Week 2

Continuity (continuation)

Definition (Uniform continuity)**.** Let (X, d_X) and (Y, d_Y) be metric spaces and $f : X \to Y$. We say that f is *uniformly continuous* on X if

$$\forall \varepsilon > 0\ \exists \delta > 0 \text{ s.t. } \big(d_X(x, x') < \delta \implies d_Y\big(f(x), f(x')\big) < \varepsilon\big) \quad \text{for all } x, x' \in X.$$

Equivalently: the same $\delta = \delta(\varepsilon)$ works *simultaneously* for all pairs x, x' in X.

Remark (Negation (quantifiers made explicit))**.** The negation of uniform continuity is:

$$\exists\, \varepsilon_0 > 0 \text{ such that } \forall\, \delta > 0\ \exists\, x, x' \in X \text{ with } d_X(x, x') < \delta \text{ and } d_Y\big(f(x), f(x')\big) \geq \varepsilon_0.$$

Remark (Equivalent "constructive" failure)**.** Equivalently, f fails to be uniformly continuous iff there exist $\varepsilon_0 > 0$ and two sequences $\{x_n\}, \{x'_n\} \subset X$ such that

$$d_X(x_n, x'_n) \xrightarrow[n\to\infty]{} 0 \quad \text{but} \quad d_Y\big(f(x_n), f(x'_n)\big) \geq \varepsilon_0 \text{ for all } n.$$

This rephrasing is often handier in proofs.

Intuition. Uniform continuity asks for a δ that works *globally* (same δ for the whole domain) once ε is fixed. In ordinary continuity at a point x^*, the admissible δ may depend on both ε *and* the base point x^*. The negation highlights that, if uniform continuity fails, you can zoom in ($\delta \downarrow 0$) and still find pairs x, x' arbitrarily close whose images stay separated by at least some fixed ε_0.

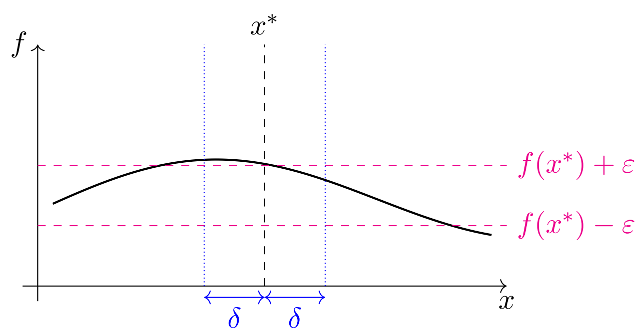

Figure 1: Uniform continuity: one $\delta(\varepsilon)$ works *everywhere* in the domain.

Remark (On the dependence of x, x' in the negation)**.** In the negation, the "bad" points x, x' may (and typically do) depend on ε_0 and on the chosen δ. There is no single pair that witnesses failure for all δ; instead, you can find a violating pair *for every* $\delta > 0$.

Remark (Two sequence facts used repeatedly)**.** Let $\{a_n\} \subset \mathbb{R}$.

(a) If $a_n \to a^*$, then $\liminf_{n\to\infty} a_n = \limsup_{n\to\infty} a_n = a^*$.

(b) If $\liminf_{n\to\infty} a_n = \limsup_{n\to\infty} a_n \in \mathbb{R}$, then $\{a_n\}$ converges and its limit equals that common value.

Theorem 1 (Intermediate Value Theorem)**.** *Let $f : [\underline{x}, \overline{x}] \to \mathbb{R}$ be continuous. For any y^* between $f(\underline{x})$ and $f(\overline{x})$ (i.e., either $f(\underline{x}) \le y^* \le f(\overline{x})$ or $f(\overline{x}) \le y^* \le f(\underline{x})$), there exists $x^* \in [\underline{x}, \overline{x}]$ such that $f(x^*) = y^*$.*

Intuition. A continuous graph on a closed interval cannot "jump over" a horizontal level y^*: if the endpoint values lie on opposite sides of y^* (or one equals it), the graph must cross the line $y = y^*$ somewhere in between.

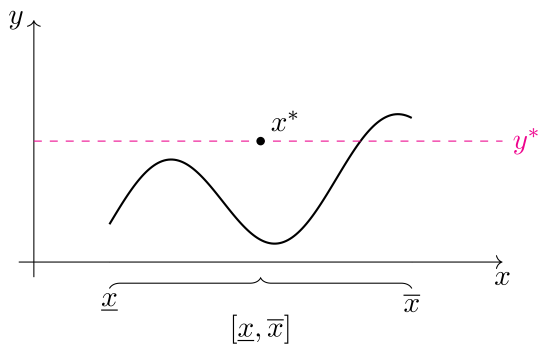

Figure 2: IVT: a continuous function on $[\underline{x}, \overline{x}]$ crosses every intermediate level.

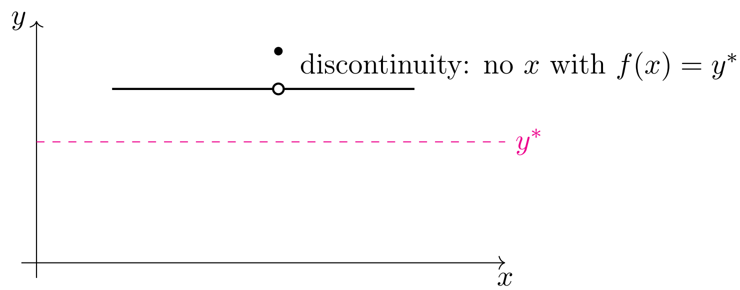

Figure 3: If f is not continuous, an intermediate value y^* may fail to be attained.

Remark (Restriction trick for IVT)**.** Let $f : \mathbb{R} \to \mathbb{R}$ be continuous and let $[\underline{x}, \overline{x}] \subset \mathbb{R}$. Consider the restriction

$$g := f\big|_{[\underline{x}, \overline{x}]} : [\underline{x}, \overline{x}] \to \mathbb{R}, \qquad g(x) = f(x).$$

If $y^* \in \mathbb{R}$ is such that either $f(\underline{x}) \le y^* \le f(\overline{x})$ or $f(\overline{x}) \le y^* \le f(\underline{x})$, then, since g is continuous on the closed interval $[\underline{x}, \overline{x}]$, the Intermediate Value Theorem applied to g yields $\exists\, x^* \in [\underline{x}, \overline{x}]$ with $f(x^*) = y^*$.

Intuition. You do not need any global property of f beyond continuity: restricting a continuous f to $[\underline{x}, \overline{x}]$ keeps it continuous there, so IVT applies to the *restricted* function g.

Theorem 2 (Extreme Value (Weierstrass))**.** *Let (X, d) be a compact metric space and let $f : X \to \mathbb{R}$ be continuous. Then f attains a maximum and a minimum on X, i.e., there exist $\overline{x}, \underline{x} \in X$ such that*

$$f(\underline{x}) \le f(x) \le f(\overline{x}) \quad \textit{for all } x \in X.$$

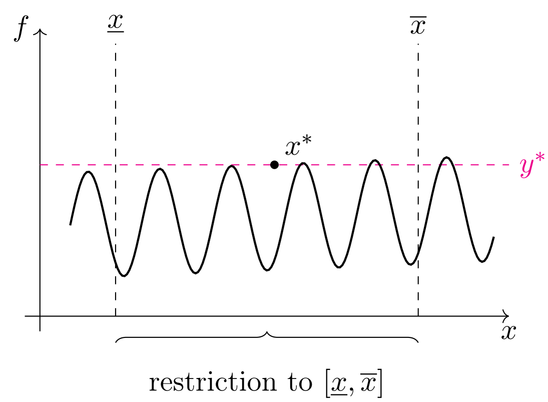

Figure 4: Apply IVT to $g = f|_{[\underline{x},\overline{x}]}$.

Definition (Attainment of extrema)**.** Fix $f : X \to \mathbb{R}$.

- f *attains a maximum* if $\exists\, \overline{x} \in X$ such that $f(\overline{x}) \geq f(x)$ for all $x \in X$.
- f *attains a minimum* if $\exists\, \underline{x} \in X$ such that $f(\underline{x}) \leq f(x)$ for all $x \in X$.

Intuition. Compactness rules out "escaping to infinity" and "missing boundary points." Continuity prevents jumps. Together they force the sup and inf to be *achieved*.

Example (Identity map and the role of the domain)**.** Let $f(x) = x$ (identity) and $X \subseteq \mathbb{R}$.

(a) If $X = \mathbb{R}$, then f has no maximum (unbounded above).

(b) If $X = [0, \frac{1}{2})$, then f has no maximum: $\sup_X f = \frac{1}{2}$ but it is not attained because $1/2 \notin X$ (domain not closed).

(c) If X is a *finite* disjoint union of closed intervals,

$$X = \bigcup_{k=1}^{K} [x_k, y_k], \qquad x_1 < y_1 < x_2 < y_2 < \cdots < x_K < y_K,$$

then X is compact (finite union of compact sets), hence by Weierstrass f attains both extrema; in fact, $\max_X f = y_K$ and $\min_X f = x_1$.

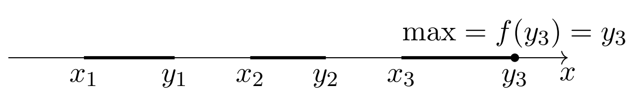

Figure 5: Finite union of closed intervals $\Rightarrow$ compact set $\Rightarrow$ extrema attained by $f(x) = x$.

Example (Compact domain without continuity: no maximum)**.** Let $s > 0$ and define $f : [0, 1] \to \mathbb{R}$ by

$$f(x) = \begin{cases} s\,x, & x \in [0, \frac{1}{2}), \\ 0, & x \in [\frac{1}{2}, 1]. \end{cases}$$

Then f is *not* continuous at $x = \frac{1}{2}$, and f has *no* maximum on $[0, 1]$: the supremum is $s/2$, but it is not attained since the left branch does not include $x = \frac{1}{2}$ and the right branch equals 0.

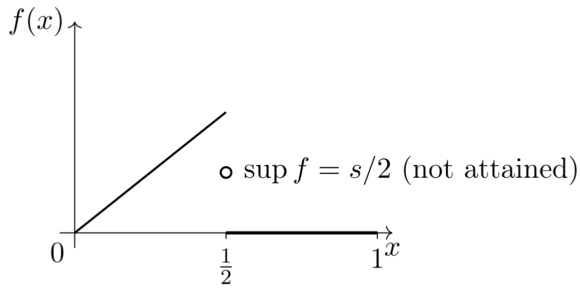

Figure 6: Discontinuity at $x = \frac{1}{2}$ breaks EVT: compact domain alone is not enough.

A useful version for constrained optimization

Let (X, d) be a metric space, $C \subseteq X$ a constraint set, and $f : X \to \mathbb{R}$.

$$\textbf{Problem:} \quad \max_{x \in C} f(x).$$

Definition (Solution / maximizer on C)**.** A *solution* (or *maximizer*) is any $x \in C$ such that $f(x) \geq f(y)$ for all $y \in C$. The set of all solutions is the *arg max*,

$$\operatorname{argmax}_{x \in C} f(x) \;\coloneqq\; \{\, x \in C : \; f(x) \geq f(y) \text{ for all } y \in C \,\}.$$

Remark. In Example with $f(x) = x$:

- If $(X, d) = (\mathbb{R}, d_u)$ and $C = \mathbb{R}$, then $\operatorname{argmax}_{x \in C} f(x) = \varnothing$.
- If $(X, d) = ([0, 1], d_u)$ and $C = [0, 1]$, then $\operatorname{argmax}_{x \in C} f(x) = \{1\}$.

Proposition. *If $f : X \to \mathbb{R}$ is continuous and $C \subseteq X$ is compact, then* $\operatorname{argmax}_{x \in C} f(x)$ *is nonempty and compact.*

Proof. Needs double check with page 4 of class notes week two.

Step 1 (restriction). Let $g : C \to \mathbb{R}$ be the restriction $g(x) = f(x)$ for $x \in C$. Then g is continuous on C.

Step 2 (existence). By Weierstrass, the continuous image $g(C) \subseteq \mathbb{R}$ is compact, hence closed and bounded, so it contains its maximum. Let $y^* = \max g(C) \in g(C)$. Then the *arg max* can be written as a level set:

$$\operatorname{argmax}_{x \in C} f(x) = \operatorname{argmax}_{x \in C} g(x) = g^{-1}(\{y^*\}),$$

which is nonempty because $y^* \in g(C)$.

Step 3 (compactness). Since singletons $\{y^*\}$ are closed in $\mathbb{R}$ and g is continuous, $g^{-1}(\{y^*\})$ is closed in C. A closed subset of a compact set is compact; hence $\operatorname{argmax}_{x \in C} f(x)$ is compact.

Alternative (sequences). Let $(x_n) \subseteq \operatorname{argmax}_{x \in C} g(x)$. Any limit point $x^* \in C$ of (x_n) satisfies, by continuity of g, $g(x_n) = y^*$ for all $n \implies g(x_n) \to g(x^*) = y^*$, hence $x^* \in \operatorname{argmax}_{x \in C} g(x)$. Therefore the arg max is closed in C and thus compact. $\square$

Lemma (Continuous image of a compact set). *Let (X, d) be compact and consider $(\mathbb{R}, d_u)$. If $f : X \to \mathbb{R}$ is continuous, then $f(X) \subseteq \mathbb{R}$ is compact.*

Lemma (Singletons are closed). *In $(\mathbb{R}, d_u)$, every singleton $\{y^*\}$ is closed.*

Lemma (Closed-set characterization of continuity). *Let (X, d_X) and (Y, d_Y) be metric spaces and $f : X \to Y$. Then f is continuous iff for every closed set $F \subseteq Y$ the preimage $f^{-1}(F) \subseteq X$ is closed.*

Function spaces

Definition (Real-valued function space). Let (X, d) be a metric space and $(\mathbb{R}, d_u)$ the real line with its usual metric. We denote by

$$\mathcal{F}(X, \mathbb{R}) \;\coloneqq\; \{\, f : X \to \mathbb{R} \,\}$$

the set of all real-valued functions on X.

