Chapter 1

Introduction

This book provides an introduction to linear functional analysis, extending techniques and results of classical linear algebra to infinite-dimensional spaces. With the development of a theory of function spaces, functional analysis yields a powerful tool for the study of linear ordinary and partial differential equations. It provides fundamental insights on the existence and uniqueness of solutions, their continuous dependence on initial or boundary data, the convergence of approximations, and on various other properties.

The following remarks highlight some key results of linear algebra and their infinite-dimensional counterparts.

1.1. Linear equations

Let A be an $n \times n$ matrix. Given a vector $\mathbf{b} \in \mathbb{R}^n$, a basic problem in linear algebra is to find a vector $\mathbf{x} \in \mathbb{R}^n$ such that

$$A\mathbf{x} = \mathbf{b}\,. \tag{1.1}$$

In the theory of linear PDEs, an analogous problem is the following. Consider a bounded open set $\Omega \subset \mathbb{R}^n$ and a linear partial differential operator of the form

$$Lu \doteq -\sum_{i,j=1}^{n} (a^{ij}(x) u_{x_i})_{x_j} + \sum_{i=1}^{n} b^i(x) u_{x_i} + c(x) u\,. \tag{1.2}$$

Given a function $f : \Omega \mapsto \mathbb{R}$, find a function u, vanishing on the boundary of Ω, such that

$$Lu = f\,. \tag{1.3}$$

There are fundamental differences between the problems (1.1) and (1.3). The matrix A yields a continuous linear transformation on the finite-dimensional space $\mathbb{R}^n$. On the other hand, the differential operator L can be regarded as an unbounded (hence discontinuous) linear operator on the infinite-dimensional space $\mathbf{L}^2(\Omega)$. In particular, the domain of L is not the entire space $\mathbf{L}^2(\Omega)$ but only a suitable subspace.

In spite of these differences, since both problems (1.1) and (1.3) are linear, there are a number of techniques from linear algebra that can be applied to (1.3) as well.

(I): Positivity

Assume that the matrix A is strictly positive definite, i.e., there exists a constant $\beta > 0$ such that

$$\langle Ax, x\rangle \ \geq\ \beta\,|x|^2 \qquad \text{for all } x \in \mathbb{R}^n.$$

Then A is invertible and the equation (1.1) has a unique solution for every $\mathbf{b} \in \mathbb{R}^n$.

This result has a direct counterpart for elliptic PDEs. Namely, assume that the operator L is strictly positive definite, in the sense that (after a formal integration by parts)

(1.4)
$$\begin{aligned}\langle Lu, u\rangle_{\mathbf{L}^2} \ &=\ \int_\Omega \left(\sum_{i,j=1}^n a^{ij}(x)u_{x_i}u_{x_j} + \sum_{i=1}^n b^i(x)u_{x_i}u + c(x)u^2\right) dx \\ &\geq\ \beta\,\|u\|^2_{H_0^1(\Omega)}\end{aligned}$$

for some constant $\beta > 0$ and all $u \in H_0^1(\Omega)$. Here $H_0^1(\Omega)$ is a space of functions which vanish on the boundary of Ω and such that

$$\|u\|_{H_0^1(\Omega)} \ \doteq\ \left(\int_\Omega |u|^2\,dx + \int_\Omega |\nabla u|^2\,dx\right)^{1/2} \ <\ \infty\,;$$

see Chapter 8 for precise definitions. If (1.4) holds, one can then prove that the problem (1.3) has a unique solution $u \in H_0^1(\Omega)$, for every given $f \in \mathbf{L}^2(\Omega)$.

A key assumption, in order for the inequality (1.4) to hold, is that the operator L should be elliptic. Namely, at each point $x \in \Omega$ the $n \times n$ matrix $(a^{ij}(x))$ should be strictly positive definite.

(II): Fredholm alternative

A well-known criterion in linear algebra states that the equation (1.1) has a unique solution for every given $\mathbf{b} \in \mathbb{R}^n$ if and only if the homogeneous equation

$$A\mathbf{x} = 0$$

has only the solution $\mathbf{x} = 0$. Of course, this holds if and only if the matrix A is invertible.

In general, continuous linear operators on an infinite-dimensional space X do not share this property. Indeed, one can construct a bounded linear operator $\Lambda : X \mapsto X$ which is one-to-one but not onto, or conversely.

Yet, the finite-dimensional theory carries over to an important class of operators, namely, those of the form $\Lambda = I - K$, where I is the identity and K is a compact operator. If Λ is in this class, then one can still prove the equivalence

$$\Lambda \text{ is one-to-one} \qquad \Longleftrightarrow \qquad \Lambda \text{ is onto.}$$

By an application of this theory it follows that, for a linear elliptic operator, the equation (1.3) has a unique solution $u \in H_0^1(\Omega)$ for every $f \in \mathbf{L}^2(\Omega)$ if and only if the homogeneous equation

$$Lu = 0$$

has only the zero solution.

(III): Diagonalization

If one can find a basis $\{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ of $\mathbb{R}^n$ consisting of eigenvectors of A, then with respect to this basis the system (1.1) takes a diagonal form and is thus easy to solve.

