

Elements of Fractal Geometry and Dynamics

Yakov Pesin

Vaughn Climenhaga

DEPARTMENT OF MATHEMATICS, PENNSYLVANIA STATE UNIVERSITY,
UNIVERSITY PARK, PENNSYLVANIA 16802

E-mail address: pesin@math.psu.edu

DEPARTMENT OF MATHEMATICS, PENNSYLVANIA STATE UNIVERSITY,
UNIVERSITY PARK, PENNSYLVANIA 16802

E-mail address: climaha@math.psu.edu

Contents

Foreword: MASS and REU at Penn State University	ix
Lecture 1	1
a. A three-fold cord	1
b. Intricate geometry and self-similarity	1
c. Things that move (or don't)	4
Lecture 2	7
a. Dynamical systems: terminology and notation	7
b. Examples	8
Lecture 3	14
a. A linear example with chaotic behaviour	14
b. The Cantor set and symbolic dynamics	17
Lecture 4	21
a. A little basic topology	21
b. The topology of symbolic space	22
c. What the coding map doesn't do	23
Lecture 5	26
a. Geometry of Cantor-like sets	26
b. More general Cantor-like sets	27
c. Making sense of it all	30
Lecture 6	32
a. The definition of Hausdorff dimension	32
Lecture 7	36
a. Properties of Hausdorff dimension	36
b. Topological dimension	38
c. Comparison of Hausdorff and topological dimension	39
Lecture 8	41
a. Some point set topology	41
b. Metrics and topologies	43
Lecture 9	48
a. Topological consequences	48
b. Hausdorff dimension of Cantor-like constructions	49
Lecture 10	51
a. Completion of the proof of Theorem 21	51
Lecture 11	53
a. A competing idea of dimension	53
b. Basic properties and relationships	55

Lecture 12	58
a. Further properties of box dimension	58
b. A counterexample	59
Lecture 13	62
a. Subadditivity, or the lack thereof	62
b. A little bit of measure theory	63
Lecture 14	65
a. Lebesgue measure and outer measures	65
b. Hausdorff measures	67
Lecture 15	69
a. Choosing an outer measure	69
b. Measures on symbolic space	70
c. Measures on Cantor sets	71
Lecture 16	73
a. Markov measures	73
b. The support of a measure	74
c. Markov measures and dynamics	76
Lecture 17	78
a. Using measures to determine dimension	78
b. Pointwise dimension	79
Lecture 18	81
a. The Non-uniform Mass Distribution Principle	81
b. Non-constant pointwise dimension	82
Lecture 19	86
a. More about the Lyapunov exponent	86
b. Fractals within fractals	87
c. Hausdorff dimension for Markov constructions	89
Lecture 20	92
a. FitzHugh-Nagumo and you	92
Lecture 21	94
a. Numerical investigations	94
b. Studying the discrete model	95
Lecture 22	97
a. Stability of fixed points	97
b. Things that don't stand still, but do at least come back	101
Lecture 23	105
a. Down the rabbit hole	105
b. Becoming one-dimensional	107
Lecture 24	110
a. Bifurcations for the logistic map	110
b. Different sorts of bifurcations	112
Lecture 25	115
a. The simple part of the bifurcation diagram	115
b. The other end of the bifurcation diagram	116
c. The centre cannot hold—escape to infinity	118

Lecture 26	120
a. Some parts of phase space are more equal than others	120
b. Windows of stability	121
c. Outside the windows of stability	123
Lecture 27	124
a. Chaos in the logistic family	124
b. Attractors for the FitzHugh–Nagumo model	124
c. The Smale–Williams solenoid	126
Lecture 28	128
a. The Smale–Williams solenoid	128
b. Quantifying the attractor	130
Lecture 29	132
a. The non-conformal case	132
b. The attractor for the FitzHugh–Nagumo map	132
c. From attractors to horseshoes	133
Lecture 30	137
a. Symbolic dynamics on the Smale horseshoe	137
b. Variations on the horseshoe map	138
Lecture 31	141
a. Transient and persistent chaos	141
b. Continuous-time systems	141
Lecture 32	146
a. The pendulum	146
b. Two-dimensional systems	148
c. The Lorenz equations	151
Lecture 33	153
a. Beyond the linear mindset	153
b. Examining the Lorenz system	154
Lecture 34	158
a. Homoclinic orbits and horseshoes	158
b. Horseshoes for the Lorenz system	160
Lecture 35	163
a. More about horseshoes in the Lorenz system	163
Lecture 36	166
a. The Lorenz attractor	166
b. The geometric Lorenz attractor	166
Lecture 37	170
a. Random fractals	170
b. Back to the Lorenz attractor, and beyond	170

Lecture 1

a. A three-fold cord. The word “fractal” is one which has wriggled its way into the popular consciousness over the past few decades, to the point where a Google search for “fractal” yields over 12 million results (at the time of this writing), more than six times as many as a search for the rather more fundamental mathematical notion of “isomorphism”. With a few clicks of a mouse, and without any need to enter the jargon-ridden world of academic publications, one may find websites devoted to fractals for kids, a blog featuring the fractal of the day, photo galleries of fractals occurring in nature, online stores selling posters brightly emblazoned with computer-generated images of fractals. . . the list goes on.

Faced with this jungle of information, we may rightly ask, echoing Paul Gauguin, “What are fractals? Where do they come from? Where do we go with them?”

The answers to the second and third questions, at least as far as we are concerned, will have to do with the other two strands of the three-fold cord holding this course together—namely, dynamical systems and chaos.¹ As an initial, naïve formulation, we may say that the combination of dynamical systems and fractals is responsible for the presence of chaotic behaviour. For our purposes, fractals will come from certain dynamical systems, and will lead us to an understanding of certain aspects of chaos.

But all in good time. We must begin by addressing the first question; “What are fractals?”

b. Intricate geometry and self-similarity. Consider an oak tree in the dead of winter, viewed from a good distance away. Its trunk rises from the ground to the point where it narrows and sends off several large boughs; each of these boughs leads away from the centre of the tree and eventually sends off smaller branches of its own. Walking closer to the tree, one sees that these branches in turn send off still smaller branches, which were not visible from further away, and more careful inspection reveals a similar branching structure all the way down to the level of tiny twigs only an inch or two long.²

The key points to observe are as follows. First, the tree has a complicated and intricate shape, which is not well captured by the more familiar geometric objects, such as lines, circles, polygons, and so on. Secondly, we see the same sort of shape on all scales—whether we view the tree from fifty yards away or from fifty inches, we will see a branching structure in which the largest branch (or trunk) in our field of view splits into smaller branches, which then divide themselves, and so on.

¹Chaos theory has, of course, also entered the popular imagination in its own right recently, thanks in part to its mention in movies such as *Jurassic Park*.

²All of this is still present in summer, of course, but the leaves get in the way of easy observation.

These features are shared by many other objects which we may think of as fractals—we see a similar picture if we consider the bronchial tree, the network of passageways leading into the lungs, which branches recursively across a wide range of scales. Or we may consider some of the works of the artist M. C. Escher, we see intricate patterns repeating at smaller and smaller scales.

Yet another striking example may be seen by looking at a high-resolution satellite image (or detailed map) of a coastline. The boundary between land and sea does not follow a nice, simple path, but rather twists and turns back and forth; each bay and peninsula is adorned with still smaller bays and peninsulas, and given a map of an unfamiliar coast, we would be hard pressed to identify the scale at which the map was printed if we were not told what it was.

The two threads connecting these examples are their complicated geometry and some sort of *self-similarity*. Recall that two geometric figures (for example, two triangles) are *similar* if one can be obtained from the other by a combination of rigid motions and rescaling. A fractal exhibits a sort of similarity with itself; if we rescale a part of the image to the size of the whole, we obtain something which looks nearly the same as the original.

We now make these notions more precise. Simple geometric shapes, such as circles, triangles, squares, etc., have boundaries which are smooth curves, or at least piecewise smooth. That is to say, we may write the boundary parametrically as

$$\vec{r}(t) = (x(t), y(t)),$$

and for the shapes we are familiar with, x and y are piecewise differentiable functions from \mathbb{R} to \mathbb{R} , so that the tangent vector $\vec{r}'(t)$ exists for all but a few isolated values of t . By contrast, we will see that a fractal “curve”, such as a coastline, is continuous everywhere but differentiable *nowhere*.

As an example of this initially rather unsightly behaviour, we consider the *von Koch curve*, defined as follows. Taking the interval $[0, 1]$, remove the middle third (just as in the construction of the usual Cantor set), and replace it with the other two sides of the equilateral triangle for which it is the base. One obtains the piecewise linear curve at the top of Figure 1; this is the basic pattern from which we will build our fractal.

Observe that the second curve in Figure 1 consists of four copies of the first, each of which has been scaled to $1/3$ its original size and then used to replace one of the four line segments in the original pattern. The new curve contains 16 line segments, each of length $1/9$. Replacing each of these segments with an appropriately scaled copy of the basic pattern, we obtain the third curve in the figure, and so on.

Each step in this construction—each curve in Figure 1—is piecewise smooth, even piecewise linear. We may consider their parametrisations f_1, f_2, f_3, \dots , each of which is a piecewise smooth map from $[0, 1]$ to \mathbb{R}^2 . It is not too difficult to show that the sequence $\{f_n\}_{n=1}^{\infty}$ converges uniformly,

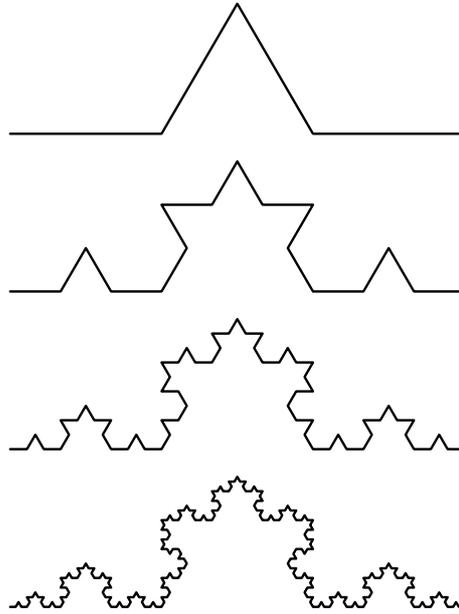


FIGURE 1. The first few steps in the construction of the von Koch curve.

and hence the limit $f: [0, 1] \rightarrow \mathbb{R}^2$ exists and is continuous. The von Koch curve is the image of this function f , the end of the limiting process whose first few steps are shown.

Although each of the functions f_n is piecewise smooth, their limit f is not differentiable anywhere, and hence the von Koch curve does not admit a tangent vector at any point, despite being continuous. This is a manifestation of the complicated and intricate geometry we referred to earlier; since the self-similarity of the curve is evident from the construction, we may justifiably call this object a fractal.

It is natural to characterise a curve by its length, and so we may ask how long the von Koch curve is. One may easily verify that the n^{th} step in the iterative procedure leading to the von Koch curve is a piecewise linear curve containing 4^n line segments, each of length $(1/3)^n$. At this stage of the iteration, then, the entire curve has length $(4/3)^n$ —but this quantity grows without bound as n goes to infinity! The only conclusion we can reach is that the von Koch curve has infinite length, despite being contained in a bounded region of the plane—we will see that this sort of behaviour is in fact quite common for fractals.

Indeed, consider the iterative procedure illustrated in Figure 2, wherein each side of the square is replaced with the zig-zag shown, which comprises four line segments of length slightly greater than $1/4$. Each of these segments is then replaced with an appropriately scaled version of the zig-zag pattern, and so on; the first few steps of the iteration are shown. Note that at each

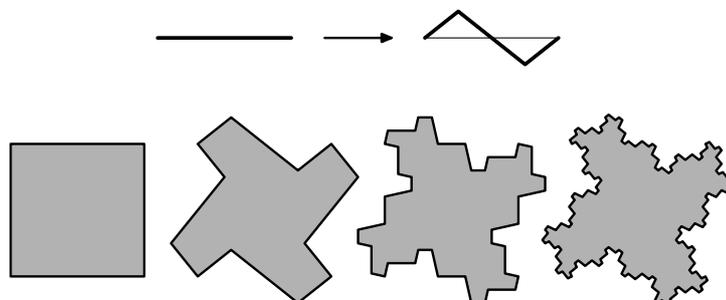


FIGURE 2. A fractal island.

step, we add exactly as much area as we remove, and so the area of each “island” is equal to 1. However, a similar calculation to the one above shows that the limiting fractal island has a coastline of infinite length, despite having unit area.

These last two examples show that the usual ways of characterising and measuring geometric objects—length, area, volume, etc.—are insufficient to deal with fractals. Both the von Koch curve and the coastline of the fractal island have infinite length, but zero area, and so we will need new tools in order to study them properly. First, though, we briefly turn our attention to the second strand of the three-fold cord, dynamical systems.

c. Things that move (or don’t). In some sense, anything that moves is a dynamical system (and for that matter, so is everything that doesn’t move). Somewhat more helpfully, we may consider any set X with a map f taking X to itself; that is, f assigns to each $x \in X$ an element $f(x) \in X$. If we think of each point in X as specifying a particular configuration of some system, then f is merely an encoding of the rule by which the system evolves from one state to the next. Some states evolve to other states under the action of f , while others may be fixed; if every point x is fixed, then f is the identity map, and nothing moves. But this is, of course, a rather trivial case.

We refer to the point $f(x)$ as the *image* of x under the action of f ; the essential feature of a dynamical system is that each image $f(x)$ is also an element of X , and thus lies in the domain of f ; that is, the map f takes X *into itself*, and so we can iterate it. Having found the image of a point x , we can then take the image of $f(x)$ in turn, which will be denoted $f^2(x) = f(f(x))$. Continuing the iteration, we obtain $f^3(x) = f(f^2(x))$, and in general, $f^{n+1}(x) = f(f^n(x))$.

In light of the key role iterative processes played in our earlier examples of fractals, the reader may feel justified in suspecting that the presence of an iterative process in this description of a dynamical system has something to do with the promised connection between the two; we will see later that this is in fact the case.

The sequence $x, f(x), f^2(x), \dots$ is referred to as the *trajectory* of x ; if we think of each iteration of the map f as specifying how the system evolves from one time step to the next, then it makes sense to think of the number of iterations n as the amount of time which has elapsed, and the trajectory is simply a list of the states through which the system passes as time goes on.³

If the map f is invertible, then we may also consider the point $f^{-1}(x)$, which is known as the *preimage* of x ; similarly, we have $f^{-2}(x), f^{-3}(x)$, and so on. Thus the trajectory is defined not just for positive values of n , but over the entire set of integers, and is a doubly infinite sequence of points in X . In fact, the notion of preimage is well defined and useful even if f is not invertible, but we will come to that later.

In and of themselves, sets are rather bland objects (with apologies to any set theorists in the audience), and so we usually consider dynamical systems defined on sets X which possess some additional structure. In particular, if we hope to have anything to do with fractals, which are geometric objects, the set X should possess some geometric structure, and so Euclidean space is a natural place to begin.

As an example, we may consider a rotation of the plane \mathbb{R}^2 by some angle α , or a translation by some vector \vec{v} , or a reflection in some line ℓ —these, together with the set of glide reflections, are all the rigid motions of the plane, and may all be thought of as dynamical systems.

We do not need to restrict ourselves to isometries—any matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ defines a dynamical system on \mathbb{R}^2 by

$$f: \quad \mathbb{R}^2 \rightarrow \mathbb{R}^2, \\ \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \mapsto \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

The value of the determinant of this matrix has ramifications for the properties of the dynamical system it defines. For example, if the determinant is equal to 1, then f is *area-preserving*; that is, the image $f(A)$ of a domain $A \subset \mathbb{R}^2$ has the same area as A itself.

We may also consider non-linear maps from the plane to itself—in fact, most of the interesting examples are of this sort. So for now, let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be any continuous map of the plane into itself. Let $A \subset \mathbb{R}^2$ be some domain—perhaps a disc, perhaps something rather more complicated—and suppose that the image of each point in A is itself in A . We say that A is *invariant* under the action of f , and so we can write $f: A \rightarrow A$. It follows that for any point $x \in A$, the entire trajectory of x lies within A .

In principle, if we have precise knowledge of the map f and the initial point x , then we can precisely compute each point $f^n(x)$ in the trajectory of

³This describes what is known as a *discrete time dynamical system*. One may also consider continuous time dynamical systems, in which the time variable may take any real value, but we will not do so just now.

x . There is no randomness in the action of f —it is entirely deterministic, and given sufficient patience and computing power, we can predict the future. Or so it seems. . .

Suppose we divide A into two subdomains A_1 and A_2 in such a way that every point $x \in A$ lies in exactly one of the two. Now instead of describing the trajectory of a point x by giving the precise location of each iterate $f^n(x)$, we may instead “blur our vision” and only record whether it lies in A_1 or A_2 . In this way we assign to a point x a sequence of 1’s and 2’s, known as a *coding* of x (really, of the trajectory of x).

It is natural to ask if we can go in the other direction—given a sequence of 1s and 2s, can we find a point x whose trajectory is coded by that sequence? If we can, is it unique, or might there be several such points?

The answer to these questions is somewhat involved, and depends heavily on the particular system f and on the choice of partition $\{A_1, A_2\}$. We will see that in many important cases, the answer to both questions is yes.

Suppose for the moment, then, that we have such a correspondence between trajectories of our dynamical system and sequences of 1’s and 2’s. Imagine taking a coin and flipping it repeatedly; after each flip, write down the number 1 if the coin comes up heads, and the number 2 if it comes up tails. In this manner we obtain a sequence of 1’s and 2’s which is entirely random, and which codes the trajectory of some point x .

This brings us to a rather jarring conclusion—the trajectory of this point x will appear to hop at random between A_1 and A_2 , just as the outcome of the coin toss hops at random between heads and tails. But we said earlier that f is wholly deterministic, with no randomness whatsoever—where, then, does this random-looking behaviour come from?

We will eventually resolve this paradox, but will first need to make the concepts involved more precise. For the time being, we merely observe that this initially unpalatable coexistence of deterministic and random behaviour is at the heart of the theory of chaos; indeed, it was to describe such situations that James Yorke first coined the somewhat controversial term “deterministic chaos”. We will see in due course how such behaviour arises from the combination of dynamical systems and fractal geometry.

Lecture 2

a. Dynamical systems: terminology and notation. Let us slow down now and take a more leisurely look at some of the concepts which will be foundational to our discussion of dynamical systems, before moving on to consider some apparently simple but ultimately extremely challenging and enlightening examples.

We begin with the n -dimensional Euclidean space \mathbb{R}^n , that is, the collection of n -element vectors, each of which is an n -tuple of real numbers. As our dynamical system f , we may consider any rule, algorithm, map, etc. which takes one element of \mathbb{R}^n and gives us back another.

It may happen that f is not defined on all of \mathbb{R}^n , but only on some domain $D \subset \mathbb{R}^n$. For instance, the rule which lets us go from x to $f(x)$ may only make sense when x is a vector of length no greater than R ; in this case, the domain of definition is the ball of radius R centred at the origin.

As mentioned last time, we say that $f(x)$ is *the* image of x , and x is a preimage of $f(x)$. The choice of article is important; while the map f must send x to a unique point $f(x)$, it is quite possible that there is some point $y \neq x$ with $f(y) = f(x)$, in which case y is also a preimage of $f(x)$.

We will also speak of the image of a set—if $A \subset D$ lies within the domain of definition, then the image of A is

$$f(A) = \{ f(x) \mid x \in A \}.$$

Of particular importance is the image of the domain D . This image $f(D)$ is also known as the *range* of f , and we will mostly be concerned with examples for which the range lies inside the domain of definition.

A subset $A \subset D$ for which $f(A) \subset A$ is said to be *invariant* under f ; if the domain D is invariant, then we can apply f again, and again, and again, *ad infinitum*, without ever leaving D . Thus we may consider not only f , but the map obtained by applying f twice—we denote this by f^2 , and write $f^2(x) = (f \circ f)(x) = f(f(x))$. Similarly, we may consider f^3, f^4 , and in general f^n , defined as

$$f^n = \underbrace{f \circ f \circ \cdots \circ f}_{n \text{ times}}.$$

Thus the ability to iterate f means that we are in fact considering a whole family of maps $f^n: D \rightarrow D$. It is immediate from the definition that

$$(1) \quad f^{n+m} = f^n \circ f^m = f^m \circ f^n$$

whenever m and n are non-negative integers. This is the so-called *semi-group property*, which allows us to translate questions about the behaviour of iterates of f to questions about the action of a particular semi-group⁴ on the domain D , leading to a more algebraic approach which is sometimes useful.

⁴We must speak of semi-groups rather than groups since f may be non-invertible.

The sequence of points $x, f(x), f^2(x), \dots$ is referred to as the *trajectory* of x (or sometimes as the *orbit*). The number n of iterations plays the role of time—the near future corresponds to small values of n , the far future to large values. While it may at first seem somewhat unnatural to think of time as moving in discrete increments, instead of a steady stream as we are accustomed to, observe that any measurement we may wish to make can only be carried out at discrete times—any two observations will be separated by a small interval, which may be vanishingly short or numbingly long, but nevertheless has the effect that our data is always collected with respect to a set of discrete time steps.

As suggested above, there are many important systems for which f is not one-to-one, and hence not invertible. Thus different initial conditions may eventually wind up following the same trajectory, and we cannot run time backwards—the past determines the future, but the future does not necessarily determine the past. Even in this time-irreversible case, we still may (and often do) consider preimages of points, and also of sets; given a set $A \subset D$, the preimage of A is

$$(2) \quad f^{-1}(A) = \{x \in D \mid f(x) \in A\},$$

which consists of all the points whose images lie in A . Note that although there may be several preimages of a point, a set only has one preimage.⁵

The definition (2) is valid whether f is invertible or not. If f is in fact invertible, with $f(D) = D$, then f^{-1} is a map in its own right. The existence of an inverse for f allows us to go backward in time, and produce the negative iterates

$$f^{-n} = \underbrace{f^{-1} \circ f^{-1} \circ \dots \circ f^{-1}}_{n \text{ times}}.$$

In this case preimages are uniquely defined, and the trajectory of x may be extended to a doubly infinite sequence $\{f^n(x)\}_{n \in \mathbb{Z}}$. Thus the group property (1) holds for *any* integers m and n , whether positive or negative.⁶

b. Examples.

b.1. *A rather unrealistic population model.* Consider a population of duck-billed platypi (or bacteria, or whatever species you fancy), whose size will be represented by a variable x . Given the size of the population at the present time, we want to predict the size of next year's population (or perhaps the next hour's, in the case of bacteria). So if there are x platypi this year, there will be $f(x)$ next year—of course, since we cannot have a negative number of platypi, we must restrict x to lie in the interval $[0, \infty)$, which will be the domain of definition for f .

⁵Of course, if we think of a point as a set with only one element, then its preimage as defined in (2) is unique as a set.

⁶And we have a true group action, rather than the action of a semi-group.

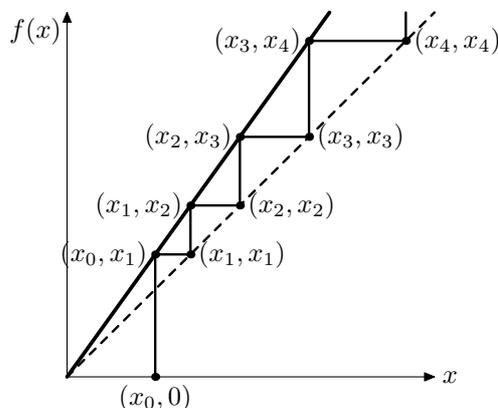


FIGURE 3. Cobweb diagram for a simple population model.

What form should f take? As a first (simplistic) approximation, we may suppose that the platypi reproduce at a constant rate, and so if there are x of them this year, there will be rx next year, where $r > 1$ is a real number, and $r - 1$ represents the proportion of newborns each year.

We would like to understand what the trajectories of the system look like for various possible starting populations. To this end, we use a *cobweb diagram*, as shown in Figure 3, which shows the graph of f . If x_0 is the initial value of x , then the next point in the trajectory is $f(x_0)$, which we denote by x_1 ; we may find this value by following the vertical line through $(x_0, 0)$, which intersects the graph of f at the point $(x_0, f(x_0)) = (x_0, x_1)$. Following the horizontal line through this point until it intersects the bisectrix $y = x$, we reach the point (x_1, x_1) , and now our x -coordinate is $x_1 = f(x_0)$, the next point in the trajectory after x_0 .

In order to find the next point in the trajectory after x_1 , we repeat this process; first follow the vertical line through (x_1, x_1) to its intersection with the graph of f , the point $(x_1, f(x_1)) = (x_1, x_2)$, and then move horizontally to (x_2, x_2) .

In general, we write $x_n = f^n(x_0)$ for the points of the trajectory, and we see that one obtains x_n from x_{n-1} by moving vertically to the graph, and then horizontally to the bisectrix. This gives a simple graphical procedure which allows us to investigate the qualitative properties of the trajectory of x_0 .

In this case, we see that for any initial population size $x_0 \neq 0$, the population size grows without bound; we say that the trajectory goes to infinity. The case $x_0 = 0$ is different, reflecting the fact that if there are no platypi to begin with, then no new ones will be born; nothing begets nothing. We say that 0 is a *fixed point* for the map f ; algebraically, a fixed point is a point x such that $f(x) = x$, and for any such point we see that the trajectory never moves.

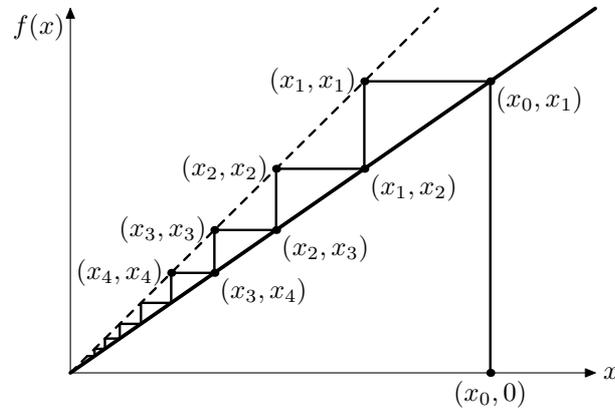


FIGURE 4. A dying population.

An important feature of this particular fixed point is that it is *unstable*—even a very small population x_0 will eventually grow to be arbitrarily large. Fixed points with this property, for which the trajectories of nearby points are driven away, are also called *repelling*.

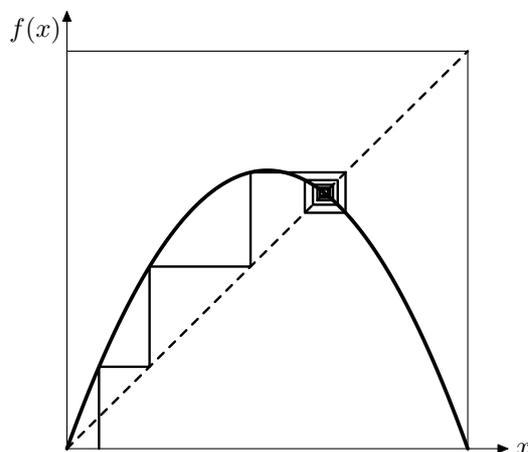
b.2. *A model which could be realistic.* Of course, as everyone knows, platypi are not immortal. *Alles Fleisch es ist wir Gras*, and our model needs to take into account the population reduction caused by death by disease, predation, etc. This will have the effect of changing the value of the parameter r , reducing it by counteracting the increase in population provided by the year's births. If it reduces it to the point where $r < 1$, then the graph of f is as shown in Figure 4, and the cobweb diagram clearly illustrates the fate of the platypus colony.

In this case, 0 is still a fixed point, but it is now *stable*—a value of x_0 near 0 will lead to a trajectory which converges to 0. Fixed points with this property, for which the trajectories of nearby points are drawn to the fixed point, are also called *attracting*.

Note that from the mathematical point of view, there is a simple relationship between the case $r > 1$ and the case $r < 1$; they are inverses. Indeed, if we write $f_r: x \mapsto rx$, then it is easy to see that $f_r^{-1} = f_{1/r}$.

b.3. *An innocent-looking model.* As any biologist or ecologist will no doubt protest quite vigorously, the preceding models are so simplistic as to be entirely unrealistic. Among other weaknesses, they fail to take into account the fact that resources are limited, and whatever river our platypi find themselves in can only support a finite population size before starvation or overcrowding leads to disaster.

To address this shortcoming, we introduce a new term into our equation. Suppose the environment determines some maximum population P , which corresponds, for instance, to the amount of resources available. Then the population cannot grow beyond P , and furthermore, if the population reaches P , all the food will be eaten and the platypi will starve, sending

FIGURE 5. The logistic map with $r = 2.8$.

the next year's population to 0. We model this situation with the formula $x \mapsto rx(P - x)$; in order to keep the equations as simple as possible, though, we rescale x so that it stands for the *proportion* of the maximum population P , and so lies between 0 and 1. Then the dynamical system in question is

$$(3) \quad \begin{aligned} f: [0, 1] &\rightarrow [0, 1], \\ x &\mapsto rx(1 - x), \end{aligned}$$

where r is a parameter encoding information about the system, such as reproduction rate, mortality rate, etc. The map in (3) is known as the *logistic map*, and its graph for the value $r = 2.8$ is shown in Figure 5, along with a typical trajectory.

Unlike the example $x \mapsto rx$ examined earlier, the logistic map displays a startling intricacy when we begin to track the behaviour of typical trajectories for various values of the parameter r . Indeed, the amount of literature on the logistic map is such that one could easily devote an entire year's course to the subject without exhausting the corpus of present knowledge, and the logistic map (along with its relatives) is still an area of active research.

Leaving behind for the time being any physical interpretations of the model, let us focus on the mathematical structure. By performing the appropriate change of coordinates $x \mapsto y$, we can show that the map $f: x \mapsto rx(1 - x)$ is equivalent to the map $g: y \mapsto y^2 + c$, where the value of the parameter c will depend on the value of r .

EXERCISE 1. Find an explicit change of coordinates which demonstrates the above correspondence. Which values of c corresponds to values of r which could occur in the model?

For large enough values of c , the graph of g lies entirely above the bisectrix, and every trajectory escapes to infinity, as shown in Figure 6. The

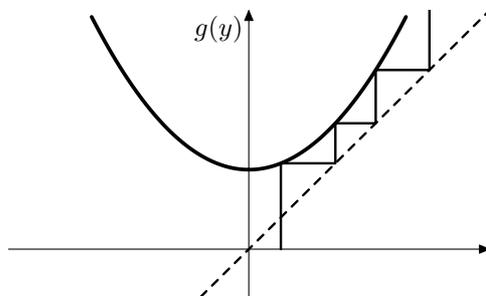


FIGURE 6. Trajectories escaping to infinity.

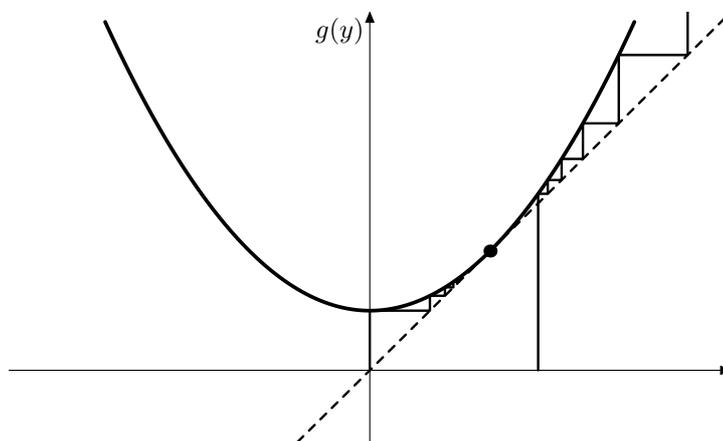


FIGURE 7. A fixed point which is neither attracting nor repelling.

parabola moves down as c decreases, and eventually, for some critical value of c , becomes tangent to the bisectrix, as shown in Figure 7.

EXERCISE 2. Find the value of c for which this occurs.

The point of tangency p is a fixed point for g . As is evident from the cobweb diagram, trajectories which start a little bit to the left of p are attracted to it, while trajectories which start just to the right are repelled and go to infinity. Thus in this case we have a fixed point which is neither an attractor nor a repeller.

Moving the initial point further to the left, one sees that for large enough negative values of x_0 , the next point in the trajectory leaps to the right of p , and then the trajectory goes to infinity. The point of transition between the two sorts of behaviour is $x_0 = -p$, which leads to $x_1 = f(x_0) = p$, and so the trajectory becomes trapped on the fixed point p .

Thus we have completely classified the asymptotic behaviour of trajectories for this particular map; points in $[-p, p]$ are attracted to the fixed point p , while all other points go to $+\infty$ under repeated iterations.

We will see, however, that the picture becomes vastly more complicated than this if we continue to decrease the parameter c .

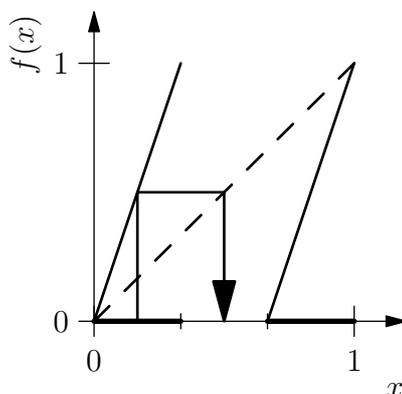


FIGURE 8. A piecewise linear map with chaotic behaviour.

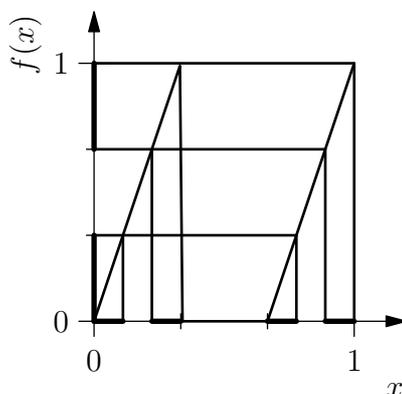
Lecture 3

a. A linear example with chaotic behaviour. Aside from being quite unrealistic, the linear population model in the previous lecture did not display any chaotic behaviour. This is actually a feature of any linear map—the theory of *Jordan normal form* (sometimes called *Jordan canonical form*), which is one of the most important results in basic linear algebra, offers a complete classification of linear maps in \mathbb{R}^n , and describes all the possible behaviours, none of which display any real complexity.

By contrast, the logistic map $f: x \mapsto x^2 + c$ displays a variety of complex behaviours as we consider different parameter values, a quality which makes it eminently worth of further study. In the previous lecture, we described its behaviour for large values of c , and took a very brief look at how that behaviour becomes more intricate as c decreases. In fact, for values of c near -2 , the logistic map exhibits fully chaotic behaviour, in the sense discussed in Lecture 1.

We will have more to say about the logistic map later—for the time being, we remark that a good deal of its complex behaviour can be attributed to its non-linearity. That same non-linearity, though, makes the map far more difficult to study—linear models are simply more tractable than non-linear ones. For this reason, we will first spend some time studying a map which is only piecewise linear, rather than fully linear, but which nevertheless displays chaotic behaviour.

Consider the map f shown in Figure 8, which is defined linearly on each of the intervals $I_1 = [0, 1/3]$ and $I_2 = [2/3, 1]$ so that the image of both intervals is $f(I_1) = f(I_2) = [0, 1]$. Thus the domain of definition of f is $D = I_1 \cup I_2$, and the range is $[0, 1]$. Notice that the range does not lie inside the domain of definition— $I_1 \cup I_2$ is not invariant for f , and so f is not defined at every point in the range. The cobweb diagram in the figure shows one iteration in the trajectory of the point $1/6$, whose image lies outside the domain of definition, and hence cannot be iterated further.

FIGURE 9. Finding the preimage $f^{-1}(D)$.

If we cannot iterate the map f , then we cannot study the dynamics, and so we must determine which points admit a second iteration. That is, what is the domain on which the map $f^2 = f \circ f$ is defined?

In order for $f^2(x_0)$ to be defined, both x_0 and $f(x_0)$ must lie in the domain of f ; that is, we must have

$$x_0 \in D \cap f^{-1}(D) = \{x \mid x \in D \text{ and } f(x) \in D\}.$$

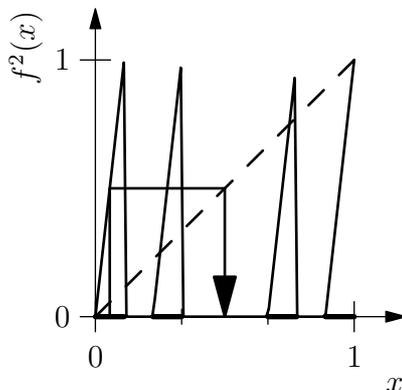
This is shown graphically in Figure 9, a sort of cobweb diagram; placing the domain D along the vertical axis, we find its preimage $f^{-1}(D)$ by following each horizontal line through D to *all* of the points where it intersects the graph of f , and then moving vertically from these intersection points to the x -axis. We see that the domain on which f^2 is defined consists of four closed intervals, each of length $1/9$. Writing these as

$$\begin{aligned} I_{11} &= \left[0, \frac{1}{9}\right], & I_{21} &= \left[\frac{2}{3}, \frac{7}{9}\right], \\ I_{12} &= \left[\frac{2}{9}, \frac{1}{3}\right], & I_{22} &= \left[\frac{8}{9}, 1\right], \end{aligned}$$

we see that $f(I_{11}) = f(I_{21}) = I_1$, that $f(I_{12}) = f(I_{22}) = I_2$, and that $f^2(I_{11}) = f^2(I_{12}) = f^2(I_{21}) = f^2(I_{22}) = [0, 1]$. Observe that the intervals $I_{i_1 i_2}$ may be defined in terms of the action of f as follows:

$$(4) \quad I_{i_1 i_2} = I_{i_1} \cap f^{-1}(I_{i_2}).$$

We obtained the domain of f by removing the (open) middle third from the interval $[0, 1]$, leaving two closed intervals of length $1/3$. From these, we obtained the domain of f^2 by removing the (open) middle third of each, leaving four closed intervals of length $1/9$. The graph of f^2 is shown in Figure 10; in this picture we already see the beginnings of self-similarity, and it is easy to see that the domain of f^3 will likewise consist of eight closed intervals of length $1/27$.

FIGURE 10. The second iterate of f .

In general, an inductive argument shows that the domain of f^n consists of 2^n closed intervals, each of length 3^{-n} . Following (4), we may denote these by

$$(5) \quad I_{i_1 i_2 \dots i_n} = I_{i_1} \cap f^{-1}(I_{i_2}) \cap \dots \cap f^{-(n-1)}(I_{i_n}),$$

where each i_k is either 1 or 2. Observe that for a fixed value of n , any two such intervals are disjoint; that is,

$$I_{i_1 \dots i_n} \cap I_{j_1 \dots j_n} = \emptyset$$

whenever $(i_1, \dots, i_n) \neq (j_1, \dots, j_n)$. Increasing n by one, we see from the construction that

$$I_{i_1 \dots i_n} = I_{i_1 \dots i_{n-1} 1} \cup I_{i_1 \dots i_{n-1} 2}.$$

Thus the domain of definition of the n^{th} iterate f^n may be written

$$D_n = \bigcup_{(i_1, \dots, i_n)} I_{i_1 \dots i_n},$$

where the union is taken over all n -tuples with values in $\{1, 2\}$. Letting n run to infinity, we see that the domain on which *every* iterate f^n is defined is

$$C = \bigcap_{n \geq 1} \bigcup_{(i_1, \dots, i_n)} I_{i_1 \dots i_n},$$

which is the standard middle-thirds Cantor set. As regards the dynamics of the map f , the key property of C is that it is the largest invariant set. In the first place, it is genuinely invariant, unlike the approximations D_n , for which we have $f(D_n) = D_{n+1} \not\subset D_n$; for the Cantor set itself, we have $f(C) = C$. Furthermore, if $A \subset D$ is any invariant set, $f(A) \subset A$, then we must have $A \subset C$.

