

# 1

## Introduction

### 1.1 Turbulence

Turbulence is the last great unsolved problem of classical physics.<sup>1</sup> Although temporarily abandoned by much of the community in favor of particle physics, the current popularity of chaos and dynamical systems theory (as well as funding problems in particle physics) is now drawing the physicists back. During the interim and up to the present, turbulence has been avidly pursued by engineers.

Turbulence has enormous intellectual fascination for physicists, engineers, and mathematicians alike. This scientific appeal stems in part from its inherent difficulty – most of the approaches that can be used on other problems in fluid mechanics are useless in turbulence. Turbulence is usually approached as a stochastic problem, yet the simplifications that can be used in statistical mechanics are not applicable – turbulence is characterized by strong dependency in space and in time, so that not much can be modeled usefully as a simple Markov process, for example. The nonlinearity of turbulence is essential – linearization destroys the problem. Many problems in fluid mechanics can be approached by supposing that the flow is irrotational – that is, that the vorticity is zero everywhere. In turbulence, the presence of vorticity is essential to the dynamics. In fact, the nonlinearity, rotationality, and the dimensionality interact dynamically to feed the turbulence – hence, to suppose that a realization of the flow is two-dimensional also destroys the problem. There is more, but this is enough to make it clear that one faces the turbulence problem stripped of the usual arsenal of techniques, reduced to hand-to-hand combat. One is forced to find unexpected chinks in its armor almost by necromancy, and to fabricate new approaches from whole cloth. This is its fascination.

At the same time, turbulence is of the greatest practical importance. The turbulent transport of heat, mass, and momentum is usually some three orders of magnitude greater than molecular transport. Turbulence is responsible for the vast majority of human energy

<sup>1</sup> Remarks of this sort have been variously attributed to Sommerfeld, Einstein, and Feynman, although no one seems to know precise references, and searches of some likely sources have been unproductive. Of course, the allegation is a matter of fact, not much in need of support by a quotation from a distinguished author. However, it would be interesting to know when the matter was first recognized. In this connection, similar sentiments were expressed by Horace Lamb in his *Hydrodynamics*, beginning in the second edition in 1895, and continuing through the sixth (and last) edition in 1932. We are indebted for this reference to Julian Hunt, citing its use by George Batchelor in his book *The Life and Legacy of G. I. Taylor*, Cambridge University Press, 1996.

consumption, in automobile and aircraft fuel, pipeline pumping charges, and so forth. It is responsible for the wind chill factor. In the atmosphere and ocean it is responsible for the transport of gases and nutrients and for the uniformization of temperature that make life on earth possible. For example, oxygen and carbon dioxide are not produced in the same places – oxygen comes largely from the equatorial rain forests and carbon dioxide is manufactured in industrial and urban centers such as New York City. Some mechanism is necessary to bring the carbon dioxide to Brazil, and the oxygen to the Big Apple. Radiation from the sun heats the surface of the earth; something is necessary to transfer the heat quickly and uniformly to the atmosphere where we can benefit from it. Without turbulence our speedy demise would be a race between frying our feet and freezing our heads, gasping in an atmosphere with too much or too little oxygen and/or carbon dioxide.

These practical aspects are, of course, responsible for most of the funding for turbulence research. It is absolutely essential as a design tool to be able to predict accurately the forces on and heat transfer from aircraft and automobiles. For regulatory purposes it is essential to be able to predict the results of siting of power plants and incinerators under various synoptic conditions. Manufacturers cry out for the ability to predict fluctuations in dopant distribution in the billets of silicon from which chips are formed. The military is concerned about the information loss in battlefield communication links induced by index of refraction fluctuations due to thermal turbulence. The list is endless.

From five centuries of observation and experiment, in many ways a reasonable physical understanding of turbulence has emerged. It is no longer a complete mystery. We can cite many simple physical arguments that shed light on common situations. When it comes to accurate predictions, however, we are in trouble. Aircraft manufacturers, for example, want accuracy corresponding roughly to the effect of adding one passenger to a Boeing 747. Automobile manufacturers want accuracy corresponding to the effect of adding one outside rear-view mirror. Regulatory agencies want assurances of comparable accuracy before going to court. Although our ability to calculate is improving constantly, we are not yet close to this level of accuracy.