Remark (Prominent examples).

(a) $f : \mathbb{R} \to \mathbb{R}$ given by $f(x) = x^2$.

(b) $f : \mathbb{R}^2 \to \mathbb{R}$ given by $f(x_1, x_2) = 10\, x_1\, x_2^2$.

(c) (Operator on a function space) $T : \mathcal{F}(X, \mathbb{R}) \to \mathcal{F}(X, \mathbb{R})$ defined by

$$(Tg)(x) \;=\; \tfrac{1}{2}\, g(x) \qquad \text{for all } x \in X.$$

Thus $T \in \mathcal{F}(\mathcal{F}(X, \mathbb{R}),\, \mathcal{F}(X, \mathbb{R}))$.

(d) (Functional) Let

$$\mathcal{I} \;\coloneqq\; \big\{\, g \in \mathcal{F}(\mathbb{R}, \mathbb{R}) : g \text{ is (Lebesgue/Riemann) integrable on } \mathbb{R} \,\big\}.$$

Define $I : \mathcal{I} \to \mathbb{R}$ by

$$I(g) \;=\; \int_{\mathbb{R}} g(x)\, dx.$$

This is a map $I \in \mathcal{F}(\mathcal{I}, \mathbb{R})$; it is *well-defined only when the integral is finite* (not $\pm\infty$).

Intuition. Items (c) and (d) emphasize two common kinds of maps involving function spaces: operators $T : \mathcal{F}(X, \mathbb{R}) \to \mathcal{F}(X, \mathbb{R})$ that *return* a new function, and functionals $I : \mathcal{F}(X, \mathbb{R}) \to \mathbb{R}$ that *return* a number from a function.

Definition (Sequence of functions). A *sequence of functions* on X is a family $\{f_m\}_{m \in \mathbb{N}} \subseteq \mathcal{F}(X, \mathbb{R})$.

Remark (What do we mean by “convergence of functions”?). Up to now we often fixed a function f and studied numeric sequences like $\{f(x_n)\}_{n \in \mathbb{N}}$ when $x_n \to x^*$; then $f(x_n) \to f(x^*)$ if f is continuous at x^*. *That is not the same question* as asking whether the *functions* f_m themselves converge to some new function f on X. In the sequel we will work inside $\mathcal{F}(X, \mathbb{R})$ and make precise notions of convergence (e.g., pointwise vs. uniform).

Pointwise convergence

Definition (Pointwise convergence)**.** Let $\{f_n\}_{n\in\mathbb{N}} \subseteq \mathcal{F}(X,\mathbb{R})$. We say that f_n *converges pointwise* to a function $f \in \mathcal{F}(X,\mathbb{R})$ if

$$\forall x \in X: \qquad \lim_{n\to\infty} f_n(x) = f(x).$$

Equivalently: for each fixed $x \in X$, the numeric sequence $\{f_n(x)\}_{n\in\mathbb{N}}$ converges in $\mathbb{R}$, and we *define* $f(x)$ to be that limit. Formally:

$$\forall x \in X\ \forall \varepsilon > 0\ \exists N = \mathbf{N}(\mathbf{x}, \varepsilon) \in \mathbb{N} \text{ such that } \forall n \geq N: \ |f_n(x) - f(x)| < \varepsilon.$$

Remark. Pointwise convergence is checked *point-by-point*. It does not control how fast the convergence occurs across different x's, and it does not preserve continuity in general.

Example (A continuous-to-discontinuous pointwise limit)**.** Let $X = [0,1]$ and $f_n : X \to \mathbb{R}$ be $f_n(x) = x^n$. Each f_n is continuous (even differentiable). For any fixed $x \in [0,1)$ we have $x^n \to 0$, while $1^n \to 1$. Hence

$$f_n \xrightarrow[\text{pointwise}]{} f, \qquad f(x) = \begin{cases} 0, & x \in [0,1), \\ 1, & x = 1, \end{cases}$$

and f is discontinuous at $x = 1$ (thus non-differentiable there).

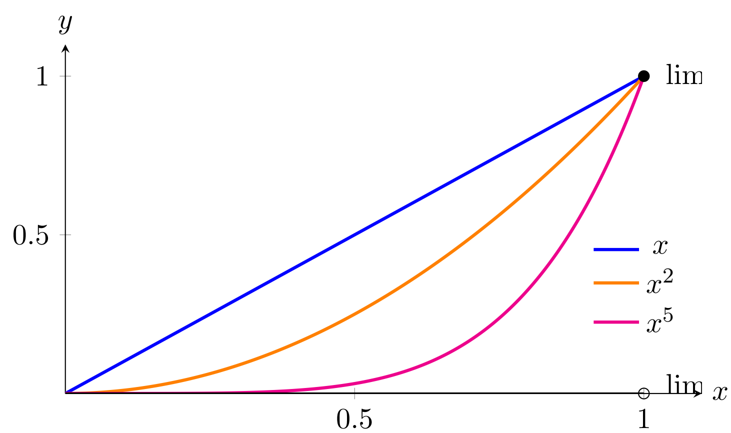

Figure 7: Pointwise limit of $f_n(x) = x^n$ on $[0,1]$: a discontinuous function.

Remark (What if the domain is $\mathbb{R}_+$?)**.** If we take $X = \mathbb{R}_+$ and $f_n(x) = x^n$, then for $x > 1$ we have $x^n \to +\infty$ (no finite limit in $\mathbb{R}$), for $x \in [0,1)$ we have $x^n \to 0$, and at $x = 1$ we have $1^n \to 1$. Therefore $\{f_n\}$ is *not* pointwise convergent as a sequence in $\mathcal{F}(\mathbb{R}_+,\mathbb{R})$. (Allowing extended reals would give a limit taking value $+\infty$ on $(1,\infty)$, which lies outside $\mathbb{R}$.)

Intuition. This example shows that a sequence of continuous (even smooth) functions may converge *pointwise* to a function that is discontinuous and non-differentiable. Pointwise convergence alone is too weak to preserve regularity properties—this motivates stronger notions (e.g., uniform convergence).

More on pointwise convergence: examples and a warning

Example (Bounded functions converging to an unbounded function)**.** Let $X = \mathbb{R}_+$ and define $f_n : X \to \mathbb{R}$ by

$$f_n(x) = \begin{cases} x, & x \le n, \\ 0, & x > n. \end{cases}$$

Each f_n is bounded (indeed $\sup_{x\in\mathbb{R}_+} |f_n(x)| \le n$). Fix $x \in \mathbb{R}_+$. Then $(f_1(x), f_2(x), \dots) = \big(0, 0, \dots, 0, \underbrace{x, x, x, \dots}_{\text{from } n \ge \lceil x \rceil}\big)$, so $\lim_{n\to\infty} f_n(x) = x$. Hence $f_n \xrightarrow{\text{pointwise}} f$ with $f(x) = x$ on $\mathbb{R}_+$, and the limit f is unbounded. Pointwise limits need not inherit boundedness.

Example (Pointwise convergence does not preserve limits/continuity)**.** Recall the sequential characterization: for $g : X \to \mathbb{R}$ and $x^* \in X$, $\lim_{x\to x^*} g(x) = y^*$ iff for every sequence (x_m) with $x_m \neq x^*$ and $x_m \to x^*$ we have $g(x_m) \to y^*$.

Let $X = [0, 1]$ and $f_n(x) = x^n$. We know $f_n \xrightarrow{\text{pointwise}} f$ where

$$f(x) = \begin{cases} 0, & x \in [0, 1), \\ 1, & x = 1. \end{cases}$$

Fix the sequence $x_m = 1 - \frac{1}{m} \uparrow 1$ with $x_m \neq 1$. For each *fixed* n,

$$\lim_{m\to\infty} f_n(x_m) = \lim_{m\to\infty} (1 - \tfrac{1}{m})^n = 1.$$

Thus all f_n have left-limit 1 at $x^* = 1$. But along the same sequence,

$$\lim_{m\to\infty} f(x_m) = \lim_{m\to\infty} 0 = 0 \neq 1.$$

Therefore the limit function f does *not* preserve that limit at x^* (indeed, f is discontinuous at 1) even though each f_n is continuous. Pointwise convergence is too weak to preserve limits/continuity.

Uniform versus pointwise

Definition (Uniform convergence)**.** Let $\{f_n\}_{n\in\mathbb{N}} \subseteq \mathcal{F}(X, \mathbb{R})$ and $f \in \mathcal{F}(X, \mathbb{R})$. We say that $f_n \to f$ *uniformly on* X if

$$\forall \varepsilon > 0 \ \exists N \in \mathbb{N} \text{ s.t. } \forall n \ge N, \ \forall x \in X : \ |f_n(x) - f(x)| < \varepsilon.$$

or in another way (more helpful for exercises sometimes):

$$\forall \varepsilon > 0 \ \exists N = \mathbf{N}(\varepsilon) \in \mathbb{N} \text{ such that } \forall x \in X \ \forall n \ge N : \ |f_n(x) - f(x)| < \varepsilon.$$

In this case, call f the **uniform limit** of $(f_n : n \in \mathbb{N})$

Intuition: The important thing here is the order of the quantifiers. In **uniform** convergence the index can be chosen as $\mathbf{N} = \mathbf{N}(\varepsilon)$ (the **same** N works for **all** x). By contrast, in **pointwise** convergence one only has $\mathbf{N} = \mathbf{N}(\mathbf{x}, \varepsilon)$. Equivalently: uniform $\iff (\forall\varepsilon)(\exists N)(\forall x)(\forall n \ge N)$.

Proposition (Uniform $\Rightarrow$ pointwise)**.** *If $f_n \to f$ uniformly on X, then $f_n \to f$ pointwise on X.*

Proof. Given $x \in X$ and $\varepsilon > 0$, choose N such that for all $n \geq N$, $\sup_{y\in X}|f_n(y) - f(y)| < \varepsilon$. In particular $|f_n(x) - f(x)| < \varepsilon$ for all $n \geq N$, hence $f_n(x) \to f(x)$. □

Remark. To analyze a sequence (f_n):

(i) First, figure out to which f it converges *pointwise.*

(ii) Then, check whether the convergence is *uniform.*

Example (Back to Example : not uniform)**.** For $f_n(x) = x^n$ on $[0, 1]$ with pointwise limit $f = \mathbf{1}_{\{1\}}$ on the endpoint, we have

$$\sup_{x\in[0,1]} |f_n(x) - f(x)| = \max\Big\{ \sup_{x\in[0,1)} x^n,\ |1^n - 1|\Big\} = \sup_{x\in[0,1)} x^n = 1 \quad \text{for every } n.$$

(Indeed, $x^n \uparrow 1$ as $x \uparrow 1$.) In particular, the sup norm does not go to 0, so the convergence is *not* uniform. A quantitative lower bound is obtained by $x_n = (1 - \frac{1}{n})$, for which $x_n^n \to e^{-1} > 0$.

Remark (Negation)**.** Failure of uniform convergence means:

$$\exists\, \varepsilon_0 > 0 \text{ s.t. } \forall N \in \mathbb{N}\ \exists n \geq N,\ \exists x \in X \text{ with } |f_n(x) - f(x)| \geq \varepsilon_0.$$

Here the "bad" x may depend on n (and on N).

Proposition (Sup-norm criterion)**.** *Let* $\|g\|_{\infty,X} := \sup_{x\in X}|g(x)|$. *Then*

$$f_n \to f \textit{ uniformly on } X \quad \iff \quad \|f_n - f\|_{\infty,X} \xrightarrow[n\to\infty]{} 0.$$

Example ($f_n(x) = x^n$ on $[0, 1]$ is not uniform)**.** The pointwise limit is $f(x) = 0$ for $x \in [0, 1)$ and $f(1) = 1$. For $\varepsilon = \frac{1}{4}$ and any N, choose $n = N$ and $x = (\frac{1}{4})^{1/n} \in (0, 1)$. Then $f(x) = 0$ and

$$|f_n(x) - f(x)| = |x^n - 0| = \tfrac{1}{4} \geq \varepsilon.$$

By the negation, $f_n \not\to f$ uniformly on $[0, 1]$. Equivalently, $\|f_n - f\|_{\infty,[0,1]} = \sup_{x\in[0,1)} x^n = 1$ for all n.

Example ($f_n = \mathbf{1}_{[0,n]} \cdot \mathrm{id}$ on $\mathbb{R}_+$ is not uniform)**.** Recall $f_n(x) = x$ for $x \leq n$ and $f_n(x) = 0$ for $x > n$, so $f_n \to f$ pointwise with $f(x) = x$. Take $\varepsilon = 1$. For any N, pick $n = N$ and $x = 2N > n$. Then $f_n(x) = 0$ while $f(x) = 2N$, hence

$$|f_n(x) - f(x)| = 2N \geq \varepsilon.$$

Thus the convergence is not uniform on $\mathbb{R}_+$.

Example ($f_n(x) = x/n$)**.**

(a) On $X = [0, K]$ the limit is $f \equiv 0$ and

$$\|f_n - f\|_{\infty,[0,K]} = \sup_{x\in[0,K]} \frac{x}{n} = \frac{K}{n} \xrightarrow[n\to\infty]{} 0,$$

so $f_n \to 0$ uniformly on $[0, K]$. In ε–N form: given $\varepsilon > 0$, take $N > \frac{K}{\varepsilon}$; then for $n \geq N$ and all $x \in [0, K]$, $|f_n(x) - 0| = \frac{x}{n} \leq \frac{K}{n} < \varepsilon$.