For a general matrix A with multiple eigenvalues, it is well known that such a basis of eigenvectors need not exists. A positive result in this direction is the following. If the $n \times n$ matrix A is *symmetric*, then one can find an orthonormal basis $\{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ of the Euclidean space $\mathbb{R}^n$ consisting of eigenvectors of A. Namely,

$$\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \begin{cases} 1 & \text{if } i = j\,, \\ 0 & \text{if } i \neq j\,, \end{cases} \qquad\qquad A\mathbf{v}_k = \lambda_k \mathbf{v}_k\,.$$

Here $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ are the corresponding eigenvalues. The solution $\mathbf{x}$ of (1.1) can now be found by computing its coefficients $c_1, \ldots, c_n$ with respect to the orthonormal basis:

$$\mathbf{x} = \sum_{k=1}^{n} c_k \mathbf{v}_k, \qquad A\mathbf{x} = \sum_{k=1}^{n} \lambda_k c_k \mathbf{v}_k = \mathbf{b} = \sum_{k=1}^{n} \langle \mathbf{b}, \mathbf{v}_k \rangle \mathbf{v}_k .$$

Notice that, thanks to the basis of eigenvectors, the problem becomes decoupled. Instead of a large system of n equations in n variables, we only need to solve n scalar equations, one for each coefficient c_k. If all eigenvalues λ_k are nonzero, we thus have the explicit formula

$$\mathbf{x} = \sum_{k=1}^{n} \frac{1}{\lambda_k} \langle \mathbf{b}, \mathbf{v}_k \rangle \mathbf{v}_k . \tag{1.5}$$

One can adopt the same approach in the analysis of the elliptic operator L in (1.2), provided that $a^{ij} = a^{ji}$ and $b^i(x) = 0$. Indeed, these conditions make the operator "symmetric". One can then find a countable orthonormal basis $\{\phi_1, \phi_2, \ldots\}$ of the space $\mathbf{L}^2(\Omega)$ consisting of functions $\phi_k \in H_0^1(\Omega)$ such that

$$\langle \phi_i, \phi_j \rangle_{\mathbf{L}^2} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j, \end{cases} \qquad L\phi_k = \lambda_k \phi_k , \tag{1.6}$$

for a suitable sequence of real eigenvalues $\lambda_k \to +\infty$. Assuming that $\lambda_k \neq 0$ for all k, the unique solution of (1.3) can now be written explicitly as

$$u = \sum_{k=1}^{\infty} \frac{1}{\lambda_k} \langle f, \phi_k \rangle_{\mathbf{L}^2} \phi_k . \tag{1.7}$$

Notice the close resemblance between the formulas (1.5) and (1.7). In essence, one only needs to replace the Euclidean inner product on $\mathbb{R}^n$ by the inner product on $\mathbf{L}^2(\Omega)$.

1.2. Evolution equations

Let A be an $n \times n$ matrix. For a given initial state $\mathbf{b} \in \mathbb{R}^n$, consider the Cauchy problem

$$\frac{d}{dt}\mathbf{x}(t) = A\mathbf{x}(t), \qquad \mathbf{x}(0) = \mathbf{b} . \tag{1.8}$$

According to linear ODE theory, this problem has a unique solution:

$$\mathbf{x}(t) = e^{tA}\mathbf{b} , \tag{1.9}$$

where

$$e^{tA} = \sum_{k=0}^{\infty} \frac{t^k A^k}{k!} . \tag{1.10}$$

Notice that the right-hand side of (1.10) is defined as a convergent series of $n \times n$ matrices. Here $A^0 = I$ is the identity matrix. The family of matrices $\{e^{tA}\,;\ t \in \mathbb{R}\}$ has the "group property", namely

$$e^{0\,A} = I\,, \qquad e^{tA}e^{sA} = e^{(t+s)A} \qquad \text{for all } t, s \in \mathbb{R}\,.$$

If A is symmetric, then it admits an orthonormal basis of eigenvectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$, with corresponding eigenvalues $\lambda_1, \ldots, \lambda_n$. In this case, the solution (1.9) can be written more explicitly as

$$e^{tA}\mathbf{b} \;=\; \sum_{k=1}^{n} e^{t\lambda_k} \langle \mathbf{b},\, \mathbf{v}_k\rangle\, \mathbf{v}_k\,.$$

The theory of linear semigroups provides an extension of these results to unbounded linear operators in infinite-dimensional spaces. In particular, it applies to parabolic evolution equations of the form

$$\text{(1.11)} \qquad \frac{d}{dt}\,u(t) \;=\; -Lu(t), \qquad u(0) \;=\; g \;\in\; \mathbf{L}^2(\Omega)\,, \qquad u = 0 \text{ on } \partial\Omega\,,$$

where L is the partial differential operator in (1.2) and $\partial\Omega$ denotes the boundary of Ω. When $a^{ij} = a^{ji}$ and $b^i(x) = 0$, the elliptic operator L is symmetric and the solution can be decomposed along the orthonormal basis $\{\phi_1, \phi_2, \ldots\}$ of the space $\mathbf{L}^2(\Omega)$ considered in (1.6). This yields the representation

$$\text{(1.12)} \qquad u(t) \;=\; S_t g \;\doteq\; \sum_{k=1}^{\infty} e^{-t\lambda_k} \langle g,\, \phi_k\rangle_{\mathbf{L}^2}\, \phi_k, \qquad t \geq 0\,.$$

Notice that the operator L is unbounded (its eigenvalues satisfy $\lambda_k \to +\infty$ as $k \to \infty$). However, the operators S_t in (1.12) are bounded for every $t \geq 0$ (but not for $t < 0$). The family of linear operators $\{S_t\,;\ t \geq 0\}$ is called a *linear semigroup*, since it has the semigroup properties

$$S_0 \;=\; I\,, \qquad S_t \circ S_s \;=\; S_{t+s} \quad \text{for all } s, t \geq 0\,.$$

Intuitively, we could think of S_t as an exponential operator: $S_t \doteq e^{-Lt}$. However, since L is unbounded, one should be aware that an exponential formula such as (1.10) is no longer valid. When the explicit formula (1.12) is not available, the operators S_t must be constructed using some different approximation method. In the finite-dimensional case, the exponential of a matrix A can be recovered by

$$\text{(1.13)} \qquad e^{tA} \;=\; \lim_{n\to\infty} \left(I - \frac{t}{n}A\right)^{-n}$$

and also by

$$\text{(1.14)} \qquad e^{tA} \;=\; \lim_{\lambda\to\infty} e^{tA_\lambda}, \qquad A_\lambda \;\doteq\; A(I - \lambda^{-1}A)^{-1}\,.$$

Remarkably, the two formulas (1.13)–(1.14) retain their validity also for a wide class of unbounded operators on infinite-dimensional spaces.