Direct numerical simulation is not a realistic possibility in most cases of practical importance. In the foreseeable future, the cost of such simulation will remain far beyond our means, and will be limited to very low Reynolds numbers and simple geometries. In any event, simulation by itself does not bring understanding.

In a given practical problem, there may be many things that one wishes to know. The most common goals of computation are the mean forces and/or the mean heat transfer at various locations in the flow. These involve knowledge of second order quantities, the mean fluxes of momentum and heat. That is, the mean flux of  $j$ -momentum through a surface with a normal in the  $i$ -direction is  $-\rho\langle u_i u_j \rangle$ , where  $u_i$  is the fluctuating turbulent velocity,  $\langle \cdot \rangle$  denotes an average, and  $\rho$  is the mean density. The flux of heat into an  $i$ -surface is  $-\rho c_p \langle u_i \theta \rangle$ , where  $c_p$  is the specific heat at constant pressure, and  $\theta$  is the fluctuation in temperature. Both involve mean values of products of no more than two fluctuating quantities. Computation of index of refraction fluctuations in the atmosphere involves knowledge of the probability densities of fluctuating quantities, but an assumption about the form of the densities, plus knowledge of the variances, is usually enough. Hence, again, second order quantities are sufficient. A similar statement can be made about

the dopant fluctuations in the silicon billet. There are more complex questions, however, that require more complex information. For example, suppose we wish to simulate the fluctuating pressure field on a panel, due to the presence of a turbulent boundary layer over the surface, perhaps to predict the spurious noise field generated on a sonar dome. This requires much more sophisticated modeling of the field.

It was in an effort to answer such deeper questions, that depend on a knowledge of the structure of the flow, that we embarked on the work described here. As we shall see below, many turbulent flows are characterized by considerable structure, and in particular by characteristic recurrent forms that are collectively called coherent structures. These are energetically dominant in many flows. We feel that, for flows in which these structures are dominant, it should be possible to build a relatively realistic, low-dimensional model of the flow by keeping only the dominant coherent structures, and simulating the effect of the smaller, less energetic, apparently incoherent part of the flow in some way. In this book we describe our tentative steps in this direction.

## 1.2 Low-dimensional models

Perhaps the first attempts to bring a dynamical systems perspective to turbulence studies were those of Landau (e.g. [204]) and Hopf [163]. They suggested that the continuous Fourier spectrum of temporal frequencies typical of turbulence might be produced via bifurcations occurring as the Reynolds number is increased (which Hopf, betraying his background, called  $\mu$  rather than  $R_e$ ). They envisaged a sequence in which at first periodic and then quasiperiodic attractors with increasing numbers of independent frequencies were created. In the language of modern dynamical systems theory, we would say that the resulting fluid flow corresponds to a *phase flow* on an  $n$ -dimensional torus in the state (or phase) space of the dynamical system. Hopf even constructed a model problem which exhibited just such a bifurcation sequence: what we might call a “route to chaos,” except that we now realize that quasiperiodic flows are not strongly chaotic, since these solutions do not depend sensitively on initial conditions. Perhaps more significantly, Hopf also proposed that “to the flows observed in the long run after the influence of the initial conditions has died down there correspond certain solutions of the Navier–Stokes equations. These solutions constitute a certain manifold  $\mathcal{M}(\mu)$  in phase space invariant under the phase flow. Presumably owing to viscosity  $\mathcal{M}(\mu)$  has a finite number  $\mathcal{N}(\mu)$  of dimensions.” ([163], p. 305.) Hopf envisaged a finite-dimensional attractor.

Some twenty years after Hopf’s paper, Ruelle and Takens [322] built on this suggestion. They observed that the quasiperiodic flows proposed by Landau and Hopf are not *structurally stable* and so would be expected to appear only in unusual circumstances. Drawing on the qualitative theory of (finite-dimensional) dynamical systems, which Anosov, Smale, Arnold, and others had extensively developed in the meantime, they gave an example of a structurally stable “strange” attractor that can appear after two or three quasiperiodic bifurcations, and so can live on a torus of only four dimensions (subsequently this was reduced to three: Newhouse *et al.* [258]). In connection with one of our themes, a footnote

in their introduction is also noteworthy: “If a viscous fluid is observed in an experimental setup which has a certain symmetry, it is important to take into account the invariance of [the dynamical system] under the corresponding symmetry group.” ([322], p. 168.) Ruelle gives an interesting account of the genesis of and tribulations encountered by their paper in [321].

