C}. \end{cases}$$

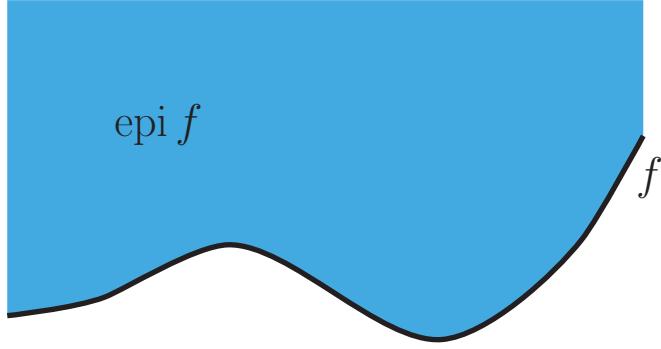
This function is certainly not continuous, but it does meet the definition of convexity. Indicator functions will appear when we start to talk about constrained optimization, as minimizing $f(\mathbf{x})$ over \mathcal{C} is the same as minimizing $f(\mathbf{x}) + \iota(\mathbf{x})$ without constraints.

The epigraph

A useful notion that we will encounter later in the course is that of the **epigraph** of a function. The epigraph of a function $f : \mathbb{R}^N \rightarrow \mathbb{R}$

is the subset of \mathbb{R}^{N+1} created by filling in the space above f :

$$\text{epi } f = \left\{ \begin{bmatrix} \mathbf{x} \\ t \end{bmatrix} \in \mathbb{R}^{N+1} : \mathbf{x} \in \text{dom } f, \quad f(\mathbf{x}) \leq t \right\}.$$



It is not hard to show that f is convex if and only if $\text{epi } f$ is a convex set. This connection should help to illustrate how even though the definitions of a convex set and convex function might initially appear quite different, they actually follow quite naturally from each other.

Examples of convex functions

Here are some standard examples for functions on \mathbb{R} :

- $f(x) = x^2$ is (strictly) convex.
- affine functions $f(x) = ax + b$ are both convex and concave for $a, b \in \mathbb{R}$.
- exponentials $f(x) = e^{ax}$ are convex for all $a \in \mathbb{R}$.
- powers x^α are:
 - convex on \mathbb{R}_+ for $\alpha \geq 1$,
 - concave for $0 \leq \alpha \leq 1$,
 - convex for $\alpha \leq 0$.
- $|x|^\alpha$ is convex on all of \mathbb{R} for $\alpha \geq 1$.
- logarithms: $\log x$ is concave on $\mathbb{R}_{++} := \{x \in \mathbb{R} : x > 0\}$.
- the entropy function $-x \log x$ is concave on \mathbb{R}_{++} .

Here are some standard examples for functions on \mathbb{R}^N :

- affine functions $f(\mathbf{x}) = \langle \mathbf{x}, \mathbf{a} \rangle + b$ are both convex and concave on all of \mathbb{R}^N .
- any valid norm $f(\mathbf{x}) = \|\mathbf{x}\|$ is convex on all of \mathbb{R}^N .
- if $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$ are both convex, then the sum $f_1(\mathbf{x}) + f_2(\mathbf{x})$ is also convex. Note, however, that it is not necessarily true that $f_1(\mathbf{x}) - f_2(\mathbf{x})$ is convex.

A useful tool for showing that a function $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is convex is the fact that f is convex if and only if the function $g_{\mathbf{v}} : \mathbb{R} \rightarrow \mathbb{R}$,

$$g_{\mathbf{v}}(t) = f(\mathbf{x} + t\mathbf{v}), \quad \text{dom } g = \{t : \mathbf{x} + t\mathbf{v} \in \text{dom } f\}$$

is convex for every $\mathbf{x} \in \text{dom } f$, $\mathbf{v} \in \mathbb{R}^N$.

Example:

Let $f(\mathbf{X}) = -\log \det \mathbf{X}$ with $\text{dom } f = \mathcal{S}_{++}^N$, where \mathcal{S}_{++}^N denotes the set of symmetric and (strictly) positive definite matrices. For any $\mathbf{X} \in \mathcal{S}_{++}^N$, we know that

$$\mathbf{X} = \mathbf{U} \boldsymbol{\Lambda} \mathbf{U}^T,$$

for some diagonal, positive $\boldsymbol{\Lambda}$, so we can define

$$\mathbf{X}^{1/2} = \mathbf{U} \boldsymbol{\Lambda}^{1/2} \mathbf{U}^T, \quad \text{and} \quad \mathbf{X}^{-1/2} = \mathbf{U} \boldsymbol{\Lambda}^{-1/2} \mathbf{U}^T.$$

Now consider any $\mathbf{V} \in \mathcal{S}^N$ and t such that $\mathbf{X} + t\mathbf{V} \in \mathcal{S}_{++}^N$:

$$\begin{aligned} g_{\mathbf{V}}(t) &= -\log \det(\mathbf{X} + t\mathbf{V}) \\ &= -\log \det(\mathbf{X}^{1/2} (\mathbf{I} + t\mathbf{X}^{-1/2} \mathbf{V} \mathbf{X}^{-1/2}) \mathbf{X}^{1/2}) \\ &= -\log \det \mathbf{X} - \log \det(\mathbf{I} + t\mathbf{X}^{-1/2} \mathbf{V} \mathbf{X}^{-1/2}) \\ &= -\log \det \mathbf{X} - \sum_{n=1}^N \log(1 + \sigma_i t), \end{aligned}$$

where the σ_i are the eigenvalues of $\mathbf{X}^{-1/2} \mathbf{V} \mathbf{X}^{-1/2}$. The function $-\log(1 + \sigma_i t)$ is convex, so the above is a sum of convex functions, which is convex.