Thus if we wish to study the dynamics of f , the “proper” domain to consider is the Cantor set C . To get a first idea of how the dynamics of $f: C \rightarrow C$ behave, consider two distinct points $x, y \in C$ which are very

close together, so that $d(x, y)$ is very small. How far apart are their images? It is not too hard to see that we have

$$d(f(x), f(y)) = 3d(x, y),$$

so that the distance between x and y is increased by a factor of 3. For higher iterates f^n , we have

$$d(f^n(x), f^n(y)) = 3^n d(x, y),$$

provided that the trajectories have stayed close up until that point; in particular, this will be true if $f^k(x)$ and $f^k(y)$ lie in the same interval I_1 or I_2 for each $1 \leq k < n$.

The significance of this result is that small errors are magnified; if x represents the true state of the system, but we instead measure it as y due to a very small experimental error, then the trajectory we predict will diverge exponentially quickly from the true trajectory. In this case, the notion of instability, which we introduced earlier for fixed points, applies to every trajectory of the system; whatever trajectory we look at, nearby trajectories will be repelled at an exponential rate. This phenomenon is known as *sensitive dependence on initial conditions*, and may be thought of as a preliminary indication of chaotic behaviour.

b. The Cantor set and symbolic dynamics. When Georg Cantor first conceived his eponymous set, he was hoping to settle the continuum hypothesis by constructing a subset of the interval whose cardinality lay strictly between that of the integers and that of the real line. While this turned out not to be the case, as we shall see, the Cantor set has nevertheless become an object of fundamental importance to a number of different areas in mathematics, not just dynamics, although it is this last incarnation of the Cantor set with which we shall be most concerned.

We start with an innocent-looking question. How big is the Cantor set? Of course, we need a notion of “bigness”. Since each step in the construction involves intervals, we might first try the notion of length. At the first iteration, we see that D_1 comprises two intervals of length $1/3$, and so has total length $2/3$. At the next iteration, four intervals of length $1/9$ give D_2 a total length of $4/9$. In general, we may easily see that D_n is a union of intervals with total length $(2/3)^n$; since this goes to 0 as $n \rightarrow \infty$, we must consider the “length” of C to be 0. Alternately, we may look at the lengths of the intervals which are removed at each step, and see that they sum to 1.

From a probabilistic point of view, this means that if we choose a point in the interval $[0, 1]$ at random, the probability of picking a point on the Cantor set is precisely zero. It would seem that length is not the proper way to measure how big C is.

Since C is not “big enough” to have positive length, we may try measuring it a different way, by counting the number of points it contains. We immediately see that it has infinitely many points, and so we next ask

whether it is countable or uncountable. To answer this question, we observe, as Cantor did, that each point $x \in C$ uniquely determines a sequence i_1, i_2, \dots , where each i_k is either 1 or 2, by the rule

$$x \in I_{i_1} \cap I_{i_1 i_2} \cap \dots \cap I_{i_1 \dots i_n} \cap \dots,$$

where all we are doing is asking which interval $I_{i_1 \dots i_n}$ contains x at each step n of the iteration. This defines a map from C to the space of symbolic sequences

$$\Sigma_2^+ = \{1, 2\}^{\mathbb{N}} = \{ (i_k)_{k=1}^{\infty} \mid i_k = 1 \text{ or } 2 \forall k \geq 1 \}.$$

Furthermore, the correspondence is bijective—given any such sequence, we see that the intersection

$$\bigcap_{n=1}^N I_{i_1 \dots i_n} = I_{i_1 \dots i_n}$$

is an interval whose length goes to 0 as $N \rightarrow \infty$, and so

$$\bigcap_{n=1}^{\infty} I_{i_1 \dots i_n} = \{x\}.$$

It follows that the sequence i_1, i_2, \dots comes from exactly one point x , and we have demonstrated that the following *coding map* is a bijection:

$$\begin{aligned} h: \quad \Sigma_2^+ &\rightarrow C, \\ \omega = (i_1, i_2, \dots) &\mapsto \bigcap_{n \geq 1} I_{i_1 \dots i_n}. \end{aligned}$$

EXERCISE 3. Using binary expansions of real numbers, show that Σ_2^+ has the same cardinality as $[0, 1]$, and hence that C does as well.

In fact, the coding map h does more than just establish a bijection between Σ_2^+ and C , which only shows that the two are the same from a set-theoretic point of view. The correspondence runs deeper than that, to an equivalence between the *dynamics* of the two sets as well.

Of course, at this point we have not put any dynamics on the set Σ_2^+ , and so we define a map $\sigma: \Sigma_2^+ \rightarrow \Sigma_2^+$ in order for the previous claim to make any sense. Recalling the definition of the sets $I_{i_1 \dots i_n}$ in (5), we see that the coding of a point $\omega = (i_n)_{n \in \mathbb{N}} \in \Sigma_2^+$ can be written

$$h(\omega) = \bigcap_{n=1}^{\infty} f^{-(n-1)}(I_{i_n}) = I_{i_1} \cap f^{-1}(I_{i_2}) \cap f^{-2}(I_{i_3}) \cap \dots,$$

and so

$$\begin{aligned} (6) \quad f(h(\omega)) &= I_{i_2} \cap f^{-1}(I_{i_3}) \cap f^{-2}(I_{i_4}) \cap \dots \\ &= h(\omega') \end{aligned}$$

where we write $\omega' = (i_2, i_3, \dots)$, and use the fact that $f(I_{i_1}) = [0, 1]$, and also that $f(f^{-1}(X)) = X$ for any set X in the range of f .⁷ The map which takes ω to ω' is particularly simple—all we have to do is drop the first symbol, i_1 , and shift all the others one position to the left. This is the *shift map*

$$\begin{aligned} \sigma: \quad \Sigma_2^+ &\rightarrow \Sigma_2^+, \\ (i_1, i_2, i_3, \dots) &\mapsto (i_2, i_3, i_4, \dots), \end{aligned}$$

with which (6) can be written in the form

$$(7) \quad f \circ h = h \circ \sigma.$$

Thus we have shown that the following diagram commutes:

$$(8) \quad \begin{array}{ccc} \Sigma_2^+ & \xrightarrow{\sigma} & \Sigma_2^+ \\ \downarrow h & & \downarrow h \\ C & \xrightarrow{f} & C \end{array}$$

That is, we may follow a point in Σ_2^+ by first applying σ and then using h to pass to a point in C , or by first passing to a point in C via h , and then applying f , and we will reach the same point whichever way we go.

The relationship between f and σ given by (7) is called a *conjugacy*, and allows us to draw conclusions about the dynamics of f based on analogous results for the dynamics of σ .

For example, we may ask how many periodic points of a given order f has; that is, how many solutions there are to the equation $f^m(x) = x$ for a fixed integer m . Two obvious periodic points are 0 and 1, which are fixed by f and are thus immediately periodic. It is not so obvious what happens for larger values of m , but we may obtain the answer relatively easily by passing to the symbolic setting. Here we see that any fixed point must have $i_2 = i_1$, and similarly $i_{n+1} = i_n$ for every n . Thus the only fixed points are $(1, 1, 1, \dots)$ and $(2, 2, 2, \dots)$, which correspond to 0 and 1, respectively. For $m = 2$, the equation $\sigma^2(\omega) = \omega$ tells us that we may choose i_1 and i_2 to be either 1 or 2, but that after that we must have

$$i_n = \begin{cases} i_1 & n \text{ odd,} \\ i_2 & n \text{ even.} \end{cases}$$

Thus there are four points with $\sigma^2(\omega) = \omega$ —in addition to the two mentioned above, we have $(1, 2, 1, 2, \dots)$ and $(2, 1, 2, 1, \dots)$.

In general, any sequence ω which repeats after m digits will satisfy $f^m(\omega) = \omega$, and since there are 2^m such sequences, we have 2^m periodic points of period m .⁸ Passing to C via the conjugacy given by h , we see that

⁷Note that the similar-looking statement $f^{-1}(f(X)) = X$ is *not* true in general.

⁸Of course, some of these will also be periodic points of lower order. We must do a little more work if we wish to count points with *primitive* period m .

f also has 2^m periodic points of period m , and so the set of periodic points is countable.

EXERCISE 4. Argue directly from the definition of f that the set of periodic points is countably infinite. Can you obtain a result on the number of periodic points with a given period without using symbolic dynamics?

Having seen that h respects the dynamics of f and σ , it is natural to ask what other aspects of the sets C and Σ_2^+ are preserved by the conjugacy. Since C lies in the interval $[0, 1]$, we have a notion of distance, and hence a definition of convergence;⁹ we can say what it means for a sequence $(x_n)_n \subset C$ to converge to a point $x \in C$.

If we define a distance function on Σ_2^+ , then we will have a notion of convergence there too, and we may ask whether h takes convergent sequences in Σ_2^+ to convergent sequences in C , and vice versa. This will expand the range of questions about the dynamics of f which can be answered by looking at the symbolic case to include questions of a topological nature—that is, questions involving convergence.

To this end, given two sequences $\omega = (i_1, i_2, \dots)$ and $\omega' = (i'_1, i'_2, \dots)$ in Σ_2^+ , define the distance between them by

$$(9) \quad d_\alpha(\omega, \omega') = \sum_{j \geq 1} \frac{|i_j - i'_j|}{\alpha^j},$$

where $\alpha > 1$ is fixed. Note that since each numerator $|i_j - i'_j|$ is either 0 or 1, this series converges absolutely. We may easily verify that $d = d_\alpha$ satisfies the axioms of a metric:

- (1) $d(\omega, \omega') \geq 0$, with equality iff $\omega = \omega'$.
- (2) $d(\omega, \omega') = d(\omega', \omega)$.
- (3) $d(\omega, \omega') \leq d(\omega, \omega'') + d(\omega'', \omega')$ (triangle inequality).

Each of these follows immediately from its counterpart for the usual distance on \mathbb{R} .

With this definition, we will see in the next lecture that if $(x_n)_n \subset \Sigma_2^+$ is any convergent sequence (with respect to d), then $(h(x_n))_n \subset C$ converges as well, and so h is *continuous*. A similar conclusion applies for h^{-1} , and so the coding map $h: \Sigma_2^+ \rightarrow C$ is a *homeomorphism*, that is, a continuous bijection with continuous inverse.

Colloquially, we may say that h respects convergence, and so we have a correspondence not only between the dynamics of $f: C \rightarrow C$ and $\sigma: \Sigma_2^+ \rightarrow \Sigma_2^+$, but between the topologies as well.

⁹More technically, C inherits a metric structure from the interval, which defines a topology in the natural way.

Lecture 4

a. A little basic topology. Let us pause to recall some of the basic definitions of point-set topology in the context of a metric space. Given a space X with a metric d , we may consider the set of all points which lie within a fixed distance $r > 0$ of a point $x \in X$:

$$B(x, r) = \{ y \in X \mid d(x, y) < r \}.$$

This is the *open ball of radius r centred at x* . We say that a set $U \subset X$ is *open* if for every $x \in U$ there exists $r > 0$ such that $B(x, r) \subset U$; that is, if a sufficiently small ball around x is contained in U for every $x \in U$.

If U_1 and U_2 are open sets, it is easy to verify that their union $U_1 \cup U_2$ is open as well—indeed, this holds for the union of *any* collection of open sets, no matter how large. One may also check that the intersection $U_1 \cap U_2$ is open as well, and that this property carries over to *finite* intersections $U_1 \cap \dots \cap U_n$, but not to infinite intersections, as the example $U_n = (-1/n, 1/n) \subset \mathbb{R}$ illustrates.

Given any set $E \subset X$, we may consider the *r -neighbourhood of E* , defined as

$$\bigcup_{x \in E} B(x, r);$$

this is just the set of all points in X which lie within r of some point in E , and is in some sense a “fattening” of the set E . For example, if E is the ball $B(0, a) \subset \mathbb{R}^n$, then the r -neighbourhood of E is just the larger ball $B(0, a + r)$. If E is a set with a more complicated geometry, such as the Cantor set, then its r -neighbourhoods will in some sense have a simpler geometric structure than E itself.

We say that a set E is *closed* if for every sequence $(x_n)_{n \in \mathbb{N}} \subset E$ which converges to some point $x \in X$, we have in fact $x \in E$. This is often expressed as the statement that E contains its limit points, and we leave as an exercise the fact that this is equivalent to the property that the complement of E , denoted $X \setminus E = E^c$, is open.

From this last statement, or from the definition, it follows that arbitrary intersections of closed sets are closed, as are finite unions of closed sets. Infinite unions of closed sets may not be closed—consider $E_n = [1/n, 1]$.

We need a notion of when two metric spaces (X, d) and (Y, ρ) have the same topological structure, and this relies on the idea of a continuous map. Given $h: X \rightarrow Y$ and $x \in X$, we say that h is *continuous at x* if for all $\varepsilon > 0$ there exists $\delta > 0$ such that $\rho(h(x), h(y)) < \varepsilon$ whenever $d(x, y) < \delta$. A map which is continuous at each point $x \in X$ is referred to simply as *continuous*. If h is one-to-one and onto—that is, if it is a bijection—we may consider the inverse map $h^{-1}: Y \rightarrow X$. If both h and h^{-1} are continuous, then h is a *homeomorphism*; such an h gives a correspondence between open sets in X and open sets in Y , and similarly for closed sets, so the two spaces have equivalent topologies.

b. The topology of symbolic space. What are the open and closed sets in the symbolic space Σ_2^+ ? We defined a metric (9) on Σ_2^+ , and so all the notions in the previous section make sense, but what do the open and closed sets actually look like?

For the sake of this discussion, fix $\alpha > 2$, and consider the ball $B(\omega, r)$ centred at a point $\omega = (i_1, i_2, \dots) \in \Sigma_2^+$ with radius $r = 1/\alpha > 0$. How do we tell if another point $\omega' = (i'_1, i'_2, \dots)$ is in $B(\omega, r)$? The distance between the two points is

$$d_\alpha(\omega, \omega') = \sum_{j=1}^{\infty} \frac{|i_j - i'_j|}{\alpha^j},$$

and we see immediately that if $i_1 \neq i'_1$, the first term alone means that the sum is $\geq 1/\alpha$. Conversely, if $i_1 = i'_1$, then the first term in the sum vanishes, and the distance is at most

$$\sum_{j=2}^{\infty} \frac{1}{\alpha^j} = \frac{1}{\alpha} \frac{1}{\alpha - 1} < \frac{1}{\alpha},$$

where the last inequality uses the fact that $\alpha > 2$. Thus we see that

$$B(\omega, r) = \{ \omega' \in \Sigma_2^+ \mid i'_1 = i_1 \}.$$

There are exactly two possibilities for i_1 , and so there are exactly two possible sets of this form:

$$\begin{aligned} C_1 &= \{ \omega' = (1, i'_2, i'_3, \dots) \}, \\ C_2 &= \{ \omega' = (2, i'_2, i'_3, \dots) \}. \end{aligned}$$

We refer to C_1 and C_2 as *cylinders of length 1*; each contains all sequences in Σ_2^+ for which the first term matches a particular specification. If we demand that the first n terms follow a particular pattern, we obtain a *cylinder of length n* :

$$(A) \quad C_{i_1 \dots i_n} = \{ \omega' \in \Sigma_2^+ \mid i'_k = i_k \ \forall 1 \leq k \leq n \}.$$

Following the above argument, we see that these are exactly the balls of radius $1/\alpha^n$, provided $\alpha > 2$. There is a one-to-one correspondence between cylinders of length n and n -tuples with entries in $\{1, 2\}$.

In fact, the first part of the argument goes through whatever value of $\alpha > 1$ we use for the metric, and we have the following:

PROPOSITION 1. *Cylinders are open.*

PROOF. Given a cylinder $C_{i_1 \dots i_n}$ and a point $\omega \in C$, we may choose any $r < 1/\alpha^n$, and then we see, as above, that $d(\omega, \omega') \geq r$ unless all the terms with $j \leq n$ vanish; that is, unless $i_j = i'_j$ for all $1 \leq j \leq n$. Thus $d(\omega, \omega') < r$ implies $\omega' \in C_{i_1 \dots i_n}$, and so $B(\omega, r) \subset C_{i_1 \dots i_n}$. \square

That's not the end of the story, though. . .

PROPOSITION 2. *Cylinders are closed.*

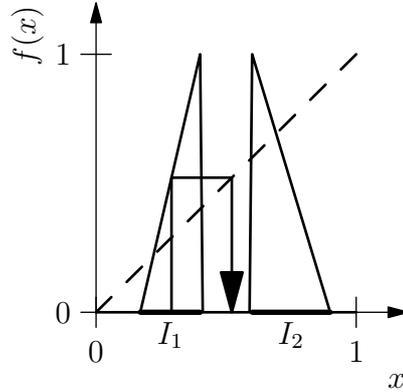


FIGURE 11. Another piecewise linear map with chaotic behaviour.

PROOF. Let $C_{i_1 \dots i_n}$ be a cylinder in Σ_2^+ , and suppose $(\omega^{(m)})_m \subset C_{i_1 \dots i_n}$ is a sequence which converges to $\omega^{(0)} \in \Sigma_2^+$ as $m \rightarrow \infty$. Then $d(\omega^{(m)}, \omega^{(0)}) \rightarrow 0$, and in particular, each term in the sum (9) must go to 0. Thus $i_j^{(m)} \rightarrow i_j^{(0)}$ for every $j \geq 1$, and since $i_j^{(m)} = i_j$ for every $1 \leq j \leq n$ and all m , we have $i_j^{(0)} = i_j$ for $1 \leq j \leq n$, and so $\omega^{(0)} \in C_{i_1 \dots i_n}$. It follows that $C_{i_1 \dots i_n}$ is closed. \square

Thus cylinders are both open *and* closed, a somewhat unfamiliar phenomenon if our only experience is with the topology of \mathbb{R} . The feature of the topology of Σ_2^+ which permits this behaviour is the fact that the cylinders of a given length are all disjoint, and their union is the whole space—we say that they *partition* Σ_2^+ . This gives an alternate proof of the second proposition above, once the first is known; the complement of an n -cylinder $C_{i_1 \dots i_n}$ is a union of $2^n - 1$ n -cylinders, each of which is open, and hence $C_{i_1 \dots i_n}$ is closed.

We can now show that the topologies of the Cantor set C and the symbolic space Σ_2^+ are equivalent:

THEOREM 3. *The coding map $h: \Sigma_2^+ \rightarrow C$ is a homeomorphism.*

PROOF. Recall that h is a bijection, and is defined by the inclusion

$$x \in I_{i_1} \cap I_{i_1 i_2} \cap \dots,$$

and so we see that $h(C_{i_1 \dots i_n}) = I_{i_1 \dots i_n}$ for every cylinder in Σ_2^+ . Since the sets $I_{i_1 \dots i_n}$ are all closed in C (being closed intervals), we have shown that h and h^{-1} both take closed sets to closed sets, which suffices to show that h is a homeomorphism. \square

c. What the coding map doesn't do. In the previous lecture, we saw that the coding map respects the dynamics of the two systems $f: C \rightarrow C$ and $\sigma: \Sigma_2^+ \rightarrow \Sigma_2^+$; we have now seen that it respects topology as well. So

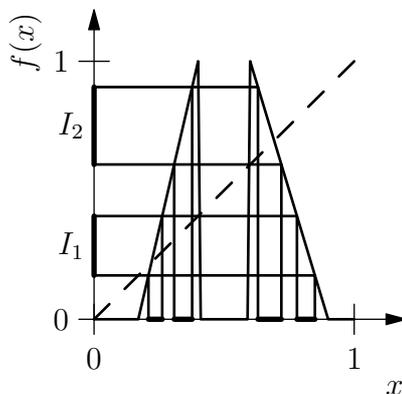


FIGURE 12. The domain of definition of f^2 .

from either a dynamical or topological point of view, we may as well study whichever is better suited to the problem at hand, knowing that our results will be valid for the other as well.

However, the two systems are not equivalent in every aspect. In the first place, C lies on the real line, from which it inherits both an ordering and a differentiable structure; it makes sense to think of a point $z \in C$ as lying *between* two others x and y , and since we can add and subtract elements of C , we can define the derivative of a map f as the limit of the ratio $(f(y) - f(x))/(y - x)$. In contrast, Σ_2^+ has no notion of “between-ness”, nor have we defined a way add or subtract sequences, and so the definition of derivative is meaningless.

A further hint that Σ_2^+ might not capture quite everything there is to know about the Cantor set C comes from considering more general dynamical systems defined in the interval. Fix two disjoint closed intervals $I_1, I_2 \subset [0, 1]$, and define a piecewise linear map $f: I_1 \cup I_2 \rightarrow [0, 1]$ as shown in Figure 11, so that $f(I_1) = f(I_2) = [0, 1]$ (note that for our purposes, each branch of f may be either increasing or decreasing).

If we try to iterate f more than once, we run into the same problem as before; some points in I_1 or I_2 have images which do not lie in either interval, and so cannot be iterated again. This leads us down exactly the same path as in the previous lecture; the domain of definition of f^2 is a union of four intervals, as shown in Figure 12, and so on for f^3, f^4, \dots . The only difference in this case is that the intervals may be of varying lengths, but the combinatorial and topological structure is identical to that in the previous analysis, and we again get *a* Cantor set (rather than *the* Cantor set), for which we have a coding map and symbolic dynamics just as before.

Thus we see that Σ_2^+ models not just the dynamics of our original map on the middle-thirds Cantor set, but the dynamics of *any* map defined in this fashion (note that the original definition is just a special case of this one, with $I_1 = [0, 1/3]$ and $I_2 = [2/3, 1]$). Indeed, all these maps (and

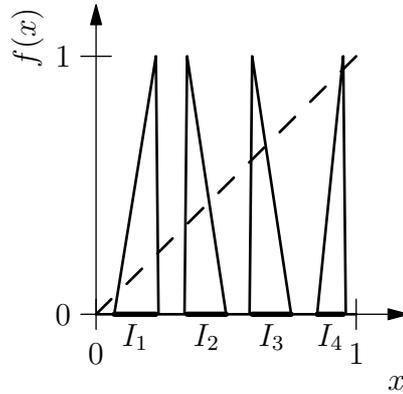


FIGURE 13. A piecewise linear map with 4 branches.

Cantor sets) have the same dynamical and topological structure—however, as we observed above, Σ_2^+ does not capture certain metric properties of the system, which vary depending on the choice of I_1 and I_2 .

Indeed, we may consider maps with more than two branches, as shown in Figure 13. Once again, we build the biggest invariant subset, a repeller C with a Cantor-like structure. The only change in this case is that because there are more intervals at each step, the *alphabet* for the symbol space is larger— $\{1, \dots, p\}$ instead of $\{1, 2\}$, where p is the number of branches of the map. Then writing

$$\Sigma_p^+ = \{1, \dots, p\}^{\mathbb{N}} = \{ \omega = (i_1, i_2, \dots) \mid i_j \in \{1, \dots, p\} \},$$

we again have a coding map $h: \Sigma_p^+ \rightarrow C$ which is a homeomorphism and which respects the dynamics. As before, different Cantor sets have the same coding, and so some structure is certainly lost in passing to the symbolic point of view.

In the end, the fact that we cannot completely restore the set C from knowledge of Σ_p^+ will not cause us to lose much sleep, because the crucial information is preserved. We will see rather more complicated examples which are modeled by these same symbolic dynamics, which turn out to contain the essence of the chaotic behaviour.

Lecture 5

a. Geometry of Cantor-like sets. Let us return for the time being to the Cantor-like set generated by a piecewise linear map defined on just two disjoint closed intervals, I_1 and I_2 , and focus on the more geometrical aspects of the situation.

The Cantor-like set C is defined as

$$(10) \quad C = \bigcap_{n \geq 1} \left(\bigcup_{(i_1, \dots, i_n)} I_{i_1 \dots i_n} \right),$$

How long are the intervals $I_{i_1 \dots i_n}$? Beginning with $n = 1$, let λ_1 and λ_2 denote the lengths of I_1 and I_2 , respectively—then writing $|I_i|$ for the length of the interval I_i , we have $|I_i| = \lambda_i$ for $i = 1, 2$.

For $n = 2$, we examine the four intervals shown in Figure 10, and recall that the ratio $|f(I_{i_1 i_2})|/|I_{i_1 i_2}|$ is given by the slope of f in $I_{i_1 i_2}$, which the comments above show to be $1/\lambda_{i_1}$. Thus the intervals $I_{i_1 i_2}$ have lengths given by

$$\begin{aligned} |I_{11}| &= \lambda_1^2, & |I_{21}| &= \lambda_2 \lambda_1, \\ |I_{12}| &= \lambda_1 \lambda_2, & |I_{22}| &= \lambda_2^2, \end{aligned}$$

where we use the fact that $f(I_{i_1 i_2}) = I_{i_2}$.

This generalises immediately to a formula for all values of n :

$$(11) \quad |I_{i_1 \dots i_n}| = \lambda_{i_1} \cdots \lambda_{i_n}.$$

Now that we know how long the intervals $I_{i_1 \dots i_n}$ ought to be, we can try to carry out the construction of the set C (or one like it) without reference to the dynamics of f . To this end, consider the following geometric construction:

- (1) Begin by choosing two disjoint closed intervals $I_1, I_2 \subset [0, 1]$ and two *ratio coefficients* $\lambda_1, \lambda_2 > 0$ with $\lambda_1 + \lambda_2 < 1$.
- (2) Put two closed intervals inside I_1 , whose lengths are $\lambda_1|I_1|$ and $\lambda_2|I_1|$; denote these by I_{11} and I_{12} , respectively.
- (3) Construct $I_{21}, I_{22} \subset I_2$ in a similar manner.
- (4) Repeat steps (2) and (3) within each of the intervals $I_{i_1 i_2}$ to construct eight disjoint closed intervals $I_{i_1 i_2 i_3}$; iterate this procedure to produce intervals $I_{i_1 \dots i_n}$ with length given by

$$|I_{i_1 \dots i_n}| = |I_{i_1}| \prod_{j=2}^n \lambda_{i_j}.$$

- (5) Construct a Cantor-like set C as the limit of this iterative procedure, as in (10).

The only difference between this construction and the earlier construction via the dynamics of f is that in the initial procedure, the position of the intervals $I_{i_1 \dots i_n}$ was determined by the dynamics, whereas here they

are free to be placed anywhere within $I_{i_1 \dots i_{n-1}}$. We can construct a coding map $h: \Sigma_2^+ \rightarrow C$ exactly as we did before, by considering the intersection $\bigcap_{n \geq 1} I_{i_1 \dots i_n}$, and we get the following diagram:

$$\begin{array}{ccc} \Sigma_2^+ & \xrightarrow{\sigma} & \Sigma_2^+ \\ \downarrow h & & \downarrow h \\ C & & C \end{array}$$

The lack of any dynamics on $C \subset [0, 1]$ means that the diagram does not close as it did in the commutative diagram in Lecture 3. We can, however, define a map $f: C \rightarrow C$ so that the diagram closes and commutes; simply take $f = h \circ \sigma \circ h^{-1}$.

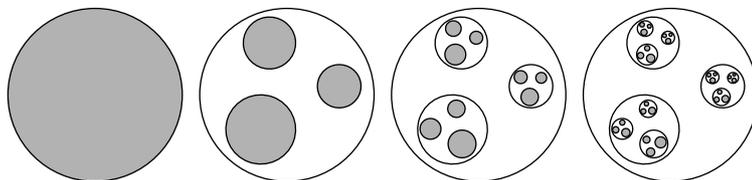
We have much less information about this artificially constructed map f than we had before, when we began with f and used it to construct C . All we can say in this case is that f is continuous, since each of σ , h , and h^{-1} are, and that it has all the dynamical properties of the shift map σ which we discussed before, such as a dense set of periodic orbits.

The geometric construction outlined above may just as well be carried out with more than two intervals; if we begin with disjoint intervals $I_1, \dots, I_p \subset [0, 1]$ and ratio coefficients $\lambda_1, \dots, \lambda_p > 0$ with $\sum_i \lambda_i < 1$, then we may build $I_{i_1 \dots i_n}$ as before, with length given by (11), and define a Cantor-like set C by (10). As in the case $p = 2$, this generalises the construction in the previous lecture of a Cantor-like set as the maximal invariant set (the repeller) of a piecewise linear map with p branches, by “forgetting” about the map f and allowing the intervals $I_{i_1 \dots i_n}$ to be placed arbitrarily within $I_{i_1 \dots i_{n-1}}$. The coding map $h: \Sigma_p^+ \rightarrow C$ is defined as before, and we may again place some dynamics on C by the formula $f = h \circ \sigma \circ h^{-1}$.

These examples illustrate the use of dynamical systems and geometric constructions as tools to study each other, which will be a prominent theme of this course. Many dynamical systems can be better understood by examining the appropriate geometric construction, and similarly, many geometric constructions are best viewed as arising from a particular dynamical system.

At this point, however, we have not yet developed the proper tools to study the geometric properties of the various Cantor-like sets we have encountered. Each has the power of the continuum (that is, it can be put into a bijective correspondence with the set of real numbers), and yet has zero “length”, in a sense which can be made precise. If we wish to use them as a tool to study the associated maps, we must somehow characterise them, but do not yet have the means to do so.

b. More general Cantor-like sets. Before examining possible ways of characterising Cantor-like sets, let us stretch our legs a bit and examine some of the other creatures in the zoo. So far we have been planted firmly in front of the cage labeled “one-dimensional constructions”, but there is no reason why we could not consider examples in higher dimensions as well.

FIGURE 14. A Cantor-like construction in \mathbb{R}^2 .

To this end, let D_1, \dots, D_p be disjoint closed discs contained in the unit disc in \mathbb{R}^2 ; the case $p = 3$ is illustrated in Figure 14. Choose ratio coefficients $\lambda_1, \dots, \lambda_p$, and carry out an iterative procedure as before; within the disc D_{i_1} , place disjoint discs $D_{i_1 i_2}$ whose diameters are $\text{diam}(D_{i_1})\lambda_{i_2}$, and so on. Taking the union over all discs corresponding to words of length n , and then taking the intersection over all $n \geq 1$, we obtain a Cantor-like set as before.

Of course, there is nothing special about discs, or about two dimensions, in this construction. The same procedure goes through for any domain in \mathbb{R}^2 , or indeed in any \mathbb{R}^n , and the end result will always be homeomorphic to Σ_p^+ . Thus we see that all these various Cantor-like sets have the same topology, despite our feeling that they must be somehow different geometrically; this reinforces our earlier point that Σ_p^+ carries no information about the geometry of the Cantor-like sets it models.

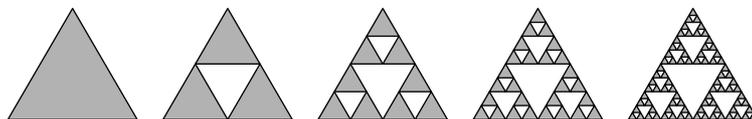


FIGURE 15. Constructing the Sierpiński gasket.

A similar construction, which has a little more built-in regularity, may be carried out by dividing an equilateral triangle into four smaller triangles, each similar to the first and congruent to each other, removing the middle triangle, and then iterating the procedure on the remaining three. The first few steps of the process are shown in Figure 15.

The fractal set C obtained as the limit of this procedure is known as the *Sierpiński gasket*¹⁰, and shares many properties with the Cantor-like sets we have been discussing, although it differs from them by being connected; that is, it has no proper non-trivial subsets which are both open and closed.

Another interpretation of this set comes from the following algorithm, which Manfred Schroeder refers to as “Sir Pinski’s game” in his book *Fractals, Chaos, Power Laws*. Given a point x inside the equilateral triangle, we define $f(x)$ by first finding the nearest vertex of the triangle to x , and then doubling the distance from x to that vertex, as shown in Figure 16.

¹⁰Or as the Sierpiński triangle, sieve, dust, carpet, or any number of other descriptive terms.

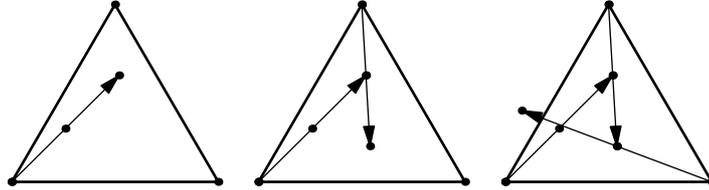


FIGURE 16. Sir Pinski's game.

Repeating the process takes us to the point $f^2(x)$, and so on, until we leave the triangle, at which point the trajectory will go off to infinity. The game, then, is to choose an initial point x whose trajectory remains in the triangle for as long as possible.

The reader may verify that the winning points, whose trajectory never leaves the triangle, are precisely the points in the Sierpiński gasket. The map f is of the same sort as we encountered earlier, when we looked at piecewise linear maps on the interval; in this case, f is a piecewise affine map on the plane, and the Sierpiński gasket is the maximal invariant set, the repeller, for f .

We can produce a coding map $h: \Sigma_3^+ \rightarrow C$ by labeling the triangles at each step of the iteration with the appropriate sequence of 1's, 2's, and 3's, and associating to each infinite sequence in Σ_3^+ the corresponding infinite intersection of nested triangles, which is just a single point. As in the one-dimensional case, the coding map completes the commutative diagram in (8).

The careful reader will by this point be howling in protest that the map f is not well defined everywhere. Indeed it is not; what are we to do with the points in C which are equidistant from two vertices of the triangle? f is supposed to double the distance from the nearest vertex, but what if that vertex is not unique? There are only 3 points in C which encounter this problem, but there are 9 more which are mapped into one of those 3, and in general, there are 3^n points in C for which f^n is not uniquely defined.

This technicality arises since C is not totally disconnected, as the Cantor-like sets we constructed earlier were. It is the same difficulty one encounters when dealing with decimal representations of the real numbers—such representations are not unique for certain numbers, namely those whose decimal expansion terminates. In this case, it is the points whose trajectory under f eventually stabilises at one of the vertices which cause trouble.

For the time being, we choose to ignore this troublesome quirk of the Sierpiński gasket, and divert prying eyes elsewhere by bedazzling the reader with a higher-dimensional version of the same thing. Instead of a triangle, begin with a tetrahedron in \mathbb{R}^3 , and remove the middle of the five congruent tetrahedra into which it decomposes. Iterating this procedure, one obtains a fractal sometimes known as the *Sierpiński sponge*.

All our examples up to this point have been linear—the building blocks at any given step of the construction are just scaled-down copies of those

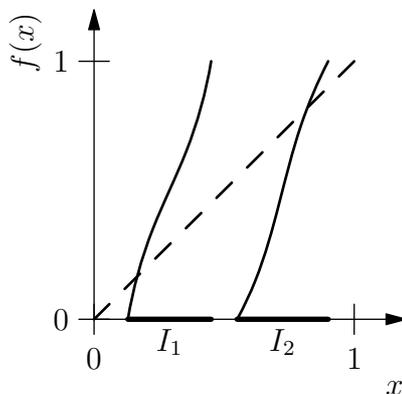


FIGURE 17. A non-linear interval map.

at the previous step. This will not always be the case in the examples of interest—most of the truly interesting phenomena in the real world, after all, are not particularly linear. Thus we may return to the one-dimensional setting and consider the sort of map $f: I_1 \cup I_2 \rightarrow [0, 1]$ shown in Figure 17, which is piecewise continuously differentiable, with $|f'(x)| \geq a > 1$, where a is a fixed real number.

Nearly everything from our previous discussion of piecewise linear maps goes through—we can follow a Cantor-like construction to obtain a repeller C for f , which will be a maximal invariant set, and which is homeomorphic to symbolic space via the coding map $h: \Sigma_2^+ \rightarrow C$, which also conjugates the dynamics of σ and f . The only thing that fails is the formula (11) for the lengths of the intervals $I_{i_1 \dots i_n}$ —a new formula can be found, but it is rather more complicated.

Or we could modify the two-dimensional construction in Figure 14; rather than shrinking the discs equally in all directions, we could contract by a factor of λ in one direction, and μ in another, so that the building blocks at successive iterations are increasingly eccentric ellipses. The topological characterisation in terms of Σ_p^+ still goes through, but the geometry is patently different from what came before.

c. Making sense of it all. Up to this point, we have been behaving like Adam, merely wandering around the Garden and naming all the animals; now we must become Linnaeus, and make some attempt at *classifying* the fractal fauna we find around us. For the various fractal sets we have described are in fact different from each other geometrically, but our usual measuring sticks are not properly equipped to distinguish them. Every example we have encountered is uncountable, and so cardinality alone is insufficient. If we try to measure “length”, the Cantor-like sets all have length zero, while the fractal curves (the von Koch curve, the fractal coastlines) have infinite length. Trying a two-dimensional notion of size, one may see that the latter have zero area; we seem to be at an impasse.

The way out of our predicament is provided by the idea of *fractal dimension*, a notion which is at once tremendously important and frustratingly elusive. Its importance will become apparent when we see how readily it lets us make sense of the thicket of examples we have presented thus far; its elusiveness is due to the fact that, put starkly, it is not defined!

Let us explain this last statement. The concept of fractal dimension was popularised by Benoit Mandelbrot in his landmark 1977 book *The Fractal Geometry of Nature*. In that work, he found the fractal dimension of a wide variety of examples, demonstrating that unlike our usual idea of dimension, it does not have to be an integer; however, the book is more concerned with exhibiting the utility of fractals as tools in various scientific contexts than it is with presenting precise mathematical definitions. Thus while this book introduced fractals to the broader scientific community, which has found uses for them in a dazzlingly wide variety of scientific models (to say nothing of the fractal artwork which has sprung up in the decades following Mandelbrot's work), it does not contain a single unifying definition of just what the fractal dimension of a set actually *is*.

In fact, the “proper” notion of fractal dimension predates Mandelbrot by over half a century, and was first introduced in 1919 by Felix Hausdorff. Now referred to as the *Hausdorff dimension*, it is one of the fundamental geometric characteristics of a set, and we will see that it lets us distinguish between the various Cantor-like sets we have met.

However, we will also see that the Hausdorff dimension is rather difficult to compute, and that there are dozens of alternative definitions of fractal dimension, which are often more tractable. Upper box dimension, lower box dimension, correlation dimension, information dimension, Lyapunov dimension—these and many others are available to us when we seek to classify fractal sets, and are all candidates for the title of “fractal dimension”. The situation is made quite messy by the fact that these various dimensional quantities may not coincide—indeed, we will see concrete examples for which they are different.

One of the primary goals of this course, then, will be to clarify the relations between the various notions of fractal dimension. These quantities are fundamental characteristics of many dynamical systems, giving a geometrical characterisation of chaotic behaviour, and so it is important to understand how they fit together. Mercifully, we will see that in many important cases, all the reasonable definitions of fractal dimension coincide, and the mess cleans itself up; this goes some way towards explaining Mandelbrot's omission of a precise definition.

Lecture 6

a. The definition of Hausdorff dimension. In this lecture, we define the notion of Hausdorff dimension for a set $Z \subset \mathbb{R}^n$; the definition, however, requires some work to set up, and so we first take some time to go through the necessary preliminaries.