In none of this work was a clear connection made between a particular fluid flow modeled by the Navier–Stokes equations with specific boundary conditions, and the “abstract” dynamical systems which exhibited quasiperiodic or strange attractors. However, unknown to Ruelle, Takens, and virtually all other mathematicians and physicists in 1971, Lorenz [217] had provided an example almost ten years before. A meteorologist and a former student of the dynamical systems pioneer, George Birkhoff, Lorenz was interested in the problem of weather prediction. He took a drastic truncation to three Fourier modes of the coupled Navier–Stokes and heat equations for Boussinesq convection in a two-dimensional layer (a Rayleigh–Bénard problem) and investigated them numerically and analytically. He found strong evidence for a strange attractor, unfortunately far beyond the Rayleigh number range in which his truncation was reasonable. Nonetheless, after its general discovery in the early 1970s, due largely to the mathematician Jim Yorke, Lorenz’s paper has had an enormous influence. In Chapter 6 we give a sketch of what is now called the Lorenz attractor.

The events which first began to persuade fluid dynamicists that low-dimensional models and strange attractors might have some practical interest for them were probably the experiments of Gollub, Swinney, and their colleagues in the mid 1970s (see Swinney and Gollub [361]). Working with small, closed fluid systems, and especially with the Taylor–Couette flow between counter-rotating cylinders and thermal convection in small boxes, they found striking experimental evidence of sequences of bifurcations leading to “low-dimensional” chaos as the Taylor and Rayleigh numbers respectively were raised modestly above the initial onset of linear instability. Power spectra displaying jumps from two or three frequency quasiperiodic motions to broad band chaos were measured. Subsequently it became possible to link some of these results tightly with bifurcation analyses of the governing equations, particularly in the Taylor–Couette problem (see Golubitsky and Stewart [135], Golubitsky and Langford [133], Golubitsky *et al.* ([136], case study 6), Iooss *et al.* [75, 95, 173, 174], and Laure and Demay [206], for example). There is an enormous literature on this system: in his 1994 review, Tagg [363] estimates nearly 2000 papers while citing some 350 himself. Chossat and Iooss have published a book on the mathematical aspects of the problem [76].

In some cases, previously unknown classes of solution were predicted which were subsequently observed experimentally (e.g. Andereck *et al.* [6], Tagg *et al.* [364]). Again, the symmetries of the experimental apparatus were crucial in this. It is probably fair to say that the tools and viewpoint of dynamical systems theory are now acknowledged to have a useful rôle to play in the study of such closed fluid systems, in which relatively few spatial modes are active. These methods, including invariant manifold techniques, bifurcation theory, and the unfolding of degenerate singularities, have joined more classical asymptotic and perturbation methods for the study of hydrodynamic stability and “weakly” nonlinear or pre-turbulent interactions.

In this book we want to take a further tentative step. We propose that low-dimensional dynamical systems can also provide models for, and hence contribute to the understanding of, certain fully developed, open turbulent flows. As we remarked in the Preface, our “low” is not so low in dynamical systems terms: we are thinking of sets of 10–100 ordinary differential equations (ODEs). But in the fluid mechanical context this is *very* low and we clearly cannot expect to reproduce fine-scale spatial details of the flow. Consequently, to be sure of capturing the key behaviors, we will have to pay particular attention to the manner in which the fluid velocity field in physical space is represented in the phase space of the dynamical system. We shall focus on flows with predominant coherent structures, and use the proper orthogonal or Karhunen–Loève decomposition (POD) to extract, from experimental or simulated ensembles of data, those “modes” or empirical eigenfunctions that carry the greatest kinetic energy on average. This procedure will provide us with the basis for a sequence of subspaces, of increasing dimension, onto which the Navier–Stokes (or other) equations can be projected by Galerkin’s method to yield sets of ODEs. In this procedure we represent the fluid velocity field by a superposition of the empirical spatial modes multiplied by (as yet unknown) time-dependent coefficients. Substituting this representation into the governing equations and taking the inner product with each basis function in turn yields a set of nonlinear ODEs for the modal coefficients. These entirely deterministic dynamical systems will be the foundations for our low-dimensional models.