(b) On $X = \mathbb{R}_+$ the convergence to 0 is *not* uniform. Indeed, fix $\varepsilon > 0$. For any N choose $n = N$ and $x = N\varepsilon$. Then $|f_n(x) - 0| = \frac{x}{n} = \varepsilon$, so the sup-norm never falls below ε.

Intuition. Uniformity requires that a single N work simultaneously for all x in the domain. In unbounded domains (such as $\mathbb{R}_+$), it is typical that we can "push" x toward the region where the error becomes large again, breaking uniformity.

Restriction and preservation properties

Proposition (Restriction preserves uniform convergence)**.** *Let $(f_n)_{n\in\mathbb{N}} \subset \mathcal{F}(X, \mathbb{R})$ converge uniformly to f on X. If $E \subseteq X$ and we denote by $f_n|_E$ and $f|_E$ the restrictions to E, then $(f_n|_E)$ converges uniformly to $f|_E$ on E.*

Proof. Done in a PS I think. □

Remark (Uniform continuity: what is preserved?)**.**

(a) **Heine–Cantor.** If X is compact and $f : X \to \mathbb{R}$ is continuous, then f is uniformly continuous.

(b) **Linear operations.** If $f, g : X \to \mathbb{R}$ are uniformly continuous and $c \in \mathbb{R}$, then $f \pm g$ and $c\,f$ are uniformly continuous.

(c) **Absolute value, max, min.** If f, g are uniformly continuous, then $|f|$, $\max\{f, g\}$, and $\min\{f, g\}$ are uniformly continuous.

(d) **Composition.** If $f : (X, d_X) \to (Y, d_Y)$ and $g : (Y, d_Y) \to (Z, d_Z)$ are uniformly continuous, then $g \circ f$ is uniformly continuous on X.

(e) **Product: not preserved in general.** Even if f and g are uniformly continuous, the product $f \cdot g$ may fail to be uniformly continuous on non-compact domains.

(f) **Division: not preserved in general.** Even if f and $g \neq 0$ are uniformly continuous, the product $f \cdot g$ may fail to be uniformly continuous on non-compact domains.

Example (Product counterexample on $\mathbb{R}_+$)**.** Let $g(x) = x$ on $\mathbb{R}_+$. Then g is 1-Lipschitz (hence uniformly continuous). But $f = g \cdot g = x^2$ is *not* uniformly continuous on $\mathbb{R}_+$: take $x_n = n$, $y_n = n + \frac{1}{2n}$; then $|x_n - y_n| = \frac{1}{2n} \to 0$ while $|f(x_n) - f(y_n)| = |y_n^2 - x_n^2| = (x_n + y_n)|y_n - x_n| \geq 2n \cdot \frac{1}{2n} = 1$.

Uniform limits of continuous functions

Proposition. *Let $(f_n)_{n\in\mathbb{N}} \subset \mathcal{F}(X, \mathbb{R})$ converge uniformly to f on X.*

(a) If for each n the function f_n is continuous at a point $x^* \in X$, then f is continuous at x^*.

(b) If each f_n is continuous on X, then f is continuous on X.

Proof. (a) Fix $\varepsilon > 0$. By uniform convergence choose N with $\|f_N - f\|_{\infty,X} < \varepsilon/3$. Since f_N is continuous at x^*, there exists $\delta > 0$ such that $d_X(x, x^*) < \delta$ implies $|f_N(x) - f_N(x^*)| < \varepsilon/3$. Then, for $d_X(x, x^*) < \delta$,

$$\begin{aligned} |f(x) - f(x^*)| &= |f(x) - f_N(x) + f_N(x) - f_N(x^*) + f_N(x^*) - f(x^*)| \\ &\leq |f(x) - f_N(x)| + |f_N(x) - f_N(x^*)| + |f_N(x^*) - f(x^*)| < \varepsilon. \end{aligned}$$

(b) Apply (a) at each $x^* \in X$. □

Uniform limits live in the bounded space and the sup metric

Definition (Bounded function space and sup metric)**.** Let

$$\mathcal{F}^B(X,\mathbb{R}) \;\coloneqq\; \{\, f \in \mathcal{F}(X,\mathbb{R}) : \; \|f\|_{\infty,X} \coloneqq \sup_{x\in X} |f(x)| < \infty \,\}.$$

On $\mathcal{F}^B(X,\mathbb{R})$ define the metric

$$d_\infty(f,g) \;\coloneqq\; \sup_{x\in X} |f(x) - g(x)| = \|f-g\|_{\infty,X}.$$

Proposition (Uniform limit of bounded functions is bounded)**.** *Let $(f_n)_{n\in\mathbb{N}} \subset \mathcal{F}^B(X,\mathbb{R})$ converge uniformly to $f \in \mathcal{F}(X,\mathbb{R})$. Then $f \in \mathcal{F}^B(X,\mathbb{R})$.*

Proof. Beyond the scope. □

Proposition (Uniform $\iff$ d_∞-convergence)**.** *Let $(f_n)_{n\in\mathbb{N}} \subset \mathcal{F}^B(X,\mathbb{R})$ and let $f \in \mathcal{F}^B(X,\mathbb{R})$. Then the following are equivalent:*

(a) $f_n \to f$ uniformly on X;

(b) $d_\infty(f_n, f) = \|f_n - f\|_{\infty,X} \to 0$ as $n \to \infty$.

Proof. Beyond the scope. □

Intuition. The metric d_∞ "measures" the worst error uniformly across the entire domain. Therefore, uniform convergence and convergence in d_∞ are the same thing; and if any (in fact, all) of the f_n are bounded, the uniform limit is also bounded.

Vector spaces and norms

Intuition. A vector space is an abstract set of objects equipped with two operations (vector addition and scalar multiplication) that obey precise rules. In the background there is a *field* of scalars (in this course we usually take $\mathbb{R}$, but any field works).

Definition (Vector space)**.** Let $\mathbb{F}$ be a field (typically $\mathbb{R}$). A *vector space over* $\mathbb{F}$ is a triple $(V, +, \cdot)$ where:

- V is a set (its elements are called *vectors*);
- $+$ is a binary operation $V \times V \to V$ (vector addition);
- $\cdot$ is an operation $\mathbb{F} \times V \to V$ (scalar multiplication, $(\alpha, v) \mapsto \alpha v$);

satisfying the following axioms for all $u, v, w \in V$ and all $\alpha, \beta \in \mathbb{F}$:

(1) **Addition is commutative:** $v + w = w + v$.

(2) **Addition is associative:** $u + (v + w) = (u + v) + w$.

(3) **Additive identity:** there exists $0 \in V$ (the *zero vector*) such that $0 + v = v + 0 = v$.

(4) **Additive inverses:** for each $v \in V$ there exists $w \in V$ (denoted $-v$) with $v + w = 0$.

(5) **Multiplicative identity:** $1 \cdot v = v$.

(6) **Compatibility with field multiplication:** $\alpha(\beta v) = (\alpha\beta)v$.

(7) **Distributivity over vector addition:** $\alpha(v + w) = \alpha v + \alpha w$.

(8) **Distributivity over scalar addition:** $(\alpha + \beta)v = \alpha v + \beta v$.

Example (A vector space of bounded functions)**.** Fix a set X. Let

$$\mathcal{F}^B(X, \mathbb{R}) = \{f : X \to \mathbb{R} : \ f \text{ is bounded}\}.$$

Define the operations pointwise:

$$(f + g)(x) = f(x) + g(x), \qquad (\alpha f)(x) = \alpha\, f(x) \quad (\alpha \in \mathbb{R}).$$

Let $0 \in \mathcal{F}^B(X, \mathbb{R})$ be the zero function $0(x) = 0$ and for $f \in \mathcal{F}^B(X, \mathbb{R})$ define its additive inverse by $(-f)(x) = -f(x)$. Then $(\mathcal{F}^B(X, \mathbb{R}), +, \cdot)$ is a vector space over $\mathbb{R}$.

Remark. A vector space axiomatizes *addition of vectors* and *scalar multiplication.* It does *not* prescribe a rule to multiply two vectors with each other. Any notion of "multiplying vectors" (dot product, cross product, matrix product, convolution, etc.) is *extra structure* that depends on the environment we are working in.

Definition (Dot product on $\mathbb{R}^I$)**.** For $I \in \mathbb{N}$, the *dot product* on $\mathbb{R}^I$ is

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle \ = \ \sum_{i=1}^{I} x_i\, y_i \in \mathbb{R}, \qquad \boldsymbol{x} = (x_1, \ldots, x_I), \ \boldsymbol{y} = (y_1, \ldots, y_I) \in \mathbb{R}^I.$$

Note that $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$ is a *scalar*, not a vector.

Lemma (Basic properties of the dot product)**.** *For all $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathbb{R}^I$ and all $\alpha \in \mathbb{R}$:*

(a) ***Nonnegativity and definiteness:*** $\langle \boldsymbol{x}, \boldsymbol{x} \rangle \geq 0$, *and* $\langle \boldsymbol{x}, \boldsymbol{x} \rangle = 0$ *iff* $\boldsymbol{x} = \mathbf{0}$.

(b) ***Symmetry:*** $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \langle \boldsymbol{y}, \boldsymbol{x} \rangle$.

(c) ***Additivity in the second entry:*** $\langle \boldsymbol{x}, \boldsymbol{y} + \boldsymbol{z} \rangle = \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{x}, \boldsymbol{z} \rangle$.

(d) ***Homogeneity in the second entry:*** $\langle \boldsymbol{x}, \alpha \boldsymbol{y} \rangle = \alpha \langle \boldsymbol{x}, \boldsymbol{y} \rangle$.

(e) Equivalently, linearity in the first entry also holds by symmetry:

$$\langle \boldsymbol{x} + \boldsymbol{y}, \boldsymbol{x} + \boldsymbol{y} \rangle = \langle \boldsymbol{x}, \boldsymbol{x} \rangle + 2 \cdot \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{y}, \boldsymbol{y} \rangle.$$

Sketch. (a) $\langle \boldsymbol{x}, \boldsymbol{x} \rangle = \sum_{i=1}^I x_i^2 \geq 0$ and it is zero only when each $x_i = 0$. (b) $\sum_i x_i y_i = \sum_i y_i x_i$. (c) Expand componentwise: $\sum_i x_i(y_i + z_i) = \sum_i x_i y_i + \sum_i x_i z_i$. (d) Likewise, $\sum_i x_i(\alpha y_i) = \alpha \sum_i x_i y_i$. □

Subspaces

Remark. If $W \subseteq V$, then $(W, +, \cdot)$ need not be a vector space by itself. Here we are talking about *subspaces.*

Proposition (Subspace test)**.** *Let $W \subseteq V$. Then $(W, +, \cdot)$ is a vector space* iff *for every $v, w \in W$ and every scalars $\lambda_1, \lambda_2 \in \mathbb{R}$,*

$$\lambda_1 v + \lambda_2 w \in W.$$

Corollary. Under the subspace test, $0 \in W$ (e.g., take $\lambda_1 = \lambda_2 = 0$).

Remark (Notation)**.** When the operations $+$ and $\cdot$ are clear, we simply write V for the vector space $(V, +, \cdot)$.

Definition (Norm)**.** Fix a vector space V. A *norm* on V is a function $\| \cdot \| : V \to \mathbb{R}$ such that, for all $v, w \in V$ and $\alpha \in \mathbb{R}$,

1. (Nonnegativity) $\|v\| \geq 0$.
2. (Definiteness) $\|v\| = 0$ iff $v = 0$.
3. (Homogeneity) $\|\alpha v\| = |\alpha| \, \|v\|$.
4. (Triangle inequality) $\|v + w\| \leq \|v\| + \|w\|$.

The Euclidean norm on $\mathbb{R}^I$

Let $x = (x_1, \ldots, x_I) \in \mathbb{R}^I$. With the standard inner product

$$x \cdot y \; = \; \sum_{i=1}^I x_i y_i,$$

the Euclidean norm is

$$\|x\|_2 \; = \; \sqrt{x \cdot x} \; = \; \Big(\sum_{i=1}^I x_i^2 \Big)^{1/2}.$$

Norm properties.

(1) $\|x\|_2 \geq 0$ for all x (sum of squares).

(2) $\|x\|_2 = 0 \iff x = 0$.

(3) Homogeneity: for any scalar $\lambda \in \mathbb{R}$,

$$\|\lambda x\|_2 = \sqrt{(\lambda x) \cdot (\lambda x)} = \sqrt{\lambda^2\, x \cdot x} = |\lambda|\, \|x\|_2.$$

(4) Triangle inequality: for any $v, w \in \mathbb{R}^I$,

$$\|v + w\|_2 \leq \|v\|_2 + \|w\|_2.$$

A direct coordinate expansion gives

$$\|v + w\|_2^2 = (v + w) \cdot (v + w) = \|v\|_2^2 + 2\, v \cdot w + \|w\|_2^2,$$

which does not by itself yield the inequality unless one controls $v \cdot w$. Using Cauchy–Schwarz, $|v \cdot w| \leq \|v\|_2 \|w\|_2$, hence

$$\|v + w\|_2^2 \leq \|v\|_2^2 + 2\|v\|_2\|w\|_2 + \|w\|_2^2 = (\|v\|_2 + \|w\|_2)^2,$$

and taking square roots proves (4).