The hyperbolic initial value problem

$$(1.15)\qquad u_{tt} + Lu = 0, \qquad \begin{cases} u(0) = g\,, \\ u_t(0) = h\,, \end{cases} \qquad u = 0 \text{ on } \partial\Omega\,,$$

can also be treated by similar methods.

The finite-dimensional counterpart of (1.15) is the system of second-order linear equations

$$(1.16)\qquad \frac{d^2}{dt^2}\mathbf{x}(t) + A\mathbf{x}(t) = 0, \qquad \mathbf{x}(0) = \mathbf{a}\,, \quad \frac{d}{dt}\mathbf{x}(0) = \mathbf{b}\,.$$

Here $\mathbf{x}\,, \mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ and A is an $n \times n$ matrix. Denoting time derivatives by an upper dot and setting $\mathbf{y} \doteq \dot{\mathbf{x}}$, (1.16) can be written as a first-order system:

$$(1.17)\qquad \begin{pmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{y}} \end{pmatrix} = \begin{pmatrix} 0 & I \\ -A & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix}, \qquad \begin{pmatrix} \mathbf{x}(0) \\ \mathbf{y}(0) \end{pmatrix} = \begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix}.$$

The same results valid for first-order linear ODEs can thus be applied here. If A is symmetric, then it has an orthonormal basis of eigenvectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ with corresponding eigenvalues $\lambda_1, \ldots, \lambda_n$. In this case, the solution of (1.16) can be written as

$$(1.18)\qquad \mathbf{x}(t) = \sum_{k=1}^{n} c_k(t)\mathbf{v}_k\,.$$

Each coefficient $c_k(\cdot)$ can be independently computed, by solving the second-order scalar ODE

$$\frac{d^2}{dt^2}c_k(t) + \lambda_k c_k(t) = 0\,, \qquad c_k(0) = \langle \mathbf{a}, v_k\rangle\,, \qquad \frac{d}{dt}c_k(0) = \langle \mathbf{b}, v_k\rangle\,.$$

Returning to the problem (1.15), if the elliptic operator L is symmetric, then the solution can again be decomposed along the orthonormal basis $\{\phi_1, \phi_2, \ldots\}$ of the space $\mathbf{L}^2(\Omega)$ considered in (1.6). This yields the entirely similar representation

$$(1.19)\qquad u(t) = \sum_{k=1}^{\infty} c_k(t)\,\phi_k \qquad t \geq 0\,,$$

where each function c_k is determined by the equations

$$\frac{d^2}{dt^2}c_k(t) + \lambda_k c_k(t) = 0\,, \qquad c_k(0) = \langle g, \phi_k\rangle_{\mathbf{L}^2}\,, \qquad \frac{d}{dt}c_k(0) = \langle h, \phi_k\rangle_{\mathbf{L}^2}\,.$$

1.3. Function spaces

In functional analysis, a key idea is to regard functions $f : \mathbb{R}^n \mapsto \mathbb{R}$ as points in an abstract vector space. All the information about a function is condensed in one single number $\|f\|$, which we call the *norm* of f. Typically, the norm measures the "size" of f and of its partial derivatives up to some order k. It is remarkable that so many results can be achieved in such an economical way, relying only on this single concept, coupled with the structure of vector space. This accounts for the wide success of functional analytic methods.

Toward all applications of functional analysis to integral or differential equations, one needs to develop a theory of function spaces. In this direction, it is natural to consider the spaces $\mathcal{C}^k$ of functions with bounded continuous partial derivatives up to order k. The "size" of a function $f \in \mathcal{C}^k(\mathbb{R}^n)$ is here measured by the norm

$$\|f\|_{\mathcal{C}^k} \doteq \max_{\alpha_1+\cdots+\alpha_n\leq k} \sup_{x\in\mathbb{R}^n} \left|\partial_{x_1}^{\alpha_1}\cdots\partial_{x_n}^{\alpha_n} f(x)\right|.$$

The spaces $\mathcal{C}^k$, however, are not always appropriate for the study of PDEs. Indeed, from physical or geometrical considerations one can often provide estimates not on the maximum value of a solution and its derivatives, but on their $\mathbf{L}^p$ norm, for some $p \geq 1$. This motivates the introduction of the Sobolev spaces $W^{k,p}$, containing all functions whose derivatives up to order k lie in $\mathbf{L}^p$. The "size" of a function $f \in W^{k,p}(\mathbb{R}^n)$ is now measured by the norm

$$\|f\|_{W^{k,p}} \doteq \left(\sum_{\alpha_1+\cdots+\alpha_n\leq k} \int_{\mathbb{R}^n} \left|\partial_{x_1}^{\alpha_1}\cdots\partial_{x_n}^{\alpha_n} f(x)\right|^p dx\right)^{1/p}.$$

Because of their fundamental role in PDE theory, all of Chapter 8 will be devoted to the study of Sobolev spaces.

1.4. Compactness

When solving an equation, if an explicit formula for the solution is not available, a common procedure relies on three steps:

(i) Construct a sequence of approximate solutions $(u_n)_{n\geq 1}$.