Our main goal is not to reproduce accurately the results of a direct numerical simulation with fewer, more efficient modes. The fact that such empirical basis functions are adapted to a particular flow geometry and Reynolds number, and are only available at the end of extensive data collection and computation, probably makes them a poor choice for efficient simulations in any case. Rather we are interested in *understanding* the fundamental mechanisms of turbulence generation in “simple” flows such as shear layers, jets, wakes, and boundary layers. In this quest for understanding we often want to reduce the dimension of our models to a minimum. Thus, even with the optimal bases of the POD, our truncations are typically so severe that a bare projection is unsatisfactory and some sort of additional modeling is needed to account for neglected modes and/or spatial locations. Such modeling might include relatively simple “eddy viscosity” energy transfer of the sort proposed by Heisenberg and Smagorinsky (see Batchelor [32] or Tennekes and Lumley [368]) as well as models to account for slow variations of the mean shear which drives the turbulent fluctuations in flows such as boundary and shear layers, due to the turbulence itself. Ideally, and in greater generality, we envisage the introduction of a probabilistic element to our deterministic ODEs to reproduce the conditional probability measures that describe the activity of the neglected modes as a function of the state of those modes included in the model. Very little appears to be known about this issue, but we have encountered and partially resolved a crude version of it in our treatment of the outer part of the boundary layer in models of the wall region.

After determination of a “good” subspace, projection of the governing equations, and modeling to account for neglected modes, we have a set of ODEs, for an understanding of which we can appeal to the methods of dynamical systems theory along with other, more widely known mathematical tools. If done properly, the projection and modeling preserve the underlying symmetries of the fluid flow and of the governing equations and

boundary conditions. Such symmetries may include spanwise translations and reflections for a shear layer or a boundary layer on an infinite flat plate, and rotations and reflections for a circular jet or wake. Thus the ODEs will exhibit a corresponding symmetry: in the language of dynamical systems theory, they will be *equivariant under some group*  $\Gamma$ , and we have to take this into account in studying the bifurcations and other dynamical behavior of the system. Behavior that is structurally unstable and hence rare in general may be stable and relatively prevalent for such  $\Gamma$ -equivariant systems. We have already mentioned heteroclinic attractors in the Preface, and the reader will find several more examples later in the book.

The result of our dynamical systems analyses of the low-dimensional models is a (partial) understanding of the structure of solutions in phase space and in particular of attracting sets and how they change through bifurcations as external and modeling parameters are varied. The final tasks are to map those results back into physical space, reconstructing the space–time velocity field of the fluid flow from the empirical basis functions and their time-dependent coefficients, to compare the resulting instantaneous and averaged quantities with experiments, and to translate the understanding achieved in state space into insights about the fluid flow itself.

This is the general strategy we propose: find good basis functions for the turbulent flow in question, model to account for neglected effects, project the governing partial differential equations onto a low-dimensional subspace spanned by the most energetic modes, analyze the resulting low-dimensional model, and finally return to the physical domain to interpret that analysis. As we see in Chapter 2, not all turbulent fluid flows are energetically dominated by coherent structures, and so the approach we describe here is far from offering a complete solution to “the problem of turbulence.” We believe, nonetheless, that it provides one more approach and a set of new tools, or even weapons, for the unequal combat referred to in the introduction to this chapter.

### 1.3 The contents of this book

As noted in the Preface, the book has four parts. The first two, which constitute well over half the book, are fairly general in nature. We introduce key ideas from fluid mechanics, turbulence theory, and dimension reduction methods in the first five chapters, and from dynamical systems theory in the following four. Turbulence experts can probably skip pieces of Part One, and dynamicists can certainly skip most of Part Two, but in both places readers may find new viewpoints recommended and unfamiliar connections drawn. We hope that these parts of the book will be of fairly lasting and general interest. The remaining two parts are more specific and more speculative, for in them we focus on our own work on the turbulent boundary layer and on other attempts to derive low-dimensional models for turbulent and transition flows. We offer our own work mainly in the spirit of an extended example, since it allows us to discuss and illustrate difficulties and limitations as well as successes of the approach. We are far from claiming a complete understanding of boundary layer turbulence via our models, but we hope that the reader who accompanies us to the end of Chapter 13 will agree that some new things have been said.