Given a set $Z \subset \mathbb{R}^n$, we consider a collection $\mathcal{U} = \{U_i\}$ of open sets in \mathbb{R}^n which cover Z ; that is, for which

$$\bigcup_i U_i \supset Z.$$

Such a collection is known as an *open cover*; we will usually simply refer to a *cover*, with the implicit assumption that all elements of the cover are open. The picture to keep in mind is a collection of open balls (whose radius may vary), although of course more general open sets are allowed.

We denote the diameter of a set U_i by

$$\text{diam } U_i = \sup\{d(x, y) \mid x, y \in U_i\},$$

and the diameter of a cover by

$$\text{diam } \mathcal{U} = \sup_{U_i \in \mathcal{U}} \text{diam } U_i.$$

Fix $\varepsilon > 0$. If $\text{diam } \mathcal{U} \leq \varepsilon$, that is, if every $U_i \in \mathcal{U}$ has $\text{diam } U_i \leq \varepsilon$, then we say that \mathcal{U} is an ε -*cover*.

We will only consider covers with at most countably many elements. The reason for this is that we want to use the cover \mathcal{U} to “measure” the set Z , by assigning each element of the cover a certain (positive) weight, and then summing these weights over all elements of the cover. Clearly, though, an uncountable collection of positive numbers cannot have a finite sum.

In what follows, we consider a fixed $\alpha \geq 0$, $\varepsilon > 0$, and an ε -cover \mathcal{U} . We define (somewhat arbitrarily, it may seem) a “potential” for the entire cover \mathcal{U} by the formula

$$(12) \quad \sum_i (\text{diam } U_i)^\alpha.$$

The reader may justifiably feel that we have just pulled a rabbit from a hat without any explanation of where it came from—that will come in due course. For the time being, observe merely that (12) gives us a way of assigning a number, the “potential”, to any ε -cover \mathcal{U} . We want to use this number to characterise the set Z ; the difficulty, of course, is that the set Z admits many open covers. So how are we to decide which one to use? Which of the many possible numbers obtained from (12) properly measures Z ?

By adding unnecessary extra sets to our cover (repeating a set several times, for example), we can make the quantity in (12) arbitrarily large; thus it seems that large values of the potential are somehow to be disregarded, and we should look for the cover which minimises (12). Since an optimal

cover may not exist (the minimum may not be achieved), we consider the greatest lower bound of such quantities, and write

$$(13) \quad m(Z, \alpha, \varepsilon) = \inf \left\{ \sum_i (\text{diam } U_i)^\alpha \mid \mathcal{U} \text{ is an } \varepsilon\text{-cover of } Z \right\}.$$

Now we have a function m which depends on the set Z and the parameters α and ε . Observe that given $\varepsilon_1 > \varepsilon_2 > 0$, any ε_2 -cover is also an ε_1 -cover. Thus the set of covers over which the infimum in (13) is taken for ε_2 is a subset of the set of covers for ε_1 , and it is then immediate that

$$m(Z, \alpha, \varepsilon_2) \geq m(Z, \alpha, \varepsilon_1).$$

This shows that $m(Z, \alpha, \varepsilon)$ is monotonic as a function of ε , and hence the limit

$$(14) \quad m(Z, \alpha) = \lim_{\varepsilon \rightarrow 0} m(Z, \alpha, \varepsilon)$$

exists, although it may be ∞ , and indeed often is, as we shall see.

Now for a particular value of $\alpha \geq 0$, we have a *set function* $m(\cdot, \alpha)$; this is a real-valued function defined on the space of all subsets of \mathbb{R}^n , which assigns to a subset $Z \subset \mathbb{R}^n$ the value $m(Z, \alpha)$ defined as above. The next proposition summarises its most important properties.

PROPOSITION 4. *The set function $m(\cdot, \alpha): Z \mapsto m(Z, \alpha)$ satisfies the following properties:*

- (1) *Normalisation:* $m(\emptyset, \alpha) = 0$ for all $\alpha > 0$.
- (2) *Monotonicity:* $m(Z_1, \alpha) \leq m(Z_2, \alpha)$ whenever $Z_1 \subset Z_2$.
- (3) *Subadditivity:* $m(\bigcup_j Z_j, \alpha) \leq \sum_j m(Z_j, \alpha)$ for any finite or countable collection of subsets Z_j .

PROOF. (1) follows immediately upon observing that any open set, of any diameter, covers \emptyset .

(2) uses the same idea as in the proof of monotonicity of $m(Z, \alpha, \cdot)$; an ε -cover of Z_2 is an ε -cover of Z_1 , and hence the infimum in $m(Z_2, \alpha, \varepsilon)$ is being taken over a smaller set.

(3) is slightly more involved, and requires the following technical lemma, which is an immediate consequence of the definitions of limit and infimum.

LEMMA 5. *Fix $Z \subset \mathbb{R}^2$ and $\alpha \geq 0$ such that $m(Z, \alpha) < \infty$. For every $\delta > 0$ and $\varepsilon > 0$, there exists an open ε -cover $\mathcal{U} = \{U_i\}$ of Z such that $|m(Z, \alpha) - \sum_i (\text{diam } U_i)^\alpha| \leq \delta$.*

To prove (3), we first observe that if any of the values $m(Z_j, \alpha)$ is infinite, then their sum is infinite, and the inequality is trivial. Thus we assume they are all finite; fixing $\delta > 0$ and writing $Z = \bigcup_j Z_j$, we may apply the lemma to each Z_j to obtain ε -covers $\mathcal{U}_j = \{U_{ji}\}_i$ (for arbitrarily small $\varepsilon > 0$) such that

$$|m(Z_j, \alpha) - \sum_i (\text{diam } U_{ji})^\alpha| \leq \frac{\delta}{2^j}.$$

We see that $\mathcal{U} = \bigcup_j \mathcal{U}_j = \{U_{ji}\}_{i,j}$ is an open cover of Z , and since each of the U_{ji} has diameter $\leq \varepsilon$, it is actually an ε -cover. Thus

$$\begin{aligned} m(Z, \alpha, \varepsilon) &\leq \sum_{i,j} (\text{diam } U_{ji})^\alpha \\ &= \sum_j \left(\sum_i (\text{diam } U_{ji})^\alpha \right) \\ &\leq \sum_j \left(m(Z_j, \alpha) + \frac{\delta}{2^j} \right) \\ &= \sum_j m(Z_j, \alpha) + \delta, \end{aligned}$$

This holds for all $\delta > 0$ and for all $\varepsilon > 0$; hence (3) holds. \square

So far the parameter α has been listed among the *dramatis personae*, but has done little more than linger at the edge of the stage, constant and unchanging. Its appearance on centre stage will finally bring us to the definition of Hausdorff dimension.

To that end, then, let us consider $m(Z, \cdot): [0, +\infty) \rightarrow [0, +\infty]$ as a function of α . As they say, a picture is worth a thousand words, and so we try to draw its graph.

There are three possibilities for the value of $m(Z, \alpha)$ at any given α —it may be 0, it may be ∞ , or it may be finite. The former two are not particularly interesting; after all, our main grievance with the ideas of cardinality, length, area, etc. as tools for classifying fractals was that they always returned answers which were either 0 or ∞ .

The third possibility, that $0 < m(Z, \alpha) < \infty$ for a particular value of α , turns out to have rather drastic consequences, as we see in the following two propositions.

PROPOSITION 6. *If $\alpha \geq 0$ is such that $m(Z, \alpha) < \infty$, then $m(Z, \beta) = 0$ for every $\beta > \alpha$.*

PROOF. A straightforward computation shows that

$$\begin{aligned} m(Z, \beta, \varepsilon) &= \inf_{\mathcal{U}} \sum_i (\text{diam } U_i)^\beta \\ &= \inf_{\mathcal{U}} \sum_i (\text{diam } U_i)^{\beta-\alpha} (\text{diam } U_i)^\alpha \\ &\leq \inf_{\mathcal{U}} \sum_i \varepsilon^{\beta-\alpha} (\text{diam } U_i)^\alpha \\ &= \varepsilon^{\beta-\alpha} m(Z, \alpha, \varepsilon), \end{aligned}$$

and since $\beta - \alpha > 0$, we have $\varepsilon^{\beta-\alpha} \rightarrow 0$. Since $m(Z, \alpha, \varepsilon) \leq m(Z, \alpha) < \infty$, this implies that $m(Z, \beta, \varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$. \square

As an immediate consequence of this proposition, we have the following dual statement:

PROPOSITION 7. *If $\alpha \geq 0$ is such that $m(Z, \alpha) > 0$, then $m(Z, \beta) = \infty$ for every $\beta < \alpha$.*

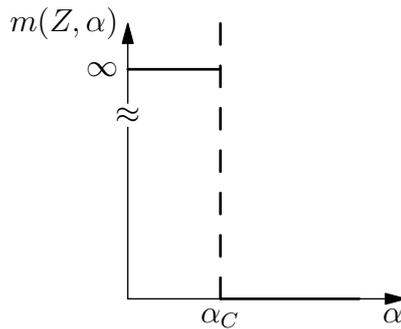


FIGURE 18. The graph of $m(Z, \cdot)$.

It follows from these propositions that the graph of $m(Z, \cdot)$ is as shown in Figure 18; below some critical value α_C , the function takes infinite values, and for $\alpha > \alpha_C$, we have $m(Z, \alpha) = 0$. Thus the function $m(Z, \cdot)$ is entirely determined by the location of α_C and the value of $m(Z, \alpha_C)$; the latter may lie anywhere in $[0, \infty]$, while the former may take any value in $[0, \infty)$. Just which values they take, of course, depends on the set Z —this is the whole point.

We are now in a position to complete our definition—the Hausdorff dimension of a set Z , denoted $\dim_H Z$, is the critical value α_C at which the function $m(Z, \cdot)$ passes from ∞ to 0. Thus, we have

$$\begin{aligned} \dim_H Z &= \sup\{\alpha \in [0, \infty) \mid m(Z, \alpha) = \infty\} \\ &= \inf\{\alpha \in [0, \infty) \mid m(Z, \alpha) = 0\}. \end{aligned}$$

So we have the definition! But what does it mean? In what sense is this the “dimension” of the set Z ? Does it agree with our usual intuitive ideas of dimension? What properties does it have? How do we actually compute it for specific examples? What does it have to do with fractals? Where in the world does the function $m(Z, \alpha, \varepsilon)$ come from? We will address these questions in the next lecture, and see that they do in fact have satisfactory answers.

Lecture 7

a. Properties of Hausdorff dimension. Now that we have a definition of Hausdorff dimension and have seen what's under the hood, we take this new notion out for a test drive, and see how it behaves. Some important properties of Hausdorff dimension can be deduced as immediate consequences of the corresponding properties of the set function $m(\cdot, \alpha)$ given in Proposition 4:

PROPOSITION 8. *The Hausdorff dimension has the following basic properties:*

- (1) $\dim_H \emptyset = 0$.
- (2) $\dim_H Z_1 \leq \dim_H Z_2$ whenever $Z_1 \subset Z_2$.
- (3) $\dim_H(\bigcup_i Z_i) = \sup_i \dim_H Z_i$, where $\{Z_i\}$ is any countable collection of subsets of \mathbb{R}^n .

A singleton set $Z = \{x\}$ has $m(Z, \alpha, \varepsilon) = 0$ for all $\alpha > 0$, $\varepsilon > 0$, and so applying the third property above, we obtain

COROLLARY 9. *If Z is countable, then $\dim_H Z = 0$.*

Thus the set of rational numbers has Hausdorff dimension zero, despite being dense in the interval, and hence fairly “large” in the topological sense.

So points have zero Hausdorff dimension, which we would expect. What about lines and planes? Do they have the “correct” Hausdorff dimension? Before answering this question, we state two lemmas which codify a common technique for giving upper and lower bounds on the Hausdorff dimension.

LEMMA 10. *We have that $\dim_H Z \leq \alpha$, if there exists $C > 0$ such that for all $\varepsilon > 0$, there exists an ε -cover $\mathcal{U} = \{U_i\}$ with $\sum_i (\text{diam } U_i)^\alpha \leq C$.*

PROOF. The condition given guarantees that $m(Z, \alpha, \varepsilon) \leq C$ for all $\varepsilon > 0$; hence $m(Z, \alpha) < \infty$, and the result follows. \square

LEMMA 11. *We have that $\dim_H Z \geq \alpha$, if there exists $C > 0$ and $\varepsilon > 0$ such that $\sum_i (\text{diam } U_i)^\alpha \geq C$ for all ε -covers $\mathcal{U} = \{U_i\}$.*

PROOF. The condition given guarantees that $m(Z, \alpha, \varepsilon) \geq C$ for some $\varepsilon > 0$, hence $m(Z, \alpha) > 0$, and the result follows. \square

We now apply these lemmas to the case of the real line.

PROPOSITION 12. $\dim_H \mathbb{R} = 1$.

PROOF. Let $Z = [0, 1]$ be the unit interval in \mathbb{R} . By part (3) of Proposition 8, it suffices to show that $\dim_H Z = 1$. We do this by using Lemma 10 to show that $\dim_H Z \leq 1$, and then using Lemma 11 to show that $\dim_H Z \geq 1$.

To satisfy the condition of Lemma 10, we consider $\varepsilon > 0$ and choose an integer n such that $1/n \leq \varepsilon$. Consider the intervals $(i/3n, (i+1)/3n)$ for $i = 0, \dots, 3n-1$; these cover every point of $[0, 1]$ except the endpoints $i/3n$. If we extend each interval to include the two beside it, then we get

the intervals $U_i = ((i-1)/3n, (i+2)/3n)$, each of which has length $1/n \leq \varepsilon$, and so $\mathcal{U} = \{U_i\}$ is an ε -cover of Z . It has $3n$ elements, and so we see that

$$\sum_i (\text{diam } U_i) = 3n \cdot \frac{1}{n} = 3.$$

Thus Lemma 10 applies with $C = 3$ and $\alpha = 1$, and we have $\dim_H Z \leq 1$.

The other direction is rather harder, reflecting the general fact that upper bounds on the Hausdorff dimension are usually easier to obtain than lower bounds. The reason for this is that in order to apply Lemma 10, we only needed to construct a single “good” partition for each $\varepsilon > 0$; in order to get a lower bound by applying Lemma 11, we need to deal with *every* ε -cover for some sufficiently small ε .

While we cannot apply Lemma 11 directly to our present problem with $\alpha = 1$, we can reach the same result by showing that the lemma applies for every $\alpha < 1$. Indeed, then we will have $\dim_H Z \geq \alpha$ for all $\alpha < 1$, which of course implies that $\dim_H Z \geq 1$, as desired.

To that end, fix $\alpha < 1$. We wish to find $\varepsilon > 0$ such that $\sum_i (\text{diam } U_i)^\alpha > 1$ for every ε -cover \mathcal{U} . For any such \mathcal{U} , we have

$$\begin{aligned} \sum_i (\text{diam } U_i)^\alpha &= \sum_i (\text{diam } U_i) (\text{diam } U_i)^{\alpha-1} \\ &\geq \left(\sum_i \text{diam } U_i \right) \varepsilon^{\alpha-1}. \end{aligned}$$

It is relatively easy to see that since the sets U_i cover $[0, 1]$, we must have $\sum_i \text{diam } U_i \geq 1$. Indeed, if we write $a_i = \inf U_i$ and $b_i = \sup U_i$, then we have $U_i \subset (a_i, b_i)$, with $\text{diam } U_i = b_i - a_i$ (where we take $[0, b_i)$ and $(a_j, 1]$ in the cases $a_i = 0$, $b_j = 1$), and so the argument reduces to the case where each U_i is an interval.¹¹ By compactness of $[0, 1]$, we may find i_1, \dots, i_n such that

$$0 = a_{i_1} < a_{i_2} < b_{i_1} < a_{i_3} < b_{i_2} < \dots < a_{i_n} < b_{i_{n-1}} < b_{i_n} = 1,$$

¹¹More generally, we may observe that each U_i is a finite or countable union of intervals, by letting $I_i(x)$ denote the largest open interval containing x which is a subset of U_i , and verifying that $I_i(x)$ and $I_i(y)$ either coincide or are disjoint. There are at most countably many different $I_i(x)$ since each contains a distinct rational number.

at which point we see that

$$\begin{aligned}
 \sum_i \operatorname{diam} U_i &= \sum_i b_i - a_i \\
 &\geq \sum_{k=1}^n b_{i_k} - a_{i_k} \\
 &> \left(\sum_{k=1}^{n-1} a_{i_{k+1}} - a_{i_k} \right) + b_{i_n} - a_{i_n} \\
 &= b_{i_n} - a_{i_1} = 1.
 \end{aligned}$$

Thus $\sum_i (\operatorname{diam} U_i)^\alpha \geq \varepsilon^{\alpha-1}$. Since $\alpha < 1$, we have $\varepsilon^{\alpha-1} > 1$ for sufficiently small $\varepsilon > 0$. Then Lemma 11 applies, showing that $\dim_H Z \geq \alpha$, and since $\alpha < 1$ was arbitrary, we have $\dim_H Z \geq 1$, which completes the proof. \square

This result shows that Hausdorff dimension gives the result we would expect for a line; in fact, a more general argument along the same lines shows that if $Z \subset \mathbb{R}^n$ is any open set (indeed, any set with non-empty interior), then $\dim_H Z = n$. In particular, $\dim_H \mathbb{R}^n = n$, and so Hausdorff dimension agrees with our usual definition of dimension for lines, planes, and so on.

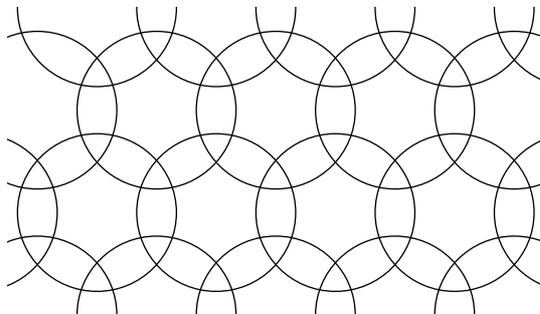
But what is our usual definition of dimension? Of course we know that the n -dimensional Euclidean space \mathbb{R}^n has, well, dimension n , but why? What is it about this space that makes it n -dimensional?

b. Topological dimension. Consider an open cover \mathcal{U} of \mathbb{R}^n . Fix a point $x \in \mathbb{R}^n$, and count the number of elements of the cover which contain x ; we call this the *multiplicity of \mathcal{U} at x* , and denote it by $M(\mathcal{U}, x)$. The quantity $M(\mathcal{U}) = \sup_x M(\mathcal{U}, x)$ is called the *multiplicity of \mathcal{U}* ; as with our earlier definition of $m(Z, \alpha, \varepsilon)$ via covers, we can make $M(\mathcal{U})$ as large as we like (even infinite) by choosing a cover with “too many” elements, and so we want a cover for which $M(\mathcal{U})$ is minimal. Thus we define $M(\mathbb{R}^n) = \inf_{\mathcal{U}} M(\mathcal{U})$, and investigate how this quantity is connected to the dimension of \mathbb{R}^n .

In the case $n = 1$, we are just covering the line, and it is easy to construct a cover \mathcal{U} with $M(\mathcal{U}) = 2$. Since we must have intersections between the elements of the cover (otherwise \mathbb{R} would be disconnected), we see that $M(\mathbb{R}) = 2$.

In the plane, one may construct a cover by open discs which has multiplicity 3, as shown in Figure 19, and it is not too hard to show that this is optimal, so $M(\mathbb{R}^2) = 3$. Similarly, the minimal multiplicity for a cover of \mathbb{R}^n is $n + 1$, which connects $M(\mathbb{R}^n)$ to the dimension of \mathbb{R}^n . All that is required for the definition of M , however, is a topological space, and so we may make the following definition.

DEFINITION 13. The *topological dimension* of a topological space X is the quantity $M(X) - 1$; that is, the dimension is one less than the multiplicity of an optimal cover.

FIGURE 19. A cover of \mathbb{R}^2 with multiplicity 3.

The discussion at the beginning of this lecture shows that Hausdorff dimension and topological dimension agree when it comes to the Euclidean spaces \mathbb{R}^n ; is this always the case? Have we just given two rather different definitions of the same quantity?

c. Comparison of Hausdorff and topological dimension. One difference in the two definitions is immediately apparent; the topological dimension is always an integer, while the Hausdorff dimension can take any non-negative real value, as we will eventually see.

Another difference becomes apparent if we look at what notions are used in the definitions; the topological dimension can be defined for any topological space, whether or not it has a metric, while the Hausdorff dimension requires a metric for its definition. Of course, a single topological space can be equipped with multiple metrics; for example, the usual metric on \mathbb{R}^n is given by Pythagoras' formula

$$(15) \quad d(x, y) = \sqrt{\sum_i (x_i - y_i)^2},$$

but other metrics may be introduced by the formulae

$$(16) \quad \rho(x, y) = \sum_i |x_i - y_i|,$$

$$(17) \quad \sigma(x, y) = \max_i |x_i - y_i|.$$

We say that two metrics d and ρ on a space X are *equivalent* if they induce the same topology; that is, if the identity map from (X, d) to (X, ρ) is a homeomorphism. All three of the metrics just introduced for \mathbb{R}^n are equivalent, and so they lead to the same topological dimension. Do they also lead to the same Hausdorff dimension?

This question cannot be answered by topology alone; we need a stronger relationship between the metrics. We say that d and ρ are *strongly equivalent* if there exists a positive constant C such that

$$(18) \quad \frac{1}{C}\rho(x, y) \leq d(x, y) \leq C\rho(x, y)$$

for all x, y in the space.

EXERCISE 5. Show that d , ρ , and σ as defined above on \mathbb{R}^n are all strongly equivalent.

So what happens to the Hausdorff dimension under strongly equivalent metrics? Recalling the definition of $m(Z, \alpha, \varepsilon)$, we see that passing to a strongly equivalent metric changes each term $(\text{diam } U_i)^\alpha$ by at most a factor of C ; hence $m(Z, \alpha, \varepsilon)$ and $m(Z, \alpha)$ are changed by at most a factor of C . This does not affect whether or not $m(Z, \alpha)$ is finite or infinite for a particular α , so the critical value α_C , and hence the Hausdorff dimension, remains unchanged.

Lecture 8

a. Some point set topology. We have, of necessity, been dabbling in topological notions at various points in the preceding lectures; for the benefit of the reader who does not have much background in point set topology, or who may be aided by a brief review of the relevant definitions and results, we give a (rather hasty) sketch of the basic outlines of the subject in this lecture.

There are various senses in which two metrics d_1 and d_2 on \mathbb{R}^n (or more generally, on any metric space X) may be said to be “the same”; we give three such notions, from weakest to strongest.

DEFINITION 14. d_1 and d_2 are *equivalent* (denoted $d_1 \sim d_2$) if the identity map $\text{Id}: (X, d_1) \rightarrow (X, d_2)$ is a homeomorphism; that is, if for every $x \in X$ and every $\varepsilon > 0$ there exists $\delta > 0$ such that $d_1(x, y) \leq \delta$ implies $d_2(x, y) \leq \varepsilon$, and $d_2(x, y) \leq \delta$ implies $d_1(x, y) \leq \varepsilon$.

d_1 and d_2 are *uniformly equivalent* if the identity map and its inverse are both *uniformly* continuous; that is, if for every $\varepsilon > 0$ there exists $\delta > 0$ (independent of x, y) such that the implications in the previous paragraph hold.

d_1 and d_2 are *strongly equivalent* if there exists $C > 0$ such that for all x, y , we have

$$C^{-1}d_2(x, y) \leq d_1(x, y) \leq Cd_2(x, y).$$

Strong equivalence implies uniform equivalence, which in turn implies equivalence, but neither of the reverse implications holds in general. For the time being, we focus on the weakest notion, that of plain old unadorned equivalence.

The statement that d_1 and d_2 are equivalent may be rephrased as the statement that every d_1 -ball contains a d_2 -ball, and vice versa. In particular, since an open set U is one which contains a sufficiently small ball around every point in U , we see that the metrics d_1 and d_2 define precisely the same collection of open sets; that is, they define the same topology.

DEFINITION 15. Let X be a set, and \mathcal{T} a collection of subsets of X such that:

- (1) $\emptyset, X \in \mathcal{T}$.
- (2) If $U, V \in \mathcal{T}$ then $U \cap V \in \mathcal{T}$.
- (3) If $\{U_\alpha\} \subset \mathcal{T}$ then $\bigcup_\alpha U_\alpha \in \mathcal{T}$, where the union may be taken over any collection of sets in \mathcal{T} , countable or not.

Then the pair (X, \mathcal{T}) is a *topological space*, and \mathcal{T} is the *topology* on X . The sets in \mathcal{T} are referred to as *open sets*; given a point $x \in X$, an open set containing x is a *neighbourhood* of x .

This definition codifies these three properties of open sets in \mathbb{R}^n as the only axioms which need to be satisfied in order for topological concepts such as convergence, compactness, and continuity to be used. Once we know which sets are open, all these concepts are in play.

It also suffices to know which sets are *closed*; we say that $A \subset X$ is closed if its complement $X \setminus A$ is open. There are various equivalent definitions of closed sets. For example, $x \in X$ is an *accumulation point* for $A \subset X$ if every neighbourhood of x contains a point in A . Then the *closure* of A , denoted \bar{A} , is the set of all accumulation points of A , and A is closed if and only if $A = \bar{A}$.

Conspicuously lacking from the definition of a topological space is any notion of *size*. We have no way to say that a particular neighbourhood of x is “small”, or to compare neighbourhoods of different points. Clearly, then, many concepts from calculus and real analysis cannot be stated in the context of a general topological space; we can go a long way and obtain many analytic results using only topological methods, but we cannot do everything. The definition of a topological space is general enough to permit some behaviour which is rather pathological from the point of view of the standard topology on \mathbb{R}^n .



FIGURE 20. Vacationing in France.

As an example of the sort of situation that can occur, let X be the unit disc in \mathbb{R}^2 , and introduce a new metric d by the formula

$$d(x, y) = \begin{cases} \|x - y\| & \text{if } x \text{ is a scalar multiple of } y, \\ \|x\| + \|y\| & \text{otherwise,} \end{cases}$$

where $\|\cdot\|$ denotes the usual norm on \mathbb{R}^2 , $\|x\| = \sqrt{x_1^2 + x_2^2}$. This is not quite so unnatural an example to consider as it may appear; suppose you are on holiday in France, and wish to take the high-speed train (the TGV) from Marseille, on the Mediterranean, to Nantes, near the Atlantic (see Figure 20). Because there is no direct TGV line between the two cities, you must go via Paris (which corresponds to the origin in this example), and so

as far as your travels are concerned, the distance from Marseille to Nantes is found by adding together the distance from Marseille to Paris ($\|x\|$) and the distance from Paris to Nantes ($\|y\|$).

On the other hand, if you wish to go from Marseille to Lyon, then there is no need to go all the way to Paris first, since the two cities lie on the same branch of the rail system, and so the distance is given by the usual formula ($\|x - y\|$).

The new metric induces a rather different topology on the disc than the one we are used to—for example, given any $x \neq 0$ in the disc, the interval

$$((1 - \varepsilon)x, (1 + \varepsilon)x) = \{rx \mid 1 - \varepsilon < r < 1 + \varepsilon\}$$

is open for every $\varepsilon > 0$, although these sets were neither open nor closed in the usual topology. A striking distinction between the two spaces is given by the notion of *separability*—a topological space is *separable* if it has a countable dense subset. With the usual metric, the disc is a separable metric space (consider points with rational coordinates), but in this new metric, there is no countable dense subset, and so the space is not separable.

Another metric space which has unusual properties vis-à-vis \mathbb{R}^n is the symbolic space Σ_p^+ , equipped with the metric d_α defined in (9). Like \mathbb{R}^n , this space is separable, but recalling the characterisation of open and closed sets in terms of cylinders, we quickly see that open sets are also closed, and vice versa! A set which is both open and closed is said to *disconnect* a topological space; a *connected* space is one which cannot be disconnected. Both the symbolic space Σ_p^+ and the Cantor set C are *totally disconnected*; they have no non-trivial connected subsets. This is indeed different from the connectedness that we see in \mathbb{R}^n .

b. Metrics and topologies. All metric spaces carry a topology which is induced by the metric, and so far, all our examples of topological spaces have been of this form. One may rightly ask, then, if every example arises this way; given a topological space (X, \mathcal{T}) , can we always find a metric d on X such that the sets in \mathcal{T} are precisely those sets which are unions of d -balls? Such a space is called *metrisable*, and so we may ask, are all topological spaces metrisable?

It turns out that the answer is “no”—some topologies do not come from metrics. But which ones? Given a particular topology, how can we tell whether or not it comes from a metric?

If (X, d) is a metric space, and $A, B \subset X$ are any closed sets, then we can find open sets $U, V \subset X$ such that $U \supset A$, $V \supset B$, and $U \cap V = \emptyset$. A topological space with this property is called *normal*—there are a number of related properties, known as *separation axioms*, which more or less describe what sort of sets (single points, closed sets, etc.) can be separated by open sets in the manner described above. For example, if the above property holds when A and B are replaced by single points $\{x\}$ and $\{y\}$, then the space is

called *Hausdorff*; it is easy to see that metric spaces have this property as well.

Neither the Hausdorff property nor the normal property follows immediately from the definition of a topological space. For example, the rather small collection of subsets $\mathcal{T} = \{\emptyset, X\}$ defines the *trivial topology* on a space X , which is not Hausdorff, and hence not metrisable. Or we might take two arbitrary sets $A, B \subset X$, and consider the smallest topology with respect to which both A and B are open:

$$\mathcal{T} = \{\emptyset, A, B, A \cap B, A \cup B, X\}.$$

(Note that $A \cap B$ may coincide with \emptyset , and $A \cup B$ may coincide with X .) This too is not Hausdorff, and hence not metrisable.

On the other hand, if a topological space *is* both normal and Hausdorff, then this is almost enough to make it metrisable.

In order to state what more is needed, we recall that in any metrisable topology (X, \mathcal{T}) , an open set can always be written as a union of balls $B(x, r)$. Thus although there are many open sets which are not of the form $B(x, r)$, the collection of balls is sufficient to generate the topology. More generally, any collection of open sets $\mathcal{B} \subset \mathcal{T}$ with the property that any element of \mathcal{T} can be written as a union of elements of \mathcal{B} is known as a *base* (or *basis*) of the topology.¹²

This is a global concept—it applies to the entire topological space. The local version is that of a *neighbourhood base*, which is a collection $\mathcal{B}(x)$ of neighbourhoods of $x \in X$ such that any neighbourhood of x contains some element of $\mathcal{B}(x)$. In a metric space, the collection

$$\mathcal{B}(x) = \{B(x, r) \mid r > 0\}$$

is a neighbourhood base at each $x \in X$. In fact, it suffices to consider any sequence of values for r which goes to 0; in particular, $\{B(x, 1/n) \mid n \in \mathbb{N}\}$ is a *countable* neighbourhood base. A topological space which admits a countable neighbourhood base at each point is said to be *first countable*, and so we see that metric spaces are first countable.

In the familiar case of \mathbb{R}^n , we may consider all balls of rational radius centred at points with rational coordinates—this forms a countable *base* for the topology, not just a neighbourhood base. A topological space with a countable base is called *second countable*, and one sees immediately that every *separable* metric space is second countable.

This flurry of definitions allows us to state (without proof) one of the most important results in basic point set topology, which gives a nearly complete answer to the question of which topological spaces are metrisable.

THEOREM 16 (Urysohn's metrisation theorem). *If (X, \mathcal{T}) is a second countable, normal, Hausdorff topological space, then X is metrisable.*

¹²Equivalently, a base may be characterised by the requirement that every open set contain a member of the base.

Actually, slightly stronger versions of this theorem are available, but this will be enough for our purposes. In a nutshell, then, the moral of the story is that *all topological spaces are metrisable, except for rather weird cases with which we will not concern ourselves.*

So we are interested in topological spaces whose topologies come from some metric. But which metric? As we saw in Example 5, the three metrics in (15)–(17) all lead to the same topology on \mathbb{R}^n , and in general, a single topology on a space X may be induced by many different, but equivalent, metrics.

With this in mind, we return to our discussion of the topological and Hausdorff dimensions of a set $Z \subset \mathbb{R}^n$ with metric ρ , denoting the latter by $\dim_{H,\rho} Z$ to indicate the dependence on the metric. We continue to write d for the standard metric given by (15).

The relationship between the two notions of dimension is given by the following deep theorem due to Hausdorff; this pearl of dimension theory was the motivation for his introduction of the notion of Hausdorff dimension.

THEOREM 17 (Hausdorff's theorem). *Given a set $Z \subset \mathbb{R}^n$, the topological dimension $\dim Z$ and the Hausdorff dimensions $\dim_{H,\rho} Z$ are related by the following variational principle:*

$$(19) \quad \dim Z = \inf_{\rho \sim d} \dim_{H,\rho} Z.$$

That is, the topological dimension is the infimum of the possible Hausdorff dimensions, taken over all metrics ρ which are equivalent to the standard metric d .

In fact, we will eventually see an even better result than this. For the time being, let us consider a continuous map $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$, and examine the relationship between $\dim_H Z$ and $\dim_H f(Z)$, which on the face of it is a different matter than what we have been discussing.

A simple example where Z and $f(Z)$ have different Hausdorff dimensions is the case where f is a projection to a subspace, and so $\dim_H f(Z) < \dim_H Z$ for an appropriate choice of Z . In general, the Hausdorff dimension can also increase under the action of f ; we will see later that the von Koch curve has Hausdorff dimension strictly greater than 1, despite being the homeomorphic image of the unit interval.

If f is Lipschitz, however, the story is different. We first observe that in the definition of Hausdorff dimension, it suffices to consider covers by open balls $B(x, r)$ with $r \leq \varepsilon$.

PROPOSITION 18. *Define a function m_B by*

$$m_B(Z, \alpha, \varepsilon) = \inf \left\{ \sum_i (2r_i)^\alpha \mid \mathcal{U} = \{B(x_i, r_i)\} \text{ is an } \varepsilon\text{-cover of } Z \right\},$$

for any $Z \subset \mathbb{R}^n$, $\alpha \geq 0$, and $\varepsilon > 0$. As before, define $m_B(Z, \alpha) = \lim_{\varepsilon \rightarrow 0} m_B(Z, \alpha, \varepsilon)$; then we have

$$\begin{aligned} \dim_H(Z) &= \inf\{\alpha \geq 0 \mid m_B(Z, \alpha) = 0\} \\ &= \sup\{\alpha \geq 0 \mid m_B(Z, \alpha) = \infty\}. \end{aligned}$$

PROOF. We are taking the infimum over a smaller collection of covers than in (13), and so the inequality $m_B \geq m$ is automatic. Now given an ε -cover \mathcal{U} of Z by arbitrary open sets of small diameter, observe that for each $U_i \in \mathcal{U}$, we may take $x_i \in U_i$ and $r_i = \text{diam } U_i$, and then we have

$$U_i \subset B(x_i, r_i),$$

from which it follows that

$$2^\alpha m_B(Z, \alpha, 2\varepsilon) \leq m(Z, \alpha, \varepsilon) \leq m_B(Z, \alpha, \varepsilon).$$

Taking the limit as $\varepsilon \rightarrow 0$, we have

$$2^\alpha m_B(Z, \alpha) \leq m(Z, \alpha) \leq m_B(Z, \alpha),$$

and so the critical value α_C is the same for both m and m_B . \square

We now address the case where f is Lipschitz.

PROPOSITION 19. *If $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz, then $\dim_H f(Z) \leq \dim_H Z$ for every $Z \subset \mathbb{R}^n$.*

PROOF. Let $L > 0$ be such that $d(f(x), f(y)) \leq Ld(x, y)$. Then if $\mathcal{U} = \{U_i\}$ is any ε -cover of Z , we have $\text{diam } f(U_i) \leq L \text{diam } U_i$, and so $f(\mathcal{U}) = \{f(U_i)\}$ is an $L\varepsilon$ -cover of $f(Z)$, for which

$$\sum_i (\text{diam } f(U_i))^\alpha \leq L^\alpha \sum_i (\text{diam } U_i)^\alpha.$$

It follows that $m(f(Z), \alpha, L\varepsilon) \leq L^\alpha m(Z, \alpha, \varepsilon)$, whence $m(f(Z), \alpha) \leq L^\alpha m(Z, \alpha)$. Thus if $m(Z, \alpha)$ is finite, so is $m(f(Z), \alpha)$, which implies that $\dim_H f(Z) \leq \dim_H Z$. \square

A bijection f such that both f and f^{-1} are Lipschitz is known as a *bi-Lipschitz* function. It follows from Proposition 19 that $\dim_H f(Z) = \dim_H Z$ whenever f is bi-Lipschitz.

All this is actually quite related to our previous discussion of the dependence of Hausdorff dimension on the metrics. If f is a homeomorphism and ρ is a metric on \mathbb{R}^n , then the formula $\rho_f(x, y) = \rho(f(x), f(y))$ defines a new metric ρ_f on \mathbb{R}^n , and we see that

$$\dim_{H, \rho_f} Z = \dim_H f(Z).$$

With this correspondence between bijections and changes of metric, we see that the bi-Lipschitz functions are precisely those for which ρ and ρ_f are strongly equivalent; indeed, the definition of strong equivalence boils down to the statement that the identity map is bi-Lipschitz. Thus Hausdorff

dimension is preserved by strong equivalence, but not by topological equivalence; the latter preserves topological dimension, which by Theorem 17 is the infimum of the possible Hausdorff dimensions.

Lecture 9

a. Topological consequences. Even though the Hausdorff dimension depends on the metric, we can still occasionally use it to deduce some purely topological information.

THEOREM 20. *If $Z \subset \mathbb{R}^n$ has $\dim_H Z < 1$, then Z is totally disconnected.*

Before we prove this statement, recall that Z is *connected* if it cannot be written as a union of non-empty disjoint open sets. There is a stronger notion of *path connectedness*, which requires that we be able to connect any two points x, y by a path (that is, a continuous image of $[0, 1]$), but the former definition is more important for our purposes. Z is *totally disconnected* if for every $x, y \in Z$, $x \neq y$, we can write Z as a union of two non-empty disjoint open sets U_1, U_2 such that $x \in U_1$, $y \in U_2$.

Any Cantor-like set is totally disconnected, as is the symbolic space Σ_p^+ ; this follows since cylinders are both open and closed.

In order to prove Theorem 20, we first need to say a few words about *Lebesgue measure* on $[0, 1]$. This is a generalisation of the notion of “length” to sets which are not intervals. Given a set $Z \subset [0, 1]$, we will write $\text{Leb}(Z)$ for the Lebesgue measure of Z ; if Z is an interval $[a, b]$, then the Lebesgue measure, is just the length,

$$(20) \quad \text{Leb}([a, b]) = b - a.$$

Endpoints do not contribute to the Lebesgue measure of an interval; (20) also applies to intervals of the form (a, b) , $(a, b]$, and $[a, b)$.

If I_1 and I_2 are two intervals, they may either overlap or be disjoint. If they overlap, then their union is again an interval, and so $\text{Leb}(I_1 \cup I_2)$ is given by (20). If they are disjoint, we define

$$\text{Leb}(I_1 \cup I_2) = \text{Leb}(I_1) + \text{Leb}(I_2).$$

In fact, we define Lebesgue measure this way for any set Z which is a union of countably many disjoint intervals:

$$(21) \quad \text{Leb} \left(\bigcup_i I_i \right) = \sum_i \text{Leb}(I_i).$$

Each partial sum over finitely many disjoint intervals is less than 1, and so the infinite sum converges.