Amanda's proof (triangle inequality ⇔ Cauchy–Schwarz).

$$\begin{aligned}
\|x + y\|_2 \leq \|x\|_2 + \|y\|_2 &\iff \|x + y\|_2^2 \leq (\|x\|_2 + \|y\|_2)^2 \\
&\iff (x + y) \cdot (x + y) \leq \|x\|_2^2 + 2\|x\|_2\|y\|_2 + \|y\|_2^2 \\
&\iff x \cdot x + 2\, x \cdot y + y \cdot y \leq \|x\|_2^2 + 2\|x\|_2\|y\|_2 + \|y\|_2^2 \\
&\iff x \cdot y \leq \|x\|_2\|y\|_2,
\end{aligned}$$

which is precisely the Cauchy–Schwarz inequality

$$\sum_{i=1}^{I} x_i y_i \;\leq\; \Big(\sum_{i=1}^{I} x_i^2\Big)^{1/2} \Big(\sum_{i=1}^{I} y_i^2\Big)^{1/2}.$$

Thus Cauchy–Schwarz implies (and is equivalent to, through the chain above) the triangle inequality for $\|\cdot\|_2$.

The sup–norm on bounded functions. Let $V = \mathcal{F}^B(X, \mathbb{R})$ be the vector space of all bounded real–valued functions on X. For $f \in V$ define the sup–norm

$$\|f\|_\infty \;:=\; \sup_{x \in X} |f(x)|.$$

Axioms (1) nonnegativity and (2) $\|f\|_\infty = 0 \iff f \equiv 0$ are immediate. For the other two axioms:
(3) Absolute homogeneity. For every $\lambda \in \mathbb{R}$,

$$\|\lambda f\|_\infty = \sup_{x \in X} |\lambda f(x)| = |\lambda| \sup_{x \in X} |f(x)| = |\lambda|\, \|f\|_\infty.$$

(4) Triangle inequality. For $f, g \in V$,

$$\|f + g\|_\infty = \sup_{x \in X} |f(x) + g(x)| \leq \sup_{x \in X} \big(|f(x)| + |g(x)|\big) \leq \sup_{x \in X} |f(x)| + \sup_{x \in X} |g(x)| = \|f\|_\infty + \|g\|_\infty.$$

Hence $\|\cdot\|_\infty$ is a norm on $\mathcal{F}^B(X, \mathbb{R})$.

Remark 4 (Every norm induces a metric). If $(V, \|\cdot\|)$ is a normed vector space, then

$$d(v,w) \;:=\; \|v-w\|$$

defines a metric on V. (Nonnegativity and symmetry are clear, $d(v,w)=0 \iff v=w$ by the norm axiom, and the triangle inequality follows from $\|v-w\| \le \|v-z\| + \|z-w\|$.)

Examples.

- On $\mathbb{R}^t$ with the Euclidean norm $\|\cdot\|_2$: $d_{\text{euc}}(v,w) = \|v-w\|_2$.
- On $\mathcal{F}^B(X,\mathbb{R})$ with the sup–norm: $d_\infty(f,g) = \|f-g\|_\infty$.

Key idea to remember: a *norm* measures the *length* of a vector; a *metric* measures the *distance* between two points.

Vectors in $\mathbb{R}^2$ as points and as displacements

- We begin by focusing on the vector space $\mathbb{R}^2$.
- Pictorially, a vector in $\mathbb{R}^2$ can be seen in two (equivalent) ways:
  1. as a *point* (x_1, x_2) in the plane, or
  2. as a *displacement* (an arrow) with a length and a direction.

Let $\mathbf{x} = (x_1, x_2)$ and $\mathbf{y} = (y_1, y_2)$. The displacement that carries $\mathbf{x}$ to $\mathbf{y}$ is

$$\mathbf{z} = \mathbf{y} - \mathbf{x} = (y_1 - x_1,\ y_2 - x_2).$$

Equivalently, $\mathbf{x} + \mathbf{z} = \mathbf{y}$.

Key idea. Thinking of vectors as displacements, two arrows are *equivalent* if they have the same length and the same direction. Thus the arrow drawn from the origin to $\mathbf{z}$ is equivalent to the arrow drawn from $\mathbf{x}$ to $\mathbf{y}$. The arrow $\mathbf{w} = \mathbf{x} - \mathbf{y} = -(\mathbf{y} - \mathbf{x})$ has the same length but the *opposite* direction, so it is not equivalent to $\mathbf{z}$.

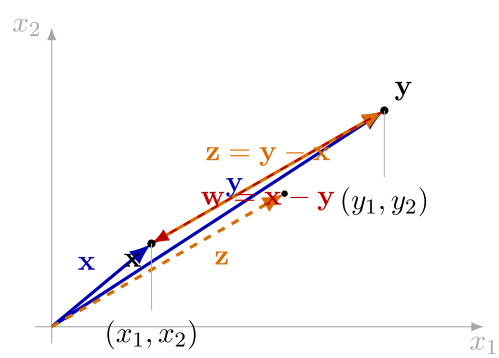

Figure 8: Vectors as points and as displacements: $\mathbf{x} + \mathbf{z} = \mathbf{y}$ with $\mathbf{z} = \mathbf{y} - \mathbf{x}$. The dashed orange arrow is equivalent to the solid orange arrow (same length and direction).

Two viewpoints on $\mathbb{R}^2$

- *As a set of points:* elements are (x_1, x_2).
- *As a set of displacements:* elements are arrows with length and direction; two displacements are equivalent iff they share the same length and direction.

Remarks on notation

- We often write a vector $\mathbf{x}$ in coordinates as $\mathbf{x} = (x_1, \ldots, x_J)$. When it is helpful to stress "column" form we use
$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_J \end{pmatrix} \in \mathbb{R}^{J\times 1}.$$
- For a scalar $\lambda \in \mathbb{R}$ we may use the constant vector
$$\bar{\lambda} = (\lambda, \ldots, \lambda).$$
- The zero vector is denoted by $\mathbf{0} = (0, \ldots, 0)$.

Scaling a vector

Let $\mathbf{x} \in \mathbb{R}^j$ and $\lambda \in \mathbb{R} \setminus \{0\}$. Scaling by λ produces the vector $\lambda\mathbf{x}$:

$$\begin{cases} \lambda > 1: & \text{stretches } \mathbf{x} \text{ in the same direction,} \\ 0 < \lambda < 1: & \text{shrinks } \mathbf{x} \text{ in the same direction,} \\ \lambda < 0: & \text{reverses direction and scales by } |\lambda|. \end{cases}$$

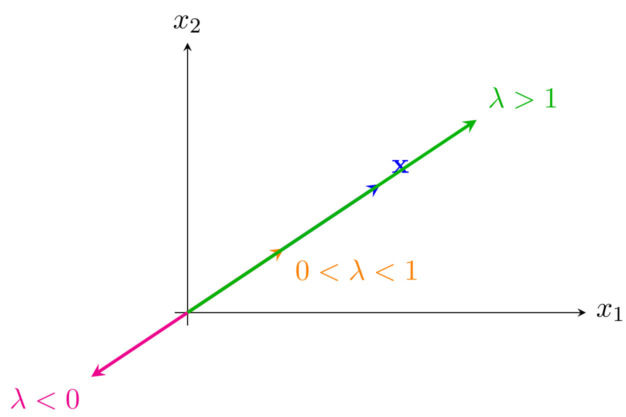

Inner product and the angle between vectors

In $\mathbb{R}^j$ with the standard (Euclidean) inner product and norm,

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{j} x_i y_i, \qquad \|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}.$$

By the law of cosines applied to the triangle with sides $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{y}-\mathbf{x}$,

$$\|\mathbf{y}-\mathbf{x}\|_2^2 = \|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 - 2\,\|\mathbf{x}\|_2\,\|\mathbf{y}\|_2\,\cos\theta.$$

On the other hand,

$$\|\mathbf{y}-\mathbf{x}\|_2^2 = \langle \mathbf{y}-\mathbf{x}, \mathbf{y}-\mathbf{x}\rangle = \|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 - 2\,\langle \mathbf{x}, \mathbf{y}\rangle\,.$$

Equating the two expressions gives the fundamental link between angle and inner product:

$$\boxed{\cos\theta = \frac{\langle \mathbf{x}, \mathbf{y}\rangle}{\|\mathbf{x}\|_2\,\|\mathbf{y}\|_2}}$$

Hence $\langle \mathbf{x}, \mathbf{y}\rangle = 0 \iff \cos\theta = 0 \iff \theta = 90^\circ$ (the vectors are perpendicular).

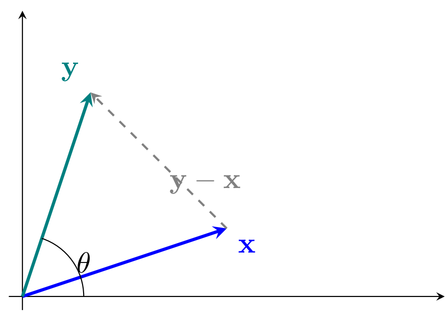

Orthogonality and orthonormality

- Two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^j$ are *orthogonal* if $\langle \mathbf{x}, \mathbf{y}\rangle = 0$.
- A vector $\mathbf{x}$ is orthogonal to a set $X \subseteq \mathbb{R}^j$ if $\langle \mathbf{x}, \mathbf{y}\rangle = 0$ for every $\mathbf{y} \in X$.
- The zero vector satisfies $\langle \mathbf{0}, \mathbf{y}\rangle = 0$ for all $\mathbf{y}$, so $\mathbf{0}$ is orthogonal to every vector.
- A set $\{\mathbf{u}_1, \ldots, \mathbf{u}_m\} \subset \mathbb{R}^j$ is *orthonormal* if the vectors are pairwise orthogonal and each has Euclidean norm 1:
$$\langle \mathbf{u}_i, \mathbf{u}_k\rangle = 0 \ (i \neq k), \qquad \|\mathbf{u}_i\|_2 = 1.$$

Example. Let $\mathbf{x} = (3,1)$ and $\mathbf{y} = (-1,3)$. Then

$$\langle \mathbf{x}, \mathbf{y}\rangle = 3(-1) + 1\cdot 3 = -3 + 3 = 0 \quad \Rightarrow \quad \mathbf{x} \perp \mathbf{y}.$$

If we scale $\mathbf{y}$ by $\frac{1}{2}$, $\mathbf{z} := \frac{1}{2}\mathbf{y} = \left(-\frac{1}{2}, \frac{3}{2}\right)$, then

$$\langle \mathbf{x}, \mathbf{z}\rangle = \tfrac{1}{2}\,\langle \mathbf{x}, \mathbf{y}\rangle = 0,$$

so $\mathbf{x} \perp \mathbf{z}$ as well. (For comparison, $\mathbf{w} = (-2,3)$ gives $\langle \mathbf{x}, \mathbf{w}\rangle = 3(-2) + 1\cdot 3 = -3 \neq 0$, hence not orthogonal.)

Example. Let $\mathbf{x} = (1,0)$ and $\mathbf{y} = (0,2)$. Then $\langle \mathbf{x}, \mathbf{y}\rangle = 0$ so they are orthogonal. However,

$$\|\mathbf{x}\|_2 = 1, \qquad \|\mathbf{y}\|_2 = 2,$$

so the pair is *not* orthonormal (the second vector does not have unit length).

Linear combinations and span

Definition (Linear combination). Let V be a vector space and let $X = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\} \subset V$ be finite. A vector $\mathbf{v} \in V$ is a *linear combination* of X if there exist scalars $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$ such that

$$\mathbf{v} = \sum_{j=1}^{k} \lambda_j \, \mathbf{v}_j.$$

Coordinate form in $\mathbb{R}^J$. Write $\mathbf{x}^1, \ldots, \mathbf{x}^k \in \mathbb{R}^J$ with $\mathbf{x}^r = (x_{r1}, \ldots, x_{rJ})$. A vector $\mathbf{x} = (x_1, \ldots, x_J) \in \mathbb{R}^J$ is a linear combination of $\{\mathbf{x}^1, \ldots, \mathbf{x}^k\}$ iff there exist $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$ with

$$x_i = \sum_{r=1}^{k} \lambda_r \, x_{ri} \qquad \text{for each } i = 1, \ldots, J.$$

Definition (Span). For $W \subset V$ define the *span* of W by

$$\operatorname{span}(W) = \left\{ \sum_{j=1}^{m} \lambda_j \, \mathbf{w}_j : \; \mathbf{w}_j \in W, \; \lambda_j \in \mathbb{R}, \; m \in \mathbb{N} \right\}.$$

Equivalently, $\mathbf{v} \in \operatorname{span}(W)$ iff there exists a *finite* subset $\widehat{W} \subset W$ such that $\mathbf{v}$ is a linear combination of vectors in $\widehat{W}$.

Remark. $\mathbf{0} \in V$ is a linear combination of any finite set (take all coefficients $\lambda_j = 0$). For two vectors $\mathbf{x}_1, \mathbf{x}_2 \in V$, any vector of the form $\mathbf{x}_3 = \lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2$ is reached by scaling and then adding.

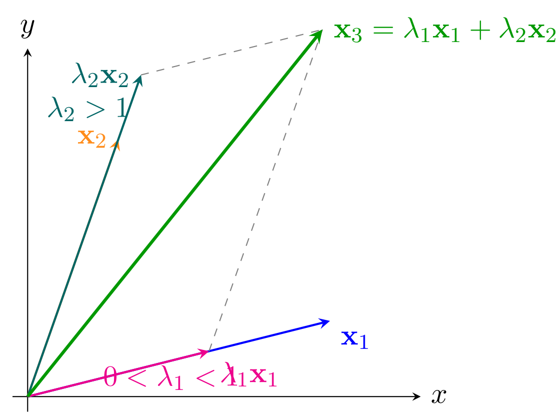

Basic facts. For any $W \subset V$,

$$W \subset \operatorname{span}(W).$$

Example. If $W = \{(1,0,0), (0,1,0)\} \subset \mathbb{R}^3$, then

$$\operatorname{span}(W) = \{(z_1, z_2, 0) : z_1, z_2 \in \mathbb{R}\} \qquad \text{(the } xy\text{-plane)}.$$

Also, is true that $\{\,\mathbf{0}\,\} = \operatorname{span}(\{\mathbf{0}\})$ and $V = \operatorname{span}(V)$.

