Introduction

Since Poincaré initiated the qualitative study of solutions of differential equations at the end of the nineteenth century, the theory of dynamical systems has had an important rôle to play in our understanding of many physical models.

However, until relatively recently the application of such ideas was mainly restricted to finite-dimensional systems, such as those that arise in the study of ordinary differential equations (ODEs) or iterated low-dimensional maps. It is only within the past two decades that similar techniques have been systematically applied to the infinite-dimensional systems that arise from partial differential equations (PDEs).

This book develops the dynamical systems approach to a certain class of PDEs, dissipative parabolic equations, and investigates their asymptotic behaviour by means of an object called the global attractor.

Provided that the function f is sufficiently smooth (see Chapter 2), the solutions of a finite set of coupled ODEs,

$$\dot{x} = f(x), \qquad x \in \mathbb{R}^m, \tag{1}$$

($\dot{x}$ is short for dx/dt) give rise to a dynamical system on the finite-dimensional phase space $\mathbb{R}^m$. If $x(t; x_0)$ is the solution of (1) with $x(0) = x_0$, then we can define a solution operator $T(t) : \mathbb{R}^m \to \mathbb{R}^m$ by

$$T(t)x_0 = x(t; x_0).$$

The dynamical system generated by (1) is specified by the pair

$$(\mathbb{R}^m, \{T(t)\}_{t\in\mathbb{R}}) \tag{2}$$

of the phase space $\mathbb{R}^m$ and the family of solution operators $\{T(t)\}_{t\in\mathbb{R}}$. We sometimes abbreviate (2) to $(\mathbb{R}^m, T(t))$.

Although the inclusion of the phase space in the definition might seem somewhat pedantic, in the case of PDEs the specification of an appropriate phase space is a large part of the resolution of the problem. Since this phase space will be an infinite-dimensional space of functions, we will need various tools from functional analysis, and these are developed in Chapters 1–5.

We illustrate some ideas with the simple example of the heat equation on a one-dimensional domain,

$$u_t = u_{xx}, \qquad u(0) = u(\pi) = 0. \tag{3}$$

(The notation u_t is shorthand for $\partial u/\partial t$, and u_{xx} is shorthand for $\partial^2 u/\partial x^2$.)

To look at this equation a little more closely, we will expand the solution $u(x,t)$ of (3) in terms of a Fourier series on $[0,\pi]$. The boundary conditions mean that we need only sine terms, so we can write

$$u(x,t) = \sum_{n=1}^{\infty} c_n \sin nx. \tag{4}$$

Substituting this expansion into (3) and comparing coefficients leads to an infinite set of ODEs for the $\{c_n\}$:

$$\frac{dc_n}{dt} = -n^2 c_n. \tag{5}$$

This is one way in which we can think of (3) as giving rise to an "infinite-dimensional" problem.

However, it is not immediately clear what constraints we ought to put on the variables $\{c_n\}$ to specify our phase space completely. One possible choice would be to restrict to solutions for which the energy

$$E = \int_0^\pi |u(x)|^2 \, dx$$

is finite. Since the functions $\{\sin nx\}$ are orthogonal, we can calculate E in terms of the $\{c_n\}$:

$$E = \frac{\pi}{2} \sum_{n=1}^{\infty} |c_n|^2.$$

In this way finite-energy solutions correspond to choices of Fourier coefficients that are square summable:

$$\sum_{n=1}^{\infty} |c_n|^2 < \infty. \tag{6}$$

Although this choice of "finite-energy" solutions is a natural one, it does pose a problem. If we are prepared to treat any function $u(x)$ that corresponds to a set of coefficients satisfying (6), we end up with the Lebesgue space of square integrable functions $L^2(0, \pi)$,

$$L^2(0, \pi) = \left\{ u : \int_0^{\pi} |u(x)|^2 \, dx < \infty \right\}.$$

Now, there are functions in $L^2(0, \pi)$ (or, equivalently, functions whose Fourier coefficients satisfy (6)) that are not continuous, let alone twice differentiable. Although (5) makes sense for the coefficients, the original equation (3) does not make sense for the corresponding function u given by (4).