In order to generalise (21) to sets which are not countable unions of intervals, we cover them with such unions. Paralleling the definition of $m(Z, \alpha)$, we define the Lebesgue measure of a set $Z \subset [0, 1]$ by

$$(22) \quad \text{Leb}(Z) = \inf \left\{ \sum_i \text{Leb}(I_i) \mid \{I_i\} \text{ is a cover of } Z \text{ by open intervals} \right\}.$$

Indeed, it follows from (22) that $\text{Leb}(Z) = m(Z, 1)$, and so the Hausdorff function $m(\cdot, 1)$ is just Lebesgue measure. In particular, we see if a set Z

has Hausdorff dimension strictly less than 1, then Z admits interval covers of arbitrarily small total length, hence $\text{Leb}(Z) = 0$.

PROOF OF THEOREM 20. To show that Z is totally disconnected, we consider arbitrary points $x, y \in Z$, and produce two open sets $U, V \subset Z$ such that $x \in U$, $y \in V$, $U \cap V = \emptyset$, and $U \cup V = Z$. To this end, define a function $f: Z \rightarrow \mathbb{R}^+$ by $f(z) = d(x, z)$; that is, f measures the distance from x to z .

Clearly, f is a Lipschitz function, and so Proposition 19 gives the bound

$$\dim_H f(Z) \leq \dim_H Z < 1,$$

and it follows from our earlier remarks that $\text{Leb}(f(Z)) = 0$. But this implies that $f(Z)^c$, the complement of $f(Z)$, is dense; indeed, if it were not dense, then $f(Z)$ would contain an interval, and hence have positive Lebesgue measure.

Thus we may find $r \in f(Z)^c$ such that $0 < r < f(y)$, and we define two open sets by

$$U = f^{-1}([0, r)) = \{z \in Z \mid d(x, z) < r\},$$

$$V = f^{-1}((r, \infty)) = \{z \in Z \mid d(x, z) > r\}.$$

Obviously U and V are disjoint; furthermore, since $r \notin f(Z)$, we see that $U \cup V = Z$. Finally, $x \in U$ and $y \in V$, so x and y do not lie in the same connected component of Z . Since x and y were arbitrary, we are done. \square

b. Hausdorff dimension of Cantor-like constructions. We have spent some time discussing Cantor-like sets C obtained through an iterative construction as in (10), but have not yet determined their Hausdorff dimension. The construction is a fairly general one—we may use any number p of basic intervals at each step of the iteration, we may choose any positive numbers $\lambda_1, \dots, \lambda_p$ whose sum is less than 1 as the ratio coefficients, and we may place the basic intervals $I_{i_1 \dots i_{n+1}}$ anywhere we like within the basic interval $I_{i_1 \dots i_n}$ from the previous step, provided they are disjoint. How do these choices affect the Hausdorff dimension $\dim_H C$? Does it matter where we put the intervals? What is the dependence on λ_i ?

These questions were first asked by Abram Besicovitch, one of the founders of the dimension theory of fractals, to his students in a seminar he organised at Cambridge upon his arrival there from Russia in 1927. After various arguments had been put forth to the effect that the Hausdorff dimensions ought to depend on the spacing between the intervals, one of those students, Patrick Moran, proved that the spacing actually has no effect on the Hausdorff dimension, by establishing the following result:

THEOREM 21. *If C is any Cantor-like set given by (10), with ratio coefficients $\lambda_1, \dots, \lambda_p$, then its Hausdorff dimension $\dim_H C$ is the unique value of t which satisfies the equation*

$$(23) \quad \lambda_1^t + \lambda_2^t + \dots + \lambda_p^t = 1.$$

PROOF. First we must verify that (23) does in fact have a unique solution. The function defined by the left hand side is continuous, takes the value p at $t = 0$, and is equal to $\sum_i \lambda_i < 1$ at $t = 1$, hence by the Intermediate Value Theorem, there exists some $t \in (0, 1)$ such that the function is equal to 1. Furthermore, a simple computation of the derivative shows that the function is strictly decreasing, and so the solution t is unique.

From now on, then, t shall denote the unique value for which (23) holds. As in the proof of Proposition 12, there are two parts to the arguments; first we show that $\dim_H C \leq t$, then that $\dim_H C \geq t$.

We begin by obtaining the upper bound, which uses Lemma 10. So for every $\varepsilon > 0$, we must find a “good” ε -cover, for which the quantity $\sum_i (\text{diam } U_i)^t$ does not depend on ε .

We claim that the cover by basic intervals at an appropriate step n of the iteration is the desired one. Writing $\lambda_{\max} = \max\{\lambda_1, \dots, \lambda_p\}$, we see that

$$|I_{i_1 \dots i_n}| = |I_{i_1}| \prod_{j=2}^n \lambda_{i_j} \leq |I_{i_1}| \lambda_{\max}^{n-1} \leq \lambda_{\max}^{n-1},$$

and so if we fix n such that $\lambda_{\max}^{n-1} < \varepsilon$, we may consider the ε -cover

$$\mathcal{U} = \{I_{i_1 \dots i_n} \mid 1 \leq i_j \leq p \forall 1 \leq j \leq n\}.$$

It follows that

$$\begin{aligned} m(C, t, \varepsilon) &\leq \sum_{(i_1, \dots, i_n)} |I_{i_1 \dots i_n}|^t \\ &= \sum_{(i_1, \dots, i_{n-1})} |I_{i_1 \dots i_{n-1} 1}|^t + \dots + |I_{i_1 \dots i_{n-1} p}|^t \\ &= \sum_{(i_1, \dots, i_{n-1})} |I_{i_1 \dots i_{n-1}}|^t (\lambda_1^t + \dots + \lambda_p^t) \\ &= \sum_{(i_1, \dots, i_{n-1})} |I_{i_1 \dots i_{n-1}}|^t \\ &= \dots \\ &= \sum_{i_1} |I_{i_1}|^t, \end{aligned}$$

but this last quantity is a constant, independent of ε , and so Lemma 10 applies: $m(C, t) \leq \sum_i |I_i|^t < \infty$, therefore $\dim_H C \leq t$.

As usual, the proof that $\dim_H C \geq t$ is harder, and we take this up in the next lecture. \square

Lecture 10

a. Completion of the proof of Theorem 21. In order to show that $\dim_H C \geq t$, we want to apply Lemma 11 by showing that for a sufficiently small $\varepsilon > 0$, every ε -cover has

$$(24) \quad \sum_i (\text{diam } U_i)^t \geq K > 0,$$

where K is a constant chosen independently of the cover. The plan of attack is to show that it suffices to establish (24) for any ε -cover *by basic intervals*, and then to apply the following lemma, which gives the necessary bound in that particular case.

LEMMA 22. *If every element of the open cover \mathcal{U} is a basic interval, then*

$$(25) \quad \sum_i (\text{diam } U_i)^t \geq \sum_{j=1}^p |I_j|^t,$$

and the quantity on the right is independent of the cover.

PROOF. Because C is compact, \mathcal{U} has a finite subcover; restricting the sum in (25) to the elements of this subcover will not increase the sum, and so without loss of generality we may assume that \mathcal{U} is finite. Indeed, we take \mathcal{U} to be minimal in the sense that no proper subcollection of the open sets U_i forms a cover of C .

Given a basic interval $I_{i_1 \dots i_n}$, we refer to n as the *depth*. Because \mathcal{U} is finite, there exists some n such that the depth of each basic interval in \mathcal{U} is at most n .

Let $I_{i_1 \dots i_n}$ be a basic interval of maximal depth in \mathcal{U} . Since \mathcal{U} is minimal, it does not contain the basic interval $I_{i_1 \dots i_{n-1}}$ (otherwise we could eliminate $I_{i_1 \dots i_n}$ and obtain a proper subcover). It follows that each of the basic intervals $I_{i_1 \dots i_{n-1}j}$ for $j = 1, \dots, p$ is contained in \mathcal{U} .

Thus the sum in (25) contains the partial sum

$$|I_{i_1 \dots i_{n-1}1}|^t + \dots + |I_{i_1 \dots i_{n-1}p}|^t.$$

By the formula (11) for the lengths of the basic intervals and the definition of t , this is equal to

$$\begin{aligned} (|I_{i_1 \dots i_{n-1}}| \lambda_1)^t + \dots + (|I_{i_1 \dots i_{n-1}}| \lambda_p)^t &= |I_{i_1 \dots i_{n-1}}|^t (\lambda_1^t + \dots + \lambda_p^t) \\ &= |I_{i_1 \dots i_{n-1}}|^t, \end{aligned}$$

and it follows that the sum in (25) is not changed if we replace all the basic intervals of depth n by the corresponding intervals of depth $n-1$. The result follows by induction. \square

We now show that the case of an arbitrary ε -cover can be reduced to the case of a cover by basic intervals, and find the lower bound K in terms of $\sum_j |I_j|^t$.

To this end, for each $r > 0$, we consider the collection $\mathcal{V}(r)$ of basic intervals $I_{i_1 \dots i_n}$ whose lengths satisfy

$$(26) \quad \lambda_{\min} r \leq |I_{i_1 \dots i_n}| \leq \frac{r}{\lambda_{\min}},$$

where $\lambda_{\min} = \min_i \lambda_i < 1$. It follows that $\mathcal{V}(r)$ is a cover of C with uniformly bounded multiplicity; that is, each point $x \in C$ is covered by at most $M = 2 \log \lambda_{\min} / \log \lambda_{\max}$ elements of \mathcal{V} .

Thus $\mathcal{V}(r)$ has the following crucial property: if $U \subset \mathbb{R}$ has $\text{diam } U = r$, then U intersects at most $M' = M / \lambda_{\min}$ elements of $\mathcal{V}(r)$. The exact value of the constant M' is unimportant; what matters is that it is independent of r .

Now let \mathcal{U} be any ε -cover of C ; for each U_i , write $r_i = \text{diam } U_i$, and let $U_{i,1}, \dots, U_{i,m(i)}$ be the basic intervals in $\mathcal{V}(r_i)$ which intersect U_i . It follows from the above remarks that $m(i) \leq M'$; furthermore, we see from (26) that

$$\text{diam } U_{i,j} \leq \frac{\text{diam } U_i}{\lambda_{\min}},$$

whence we obtain the bound

$$\sum_{j=1}^{m(i)} (\text{diam } U_{i,j})^t \leq \frac{M'}{\lambda_{\min}^t} (\text{diam } U_i)^t.$$

Summing over all the elements of \mathcal{U} yields

$$\sum_i (\text{diam } U_i)^t \geq \left(\frac{\lambda_{\min}^t}{M'} \right) \sum_i \sum_{j=1}^{m(i)} (\text{diam } U_{i,j})^t,$$

and since $\{U_{i,j}\}$ is a cover of C by basic intervals, we may apply Lemma 22 to obtain

$$\sum_i (\text{diam } U_i)^t \geq \left(\frac{\lambda_{\min}^t}{M'} \right) \sum_{j=1}^p |I_p|^t = K > 0,$$

which completes the proof. \square

The preceding proof is easily the most intricate argument we have come across so far; the reader is encouraged to test his or her understanding of the method by considering a Cantor-like construction in the plane, using discs instead of intervals (see Figure 14). If the ratio coefficients are $\lambda_1, \dots, \lambda_p$, then the set resulting from the iterative construction again has Hausdorff dimension given as the solution t of (23). The proof is essentially a verbatim copy of the one we have just given, but the reader will be well served by following it through step-by-step and writing down a formal argument.

Lecture 11

a. A competing idea of dimension. In the definition of the Hausdorff function $m(Z, \alpha)$, we considered covers whose sets have diameter less than or equal to ε ; within a particular cover, we might find sets on many different scales, some of which could have diameter much smaller than ε .

An alternate approach is to restrict our attention to covers by sets on the same (small) scale; consider the function

$$(27) \quad r(Z, \alpha, \varepsilon) = \inf \left\{ \sum_i (\text{diam } U_i)^\alpha \mid U_i \text{ open, } \bigcup_i U_i \supset Z, \text{diam } U_i = \varepsilon \right\}.$$

This differs from the definition (13) of $m(Z, \alpha, \varepsilon)$ in only a single symbol; \leq is replaced by $=$. One effect of this change is immediate; the argument that $m(Z, \alpha, \varepsilon)$ depends monotonically on ε does not apply to $r(Z, \alpha, \varepsilon)$, since the collections of admissible covers for two different values of ε are disjoint!

As a result of this change, we have no *a priori* guarantee that the limit of $r(Z, \alpha, \varepsilon)$ as $\varepsilon \rightarrow 0$ exists; indeed, there are many examples for which it does not. To deal with this difficulty, we need the concept of *upper and lower limits*.

DEFINITION 23. Given a sequence $(x_n) \subset \mathbb{R}$, a point $x \in \mathbb{R}$ is an *accumulation point* of (x_n) if there exists a subsequence (x_{n_k}) which converges to x . The *lower limit* of (x_n) is

$$\liminf_{n \rightarrow \infty} x_n = \inf \{ x \mid x \text{ is an accumulation point of } (x_n) \},$$

and the *upper limit* is

$$\limsup_{n \rightarrow \infty} x_n = \sup \{ x \mid x \text{ is an accumulation point of } (x_n) \}.$$

The lower and upper limits are sometimes denoted by \liminf and \limsup , respectively, and one may hear the terms *infimum (supremum) limit*, or possibly *limit inferior (superior)*.

EXAMPLE 24. Define two sequences by $x_n = 1/n$ and $y_n = 1 - 1/n$, and interweave them:

$$\begin{aligned} z_{2n-1} &= x_n, \\ z_{2n} &= y_n. \end{aligned}$$

Then the sequence (z_n) does not converge, but it has subsequences which do. The set of accumulation points of (z_n) is $\{0, 1\}$, and we see that

$$\begin{aligned} \liminf_{n \rightarrow \infty} z_n &= 0, \\ \limsup_{n \rightarrow \infty} z_n &= 1. \end{aligned}$$

EXERCISE 6. Show that the lower and upper limits may equivalently be defined by

$$\begin{aligned}\underline{\lim}_{n \rightarrow \infty} x_n &= \lim_{n \rightarrow \infty} \left(\inf_{m \geq n} x_m \right), \\ \overline{\lim}_{n \rightarrow \infty} x_n &= \lim_{n \rightarrow \infty} \left(\sup_{m \geq n} x_m \right).\end{aligned}$$

Furthermore, show that $x = \underline{\lim} x_n$ if and only if both of the following hold:

- (1) For every $\varepsilon > 0$, there exists N such that $x_n \geq x - \varepsilon \forall n \geq N$.
- (2) There exists a subsequence (x_{n_k}) of (x_n) which converges to x .

Similarly, show that $x = \overline{\lim} x_n$ if and only if:

- (1) For every $\varepsilon > 0$, there exists N such that $x_n \leq x + \varepsilon \forall n \geq N$.
- (2) There exists a subsequence (x_{n_k}) of (x_n) which converges to x .

An immediate consequence of the definition is that $\underline{\lim} x_n = \overline{\lim} x_n$ if and only if $\lim x_n$ itself exists, in which case it is equal to the common value, and is the only accumulation point.

We have given the definition of the lower and upper limits for a discrete index (n), but it goes through equally well in the case of a continuous index (such as ε). Thus returning to the function $r(Z, \alpha, \varepsilon)$, we define

$$\begin{aligned}\underline{r}(Z, \alpha) &= \underline{\lim}_{\varepsilon \rightarrow 0} r(Z, \alpha, \varepsilon), \\ \overline{r}(Z, \alpha) &= \overline{\lim}_{\varepsilon \rightarrow 0} r(Z, \alpha, \varepsilon).\end{aligned}$$

The following partial analogue of Proposition 4 is immediate from the definitions:

PROPOSITION 25. *The set functions $\underline{r}(\cdot, \alpha)$ and $\overline{r}(\cdot, \alpha)$ satisfy the following properties:*

- (1) *Normalisation:* $\underline{r}(\emptyset, \alpha) = \overline{r}(\emptyset, \alpha) = 0$ for all $\alpha > 0$.
- (2) *Monotonicity:* $\underline{r}(Z_1, \alpha) \leq \underline{r}(Z_2, \alpha)$ and $\overline{r}(Z_1, \alpha) \leq \overline{r}(Z_2, \alpha)$ whenever $Z_1 \subset Z_2$.

Conspicuously absent from Proposition 25 is the subadditivity property which held for $m(\cdot, \alpha)$. The proof of that property relied on the construction of an ε -cover as a union of covers of arbitrarily small diameter; because the definition of $r(Z, \alpha, \varepsilon)$ does not allow us to use sets of diameter less than ε , the proof does not go through here. While we will see in the next lecture that $\overline{r}(Z, \alpha, \varepsilon)$ is subadditive for *finite* collections $\{Z_i\}$, no such result holds for countable collections, or for $\underline{r}(Z, \alpha, \varepsilon)$. The consequences of this will become apparent shortly.

As functions of α , $\underline{r}(Z, \alpha)$ and $\overline{r}(Z, \alpha)$ have similar properties to $m(Z, \alpha)$; there are critical values $\underline{\alpha}_C, \overline{\alpha}_C$ below which the value of the function is ∞ , and above which it is 0. The critical value of $m(Z, \alpha)$ determines the Hausdorff dimension of Z , and the critical values of \underline{r} and \overline{r} are also dimensional

quantities, known as the *lower box dimension* and *upper box dimension*, respectively,¹³ and denoted $\underline{\dim}_B Z$ and $\overline{\dim}_B Z$. As with α_C , we have

$$\begin{aligned}\underline{\dim}_B Z &= \underline{\alpha}_C = \inf\{\alpha > 0 \mid \underline{r}(Z, \alpha) = 0\} \\ &= \sup\{\alpha > 0 \mid \underline{r}(Z, \alpha) = \infty\}, \\ \overline{\dim}_B Z &= \overline{\alpha}_C = \inf\{\alpha > 0 \mid \overline{r}(Z, \alpha) = 0\} \\ &= \sup\{\alpha > 0 \mid \overline{r}(Z, \alpha) = \infty\}.\end{aligned}$$

b. Basic properties and relationships. As an immediate consequence of Proposition 25, we have the following analogue of Proposition 8:

PROPOSITION 26. *The upper and lower box dimensions have the following basic properties:*

- (1) $\underline{\dim}_B \emptyset = \overline{\dim}_B \emptyset = 0$.
- (2) $\underline{\dim}_B Z_1 \leq \underline{\dim}_B Z_2$ and $\overline{\dim}_B Z_1 \leq \overline{\dim}_B Z_2$ whenever $Z_1 \subset Z_2$.
- (3) $\underline{\dim}_B(\bigcup_i Z_i) \geq \sup_i \underline{\dim}_B Z_i$ and $\overline{\dim}_B(\bigcup_i Z_i) \geq \sup_i \overline{\dim}_B Z_i$, where $\{Z_i\}$ is any countable collection of subsets of \mathbb{R}^n .

Property (3) is weaker than its analogue in Proposition 8 because of the failure of countable subadditivity for the lower and upper box dimensions; we will see an example where the inequality is strict.

First, we examine the relationship between the Hausdorff dimension and the two box dimensions. It follows immediately from the definitions that

$$m(Z, \alpha) \leq \underline{r}(Z, \alpha) \leq \overline{r}(Z, \alpha)$$

for any $Z \subset \mathbb{R}^n$ and $\alpha > 0$, and thus we have the relations

$$(28) \quad \dim_H Z \leq \underline{\dim}_B Z \leq \overline{\dim}_B Z.$$

One of our primary goals will be to establish conditions on Z under which all three dimensional quantities in (28) coincide. When this occurs, we may refer to the common value as the *fractal dimension* without fear of ambiguity, and we will see that the three quantities agree for a wide range of examples, including some rather complicated sets. However, we will also see relatively simple examples of sets Z for which the inequalities in (28) become strict; the challenge is to develop some criteria which may let us know what sort of behaviour to expect for a particular Z .

First, though, let us address the definition (27) of $r(Z, \alpha, \varepsilon)$, which played a key role in the description of the lower and upper box dimensions. Since we restrict our attention to covers in which *every* set has diameter ε , every term in the sum $\sum_i (\text{diam } U_i)^\alpha$ is the same! It seems rather silly, then, to continue writing it as a sum, and indeed (27) is equivalent to

$$r(Z, \alpha, \varepsilon) = \left\{ \varepsilon^\alpha N(\mathcal{U}) \mid U_i \text{ open, } \bigcup_i U_i \supset Z, \text{diam } U_i = \varepsilon \right\},$$

¹³In the literature, one may also find the box dimensions referred to as *box counting dimensions*, *entropy dimensions*, or *capacities*.

where $N(\mathcal{U})$ denotes the number of elements in the cover \mathcal{U} . Writing $N(\varepsilon)$ (or sometimes $N(Z, \varepsilon)$) for the smallest value of $N(\mathcal{U})$, where \mathcal{U} is any open cover by sets of diameter ε , we see that

$$(29) \quad r(Z, \alpha, \varepsilon) = \varepsilon^\alpha N(\varepsilon).$$

EXERCISE 7. Show that the lower and upper box dimensions may be characterised by

$$(30) \quad \underline{\dim}_B Z = \underline{\alpha}_C = \liminf_{\varepsilon \rightarrow 0} \frac{\log N(\varepsilon)}{\log 1/\varepsilon},$$

$$(31) \quad \overline{\dim}_B Z = \overline{\alpha}_C = \limsup_{\varepsilon \rightarrow 0} \frac{\log N(\varepsilon)}{\log 1/\varepsilon}.$$

Proposition 18 showed that defining $m(Z, \alpha, \varepsilon)$ in terms of open *balls* rather than open *sets* does not change the critical value α_C , and hence leads to an equivalent definition of the Hausdorff dimension. The same is true for the lower and upper box dimensions, and the proof goes through verbatim; hence we may also use $N_B(\varepsilon)$, the smallest cardinality of a cover of Z by balls of diameter ε , in (30) and (31).

Because $N(\varepsilon)$ (and $N_B(\varepsilon)$) must be finite in order for the definition of lower and upper box dimension to make any sense, we restrict our attention to *bounded* subsets of \mathbb{R}^n . Furthermore, we may in fact consider only *compact* subsets Z , thanks to the following fact.

PROPOSITION 27. *The box dimension (lower or upper) of a set Z is the same as the box dimension (lower or upper) of its closure.*

PROOF. Recall that the closure \bar{Z} of a set $Z \subset \mathbb{R}^n$ is just the union of Z and its accumulation points. If $\{B(x_i, \varepsilon/2)\}$ covers Z , then $\{B(x_i, \varepsilon)\}$ covers \bar{Z} , and so

$$N(\bar{Z}, 2\varepsilon) \leq N(Z, \varepsilon) \leq N(\bar{Z}, \varepsilon),$$

from which (30) and (31) show that $\underline{\dim}_B Z = \underline{\dim}_B \bar{Z}$, and also $\overline{\dim}_B Z = \overline{\dim}_B \bar{Z}$. \square

Thus it suffices to consider subsets of \mathbb{R}^n which are both closed and bounded, hence compact.

In particular, Proposition 27 lets us give our first example of a set for which the competing ideas of dimension do not agree, and the inequality in (28) becomes strict.

EXAMPLE 28. Let $Z = \mathbb{Q} \cap [0, 1]$ be the set of all rational numbers in the unit interval. Then Z is countable, which implies that $\dim_H Z = 0$ by Corollary 9. However, Z is dense in $[0, 1]$, and so

$$\begin{aligned} \underline{\dim}_B Z &= \underline{\dim}_B [0, 1] = 1, \\ \overline{\dim}_B Z &= \overline{\dim}_B [0, 1] = 1. \end{aligned}$$

There are many other examples of sets for which the Hausdorff and box dimensions do not agree; indeed, in the next lecture we will see how to construct, for any $0 < \alpha \leq \beta < 1$, a countable closed set $Z \subset [0, 1]$ such that $\dim_H Z = 0$, $\underline{\dim}_B Z = \alpha$, and $\overline{\dim}_B Z = \beta$. In fact, Z will have a particularly simple form; it will be a sequence converging to 0.

Despite the fact that such simple sets can display this rather unpleasant behaviour, there is a large class of examples for which $\dim_H Z = \underline{\dim}_B Z = \overline{\dim}_B Z$; in particular, the three quantities coincide for the “standard” geometric objects, such as line segments, discs, cubes, etc.

For general subsets $Z \subset \mathbb{R}^n$, even though the inequalities in (28) may become strict, it is still possible to prove the following (rather deep) theorem, due to Lev Pontryagin and Lev Shnirel’man in 1932, which is a stronger version of Hausdorff’s theorem (Theorem 17). As with the Hausdorff dimension, the lower and upper box dimensions depend on the choice of metric; we will see in the next lecture that passing to a strongly equivalent metric preserves these dimensions, but they may change if the new metric is merely equivalent.

THEOREM 29 (Pontryagin-Shnirel’man). *Given a set $Z \subset \mathbb{R}^n$, the topological dimension $\dim Z$ and the lower box dimensions $\underline{\dim}_{B,\rho} Z$ are related by the following variational principle:*

$$(32) \quad \dim Z = \inf_{\rho \sim d} \underline{\dim}_{B,\rho} Z.$$

That is, the topological dimension is the infimum of the possible lower box dimensions, taken over all metrics ρ which are equivalent to the standard metric d .

No analogue of this result holds for the upper box dimension, and so this theorem is in some sense the best result possible. Pontryagin and Shnirel’man were the first to introduce the concept of lower box dimension, which they called the *metric order*, a piece of terminology which has fallen by the wayside. Perhaps because no such result holds for the upper box dimension, they did not consider that quantity.

Lecture 12

a. Further properties of box dimension. Given a continuous map $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$, how are the box dimensions of Z related to the box dimensions of $f(Z)$? If f is Lipschitz, then we have the following analogue of Proposition 19.

PROPOSITION 30. *If $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz, then $\underline{\dim}_B f(Z) \leq \underline{\dim}_B Z$ and $\overline{\dim}_B f(Z) \leq \overline{\dim}_B Z$ for every $Z \subset \mathbb{R}^n$.*

PROOF. Given $Z \subset \mathbb{R}^n$, let $N(Z, \varepsilon)$ denote the cardinality of a minimal cover of Z by open sets of diameter ε . Now if $\mathcal{U} = \{U_i\}$ is any such cover of Z , then $\text{diam } f(U_i) \leq L \text{diam } U_i$, where L is a Lipschitz constant for f , and so $f(\mathcal{U}) = \{f(U_i)\}$ can be “fattened up” into an open cover of $f(Z)$ by sets of diameter precisely $L\varepsilon$; it follows that $N(f(Z), L\varepsilon) \leq N(Z, \varepsilon)$, and so

$$\frac{\log N(f(Z), L\varepsilon)}{\log L + \log 1/L\varepsilon} \leq \frac{\log N(Z, \varepsilon)}{\log 1/\varepsilon}.$$

Taking lower and upper limits as $\varepsilon \rightarrow 0$ gives the result. \square

It follows that if f is bi-Lipschitz, then Z and $f(Z)$ have the same box dimensions. In particular, the lower and upper box dimensions do not change if we pass to a strongly equivalent metric.

We have seen many similarities between the Hausdorff dimension and the box dimension, and have also seen one example (the rational numbers) for which they differ. Before giving a further class of such examples, we show that for a very important family of sets, the Cantor-like constructions we have seen so far, all three quantities coincide.

THEOREM 31. *Let C be the limit of a Cantor-like construction (10), where the diameters of the basic intervals $I_{i_1 \dots i_n}$ are given by (11). Then $\dim_H C = \underline{\dim}_B C = \overline{\dim}_B C$.*

PROOF. From the inequalities in (28), it will suffice to show that $\overline{\dim}_B C \leq \dim_H C$. By Moran’s theorem (Theorem 21), $\dim_H C$ is the unique solution of (23), and so we show that $\overline{\dim}_B C \leq t$, where $\sum_i \lambda_i^t = 1$.

Given $\varepsilon > 0$, we want to estimate $N(Z, \varepsilon)$ from above, and so we produce a cover of C by open sets of diameter ε as follows. Fix $x \in C$, and let $n(x)$ be the unique integer such that

$$|I_{i_1 \dots i_{n(x)-1}| \geq \varepsilon > |I_{i_1 \dots i_{n(x)}}|.$$

Thus for each x we may choose an open interval $U(x)$ of length ε which contains $I_{i_1 \dots i_{n(x)}}$. Because C is compact, it can be covered by a finite collection $\{I_{i_1 \dots i_{n(x_i)}}\}_{i=1}^N$. Without loss of generality, we may take this collection to be

- (1) (b_n) tends to infinity monotonically as $n \rightarrow \infty$.
- (2) $\sum_{n=1}^{\infty} a_n b_n < 1$.
- (3) The exponential growth rate of the partial sums $S_n = \sum_{k=1}^n b_k$ is given by

$$\underline{\lim}_{n \rightarrow \infty} \frac{\log S_n}{\log 1/a_n} = \alpha, \quad \overline{\lim}_{n \rightarrow \infty} \frac{\log S_n}{\log 1/a_n} = \beta.$$

- (4) The “tail” $[T_n, T]$ of the set A is not too long: there exists a constant C such that

$$\frac{T - T_n}{a_n S_n} \leq C$$

for all n .

The summability property (2) guarantees that A is bounded. The significance of the partial sums S_n is that they let us estimate $N(\varepsilon)$, the cardinality of an optimal cover by open sets of diameter ε . Indeed, for $\varepsilon = a^n$, we see that any such cover must contain at least S_n sets, since each set can contain at most one of the points $T_k + ma_k$ for $k \leq n$, $m \leq b_k$, and S_n is the number of such points. Furthermore, we may cover the interval $[T_n, T]$ with $(T - T_n)/\varepsilon$ intervals of length ε , and by property (4), this shows that $N(\varepsilon) \leq (C + 1)S_n$. The result for the lower and upper box dimensions will then follow immediately from property (3).

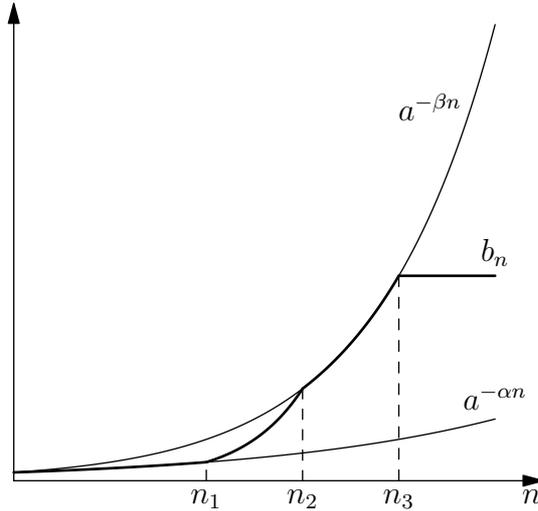


FIGURE 21. Building a sequence (b_n) .

It remains only to produce a sequence with these properties. To this end, we follow a four-step recursive procedure, illustrated in Figure 21. At first, b_n is just the integer part of $a^{-\alpha n}$, and so S_n also grows at the same rate as $a^{-\alpha n}$. It follows that the quantity $(\log S_n)/(\log 1/a_n)$ converges to α

as n grows, and so we may choose n_1 such that

$$(35) \quad \left| \frac{\log S_{n_1}}{\log 1/a_{n_1}} - \alpha \right| < \frac{1}{2}.$$

Now we would like to let b_n follow the function $a^{-\beta n}$ for a while, to approximate the desired upper limit, but if we jump directly to the graph of $a^{-\beta n}$ at n_1 , we will find that the sequence we eventually produce fails property (4), and so we must be slightly more careful. Thus for $n > n_1$, we let b_n grow exponentially until it reaches the upper function at n_2 , at which point we have b_n follow $a^{-\beta n}$ until

$$(36) \quad \left| \frac{\log S_{n_3}}{\log 1/a_{n_3}} - \beta \right| < \frac{1}{2}.$$

Finally, for $n > n_3$, we leave b_n constant until it is once again equal to $a^{-\alpha n}$ at n_4 (which is somewhere off the right edge of the graph in Figure 21). Then we iterate all four steps of this procedure, replacing the bound $1/2$ in (35) and (36) with $(1/2)^k$ at the k^{th} iteration. One may then verify that the sequence (b_n) so constructed has properties (1)–(4), and thus the set A given in (34) has the dimensions claimed. \square

In some sense, this is the simplest possible counterexample to equality in (28); any simpler set is just finite, and then all three quantities are immediately 0.

Note that Theorem 32 does *not* provide a set A with $\dim_H A = 0$, $0 < \underline{\dim}_B A < 1$, and $\overline{\dim}_B A = 1$; this case, which requires a different construction, is left as an exercise for the reader.

We will use this construction again in the next lecture, when we examine another difference between the lower and upper box dimensions; namely, their behaviour under finite unions.

Lecture 13

a. Subadditivity, or the lack thereof. An important property of the dimensional quantities we have seen so far is their behaviour under taking countable unions. We saw in Proposition 8 that $\dim_H(\bigcup_i Z_i) = \sup_i \dim_H Z_i$, while the best that Proposition 26 could offer for the lower and upper box dimensions was an inequality. Indeed, the example of the rational numbers in the unit interval showed that the box dimensions can increase when we take a countable union; taking Z_i to be singleton sets which exhaust $\mathbb{Q} \cap [0, 1]$, we saw that

$$\begin{aligned} \underline{\dim}_B \left(\bigcup_i Z_i \right) &= 1 > 0 = \sup_i \underline{\dim}_B Z_i, \\ \overline{\dim}_B \left(\bigcup_i Z_i \right) &= 1 > 0 = \sup_i \overline{\dim}_B Z_i. \end{aligned}$$

The reason for the discrepancy between the two sorts of dimensions is the fact that while $m(Z, \alpha)$ is countably subadditive (Property (3) in 4), neither $\underline{r}(Z, \alpha)$ nor $\overline{r}(Z, \alpha)$ has this property. The best we can hope for in this case is *finite* subadditivity:

EXERCISE 8. Show that for any finite collection of sets $\{Z_i\}_{i=1}^k$ in \mathbb{R}^n , we have

$$(37) \quad \overline{r} \left(\bigcup_{i=1}^k Z_i \right) \leq \sum_{i=1}^k \overline{r}(Z_i, \alpha),$$

and that consequently

$$(38) \quad \overline{\dim}_B \bigcup_{i=1}^k Z_i = \max_{i=1}^k \overline{\dim}_B Z_i.$$

However, even this weakened property only holds for the *upper* box dimension; for the set function $\underline{r}(Z, \alpha)$ associated to the lower box dimension, we have no subadditivity property at all! This is the second case we have seen in which the lower and upper box dimensions behave in fundamentally different ways; the first was the Pontryagin-Shnirel'man theorem for the lower box dimension, which had no analogue for the upper box dimension.

The following example demonstrates the failure of even finite subadditivity for $\underline{r}(Z, \alpha)$, by giving two sets Z_1, Z_2 such that $\underline{\dim}_B Z_1 \cup Z_2 > \max\{\underline{\dim}_B Z_1, \underline{\dim}_B Z_2\}$.

EXAMPLE 33. Fixing $0 < \alpha < \beta < 1$, we may use the construction in the proof of Theorem 32 to find $Z_1 \subset [0, 1]$ and $Z_2 \subset [2, 3]$ such that $\dim_H Z_1 = \dim_H Z_2 = 0$, $\underline{\dim}_B Z_1 = \underline{\dim}_B Z_2 = \alpha$, and $\overline{\dim}_B Z_1 = \overline{\dim}_B Z_2 = \beta$. We wish to modify the construction slightly to ensure that $\underline{\dim}_B Z_1 \cup Z_2 > \alpha$.

The key idea in the construction was for b_n to follow first one exponential curve, then the other, transferring at the appropriate indices n_i . In our

present case, we wish to start $b_n^{(1)}$ (the sequence defining Z_1) on the lower exponential curve, and $b_n^{(2)}$ (defining Z_2) on the upper curve; furthermore, we define a *single* sequence n_i of indices for *both* $b_n^{(1)}$ and $b_n^{(2)}$ by requiring that the estimates on $S_{n_i}^{(1)}$ and $S_{n_i}^{(2)}$ are both within the range given by (35) and (36).

Defining Z_1 and Z_2 in this way, we see that if $\varepsilon > 0$ is such that $(\log N(Z_1, \varepsilon))/(\log 1/\varepsilon)$ is near α , then $(\log N(Z_2, \varepsilon))/(\log 1/\varepsilon)$ will be near β , and vice versa. Since $N(Z_1 \cup Z_2, \varepsilon) \geq N(Z_1, \varepsilon) + N(Z_2, \varepsilon)$, it can then be shown that $(\log N(Z_1 \cup Z_2, \varepsilon))/(\log 1/\varepsilon)$ is bounded away from α , and hence $\underline{\dim}_B Z_1 \cup Z_2 > \alpha = \underline{\dim}_B Z_1 = \underline{\dim}_B Z_2$.

b. A little bit of measure theory. As we have seen, computing the Hausdorff dimension of a set can be very difficult, even if the set is geometrically quite regular, as was the case for the Cantor-like sets we have considered so far. For sets whose self-similarity is not quite so regular, the situation becomes even worse; for example, consider the non-linear map shown in Figure 17. This map generates a repelling Cantor set C through the same construction that we carried out for the linear map in Figure 8, and so we may ask what the Hausdorff dimension of C is. C is the limit set of a Cantor-like construction in which the ratio coefficients are no longer constant, but may change at each step of the iteration; this occurs because the map is no longer linear, and so how much contraction occurs at each step now depends on which point we consider, not just on which interval it is in.

How does the Hausdorff dimension respond to this change in the construction? The problems which arise at this stage are much more difficult than those we encountered in proving Theorem 21, and many answers are unknown.

There are various lines of attack that we might pursue at this point. For example, we might write $\lambda_i^{(n)}$ for the ratio coefficients at the n^{th} step of the iteration, and try to figure out what happens if the $\lambda_i^{(n)}$ converge to some fixed coefficients λ_i as $n \rightarrow \infty$. If we can deal with this more general case, which seems more tractable than consider completely arbitrary $\lambda_i^{(n)}$, then we will have successfully expanded the class of examples for which we have some answers.

In order to go further, though, we will need to expand our toolkit. At some point it will become necessary to distinguish between “bad” points of C , which we may want to ignore, and “good” points, which we may be able to deal with, and to show that there are in some sense “more” of the latter. In order to make all this precise, we need the idea of a *measure*. Without further ado, then, we have the following definition.

DEFINITION 34. Let X be any set (which will be called our *space*), and let \mathcal{A} be a collection of subsets of X . \mathcal{A} is called an *algebra* if

- (1) $\emptyset, X \in \mathcal{A}$.

- (2) $A \in \mathcal{A}$ implies $X \setminus A \in \mathcal{A}$.
 (3) $A_1, A_2 \in \mathcal{A}$ implies $A_1 \cup A_2 \in \mathcal{A}$.

Property (3) immediately implies that \mathcal{A} is closed under finite unions: $\bigcup_{i=1}^n A_i \in \mathcal{A}$ whenever $A_1, \dots, A_n \in \mathcal{A}$. If in addition \mathcal{A} is closed under *countable* unions, that is, if

- (4) $A_n \in \mathcal{A} \forall n \in \mathbb{N}$ implies $\bigcup_{n=1}^{\infty} A_n \in \mathcal{A}$,

then \mathcal{A} is a σ -algebra. The elements of \mathcal{A} are referred to as *measurable* sets.