In the remaining four chapters of this first part, we give some background on turbulence, describe coherent structures from an experimental viewpoint and summarize some of the major findings relevant to shear dominated flows. In Chapter 2 we sketch some experimental methods by which coherent structures in developed turbulent flows may be found and characterized, and describe their relation to instabilities of simpler laminar and transition flows. We also review the “classical” approach to turbulent flows, via the averaged Navier–Stokes equations and careful order-of-magnitude and scaling estimates. We discuss in some detail the cases of turbulent mixing and boundary layers (the main illustrative application of our approach is to the latter). We close with a brief preview of how coherent structures might appear as attractors in simple dynamical systems.

Chapter 3 is devoted to the proper orthogonal decomposition. We provide the basic mathematical results, with enough elements of their proofs to illustrate both the scope and limitations of representing turbulent fields by finite- (low-) dimensional projections. We pay particular attention to the influence both of symmetries, in the physical flow and in the particular data ensembles used, and of the ensemble averaging on which the method is founded, on the sets of basis functions that it produces. We also describe the relation of these empirical modes to certain other statistically based techniques for prediction and analysis of turbulence, such as stochastic estimation.

In Chapter 4 we discuss the Galerkin method, and show how the Navier–Stokes equations, or in general any evolution equation, may be projected onto a finite-dimensional subspace, and in particular onto a subspace spanned by empirical modes, to produce a finite set of ODEs. We also discuss various “modeling” issues such as those mentioned in Section 1.2.

Chapter 5 introduces recent work on the balanced proper orthogonal decomposition, and on the balanced truncation method for linear systems on which it is based. This method is particularly useful for systems with control inputs based on observation of specific flow quantities, and we provide examples to illustrate its superiority to the proper orthogonal decomposition in such cases.

The second part of the book is a mini-treatise on dynamical systems theory. Since we are concerned only with low-dimensional models, we restrict ourselves to finite-dimensional ODEs and iterated maps. In Chapter 6 we sketch the main ideas and tools, including linearization, invariant manifolds, structural stability, the center manifold theorem, normal forms, and local and global bifurcation theory. We end the chapter with a discussion of attractors, the main example being the strange attractor of Lorenz. Throughout this and the remaining chapters in this second part, we illustrate the theory with many simple and *very* low-dimensional examples.

Chapter 7 deals with symmetries, bifurcations, and local and global dynamical behavior of equivariant ODEs, leading up to an important example derived from spatial translation and reflection invariance, which can be understood in the context of Fourier mode representations of traveling waves. This example, an  $O(2)$ -equivariant normal form for the interaction of wavenumbers in the ratio 1:2, is a four-dimensional ODE possessing heteroclinic attractors, which, while not strange, have an interesting structure which seems relevant in models of many fluid flows with symmetries. The chapter ends with a

brief description of how the POD method can be extended to provide empirical modes that represent uniformly translating structures (traveling modes).

In Chapter 8 we exercise our new methods on a simple model problem: the one-space-dimensional Kuramoto–Sivashinsky partial differential equation. We find the  $O(2)$ -equivariant normal form of Chapter 7 buried in this system. In the final chapter of this part, Chapter 9, we consider the effects of stochastic and other symmetry-breaking perturbations on systems with heteroclinic cycles.

The third part of the book is devoted to a description of attempts by ourselves and our students and colleagues, to apply our strategy to the wall region of the turbulent boundary layer. Most of Chapter 10 contains discussions of the Galerkin projection and modeling issues, introduced in Chapter 4, in the specific case of the wall region. We describe the choice of specific subspaces and the resulting hierarchy of nested systems of increasing dimension that results as more modes are included. The chapter contains a description of the various symmetries that the low-dimensional model ODEs inherit from those of the boundary layer itself, and ends with extensive discussions of the validity of the models for the mean flow and losses to neglected modes.