If we are not going to exclude certain choices of $\{c_n\}$ that satisfy (6) but do not correspond to twice differentiable functions (which would produce a very convoluted definition of our phase space), then we need a way of understanding (3) even if $u(x, t)$ is not twice differentiable. The idea, loosely speaking, is that if

$$u = \sum_{n=1}^{\infty} c_n \sin nx$$

then we can *define* u_{xx} by

$$u_{xx} = \sum_{n=1}^{\infty} -n^2 c_n \sin nx \tag{7}$$

even if the series in (7) does not converge to a function in any classical sense. The distribution derivative and the Sobolev spaces of functions whose derivatives are in L^2, covered in Chapter 5, offer one way of doing this rigorously.

Related is the concept of a weak solution, which essentially allows the classical derivatives we expect in (3) to be replaced with such "generalised derivatives" (they are actually weaker than this!). This idea is introduced in Chapters 6 and 7, which deal with Poisson's equation,

$$-\Delta u = f, \qquad u|_{\partial\Omega} = 0$$

("weak in space") and a linear parabolic equation,

$$\frac{\partial u}{\partial t} + \Delta u = f(x,t), \qquad u|_{\partial\Omega} = 0$$

("weak in space and time").

We can interpret (3) as a dynamical system for the coefficients $\{c_n\}$, but we can also interpret it as a dynamical system on $L^2(0,\pi)$. In this case we treat the spatial and temporal dependence of $u(x,t)$ in fundamentally different ways. If we take a "snapshot" of $u(x,t)$ at a particular time τ, then the result is a function of x alone, $u(x,\tau)$. If this is an element of L^2 for each τ then the evolution of $u(x,t)$ in time traces out a trajectory in L^2. In later chapters we make almost exclusive use of this interpretation, suppressing the x dependence and writing things such as "$u(\tau) \in L^2$" as a convenient shorthand. This gives a second, and more useful, sense in which we can understand (3) as an infinite-dimensional problem.

If we are going to use (3) to define a dynamical system on $L^2(0,\pi)$, then we need to obtain existence and uniqueness of solutions, and in particular we have to be sure that if the initial condition $u(x,0)$ is an element of $L^2(0,\pi)$ then so is $u(x,t)$. To check this for (3) is simple, since we can use (5) to find the solution $u(x,t)$ explicitly. If

$$u(x,0) = \sum_{n=1}^{\infty} c_n \sin nx,$$

then (5) gives

$$u(x,t) = \sum_{n=1}^{\infty} c_n e^{-n^2 t} \sin nx.$$

Assuming that $u(x,0) \in L^2$, so that $\sum_{n=1}^{\infty} |c_n|^2\,dx < \infty$, we want to check that

$$\sum_{n=1}^{\infty} |c_n|^2 e^{-2n^2 t} < \infty. \tag{8}$$

We now come across another potential problem. It is easy to see that if $t > 0$ then the sum in (8) is finite. In fact, we have

$$\sum_{n=1}^{\infty} |c_n|^2 e^{-2n^2 t} < e^{-2t} \sum_{n=1}^{\infty} |c_n|^2,$$

and so the energy (which is proportional to this expression) decays to zero as $t \to \infty$. [Since

$$\frac{\pi}{2}\sum_{n=1}^{\infty}|c_n(t)|^2 = \int_0^{\pi}|u(x,t)|^2\,dx \equiv \|u(x,t)\|_{L^2}^2,$$

we could also write this in terms of the function $u(x,t)$ as

$$\|u(x,t)\|_{L^2} \le e^{-t}\|u(x,0)\|_{L^2}, \tag{9}$$

showing that $u \to 0$ in the sense of L^2 convergence.]

However, if we consider the case $t < 0$, then the exponentials in the sum, which were responsible for the dissipation of energy as $t \to \infty$, now cause the coefficients to increase dramatically with n. So dramatically, in fact, that if, for example

$$u(x,0) = \sum_{n=1}^{\infty}\frac{1}{n}\sin nx$$

(this has finite energy $\pi^3/12$), then

$$\sum_{n=1}^{\infty}\frac{1}{n^2}e^{-2n^2 t} \tag{10}$$

does not converge for any $t < 0$. The dissipation, which leads to "nice" behaviour as $t \to \infty$, is directly responsible for the blowup of solutions in $t < 0$.