A set function $m: \mathcal{A} \rightarrow [0, \infty]$, which assigns to each measurable set a non-negative number (or possibly ∞), is a *measure* if it satisfies the following properties:

- (1) $m(\emptyset) = 0$.
 (2) Monotonicity: $m(A_1) \leq m(A_2)$ whenever $A_1 \subset A_2$ and both are measurable.
 (3) σ -additivity: $m(\bigcup_i A_i) = \sum_i m(A_i)$ whenever $\{A_i\} \subset \mathcal{A}$ is a countable collection of disjoint measurable sets.

The triple (X, \mathcal{A}, m) is known as a *measure space*.

Before moving on to meatier applications, we give a few very basic examples of measure spaces.

EXAMPLE 35. Let X be any set, and let $\mathcal{A} = \{\emptyset, X\}$. Then \mathcal{A} is a σ -algebra, and we may define a measure m by $m(\emptyset) = 0$, $m(X) = \infty$. This, of course, is a completely trivial example.

EXAMPLE 36. Let X be any set, and let $\mathcal{A} = 2^X$ be the power set of X ; that is, the collection of all subsets of X . \mathcal{A} is again a σ -algebra (in fact, it is the largest possible σ -algebra on X), and we define a measure m by

$$m(A) = \begin{cases} \text{card}(A) & A \text{ finite,} \\ \infty & \text{otherwise.} \end{cases}$$

So if A is a finite set, m counts the number of points in A , otherwise it gives ∞ ; this is known as the *counting measure* on X .

EXAMPLE 37. Let X and \mathcal{A} be as in the previous example, and fix $x \in X$. Define m by

$$m(A) = \begin{cases} 1 & x \in A, \\ 0 & x \notin A. \end{cases}$$

Thus m simply measures whether or not A contains the point x ; this is known as the *point measure* sitting on x .

Lecture 14

a. Lebesgue measure and outer measures. We now introduce a more complicated example of a measure space; Lebesgue measure on \mathbb{R}^n , which we first mentioned in the proof of Theorem 20. In the case $n = 1$, which we will consider first, this is a generalisation of the idea of “length” which applies to a broader class of sets than merely intervals.

The full construction of Lebesgue measure is one of the primary parts of measure theory, and a complete treatment of all the details requires most of a graduate-level course, so our discussion here will necessarily be somewhat abbreviated, and we will omit proofs. (The interested reader is referred to *Measure Theory* by Paul Halmos, or *Real Analysis* by H. L. Royden, for a complete exposition.)

In order to construct a measure space (X, \mathcal{A}, m) , we must do two things. First, we must produce a σ -algebra \mathcal{A} , which we would like to make as large as possible. Second, we must figure out how to define a set function m which satisfies the three properties of a measure. In particular, it is far from clear how to guarantee that whatever function m we construct is σ -additive, especially if the collection \mathcal{A} of sets for which this must be checked is very large.

There is a standard procedure in measure theory which deals with both these challenges; for this we need the notion of *outer measure*.

DEFINITION 38. A set function $m^*: 2^X \rightarrow [0, \infty]$ is an *outer measure* if it satisfies

- (1) $m^*(\emptyset) = 0$.
- (2) $m^*(A_1) \leq m^*(A_2)$ whenever $A_1 \subset A_2$.
- (3) $m^*(\bigcup_i A_i) \leq \sum_i m^*(A_i)$ for any countable collection of subsets $A_i \subset X$.

The first two properties are exactly those which we required for a measure; however, the third property of a measure, σ -additivity, has been replaced here with the weaker property of σ -subadditivity. So in that regard, the notion of outer measure is weaker than the notion of measure; however, we require the outer measure to be defined for *every* subset of X .

In fact, we have already seen an example of an outer measure; the three properties above are exactly what we proved in Proposition 4 for the set function $m(\cdot, \alpha)$. Thus Proposition 4 may be rephrased as the statement that $m(\cdot, \alpha)$ is an outer measure on \mathbb{R}^n ; from now on, we will write it as $m_H(\cdot, \alpha)$ so as to stress the origin of this particular outer measure, and to avoid confusion with our notation m for a measure.

Once we have an outer measure, there is a canonical way to produce both a σ -algebra \mathcal{A} and a measure m . The key step is the following definition.

DEFINITION 39. Given an outer measure $m^*: 2^X \rightarrow [0, \infty]$, we say that $E \subset X$ is *measurable* if for every $A \subset X$, we have

$$(39) \quad m^*(A) = m^*(A \cap E) + m^*(A \cap E^c).$$

As we need to gain σ -additivity in order to have a measure, this is a reasonable definition to make; after all, (39) is just a very particular case of finite additivity. Indeed, it turns out to be enough.

THEOREM 40. *Let \mathcal{A} be the collection of measurable subsets of X for an outer measure m^* , and let m be the restriction of m^* to \mathcal{A} ; that is, $m(A) = m^*(A)$ for $A \in \mathcal{A}$. Then*

- (1) \mathcal{A} is a σ -algebra.
- (2) m is a measure.

The procedure of passing from an outer measure to a measure and a σ -algebra of measurable sets is a completely general one, which works for any (X, μ^*) .

Restricting our attention now to the case $X = \mathbb{R}$, we follow the above recipe by first defining an outer measure on \mathbb{R} as follows:

$$(40) \quad m^*(A) = \inf \left\{ \sum_k \ell(I_k) \mid \bigcup_k I_k \supset A, I_k \in \mathcal{C} \forall k \right\},$$

where \mathcal{C} is the class of all open intervals, and $\ell: \mathcal{C} \rightarrow [0, \infty]$ is the set function $\ell(I_k) = \text{diam}(I_k)$. This is a particular example of a *Carathéodory construction*, an important procedure for building an outer measure, and hence a measure, using a set function on some family of subsets of X .

Observe that (40) is very reminiscent of our definition of $m_H(A, 1)$; the only difference between the two is that the latter requires that we consider covers whose diameter becomes arbitrarily small, whereas m^* is not concerned with the size of the elements of the cover. In fact, this makes no difference to the actual value of the set function.

PROPOSITION 41. $m^*(A) = m_H(A, 1)$ for every $A \subset \mathbb{R}$.

PROOF. It follows from the definitions that $m^*(A) \leq m_H(A, 1, \varepsilon)$ for every $\varepsilon > 0$, since the infimum in the latter is taken over a smaller collection of covers, and hence $m^*(A) \leq m_H(A, 1)$.

To prove the reverse inequality, fix $\gamma > 0$, and observe that there exists $\varepsilon > 0$ such that

$$m_H(A, 1) \leq m_H(A, 1, \varepsilon) + \gamma,$$

and that there exists a cover $\{I_k\}$ of A by open intervals such that

$$m^*(A) \geq \sum_k |I_k| - \gamma.$$

If the interval I_k has length greater than ε , we can cover it with intervals $J_{k,i}$ of length less than ε in such a way that

$$\sum_i |J_{k,i}| \leq |I_k| + \frac{\gamma}{2^k}.$$

It follows that

$$m^*(A) \geq \sum_{k,i} |J_{k,i}| - 2\gamma \geq m_H(A, 1, \varepsilon) - 2\gamma \geq m_H(A, 1) - 3\gamma,$$

and since $\gamma > 0$ was arbitrary, the result follows. \square

COROLLARY 42. *m^* is an outer measure.*

It follows that m^* defines a σ -algebra $\mathcal{A} \subset 2^{\mathbb{R}}$ and a measure $m: \mathcal{A} \rightarrow [0, \infty]$. One may check that every interval I (whether open, closed, or neither) is measurable, and that $m(I) = |I|$, so that m really is an extension of length.

Indeed, most of the sets we usually encounter are measurable; \mathcal{A} contains all open sets, all closed sets, and all countable sets, along with a great deal more. In particular, the Cantor set C is closed, and hence measurable; one may show that $m(C) = 0$, and this makes precise our statement in Lecture 3 that “the length of the Cantor set is zero”.

Lebesgue measure can be defined not just on the real line, but on any \mathbb{R}^n ; simply replace the open intervals in the above procedure by open balls. One may once again prove that $m^*(A) = m_H(A, n)$, although the argument is slightly more difficult; it then follows that m^* is an outer measure, and it defines a measure m which agrees with the usual idea of n -dimensional volume for familiar geometric shapes such as balls and cubes.

b. Hausdorff measures. The fact that $m_H(\cdot, \alpha)$ is an outer measure is not restricted to integer values of α , but holds true for any $\alpha > 0$. We call this the *Hausdorff outer measure*, and the measure it induces is called *Hausdorff measure*. In this manner, we obtain quite a large collection of measures sitting on \mathbb{R}^n ; in fact, a one-parameter family of them, indexed by α .

Consider the middle-thirds Cantor set $C \subset [0, 1]$. What is its Hausdorff measure? We know from the definition of Hausdorff dimension that

$$(41) \quad m_H(C, \alpha) = \begin{cases} \infty & \alpha < \dim_H C, \\ 0 & \alpha > \dim_H C; \end{cases}$$

the result $m(C) = m_H(C, 1) = 0$ is just a particular case of this more general fact. It follows from Moran’s equation (23) that the critical value α_C is

$$\alpha_C = \frac{\log 2}{\log 3}.$$

EXERCISE 9. Prove that $m_H(C, \alpha_C) = 1$.

The result of Exercise 9 is valid for any of the Cantor-like sets we have considered; the key fact is the relationship (23) between the ratio coefficients, which lets us move between different levels of the construction without changing the potential of a cover by basic intervals. Like Goldilocks, we find that other values of α are either too big or too small for $m_H(C, \alpha)$ to measure C properly (41), but that α_C is “just right”. For this particular

choice of α , we get a Hausdorff measure which sits *on the Cantor set*, and which will be of great utility to us later on.

In the meantime, one final property of the Hausdorff measures is worth noting. Both the Hausdorff measures and the outer measures they are induced from are *translation invariant*; that is, $m_H(A, \alpha) = m_H(A + x, \alpha)$ for every $A \subset \mathbb{R}^n$, $x \in \mathbb{R}^n$, and $\alpha > 0$, where $A + x$ is the image of A under a translation by the vector x . In particular, this is true for Lebesgue measure. Note that for example, the point measure m_x defined in the previous lecture is not translation invariant.

In the next lecture, we will describe another way of building measures on the Cantor set, by utilising the symbolic space Σ_p^+ , building measures there, and then transferring them to the Cantor set via the coding map h .

Lecture 15

a. Choosing an outer measure. The Carathéodory construction of an outer measure described in (40) is actually quite a common one. In the construction of the Lebesgue measure, we took \mathcal{C} to be the class of open intervals, and ℓ to be the length function, $\ell((a, b)) = b - a$. What would happen if we chose a different set function ℓ ? We can still define m^* , determine which sets are measurable, and obtain a set function m ; but will it be a measure?

We examine the possible outcomes by considering three candidate set functions:

$$\begin{aligned}\ell_1((a, b)) &= e^{b-a}, \\ \ell_2((a, b)) &= (b - a)^2, \\ \ell_3((a, b)) &= \sqrt{b - a}.\end{aligned}$$

The proof that m^* as defined in (40) is monotonic and σ -subadditive does not rely on the particular form of ℓ , and so these properties hold whatever set function we begin with. However, we see that in order to have $m^*(\emptyset) = 0$, we must be able to find sets $U \in \mathcal{C}$ for which $\ell(U)$ is arbitrarily small. Hence if we take ℓ_1 as our set function, m^* will not be an outer measure.

Thus we must demand that $\ell(U)$ take arbitrarily small values, and both ℓ_2 and ℓ_3 satisfy this requirement. It follows that these set functions will define outer measures m_2^* and m_3^* , respectively.

EXERCISE 10. Show that $m_2^*(A) = 0$ for every $A \subset \mathbb{R}$; in particular, show that every subset of X is measurable, but $m_2^*(A) \neq \ell(A)$ for $A \in \mathcal{C}$.

Given that the whole point of introducing an outer measure using ℓ was to extend the definition of ℓ beyond the elements of \mathcal{C} , the result of Exercise 10 is rather undesirable; m_2^* does not agree with ℓ on the class of intervals!

To avoid this behaviour, we consider only set functions ℓ such that the outer measure they induce agrees with ℓ on the members of \mathcal{C} . It is not difficult to show that ℓ_3 has this property—however, there is now another problem. In order for the result of the Carathéodory construction to be a genuine extension of the initial set function, the σ -algebra of measurable sets should contain the collection of intervals, but this is not the case for ℓ_3 . Indeed, given $a < c < b$, one immediately sees that

$$m_3^*((a, b)) = \sqrt{b - a} \neq \sqrt{c - a} + \sqrt{b - c} = m_3^*((a, c)) + m_3^*((c, b)),$$

and so intervals are non-measurable!

The upshot of all this is that if we want the measure induced by ℓ to be “sensible”—that is, if we want it to agree with ℓ on intervals, and if we want intervals to be measurable—then we must choose ℓ to be σ -additive on the class of intervals. This is a very restrictive condition; indeed, if we include the further requirement that ℓ be translation invariant, the only

possibility for ℓ is a constant multiple of the length function, leading to Lebesgue measure or a scalar multiple thereof.

b. Measures on symbolic space. We now move to another context and discuss measures on the symbolic space Σ_k^+ . As before, we follow the Carathéodory approach, taking for the collection \mathcal{C} of basic sets the cylinders in Σ_k^+ given by (A) on page 22.

Now we define a σ -additive set function m^* on cylinders; choose positive numbers p_1, \dots, p_k such that $p_1 + \dots + p_k = 1$, and let m^* be given by

$$(42) \quad m^*(C_{i_1 \dots i_n}) = p_{i_1} p_{i_2} \dots p_{i_n}.$$

EXAMPLE 43. In the case $k = 2$, the condition on the numbers p_i reduces to $p_1 + p_2 = 1$, and we often write $p = p_1$, $q = p_2 = 1 - p$. Then

$$(43) \quad m^*(C_{i_1 \dots i_n}) = p^j q^{n-j},$$

where j is the number of times 1 appears in the sequence i_1, \dots, i_n .

If we repeatedly toss a weighted coin such that the probability of heads appearing on any given toss is p , and the probability of tails is q , then $m^*(C_{i_1 \dots i_n})$ is the probability that the first n tosses give the result i_1, \dots, i_n , where $i_j = 1$ denotes heads, and $i_j = 2$ denotes tails.

Now we want to build a measure m on Σ_k^+ from the set function m^* . In light of the discussion in the previous section, we hope to obtain a measure for which cylinders are measurable, and which agrees with the set function on cylinders. It can be shown that these will follow once we show that m^* is additive on the class of cylinders; that is, that if C_1, \dots, C_n are cylinders whose union $C_1 \cup \dots \cup C_n$ is also a cylinder, then

$$(44) \quad m^*(C_1 \cup \dots \cup C_n) = m^*(C_1) + \dots + m^*(C_n).$$

This in turn follows from the formula

$$(45) \quad m^*(C_{i_1 \dots i_n}) = m^*(C_{i_1 \dots i_{n-1} 1}) + \dots + m^*(C_{i_1 \dots i_{n-1} k}),$$

which is a direct consequence of (42). The proof that (45) implies (44) is exactly the same argument that we used in the proof of Moran's theorem, with p_1, \dots, p_k taking the place of $\lambda_1^t, \dots, \lambda_k^t$.

The measure m which is built from the set function m^* in (42) is called a *Bernoulli measure*. In fact, there are many such measures, since we may choose any positive parameters p_1, \dots, p_k which sum to 1.

Bernoulli measures all have the following interesting property:

PROPOSITION 44. *Let m be a Bernoulli measure on Σ_k^+ , and $C_{i_1 \dots i_n}$ any n -cylinder. Then*

$$(46) \quad m(\sigma^{-1}(C_{i_1 \dots i_n})) = m(C_{i_1 \dots i_n}).$$

PROOF. Recall that the preimage of the cylinder $C_{i_1 \dots i_n}$ is

$$\begin{aligned} \sigma^{-1}(C_{i_1 \dots i_n}) &= \{ \omega \in \Sigma_k^+ \mid \sigma \omega \in C_{i_1 \dots i_n} \} \\ &= C_{1i_1 \dots i_n} \cup \dots \cup C_{ki_1 \dots i_n}. \end{aligned}$$

Then it follows from (42) that

$$\begin{aligned} m(\sigma^{-1}(C_{i_1 \dots i_n})) &= p_1 m(C_{i_1 \dots i_n}) + \dots + p_k m(C_{i_1 \dots i_n}) \\ &= (p_1 + \dots + p_k) m(C_{i_1 \dots i_n}) \\ &= m(C_{i_1 \dots i_n}). \end{aligned} \quad \square$$

It follows from Proposition 44 that $m(\sigma^{-1}E) = m(E)$ for *any* measurable set $E \subset \Sigma_k^+$; we say that m is an *invariant measure with respect to σ* , or sometimes that it is a *shift-invariant measure*.

c. Measures on Cantor sets. Let C be a Cantor-like set, $h: \Sigma_k^+ \rightarrow C$ the coding map and recall the conjugacy between the shift σ and the map $f: C \rightarrow C$ which is illustrated in the commutative diagram (8). We may use this conjugacy to transfer a measure m on Σ_k^+ to a measure μ on C , as follows.

The measurable subsets of C will be precisely those sets A whose preimages $h^{-1}(A)$ are measurable; we may define the measure of such sets as

$$(47) \quad \mu(A) = m(h^{-1}(A)).$$

After having built one finite measure on C , the Hausdorff measure $m_H(\cdot, \alpha)$, we now have a whole family of finite measures on C which come from the Bernoulli measures on Σ_k^+ . The Hausdorff measure corresponds to the particular choice $p_i = \lambda_i^\alpha$; all the other measures are new.

All of these measures are f -invariant; this follows since

$$\begin{aligned} \mu(f^{-1}(A)) &= m(h^{-1}f^{-1}(A)) \\ &= m(\sigma^{-1}h^{-1}(A)) \\ &= m(h^{-1}(A)) \\ &= \mu(A). \end{aligned}$$

Thus we have one class of f -invariant measures on the Cantor set; the Bernoulli measures. We could consider many other classes of measures as well, such as the point measures. However, as we are interested in studying the dynamics of f , only the invariant measures are really important, and a point measure m_x is only invariant if x is a fixed point for f .

Eventually, we will see an even larger class of invariant measures on the Cantor set, the so-called *Markov measures*, which will be introduced the same way as the Bernoulli measures, via a particular set function on cylinders in symbolic space. The idea motivating their introduction is simple enough; the coin flips described in Example 43 are independent events, as the outcome of a given flip does not depend on the outcomes which came before it. But what if this was not the case? What if the process had some sort of memory, so that the future depends somehow on the past? For example, we could consider a bag containing two colours of marbles, white and black; we pull out a marble at random, and write 1 if it is white, and 2 if it is black. Repeating this, without replacing our earlier picks, gives us

a sequence of 1's and 2's, but now the probabilities at each step depend on the previous results. A process such as this which exhibits finite memory is known as a *Markov process*.

Lecture 16

a. Markov measures. We now turn our attention to another class of measures on Σ_k^+ , called *Markov measures*. We fix k positive numbers p_1, \dots, p_k such that $\sum_i p_i = 1$, and a $k \times k$ matrix $P = (p_{ij})$ with non-negative entries such that

$$(48) \quad \sum_{j=1}^k p_{ij} = 1 \text{ for all } i,$$

called a *stochastic matrix*.

Now we construct a σ -additive set function on cylinders by

$$(49) \quad m(C_{i_1 \dots i_n}) = p_{i_1} p_{i_1 i_2} p_{i_2 i_3} \cdots p_{i_{n-1} i_n},$$

and follow the usual Carathéodory construction to obtain a measure m on Σ_k^+ ; this is the *Markov measure* associated to the probability vector $p = (p_1, \dots, p_k)$ and the stochastic matrix P .

EXERCISE 11. Use the property (48) of a stochastic matrix P to show that the set function defined in (49) is σ -additive on cylinders.

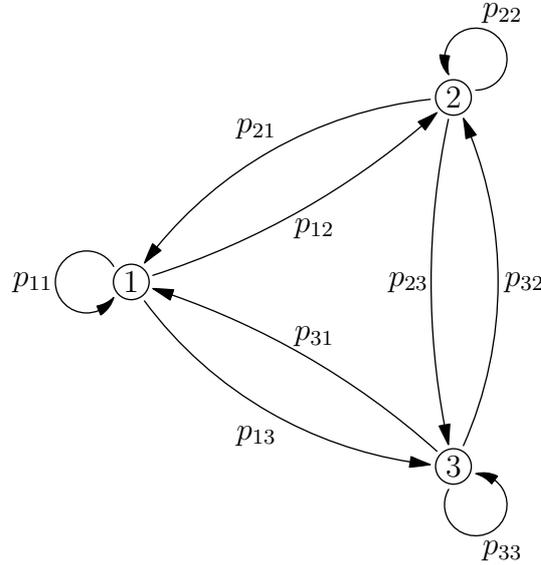
EXERCISE 12. Show that the Markov measure generated by p and P is shift-invariant.

One interpretation of (49) is as follows: consider a process which may give any one of k results at a given time $n \in \mathbb{N}$, and write i_n for the result observed. Then each sequence $\omega = (i_1, i_2, \dots) \in \Sigma_k^+$ represents a particular instance of the process, in which the results i_1, i_2, \dots were observed. As was the case with Bernoulli measures, the Markov measure $m(C_{i_1 \dots i_n})$ of a cylinder gives the probability that the first n results are i_1, \dots, i_n .

Alternatively, we may think of i_n as representing the state of a system at time n . An element p_i of the probability vector gives the probability of beginning in state i , and an entry p_{ij} of the stochastic matrix P gives the probability of going from state i to state j at any given time-step; that is, it gives the conditional probability of being in state j at time $n + 1$, given that the state at time n is i .

We may represent this graphically, as shown in Figure 22 for the case $k = 3$. The vertices of the graph represent the states of the system, and to the edge from vertex i to vertex j is associated a *transition probability* p_{ij} . The condition that the matrix P be stochastic may be rephrased as the condition that the outgoing transition probabilities from any given vertex sum to 1; that is, if we are at vertex i now, then at the next time step, we will be *somewhere*. We may be at vertex i again (with probability p_{ii}), but we will not simply vanish from the graph.

In this interpretation, the symbolic space Σ_k^+ may be thought of as the set of all possible paths along edges of the graph, and the measure of a cylinder $C_{i_1 \dots i_n}$ is the probability that a randomly chosen itinerary on the graph begins with i_1, \dots, i_n .

FIGURE 22. A Markov measure on Σ_3^+ .

Now what happens if one of the transition probabilities is 0? If $p_{ij} = 0$, then the probability of going from vertex i to vertex j is 0, and so we may as well erase this edge from the graph. What does this do to the set of all possible paths, which was in one-to-one correspondence with symbolic space Σ_k^+ ?

Before answering this question, we introduce another definition from measure theory.

b. The support of a measure. Fix a point $x \in \mathbb{R}^n$, and consider the point measure m_x on \mathbb{R}^n defined in Exercise 37. If $E \subset \mathbb{R}^n$ does not contain x , then as far as the measure is concerned, the points in E may be neglected without losing anything important, since $m_x(E) = 0$. Thus from the measure theoretic point of view, the measure space (\mathbb{R}^n, m_x) is the same as the measure space $(\{x\}, m_x)$, since the only difference between the two is a set of zero measure.

In the same vein, given any measure m on \mathbb{R}^n , there is a canonical way to decompose \mathbb{R}^n into two parts, one of which may be discarded, since it has measure zero with respect to m , and the other of which carries all the information about the measure. The latter set is called the *support* of m , and is the smallest closed set of full measure. This definition goes through not just for \mathbb{R}^n , but for any topological space X equipped with a measure; we write

$$(50) \quad \text{supp } m = \{x \in X \mid m(E) > 0 \forall E \ni x, E \text{ open}\}.$$

The measure m sits on the support $\text{supp } m$ in the sense that an open set has positive measure if it intersects the support non-trivially, and zero measure otherwise.

EXAMPLE 45. Let X be the unit interval $[0, 1]$. If m is Lebesgue measure, then $m((a, b)) = b - a > 0$ for all $a < b$, so open intervals have positive measure. Since open sets are countable unions of open intervals, we see that all open sets have positive measure, hence $\text{supp } m = [0, 1]$.

For the point measure m_x , the above discussion may be summarised as the statement that $\text{supp } m_x = \{x\}$.

EXAMPLE 46. Let X be the symbolic space Σ_k^+ . If m is the Bernoulli measure given by a probability vector with positive entries, we see from (42) that cylinders have positive measure; since any open set is a union of cylinders, it follows that the support of m is the entire space Σ_k^+ .

Similarly, if m is a Markov measure for which all entries of both the probability vector and the stochastic matrix are positive, then all cylinders have positive measure by (49), and we again have $\text{supp } m = \Sigma_k^+$.

Returning to our question from the previous section, we may ask what $\text{supp } m$ looks like if one or more of the entries p_{ij} are equal to 0. In this case, we see that the measure of a cylinder $C_{i_1 \dots i_n}$ is positive precisely when all the entries $p_{i_j i_{j+1}}$ are positive for $j = 1, \dots, n-1$. In particular, if any of the entries in the stochastic matrix vanish, then some cylinders have zero measure; since cylinders are open, this means that the support of m is *not* the whole space Σ_k^+ .

To describe $\text{supp } m$, we introduce a $k \times k$ matrix whose entries are either 0 or 1, which keeps track of which entries of P vanish; let $A = (a_{ij})$ be defined by

$$a_{ij} = \begin{cases} 0 & p_{ij} = 0, \\ 1 & p_{ij} > 0. \end{cases}$$

A is called a *transition matrix*; it records which transitions $i_j \rightarrow i_{j+1}$ have a non-zero probability of occurring. We say that a sequence $\omega = (i_1, i_2, \dots)$ is *admissible* if $a_{i_j i_{j+1}} = 1$ for all j ; that is, if we only follow edges in the graph which carry a non-zero probability.

Note that a cylinder has positive measure precisely when it contains an admissible sequence, and so we consider the set of admissible sequences,

$$(51) \quad \Sigma_A^+ = \{\omega = (i_1, i_2, \dots) \in \Sigma_k^+ \mid a_{i_j i_{j+1}} = 1 \ \forall j\}.$$

If $E \subset \Sigma_k^+$ is open and intersects Σ_A^+ non-trivially, there exists a cylinder $C_{i_1 \dots i_n} \subset E \cap \Sigma_A^+$, and we see that $m(C_{i_1 \dots i_n}) > 0$ since $p_{i_j i_{j+1}} > 0$ for all $j = 1, \dots, n-1$; it follows that $m(E) \geq m(C_{i_1 \dots i_n}) > 0$, and thus $\text{supp } m = \Sigma_A^+$.

In Lecture 15(c), we saw how to use the coding map h to go from a measure m on symbolic space Σ_k^+ to a measure $\mu = m \circ h^{-1}$ on a Cantor set $C \subset [0, 1]$. Let m be a Markov measure whose support is not the whole

space; then some cylinders in Σ_k^+ have zero measure, and hence their images, which are open sets in C , have zero measure as well.

In particular, we have $\text{supp } \mu \subsetneq C$, where the points which are removed from C are precisely those whose codings are not admissible; that is, at each step of the construction we remove certain basic intervals which would introduce an inadmissible coding. The result of this *Markov construction* is another Cantor-like set, $\text{supp } \mu$, which we would like to gain some information about. For example, what is the Hausdorff dimension of the support of μ ? Because we have deleted basic intervals at each step of the construction, we can no longer apply Moran's theorem, and the question becomes much more complicated. A version of Moran's argument still goes through, but it must be modified substantially, and becomes much more subtle in this case.

c. Markov measures and dynamics. *A priori*, the Markov constructions described in the previous section seem rather artificial and contrived; it may not be obvious just why such constructions would be important, or where they would appear.

In fact, they are tremendously important in dynamics; as a very simple example, consider a piecewise linear map of the interval $[0, 1]$ of the sort shown in Figure 23.

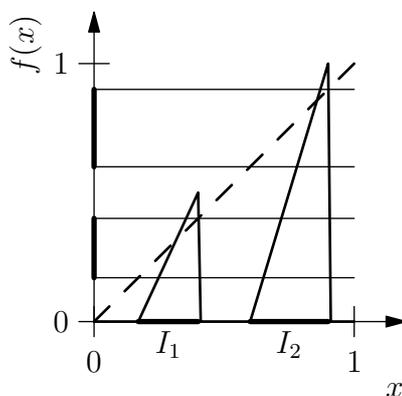


FIGURE 23. A map which is modelled by a subshift.

This map generalises the sort of map we saw in Figure 11; before, we demanded that the image of each interval I_i be the entire interval $[0, 1]$, whereas now we demand only that $f(I_i)$ be a subinterval which contains every I_i it intersects. As before, we may consider the set of points whose trajectories remain within the domain of definition of f ; this is a repelling Cantor set, which can be modelled by subsets Σ_A^+ of the symbolic space Σ_k^+ , but not by the entire space. For the map shown in Figure 23, we see that points in I_2 can be mapped to either I_1 or I_2 , but points in I_1 can only be mapped to I_1 . Thus the transition matrix is $A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$.

It is obvious that if $\omega \in \Sigma_A^+$ is an admissible sequence, then its image $\sigma(\omega)$ under the shift will be admissible as well. Thus $\sigma(\Sigma_A^+) \subset \Sigma_A^+$, and we say that Σ_A^+ is *invariant* under the action of σ . We may consider the restriction of the shift σ to the domain Σ_A^+ , which is called a *subshift of finite type*; for the sake of clarity, the shift $\sigma: \Sigma_k^+ \rightarrow \Sigma_k^+$ is often referred to as the *full shift*.

Subshifts of finite type are a very important class of models in dynamics, whose applications stretch far beyond the simple interval maps we have mentioned so far. Together with the Markov measures sitting on them, they provide a powerful tool with which to study questions in dimension theory and to examine the stochastic and chaotic properties of certain dynamical systems.

Lecture 17

a. Using measures to determine dimension. The measures we have introduced so far are connected to dynamics by the fact that Bernoulli and Markov measures on symbolic space are shift-invariant, and their projections to the Cantor set are invariant under the action of f . Now we provide some justification for our claim that these measures can also tell us something about the dimension of certain sets.

A finite measure is often called a *mass distribution*, since we may think of it as describing how a specific amount of mass is distributed over a space. If the measure is normalised so that $\mu(X) = 1$, we often call μ a *probability measure*, since in this case subsets of X may be thought of as events in a probabilistic process.

The following theorem shows that existence of finite measures, or mass distributions, with particular scaling properties has consequences for the Hausdorff dimension of sets with positive measure.

THEOREM 47 ((Uniform) Mass Distribution Principle). *Suppose μ is a finite measure on \mathbb{R}^n , and that there exist $\alpha, K, \delta > 0$ such that for every ball $U = B(x, r)$ with radius $r \leq \delta$, we have*

$$(52) \quad \mu(U) \leq Kr^\alpha.$$

Then if $E \subset \mathbb{R}^n$ is measurable and $\mu(E) > 0$, we have $\dim_H E \geq \alpha$.

PROOF. For any $0 < \varepsilon < \delta$ and any cover $\mathcal{U} = \{U_i\}$ of E by balls of diameter less than ε , we have

$$\sum_i (\text{diam } U_i)^\alpha \geq \sum_i \frac{\mu(U_i)}{K} \geq \frac{\mu(\bigcup_i U_i)}{K} \geq \frac{\mu(E)}{K} > 0,$$

and hence $m_H(Z, \alpha) \geq \mu(E)/K > 0$, so $\dim_H E \geq \alpha$. \square

Theorem 47 gives a lower bound for the Hausdorff dimension, which is typically the more difficult bound to obtain. Indeed, finding a measure μ which satisfies the hypothesis of the theorem can be quite hard, and so we will eventually see how to weaken the assumption on μ . First, though, we consider the particular case of a Cantor-like construction in $[0, 1]$ with ratio coefficients $\lambda_1, \dots, \lambda_k > 0$, $\sum_i \lambda_i < 1$, and describe a measure μ for which the Mass Distribution Principle may be applied.

As we showed in the proof of Moran's theorem, if $\mathcal{V}(r)$ denotes the collection of basic intervals whose diameter is between $\lambda_{\min}r$ and r/λ_{\min} , then each ball of radius r intersects at most M' elements of $\mathcal{V}(r)$, where M' is a constant independent of r . In particular, $B(x, r)$ can be covered by at most M' elements of $\mathcal{V}(r)$.

If we can show that (52) holds for basic intervals, then for each element I of $\mathcal{V}(r)$ we will have $\mu(I) \leq K(\text{diam } I)^\alpha \leq K'r^\alpha$, where $K' = K/\lambda_{\min}^\alpha$, and hence $B(x, r)^\alpha \leq M'K'r^\alpha$. Thus it suffices to consider basic intervals.

Now for a Bernoulli measure μ given by a probability vector (p_1, \dots, p_k) , we have $\mu(I_{i_1 \dots i_n}) = p_{i_1} \cdots p_{i_n}$, and $(\text{diam } I_{i_1 \dots i_n})^\alpha = \lambda_{i_1}^\alpha \cdots \lambda_{i_n}^\alpha (|I_{i_1}|^\alpha / \lambda_{i_1}^\alpha)$. Thus in order to have $\mu(I_{i_1 \dots i_n}) / (\text{diam } I_{i_1 \dots i_n})^\alpha$ bounded by a constant independent of (i_1, \dots, i_n) , we take $p_i = \lambda_i^\alpha$, where α is the root of Moran's equation (23). This defines a Bernoulli measure μ which satisfies (52), and so the Mass Distribution Principle gives the lower bound $\dim_H C \geq \alpha$, since $\mu(C) = 1 > 0$.

Of course in this case we already knew the answer, but this illustrates the principle of using measures to obtain bounds on dimensional quantities which might otherwise be difficult to gain information about. Certainly the calculations in the proof of Theorem 47 are simpler than those involved in computing the potential associated to a cover by basic intervals, and showing that we can move between various levels of the construction without changing the potential. Further, the present argument shows that *any* subset of C which is given positive measure by the Bernoulli measure μ must have Hausdorff dimension equal to α , which we could not have deduced directly from Moran's theorem.

However, Cantor-like sets are somewhat unique in admitting measures satisfying the conditions of the Mass Distribution Principle. Another example is Lebesgue measure on \mathbb{R}^n , but further progress beyond these two rather restrictive cases requires us to introduce a weaker condition on the measure, which still gives us information about the dimension of certain sets.

b. Pointwise dimension. In the statement of Theorem 47, the constants α, K, δ appearing in (52) were to be independent of both scale and position; that is, they could not depend on either the centre x or the radius r of the ball $U = B(x, r)$. We now consider a slightly more general condition, in which the constants are still independent of r , but may vary between different positions in \mathbb{R}^n , i.e., may depend on x .

To this end, we define the *pointwise dimension of μ at x* as the limit

$$(53) \quad d_\mu(x) = \lim_{r \rightarrow 0} \frac{\log \mu(B(x, r))}{\log r},$$

if the limit exists. Of course, the limit may not exist at every point x , and so the pointwise dimension may not be defined for all $x \in \mathbb{R}^n$. We emphasise that the pointwise dimension is a property of the *measure* μ , rather than of any particular set; this is in contrast to previous dimensional quantities we have seen, which were all properties of sets.

The existence of the pointwise dimension at a point x leads to the following estimates: for every $\varepsilon > 0$, there exists $\delta > 0$ such that for every $0 < r < \delta$,

$$d_\mu(x) - \varepsilon \leq \frac{\log \mu(B(x, r))}{\log r} \leq d_\mu(x) + \varepsilon.$$

Since $r < 1$, the function $t \mapsto r^t$ is decreasing in t , and so

$$r^{d_\mu(x) - \varepsilon} \geq \mu(B(x, r)) \geq r^{d_\mu(x) + \varepsilon},$$

which gives us bounds reminiscent of (52). However, the scale δ for which this estimate holds may vary from point to point, and $d_\mu(x)$ itself may also vary, and may not exist everywhere. Despite this, we will eventually see that the pointwise dimension provides a useful tool for gaining information about the dimension of a set by considering an appropriate measure.

EXAMPLE 48. Consider a piecewise linear map $f: I_1 \cup I_2 \rightarrow [0, 1]$ as in Figure 8, where $|I_1| = |I_2| = \lambda < 1/2$. The repeller of f is a Cantor-like set C modelled on Σ_2^+ , and both ratio coefficients in the construction of C are equal to λ .

Let m be a Bernoulli measure on Σ_2^+ with probability vector (p, q) , where $p, q > 0$, $p + q = 1$, and let μ be the corresponding Bernoulli measure on C . Given $x \in C$, what is the pointwise dimension $d_\mu(x)$?

Say $x = h(\omega) = h(i_1, i_2, \dots)$; the pointwise dimension is determined by the measure of the balls $B(x, r)$ centred at x , and for appropriate values of r , these are just the basic intervals $I_{i_1 \dots i_n}$. We have $|I_{i_1 \dots i_n}| = \lambda^n$, and so $\log r = n \log \lambda$ (up to some constant difference which will vanish in the limit). Furthermore, $\mu(I_{i_1 \dots i_n}) = p^k q^{n-k}$, where k is the number of times the symbol 1 appears in the string (i_1, \dots, i_n) .

In the simplest case, we have $p = q = 1/2$, and so $\mu(I_{i_1 \dots i_n}) = (1/2)^n$. Thus the pointwise dimension is given by

$$(54) \quad d_\mu(x) = \lim_{n \rightarrow \infty} \frac{\log(\mu(I_{i_1 \dots i_n}))}{n \log \lambda} = \frac{\log 2}{-\log \lambda} = \dim_H C.$$

Observe that in the example above, the Cantor set C is just the support of the measure μ ; this suggests some connection between the pointwise dimension of a measure and the Hausdorff dimension of its support. However, the relationship is not always this simple. For most choices of p and q , the pointwise dimension will not exist everywhere, and will not always equal $\dim_H C$ where it does exist. Indeed, consider the points $x = h(1, 1, 1, \dots)$ and $y = h(2, 2, 2, \dots)$ in C ; an easy calculation shows that

$$d_\mu(x) = \frac{\log p}{\log \lambda}, \quad d_\mu(y) = \frac{\log q}{\log \lambda},$$

and the two are not equal except in the special case $p = q = 1/2$.

Thus the true relationship between the pointwise dimension of a measure and the Hausdorff dimension of its support is somewhat more subtle, and we will take this up in the next lecture.

Lecture 18

a. The Non-uniform Mass Distribution Principle. Example 48 showed that the pointwise dimension of a measure can give us some information about the Hausdorff dimension of a set. Of course, things are rarely as simple as in that example, where the pointwise dimension exists at each point, and is equal at *every* point to the Hausdorff dimension. In general, the limit in (53) may not exist, and so we may consider instead the lower and upper limits, which always exist; these are referred to as the *lower and upper pointwise dimensions* of the measure μ at the point x , and denoted by $\underline{d}_\mu(x)$ and $\overline{d}_\mu(x)$, respectively.

We have $\underline{d}_\mu(x) \leq \overline{d}_\mu(x)$ for any measure μ and any point x ; the two coincide if and only if the limit in (53) exists, in which case their common value is the pointwise dimension.