Example (Collinear pair)**.** Let $\mathbf{x} = (4,2)$ and $\mathbf{y} = (2,1)$. Then $\mathbf{y} = \frac{1}{2}\,\mathbf{x}$, so

$$\tfrac{1}{2}\,\mathbf{x} - \mathbf{y} = \mathbf{0}.$$

Hence $\{\mathbf{x},\mathbf{y}\}$ is linearly dependent and

$$\operatorname{span}\{\mathbf{x},\mathbf{y}\} = \{\, \lambda\mathbf{x} : \lambda \in \mathbb{R} \,\} = \{\, \mu\mathbf{y} : \mu \in \mathbb{R} \,\},$$

the line through the origin with direction $\mathbf{x}$ (equivalently, $\mathbf{y}$).

Example (Not a rescaling)**.** Let $\mathbf{x} = (4,2)$ and $\mathbf{y} = (2,\frac{1}{2})$. If $\mathbf{y}$ were a scalar multiple of $\mathbf{x}$, we would have

$$(2,\tfrac{1}{2}) = t(4,2) \quad\Longrightarrow\quad 2 = 4t \ \text{ and } \ \tfrac{1}{2} = 2t,$$

which forces $t = \frac{1}{2}$ and $t = \frac{1}{4}$, a contradiction. Thus $\mathbf{y}$ is not a rescaling of $\mathbf{x}$, and in particular $\mathbf{x}$ is not a linear combination of the single vector $\mathbf{y}$. (Here $\operatorname{span}\{\mathbf{y}\} = \{\, s(2,\frac{1}{2}) : s \in \mathbb{R} \,\}$ is a different line.)

Example (All linear combinations of $\{\mathbf{x},\mathbf{y}\}$)**.** With $\mathbf{x} = (4,2)$ and $\mathbf{y} = (2,\frac{1}{2})$, a vector $\mathbf{z}$ is a linear combination of $\{\mathbf{x},\mathbf{y}\}$ iff there exist $\lambda_x, \lambda_y \in \mathbb{R}$ such that

$$\mathbf{z} = \lambda_x\mathbf{x} + \lambda_y\mathbf{y} = \big(4\lambda_x + 2\lambda_y,\ 2\lambda_x + \tfrac{1}{2}\lambda_y\big).$$

Since $\det\begin{pmatrix} 4 & 2 \\ 2 & \frac{1}{2} \end{pmatrix} = 4\cdot\frac{1}{2} - 2\cdot 2 = -2 \neq 0$, the pair $\{\mathbf{x},\mathbf{y}\}$ is linearly independent and therefore $\operatorname{span}\{\mathbf{x},\mathbf{y}\} = \mathbb{R}^2$. Equivalently, given $\mathbf{z} = (u,v) \in \mathbb{R}^2$ there are unique scalars

$$\lambda_x = v - \frac{u}{4}, \qquad \lambda_y = u - 2v$$

such that $\mathbf{z} = \lambda_x\mathbf{x} + \lambda_y\mathbf{y}$.

Remark (Nontrivial relation $\Rightarrow$ one vector depends on the others)**.** In general, if $\lambda_x\mathbf{x} + \lambda_y\mathbf{y} + \lambda_z\mathbf{z} = \mathbf{0}$ with at least one coefficient nonzero, then one of the vectors is a linear combination of the other two. For instance, if $\lambda_z \neq 0$, then

$$\mathbf{z} = -\frac{\lambda_x}{\lambda_z}\,\mathbf{x} - \frac{\lambda_y}{\lambda_z}\,\mathbf{y}.$$

Definition (Linear dependence and independence)**.** Let V be a vector space and let $W = \{v_1, \ldots, v_k\} \subset V$ be a finite set.

(i) W is *linearly dependent* if there exist scalars $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$, not all zero, such that

$$\sum_{i=1}^{k} \lambda_i\, v_i = \mathbf{0}.$$

(ii) W is *linearly independent* if for every choice of scalars $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$,

$$\sum_{i=1}^{k} \lambda_i\, v_i = \mathbf{0} \quad\Longrightarrow\quad \lambda_1 = \cdots = \lambda_k = 0.$$

Remark. Intuition: a linear dependence is a *non-trivial* way to obtain the zero vector; equivalently, we have two different ways to "get $\mathbf{0}$", so the set should be called dependent.

Example. Let $W \subset V$ be finite and suppose $\mathbf{0} \in W$. Write $W = \{v_1, \ldots, v_k, \mathbf{0}\}$. Then

$$1 \cdot \mathbf{0} + \sum_{i=1}^{k} 0 \cdot v_i \;=\; \mathbf{0}$$

is a non-trivial linear relation among the vectors in W. Hence W is linearly dependent.

Proposition. *Let $W \subseteq V$ be a finite set of vectors.*

(a) If W contains a linearly dependent subset, then W is linearly dependent.

(b) If W is linearly independent, then every nonempty subset of W is linearly independent.

Proof. (a) Let $\widehat{W} = \{v_1, \ldots, v_j\} \subseteq W$ be linearly dependent. Then there exist scalars $\lambda_1, \ldots, \lambda_j$, not all zero, with

$$\sum_{i=1}^{j} \lambda_i \, v_i = \mathbf{0}.$$

Write $W = \{v_1, \ldots, v_j, v_{j+1}, \ldots, v_\ell\}$ and extend the list of scalars by $\lambda_{j+1} = \cdots = \lambda_\ell = 0$. Then

$$\sum_{i=1}^{\ell} \lambda_i \, v_i = \mathbf{0}$$

with some $\lambda_i \neq 0$, so W is linearly dependent.

(b) This is the contrapositive of (a). If a nonempty subset of W were linearly dependent, then (a) would force W to be dependent as well, contrary to the hypothesis. Hence every nonempty subset is linearly independent. □

Remark (Convention)**.** The empty set $\varnothing$ is considered *linearly independent.*

Proposition. *Let $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_k\} \subset \mathbb{R}^t$ be a set of nonzero, pairwise orthogonal vectors (i.e., $\langle \mathbf{x}_i, \mathbf{x}_j \rangle = 0$ for $i \neq j$). Then $\{\mathbf{x}_1, \ldots, \mathbf{x}_k\}$ is linearly independent.*

Proof. Suppose $\sum_{i=1}^{k} \lambda_i \, \mathbf{x}_i = \mathbf{0}$. Fix $j \in \{1, \ldots, k\}$ and take the inner product with $\mathbf{x}_j$:

$$0 = \Big\langle \mathbf{x}_j, \sum_{i=1}^{k} \lambda_i \, \mathbf{x}_i \Big\rangle = \sum_{i=1}^{k} \lambda_i \, \langle \mathbf{x}_j, \mathbf{x}_i \rangle = \lambda_j \, \langle \mathbf{x}_j, \mathbf{x}_j \rangle = \lambda_j \, \|\mathbf{x}_j\|^2.$$

Since $\mathbf{x}_j \neq \mathbf{0}$, we have $\|\mathbf{x}_j\|^2 > 0$, hence $\lambda_j = 0$. Because j was arbitrary, all $\lambda_i = 0$, so the set is linearly independent. □

Example (Back to Example 7)**.** Take $\mathbf{x} = (4, 2)$ and $\mathbf{y} = (2, \frac{1}{2})$ in $\mathbb{R}^2$. They are *not* orthogonal because

$$\langle \mathbf{x}, \mathbf{y} \rangle = 4 \cdot 2 + 2 \cdot \tfrac{1}{2} = 8 + 1 \neq 0.$$

Nevertheless they are linearly independent: if $\alpha \mathbf{x} + \beta \mathbf{y} = \mathbf{0}$ then

$$\begin{cases} 4\alpha + 2\beta = 0, \\ 2\alpha + \frac{1}{2}\beta = 0. \end{cases}$$

Multiplying the second equation by 4 gives $8\alpha + 2\beta = 0$; subtracting 2 times the first equation yields $-2\beta = 0$, so $\beta = 0$, and then $\alpha = 0$. Hence $\{\mathbf{x}, \mathbf{y}\}$ is LI. This shows orthogonality is *sufficient* but not *necessary* for linear independence.

Definition (Basis)**.** Let $W \le V$ be a vector subspace. A finite set $B = \{v_1, \dots, v_k\} \subset W$ is a *basis* of W iff

1. B is linearly independent, and
2. $\mathrm{span}(B) = W$.

Equivalently, a basis is a minimal generating set of W (no vector of B lies in the span of the others).

Example (Standard basis of $\mathbb{R}^J$)**.** Define $e_i = (0, \dots, 0, 1, 0, \dots, 0) \in \mathbb{R}^J$ with 1 in the i-th position. Then $E = \{e_1, \dots, e_J\}$ is a basis of $\mathbb{R}^J$:

- **LI:** If $\sum_{i=1}^J \lambda_i e_i = \mathbf{0}$, then every coordinate equals 0, hence $\lambda_i = 0$ for all i.
- **Spanning:** For $x = (x_1, \dots, x_J) \in \mathbb{R}^J$ we have $x = \sum_{i=1}^J x_i e_i$.

Example (A non-standard basis of $\mathbb{R}^2$)**.** Let $\mathbf{x} = (4, 2)$ and $\mathbf{y} = (2, \frac{1}{2})$. The matrix with columns $\mathbf{x}, \mathbf{y}$ is

$$A = \begin{pmatrix} 4 & 2 \\ 2 & \frac{1}{2} \end{pmatrix}, \qquad \det A = 4 \cdot \tfrac{1}{2} - 2 \cdot 2 = -2 \neq 0.$$

Hence $\{\mathbf{x}, \mathbf{y}\}$ is linearly independent and therefore a basis of $\mathbb{R}^2$; consequently $\mathrm{span}\{\mathbf{x}, \mathbf{y}\} = \mathbb{R}^2$. (Another basis is the standard one $\{e_1, e_2\}$; both bases have two elements.)

Invariance of basis size If B and B' are both bases of the same subspace W, then $|B| = |B'|$ (they have the same number of elements). This number is called the *dimension* of W and is denoted $\dim W$.

Example (A 2-dimensional subspace of $\mathbb{R}^3$)**.** Let

$$W = \{(z_1, z_2, z_3) \in \mathbb{R}^3 : \; z_3 = 0\}.$$

Then $W = \mathrm{span}\{(1, 0, 0), (0, 1, 0)\}$ and these two vectors are LI, so $\{(1, 0, 0), (0, 1, 0)\}$ is a basis of W. Hence $\dim W = 2$ (the xy-plane).

Linear transformation and matrices

Definition (Linear transformation). A map $L : V \to W$ between vector spaces is *linear* if, for all $\mathbf{u}, \mathbf{v} \in V$ and all $\lambda \in \mathbb{R}$,

$$\text{(Additivity)} \quad L(\mathbf{u} + \mathbf{v}) = L(\mathbf{u}) + L(\mathbf{v}), \qquad \text{(Homogeneity)} \quad L(\lambda\, \mathbf{v}) = \lambda\, L(\mathbf{v}).$$

Characterizing linear maps $L : \mathbb{R}^t \to \mathbb{R}$

Proposition. *A function $L : \mathbb{R}^t \to \mathbb{R}$ is linear iff there exists a vector $\mathbf{v} = (v_1, \dots, v_t) \in \mathbb{R}^t$ such that, for every $\mathbf{x} = (x_1, \dots, x_t) \in \mathbb{R}^t$,*

$$L(\mathbf{x}) = \mathbf{v} \cdot \mathbf{x} = \sum_{i=1}^t v_i\, x_i.$$

In particular, the representing vector is unique and equals $\mathbf{v} = \big(L(\mathbf{e}_1), \dots, L(\mathbf{e}_t)\big)$, where $\{\mathbf{e}_1, \dots, \mathbf{e}_t\}$ is the standard basis of $\mathbb{R}^t$.

Proof. ($\Rightarrow$) Assume L is linear. Any $\mathbf{x} \in \mathbb{R}^t$ can be written as $\mathbf{x} = \sum_{i=1}^t x_i \, \mathbf{e}_i$. By linearity,

$$L(\mathbf{x}) = L\left(\sum_{i=1}^t x_i \, \mathbf{e}_i\right) = \sum_{i=1}^t x_i \, L(\mathbf{e}_i).$$

Set $v_i := L(\mathbf{e}_i)$ and $\mathbf{v} = (v_1, \ldots, v_t)$; then $L(\mathbf{x}) = \sum_{i=1}^t v_i x_i = \mathbf{v} \cdot \mathbf{x}$.