(ii) Extract a convergent subsequence $u_{n_j} \to \bar{u}$.

(iii) Prove that the limit $\bar{u}$ is a solution.

When we reach step (ii), a major difference between the Euclidean space $\mathbb{R}^N$ and abstract function spaces is encountered. Namely, in $\mathbb{R}^N$ all closed bounded sets are compact. Otherwise stated, in $\mathbb{R}^N$ the Bolzano-Weierstrass theorem holds:

- *From every bounded sequence $(u_n)_{n\geq 1}$ one can extract a convergent subsequence.*

As proved in Chapter 2 (see Theorem 2.22), this crucial property is valid in every finite-dimensional normed space but fails in every infinite-dimensional one. In a space of functions, showing that a sequence of approximate solutions is bounded, i.e., $\|u_n\| \leq C$ for some constant C and all $n \geq 1$, does not guarantee the existence of a convergent subsequence. To overcome this fundamental difficulty, two main approaches can be adopted.

(i) Introduce a weaker notion of convergence. Prove that every bounded sequence (also in an infinite-dimensional space) has a subsequence which converges in this weaker sense.
A key result in this direction, the Banach-Alaoglu theorem, will be proved at the end of Chapter 2. Weak convergence in Hilbert spaces is discussed in Chapter 5.

(ii) Consider two distinct norms, say $\|u\|_{weak} \leq \|u\|_{\text{strong}}$, with the following property. If a sequence $(u_n)_{n\geq 1}$ is bounded in the strong norm, i.e., $\|u_n\|_{\text{strong}} \leq C$, then there exists a subsequence that converges in the weak norm: $\|u_{n_j} - \bar{u}\|_{\text{weak}} \to 0$, for some limit $\bar{u}$.
Ascoli's theorem, proved in Chapter 3, and the Rellich-Kondrahov compact embedding theorem, proved in Chapter 8, yield different settings where this approach can be implemented.

A large portion of the analysis of partial differential equations ultimately relies on the derivation of a priori estimates. It is the nature of the problem at hand that dictates what kind of a priori bounds one can expect, and hence in which function spaces the solution can be found. This motivates the variety of function spaces which are currently encountered in literature.

While the techniques of functional analysis are very general and yield results of fundamental nature in an intuitive and economical way, one should be aware that only some aspects of PDE theory can be approached by functional analytic methods alone. Typically, the solutions constructed by these abstract methods lie in a Sobolev space of functions that possess just the minimum amount of regularity needed to make sense of the equations. For several elliptic and parabolic equations, it is known that solutions enjoy a much higher regularity. However, this regularity can only be established by a more detailed analysis. Further properties, such as the maximum principle for elliptic or parabolic equations and the finite propagation speed for hyperbolic equations, also require additional techniques, specifically designed for PDEs. For these issues, which are not within the scope of the present lecture notes, we refer to the monographs [**E, GT, McO, P, PW, RR, S, T**].

Chapter 3

Spaces of Continuous Functions

3.1. Bounded continuous functions

Let E be a metric space. By $\mathcal{C}(E)$ we denote the space of all continuous real-valued (possibly unbounded) functions $f : E \mapsto \mathbb{R}$. In general, this space does not have a natural norm. For this reason, we shall also consider the space $\mathcal{BC}(E)$ of all bounded continuous functions $f : E \mapsto \mathbb{R}$, with norm

$$\|f\| \doteq \sup_{x\in E} |f(x)|\,. \tag{3.1}$$

Most of this chapter will be concerned with the case where E is compact. In this case every continuous function $f : E \mapsto \mathbb{R}$ is necessarily bounded, hence $\mathcal{C}(E) = \mathcal{BC}(E)$.

Lemma 3.1. *$\mathcal{BC}(E)$ is a Banach space.*

Proof. 1. Let $(f_n)_{n\geq 1}$ be a Cauchy sequence in $\mathcal{BC}(E)$. Then for every fixed $x \in E$ the sequence of numbers $f_n(x)$ is Cauchy and hence converges to some limit, which we call $f(x)$.

2. By assumption, for every $\varepsilon > 0$ there exists N large enough so that

$$\sup_{x\in E} |f_n(x) - f_m(x)| \;\leq\; \varepsilon \qquad \text{for all } n, m \geq N\,.$$

Letting $m \to \infty$, since $f_m(x) \to f(x)$ we obtain

$$\sup_{n\geq N}\ \sup_{x\in E}\ |f_n(x) - f(x)| \;\leq\; \varepsilon\,.$$

In turn, this implies

$$\sup_{n\geq N} \|f_n - f\| \;\leq\; \varepsilon\,, \qquad \sup_{x\in E} |f(x)| \;\leq\; \varepsilon + \sup_{x\in E} |f_N(x)| \;<\; \infty\,.$$

Since $\varepsilon > 0$ was arbitrary, the first inequality shows the convergence $\|f_n - f\| \to 0$. The second inequality shows that f is bounded.