The following result allows us to weaken the hypotheses of Theorem 47:

THEOREM 49 (Non-uniform Mass Distribution Principle). *Suppose μ is a finite measure on \mathbb{R}^n , that $E \subset \mathbb{R}^n$ has positive measure ($\mu(E) > 0$), and that there exists $\alpha > 0$ such that*

$$(55) \quad \underline{d}_\mu(x) \geq \alpha$$

for almost every $x \in E$. Then $\dim_H E \geq \alpha$.

PROOF. Given $\varepsilon > 0$, we show that $m_H(E, \alpha - \varepsilon) > 0$, as follows. For almost every point $x \in E$, we have

$$\alpha \leq \underline{d}_\mu(x) = \liminf_{r \rightarrow 0} \frac{\log \mu(B(x, r))}{\log r},$$

and so there exists $\delta > 0$ such that for all $0 < r < \delta$,

$$\alpha - \varepsilon \leq \frac{\log \mu(B(x, r))}{\log r},$$

which leads to the inequality

$$(56) \quad r^{\alpha - \varepsilon} \geq \mu(B(x, r)).$$

We would like to argue that by picking δ small enough, we can use (56) to proceed exactly as in the proof of Theorem 47; however, since δ may depend on x , we must be slightly more careful.

Let $\delta_n = 1/n$, and write

$$E_n = \{x \in E \mid (56) \text{ holds for all } 0 < r < \delta_n\}.$$

Because (55) holds for almost every $x \in E$, we have $\mu\left(\bigcup_{n \geq 1} E_n\right) = 1$, and so there exists n such that $\mu(E_n) > 0$.

It suffices to show that $m_H(E, \alpha - \varepsilon) \geq m_H(E_n, \alpha - \varepsilon) > 0$. Indeed, for any open cover $\{B(x_i, r_i)\}$ of E_n by balls of radius $r_i \leq \delta_n$, we have

$$\begin{aligned} \sum_i \text{diam}(B(x_i, r_i))^{\alpha - \varepsilon} &= 2^{\alpha - \varepsilon} \sum_i r_i^{\alpha - \varepsilon} \\ &\geq 2^{\alpha - \varepsilon} \sum_i \mu(B(x_i, r_i)) \\ &\geq 2^{\alpha - \varepsilon} \mu\left(\bigcup_i B(x_i, r_i)\right) \\ &\geq 2^{\alpha - \varepsilon} \mu(E_n). \end{aligned}$$

It follows that $m_H(E_n, \alpha - \varepsilon) \geq 2^{\alpha - \varepsilon} \mu(E_n) > 0$, and so $\dim_H E_n \geq \alpha - \varepsilon$. By the monotonicity property of Hausdorff dimension, $\dim_H E \geq \alpha - \varepsilon$, and since $\varepsilon > 0$ is arbitrary, we have $\dim_H E \geq \alpha$. \square

b. Non-constant pointwise dimension. Let us return to the setting of Example 48, a Cantor-like repeller C in the interval for a piecewise linear map f as in Figure 8, where the two basic intervals I_1 and I_2 have equal length λ . Consider the Bernoulli measure μ on C which is generated by the probability vector (p, q) , so that $\mu(I_1) = p$, $\mu(I_2) = q$. We want to compute the pointwise dimension of μ at a point $x \in C$, and relate this to the Hausdorff dimension of C , which as we already know from Moran's theorem, is $\dim_H C = -\log 2 / \log \lambda$.

To begin with, recall that we have a one-to-one correspondence between elements of the Cantor set C and elements of the symbolic space Σ_2^+ , and so we consider the sequence $(i_1, i_2, \dots) = h^{-1}(x)$ which gives the coding of the point x .

In order to compute the ratio in (53), we may replace $B(x, r)$ with $I_{i_1 \dots i_n}$, the basic interval containing x , and r with $|I_{i_1 \dots i_n}| = \lambda^n$, since the error term between the two ratios vanishes in the limit. From the definition of a Bernoulli measure, we have $\mu(I_{i_1 \dots i_n}) = p^{k_n} q^{n - k_n}$, where k_n is the number of times the symbol "1" appears in the sequence i_1, \dots, i_n .

Proceeding naïvely and ignoring any questions regarding existence of limits, we see that

$$\begin{aligned} d_\mu(x) &= \lim_{n \rightarrow \infty} \frac{\log \mu(I_{i_1 \dots i_n})}{\log |I_{i_1 \dots i_n}|} \\ &= \lim_{n \rightarrow \infty} \frac{\log(p^{k_n} q^{n - k_n})}{\log \lambda^n} \\ (57) \quad &= \lim_{n \rightarrow \infty} \frac{k_n \log p + (n - k_n) \log q}{n \log \lambda} \\ &= \frac{(\lim_{n \rightarrow \infty} \frac{k_n}{n}) \log p + (1 - \lim_{n \rightarrow \infty} \frac{k_n}{n}) \log q}{\log \lambda}. \end{aligned}$$

Thus everything hinges on the value (and existence) of the limit

$$(58) \quad \alpha(x) = \lim_{n \rightarrow \infty} \frac{k_n(x)}{n}.$$

The ratio k_n/n is simply the proportion of ones in the first n entries of (i_1, i_2, \dots) , and so the limit (if it exists) is the asymptotic frequency of ones. But what is this value, and when does it exist?

Consider first the simple case $p = q = 1/2$. Then we may think of μ as giving the probabilities of particular sequences of outcomes of a random process, where at each step, we have an equal probability of choosing a 1 or a 2 for the sequence, as if we were flipping a fair coin. Intuitively, we expect such a process to yield an approximately equal number of ones and twos in the long run, and so we expect to see the ratio k_n/n converge to $1/2$ as n goes to infinity.

Of course, we can construct particular sequences for which this is not the case; for $x = h(1, 1, 1, \dots)$, we have $k_n = n$ for all n , and so $\alpha(x) = 1$. Similarly, $h(2, 2, 2, \dots)$ gives a limiting value of 0, and in fact, given any $\alpha_0 \in [0, 1]$, it is not hard to construct an $x \in C$ for which $\alpha(x) = \alpha_0$; there are also many points for which the limit does not exist. So we can only expect the limit to be equal to $1/2$ for certain “good” points x , which we hope are in some way typical. But typical in what sense?

Observe that a point x with periodic coding $(i_1, i_2, \dots, i_N, i_1, i_2, \dots)$ has an asymptotic frequency of ones which is equal to k_N/N ; for most such points, this is not equal to $1/2$, and since such points are dense in C , our set of “bad” points (those which are somehow atypical) is topologically quite large.

The precise statement, which we do not prove here, is that $\alpha(x) = 1/2$ *almost everywhere*, or for *almost every* x with respect to μ , meaning that the set

$$E = \left\{ x \in C \mid \alpha(x) \neq \frac{1}{2} \right\}$$

has measure zero, $\mu(E) = 0$. This is a very common condition in measure theory, where many properties can be shown to hold almost everywhere, but not necessarily at all points.

Now we can complete the calculation in (57) for the case $p = q = 1/2$ to obtain

$$d_\mu(x) = \frac{\log 2}{-\log \lambda} = \dim_H C$$

for almost every $x \in C$. In and of itself, this is not very helpful; after all, we already saw in (54) that this is true for *every* $x \in C$. However, that result did not generalise to other values of p and q , while this one does; it can be shown that for any Bernoulli measure on C with $p, q \geq 0$, $p + q = 1$, we have $\alpha(x) = \lim_{n \rightarrow \infty} k_n(x)/n = p$ almost everywhere. In particular, for almost

every $x \in C$, (57) gives

$$(59) \quad d_\mu(x) = \frac{p \log p + q \log q}{\log \lambda}.$$

It follows from Theorem 49 that this quantity gives a lower bound on the Hausdorff dimension $\dim_H C$, and indeed one can check that the right hand side of (59) takes its maximum value when $p = q = 1/2$, and we have $d_\mu(x) = \dim_H C$ almost everywhere.

The numerator in (59) is related to a very general concept in dynamical systems, for which we will borrow the notation and terminology without getting into the specifics of the general definition.

DEFINITION 50. Let μ be a Bernoulli measure on C with probability vector (p, q) . The *entropy* of the map $f: C \rightarrow C$ with respect to μ is

$$h_\mu(f) = -(p \log p + q \log q).$$

Now suppose that the basic intervals I_1, I_2 in the definition of f have different lengths $\lambda_1 \neq \lambda_2$. Then the length of the basic interval $I_{i_1 \dots i_n}$ depends on $k_n(x)$, and so in the calculation of the pointwise dimension $d_\mu(x)$, the denominator in (57) must be replaced by

$$\log |I_{i_1 \dots i_n}| = \log(\lambda_1^{k_n} \lambda_2^{n-k_n}) = k_n \log \lambda_1 + (n - k_n) \log \lambda_2.$$

Thus we have, for almost every $x \in C$,

$$d_\mu(x) = \frac{p \log p + q \log q}{p \log \lambda_1 + q \log \lambda_2},$$

which we rewrite as

$$(60) \quad d_\mu(x) = \frac{h_\mu(f)}{\lambda_\mu(f)},$$

where $\lambda_\mu(f) = -(p \log \lambda_1 + q \log \lambda_2)$ is the *Lyapunov exponent* of f with respect to μ .

Like the entropy, the Lyapunov exponent can be defined for a much broader class of maps than the piecewise linear ones we have been considering. In particular, both $h_\mu(f)$ and $\lambda_\mu(f)$ can be defined for non-linear maps f such as the one in Figure 17, and any invariant measure μ . If C is the maximal repelling Cantor set for f , then the ratio coefficients change at each step in the construction of C , and so Moran's theorem is of no use to us in determining the Hausdorff dimension of C . However, one can show that (60) holds for almost every $x \in C$, and so by Theorem 49, we have

$$(61) \quad \dim_H C \geq \frac{h_\mu(f)}{\lambda_\mu(f)}.$$

That is, the Hausdorff dimension of the Cantor set is bounded below by the ratio of the entropy and the Lyapunov exponent for any invariant measure μ . Since this ratio depends on the particular choice of measure, (61) is really a whole family of lower bounds on $\dim_H C$, and it is natural to ask if one

can find an invariant measure μ for which equality is achieved, and the ratio between the entropy and the Lyapunov exponent gives exactly the Hausdorff dimension.

It turns out that this is in fact possible, and that one can use the entropy and the Lyapunov exponent to determine the Hausdorff dimension for a much broader class of Cantor-like sets than can be dealt with by Moran's theorem. The description given here is only the barest sketch of the approach, and the full argument is rather involved; however, this provides a powerful approach to computing Hausdorff dimension in the general case.

Lecture 19

a. More about the Lyapunov exponent. We now take a closer look at the Lyapunov exponent $\lambda_\mu(f)$. In particular, we examine the interpretation of this quantity in terms of the dynamics of the map f , and give a more general definition which applies to any map f on the real line.

Consider a continuously differentiable map $f: E \rightarrow \mathbb{R}$, where $E \subset \mathbb{R}$ is the domain of definition of f , and let $x, y \in E$ be two points which are close together. We want to compare their trajectories, to see how the distance between $f^n(x)$ and $f^n(y)$ varies with n .

Taking the Taylor expansion of f around x gives

$$f(y) = f(x) + f'(x)(y - x) + o(y - x),$$

and so the distance between $f(x)$ and $f(y)$ is given by

$$d(f(x), f(y)) = |f(x) - f(y)| \approx |f'(x)| d(x, y),$$

where the error term is negligible with respect to $d(x, y)$. Thus the distance between the trajectories is multiplied by a factor of approximately $|f'(x)|$ when we pass from x and y to their images under f .

Passing to the second iterates, we have

$$\begin{aligned} d(f^2(x), f^2(y)) &= d(f(f(x)), f(f(y))) \\ &\approx |f'(f(x))| d(f(x), f(y)) \\ &\approx |f'(f(x))f'(x)| d(x, y), \end{aligned}$$

and in general, after n iterations the estimate is

$$(62) \quad d(f^n(x), f^n(y)) \approx \left(\prod_{i=0}^{n-1} |f'(f^i(x))| \right) d(x, y).$$

The error term in (62) is of order $o(d(f^{n-1}(x), f^{n-1}(y)))$, and so if the trajectories become far enough apart, the error term becomes large and the estimate is no longer valid. However, the closer together we choose x and y to be, the longer it takes for this to happen, and so we can make the estimate valid for as long as we like by choosing x and y close enough together.

We are interested, then, in the behaviour of $d_n(x) = \prod_{i=0}^{n-1} |f'(f^i(x))|$, which gives the amount of expansion in a neighbourhood of x after n iterations. In the case when f is piecewise linear as in Figure 11, the derivative f' takes two values:

$$(63) \quad f'(x) = \begin{cases} \lambda_1^{-1} & x \in I_1, \\ \lambda_2^{-1} & x \in I_2. \end{cases}$$

Thus the rate of growth of $d_n(x)$ depends on whether the iterates $f^i(x)$ are in I_1 or in I_2 . We would like to have some information on this rate of growth by finding some real number λ for which the product asymptotically behaves like $e^{\lambda n}$. This is made precise by the following definition.

DEFINITION 51. The *Lyapunov exponent* of the map f at the point x is

$$(64) \quad \lambda(x) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\prod_{i=0}^{n-1} |f'(f^i(x))| \right),$$

if the limit exists.

Thus the Lyapunov exponent tells us how quickly the distance between two nearby points grows under repeated iterations of f ; it may be thought of as the rate of expansion of the map.

PROPOSITION 52. *Let $f: I_1 \cup I_2 \rightarrow [0, 1]$ be piecewise linear, as in Figure 11, with $|I_1| = \lambda_1$ and $|I_2| = \lambda_2$, and let μ be a Bernoulli measure on C with probability vector (p, q) . Then the Lyapunov exponent exists μ -almost everywhere, and $\lambda(x) = \lambda_\mu(f) = -(p \log \lambda_1 + q \log \lambda_2)$.*

PROOF. The basic interval $I_{i_1 \dots i_n}$ consists of precisely those points x for which $f^{j-1}(x) \in I_{i_j}$ for each $j = 1, \dots, n$. Thus for $x \in I_{i_1 \dots i_n}$,

$$(65) \quad d_n(x) = \prod_{i=0}^{n-1} |f'(f^i(x))| = \prod_{j=1}^n \lambda_{i_j}^{-1} = \lambda_1^{-k_n} \lambda_2^{-(n-k_n)},$$

where k_n is the number of times the symbol “1” appears in i_1, i_2, \dots, i_n . Recalling that $k_n(x)/n \rightarrow p$ for μ -almost every point x , we may compute the Lyapunov exponent $\lambda(x)$ for such points:

$$\begin{aligned} \lambda(x) &= \lim_{n \rightarrow \infty} \frac{1}{n} \log d_n(x) \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\lambda_1^{-k_n(x)} \lambda_2^{-(n-k_n(x))} \right) \\ &= \lim_{n \rightarrow \infty} - \left(\frac{k_n(x)}{n} \log \lambda_1 + \frac{n - k_n(x)}{n} \log \lambda_2 \right) \\ &= -(p \log \lambda_1 + (1 - p) \log \lambda_2) = \lambda_\mu(f). \quad \square \end{aligned}$$

b. Fractals within fractals. The definition of the Lyapunov exponent in the previous section is more general than the definition given earlier after (60); it extends that definition in the sense made precise in Proposition 52.

In fact, there are many important examples of maps f and invariant measures μ for which the first part of Proposition 52 holds, and the Lyapunov exponent $\lambda(x)$ exists and is equal to a constant almost everywhere. Included in this class are non-linear uniformly expanding maps $f: I_1 \cup I_2 \rightarrow [0, 1]$ of the sort shown in Figure 17, where f is continuously differentiable on I_1 and I_2 , is uniformly expanding ($|f'(x)| \geq a > 1$ for some a which is independent of x), and $f(I_1) = f(I_2) = [0, 1]$, and a broad class of invariant measures for such maps.

When $\lambda(x)$ is constant μ -almost everywhere, we denote the constant value by $\lambda_\mu(f)$. The definition of $h_\mu(f)$ can also be extended to measures

which are not Bernoulli, and then it can be shown in many cases that the pointwise dimension exists and is given by (60) almost everywhere.

In such cases, the Non-uniform Mass Distribution Principle (Theorem 49) then gives the lower bound $\dim_H C \geq d_\mu$; in fact, this is a whole family of lower bounds, one for each measure μ . It is natural to ask if we can find an “optimal” measure μ which gives the greatest of all such bounds; that is, does there exist a measure μ such that the pointwise dimension d_μ is maximal?

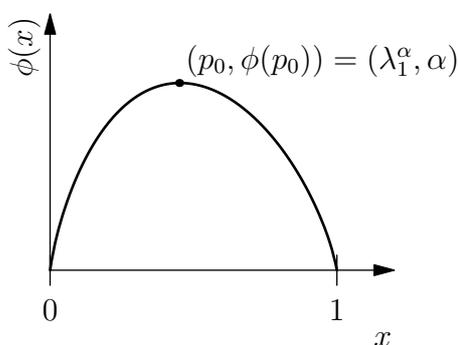


FIGURE 24. Pointwise dimension of the Bernoulli measures μ_p .

Let us consider this question in the particular setting of a piecewise linear map f on two intervals, with the family of Bernoulli measures. Write μ_p for the Bernoulli measure with probability vector $(p, 1-p)$, and define a function $\phi: [0, 1] \rightarrow \mathbb{R}$ by

$$\phi(p) = d_{\mu_p} = \frac{p \log p + (1-p) \log(1-p)}{p \log \lambda_1 + (1-p) \log \lambda_2};$$

the graph of a typical ϕ is shown in Figure 24. It follows from (60) and (61) that $\dim_H C \geq \phi(p)$ for all $p \in [0, 1]$; in order to find the best bound, one solves for the critical point p_0 at which $\phi'(p_0) = 0$ and the function ϕ achieves its maximum.

EXERCISE 13. Show that the function ϕ achieves its maximum at the point $p_0 = \lambda_1^\alpha$, where α is the unique solution of Moran’s equation $\lambda_1^\alpha + \lambda_2^\alpha = 1$, and that $\phi(p_0) = \alpha$.

The result of Exercise 13 shows that the bound $\dim_H C \geq \phi(p_0)$ is optimal, since we have $\dim_H C = \alpha = \phi(p_0) = \sup_p \phi(p)$. The probability vector associated to the Bernoulli measure μ_{p_0} is $(p_0, 1-p_0) = (\lambda_1^\alpha, \lambda_2^\alpha)$, and μ_{p_0} itself turns out to be the Hausdorff measure $m_H(\cdot, \alpha)$.

Let us step back a moment to take stock of all this, and survey the picture we have just painted. We begin with a Cantor set C , whose structure we hope to understand by examining various invariant measures for the map f . The Lebesgue measure of C is 0, so the measures which we use are very

different from Lebesgue measure; to use the language of measure theory, they are *singular*. To each $p \in [0, 1]$ we can associate a Bernoulli measure μ_p ; upon throwing away a null set (a set of measure zero with respect to μ_p), we are left with a set $C_p \subset C$ of full measure, that is, $\mu_p(C_p) = 1$, on which the pointwise dimension $d_{\mu_p}(x)$ is constant.

The fact that pointwise dimension exists and is constant everywhere on C_p means that the measure of a small ball whose centre is in C_p scales as

$$(66) \quad \mu_p(B(x, r)) \approx r^{d_{\mu_p}} = r^{\phi(p)}.$$

This exhibits a sort of measure-theoretic self-similarity on C_p , which is an analogue of the geometric self-similarity possessed by certain very regular fractals such as the middle-thirds Cantor set and the Sierpiński gasket, where the geometric structure appears the same at every scale. Such geometric regularity is not present in the Cantor-like sets generated by non-linear expanding maps, but (66) displays the measure-theoretic regularity of μ_p across all scales.

One may also recall the relationship between existence of the pointwise dimension at x and the convergence of the ratio $k_n(x)/n$, which measures the relative amount of time the orbit of x spends in I_1 . While this ratio approaches many different limits depending on x (and sometimes fails to converge at all), the limit exists and is equal to p for every $x \in C_p$. Thus C_p also has a very regular structure in terms of how the trajectories of points in C_p apportion their time between I_1 and I_2 .

These are some of the ways in which the set C_p is a better fractal set than C itself, since all the points in C_p have more or less the same “average” behaviour. For a given Bernoulli measure μ_p , C_p is the set of all points which are “typical” with respect to μ_p , and so C_p depends on the choice of p . In fact, if $\lambda_1 \neq \lambda_2$, then one can show that these sets are pairwise disjoint; $C_p \cap C_{p'} = \emptyset$ for $p \neq p'$. Thus $\{C_p \mid 0 \leq p \leq 1\}$ gives a whole *family* of “fractals within fractals”, illustrating the deeply complicated *multifractal* structure of the Cantor set C . If we let C' be the set containing all points of C which do not lie in any of the sets C_p , then we can write the *multifractal decomposition* of C :

$$C = \left(\bigcup_{p \in [0,1]} C_p \right) \cup C'.$$

Furthermore, each C_p is f -invariant, and supports the measure μ_p in the sense that $\mu_p(C_p) = 1$. This is the beginning of what is known as *multifractal analysis*.

c. Hausdorff dimension for Markov constructions. Returning to the case where f is linear on each of I_1 and I_2 , suppose further that $\lambda_1 = \lambda_2 = \lambda$, and let C be the repelling Cantor set for f .

So far we have only considered Bernoulli measures on C , but there are other invariant measures as well. In Lecture 16 we described the class of

Markov measures on Σ_2^+ , which are determined by a probability vector (p, q) and a stochastic matrix $P = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix}$; the coding map $h: \Sigma_2^+ \rightarrow C$ lets us turn each of these measures into a measure μ on C .

Now suppose we have a Markov measure for which $p_{11} = 0$. Because the entries in each row of P must sum to 1, we have $p_{12} = 1$, and so the stochastic matrix has the form $P = \begin{pmatrix} 0 & 1 \\ a & 1-a \end{pmatrix}$, where $a = p_{21} \in [0, 1]$. If we interpret the entries of P as giving the probability of a transition from one state to another, then the fact that p_{11} vanishes means that every time we are in state 1, we must immediately return to state 2 at the next time step, with zero probability of remaining in state 1. The associated transition matrix is $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, and we recall from Lecture 16 that the support of m in Σ_2^+ is the set Σ_A^+ of admissible sequences, defined in (51).

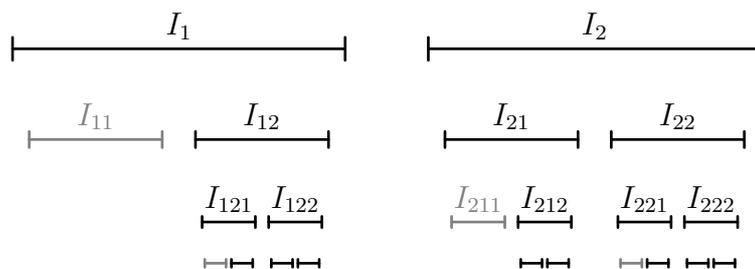


FIGURE 25. A Markov construction.

The support of the Markov measure μ on C is the image $h(\Sigma_A^+)$ of the set of admissible sequences under the coding map, whose construction is shown in Figure 25; the only difference between this and the usual construction of a Cantor-like set is that now we erase any basic interval whose associated sequence is not admissible. So at the second step, I_{11} is erased, since every $x \in I_{11}$ has a coding which begins with $1, 1, \dots$, and hence is not admissible, so $x \notin \text{supp } \mu$. Similarly, I_{211} must be erased at the third step, and I_{1211} and I_{2211} are taken away at the fourth.

The result of this *Markov construction* is a Cantor-like set $C_M \subset C$. As always, we would like to characterise C_M by determining its Hausdorff dimension; the way to do this is to use the Non-uniform Mass Distribution Principle, Theorem 49.

This requires us to exhibit an invariant measure (or measures) on C_M . We constructed C_M as the support of one particular such measure, the Markov measure with probability vector $(p, 1-p)$ and stochastic matrix $\begin{pmatrix} 0 & 1 \\ a & 1-a \end{pmatrix}$. In fact, since the construction of C_M did not depend on the value of a or p , it is the support of *any* such Markov measure, and so we have a two-parameter family of measures $\mu_{p,a}$ for which $\mu_{p,a}(C_M) = 1$.

It can be shown that for such Markov measures, the pointwise dimension $d_{\mu_{p,a}}$ exists almost everywhere and is given by (60) as the ratio of the entropy

to the Lyapunov exponent, where the entropy of a Markov measure is

$$h_\mu(f) = -p(p_{11} \log p_{11} + p_{12} \log p_{12}) - q(p_{21} \log p_{21} + p_{22} \log p_{22}).$$

For the measure $\mu_{p,a}$, the first half of this expression vanishes, and we have

$$h_{\mu_{p,a}}(f) = -(1-p)(a \log a + (1-a) \log(1-a)).$$

Observing that $\lambda(x) = -\log \lambda$ since $f'(x) = \lambda^{-1}$ for all $x \in C_M$, we see that the pointwise dimension is equal $\mu_{p,a}$ -almost everywhere to

$$\phi(p, a) = \frac{(1-p)(a \log a + (1-a) \log(1-a))}{\log \lambda}.$$

As before, Theorem 49 implies that $\dim_H C_M \geq \phi(p, a)$ for all $p, a \in [0, 1]$, and so we want to maximise ϕ in order to find the best bound.

Finding the critical points by setting $\frac{\partial \phi}{\partial p} = \frac{\partial \phi}{\partial a} = 0$, one can show that ϕ achieves its maximum value when $p = 0$, $a = 1/2$, and we obtain

$$\dim_H C_M \geq \phi\left(0, \frac{1}{2}\right) = \frac{\log 2}{\log \lambda}.$$

Observing that $C_M \subset C$ and recalling Moran's theorem, we also see that

$$\dim_H C_M \leq \dim_H C = \frac{\log 2}{\log \lambda},$$

and so we actually have equality, $\dim_H C_M = \log 2 / \log \lambda$. Thus despite the fact that we discarded what seemed like relatively large chunks of C in the construction of C_M , we did not alter the Hausdorff dimension.

Lecture 20

a. FitzHugh-Nagumo and you. And now, as they say, for something completely different. Setting aside our discussion of Cantor-like sets for the moment, we turn our attention to a model from biology, which attempts to describe the propagation of an impulse through a neuron. The (functional) structure of a neuron is shown in Figure 26; the neuron receives signals along its *dendrites*, of which it may have up to several hundred, and after processing these signals in the *soma*, sends a single output along its *axon*.

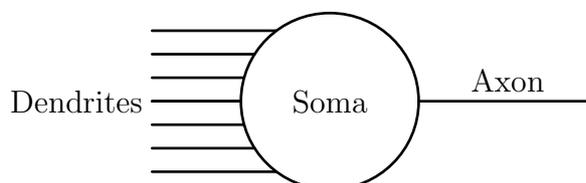


FIGURE 26. A schematic diagram of a neuron.

An axon can be thought of as a cable which transmits an electrical signal via the flow of ions (usually sodium and potassium). We model the axon as lying on the x -axis, and write $u(x, t)$ for the electric potential at point x and time t ; the goal is to obtain an equation which will describe how u changes as time passes.

A naïve idea of using the fundamental laws of physics which govern the interaction and motion of single particles (i.e. molecules) is readily seen to be a rather preposterous line of attack. This is due to the sheer scale of the system, which has many levels of structure lying between the macroscopic description we are interested in and the microscopic level of elementary particles at which the fundamental laws apply. With billions upon billions of ions moving in each axon, it is foolhardy to try to track each individual particle.

Therefore, we should resort to a *phenomenological* approach, which uses a great deal of knowledge about the system and relates various empirical observations to each other in a way which is consistent with the fundamental theory, but is not directly derived from it.

In the attempt to describe the propagation of an impulse along an axon, one of the simplest models which still captures some of the essential qualitative features of the process is the *FitzHugh-Nagumo model*, which was first suggested by Richard FitzHugh in 1961, and was investigated by means of electric circuits the following year by Jin-Ichi Nagumo. It is a simplified version of the Hodgkin-Huxley model, which is designed to describe in a detailed manner the activation and deactivation dynamics of a spiking neuron.

We first give the equations, then explain what they mean.

$$(67) \quad \frac{\partial u_1}{\partial t} = -au_1(u_1 - \theta)(u_1 - 1) - bu_2 + \kappa_1 \frac{\partial^2 u_1}{\partial x^2},$$

$$(68) \quad \frac{\partial u_2}{\partial t} = cu_1 - du_2 + \kappa_2 \frac{\partial^2 u_2}{\partial x^2}.$$

Here $u_1(x, t)$ is the potential $u(x, t)$, and u_2 is a quantity related to the derivative $\partial u_1 / \partial t$. Experimental observation is required to determine the parameters $a, b, c, d, \kappa_1, \kappa_2$.

Where do these equations come from? The behaviour of a neuron is typical for a type of circuit known as a *spike generator*; when an external stimulus reaches the neuron (via one or several of its dendrites) and exceeds a certain threshold, the neuron reacts by generating and transmitting an output signal (via its axon) before it relaxes back to its rest state. The equations that govern such a process are known in physics as the *Bonhoeffer-van der Pol model*:

$$(69) \quad \dot{v} = v - v^3 - w + I,$$

$$(70) \quad \dot{w} = av - b - cw.$$

This system is based on the classical *van der Pol oscillator*, named after the Dutch physicist Balthasar van der Pol, which describes a non-conservative oscillator with non-linear damping:

$$(71) \quad \ddot{v} + \mu(v^2 - 1)\dot{v} + v = 0.$$

Observe that (69) and (70) are together equivalent to (71) in the particular case $b = c = 0$.

To obtain the equations (67), (68) for the FitzHugh-Nagumo model, one takes (69), (70) and adds the terms $\kappa_i \frac{\partial^2 u_i}{\partial x^2}$, which account for the diffusion of ions during propagation of the impulse. Since the effect of diffusion is very small, the numbers κ_1 and κ_2 are of higher order than the other parameters.

Equations (67) and (68) are still too complicated to be solved analytically; however, their study can be carried out numerically, with the help of computer simulations. Once this is done, we are left with a whole slew of numerical predictions which follow from the model, and which must be compared somehow to observations of real, honest-to-goodness biological neurons. It is in this process that fractal geometry will once again have a role to play.

Lecture 21

a. Numerical investigations. In the previous lecture, we introduced the FitzHugh–Nagumo equations (67)–(68) for modelling the propagation of an impulse along an axon. The terms in these equations can be split into two groups; the main terms, which are polynomial expressions in u_1 and u_2 , and the diffusion terms, which depend on the second derivatives of u_1 and u_2 with respect to x . The latter terms are small relative to the others, and so as a first step to understanding the behaviour of the model, we drop the diffusion terms and deal with the simplified equations

$$(72) \quad \frac{\partial u_1}{\partial t} = -au_1(u_1 - \theta)(u_1 - 1) - bu_2,$$

$$(73) \quad \frac{\partial u_2}{\partial t} = cu_1 - du_2.$$

Given initial conditions (u_1^0, u_2^0) , it follows from the existence and uniqueness theorem for ODEs that there exists a unique solution of (72)–(73) beginning at (u_1^0, u_2^0) . However, these equations are analytically intractable, and we cannot generally write down that solution in closed form. Thus we must resort to a numerical analysis, and rely on a computer to compute the trajectory of u_1 and u_2 .

But what will the computer do? It will approximate the solution by discretising the problem: fixing some small value of $h > 0$, the time derivative $\partial u_i / \partial t$ is approximately equal to

$$\frac{u_i(t+h) - u_i(t)}{h},$$

and so (72)–(73) are approximated by the equations

$$u_1(t+h) = u_1(t) - ah u_1(t)(u_1(t) - \theta)(u_1(t) - 1) - bh u_2(t),$$

$$u_2(t+h) = u_2(t)(1 - dh) + ch u_1(t).$$

This gives an iterative procedure to find u_1 and u_2 ; writing $u_i^n = u_i(nh)$, we have

$$(74) \quad u_1^{n+1} = u_1^n - Au_1^n(u_1^n - \theta)(u_1^n - 1) - \alpha u_2^n,$$

$$(75) \quad u_2^{n+1} = \beta u_1^n + \gamma u_2^n,$$

where $\alpha = bh$, $\beta = ch$, $\gamma = 1 - dh$, and $A = ah$. For our purposes we shall assume that α and β are small, θ is near $1/2$, and γ is near 1 . A is referred to as the *leading parameter*, and a question of interest to us is how the behaviour of the trajectories of (74)–(75) changes as A varies.

First, though, we observe that the result of iterating the discrete system (74)–(75) is a sequence of points $(u_1^n, u_2^n) \in \mathbb{R}^2$; the idea is that these points lie near the curve $(u_1(t), u_2(t))$ which defines the solution of (72)–(73), but there will of course be some small error term. For many of the parameter values in which we are interested, this error term increases relatively quickly, so that the asymptotic behaviour of solutions of the two

systems may be quite different. In particular, the behaviour of the discrete system (which the computer calculates) is not necessarily a reliable guide to the behaviour of the continuous system.

One response to this difficulty is to use a more sophisticated numerical method—the approach described is known as *Euler's method*, and more advanced methods are possible, which yield better approximations. However, these are still approximations, and over time, some accumulation of the error term is unavoidable. Thus it is not clear just how one is to examine the asymptotic behaviour of even the simplified FitzHugh–Nagumo model (72)–(73).

At this point we recall that what we are really interested in studying is the physical phenomenon itself; thus the important question is not how well the discrete model approximates the continuous model, but how well it approximates the physical propagation of an impulse through a neuron. It is possible to add the discrete analogue of the diffusion terms $\kappa_i \partial^2 u_i / \partial x^2$ to the discrete model (74)–(75), and then compare the trajectories of the full discrete system to the observed data. If this turns out to approximate the real system as well as or better than the continuous model, then we may simply consider the discrete model, rather than the continuous one.

The biological literature is divided over whether or not this is the case, with both the discrete and the continuous FitzHugh–Nagumo models having their advocates. For the time being, we will leave questions of biological applicability in the background, and focus our attention on the mathematical characteristics of the simplified discrete model (74)–(75), which occurs in other situations as well—various problems in physics, chemistry, etc., lead to equations of this form, and so it is worth understanding the general behaviour of this system.

b. Studying the discrete model. The sequence of points (u_1^n, u_2^n) produced by the simplified discrete model (72)–(73) is exactly the trajectory of the point (u_1^0, u_2^0) under the action of the map $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $f(u_1, u_2) = (f_1(u_1, u_2), f_2(u_1, u_2))$, where

$$(76) \quad \begin{aligned} f_1(u_1, u_2) &= u_1 - Au_1(u_1 - \theta)(u_1 - 1) - \alpha u_2, \\ f_2(u_1, u_2) &= \beta u_1 + \gamma u_2. \end{aligned}$$

We examine the dynamics of f for different values of A , with the parameters $\alpha, \beta, \gamma, \theta$ fixed; thus we are actually studying a *family* of maps, one for each value of A .

The first step in our analysis of the map (76) is to find the fixed points of f by solving the equation $f(u_1, u_2) = (u_1, u_2)$, which leads to the following equations for the two coordinates:

$$\begin{aligned} u_1 &= u_1 - Au_1(u_1 - \theta)(u_1 - 1) - \alpha u_2, \\ u_2 &= \beta u_1 + \gamma u_2. \end{aligned}$$

From the second of these, we obtain

$$(77) \quad u_2 = \frac{\beta u_1}{1 - \gamma},$$

and then the first becomes

$$A u_1 (u_1 - \theta)(u_1 - 1) + \frac{\alpha \beta}{1 - \gamma} u_1 = 0.$$

Thus we have a fixed point at the origin, where $u_1 = u_2 = 0$, and any other fixed point must satisfy

$$A(u_1 - \theta)(u_1 - 1) + \frac{\alpha \beta}{1 - \gamma} = 0.$$

Solving this quadratic equation yields

$$(78) \quad u_1 = \frac{1}{2} \left(\theta + 1 \pm \sqrt{(\theta - 1)^2 - \frac{4\alpha\beta}{A(1 - \gamma)}} \right).$$

The discriminant is non-negative if and only if

$$A \geq A_0 = \frac{4\alpha\beta}{(1 - \gamma)(1 - \theta)^2},$$

and so we see that for $0 < A < A_0$, the origin is the only fixed point of f . For $A = A_0$, there is exactly one more fixed point, and for $A > A_0$, there are two more, given by (78) and (77); we will denote these by p_1 and p_2 .

Thus A_0 marks the boundary between two qualitatively different sorts of behaviour. Imagine a tuning knob which controls the parameter A ; if we begin with the knob turned so that $0 < A < A_0$, then the system has only one fixed point, and this general structure persists for a little while as we turn the knob and increase A . However, when we turn the knob far enough that A reaches A_0 , the system undergoes a *bifurcation*, and two new fixed points appear.

In the next lecture, we will investigate the stability of the fixed points 0 , p_1 , and p_2 , which tells us how orbits near these fixed points behave.

Lecture 22

a. Stability of fixed points. In the previous lecture, we began examining the map $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by (76), paying particular attention to how the behaviour of the dynamics of f changes as the parameter A varies. We found a parameter value A_0 such that for $0 < A < A_0$, the system has only one fixed point, which is at the origin, while for $A > A_0$, there are two additional fixed points p_1 and p_2 , which are given by (78) and (77).

The next step in the analysis of f is to determine the stability of these fixed points; we want to know whether they attract or repel nearby trajectories, and how that behaviour depends on the value of the parameter A . The tool which we use for this purpose is the Jacobian derivative

$$(79) \quad Df = \begin{pmatrix} \frac{\partial f_1}{\partial u_1} & \frac{\partial f_1}{\partial u_2} \\ \frac{\partial f_2}{\partial u_1} & \frac{\partial f_2}{\partial u_2} \end{pmatrix} = \begin{pmatrix} 1 - A\theta + 2A(1 + \theta)u_1 - 3Au_1^2 & -\alpha \\ \beta & \gamma \end{pmatrix},$$

which describes the behaviour of f in a neighbourhood of a fixed point. Specifically, we recall the following result from calculus:

PROPOSITION 53. *Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be continuously differentiable in a neighbourhood of p_0 . Then we have*

$$(80) \quad f(p) = f(p_0) + Df(p_0)(p - p_0) + o(p - p_0).$$

Heuristically, this says that in a small neighbourhood of p_0 , f behaves like a linear map, with an error term that is small relative to $p - p_0$. In particular, if p_0 is a fixed point, then its stability is determined by the eigenvalues of $Df(p_0)$; this is made precise by the Hartman–Grobman theorem, which states that as long as none of the eigenvalues lie on the unit circle (that is, $|\lambda| \neq 1$ for all eigenvalues λ of $Df(p_0)$), then the restriction of f to a neighbourhood of p_0 is topologically conjugate to the linear map $Df(p_0)$.

Rather than proving the full Hartman–Grobman theorem here, we content ourselves with a particular case.