($\Leftarrow$) Conversely, suppose $L(\mathbf{x}) = \mathbf{v} \cdot \mathbf{x}$ for some fixed $\mathbf{v}$. Then for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^t$ and $\lambda \in \mathbb{R}$,

$$L(\mathbf{x} + \mathbf{y}) = \mathbf{v} \cdot \mathbf{x} + \mathbf{y} = \mathbf{v} \cdot \mathbf{x} + \mathbf{v} \cdot \mathbf{y} = L(\mathbf{x}) + L(\mathbf{y}),$$

$$L(\lambda \mathbf{x}) = \mathbf{v} \cdot \lambda \mathbf{x} = \lambda \, \mathbf{v} \cdot \mathbf{x} = \lambda \, L(\mathbf{x}),$$

so L is linear. Uniqueness of $\mathbf{v}$ follows by evaluating at the basis vectors $\mathbf{e}_i$. □

Checks. If $L(\mathbf{x}) = \mathbf{v} \cdot \mathbf{x}$, then

$$L(\mathbf{x} + \mathbf{y}) = \mathbf{v} \cdot \mathbf{x} + \mathbf{y} = \mathbf{v} \cdot \mathbf{x} + \mathbf{v} \cdot \mathbf{y} = L(\mathbf{x}) + L(\mathbf{y}), \qquad L(\lambda \mathbf{x}) = \mathbf{v} \cdot \lambda \mathbf{x} = \lambda \, \mathbf{v} \cdot \mathbf{x} = \lambda L(\mathbf{x}).$$

Consequences.

- A linear functional $L : \mathbb{R}^t \to \mathbb{R}$ is completely determined by its values on the standard basis: $v_i = L(\mathbf{e}_i)$.
- In coordinates, $L(x_1, \ldots, x_t) = \sum_{i=1}^t v_i x_i$.

Characterizing linear transformations $L : \mathbb{R}^I \to \mathbb{R}^J$

Let $L : \mathbb{R}^I \to \mathbb{R}^J$ be linear. For each $j = 1, \ldots, J$ define the linear functional $L_j : \mathbb{R}^I \to \mathbb{R}$ by $L(\mathbf{x}) = \big(L_1(\mathbf{x}), \ldots, L_J(\mathbf{x})\big)$. For every j set

$$\mathbf{v}_j \coloneqq \big(L_j(\mathbf{e}_1), \, L_j(\mathbf{e}_2), \, \ldots, \, L_j(\mathbf{e}_I)\big) \in \mathbb{R}^I.$$

Then, for all $\mathbf{x} \in \mathbb{R}^I$,

$$L_j(\mathbf{x}) = \langle \mathbf{v}_j, \, \mathbf{x} \rangle \qquad \text{and} \qquad L(\mathbf{x}) = \big(\langle \mathbf{v}_1, \, \mathbf{x} \rangle, \, \ldots, \, \langle \mathbf{v}_J, \, \mathbf{x} \rangle \big) \in \mathbb{R}^J.$$

Matrix viewpoint. When we regard inputs/outputs as column vectors,

$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_I \end{pmatrix}, \qquad L(\mathbf{x}) = \begin{pmatrix} \langle \mathbf{v}_1, \, \mathbf{x} \rangle \\ \vdots \\ \langle \mathbf{v}_J, \, \mathbf{x} \rangle \end{pmatrix} = \underbrace{\begin{pmatrix} \mathbf{v}_1^\top \\ \vdots \\ \mathbf{v}_J^\top \end{pmatrix}}_{A \in \mathbb{R}^{J \times I}} \mathbf{x}.$$

Thus every linear map $L : \mathbb{R}^I \to \mathbb{R}^J$ is represented by the (unique) matrix A whose j-th row is $\mathbf{v}_j^\top$, and $L(\mathbf{x}) = A\mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^I$.

Example (From linear forms to the matrix of L)**.** Let $L : \mathbb{R}^3 \to \mathbb{R}^2$ be

$$L(x_1, x_2, x_3) = \big(\underbrace{x_1 + 2x_2 + 3x_3}_{\ell_1(\mathbf{x})}, \, \underbrace{4x_1 + 5x_2 + 6x_3}_{\ell_2(\mathbf{x})} \big), \qquad \mathbf{x} = (x_1, x_2, x_3).$$

Define $\mathbf{v}_1 = (1, 2, 3)$ and $\mathbf{v}_2 = (4, 5, 6)$. Then

$$\ell_1(\mathbf{x}) = \langle \mathbf{v}_1, \mathbf{x} \rangle, \qquad \ell_2(\mathbf{x}) = \langle \mathbf{v}_2, \mathbf{x} \rangle,$$

so $L(\mathbf{x}) = (\langle \mathbf{v}_1, \mathbf{x} \rangle, \langle \mathbf{v}_2, \mathbf{x} \rangle)$.

Viewing L as a $J \times I = 2 \times 3$ matrix,

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} = \begin{pmatrix} \mathbf{v}_1^\top \\ \mathbf{v}_2^\top \end{pmatrix}, \qquad A\mathbf{x} = \begin{pmatrix} x_1 + 2x_2 + 3x_3 \\ 4x_1 + 5x_2 + 6x_3 \end{pmatrix} = L(\mathbf{x}).$$

Equivalently, if $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ is the standard basis of $\mathbb{R}^3$, the *columns* of A are $L(\mathbf{e}_1)$, $L(\mathbf{e}_2)$, $L(\mathbf{e}_3)$, and

$$L(\mathbf{x}) = \sum_{i=1}^{3} x_i \, L(\mathbf{e}_i).$$

Proposition (Standard matrix of a linear map)**.** *Let $L : \mathbb{R}^I \to \mathbb{R}^J$ be linear. Its (unique) standard matrix $A_L \in \mathbb{R}^{J \times I}$ is*

$$A_L = \begin{bmatrix} L(\mathbf{e}_1) & L(\mathbf{e}_2) & \cdots & L(\mathbf{e}_I) \end{bmatrix},$$

so that for every $\mathbf{x} \in \mathbb{R}^I$,

$$L(\mathbf{x}) = A_L \, \mathbf{x} \qquad \text{and} \qquad L(\mathbf{x}) = \sum_{i=1}^{I} x_i \, L(\mathbf{e}_i).$$

Conversely, given any $A \in \mathbb{R}^{J \times I}$ there is a (unique) linear map $L_A : \mathbb{R}^I \to \mathbb{R}^J$ defined by $L_A(\mathbf{x}) = A\mathbf{x}$ for all $\mathbf{x}$.

Example (Dropping a coordinate vs. projection matrix)**.** Take $L : \mathbb{R}^3 \to \mathbb{R}^2$ given by $L(x_1, x_2, x_3) = (x_1, x_2) = (\ell_1(\mathbf{x}), \ell_2(\mathbf{x}))$, where $\ell_1(\mathbf{x}) = x_1$ and $\ell_2(\mathbf{x}) = x_2$. Then

$$L(\mathbf{e}_1) = (1, 0)^\top, \quad L(\mathbf{e}_2) = (0, 1)^\top, \quad L(\mathbf{e}_3) = (0, 0)^\top,$$

and the standard matrix is

$$A_L = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \qquad A_L \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = L(\mathbf{x}).$$

Remark. A_L is *not* a "projection matrix" in the usual sense because it is not square. The (orthogonal) projection in $\mathbb{R}^3$ onto the subspace $\operatorname{span}\{\mathbf{e}_1, \mathbf{e}_2\}$ is the 3×3 matrix

$$P = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad P^2 = P \text{ (idempotent)}, \quad P^\top = P \text{ (symmetric)}.$$

Example (Identity map)**.** Let $L : \mathbb{R}^I \to \mathbb{R}^I$ be the identity: $L(\mathbf{x}) = \mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^I$. Then for the standard basis $(\mathbf{e}_1, \ldots, \mathbf{e}_I)$ we have $L(\mathbf{e}_i) = \mathbf{e}_i$, and the standard matrix is the identity matrix

$$A_L = I_I = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}.$$

Remark (Square case)**.** Very often we consider $L:\mathbb{R}^I \to \mathbb{R}^I$. In that case the standard matrix A_L is *square*, i.e., $A_L \in \mathbb{R}^{I\times I}$.

Definition (Kernel and rank)**.** Let $L : V \to W$ be linear between vector spaces.

- The *kernel* (null space) of L is

$$\ker L \;=\; \{\, v \in V : L(v) = 0 \,\} \;=\; L^{-1}(\{0\}).$$

- The *rank* of L is

$$\operatorname{rank} L \;=\; \dim\big(L(V)\big),$$

  i.e., the dimension of the image (a subspace of W). Equivalently: the number of elements in a basis of $L(V)$.

Proposition (Column dimension equals rank)**.** *Let $L : \mathbb{R}^I \to \mathbb{R}^J$ be linear and let $A_L \in \mathbb{R}^{J\times I}$ be its standard matrix. Then*

$$\operatorname{rank} L \;=\; \dim\big(L(\mathbb{R}^I)\big) \;=\; \dim\big(\operatorname{col}(A_L)\big) \;=\; \#\{\textit{linearly independent columns of } A_L\}.$$

Corollary (Image and kernel are subspaces)**.** For any linear map $L : V \to W$,

$$L(V) \subseteq W \quad \text{is a vector subspace of } W, \qquad \ker L \subseteq V \quad \text{is a vector subspace of } V.$$

Proposition. *Let $L : V \to W$ be a linear transformation.*

(1) $L(0) = 0$ in W.

(2) If L is invertible, then the inverse map $L^{-1} : W \to V$ is a linear transformation.

(3) L is injective (one-to-one) if and only if $\ker L = \{0\}$.

(4) Let $\{v_1, \dots, v_k\}$ be a basis of V. Then

$$L : V \to W \textit{ is surjective} \iff \operatorname{span}\big(L(v_1), \dots, L(v_k)\big) = W.$$

(5) If $\{v_1, \dots, v_k\}$ are linearly dependent in V, then $\big(L(v_1), \dots, L(v_k)\big)$ is linearly dependente in W.

Intuition. About point (3):

You need *the set of v that map to* 0 is only 0. If some nonzero u satisfies $L(u) = 0$, then $L(u) = L(0)$: two different inputs ($u \neq 0$) collapse to the same output, so L is not injective. Conversely, if $Lx = Ly$ with $x \neq y$, then $L(x - y) = 0$ with $x - y \neq 0$, meaning that the kernel is not trivial.

About point (4):

Every $x \in V$ is written as a linear combination $x = \sum_i \lambda_i v_i$; by linearity, $L(x) = \sum_i \lambda_i L(v_i)$. Then L "reaches" a $w \in W$ if and only if that w is a linear combination of the images $L(v_i)$. Therefore, L is surjective exactly when $\{L(v_i)\}$ generates W.

Proposition. *Fix a linear map $L : \mathbb{R}^I \to \mathbb{R}^J$ and its standard matrix $A_L \in \mathbb{R}^{J\times I}$.*

(1) L is injective iff *the homogeneous system $A_L\,\mathbf{z} = \mathbf{0}$ has the* unique *solution $\mathbf{z} = \mathbf{0}$ (i.e.,* $\ker A_L = \{\mathbf{0}\}$*).*

(2) L is surjective iff *the columns of A_L span $\mathbb{R}^J$:*

$$\operatorname{span}\big(A_L[*,1],\dots,A_L[*,I]\big) = \mathbb{R}^J.$$

(Matrix translation of Prop. (4)).

(3) L is bijective iff *$I = J$ and there exists $B \in \mathbb{R}^{I\times I}$ with*

$$A_L\,B = B\,A_L = I_{I\times I} \quad \textit{(the identity matrix).}$$

Proposition (Rank–Nullity)**.** *Let $L : \mathbb{R}^I \to \mathbb{R}^J$ be linear. Then*

$$\dim(\ker L) \;+\; \operatorname{rank} L \;=\; I.$$

Equivalently, for the standard matrix A_L,

$$\textit{nullity}(A_L) \;+\; \textit{rank}(A_L) \;=\; \#\textit{columns of } A_L \;(= I).$$

Consequences (useful translations):

- *L injective* $\iff$ $\operatorname{rank} L = I$ *(i.e., columns linearly independent).*
- *L surjective* $\iff$ $\operatorname{rank} L = J$.
- *If $I = J$, then L bijective* $\iff$ $\operatorname{rank} L = I = J$ $\iff$ *A_L is invertible.*

Proposition (Square matrix)**.** *Fix a linear map $L : \mathbb{R}^I \to \mathbb{R}^I$ and its associated square matrix $A_L \in \mathbb{R}^{I\times I}$.*

(1) L is injective if and only if it is surjective.

(2) A_L is invertible if and only if there is a unique $\mathbf{z} \in \mathbb{R}^I$ *with $A_L\mathbf{z} = \mathbf{0}$ (equivalently* $\ker L = \{\mathbf{0}\}$*; i.e.,* injective*).*

Remark.

- Same input/output dimension $\Rightarrow$ injective $\Leftrightarrow$ surjective.
- Practical takeaway: **for a square matrix, to check invertibility it suffices to check injectivity** (or, equivalently, surjectivity).

Remark. Let $L : \mathbb{R}^I \to \mathbb{R}^I$ induce a *singular* square matrix A_L.

- Then $\ker L$ contains at least two points ($\mathbf{0}$ and some nonzero vector).
- $\ker L$ is a vector subspace (it *need not* be all of $\mathbb{R}^I$).
- Hence $\ker L$ is infinite; in particular, there are infinitely many $\mathbf{x} \in \mathbb{R}^I$ with $A_L\mathbf{x} = \mathbf{0}$.

Proposition. *Let A_L be singular. For any $\mathbf{y} \in \mathbb{R}^I$, the system $A_L\mathbf{x} = \mathbf{y}$ has either* no solutions *or* infinitely many solutions.

Idea only. If $\mathbf{x}^*$ solves $A_L\mathbf{x} = \mathbf{y}$ and $\mathbf{x}^{**} \in \ker A_L$ (exists since A_L is singular), then $A_L(\mathbf{x}^* + \mathbf{x}^{**}) = \mathbf{y} + \mathbf{0} = \mathbf{y}$. Since $\ker A_L$ is infinite, you obtain infinite solutions by adding different vectors from the kernel. If, on the other hand, there is no first solution, the system simply has no solutions. $\square$

Remark (Invertible case)**.** If A is invertible (i.e., non-singular), then L is bijective and for *each* $\mathbf{y}$ there exists a *unique* $\mathbf{x}$ with $A\mathbf{x} = \mathbf{y}$.