Operations that preserve convexity

There are a number of useful operations that we can perform on a convex function while preserving convexity. Some examples include:

- **Positive weighted sum:** A **positive** weighted sum of convex functions is also convex, i.e., if f_1, \dots, f_m are convex and $w_1, \dots, w_m \geq 0$, then $w_1 f_1 + \dots + w_m f_m$ is also convex.
- **Composition with an affine function:** If $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is convex, then $g : \mathbb{R}^D \rightarrow \mathbb{R}$ defined by

$$g(\mathbf{x}) = f(\mathbf{A}\mathbf{x} + \mathbf{b}),$$

where $\mathbf{A} \in \mathbb{R}^{N \times D}$ and $\mathbf{b} \in \mathbb{R}^N$, is convex.

- **Composition with scalar functions:** Consider the function $f(\mathbf{x}) = h(g(\mathbf{x}))$, where $g : \mathbb{R}^N \rightarrow \mathbb{R}$ and $h : \mathbb{R} \rightarrow \mathbb{R}$.
 - f is convex if g is convex and h is convex and non-decreasing.
Example: $e^{g(\mathbf{x})}$ is convex if g is convex.
 - f is convex if g is concave and h is convex and non-increasing.
Example: $\frac{1}{g(\mathbf{x})}$ is convex if g is concave and positive.
- **Max of convex functions:** If f_1 and f_2 are convex, then $f(\mathbf{x}) = \max(f_1(\mathbf{x}), f_2(\mathbf{x}))$ is convex.

Technical Details: Basic topology in \mathbb{R}^N

Here we provide a brief review of basic topological concepts in \mathbb{R}^N . Our discussion will take place using the standard Euclidean distance measure (i.e., ℓ_2 norm), but all of these definitions can be generalized to other metrics. An excellent source for this material is [Rud76].

A recurring theme in this course relates to the convergence of an iterative algorithm. We say that a sequence of vectors $\{\mathbf{x}_k, k = 1, 2, \dots\}$ **converges** to $\hat{\mathbf{x}}$ if

$$\|\mathbf{x}_k - \hat{\mathbf{x}}\|_2 \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

More precisely, this means that for every $\epsilon > 0$, there exists an n_ϵ such that

$$\|\mathbf{x}_k - \hat{\mathbf{x}}\|_2 \leq \epsilon \quad \text{for all } k \geq n_\epsilon.$$

It is easy to show that a sequence of vectors converge if and only if their individual components converge point-by-point.

A set \mathcal{X} is **open** if we can draw a small ball around every point in \mathcal{X} which is also entirely contained in \mathcal{X} . More precisely, let $\mathcal{B}(\mathbf{x}, \epsilon)$ be the set of all points within ϵ of \mathbf{x} :

$$\mathcal{B}(\mathbf{x}, \epsilon) = \{\mathbf{y} \in \mathbb{R}^N : \|\mathbf{x} - \mathbf{y}\|_2 \leq \epsilon\}.$$

Then \mathcal{X} is open if for every $\mathbf{x} \in \mathcal{X}$, there exists an $\epsilon_x > 0$ such that $\mathcal{B}(\mathbf{x}, \epsilon_x) \subset \mathcal{X}$. The standard example here is open intervals of the real line, e.g. $(0, 1)$.

There are many ways to define **closed** sets. The easiest is that a set \mathcal{X} is closed if its complement is open. A more illuminating (and equivalent) definition is that \mathcal{X} is closed if it contains all of its limit

points. A vector $\widehat{\mathbf{x}}$ is a **limit point** of \mathcal{X} if there exists a sequence of vectors $\{\mathbf{x}_k\} \subset \mathcal{X}$ that converge to $\widehat{\mathbf{x}}$.

The **closure** of a general set \mathcal{X} , denoted $\text{cl}(\mathcal{X})$, is the set of all limit points of \mathcal{X} . Note that every $\mathbf{x} \in \mathcal{X}$ is trivially a limit point (take the sequence $\mathbf{x}_k = \mathbf{x}$), so $\mathcal{X} \subset \text{cl}(\mathcal{X})$. By construction, $\text{cl}(\mathcal{X})$ is the smallest closed set that contains \mathcal{X} .

Related to the definition of open and closed sets are the technical definitions of boundary and interior. The **interior** of a set \mathcal{X} is the collection of points around which we can place a ball of finite width which remains in the set:

$$\text{int}(\mathcal{X}) = \{\mathbf{x} \in \mathcal{X} : \exists \epsilon > 0 \text{ such that } \mathcal{B}(\mathbf{x}, \epsilon) \subset \mathcal{X}\}.$$

The **boundary** of \mathcal{X} is the set of points in $\text{cl}(\mathcal{X})$ that are not in the interior:

$$\text{bd}(\mathcal{X}) = \text{cl}(\mathcal{X}) \setminus \text{int}(\mathcal{X}).$$

Another (equivalent) way of defining this is the set of points that are in both the closure of \mathcal{X} and the closure of its complement \mathcal{X}^c . Note that if the set is not closed, there may be boundary points that are not in the set itself.

The set \mathcal{X} is **bounded** if we can find a uniform upper bound on the distance between two points it contains; this upper bound is commonly referred to as the **diameter** of the set:

$$\text{diam } \mathcal{X} = \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|_2.$$

The set $\mathcal{X} \subset \mathbb{R}^N$ is **compact** if it is closed and bounded. A key fact about compact sets is that every sequence has a convergent subsequence — this is known as the Bolzano-Weierstrauss theorem.

References

- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [Rud76] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 1976.