PROPOSITION 54. *Suppose that $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ has a fixed point at p_0 , and that $\alpha < 1$ is such that $|\lambda| < \alpha$ for all eigenvalues λ of $Df(p_0)$. Then there exists $\varepsilon > 0$ such that for all $p \in \mathbb{R}^2$ with $\|p - p_0\| < \varepsilon$, we have $\lim_{k \rightarrow \infty} f^k(p) = p_0$.*

PROOF. Using (80) and the fact that p_0 is fixed, we have

$$f(p) - p_0 = Df(p_0)(p - p_0) + r(p),$$

where $r(p)$ is the error term, which goes to zero more quickly than $p - p_0$ does. Because all the eigenvalues have absolute value less than α , there exists a norm $\|\cdot\|$ on \mathbb{R}^2 such that

$$\|Df(p_0)\| = \sup_{\substack{x \in \mathbb{R}^2 \\ \|x\|=1}} \|Df(p_0)x\| < \alpha,$$

and since

$$\lim_{p \rightarrow p_0} \frac{\|r(p)\|}{\|p - p_0\|} = 0,$$

we can find $\varepsilon > 0$ such that $\|Df(p_0)\| + \|r(p)\|/\|p - p_0\| < \alpha$ for all p with $\|p - p_0\| < \varepsilon$. In particular, we have

$$\begin{aligned} \|f(p) - p_0\| &\leq \|Df(p_0)(p - p_0)\| + \|r(p)\| \\ &\leq \left(\|Df(p_0)\| + \frac{\|r(p)\|}{\|p - p_0\|} \right) (\|p - p_0\|) \\ &< \alpha \|p - p_0\|, \end{aligned}$$

whence $\|f^k(p) - p_0\| < \alpha^k \|p - p_0\|$ by induction. \square

The case where one eigenvalue lies inside the unit circle and one lies outside is harder to deal with, because one must first prove the existence of *stable and unstable manifolds*. These are curves $\gamma^s(t)$ and $\gamma^u(t)$ through p_0 such that:

- (1) γ^s and γ^u are f -invariant.
- (2) The tangent vector to γ^s at p_0 points in the direction of the eigenvector corresponding to the eigenvalue inside the unit circle, and similarly, the tangent to γ^u corresponds to the other eigenvector.
- (3) Trajectories on γ^s approach p_0 , while trajectories on γ^u are repelled from p_0 .

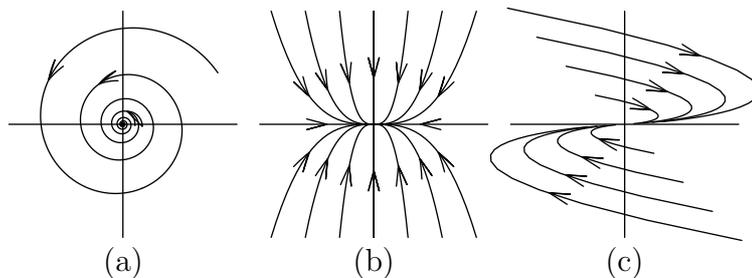


FIGURE 27. Trajectories near an attracting fixed point.

We briefly recall the possible qualitative behaviours for a linear map $A: \mathbb{R}^2 \rightarrow \mathbb{R}^2$; let λ and μ be the eigenvalues of A , with $|\lambda| \leq |\mu|$. Assuming neither λ or μ lies on the unit circle, there are three possibilities for the fixed point 0; when $A = Df(p_0)$ for a fixed point p_0 , these corresponds to three different possibilities for the stability of p_0 .

- (1) $|\lambda| \leq |\mu| < 1$: All trajectories of A converge to 0, and so p_0 is an attracting fixed point f , also called a *node*. The manner in which trajectories converge depends on λ and μ . If λ and μ are complex, then trajectories move along a logarithmic spiral, as shown in Figure 27(a); if λ and μ are real, then trajectories move along

the curves shown in Figure 27(b) (if $Df(p_0)$ is diagonalisable) or in Figure 27(c) (if $Df(p_0)$ is similar to $\begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$).

- (2) $|\lambda| < 1 < |\mu|$: p_0 is a *hyperbolic* fixed point, also called a *saddle*, shown in Figure 28. From one direction (the stable direction, horizontal in the figure), which corresponds to the eigenline for λ , trajectories approach p_0 as $n \rightarrow +\infty$, while from another direction (the unstable direction, vertical in the figure), corresponding to the eigenline for μ , the *backwards* trajectories approach p_0 as $n \rightarrow -\infty$. All other trajectories follow hyperbola-like paths, at first moving closer to p_0 , and then moving away.
- (3) $1 < |\lambda| \leq |\mu|$: All trajectories of A move away from 0, and so p_0 is a repelling fixed point. Trajectories move along one of the curves in Figure 27, but in the opposite direction.

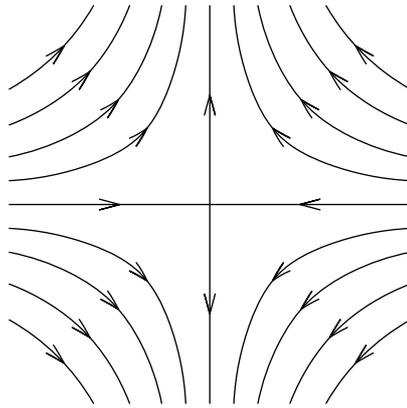


FIGURE 28. Trajectories near a hyperbolic fixed point.

This approach allows us to classify the fixed points of the FitzHugh–Nagumo map f for various values of the parameter A . However, finding the eigenvalues of the matrix Df given by (79) involves a rather complicated computation, which we would prefer to avoid.

Recall that the eigenvalues of a matrix are the roots of its characteristic polynomial, and so they depend continuously on the coefficients of that polynomial. Those coefficients in turn depend continuously on the entries of the matrix (being sums of products of those entries), and so λ and μ depend continuously on the entries of Df .

In particular, given $\varepsilon > 0$, there exists δ such that for $\alpha, \beta < \delta$, the eigenvalues of Df are within ε of the eigenvalues of the diagonal matrix

$$(81) \quad T = T(u_1, u_2) = \begin{pmatrix} 1 - A\theta + 2A(1 + \theta)u_1 - 3Au_1^2 & 0 \\ 0 & \gamma \end{pmatrix}.$$

Since T is diagonal, we can read off its eigenvalues directly. One eigenvalue is γ , and for the parameters we consider, $\gamma < 1$; thus there is always at

least one stable (contracting) direction. The other eigenvalue depends on parameters and the fixed point, and so we consider these points separately.

For all values of A , the map f has a fixed point at the origin, and so the second eigenvalue of T is $1 - A\theta$. We have $|1 - A\theta| < 1$ if and only if $0 < A < 2/\theta$, and since λ and μ are close to $1 - A\theta$ and γ , there exists A_1 near $2/\theta$ such that the behaviour of f near the origin is as follows:

- (1) For $0 < A < A_1$, both eigenvalues of $Df(0,0)$ lie inside the unit circle, hence the origin is an attracting fixed point.
- (2) For $A > A_1$, one eigenvalue of $Df(0,0)$ lies inside the unit circle (near γ), and the other lies outside the unit circle (near $1 - A\theta$), and so the origin is a hyperbolic fixed point for f .

One can show that $A_1 > A_0$, and hence by the time A reaches A_1 , there are two more fixed points to keep track of, p_1 and p_2 .

To determine the behaviour of f near the fixed points p_1 and p_2 , we first need to determine their location. Notice that the fixed points of f all lie on the line with equation (77); indeed, this line consists of precisely those points (u_1, u_2) for which $f_2(u_1, u_2) = u_2$. Similarly, the fixed points all lie on the cubic polynomial with equation

$$u_2 = -\frac{A}{\alpha}u_1(u_1 - \theta)(u_1 - 1);$$

this curve contains precisely those points for which $f_1(u_1, u_2) = u_1$. Thus the fixed points of f are the points where this curve intersects the line (77), as shown in Figure 29.

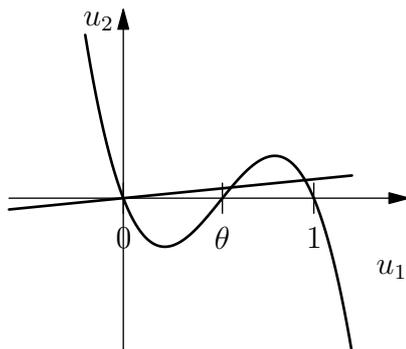


FIGURE 29. Finding the fixed points of f .

Because β is small, the slope of this line is nearly horizontal, and so these points of intersection are very near the points where the cubic intersects the x -axis, which occurs at $u_1 = 0$, $u_1 = \theta$, and $u_1 = 1$. Thus p_1 and p_2 are approximately $(\theta, 0)$ and $(1, 0)$, respectively; we could also have seen this by looking at the form (78) takes when α and β vanish.

Now we can use the approximation (81) for Df to estimate the eigenvalues of $Df(p_1)$ and $Df(p_2)$. At $(u_1, u_2) = (\theta, 0)$, the matrix T in (81) has

$\lambda_1 = \gamma$ as one eigenvalue (as always), and the other eigenvalue is

$$\lambda_2 = 1 - A\theta + 2A(1 + \theta)\theta - 3A\theta^2 = 1 + A\theta - A\theta^2.$$

Because $0 < \theta < 1$, we have $\lambda_2 = 1 + A\theta(1 - \theta) > 1$, and since $\lambda_1 = \gamma < 1$, we see that p_1 is a hyperbolic fixed point for every value of the parameter A , with one stable and one unstable direction.

The situation is different at p_2 , where we have the estimate

$$\lambda_2 = 1 - A\theta + 2A(1 + \theta) - 3A = 1 + A\theta - A,$$

and so $|\lambda_2| < 1$ if and only if $0 < A < 2/(1 - \theta)$. In particular, there is a critical value $A'_1 \approx 2/(1 - \theta)$ such that for $A_0 < A < A'_1$, the fixed point p_2 is attracting, while for $A > A'_1$, it is a saddle.

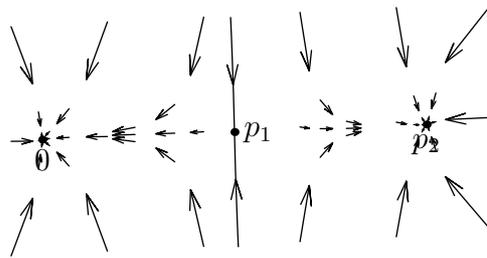


FIGURE 30. Orbits of f for $A_0 < A < \min\{A_1, A'_1\}$.

b. Things that don't stand still, but do at least come back. The ultimate goal of all the analysis in which we are presently embroiled is to understand the dynamics of f by classifying the possible trajectories, and to describe how the dynamics changes as A varies. By knowing the fixed points of f and their stability, we gain information about the *local* behaviour of the system; that is, how trajectories behave near the fixed points.

In general, it may be rather difficult to use this information to piece together a *global* picture, which describes how trajectories behave *everywhere*. For the time being, we omit the details of this particular jigsaw puzzle, and simply assert that for the map f in the parameter ranges we consider, there are two sorts of trajectories: unbounded trajectories, which diverge to ∞ , and bounded trajectories, which converge to a fixed point (or possibly, as we will soon see, to a periodic orbit).

Thus for parameter values $A_0 < A < \min\{A_1, A'_1\}$, we have three fixed points; p_1 is a saddle, while 0 and p_2 are stable. The bounded trajectories of f are as shown in Figure 30, and are of three sorts. There is a curve through p_1 , called the *stable manifold* (or *stable separatrix*) of p_1 , which comprises all points in \mathbb{R}^2 whose trajectory goes to p_1 as $n \rightarrow \infty$. If a bounded trajectory $\{f^n(x)\}$ begins to the left of this curve, then $f^n(x) \rightarrow 0$; if x lies to the right of the curve, then $f^n(x) \rightarrow p_2$.

Note that whether $A_1 < A'_1$ or $A'_1 < A_1$ depends on the values of the parameters θ , α , β , and γ ; for the sake of concreteness, we will assume that $A_1 < A'_1$, so that 0 changes behaviour before p_2 does.

In fact, A_1 is not the first time that the qualitative behaviour of trajectories near 0 changes. For some value $A = \tilde{A}_1 < A_1$, the eigenvalue $\lambda_2 \approx 1 - A\theta$ is equal to 0. This eigenvalue governs the behaviour of points which lie just to the left or right of 0; when it is positive, for $A < \tilde{A}_1$, points to the left remain on the left, and points to the right remain on the right. For $A > \tilde{A}_1$, on the other hand, we have $\lambda_2 < 0$, and so points lying just to the left of 0 are mapped to points lying just to the right, and vice versa; in some sense, the orientation of the fixed point reverses at \tilde{A}_1 .

Now as A approaches A_1 from below, the eigenvalue λ_2 approaches -1 from above. As long as the eigenvalue is greater than -1 , the image $f(x)$ is closer to 0 than x itself is, and so $f^n(x) \rightarrow 0$; however, the rate of convergence becomes slower and slower as λ_2 approaches -1 . Finally, when $\lambda_2 < -1$, nearby points (in the horizontal direction) are mapped further away by f , and the fixed point 0 is now a saddle, with one stable and one unstable direction.

This change in behaviour means that the current picture is incomplete. Some trajectories leave 0 in a more or less horizontal direction; where do they go? They cannot immediately go towards the next fixed point, p_1 , since it is also a saddle which is repelling in the horizontal direction.

What is missing? We have found all the fixed points, which represent the simplest possible orbits. The next simplest type of orbit is a periodic orbit, for which the trajectory returns to the initial point after some finite number of iterations.

In the present case, such an orbit appears around the fixed point 0 when the stability changes at A_1 ; for $A > A_1$, there exist two points q_1 and q_2 (which depend on A) such that $f(q_1) = q_2$ and $f(q_2) = q_1$. This orbit is the missing piece of the puzzle; trajectories which are repelled from 0 are attracted to the nearby periodic orbit of period 2.

This behaviour is shown in Figure ??, which for simplicity of representation actually shows the action of f^2 , the second iterate of f . Both q_1 and q_2 are mapped to themselves by f^2 , which is a manifestation of the general fact that periodic points of f with period n correspond to fixed points of f^n . The stability of the periodic orbit for f is given by the stability of the fixed point for f^n , which in this case is determined by the eigenvalues of $Df^2(q_i)$. One may easily see that it does not matter whether we compute the Jacobian at q_1 or at q_2 , since by the chain rule,

$$\begin{aligned} Df^2(q_1) &= Df(f(q_1))Df(q_1) \\ &= Df(q_2)Df(q_1) \\ &= (Df(q_1))^{-1}Df(q_1)Df(q_2)Df(q_1) \\ &= Df(q_1)^{-1}Df^2(q_2)Df(q_1); \end{aligned}$$

it follows that $Df^2(q_1)$ and $Df^2(q_2)$ are similar matrices, and hence have the same eigenvalues.

In order to compute, or even estimate, the eigenvalues of $Df^2(q_i)$, we would first need to find the points q_1, q_2 , which involves solving the equation $f^2(x, y) = (x, y)$. Even with the approximation $\alpha = \beta = 0$, this leads to a polynomial of degree nine, which is algebraically intractable.

Numerical evidence (for the two-dimensional system), along with some geometric reasoning (which we give later for its one-dimensional approximation), indicates that the eigenvalue of $Df^2(q_i)$ corresponding to the horizontal direction is slightly smaller than 1 when A is just larger than A_1 . Thus when it appears, the period-two orbit is attracting; as A increases, the eigenvalue decreases, and the orbit becomes more strongly attracting. When A increases to the point \tilde{A}_2 where the eigenvalue passes 0 and becomes negative, the orientation of trajectories near the period-two orbit changes, just as happened at \tilde{A}_1 for the fixed point 0.

Beyond \tilde{A}_2 , the rate at which trajectories converge to the period-two orbit decreases, until at some value $A = A_2$, one of the eigenvalues of $Df^2(q_i)$ passes -1 ; at this point, the orbit ceases to be attracting, and becomes hyperbolic. The map f^2 , for which q_1 and q_2 are fixed points, behaves just like f did, and spawns an attracting period-two orbit near each of the newly unstable fixed points. Together, these two orbits form a period-four orbit for the original map f , which has two points near each of q_1 and q_2 .

This behaviour continues; the period-four orbit eventually becomes unstable, at which point a period-eight orbit is born, and so on. In fact, there exists a sequence of parameter values $A_1 < A_2 < A_3 < \dots$ such that for $A_{n-1} < A < A_n$, f has a stable orbit of period 2^n in the region to the left of the stable manifold of p_1 , and unstable orbits of period 2^k for all $0 \leq k \leq n$.

Meanwhile, the same story is unfolding to the right of the stable manifold; at A'_1 the fixed point p_2 splits into an attracting period-two orbit, and we have a sequence $A'_1 < A'_2 < A'_3 < \dots$ such that for $A'_{n-1} < A < A'_n$, f has a stable orbit of period 2^n in the region to the *right* of the stable manifold of p_1 , and unstable orbits of period 2^k for all $0 \leq k \leq n$.

In the next lecture, we will introduce a graphical way of representing all this information, the *bifurcation diagram*. For the time being, we note that the bifurcation values A_n (and A'_n as well) get closer and closer together as n grows, and in fact, converge to some value A_∞ , at which we have orbits of period 2^n for *any* natural number n . In particular, there are infinitely many co-existing periodic orbits, which marks a fundamental change in the behaviour of f .

DEFINITION 55. A dynamical system $f: X \rightarrow X$ is called *Morse–Smale* if it has finitely many periodic orbits, each of which is either attracting, repelling, or a saddle, such that given any initial condition x , the trajectory $\{f^n(x)\}$ approaches one of these orbits.

For $0 < A < A_\infty$, the simplified discrete FitzHugh–Nagumo system (76) is Morse–Smale, and so the behaviour of the system is in some sense relatively simple and easy to understand. For $A \geq A_\infty$, the behaviour is much more intricate, as we will see next time.

Lecture 23

a. Down the rabbit hole. In the previous lecture, we described the sequence of period-doubling bifurcations through which the simplified discrete FitzHugh–Nagumo model (76) passes as A increases. In fact, we saw that there were two such sequences, one corresponding to the fixed point 0, for which the resulting periodic orbits lay to the left of the stable manifold of p_1 , and one corresponding to p_2 , for which the periodic orbits lay to the right of that curve.

In order to visualise these concurrent period-doubling cascades, and also to get a sense of what lies beyond them, for $A > A_\infty$, we turn to the *bifurcation diagram* of f , also sometimes called the *orbit diagram*, which is shown in Figure 31 for a particular choice of parameters $\alpha, \beta, \gamma, \theta$.

The horizontal axis of the diagram represents the parameter A , and the vertical axis represents the first coordinate u_1 of the phase space \mathbb{R}^2 ; the points on a vertical slice of the diagram are the u_1 -coordinates of the attracting part of the phase space. The diagram is generated as follows: fix a value of A , an initial condition (u_1, u_2) , which is relatively arbitrary (but should lie near enough the origin so that its orbit is bounded and to the left of the stable manifold of p_1), and then compute the iterates $f^n(u_1, u_2)$. Ignore the first few iterations,¹⁴ to give the transient part of the orbit time to die away; in the case $A < A_\infty$, for example, this gives the orbit time to converge to a periodic orbit. Then writing the subsequent iterations as $(u_1^{(n)}, u_2^{(n)})$, plot the points $(A, u_1^{(n)})$ on the diagram.¹⁵ One then chooses a new value of A and repeats the whole procedure, until enough points have been filled in to give a sense of the structure of the bifurcation diagram.

Finally, the above procedure is repeated in its entirety, with initial conditions this time chosen to the right of the stable manifold of p_1 ; this generates the top half of Figure 31.

We see that for $A < A_1$, the points which are plotted are all quite close to the fixed point; for $A_1 < A < A_2$, they are close to the period-two orbit, and so on. This reflects the fact that the long-term behaviour of any trajectory follows one of the periodic orbits, as is true of any Morse–Smale system.

For $A \geq A_\infty$, the situation is quite different. There are values of A for which the trajectory of a randomly chosen point fills out a Cantor set, or even an entire interval, and the situation is quite chaotic. However, there are also windows of stability, the most noticeable of which are the two period-three windows which occur near $A = 5.7$ and $A = 6$ in Figure 31. In these parameter ranges, we suddenly return from chaos to order; every orbit is attracted to a cycle of period 3, and the system seems relatively simple once again (although as we will see later, it is not Morse–Smale). There

¹⁴“First few” may mean several dozen, several hundred, or several thousand, depending on how refined a picture is desired.

¹⁵Again, how many of these points are plotted is in some sense a judgment call.

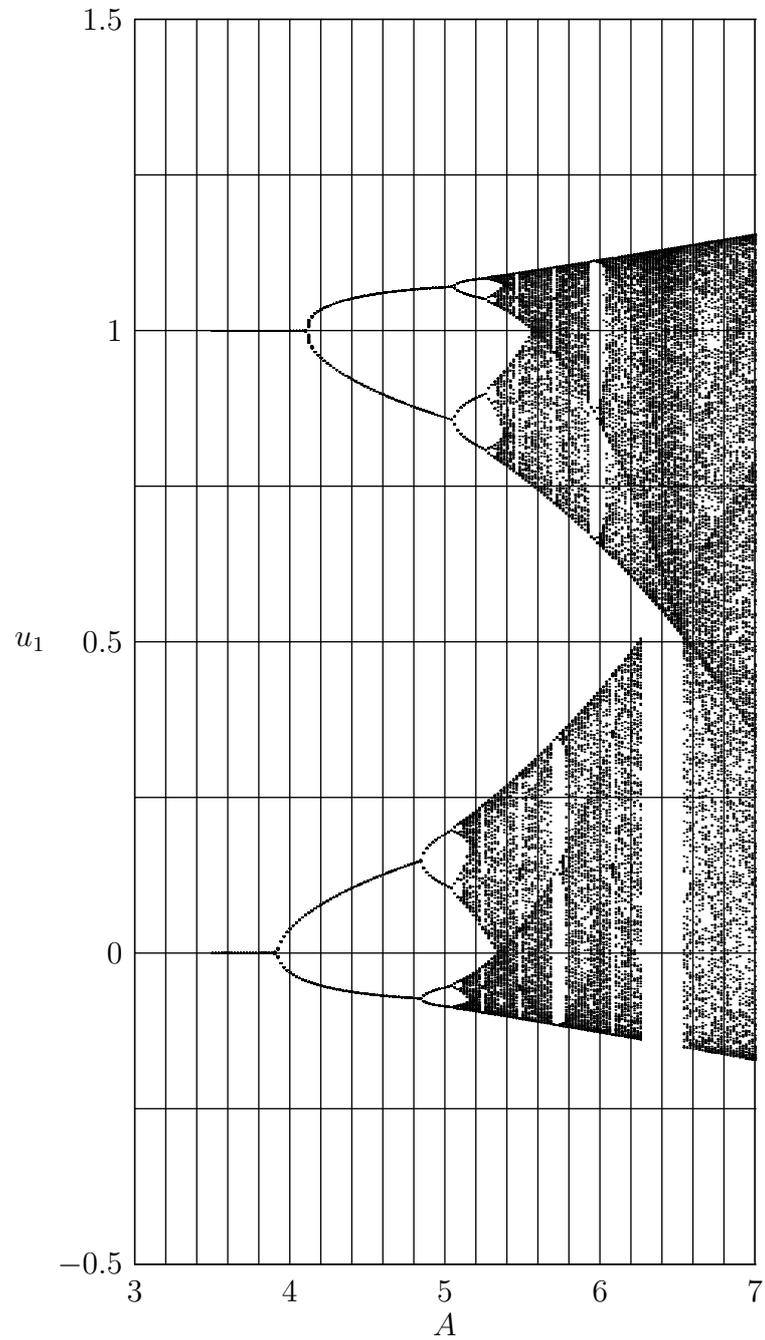


FIGURE 31. The bifurcation diagram for the discrete FitzHugh-Nagumo model (76) with $\theta = .51$, $\alpha = .01$, $\beta = .02$, and $\gamma = .8$, as A varies from 3 to 7.

are windows of stability with other orders as well, scattered throughout the bifurcation diagram.

These moments of stability are transitory. As A increases, the attracting periodic cycle of length n becomes repelling, and spawns an attracting periodic cycle of length $2n$. This too becomes repelling in its turn, shedding an attracting cycle of length $4n$, then $8n$, and so on; the whole period-doubling cascade is repeated, but over a much smaller range of the parameter A than in its original incarnation.

This reappearance of the period-doubling cascade suggests a sort of self-similarity of the bifurcation diagram, and indeed this diagram is self-similar in an asymptotic sense. We will not make this statement precise here, beyond mentioning one striking manifestation of this self-similarity. Recall that $A_n - A_{n-1}$ gives the length of the parameter interval for which the periodic orbit of length 2^n is attracting; it turns out that this length decreases exponentially in n , and so we write

$$\lambda = \lim_{n \rightarrow \infty} \frac{1}{n} \log(A_n - A_{n-1}).$$

Thus writing $\delta = e^{-\lambda}$, we have $A_n - A_{n-1} \approx C\delta^n$ for some constant C . In and of itself, this would not be cause for any special excitement. However, we can do the same calculation for the period-doubling cascade $A'_1 < A'_2 < A'_3 < \dots$, or for the period-doubling cascade in any of the windows of stability, and the truly striking result is that we get the same value of δ in every case!

Even that remarkable congruence is not the end of the story; there are in fact many one-parameter families of maps which lead to similar bifurcation diagrams, with period-doubling cascades, windows of stability, and so on. Under relatively mild conditions on the family of maps, we find the same old tale in each period-doubling cascade; the rate of decay of the lengths $A_n - A_{n-1}$ exists, and what is more, is equal to δ .

The number δ is known as *Feigenbaum's constant*, after Mitchell Feigenbaum, who was the first to discover this example of *quantitative universality*.

b. Becoming one-dimensional. Despite the various heuristic justifications which have been offered, our analysis of the FitzHugh-Nagumo model, and in particular the bifurcation diagram in Figure 31, is based primarily on numerical evidence. Many of our claims have been based on the fact that f is a small perturbation of the map

$$\tilde{f}(u_1, u_2) = (u_1 - Au_1(u_1 - \theta)(u_1 - 1), \gamma u_2),$$

and that since the u_2 -coordinate goes to 0 under repeated iteration of this map, the essential behaviour is (or ought to be) given by the one-dimensional map

$$g(x) = x - Ax(x - \theta)(x - 1).$$

A great deal of what we have said up until now has been proved for such one-dimensional maps; however, very little of that theory has been rigorously extended to the two-dimensional case. While most of the results are believed to carry over, that belief is based on numerical evidence and computer experiments rather than rigorous proofs.

Because the one-dimensional case is easier to deal with, we restrict our attention for the next little while to that case, and consider continuous maps $f: \mathbb{R} \rightarrow \mathbb{R}$. This setting allows us to prove a number of results regarding the orbit structure of the map f ; one of the most important of these is the following theorem, due to the Ukrainian mathematician Aleksandr Sarkovskii, which places surprising restrictions on which combinations of periodic orbits can exist in a given system.

THEOREM 56 (Sarkovskii's theorem). *Place a non-standard ordering on the set of positive integers \mathbb{N} as follows: given two integers m and n , we say that m precedes n (or equivalently, that n follows m), denoted $m \prec n$, if m appears before n in the following list:*

$$\begin{aligned} &3, 5, 7, 9, \dots \\ &2 \cdot 3, 2 \cdot 5, 2 \cdot 7, 2 \cdot 9, \dots \\ &2^2 \cdot 3, 2^2 \cdot 5, 2^2 \cdot 7, 2^2 \cdot 9, \dots \\ &2^3 \cdot 3, 2^3 \cdot 5, 2^3 \cdot 7, 2^3 \cdot 9, \dots \\ &\dots \\ &\dots, 2^n, 2^{n-1}, \dots, 2^2, 2, 1. \end{aligned}$$

Then if $f: \mathbb{R} \rightarrow \mathbb{R}$ has a periodic orbit of period m , it also has a periodic orbit of period n for every n which follows m , $m \prec n$.

The key ingredient in the proof of Theorem 56 (which we omit at present) is the Intermediate Value Theorem, which explains why it is crucial that the domain of f be \mathbb{R} . Indeed, if $X = \{z \in \mathbb{C} \mid |z| = 1\}$ is the unit circle, and $f: X \rightarrow X$ is the map $f(z) = e^{2\pi i/n}z$, then every periodic point of f has period n , in sharp contrast to the situation described in Sarkovskii's theorem; this example also shows that the theorem fails in higher dimensions.

In \mathbb{R} , however, the theorem gives a great deal of information about the periodic orbit structure of a continuous map f . For example, if f has a periodic point of period 2^n , then it must have periodic points of period 2^k for every $0 \leq k \leq n$, which is reminiscent of the behaviour we saw in the period-doubling cascade earlier. Similarly, if f has a periodic point whose period is not a power of two, then it must have periodic points of period 2^n for every n .

The following corollary of Theorem 56 is particularly important for historical reasons:

COROLLARY 57. *If a continuous map $f: \mathbb{R} \rightarrow \mathbb{R}$ has a period-three orbit, then it has periodic points of all orders.*

Sarkovskii proved Theorem 56 in 1964; however, as a Ukrainian mathematician behind the Iron Curtain, he had little access to the West, and so his result was not widely publicised until 1975, when Tien-Yien Li and James Yorke published a paper entitled “Period Three Implies Chaos”, in which they proved Corollary 57, which they had discovered independently of Sarkovskii’s work.

Theorem 56 also has the following corollary, which applies to any Morse–Smale system on the line:

COROLLARY 58. *If a continuous map $f: \mathbb{R} \rightarrow \mathbb{R}$ has only finitely many periodic points, then they all have a period which is a power of two.*

Sarkovskii’s theorem leads us to suspect that even within the windows of stability in the bifurcation diagram in Figure 31, there are some very complicated dynamics going on. In the next lecture, we will turn our attention to the family of logistic maps, which will yield a similar bifurcation diagram, complete with period-doubling cascades, windows of stability, and in the end, chaos.

Lecture 24

a. Bifurcations for the logistic map. We return now to the logistic map, or rather the family of logistic maps, which was first introduced in Lecture 2; for a given parameter $c \in \mathbb{R}$, the map is

$$f_c(x) = x^2 + c.$$

We have already seen that for $c > 1/4$, all trajectories of f_c go to $+\infty$; we focus on what happens as c decreases. In particular, we make precise the notion of bifurcation, which we have already discussed, and examine the types of bifurcation which occur in the logistic family.

Recall that two continuous maps $f: X \rightarrow X$ and $g: Y \rightarrow Y$ are called *topologically conjugate* if there exists a homeomorphism $\phi: Y \rightarrow X$ such that $f \circ \phi = \phi \circ g$; that is, the following diagram commutes:

$$\begin{array}{ccc} Y & \xrightarrow{g} & Y \\ \downarrow \phi & & \downarrow \phi \\ X & \xrightarrow{f} & X \end{array}$$

For example, the maps $f_c: x \mapsto x^2 + c$ and $g_\lambda: y \mapsto \lambda y(1-y)$ are topologically conjugated by the homeomorphism $\phi: y \mapsto \frac{\lambda}{2}(1-2y)$, where c and λ are related by $4c = \lambda(2-\lambda)$ in the appropriate parameter ranges.

It is often useful to think of the conjugating homeomorphism ϕ as a change of coordinates, under which f and g display the same dynamics, just as two similar matrices A and B have the same action under a suitable change of basis. We will often refer to two topologically conjugate maps as having the same qualitative behaviour.

EXERCISE 14. Show that f_c and $f_{c'}$ are topologically conjugate for any $c, c' > 1/4$.

There are many cases in which changing the value of the parameter slightly does not change the qualitative behaviour of the map; for example, Exercise 14 shows that the logistic maps f_c with $c > 1/4$ are all topologically conjugate. Another way of putting this is to say that they all have the same qualitative behaviour, even though the quantitative behaviour (in this case, how quickly orbits go to $+\infty$) varies with the map.

A bifurcation occurs when an arbitrarily small change in the value of the parameter *does* change the qualitative behaviour of the map:

DEFINITION 59. A one-parameter family of maps $F_c: X \rightarrow X$ has a *bifurcation* at c_0 if for all $\varepsilon > 0$, there exists a parameter value $c \in (c_0 - \varepsilon, c_0 + \varepsilon)$ for which F_c and F_{c_0} are not topologically conjugate.

One may easily show that the periodic orbit structure of a map is an invariant of topological conjugacy; that is, two topologically conjugate maps f and g must have the same numbers of fixed points, points of period two,

period three, etc. Thus a change in this orbit structure, such as the appearance of any new periodic orbits, immediately heralds a bifurcation in the system. Similarly, because stability is determined by where trajectories converge, a change in the type of stability of a periodic orbit also indicates a bifurcation.

For the logistic family f_c , the first bifurcation occurs at $c = 1/4$; for $c > 1/4$ there are no fixed points, while for $c = 1/4$ there is one, and for $c < 1/4$ there are two, given by

$$p_1 = \frac{1 - \sqrt{1 - 4c}}{2}, \quad p_2 = \frac{1 + \sqrt{1 - 4c}}{2}.$$

The stability of these fixed points is determined by the absolute value of the derivative, and we see that

$$f'_c(p_1) = 1 - \sqrt{1 - 4c}, \quad f'_c(p_2) = 1 + \sqrt{1 - 4c}.$$

Thus p_2 is unstable for all values of c , while p_1 is stable for a little while after the bifurcation at $c = 1/4$. As long as this state of affair persists, any trajectory which begins in the interior of the interval $I = [-p_2, p_2]$ converges to p_1 , the two trajectories which begin at the endpoints converge to p_2 , and any trajectory which begins outside of I diverges to $+\infty$.

So how long does this state of affairs persist? Observe that $f'_c(p_1)$ decreases as c decreases, and also that $f'_c(p_1) = -1$ when $c = -3/4$. Thus for $-3/4 < c < 1/4$, all trajectories in $(-p_2, p_2)$ converge to p_1 , but for $c < -3/4$, both fixed points are unstable. This implies that a bifurcation occurs at $-3/4$, since the stability of p_1 changes; to determine what trajectories behave like for $c < -3/4$, we look for periodic orbits, since there are no new fixed points.

Recall that a period-two orbit of f_c corresponds to a fixed point of f_c^2 , and so we want to solve the equation

$$f_c^2(x) = (x^2 + c)^2 + c = x,$$

which may be written as the quartic polynomial

$$x^4 + 2cx^2 - x + c^2 + c = 0.$$

This is made rather easier to solve by the observation that we already know two of the roots; the fixed points p_1 and p_2 of the original map f_c . In particular, the polynomial $f_c(x) - x$ divides $f_c^2(x) - x$, since a root of the former is obviously a root of the latter. Dividing, we obtain

$$\frac{x^4 + 2cx^2 - x + c^2 + c}{x^2 - x + c} = x^2 + x + 1 + c = 0,$$

which has solutions

$$q_1 = -\frac{1}{2} - \sqrt{-\frac{3}{4} - c}, \quad q_2 = -\frac{1}{2} + \sqrt{-\frac{3}{4} - c}.$$

These are real numbers if and only if $c \leq -3/4$; in other words, f_c has a period-two orbit if and only if the fixed point p_1 is unstable! One may easily

verify that $f_c(q_1) = q_2$ and $f_c(q_2) = q_1$, and that q_1 and q_2 lie on either side of p_1 .

The stability of this period-two orbit is given by the derivative of the map f_c^2 for which it is a fixed point; using the chain rule, we observe that

$$\begin{aligned} (f_c^2)'(q_1) &= f_c'(q_2)f_c'(q_1) \\ &= (-1 + \sqrt{-3 - 4c})(-1 - \sqrt{-3 - 4c}) \\ &= 1 - (-3 - 4c) \\ &= 4 + 4c. \end{aligned}$$

When the period-two orbit is born, at $c = -3/4$, we have $(f_c^2)'(q_1) = 1$, and this quantity decreases as c decreases, becoming equal to -1 when $c = -5/4$. Thus for $-5/4 < c < -3/4$, the period-two orbit is stable, and it is possible to show that every trajectory which begins in $(-p_2, p_2)$ asymptotically approaches the period-two orbit, in the sense that $f^{2n}(x) \rightarrow q_i$ for either $i = 1$ or $i = 2$.

This pattern of behaviour continues as c decreases further, with successive periodic orbits of length 2^n becoming unstable and spawning stable orbits of length 2^{n+1} , which become unstable in their turn, and so on *ad infinitum*. However, the algebraic approach we have been following becomes increasingly messy, as we must deal with polynomials of higher and higher degree.

b. Different sorts of bifurcations. Consider the parameter values $c_0 = 1/4$ and at $c_1 = -3/4$. At both of these values, the periodic orbit structure of the logistic map f_c changes, and so a bifurcation occurs; however, the bifurcations are of different sorts. As c decreases through c_0 , we go from having no fixed points to having two, one stable and one unstable. At c_1 , on the other hand, there is a pre-existing fixed point, which persists through the bifurcation; the change is in the stability of that fixed point and in the appearance of an attracting period-two orbit.

These two types of bifurcations are common enough to merit their own names; the bifurcation at c_0 is an example of a *saddle-node bifurcation* (sometimes called a *tangent bifurcation*), and the bifurcation at c_1 is an example of a *period-doubling bifurcation* (sometimes called a *pitchfork bifurcation*). The following definitions make these notions precise; in what follows, $F_\lambda: \mathbb{R} \rightarrow \mathbb{R}$ is any one-parameter family of continuous maps of the real line.

DEFINITION 60. F_λ has a *saddle-node (or tangent) bifurcation* at λ_0 if there exists an open interval $I \subset \mathbb{R}$ and $\varepsilon > 0$ such that the following hold:

- (1) For $\lambda_0 - \varepsilon < \lambda < \lambda_0$, the map F_λ has no fixed points in I .
- (2) For $\lambda = \lambda_0$, the map F_λ has one fixed point in I , which is neutral.
- (3) For $\lambda_0 < \lambda < \lambda_0 + \varepsilon$, the map F_λ has two fixed points in I , one attracting and one repelling.



FIGURE 32. Two examples of saddle-node bifurcations.

The conditions in the above definition imply that the fixed points appear as λ increases through λ_0 . We also say that F_λ has a saddle-node bifurcation at λ_0 if the fixed points appear as λ decreases; that is, if the family $\{F_{-\lambda}\}$ satisfies the above definition at $-\lambda_0$.

A typical picture for a saddle-node bifurcation is shown in Figure 32; note that at the value $\lambda = \lambda_0$, the graph of F_λ is tangent to the bisectrix $y = x$, hence the alternate name “tangent bifurcation”.

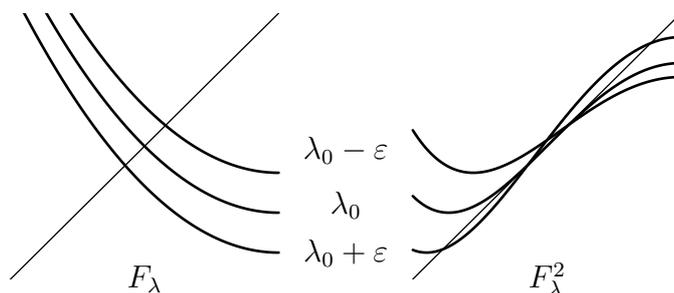


FIGURE 33. A period-doubling bifurcation.

DEFINITION 61. F_λ has a *period-doubling (or pitchfork) bifurcation* at λ_0 if there exists an open interval $I \subset \mathbb{R}$ and $\varepsilon > 0$ such that the following hold:

- (1) For all $\lambda_0 - \varepsilon < \lambda < \lambda_0 + \varepsilon$, the map F_λ has a unique fixed point p_λ in I .
- (2) For $\lambda_0 - \varepsilon < \lambda < \lambda_0$, the fixed point p_λ is attracting, and F_λ has no period-two cycle in I .
- (3) For $\lambda_0 < \lambda < \lambda_0 + \varepsilon$, the fixed point p_λ is repelling, and there exists a unique period-two cycle $\{q_\lambda^1, q_\lambda^2\}$ in I , which is attracting.
- (4) As λ converges to λ_0 , both points q_λ^i of the period-two cycle converge to the fixed point p_{λ_0} .