Determinants and invertibility

Definition. Fix a square matrix $A \in R^{1\times 1}$, then $A = [a_{11}]$ and $\det(A) = a_{11}$

Proposition. *Fix a square matrix $A \in \mathbb{R}^{I\times I}$. Then*

A *is* singular *(non-invertible)* $\iff \det(A) = 0$, A *is* non-singular *(invertible)* $\iff \det(A) \neq 0$.

Definition (Minor and cofactor)**.** Let $A \in \mathbb{R}^{I\times I}$. For $j, i \in \{1, \dots, I\}$, the M_{ji} *matrix* is obtained from A by deleting the j-th row and the i-th column; hence $M_{ji} \in \mathbb{R}^{(I-1)\times(I-1)}$.

The corresponding *cofactor-ji* is

$$c_{ji} = (-1)^{j+i} \det(M_{ji}).$$

Definition (Determinant)**.** Let $A = (a_{ji}) \in \mathbb{R}^{I\times I}$ and let $c_{ji} = (-1)^{j+i} \det(M_{ji})$ be the cofactor obtained by deleting the j-th row and i-th column. For any fixed $j \in \{1, \dots, I\}$,

$$\det(A) = \sum_{i=1}^{I} a_{ji}\, c_{ji}.$$

It does not matter which row j you choose—the final value is the same.

Example. 2×2 check. For $A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$:

$$\text{for } j = 1: \quad \det(A) = a_{11}\, a_{22} \;-\; a_{12}\, a_{21},$$

$$\text{for } j = 2: \quad \det(A) = -\, a_{21}\, a_{12} \;+\; a_{22}\, a_{11}.$$

(Same value; only the order/sign pattern changes.)

Determinant in $\mathbb{R}^3$

Let

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

Using the cofactor expansion along the first row (any fixed row/column gives the same final value),

$$\det(A) = a_{11} \det(M_{11}) - a_{12} \det(M_{12}) + a_{13} \det(M_{13}),$$

where the minors are

$$\det(M_{11}) = a_{22}a_{33} - a_{23}a_{32}, \quad \det(M_{12}) = a_{21}a_{33} - a_{31}a_{23}, \quad \det(M_{13}) = a_{21}a_{32} - a_{22}a_{31}.$$

Equivalently,

$$\boxed{\det(A) = a_{11}(a_{22}a_{33} - a_{23}a_{32}) - a_{12}(a_{21}a_{33} - a_{31}a_{23}) + a_{13}(a_{21}a_{32} - a_{22}a_{31})}$$

Definition. Properties of determinants (column-wise formulation) Throughout, $A, B, C \in \mathbb{R}^{I\times I}$ and $A[*, i]$ denotes the i-th column of A.

1. **Transpose:** $\det(A) = \det(A^\top)$.

2. **Identity:** $\det(I_I) = 1$.
3. **Duplicate column $\Rightarrow$ zero:** If $A[*, i] = A[*, \ell]$ for some $i \neq \ell$, then $\det(A) = 0$.
4. **Swap two columns flips the sign:** If

$$A[*, k] = B[*, \ell], \quad A[*, \ell] = B[*, k] \quad (k \neq \ell), \quad \text{and } A[*, i] = B[*, i] \; \forall\, i \notin \{k, \ell\},$$

then $\det(A) = -\det(B)$.
5. **Homogeneity in one column:** If

$$A[*, \ell] = c\, B[*, \ell] \quad \text{and} \quad A[*, i] = B[*, i] \; \forall\, i \neq \ell,$$

then $\det(A) = c \det(B)$.
6. **Additivity in one column:** If

$$A[*, \ell] = B[*, \ell] + C[*, \ell] \quad \text{and} \quad A[*, i] = B[*, i] = C[*, i] \; \forall\, i \neq \ell,$$

then $\det(A) = \det(B) + \det(C)$.

Also useful:

- **Multiplicativity:** $\det(AB) = \det(A)\det(B)$.
- **Triangular/diagonal:** If A is triangular, then $\det(A) = \prod_{i=1}^{I} a_{ii}$.
- **Zero column:** If some $A[*, \ell] = \mathbf{0}$, then $\det(A) = 0$ (special case of (5) with $c = 0$).

Eigenvalues and eigenvectors

Let $L : \mathbb{R}^I \to \mathbb{R}^I$ be linear. We call $\lambda \in \mathbb{R}$ an *eigenvalue* of L if there exists a nonzero vector $\mathbf{x} \neq \mathbf{0}$ such that

$$L(\mathbf{x}) = \lambda\, \mathbf{x},$$

and any such $\mathbf{x}$ is an *eigenvector* associated with λ.

Example. With

$$A = \begin{bmatrix} 2 & 0 \\ 1 & 5 \end{bmatrix}, \qquad \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix},$$

we have

$$A\mathbf{x} = \begin{bmatrix} 2x_1 \\ x_1 + 5x_2 \end{bmatrix},$$

which is typically *not* parallel to $\mathbf{x}$ (so not an eigenvector in general).

Remark. We require $\mathbf{x} \neq \mathbf{0}$ because $L(\mathbf{0}) = \mathbf{0}$ for any linear map; eigenvectors capture directions that the transformation sends to a parallel vector (possibly scaled).

Example. Projection Let

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Starting from $\mathbf{x} = (1,1,1)^\top$ gives $A\mathbf{x} = (1,1,0)^\top$ (down to the plane $z = 0$ in $\mathbb{R}^3$). There is no single λ with $A\mathbf{x} = \lambda\mathbf{x}$ for arbitrary $\mathbf{x} \in \mathbb{R}^3$, so we look inside the plane: for any $\mathbf{y} = (y_1, y_2, 0)^\top$,

$$A\mathbf{y} = \mathbf{y} \quad \Rightarrow \quad \lambda = 1 \text{ with eigenvectors in the plane } z = 0.$$

Vectors orthogonal to that plane give the other eigenvalue:

$$A \begin{bmatrix} 0 \\ 0 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = 0 \cdot \begin{bmatrix} 0 \\ 0 \\ x_3 \end{bmatrix}, \qquad \lambda = 0 \text{ with eigenvectors along the } z\text{-axis.}$$

Example. Permutation Matrix

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad A \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ x_1 \end{bmatrix} = \lambda \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

Hence:

$$\begin{cases} x_1 = x_2 \neq 0 \ \Rightarrow \ \lambda = 1, \ \text{eigenvectors} & \{(1,1)\}; \\ x_1 = -x_2 \neq 0 \Rightarrow \ \lambda = -1, \ \text{eigenvectors} & \{(1,-1)\}. \end{cases}$$

Eigenvalues and eigenvectors: general approach

- First, *solve for eigenvalues.* Then use them to obtain the corresponding eigenvectors.
- We look for nonzero vectors $\vec{x} \neq \vec{0}$ and scalars λ such that

  $$A\vec{x} = \lambda\vec{x} \iff (A - \lambda I_I)\vec{x} = \vec{0}.$$

  Denote $B := A - \lambda I_I$.
- If B were invertible, the only solution would be $\vec{x} = \vec{0}$ for every $\lambda \in \mathbb{R}$. But we want a *nonzero* solution, so B must be non-invertible (singular).
- Therefore we search for the unknown λ from the single scalar equation

  $$\det(A - \lambda I_I) = 0.$$

Corollary. A scalar $\lambda \in \mathbb{R}$ is an eigenvalue of $A \in \mathbb{R}^{I \times I}$ if and only if

$$\det(A - \lambda I_I) = 0.$$

Side note: the polynomial $p_A(\lambda) := \det(A - \lambda I_I)$ is the **characteristic polynomial**.

Example (back to the 2×2 case). For

$$A = \begin{bmatrix} 2 & 0 \\ 1 & 5 \end{bmatrix}, \qquad B = A - \lambda I_2 = \begin{bmatrix} 2-\lambda & 0 \\ 1 & 5-\lambda \end{bmatrix},$$

we get

$$\det(B) = (2-\lambda)(5-\lambda) = 0 \;\Rightarrow\; \lambda \in \{2, 5\}.$$

Basic spectral facts (no proofs)

Proposition. Fix $A \in \mathbb{R}^{I \times I}$.

1. A has I eigenvalues (counted with algebraic multiplicity), *possibly complex.*
2. If A is **symmetric** ($A = A^{\mathsf{T}}$), then all eigenvalues are real.
3. $\det(A) = \prod_{i=1}^{I} \lambda_i$ (product of the eigenvalues, with multiplicity).
4. $\sum_{i=1}^{I} \lambda_i = \operatorname{tr}(A)$, where $\operatorname{tr}(A) = \sum_{i=1}^{I} a_{ii}$.

Example.

$$A = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \quad \Longrightarrow \quad A - \lambda I_2 = \begin{bmatrix} 2-\lambda & 0 \\ 0 & 2-\lambda \end{bmatrix}, \quad \det(A - \lambda I_2) = (2-\lambda)^2 = 0,$$

so $\lambda = 2$ *twice* (double eigenvalue).

Remark (many eigenvectors for one eigenvalue)**.** If λ is an eigenvalue of $A \in \mathbb{R}^{I \times I}$ with eigenvector $\mathbf{x} \neq \mathbf{0}$, then there are *infinitely many* eigenvectors associated with λ:

$$(A - \lambda I_I)(\alpha\,\mathbf{x}) = \alpha\,(A - \lambda I_I)\mathbf{x} = \mathbf{0} \qquad \text{for any } \alpha \in \mathbb{R} \setminus \{0\}.$$

(Back to the "singular matrix" remark: if $B \in \mathbb{R}^{I \times I}$ is singular, then $B\mathbf{z} = \mathbf{0}$ has infinitely many solutions; and $B\mathbf{z} = \mathbf{y}$ has either no solutions or infinitely many.)

Proposition. *(eigenvectors for distinct eigenvalues) Let $A \in \mathbb{R}^{I \times I}$ have I* distinct *eigenvalues $\lambda_1, \ldots, \lambda_I$ with associated eigenvectors $\mathbf{x}_1, \ldots, \mathbf{x}_I$. Then $\{\mathbf{x}_1, \ldots, \mathbf{x}_I\}$ is linearly independent.*

Example. (back to the 2×2 case) Let

$$A = \begin{pmatrix} 2 & 0 \\ 1 & 5 \end{pmatrix}.$$

For $\lambda = 2$:

$$A - 2I = \begin{pmatrix} 0 & 0 \\ 1 & 3 \end{pmatrix}, \qquad (A - 2I)\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \Longrightarrow x_1 + 3x_2 = 0.$$

Eigenvectors: $\{\,(-3t,\, t) : \; t \neq 0\,\}$.

For $\lambda = 5$:

$$A - 5I = \begin{pmatrix} -3 & 0 \\ 1 & 0 \end{pmatrix}, \qquad (A - 5I)\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \implies x_1 = 0,\ x_2 \neq 0.$$

Eigenvectors: $\{\,(0,\,t) :\ t \neq 0\,\}$.

Thus each eigenvalue has infinitely many eigenvectors (different scalar multiples), and the two eigenvector *directions* $(-3, 1)$ and $(0, 1)$ are linearly independent (*note: not necessarily orthogonal*).

Example. (Jordan block / repeated eigenvalue) Let

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \qquad B(\lambda) = A - \lambda I = \begin{pmatrix} 1-\lambda & 1 \\ 0 & 1-\lambda \end{pmatrix}.$$

$$\det B(\lambda) = (1-\lambda)^2 \ \Rightarrow\ \lambda = 1 \text{ (double root)}.$$

Solve $A\mathbf{x} = \lambda\mathbf{x}$ with $\lambda = 1$:

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \implies x_1 + x_2 = x_1 \Rightarrow x_2 = 0.$$

Eigenvectors: $\{\,(t, 0) :\ t \neq 0\,\}$. Here the (algebraic) multiplicity is 2 but there is only *one* independent eigenvector direction (the eigenvectors are all scalar multiples of each other).

Check note (to check). If we take two eigenvectors that come from the *same* eigenvalue, they are (claimed to be) linearly dependent. If they come from *different* eigenvalues, they are linearly independent.

Idea of the spectral theorem

Fix a symmetric matrix $A \in \mathbb{R}^{I \times I}$.

- There exists an *orthonormal basis* of $\mathbb{R}^I$ consisting of I eigenvectors of A.
- In that ON basis, A acts by scaling each basis vector:

  $$A\,\mathbf{e}_i = \lambda_i\,\mathbf{e}_i \quad\implies\quad [A]_{\text{ON basis}} = \operatorname{diag}(\lambda_1, \ldots, \lambda_I).$$

  (Eigenvalues may repeat; e.g., $A = I$.)
- If $\lambda_i \neq \lambda_j$, the corresponding eigenvectors in this ON basis are orthogonal (hence linearly independent).
- Equivalently, A is orthogonally diagonalizable:

  $$A = Q\,\Lambda\,Q^{\mathsf{T}}, \qquad Q^{\mathsf{T}}Q = I, \quad \Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_I),$$

  where the columns of Q form the orthonormal eigenbasis.

Quadratic Forms

We now move *beyond* linear maps. Functions like

$$f : \mathbb{R} \to \mathbb{R},\ f(x) = x^2$$

or

$$f : \mathbb{R}^2 \to \mathbb{R},\ f(x_1, x_2) = x_1^2 + 2x_1x_2 + x_2^2$$

are *not* linear. Quadratic forms are a systematic way to study such "square" expressions.