3. Finally, we prove that f is continuous. Let any $x \in E$ and $\varepsilon > 0$ be given. By uniform convergence, there exists an integer N such that $|f_N(x) - f(x)| < \varepsilon/3$ for every $x \in E$. Since f_N is continuous, there exists $\delta > 0$ such that $|f_N(y) - f_N(x)| < \varepsilon/3$ whenever $d(y,x) < \delta$. Putting together the above inequalities, when $d(y,x) < \delta$ we have

$$\begin{aligned} |f(y) - f(x)| \;&\leq\; |f(y) - f_N(y)| + |f_N(y) - f_N(x)| + |f_N(x) - f(x)| \\ &<\; \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \;=\; \varepsilon\,, \end{aligned}$$

proving that f is continuous at the point x. □

Remark 3.2 (Pointwise vs. uniform convergence). The previous argument shows that, if a sequence of continuous functions f_n converges *uniformly* to a function f, then f is continuous as well. On the other hand, a sequence of continuous functions can converge *pointwise* to a discontinuous limit. For example, on the interval $E = [0,1]$, the sequence of functions $f_n(x) = x^n$ converges pointwise to the discontinuous function

$$f(x) = \begin{cases} 0 & \text{if} \quad 0 \leq x < 1\,, \\ 1 & \text{if} \quad x = 1\,. \end{cases}$$

Clearly, here the convergence is not uniform on the whole interval $[0,1]$.

The following theorem describes a case where pointwise convergence implies uniform convergence. We recall that a sequence of functions $f_n : E \mapsto \mathbb{R}$ is *increasing* if $m < n$ implies $f_m(x) \leq f_n(x)$ for all $x \in E$.

Theorem 3.3 (Dini). *Let E be a compact metric space. If $(f_n)_{n\geq 1}$ is an increasing sequence of functions in $\mathcal{C}(E)$, converging pointwise to a continuous limit function f, then $f_n \to f$ uniformly on E.*

Proof. Fix any $\varepsilon > 0$. By the assumption of pointwise convergence, for every $x \in E$ there exists an integer $N(x)$ such that $|f_{N(x)}(x) - f(x)| < \varepsilon$.

Since $f_{N(x)}$ and f are continuous, there exists an open neighborhood V_x of x such that $|f_{N(x)}(y) - f_{N(x)}(x)| < \varepsilon$ and $|f(y) - f(x)| < \varepsilon$ for every $y \in V_x$.

Since E is compact, we can cover E with finitely many of these neighborhoods, say $E \subseteq V_{x_1} \cup \cdots \cup V_{x_m}$.

Choose the integer $N = \max\{N(x_1), \ldots, N(x_m)\}$. For every $n \geq N$ and $y \in E$, assuming that $y \in V_{x_i}$ we have

$$f_{N(x_i)}(y) \;\leq\; f_N(y) \;\leq\; f_n(y) \;\leq\; f(y)\,,$$

because the sequence is increasing. Therefore

$$\begin{aligned}|f_n(y) - f(y)| \;&\leq\; |f_{N(x_i)}(y) - f(y)| \\ &\leq\; |f_{N(x_i)}(y) - f_{N(x_i)}(x_i)| + |f_{N(x_i)}(x_i) - f(x_i)| + |f(x_i) - f(y)| \\ &<\; \varepsilon + \varepsilon + \varepsilon\,.\end{aligned}$$

Since $y \in E$ and $\varepsilon > 0$ were arbitrary, this establishes the uniform convergence $f_n \to f$. □

3.2. The Stone-Weierstrass approximation theorem

Given a domain $E \subset \mathbb{R}^n$, for computational purposes it can be useful to approximate a continuous, real-valued function $f \in \mathcal{B}C(E)$ with special functions: say, polynomials, exponential functions, or trigonometric polynomials. It is thus important to understand whether every function $f \in \mathcal{B}C(E)$ can be uniformly approximated by such functions. In this section we will prove a key result in this direction.

As a preliminary, observe that the space $\mathcal{B}C(E)$ is an **algebra**. Namely, it is closed under multiplication:

if $f, g \in \mathcal{B}C(E)$, then also $fg \in \mathcal{B}C(E)$.

Moreover, the norm of the product satisfies

$$\|fg\| \;\leq\; \|f\|\,\|g\|\,.$$

We say that a subspace $\mathcal{A} \subseteq \mathcal{B}C(E)$ is a **subalgebra** if $f, g \in \mathcal{A}$ implies $fg \in \mathcal{A}$.

Lemma 3.4 (Closure of a subalgebra). *If $\mathcal{A} \subseteq \mathcal{B}C(E)$ is a subalgebra, then its closure $\overline{\mathcal{A}}$ is a subalgebra as well.*

Proof. Indeed, assume $f, g \in \overline{\mathcal{A}}$. Then there exist uniformly convergent sequences $f_n, g_n \in \mathcal{A}$ with $f_n \to f$ and $g_n \to g$. One has

$$\|fg - f_n g_n\| \;\leq\; \|fg - f_n g\| + \|f_n g - f_n g_n\| \;\leq\; \|f - f_n\|\,\|g\| + \|f_n\|\,\|g - g_n\|\,.$$

Since the sequence $(f_n)_{n\geq 1}$ is uniformly bounded, the right-hand side approaches zero as $n \to \infty$. This shows the convergence $f_n g_n \to fg$. Since $\mathcal{A}$ is an algebra, $f_n g_n \in \mathcal{A}$ for every n. Hence $fg \in \overline{\mathcal{A}}$. □

We say that a subset $\mathcal{A} \subseteq \mathcal{B}C(E)$ **separates points** if, for every couple of distinct points $x, y \in E$, there exists a function $f \in \mathcal{A}$ such that $f(x) \neq f(y)$.

Theorem 3.5 (**Stone-Weierstrass**). *Let E be a compact metric space. If $\mathcal{A}$ is a subalgebra of $\mathcal{C}(E)$ that separates points and contains the constant functions, then $\overline{\mathcal{A}} = \mathcal{C}(E)$.*

Otherwise stated, let $\mathcal{A}$ be a family of continuous, real-valued functions $f : E \mapsto \mathbb{R}$ with the following properties:

(i) If $f, g \in \mathcal{A}$ and $a, b \in \mathbb{R}$, then the linear combination $af + bg$ lies in $\mathcal{A}$.