Because of this we are forced to restrict our discussions to positive times. The best we can hope for is to be able to define a "semidynamical system," replacing the group of solution operators $\{T(t)\}_{t\in\mathbb{R}}$, which we could define in the ODE case, with a semigroup of solution operators $\{S(t)\}_{t\ge 0}$. In this way we are led to consider the semidynamical system

$$\big(L^2(0,\pi), \{S(t)\}_{t\ge 0}\big). \tag{11}$$

In contrast to the case of ODEs, for which there exists a general theory guaranteeing existence and uniqueness for a wide class of problems (this is covered in Chapter 2), no such unified approach is possible for PDEs. Each equation usually needs to be studied in its own right if we are to show that it has unique solutions, which can then be used to define a semidynamical system as in (11). We do this for two examples: a scalar reaction–diffusion equation,

$$\frac{\partial u}{\partial t} - \Delta u = f(u), \qquad u|_{\partial\Omega} = 0,$$

in Chapter 8, and the Navier–Stokes equations with periodic boundary conditions in two dimensions,

$$\frac{\partial u}{\partial t} - \nu\Delta u + (u \cdot \nabla u) + \nabla p = f(x), \qquad \nabla \cdot u = 0,$$

in Chapter 9.

The restriction that we can only define the solution for $t \geq 0$ does not interfere with one of the fundamental insights from the theory of dynamical systems, which motivates much of our approach: the complexity of the problem is significantly reduced if we are prepared to neglect any transient phenomena and concentrate only on describing the limiting behaviour as $t \to \infty$. For example, it follows from (9) that the limiting state of the system is $u(x) \equiv 0$, and a description of the "asymptotic dynamics" becomes as simple as can be.

In Chapter 10 we show that in some situations, which correspond physically to systems in which there is dissipation (friction, viscosity, etc.), the asymptotic dynamics all take place on a compact subset $\mathcal{A}$ of the original phase space. We call this set the "global attractor," and we verify that the scalar reaction–diffusion equation from Chapter 8 and the two-dimensional Navier–Stokes equations from Chapter 9 are both "dissipative" and so have a global attractor (Chapters 11 and 12). We also discuss the situation for the three-dimensional (3D) Navier–Stokes equations in some detail.

The remainder of the book concentrates on properties of this set $\mathcal{A}$ and how we can use its existence to deduce important consequences for the behaviour of the underlying PDE. For example, under certain mild conditions, it turns out that on $\mathcal{A}$ we can define $S(t)$ sensibly for all $t \in \mathbb{R}$, and so the dynamics restricted to the attractor define a standard dynamical system,

$$(\mathcal{A}, \{S(t)\}_{t\in\mathbb{R}}).$$

One of the most important properties of many of these attracting sets is that they are *finite-dimensional* subsets of the original, infinite-dimensional phase space. We prove this property for our two examples in Chapter 13, and it is the implications of this surprising result that occupy Chapters 14–16, where we investigate in what sense we can conclude that the original infinite-dimensional system is "in effect" finite-dimensional.

The final chapter consists of a long series of exercises that apply many of the techniques covered in the book to analyse the one-dimensional Kuramoto–Sivashinsky equation,

$$u_t + u_{xxxx} + u_{xx} + uu_x = 0, \qquad u(x,t) = u(x+L,t).$$

A Note on the Exercises

Each chapter ends with a series of exercises, which should be considered an integral part of the book. Indeed, several results proved in the exercises are used in the main body of the text. Some of the exercises in Chapters 14–17 are considerably more involved than those in the preceding chapters, since the material here is much more recent and is all still a focus of current research.

A full set of solutions to the exercises is available on the World Wide Web, at the following address:

`http : //www.cup.org/titles/0521635640.html`

I would welcome any comments, suggestions, or errata; these can be e-mailed to me at

`jcr@maths.warwick.ac.uk`

Errata will be posted periodically at the above web address.