In Chapter 11 we bring dynamical systems techniques to bear on the model ODEs and provide relatively complete analyses of the bifurcations and dynamical behavior of the boundary layer models. We illustrate and supplement our analyses with numerical simulations of models of various dimensions and, here and in Chapter 10, we offer a critical interpretation of the results, showing how the use of empirical basis functions can sometimes lead to paradoxical effects. The chapter includes reconstructions of the fluid velocity fields and interpretations of our findings in phase space, in terms of the turbulent flow itself.

The first of the two chapters of Part Four contains brief reviews of work by other groups in which the same general approach is taken. We do not pretend to give a complete survey of this rapidly developing field, but the examples of “laboratory” open flows that we have chosen, including jets, wakes, and transition in boundary layers, illustrate that our methods have wide applicability. Related ideas have been and continue to be used in the meteorological community (cf. [216]) – the work of Farrell and Ioannou is an interesting case in point [105–107] – and there are clearly applications to many other problems involving the dynamics of spatio-temporal patterns. In this second edition we have added references to some recent work, including new sections on time periodic flows in internal combustion engines and other applications (12.6 and 12.7).

In the closing Chapter 13 we speculate in broader terms on the place and uses of low-dimensional models among the many other approaches to turbulence. It seems clear that such models offer new *understanding* of turbulence generation involving coherent structures, and so contribute to the intellectual challenge alluded to in the second paragraph of the present chapter. Can they also be of help in answering technological questions such as those mentioned towards the end of Section 1.1? A particular interest of our own is in the use of such models in formulating strategies for the active control of turbulence and, in addition to the material in Chapter 5, we provide a brief description of our ideas at the end of Chapter 12. In Chapter 13 we also mention a number of other recent developments that are related to our story, including mathematical ideas such as inertial manifolds and other reduction methods which offer new approaches to the Navier–Stokes equations.

## 1.4 Notation and mathematical jargon

By now the reader knows that in this book we propose the application of ideas in the qualitative theory of dynamical systems to the description and analysis of turbulent flows. While qualitative theory had its beginnings in Poincaré's studies of problems in celestial mechanics about one hundred years ago [282], it was soon thereafter hijacked by pure mathematicians and only in the last ten to twenty years has it begun to see broad applications in the sciences and engineering. The explosion of interest in "chaos theory," encouraged by books such as Gleick's [128], has certainly sparked a general awareness that there are new ideas and methods out there, but we recognize that many of the basic concepts and technical issues may remain mysterious for potential users, including the intended readership of this book. Rather than try to skate over what may be unfamiliar mathematical material for some readers, we have tried to introduce it with simple examples drawn from the world of low-dimensional ordinary differential equations. By working such examples in some detail, we hope to leave our readers in a position to fill in missing steps in more complicated cases and to tackle new ones that may arise in their own research. But this is emphatically *not* a dynamical systems textbook: we do not state, much less prove, even the most basic theorems in the field, and those formal definitions that are included are given in passing, usually indicated by *italics*.

Even with an approach based on examples, so foreign to "pure" dynamical systems theorists, we cannot avoid introducing and using a modicum of mathematical jargon. Our defense of this is twofold: (1) we believe that, once learned, the symbolism of dynamical systems theory, largely drawn from topology, makes the precise description of key ideas such as invariant manifolds and attractors much simpler and more compact than is possible with the English language alone, and (2) we hope that this book might be the beginning of an exploration of the current research literature, in which case the symbolism will have to be mastered anyway. After each new excess of jargon in the text, we try to give a (rough) characterization in words, and we encourage readers who are repelled by abstract formulae to clench their teeth and read on to get to the examples and pictures.

Nonlinear analysis is built on linear analysis and, to avoid doubling the length of this book, we must assume some familiarity with the fundamental ideas of finite-dimensional linear vector spaces, spanning sets of vectors, bases, norms, inner products, linear subspaces, eigenvalue problems, and the like. Similarly, one of the major applications of this beautiful theory is to the solution of linear ordinary differential equations, and we assume a basic knowledge of that as well. Books such as Strang's *Linear Algebra and its Applications* [358] or Boyce and DiPrima's *Elementary Differential Equations and Boundary Value Problems* [56] provide the necessary background. More geometrically oriented introductions to nonlinear ordinary differential equations are Hirsch *et al.*'s *Differential Equations, Dynamical Systems and an Introduction to Chaos* [156], Arnold's *Ordinary Differential Equations* [15], and Glendinning's *Stability, Instability and Chaos* [129], all of which are written from a more mathematical viewpoint. The last of these is a good introduction to many of the modern concepts presented in Part Two of the present book.