(ii) If $f, g \in \mathcal{A}$, then the product fg lies in $\mathcal{A}$ as well.

(iii) The constant function $f(x) \equiv 1$ lies in $\mathcal{A}$.

(iv) For every two distinct points $x, y \in E$, there exists $f \in \mathcal{A}$ such that $f(x) \neq f(y)$.

Then every continuous function $f : E \mapsto \mathbb{R}$ on the compact domain E can be uniformly approximated by functions in $\mathcal{A}$.

Proof. 1. There exists a sequence of polynomials $(p_n)_{n \geq 1}$ such that $p_n(t) \to \sqrt{t}$ uniformly for $t \in [0,1]$.

To prove the above claim, the underlying idea is to construct approximate solutions to the equation $t - p^2(t) = 0$ by iteration. We thus set $p_0(t) \equiv 0$ and, by induction on $n = 0, 1, 2, \ldots$,

$$p_{n+1}(t) = p_n(t) + \frac{1}{2}(t - p_n^2(t)) \tag{3.2}$$

(see Figure 3.2.1). By induction, one checks that $p_n(t) \leq p_{n+1}(t) \leq \sqrt{t}$ for every $t \in [0,1]$. Indeed,

$$\begin{aligned} \sqrt{t} - p_{n+1}(t) &= \sqrt{t} - p_n(t) - \frac{1}{2}(t - p_n^2(t)) \\ &= (\sqrt{t} - p_n(t)) \left(1 - \frac{1}{2}(\sqrt{t} + p_n(t))\right) \geq 0. \end{aligned}$$

For every fixed $t \in [0,1]$, the sequence $p_n(t)$ is increasing and bounded above. Hence it has a unique limit, say $g(t)$. By (3.2), this limit satisfies $t - g^2(t) = 0$. Since $g(t) \geq 0$ we conclude that $g(t) = \sqrt{t}$.

Finally, by Dini's theorem, the convergence is uniform for $t \in [0,1]$.

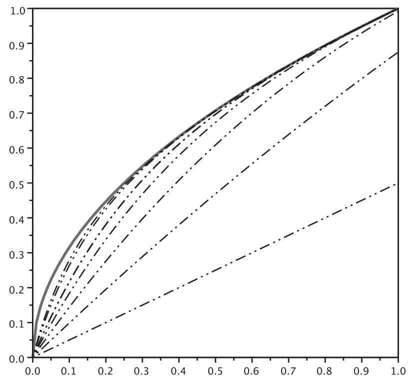

Figure 3.2.1. The first few polynomials in the sequence defined in (3.2).

2. For every function $f \in \mathcal{A}$, one has $|f| \in \overline{\mathcal{A}}$.

Indeed, let $\kappa \doteq \max_{x\in E}|f(x)|$. We can assume $\kappa \neq 0$. Then all functions $f_n(x) = p_n(f^2(x)/\kappa^2)$ lie in $\mathcal{A}$, because $\mathcal{A}$ is an algebra. Since $f^2(x)/\kappa^2 \in [0,1]$, the previous step yields the convergence

$$f_n(x) \;\to\; \sqrt{\frac{f^2(x)}{\kappa^2}} \;=\; \frac{|f(x)|}{\kappa},$$

uniformly for $x \in E$. Therefore $\frac{|f|}{\kappa} \in \overline{\mathcal{A}}$, and hence $|f| \in \overline{\mathcal{A}}$ as well.

3. We now apply the previous argument to the subalgebra $\overline{\mathcal{A}}$ and conclude that, if $f, g \in \overline{\mathcal{A}}$, then the functions

$$\max\{f,g\} \;=\; \frac{1}{2}(f+g+|f-g|)\,, \qquad \min\{f,g\} \;=\; \frac{1}{2}(f+g-|f-g|)$$

also lie in $\overline{\mathcal{A}}$.

4. For any two distinct points $y_1, y_2 \in E$ and any couple of real numbers a_1, a_2, there exists a function $f \in \mathcal{A}$ such that $f(y_1) = a_1$ and $f(y_2) = a_2$.

Indeed, by assumption there exists a continuous function $g \in \mathcal{A}$ such that $g(y_1) \neq g(y_2)$. Since $\mathcal{A}$ is an algebra and contains all constant functions, the function

$$f(x) \;\doteq\; a_1 + (a_2 - a_1)\frac{g(x) - g(y_1)}{g(y_2) - g(y_1)}$$

lies in $\mathcal{A}$ and satisfies our requirements.

5. Given any continuous function f, a point $y \in E$, and $\varepsilon > 0$, there exists a function $g_y \in \overline{\mathcal{A}}$ such that

$$g_y(y) \;=\; f(y), \qquad g_y(x) \;\le\; f(x) + \varepsilon \quad \text{for every } x \in E\,. \tag{3.3}$$

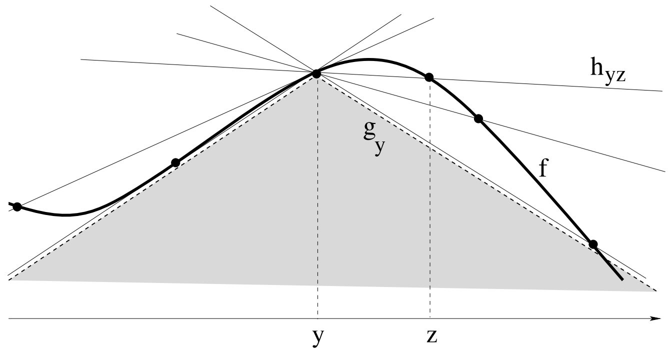

Figure 3.2.2. For a fixed y, taking the infimum of finitely many functions h_{yz} (here drawn as affine functions) we obtain a continuous function $g_y \leq f + \varepsilon$, with $g_y(y) = f(y)$.