Part I

Functional Analysis

Cambridge University Press
978-0-521-63204-1 - Infinite-Dimensional Dynamical Systems: An Introduction to Dissipative Parabolic PDEs and the Theory of Global Attractors
James C. Robinson
Excerpt
More information

1

Banach and Hilbert Spaces

The purpose of this chapter is to review some results from basic functional analysis, emphasising those aspects that will be useful in what follows. We introduce various examples of Banach spaces, in particular the spaces C^r of continuous functions and, after a brief outline of the construction of the Lebesgue integral, the L^p spaces of Lebesgue integrable functions. The final section treats some basic results concerning Hilbert spaces, of which L^2 is our most important example.

The material here is somewhat dry, but these are essential foundations. If the ideas are familiar, it is still worth paying attention to the Young and Hölder inequalities (Lemmas 1.8 and 1.9) and the technique of mollification (introduced in Section 1.3.1) by which we will prove various density results (Proposition 1.6, Theorem 1.13, and Corollary 1.14).

1.1 Banach Spaces and Some General Topology

A *norm* on a vector space X is a function $\|\cdot\| : X \to [0, \infty)$ satisfying

(i) $\|x\| = 0$ if and only if (iff) $x = 0$,
(ii) $\|\lambda x\| = |\lambda|\|x\|$ for all $x \in X$, $\lambda \in \mathbb{R}$, and
(iii) $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in X$ (the "triangle inequality").

A normed space is *complete* if every Cauchy sequence converges, and a *Banach space* is a complete normed space. Although this means that strictly a Banach space is a pair $(X, \|\cdot\|_X)$ of a space and its norm, we will usually use the more convenient notation X alone and specify the norm separately. Most common Banach spaces have a standard norm associated with them, and this is almost always the norm we will use, denoting it by $\|\cdot\|_X$. If there is any ambiguity we will specify the definition of the norm explicitly.

A subset Y of a Banach space X is *dense* if its closure in X, written $\overline{Y}$, is all of X. Equivalently, Y is dense in X if every element of X can be approximated arbitrarily closely by an element of Y, so that for any $\epsilon > 0$ there exists a $y \in Y$ such that $\|x - y\| < \epsilon$. In particular it follows that if $x \in X$ we can find a sequence $y_n \in Y$ such that

$$y_n \to x \quad \text{in} \quad X \quad \text{as} \quad n \to \infty,$$

(i.e. such that $\|y_n - x\|_X \to 0$ as $n \to \infty$). Showing that spaces of smooth functions are dense in spaces of Lebesgue integrable functions is one of the major topics of this chapter (Theorem 1.13 and Corollary 1.14).

If a Banach space X has a dense subset that is also countable then we say that X is *separable*. Occasionally the existence of such a countable dense subset is used to simplify an argument, so we will be careful in this chapter to point out any spaces that are separable. It follows easily that the finite product of separable spaces is also separable and a linear subspace of a separable space is separable (see Exercise 1.1).

Finally, the topological property of *compactness* will often be vital in our applications. Recall that a subset E of a Banach space X is *compact* if every open cover of E contains a finite subcover (see Exercise 1.2 for one application of this definition). An equivalent characterisation, which we will find much more useful, is that E is compact iff every sequence in E contains a convergent subsequence. That is, if we know that $\{x_n\} \in E$ then there exists a subsequence $x_{n_j} \to x^*$, with $x^* \in E$. Often we will be trying to solve a problem P that we cannot treat directly. Instead we will consider a sequence of easier problems P_n, which approximate P (in some sense), and for which we can find a solution x_n. If we can also show that the solutions x_n lie in a compact set, then we can hope that the limit of some convergent subsequence x_{n_j} will be a solution of our original problem P. This is the fundamental idea behind the approach we will use to prove the existence of solutions of our model PDEs in Chapters 7–9.

We will now cover various examples of Banach spaces, all of which will be needed in what follows. We start with the simplest.

1.2 The Euclidean Space $\mathbb{R}^m$

We write $x \in \mathbb{R}^m$ as $x = (x_1, \ldots, x_m)$. Then properties (i)–(iii) are easily checked for the standard Euclidean norm

$$|x| = \left(\sum_{j=1}^{m} x_j^2\right)^{1/2} \tag{1.1}$$