Definition (Quadratic form)**.** A function $Q : \mathbb{R}^I \to \mathbb{R}$ is a *quadratic form* if there exists a real matrix $A = (a_{ij}) \in \mathbb{R}^{I\times I}$ such that

$$Q(\mathbf{x}) = \sum_{i=1}^{I}\sum_{j=1}^{I} a_{ij}\, x_i x_j \ = \ \mathbf{x}^\mathsf{T} A\,\mathbf{x}, \qquad \mathbf{x} = (x_1, \dots, x_I)^\mathsf{T}.$$

Low–dimensional expansions. Writing out the sum shows explicitly where squares and cross terms come from.

- $I = 1$: $\quad Q(x_1) = a_{11}x_1^2$. In particular, $a_{11} > 0 \Rightarrow Q(\mathbf{x}) > 0$ for all $\mathbf{x} \neq \mathbf{0}$, and $a_{11} < 0 \Rightarrow Q(\mathbf{x}) < 0$ for all $\mathbf{x} \neq \mathbf{0}$.
- $I = 2$:
$$Q(x_1, x_2) = a_{11}x_1^2 + (a_{12} + a_{21})x_1x_2 + a_{22}x_2^2.$$
- $I = 3$:
$$\begin{aligned} Q(x_1, x_2, x_3) = {} & a_{11}x_1^2 + a_{22}x_2^2 + a_{33}x_3^2 \\ & + (a_{12} + a_{21})x_1x_2 + (a_{13} + a_{31})x_1x_3 + (a_{23} + a_{32})x_2x_3. \end{aligned}$$

With $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$ and $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$,

$$A\mathbf{x} = \begin{bmatrix} a_{11}x_1 + a_{12}x_2 \\ a_{21}x_1 + a_{22}x_2 \end{bmatrix} \quad \text{(linear in } \mathbf{x}\text{; we are "missing" squares)},$$

and left–multiplying by $\mathbf{x}^\mathsf{T}$ produces the quadratic terms:

$$\mathbf{x}^\mathsf{T} A\mathbf{x} = [x_1\ x_2] \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = a_{11}x_1^2 + (a_{12} + a_{21})x_1x_2 + a_{22}x_2^2.$$

Only the symmetric part matters. For any A, let $S = \frac{1}{2}(A + A^\mathsf{T})$. Then

$$\mathbf{x}^\mathsf{T} A\mathbf{x} = \mathbf{x}^\mathsf{T} S\mathbf{x} = \sum_i a_{ii}x_i^2 + \sum_{i<j} (a_{ij} + a_{ji})x_ix_j,$$

so a quadratic form is determined by the symmetric coefficients a_{ii} and $a_{ij} + a_{ji}$.

Quadratic forms: symmetric representation and definiteness

Matrix representation. Given a quadratic form $q : \mathbb{R}^I \to \mathbb{R}$, there exists a *unique* symmetric matrix $A \in \mathbb{R}^{I\times I}$ such that

$$q(\mathbf{x}) \;=\; \mathbf{x}^\top A\,\mathbf{x} \qquad \forall\,\mathbf{x} \in \mathbb{R}^I. \tag{1}$$

Conversely, every symmetric $A \in \mathbb{R}^{I\times I}$ induces a quadratic form via $q(\mathbf{x}) = \mathbf{x}^\top A\mathbf{x}$.

Definition. Definiteness (for symmetric matrices). Fix a symmetric $A \in \mathbb{R}^{I\times I}$.

1. **Positive definite (PD):** A is PD iff $\mathbf{x}^\top A\mathbf{x} > 0$ for every $\mathbf{x} \neq \mathbf{0}$.
2. **Negative definite (ND):** A is ND iff $\mathbf{x}^\top A\mathbf{x} < 0$ for every $\mathbf{x} \neq \mathbf{0}$.
3. **Positive semidefinite (PSD):** A is PSD iff $\mathbf{x}^\top A\mathbf{x} \geq 0$ for every $\mathbf{x} \in \mathbb{R}^I$.
4. **Negative semidefinite (NSD):** A is NSD iff $\mathbf{x}^\top A\mathbf{x} \leq 0$ for every $\mathbf{x} \in \mathbb{R}^I$.
5. **Indefinite:** A is indefinite if it is neither PSD nor NSD (i.e., $\exists\,\mathbf{x}, \mathbf{y} \neq \mathbf{0}$ with $\mathbf{x}^\top A\mathbf{x} > 0$ and $\mathbf{y}^\top A\mathbf{y} < 0$).

Example. Two examples:

- $A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Then $\mathbf{x}^\top A\mathbf{x} = x_1^2 + x_2^2 > 0$ for all $\mathbf{x} \neq \mathbf{0}$, hence A is PD.
- $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. Then $\mathbf{x}^\top A\mathbf{x} = x_1^2 \geq 0$ for all $\mathbf{x}$, and $\mathbf{x}^\top A\mathbf{x} = 0$ for any nonzero vector of the form $(0, x_2)$. Hence A is PSD (but not PD).

Proposition. *Eigenvalue characterization. Let $A \in \mathbb{R}^{I\times I}$ be symmetric with eigenvalues $\lambda_1, \dots, \lambda_I$.*

1. *A is PD $\iff$ $\lambda_i > 0$ for all i. A is ND $\iff$ $\lambda_i < 0$ for all i.*
2. *A is PSD $\iff$ $\lambda_i \geq 0$ for all i. A is NSD $\iff$ $\lambda_i \leq 0$ for all i.*
3. *Otherwise (mixed signs), A is indefinite.*

Why "all eigenvalues > 0" $\Rightarrow$ PD (symmetric case). Let $A \in \mathbb{R}^{I\times I}$ be *symmetric* with eigenpairs $\{(\lambda_i, \mathbf{v}_i)\}_{i=1}^I$. For an eigenvector $\mathbf{v}_i \neq \mathbf{0}$,

$$\mathbf{v}_i^\mathsf{T} A\mathbf{v}_i = \mathbf{v}_i^\mathsf{T}(\lambda_i\mathbf{v}_i) = \lambda_i\,\mathbf{v}_i^\mathsf{T}\mathbf{v}_i = \lambda_i\|\mathbf{v}_i\|^2.$$

Hence $\lambda_i > 0 \Rightarrow \mathbf{v}_i^\mathsf{T} A\mathbf{v}_i > 0$.

Key step (spectral theorem). Because A is symmetric, there exists an *orthonormal* eigenbasis $\{\mathbf{v}_i\}_{i=1}^I$ of $\mathbb{R}^I$. Any $\mathbf{z} \in \mathbb{R}^I$ can be written as $\mathbf{z} = \sum_{i=1}^I c_i\,\mathbf{v}_i$ with $c_i = \mathbf{v}_i^\mathsf{T}\mathbf{z}$. Then

$$\mathbf{z}^\mathsf{T} A\mathbf{z} = \Big(\sum_i c_i\mathbf{v}_i\Big)^\mathsf{T} A\Big(\sum_j c_j\mathbf{v}_j\Big) = \sum_{i=1}^I \lambda_i\,c_i^2 \quad \text{(orthogonality)}.$$

Therefore:

$$\lambda_i > 0\ \forall i \;\Rightarrow\; \mathbf{z}^\mathsf{T} A\mathbf{z} > 0\ \forall \mathbf{z} \neq \mathbf{0} \quad\Longrightarrow\quad A \text{ is PD.}$$

Analogously, if $\lambda_i < 0$ for all i, then A is ND; if $\lambda_i \geq 0$ (respectively ≤ 0) for all i, then A is PSD (respectively NSD).

Why symmetry is necessary (counterexample). Let

$$A = \begin{bmatrix} 1 & 10 \\ 0 & 1 \end{bmatrix} \qquad \text{(not symmetric).}$$

Its characteristic polynomial is $(1-\lambda)^2$, so the only eigenvalue is $\lambda = 1 > 0$. However the quadratic form

$$\mathbf{z}^\mathsf{T} A\mathbf{z} = \begin{bmatrix} x_1 & x_2 \end{bmatrix} \begin{bmatrix} 1 & 10 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = x_1^2 + 10x_1x_2 + x_2^2$$

is *indefinite*: $\mathbf{z} = (1,1) \Rightarrow \mathbf{z}^\mathsf{T} A\mathbf{z} = 12 > 0$ but $\mathbf{z} = (1,-1) \Rightarrow \mathbf{z}^\mathsf{T} A\mathbf{z} = -8 < 0$. Thus "all eigenvalues $> 0 \Rightarrow$ PD" fails without symmetry.

Corollary (symmetric A). Let $A \in \mathbb{R}^{I\times I}$ be symmetric.

1. If A is PD or ND, then A is invertible. (all eigenvalues are nonzero)
2. If A is PD (resp. ND), then A^{-1} is also PD (resp. ND). (eigenvalues of A^{-1} are $1/\lambda_i$)
3. If A is PSD or NSD but not PD/ND, then A is *not* invertible. (some $\lambda_i = 0$)

Principal minors: another trick (Sylvester's criteria)

For $A \in \mathbb{R}^{I\times I}$ and $r \in \{1,\dots,I\}$, let

$$A_{[r]} \;:=\; A(1{:}r,\, 1{:}r) \qquad \text{and} \qquad \Delta_r \;:=\; \det(A_{[r]})$$

be the *leading* $r \times r$ principal submatrix and its determinant (the leading principal minor of order r). More generally, if $S \subseteq \{1,\dots,I\}$ with $|S| = r$, the (general) principal submatrix is $A_{S,S}$ and its principal minor is $\det(A_{S,S})$.

Handy identity. For any $r \times r$ matrix B,

$$\det(-B) = (-1)^r \det(B).$$

Consequently, if A is PD then $-A$ is ND, because $\det((-A)_{[r]}) = (-1)^r \Delta_r$.

Proposition. *(Sylvester's criteria for symmetric matrices). Fix a symmetric matrix $A \in \mathbb{R}^{I\times I}$.*

1. ***Positive definite (PD).***

$$A \text{ is PD} \iff \Delta_r > 0 \quad \text{for all } r = 1,\dots,I.$$

2. ***Negative definite (ND).***

$$A \text{ is ND} \iff (-1)^r \Delta_r > 0 \quad \text{for all } r = 1,\dots,I,$$

i.e., each Δ_r has the sign of $(-1)^r$.

3. ***Positive semidefinite (PSD).***

$$A \text{ is PSD} \iff \det(A_{S,S}) \geq 0 \quad \text{for every principal submatrix } A_{S,S}.$$

4. ***Negative semidefinite (NSD).***

$$A \text{ is NSD} \iff (-1)^{|S|}\det(A_{S,S}) \ge 0 \quad \text{for every principal submatrix } A_{S,S},$$

equivalently: each principal minor is either 0 *or has the sign of* $(-1)^{|S|}$.

Example (2×2)**.** Two examples:

$$A = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix} \qquad \Delta_1 = \det\begin{bmatrix} 2 \end{bmatrix} = 2 > 0, \Delta_2 = \det(A) = 2 \cdot 2 - (-1)^2 = 4 - 1 = 3 > 0,$$

all leading principal minors are positive $\Rightarrow$ A is PD.

$$A = \begin{bmatrix} -2 & 1 \\ 1 & -2 \end{bmatrix} \qquad \Delta_1 = -2 \text{ (same sign as } (-1)^1), \Delta_2 = \det(A) = (-2)(-2) - 1 = 3 > 0 \text{ (same sign as } (-1)^2),$$

$(-1)^r \Delta_r > 0$ $(r = 1, 2)$ $\Rightarrow$ A is ND.

Corollary. If A is symmetric and PD (respectively ND), then A is invertible. Moreover, if A is PD (ND), then A^{-1} is also PD (ND).

If A is PSD or NSD but not PD/ND, then A is not invertible.

Lower–triangular matrices. A matrix $B \in \mathbb{R}^{I \times I}$ is *lower triangular* if $b_{ij} = 0$ for all $i < j$, i.e.

$$B = \begin{bmatrix} b_{11} & 0 & \cdots & 0 \\ b_{21} & b_{22} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ b_{I1} & \cdots & b_{I\,I-1} & b_{II} \end{bmatrix}.$$

If B is upper or lower triangular, then

$$\det(B) = \prod_{i=1}^{I} b_{ii}.$$

In particular, if every diagonal entry $b_{ii} \neq 0$ (e.g., all $b_{ii} > 0$) then B is invertible.

Definition. (Cholesky decomposition). A matrix $A \in \mathbb{R}^{I \times I}$ *admits a Cholesky decomposition* if there exists a lower–triangular B with *strictly positive* diagonal entries such that

$$A = B\,B^\top.$$

Proposition. *A symmetric matrix* $A \in \mathbb{R}^{I \times I}$ *is* ***positive definite (PD)*** iff *it admits a Cholesky decomposition.*

Proof. Sketch proof: Need to build intuition

- If $A = BB^\top$ with B lower triangular and $\operatorname{diag}(B) > 0$, then for any $x \in \mathbb{R}^I$,

$$x^\top A x = x^\top (BB^\top) x = (B^\top x)^\top (B^\top x) = \|B^\top x\|_2^2 \ge 0.$$

Let $k := B^\top x$. Since B is invertible (positive diagonal), $k = 0$ iff $x = 0$. Hence $x^\top A x > 0$ for all $x \neq 0$, i.e., A is PD.

- Conversely, if A is symmetric PD, then there exists a unique[1] lower–triangular B with positive diagonal such that $A = BB^{\mathsf{T}}$ (the Cholesky factor).

□

[1]Uniqueness holds with the convention that the diagonal of B is strictly positive.