Indeed (see Figure 3.2.2), by the previous step, for every point $z \in E$, there exists a function $h_{yz} \in \mathcal{A}$ such that $h_{yz}(y) = f(y)$ and $h_{yz}(z) = f(z)$.

Since f and h_{yz} are both continuous, there exists an open neighborhood V_z of z such that $h_{yz}(x) < f(x) + \varepsilon$ for every $x \in V_z$.

We can cover the compact set E with finitely many such neighborhoods: $E \subseteq V_{z_1} \cup \cdots \cup V_{z_m}$. Then the function

$$g_y(x) \doteq \min\left\{h_{yz_1}(x), \ldots, h_{yz_m}(x)\right\}$$

lies in $\overline{\mathcal{A}}$ and satisfies the conditions in (3.3).

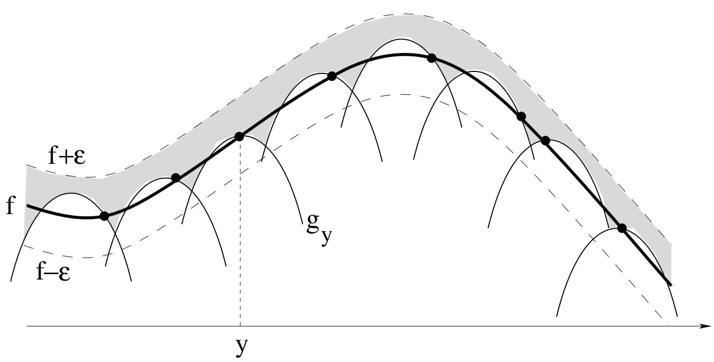

Figure 3.2.3. Taking the supremum of finitely many functions g_y we obtain a continuous function g such that $f - \varepsilon \leq g \leq f + \varepsilon$.

6. The closure of $\mathcal{A}$ is the entire space $\mathcal{C}(E)$.

Indeed, let $f \in \mathcal{C}(E)$ be any continuous function and let $\varepsilon > 0$ be given. For each $y \in E$, by the previous step there exists a function $g_y \in \overline{\mathcal{A}}$ (see

Figure 3.2.3) such that

$$g_y(y) = f(y), \qquad g_y(x) \leq f(x) + \varepsilon \quad \text{for every } x \in E\,.$$

By the continuity of f and g_y, there exists a neighborhood U_y of y such that

$$g_y(x) \geq f(x) - \varepsilon \qquad \text{for all } x \in U_y\,.$$

We now cover the compact set E with finitely many neighborhoods: $E \subseteq U_{y_1} \cup \cdots \cup U_{y_\nu}$. Then the function

$$g(x) \doteq \max\left\{g_{y_1}(x), \ldots, g_{y_\nu}(x)\right\}$$

lies in $\overline{\mathcal{A}}$ and satisfies

$$f(x) - \varepsilon \leq g(x) \leq f(x) + \varepsilon \quad \text{for all } x \in E\,.$$

□

A natural example of an algebra that satisfies all the assumptions in the Stone-Weiertrass theorem is provided by the polynomial functions.

Corollary 3.6 (Uniform approximation by polynomials). *Let E be a compact subset of $\mathbb{R}^n$. Let $\mathcal{A}$ be the family of all real-valued polynomials in the variables $(x_1, \ldots, x_n)$. Then $\mathcal{A}$ is dense in $\mathcal{C}(E)$.*

Indeed, the family of all real-valued polynomials in $(x_1, \ldots, x_n)$ is an algebra that contains the constant functions and separates points in $\mathbb{R}^n$. Hence, by Theorem 3.5 every continuous function $f : E \mapsto \mathbb{R}$ can be uniformly approximated by polynomials.

3.2.1. Complex-valued functions. A key ingredient in the proof of Theorem 3.5 was the fact that, if two real-valued functions f, g lie in a subalgebra $\mathcal{A}$, then $\max\{f, g\}$ and $\min\{f, g\}$ lie in $\overline{\mathcal{A}}$. Clearly, such a statement would be meaningless for complex-valued functions. In fact, in its original form the Stone-Weierstrass theorem is NOT valid for complex-valued functions.

In order to obtain an approximation result valid for functions $f : E \mapsto \mathbb{C}$, the main idea is to regard $\mathbb{C} = \mathbb{R} \oplus i\mathbb{R}$ as a two-dimensional space over the reals. In the following, given a compact metric space E, we shall denote by $\mathcal{C}_{\mathbb{R}}(E; \mathbb{C})$ the space of all continuous complex-valued functions on E, regarded as a vector space over the real numbers.

Theorem 3.7. *Let E be a compact metric space. Let $\mathcal{A}$ be a subalgebra of $\mathcal{C}_{\mathbb{R}}(E;\, \mathbb{C})$ that separates points and contains the constant functions. Moreover, assume that whenever $f \in \mathcal{A}$, then also the complex conjugate function $\bar{f}$ lies in $\mathcal{A}$. Then $\mathcal{A}$ is dense in $\mathcal{C}_{\mathbb{R}}(E;\, \mathbb{C})$.*

Proof. 1. By the assumptions, if $f \in \mathcal{A}$, then its real and imaginary parts

$$\mathrm{Re}(f) = \frac{f + \bar{f}}{2}, \qquad \mathrm{Im}(f) = \frac{f - \bar{f}}{2}$$

also lie in $\mathcal{A}$. Let $\mathcal{A}_0$ be the subalgebra of $\mathcal{A}$ (over the real numbers), consisting of all functions $f \in \mathcal{A}$ with real values. Applying the Stone-Weierstrass theorem to $\mathcal{A}_0$, we conclude that $\mathcal{A}_0$ is dense in $\mathcal{C}(E) = \mathcal{C}_{\mathbb{R}}(E;\, \mathbb{R})$.