Here, to prepare for the onslaught, we recall some of the mathematical notations we shall use. First there are the standard names for some commonly encountered spaces:

$\mathbb{R}^n$ :  $n$ -dimensional real Euclidean space, the elements of which are vectors  $\mathbf{x} = (x_1, x_2, \dots, x_n)$ , with each  $x_j$  a real number. The real line  $\mathbb{R}^1$  is simply written  $\mathbb{R}$ .

$\mathbb{C}^n$ :  $n$ -dimensional complex Euclidean space; as above, but each  $x_j$  is a complex number.  $\mathbb{C}^1$  is written  $\mathbb{C}$ .

We normally denote vectorial quantities by boldface letters  $\mathbf{x}$  and scalar quantities by italic letters  $x$ . Single bars  $|\cdot|$  denote the Euclidean norm or absolute value of whatever is inside them; they also denote the modulus in the case of a complex number. Other norms are generally indicated by double bars:  $\|\cdot\|$ . We occasionally use the *supremum norm*, written  $\sup|\mathbf{x}|$ , which indicates the least upper bound. If  $A$  is a subset of  $\mathbb{R}$ , the number  $M$  is an upper bound for  $A$  if  $a \leq M$  for all  $a$  in  $A$ . When  $M$  is the smallest such number, it is the *least upper bound*. The *infimum*  $\inf|\mathbf{x}|$  or *greatest lower bound* is defined analogously.

The *inner product* (also *scalar* or *dot product*) of two elements  $\mathbf{u}, \mathbf{v}$  in  $\mathbb{R}^n$  or  $\mathbb{C}^n$  is given by:

$$(\mathbf{u}, \mathbf{v}) = \sum_{i=1}^n u_i v_i^* = \mathbf{v}^* \mathbf{u},$$

where in the second expression  $*$  denotes the complex conjugate, and in the third expression  $*$  denotes complex conjugate transpose.

A set  $V$  is *open* if and only if for every point  $x \in V$  there is a neighborhood  $B_x$  of  $x$  with  $B_x$  contained in  $V$ . A set  $U$  is *closed* if and only if for each point  $y$  not in  $U$  there is a neighborhood  $B_y$  of  $y$  entirely disjoint from  $U$ . Alternatively, a set is closed if and only if it contains all its limit points. Examples are given directly below.

$[a, b]$ : the closed interval of the real line  $\mathbb{R}$ , delimited by the points  $a < b$ : all points  $x$  satisfying  $a \leq x \leq b$ . A curved parenthesis denotes that the endpoint is not included, thus  $(a, b)$  denotes the open interval ( $a < x < b$ ) and  $(a, b]$  the half open interval  $a < x \leq b$ . This notation extends to higher dimensions; thus  $[0, 1] \times [0, 1]$  or  $[0, 1]^2$  denotes the (closed) unit square in  $\mathbb{R}^2$  with corners at  $(0, 0)$ ,  $(1, 0)$ ,  $(1, 1)$ , and  $(0, 1)$ ; here  $\times$  means the direct product.

$L^2$ : the (Hilbert) space of square integrable real or complex-valued functions, an example of an infinite-dimensional inner product space. Often the domain of definition is indicated in parentheses: thus  $L^2([0, 1])$  denotes the space of functions defined over the unit interval  $0 \leq x \leq 1$ . Square integrable means that the functions  $f(x)$  belonging to  $L^2([0, 1])$  satisfy

$$\|f\| = \left[ \int_0^1 |f(x)|^2 dx \right]^{\frac{1}{2}} < \infty. \quad (1.1)$$

In general the integral is taken over the domain of definition,  $\Omega$ . The boundary of the domain is customarily written as  $\partial\Omega$ . Note that  $L^2$  is an inner product space, the inner product being defined by

$$(f, g) = \int_{\Omega} f(x)g^*(x)dx, \quad (1.2)$$