2. Given any $f \in \mathcal{C}_{\mathbb{R}}(E;\, \mathbb{C})$, we write f as a sum of its real and imaginary parts $f = \mathrm{Re}(f) + i\ \mathrm{Im}(f)$. By the previous step, there exist two sequences of real-valued functions $g_n, h_n \in \mathcal{A}_0$ such that

$$g_n \to \mathrm{Re}(f), \qquad h_n \to \mathrm{Im}(f) \tag{3.4}$$

as $n \to \infty$, uniformly on E.

Consider the sequence $f_n \doteq g_n + ih_n \in \mathcal{A}_0 + i\mathcal{A}_0 = \mathcal{A}$. By (3.4), we have the uniform convergence $f_n \to f$. Hence $\overline{A} = \mathcal{C}_{\mathbb{R}}(E;\, \mathbb{C})$. □

Example 3.8 (Complex trigonometric polynomials). Let E be the unit circumference $\{x^2 + y^2 = 1\}$ in $\mathbb{R}^2$. Points on E will be parameterized by the angle $\theta \in [0, 2\pi]$. Let $\mathcal{A}$ be the algebra of all complex trigonometric polynomials:

$$p(\theta) = \sum_{n=-N}^{N} c_n e^{in\theta}, \tag{3.5}$$

where $N \geq 0$ is any integer and the coefficients c_n are complex numbers. It is clear that $\mathcal{A}$ is an algebra, contains the constant functions, and separates points. Moreover, $p \in \mathcal{A}$ implies $\bar{p} \in \mathcal{A}$ as well. By Theorem 3.7, the family of all these complex trigonometric polynomials is dense in $\mathcal{C}_{\mathbb{R}}(E;\, \mathbb{C})$.

Relying on the previous example, we now show that a real-valued, continuous periodic function can be uniformly approximated with trigonometric polynomials of the form

$$q(x) = \sum_{k=0}^{N} \alpha_k \cos kx + \sum_{k=1}^{N} \beta_k \sin kx\,. \tag{3.6}$$

Here $N \geq 1$ is any integer, while α_n, β_n are real numbers.

Corollary 3.9 (Approximation of periodic functions by trigonometric polynomials). *Let $f : \mathbb{R} \mapsto \mathbb{R}$ be a continuous function, periodic of period 2π. Then for any $\varepsilon > 0$ there exists a trigonometric polynomial q as in* (3.6) *such that*

$$|q(x) - f(x)| \leq \varepsilon \qquad \text{for all } x \in \mathbb{R}\,. \tag{3.7}$$

Proof. By assumption, $f(x+2\pi) = f(x)$ for every $x \in \mathbb{R}$. As shown in Example 3.8, there exists a complex trigonometric polynomial p of the form (3.5) such that

$$|p(x) - f(x)| \leq \varepsilon \qquad \text{for all } x \in \mathbb{R}. \tag{3.8}$$

Consider the complex coefficients $c_n = a_n + ib_n$, with $a_n, b_n \in \mathbb{R}$. Calling $q(x) \doteq \text{Re } p(x)$ the real part of p, we compute

$$q(x) = \sum_{n=-N}^{N} a_n \cos nx - \sum_{n=-N}^{N} b_n \sin nx = \sum_{n=0}^{N} \alpha_n \cos nx + \sum_{n=1}^{N} \beta_n \sin nx,$$

with

$$\alpha_0 = a_0, \qquad \alpha_n = a_{-n} + a_n, \qquad \beta_n = b_{-n} - b_n \quad \text{for } n \geq 1.$$

Since f is real-valued, by (3.8) we have

$$|q(x) - f(x)| = |\text{Re } p(x) - \text{Re } f(x)| \leq |p(x) - f(x)| \leq \varepsilon \quad \text{for all } x \in \mathbb{R}.$$

□

Remark 3.10. If the periodic function f is even, i.e., $f(x) = f(-x)$, then it can be approximated with a trigonometric polynomial of the form (3.6) with $\beta_k = 0$ for every $k \geq 1$, that is, with a finite sum of cosine functions.

If f is odd, i.e., $f(x) = -f(-x)$, then in (3.6) one can take $\alpha_k = 0$ for every $k \geq 0$. In other words, f can be approximated by a finite sum of sine functions.

3.3. Ascoli's compactness theorem

In a finite-dimensional space, by the Bolzano-Weierstrass theorem every bounded sequence has a convergent subsequence. On the other hand, as shown in Theorem 2.22 of Chapter 2, this compactness property fails in every infinite-dimensional normed space. For example, in the space $\mathcal{C}([0,1])$ the sequences of continuous functions $f_n(x) = x^n$ or $f_n(x) = \sin nx$ are bounded but do not admit any uniformly convergent subsequence. It is thus natural to ask: what additional property of the functions f_n can guarantee the existence of a uniformly convergent subsequence? An answer is provided by Ascoli's theorem, relying on the concept of *equicontinuity*.

Let E be a metric space. A family of continuous functions $\mathcal{F} \subset \mathcal{C}(E)$ is called **equicontinuous** if, for every $x \in E$ and $\varepsilon > 0$, there exists $\delta > 0$ such that

$$d(y,x) < \delta \quad \text{implies} \quad |f(y) - f(x)| < \varepsilon \tag{3.9}$$