The above definition describes the process by which a stable fixed point becomes unstable and sheds a stable orbit of period two as λ increases through λ_0 ; as before, we also call λ_0 a period-doubling bifurcation point if

this happens as λ decreases through λ_0 . We also allow the case in which a *repelling* fixed point becomes stable and sheds an unstable period-two orbit.

Finally, we will be flexible enough with the terminology to say that F_λ has a period-doubling bifurcation at λ_0 if F_λ^n satisfies the above criteria for some n ; that is, if some periodic orbit of F_λ with length n changes stability at λ_0 and sheds a new periodic orbit of length $2n$ with the original stability properties.

A typical picture for a period-doubling bifurcation is shown in Figure 33, in which the graph of F_λ shows the fixed point becoming unstable as $F'_\lambda(p)$ passes -1 , and the graph of F_λ^2 shows the birth of a period-two orbit around the fixed point. Figure 34 shows the cobweb diagram near the fixed point for various parameter values near λ_0 .

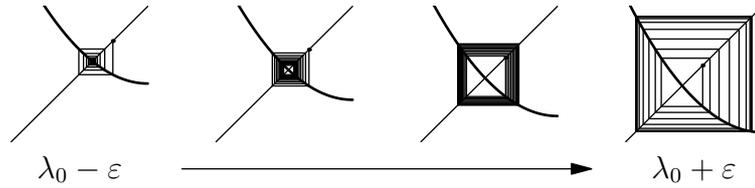


FIGURE 34. Orbits near a period-doubling bifurcation.

Lecture 25

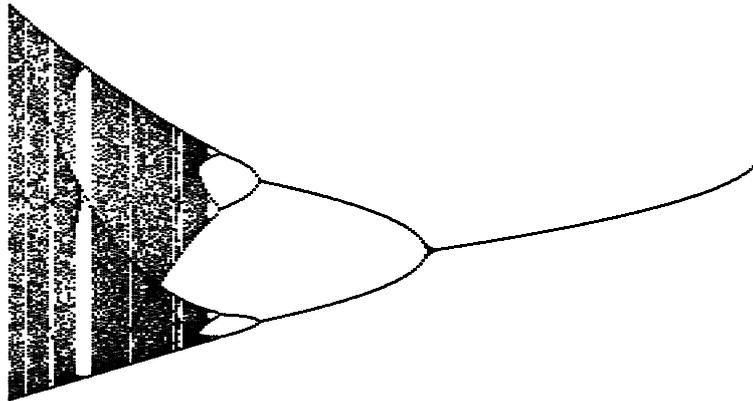


FIGURE 35. The bifurcation diagram for the family of logistic maps.

a. The simple part of the bifurcation diagram. Figure 35 shows the bifurcation diagram for the family of logistic maps $f_c: x \mapsto x^2 + c$; comparing this with Figure 31 for the FitzHugh-Nagumo model, we see many of the same qualitative features.

In the previous lecture, we examined two sorts of bifurcations; the tangent bifurcation at $c_0 = 1/4$, where two fixed points are born, one stable and one unstable, and the period-doubling bifurcation at $c_1 = -3/4$, where the stable fixed point becomes unstable and an attracting period-two orbit appears. Figure 35 shows further period-doubling bifurcations at $c_2 > c_3 > c_4 > \dots$; for $c \in (c_{n+1}, c_n)$, the map f_c has an attracting periodic orbit of length 2^n .

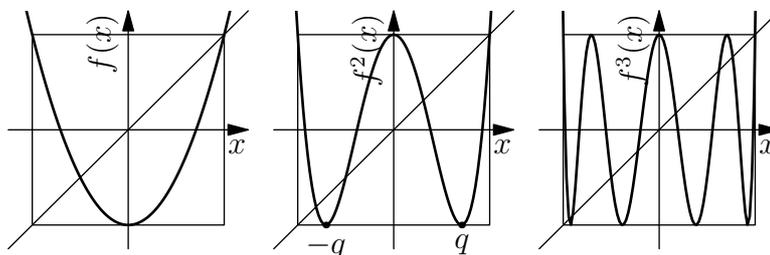
It is apparent from the diagram that the distance between successive bifurcations shrinks as n grows; indeed, it is possible to show that

$$\lim_{n \rightarrow \infty} \frac{c_n - c_{n-1}}{c_{n+1} - c_n} = \delta \approx 4.669\dots;$$

that is, the exponential rate of decay of this distance is Feigenbaum's constant, just as it was for the FitzHugh-Nagumo model.

Because $c_n - c_{n-1}$ decreases exponentially, the sequence $\{c_n\}$ converges to a limit c_∞ , the *Feigenbaum parameter*. This parameter lies at the end of the period-doubling cascade, and so for every $c < c_\infty$, the map f_c has period points of order 2^n for every natural number n . In particular, it is no longer Morse–Smale.

What do the dynamics of these maps look like? How do we describe the structure of the bifurcation diagram in the regime $c \leq c_\infty$? We begin by jumping ahead a little ways, and examining what happens for $c \leq -2$, before returning to examine the truly intricate part of the picture.

FIGURE 36. The map $f = f_{-2}$ and its iterates.

b. The other end of the bifurcation diagram. To this end, consider the parameter value $c = -2$, which is the smallest value of c shown in Figure 35. We write f for the map $f_{-2}: x \mapsto x^2 - 2$, and observe that f has fixed points at $p_1 = -1$ and $p_2 = 2$. Furthermore, if $|x| > 2$, then $f^n(x) \rightarrow +\infty$ as $n \rightarrow \infty$, and so the only interesting trajectories are those which remain within the interval $[-2, 2]$.

Observe that f achieves its minimum value at $f(0) = -2$, and that if $x \in [-2, 2]$, then $f(x) \in [-2, 2]$ as well. Thus we are really interested in the dynamics of $f: [-2, 2] \rightarrow [-2, 2]$.

f is monotonic on each of the intervals $[-2, 0]$ and $[0, 2]$; in fact, it maps each of these intervals homeomorphically to the entire interval $[-2, 2]$. Thus the action of f may be thought of as a combination of stretching and folding; the interval $[-2, 2]$ is first stretched out, and then folded in half, so that each half of the original interval has been stretched out to cover the whole thing.

Upon iterating the map f , we see that if J is an interval whose image under f covers $[-2, 2]$ once (that is, $f: J \rightarrow [-2, 2]$ is a homeomorphism), then the image of J covers $[-2, 2]$ twice under the action of f^2 . In particular, as shown in Figure 36, there are points $-q < 0 < q$ such that each of the intervals $I_1 = [-2, -q]$, $I_2 = [-q, 0]$, $I_3 = [0, q]$, and $I_4 = [q, 2]$ has $f^2(I_j) = [-2, 2]$, and f^2 is a bijection from each I_j to $[-2, 2]$. A similar observation holds for f^3 , where we have eight intervals, and for higher iterates f^n , where we have 2^n intervals which are mapped homeomorphically onto $[-2, 2]$.

EXERCISE 15. Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be continuous, and suppose that $a \leq a' < b' \leq b$ are such that $g([a', b']) \supset [a, b]$ (we say that $[a', b']$ *g-covers* $[a, b]$). Show that g has a fixed point in $[a', b']$.

It follows from Exercise 15 that the map f^n has at least 2^n fixed points, and examination of the graph of f^n shows that this number is exact. Thus for every $n \geq 1$, the map f has 2^n period points of period n (of course, some of these are also periodic points of period k for some $k < n$); this is a far cry from the limited number of periodic orbits found in the Morse–Smale case. Indeed, the exponential growth rate of the number of periodic orbits of a period n as n increases is somehow indicative of the chaotic behaviour of the map; this growth rate is in many cases related to a quantity known

as the *topological entropy*, whose positiveness is often taken as an indicator of chaos.

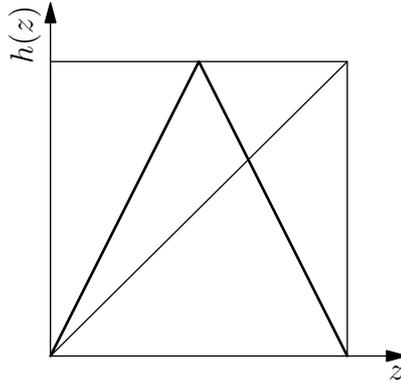


FIGURE 37. The tent map $h: [0, 1] \rightarrow [0, 1]$.

Before leaving the parameter value $c = -2$, we observe that by some felicitous alignment of the stars, the map f can be put into a rather simpler form—indeed, a piecewise linear form—via a clever change of coordinates. First, as a special case of the observation at the beginning of Lecture 24, f is conjugated to the map $g: y \mapsto 4y(1 - y)$ from $[0, 1]$ to itself by the change of coordinates

$$\begin{aligned}\phi: [0, 1] &\rightarrow [-2, 2], \\ y &\mapsto 2 - 4y.\end{aligned}$$

Then, using the further change of coordinates

$$\begin{aligned}\psi: [0, 1] &\rightarrow [0, 1], \\ z &\mapsto \sin^2\left(\frac{\pi z}{2}\right),\end{aligned}$$

it turns out that the following diagram commutes:

$$\begin{array}{ccc} [0, 1] & \xrightarrow{h} & [0, 1] \\ \downarrow \psi & & \downarrow \psi \\ [0, 1] & \xrightarrow{g} & [0, 1] \\ \downarrow \phi & & \downarrow \phi \\ [-2, 2] & \xrightarrow{f} & [-2, 2] \end{array}$$

Here $h: [0, 1] \rightarrow [0, 1]$ is the *tent map* defined by

$$(82) \quad h(z) = \begin{cases} 2z & 0 \leq z \leq 1/2, \\ 2(1 - z) & 1/2 \leq z \leq 1, \end{cases}$$

whose graph is shown in Figure 37.

A simple computation verifies the above claim:

$$\begin{aligned} g(\psi(z)) &= 4 \sin^2 \left(\frac{\pi z}{2} \right) \left(1 - \sin^2 \left(\frac{\pi z}{2} \right) \right) \\ &= 4 \sin^2 \left(\frac{\pi z}{2} \right) \cos^2 \left(\frac{\pi z}{2} \right) \\ &= \left(2 \sin \frac{\pi z}{2} \cos \frac{\pi z}{2} \right)^2 \\ &= \sin^2 \pi z, \end{aligned}$$

while for $0 \leq z \leq 1/2$,

$$\psi(h(z)) = \sin^2 \frac{\pi \cdot 2z}{2} = \sin^2 \pi z,$$

and for $1/2 \leq z \leq 1$,

$$\psi(h(z)) = \sin^2(\pi - \pi z) = \sin^2 \pi z.$$

This conjugacy allows us to answer certain questions about the non-linear map f by first answering them for the piecewise linear map h , in which context they are often more tractable.

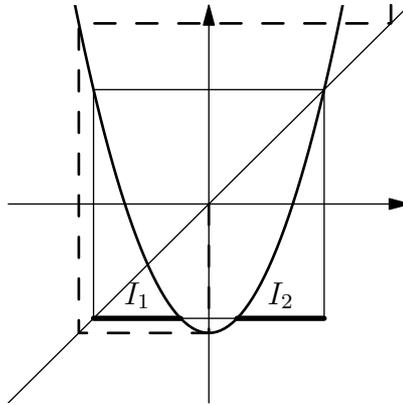


FIGURE 38. Trajectories escaping to infinity for $c < -2$.

c. The centre cannot hold—escape to infinity. What happens when $c < -2$? We begin by recalling that as before, there are two fixed points $p_1 < p_2$, and only the interval $[-p_2, p_2]$ is of interest, since any trajectory which leaves this interval diverges to infinity.

For $-2 \leq c \leq 1/4$, the interval $[-p_2, p_2]$ is invariant; any trajectory which begins there stays there, and so we only needed to exclude trajectories which began outside the interval of interest. To see if this behaviour continues for $c < -2$, we must examine $f_c(0) = c$, since this is the minimal value assumed by f_c , and indeed, $f_c([-p_2, p_2]) \subset [-p_2, p_2]$ if and only if $f_c(0) = c \geq -p_2$.

We compare these values by observing that c is the constant term in the fixed point equation $f_c(x) - x = x^2 - x + c = 0$, which is a quadratic

polynomial, and hence it is the product of the roots of that polynomial, which are p_1 and p_2 . The sum of the roots is the negative of the linear coefficient, and so $p_1 + p_2 = 1$, whence $c = p_2(1 - p_2)$. It follows that $[-p_2, p_2]$ is invariant if and only if $p_2(1 - p_2) \geq -p_2$; that is, if and only if $p_2(2 - p_2) \geq 0$.

The fixed point p_2 is always positive, and so we see that the interval is invariant for $p_2 \leq 2$, while for $p_2 > 2$, some points in the interval $[-p_2, p_2]$ have images outside that interval, and thus have a trajectory which escapes to infinity. This happens precisely when $c < -2$; thus for the parameter range we are considering now, we have the picture shown in Figure 38, where any point not in the intervals I_1 or I_2 is mapped outside of $[-p_2, p_2]$ by f_c .

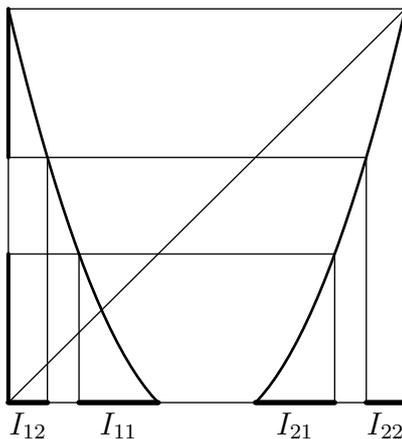


FIGURE 39. Finding an invariant set for a map with escape.

This should start to sound familiar by now; we have a map from an interval to itself, but we can only start at points whose images remain in the interval. The intervals I_1 and I_2 contain those points x for which $f_c(x) \in [-p_2, p_2]$; Figure 39 shows the construction of four intervals I_{ij} which contain all points x for which $f_c^2(x) \in [-p_2, p_2]$. Continuing in this manner, we have exactly the same type of Cantor-like construction that we examined earlier in the course, and one sees that the restriction of f_c to $I_1 \cup I_2$ is exactly the sort of map shown in Figure 17.

Thus for f_c with $c < -2$, the set of all points whose trajectories remain bounded is a Cantor-like set; outside of this set, all trajectories diverge to infinity. In particular, this set has Lebesgue measure zero, and so is “invisible” to the bifurcation diagram in Figure 35; remember that the method of constructing that diagram ensures that we only see sets in phase space which attract nearby trajectories.

Lecture 26

a. Some parts of phase space are more equal than others. In our investigations of the logistic map f_c for various values of the parameter c , we have found that some parts of the phase space \mathbb{R} are more interesting than others, in the sense that they capture the essential long-term behaviour of f_c . For example, when $c > c_\infty$, almost every trajectory tends to the stable periodic orbit, and so the points in that orbit are the most important part of phase space for f_c . When $c = -2$, the points in the interval $[-2, 2]$ are important, but those outside it are not so interesting, because their iterates tend to infinity. Finally, when $c < -2$, there is a Cantor-like set C which comprises points whose orbits remain bounded; this set C captures all the interesting long-term behaviour of f_c .

How do we formalise these ideas? How can we define what makes some sets capture interesting aspects of the dynamics, while others are somehow negligible? What properties should these “interesting” sets have?

The first important property is invariance; we want to consider a set $E \subset \mathbb{R}$ which is mapped into itself by f , so that no trajectories escape from E ; otherwise E does not contain the long-term behaviour of all the trajectories which begin in E .

Secondly, we want E to be minimal in some sense; for f_c with $c > c_\infty$, the interval $[-p_2, p_2]$ is certainly invariant, but it is too big—given an initial condition $x \in [-p_2, p_2]$, there are many open sets in $[-p_2, p_2]$ which the trajectory of x never reaches, and which are therefore of no importance in describing the long-term behaviour of that trajectory. In particular, we would like the orbit of some well-chosen point x to be dense in E .

These considerations motivate the following definition.

DEFINITION 62. Let $f: X \rightarrow X$ be continuous, and fix $x \in X$. The ω -limit set of x is

$$\omega(x) = \bigcap_{N=0}^{\infty} \overline{\bigcup_{n=N}^{\infty} \{f^n(x)\}}.$$

Equivalently, a point $y \in X$ is in the ω -limit set of x if and only if y is in the closure of the forward trajectory of every iterate of x ; that is, if and only if there exists a sequence of natural numbers $n_k \rightarrow \infty$ such that $\lim_{k \rightarrow \infty} f^{n_k}(x) = y$.

EXAMPLE 63. If x is a periodic point with $f^p(x) = x$, then $\omega(x) = \{x, f(x), \dots, f^{p-1}(x)\}$. Similarly, if x approaches a periodic orbit, that is, if there exists $y = f^p(y)$ such that $\lim_{n \rightarrow \infty} f^{np}(x) = y$, then $\omega(x) = \{y, f(y), \dots, f^{p-1}(y)\}$.

EXERCISE 16. Consider the map $f: \mathbb{C} \rightarrow \mathbb{C}$ given by $f(z) = e^{2\pi i \alpha} z$, where α is an irrational real number; f is the map which rotates the complex plane by $2\pi\alpha$ around the origin. Show that for every $z_0 \in \mathbb{C}$,

$$\omega(z_0) = \{z \in \mathbb{C} \mid |z| = |z_0|\}.$$

EXERCISE 17. Consider the shift $\sigma: \Sigma_2^+ \rightarrow \Sigma_2^+$, and show that there exists a sequence $x = (i_1, i_2, \dots) \in \Sigma_2^+$ such that $\omega(x) = \Sigma_2^+$; in particular, x has dense orbit.

Recalling the construction of the bifurcation diagram, we see that it plots approximations to the ω -limit set of a random point for each parameter value c . Does it matter which point we choose? Of course, if we choose an unstable periodic point x , then $\omega(x)$ will not capture the same information as $\omega(y)$, where y is not on the periodic orbit, and one may find other cases where $\omega(x)$ and $\omega(y)$ differ for $x \neq y$.

It turns out, though, that in many important situations, *almost every* point x (with respect to Lebesgue measure) has the same ω -limit set; that is, there exists a set $E \subset \mathbb{R}$ of Lebesgue measure zero such that $\omega(x) = \omega(y)$ for all $x, y \notin E$. For example, in the family of logistic maps, almost every initial condition x has ω -limit set equal to the unique stable periodic orbit when $c > c_\infty$, and for $c = -2$, it is possible to show that almost every initial condition x has $\omega(x) = [-2, 2]$.

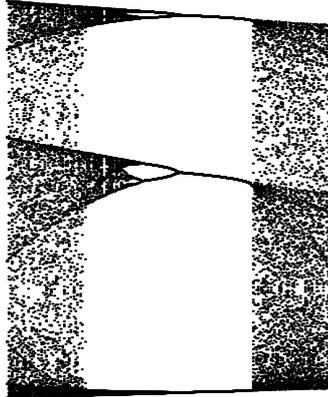


FIGURE 40. The bifurcation diagram in the period-three window.

b. Windows of stability. Let us return to the bifurcation diagram shown in Figure 35, and finally turn our attention to the truly interesting part of the picture, the parameter values $-2 \leq c \leq c_\infty$. Within this range, we find a number of *windows of stability*, intervals for c within which the map f_c suddenly has a stable periodic orbit once again, to which almost every point is attracted.

The largest and most conspicuous of these windows of stability is the period-three window between $c \approx -1.791$ and $c \approx -1.748$, which is shown in Figure 40. Notice that while the period-three orbit is stable at the beginning of this window (on the right), it eventually becomes unstable and gives birth to a stable period-six orbit, and the whole period-doubling cascade occurs here just as it did for $c > c_\infty$, leading to another occurrence of Feigenbaum's constant δ .

Because the map has a period-three orbit in this range of parameters, we can apply Corollary 57 of Sarkovskii's Theorem. In lieu of proving the theorem in its full generality, we will prove this special case, which uses the same ideas, but for which the combinatorial argument is simpler.

PROOF OF COROLLARY 57. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be continuous, and let $x_0 = f(x_2)$, $x_1 = f(x_0)$, and $x_2 = f(x_1)$ be the three points in a period-three orbit. By changing the numbering if necessary, we can assume that $x_0 < \min(x_1, x_2)$, and there are now two possibilities:

$$x_0 < x_1 < x_2 \quad \text{or} \quad x_0 < x_2 < x_1.$$

We prove the theorem in the case $x_0 < x_1 < x_2$; the proof in the other case is similar.

Let $I_1 = [x_0, x_1]$, and $I_2 = [x_1, x_2]$. It follows immediately from the Intermediate Value Theorem that

$$(83) \quad f(I_1) \supset [f(x_0), f(x_1)] = [x_1, x_2] = I_2,$$

and also that

$$(84) \quad f(I_2) \supset [f(x_2), f(x_1)] = [x_0, x_2] = I_1 \cup I_2.$$

Given $n \geq 1$, we find a periodic orbit of period n as follows. If $n = 1$, then we are after a fixed point of f , whose existence is guaranteed by the result of Exercise 15 applied to I_2 .

An easy induction argument shows that $f^n(I_1) \supset I_2$ for all $n \geq 1$, and hence $f^n(I_1) \supset I_1$ for all $n \geq 2$. Then the result of Exercise 15 applied to I_1 and f^n shows that f^n has a fixed point x in I_1 , which is almost enough to complete the proof. However, we want to prove the slightly stronger statement that n is the smallest period of x ; that is, that $f^k(x) \neq x$ for $1 \leq k \leq n - 1$. For this, we need the following lemma.

LEMMA 64. *If $g: \mathbb{R} \rightarrow \mathbb{R}$ is continuous and $[a', b']$ g -covers $[a, b]$ (that is, $g([a', b']) \supset [a, b]$), then there exists a subinterval $[a'', b''] \subset [a', b']$ such that $g([a'', b'']) = [a, b]$.*

PROOF. Let $E = \{x \in [a, b] \mid g(x) \leq a\}$ and $F = \{x \in [a, b] \mid g(x) \geq b\}$, and consider $a_1 = \sup E$, $a_2 = \sup F$. Without loss of generality, suppose that $a_1 < a_2$ (the proof in the other direction is similar). Then let $a'' = a_1$, $b'' = \inf E \cap [a_1, b']$, and the result follows. \square

From the lemma and (83)–(84), it follows that there exist intervals $I_{12} \subset I_1$ and $I_{21}, I_{22} \subset I_2$ such that $f(I_{ij}) = I_j$. Continuing, we find nested sequences of basic intervals $I_{i_1 \dots i_n}$ such that $f(I_{i_1 \dots i_n}) = I_{i_2 \dots i_n}$. Because the length of these intervals may not go to zero as $n \rightarrow \infty$, we cannot carry out the full Cantor-like construction and obtain a conjugacy with the subshift Σ_A^+ , $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$; however, we can observe that for the particular sequence $i_1 = 1, i_2 = i_3 = \dots = i_n = 2$, we have $f^n(I_{i_1 \dots i_n}) = I_1 \cup I_2 \supset I_{i_1 \dots i_n}$, and so there exists $x \in I_{i_1 \dots i_n}$ such that $f^n(x)$. Because $f^k(x) \in I_2$ for all $1 \leq k \leq n - 1$, n must be the minimal period of x . \square

This result shows that despite the long-term stability of orbits in the period-three window, the system is not Morse–Smale, because there are infinitely many periodic orbits. In fact, we observe what is known as *transient chaos* in this parameter regime; the trajectory of a point x may follow an unstable periodic orbit for quite some time before being repelled, at which point it may follow some other unstable orbit for a spell, and may take a very long time to actually settle down to the stable periodic orbit. Thus the trajectory we observe may initially appear chaotic by spending a long while wandering through the intricately intertwined tangle of periodic orbits before it becomes regular.

c. Outside the windows of stability. We see a similar picture to the one described above when we look at the windows of stability corresponding to periodic orbits of other lengths; in the first place, Sarkovskii’s theorem implies the existence of infinitely many periodic orbits, and prescribes which lengths must appear, and in the second place, the stable periodic orbit undergoes a period-doubling cascade as c decreases and eventually moves out of the window of stability.

Let S be the set of parameter values c for which the map f_c has a stable periodic orbit; S is the union of all the windows of stability. It has been shown that S is open (which is not too hard to prove) and dense in $[-2, 1/4]$ (which is quite hard to prove). So in some sense, the set S is quite large; indeed, topologically speaking, it is as large as it can be.

However, there are other sorts of behaviour possible for the logistic map f_c . It turns out that these are epitomised by the two cases $c = -2$ and $c = c_\infty$. In the first case, the ω -limit set $\omega(x)$ contains an interval (and in particular, has Hausdorff dimension equal to 1) for almost every x ; denote by A the set of parameter values c such that f_c has this behaviour. In the second case, we find that $\omega(x)$ is a Cantor set for almost every x (and hence has Hausdorff dimension less than 1); denote by C the corresponding set of parameter values.

How big are the sets A and C ? We have already said that S is open and dense, and these three sets are disjoint; thus A and C are both nowhere dense sets. However, a celebrated result due to Michael Jakobson shows that A has positive Lebesgue measure, and so the set of parameter values for which the trajectory of a randomly chosen point fills out an interval is non-negligible.

Lecture 27

a. Chaos in the logistic family. One of the most remarkable results to date concerning the family of logistic maps $f_c: x \mapsto x^2 + c$ is *Jakobson's theorem*, which describes a set $A \subset [-2, -2 + \varepsilon]$ of parameter values with $\text{Leb}(A) > 0$ such that for every $c \in A$, the trajectories of the map f_c are chaotic; that is, there exists a symbolic coding of the map such that all possible sequences of symbols appear as codes of actual trajectories.

While the set A for which such behaviour can be rigorously proven to occur is confined to parameter values near -2 , computer simulations suggest that this behaviour actually occurs for a non-negligible proportion of parameter values all the way up to c_∞ , the end of the period-doubling cascade.

b. Attractors for the FitzHugh–Nagumo model. The chasm between what has been proved and what is believed to be true based on numerical results is even wider when we consider the FitzHugh–Nagumo system (76). In the first place, the interval map on which the model is based is a cubic polynomial, rather than quadratic, and much less is known about such maps. Furthermore, because the map has been perturbed into two dimensions, many of the rigorous results no longer apply, particularly those which depend on properties unique to the one-dimensional case, such as the Intermediate Value Theorem.

Despite the lack of rigorous results, the empirical evidence overwhelmingly suggests that the results which have been proved in very restricted circumstances (one-dimensional quadratic maps in a limited parameter range) hold much more generally, as is suggested, for example, by the bifurcation diagram in Figure 31. We see a period-doubling cascade leading to the onset of chaos at A_∞ , beyond which there are windows of stability surrounded by maps with chaotic behaviour.

Throughout all this, the map $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ has three fixed points, $0, p_1, p_2$, two of which began life as stable fixed points, and then lost their stability at the bifurcation points A_1 and A'_1 . During the period-doubling cascade, all trajectories wind up approaching a stable periodic orbit of length 2^n ; after the onset of chaos, however, there are no stable periodic orbits to approach, except in the windows of stability. So where do the orbits go?

It turns out that one can find a rectangle R , as shown in Figure 41, which contains the fixed points $0, p_1, p_2$, and which is mapped into itself by f ; that is, $f(R) \subset R$. The picture shows how R is squeezed in the vertical direction, stretched in the horizontal direction, and then folded over on itself so that it fits in R . The ‘S’ shape of $f(R)$ is due to the fact that the coordinate function f_1 contains a cubic polynomial in u_1 .

A region R such that $f(R) \subset R$ is known as a *trapping region* for the map f ; once a trajectory enters this region, it will never leave it. Furthermore, $f(R)$ is also a trapping region; given $x \in f(R) \subset R$, we have $f(x) \in f(R)$, and hence $f^2(R) \subset f(R) \subset R$. Continuing in this way, we obtain a nested

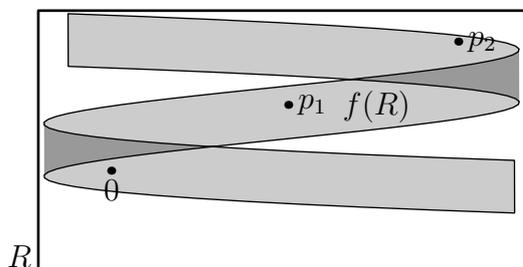


FIGURE 41. A trapping region for the FitzHugh–Nagumo map.

sequence of trapping regions:

$$R \supset f(R) \supset f^2(R) \supset \dots \supset f^n(R) \supset f^{n+1}(R) \supset \dots$$

We may “take the limit” of this sequence by taking the intersection of all these trapping regions, and obtain

$$(85) \quad \Lambda = \bigcap_{n \geq 0} f^n(R).$$

The intersection Λ is an *attractor* for the map f ; every trajectory which enters the trapping region R not only stays in R , but approaches Λ . In particular, $\omega(x) \subset \Lambda$, and Λ has the following properties:

- (1) Λ is closed and bounded, hence compact.
- (2) Λ is f -invariant; $f(\Lambda) = \Lambda$.
- (3) Λ is the largest invariant subset of R .
- (4) Λ contains the three fixed points $0, p_1, p_2$.

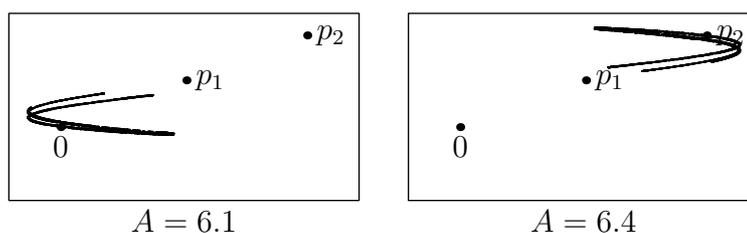
FIGURE 42. The attractor for the FitzHugh–Nagumo map with $\alpha = .01, \beta = .02, \theta = .51, \gamma = .2$, and varying A .

Figure 42 shows the attractor Λ for two different parameters of A ; notice that as A changes, the attractor may move to different locations within R . As A continues to increase, the attractor “grows”, as shown in Figure 43.

For the values of A where the attractor shown in Figures 42 and 43 appears, all three fixed points are hyperbolic; that is, they have one expanding and one contracting direction. It can be shown that for each such fixed point p there exists a smooth curve $\gamma: (-\varepsilon, \varepsilon) \rightarrow \mathbb{R}^2$ such that $\gamma(0) = p$ and $\gamma'(0)$ is the eigenvector corresponding to the larger eigenvalue, and which is

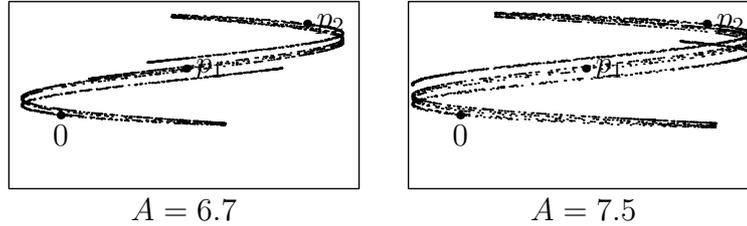


FIGURE 43. Changes in the attractor as A increases.

expanding in the following sense; writing $\Gamma_\varepsilon = \{\gamma(t) \mid -\varepsilon < t < \varepsilon\}$ for the *local unstable manifold*, the image $f(\Gamma_\varepsilon)$ is a curve which contains Γ_ε .

It follows that $f^2(\Gamma_\varepsilon) \supset f(\Gamma_\varepsilon)$, and indeed, we get a nested sequence of *increasing* curves

$$\Gamma_\varepsilon \subset f(\Gamma_\varepsilon) \subset \cdots \subset f^n(\Gamma_\varepsilon) \subset \cdots$$

Taking the union of all these, we obtain the *global unstable manifold*

$$\Gamma = \bigcup_{n \geq 0} f^n(\Gamma_\varepsilon).$$

It follows that Γ is f -invariant and contained in R ; from the properties of the attractor Λ , this implies that $\Gamma \subset \Lambda$.

In fact, it is conjectured (and widely believed) that Γ is dense in Λ ; however, no rigorous proof is known.

c. The Smale–Williams solenoid. As we have seen, the FitzHugh–Nagumo model, while very rich in intricate and interesting behaviour, is quite difficult to analyse. We thus turn our attention to simpler examples, which exhibit a similar richness of behaviour but are rather more tractable.

Our first such example is a map from the solid torus to itself. Abstractly, the solid torus is

$$P = D^2 \times S^1,$$

the direct product of a disc and a circle. It is embedded in \mathbb{R}^3 as the standard torus of revolution together with the region it encloses;

$$\rho(P) = \{(x, y, z) \in \mathbb{R}^3 \mid (\sqrt{x^2 + y^2} - 2)^2 + z^2 \leq 1\}.$$

Here $\rho: D^2 \times S^1 \rightarrow \mathbb{R}^3$ is the map given by

$$\rho(x, y, \theta) = ((2 + x) \cos \theta, (2 + x) \sin \theta, y),$$

where (x, y) are coordinates on the disc $D^2 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$, and θ is the angular coordinate on the circle $S^1 = \mathbb{R}/2\pi\mathbb{Z}$; we will use the coordinates (x, y, θ) on P , as they are more natural for our purposes.

Fixing parameters λ_1, λ_2, a such that $0 < a < 1$ and $0 < \lambda_i < \min(a, 1 - a)$, we define a map $f: P \rightarrow P$ by

$$(86) \quad f(x, y, \theta) = (\lambda_1 x + a \cos \theta, \lambda_2 y + a \sin \theta, 2\theta).$$

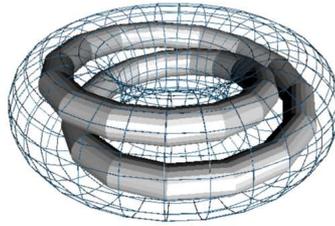


FIGURE 44. A map from the solid torus to itself.

The image of P under f is the shape shown in Figure 44. One way of visualising the way in which $f(P)$ is obtained by the action of f is to take the torus and slice it along a disc so that it becomes a tube; f squeezes this tube so that its cross-sections are no longer circles of radius 1, but ellipses with axes of length λ_1 and λ_2 , and then proceeds to stretch it along its axis by a factor of two, and finally to wrap the resulting longer, skinnier tube twice around the z -axis within the original solid torus.

In particular we have $f(P) \subset P$, and so we may repeat the procedure in the previous section, obtaining an attractor by taking the intersection of all images of P :

$$(87) \quad \Lambda = \bigcap_{n \geq 0} f^n(P).$$

The attractor Λ is known as the *Smale-Williams solenoid*, and turns out to have a fractal structure, which we will examine in the next lecture.

Lecture 28

a. The Smale–Williams solenoid. The map f from the solid torus P to itself, which was defined in (86), has an attractor Λ , determined as in (87). In order to investigate the structure of Λ , we look at a cross-section of the solid torus $P = D^2 \times S^1$, which corresponds to fixing the angular coordinate θ and considering the disc $D^2 \times \{\theta\}$.

From Figure 44, it is clear that the image $f(P)$, which is a long skinny tube wrapped twice around the z -axis, intersects this disc in two ellipses, whose axes have length λ_1 and λ_2 . The second image $f^2(P)$ is an even longer and skinnier tube which is wrapped *four* times around the z -axis, and intersects the disc in four ellipses, whose axes have lengths λ_1^2 and λ_2^2 .

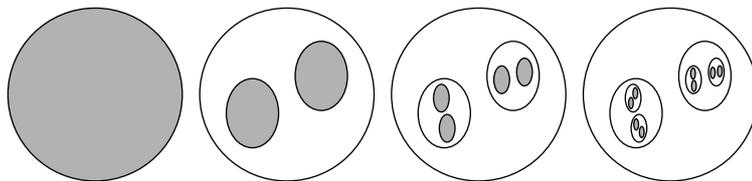


FIGURE 45. A cross-section of the Smale–Williams solenoid.

Continuing in this manner, we see that $f^n(P) \cap (D^2 \times \{\theta\})$ is the union of 2^n ellipses; the first few steps of the construction are shown in Figure 45. By now the reader should not be too shocked to discover that this is yet another example of a Cantor-like construction;¹⁶ the basic sets at each step are the ellipses just mentioned, and the cross-section $C = \Lambda \cap (D^2 \times \{\theta\})$ is the intersection of the basic sets at all levels.

Each basic set is the intersection of a tube with the disc $D^2 \times \{\theta\}$; as n increases, the diameters of the tubes decrease exponentially, and so upon passing to the limit set C , we see that each point in C is contained in precisely one curve which meets $D^2 \times \{\theta\}$ transversely. Thus in a neighbourhood of each cross-section (a slice out of the torus), the attractor is the direct product $C \times (-\varepsilon, \varepsilon)$. This product structure is only local, however; if we go far enough around the torus, the curves through different points of C may be connected.

The local product structure of the attractor Λ has more than just a geometric significance; it also helps us describe the dynamics of the map f . Through each point $A = (x, y, \theta) \in \Lambda$, we have a disc $D^2 \times \{\theta\}$ and a curve $\{(x, y)\} \times (-\varepsilon, \varepsilon)$; the former is contracting, while the latter is repelling, as follows. Given $B \in D^2 \times \{\theta\}$, we have

$$d(f(A), f(B)) = \max\{\lambda_1, \lambda_2\}d(A, B),$$

¹⁶Indeed, one could obtain the exact construction shown in Figure 14 by modifying f so that $f(P)$ wraps around the z -axis *three* times, and $\lambda_1 = \lambda_2$ depend on θ .

while for $B' \in \{(x, y)\} \times (-\varepsilon, \varepsilon)$, the orbits are driven further apart;

$$d(f(A), f(B')) = 2d(A, B').$$

Thus every point looks like a saddle; it has two stable directions (forming the disc) and one unstable direction (the curve). Notice, however, that this applies not only to fixed points, but to *any* point in Λ ; in general, then, the reference to which or from which the orbit $\{f^n(B)\}$ is attracted or repelled is not the point A itself, but its trajectory.

The set Λ is an important example of a *uniformly hyperbolic attractor*; it is hyperbolic because every point is hyperbolic (roughly, all directions are contracting or expanding), and it is uniformly so because the ratio of contraction or expansion can be bounded away from 1 independently of which point in Λ we consider. This represents a fundamentally new type of behaviour compared with the Morse–Smale systems we found for the logistic map in the period-doubling cascade, where there were only a finite number of hyperbolic points, which were all fixed. Here, by contrast, the hyperbolicity is ubiquitous.

What does this pervasive hyperbolicity mean for the dynamics of f ? If A and B are two points in Λ which do not lie on the same disc (stable manifold) or transversal curve (unstable manifold), then repeated iteration by f will decrease the distance between $f^n(A)$ and $f^n(B)$ in the stable direction (corresponding to the coordinates x and y) but will increase it in the unstable direction (corresponding to θ). In particular, the trajectory of B is repelled from the trajectory of A .

So almost every pair of trajectories moves apart under the action of f ; however, Λ is bounded, so they cannot move too far apart. Indeed, the trajectory of B is constantly being repelled from whatever trajectories it finds itself near at any given time, and eventually is repelled back towards the trajectory of A ; at this point it is once again repelled from the trajectory of A , and the whole cycle repeats itself.

This behaviour, this unending dispersal and return, is characteristic of hyperbolic dynamics. If we plot the x -coordinate of the trajectory of a point $A \in \Lambda$ as a function of n , we see a chaotic signal, without periodicity or pattern, as shown in Figure 46.

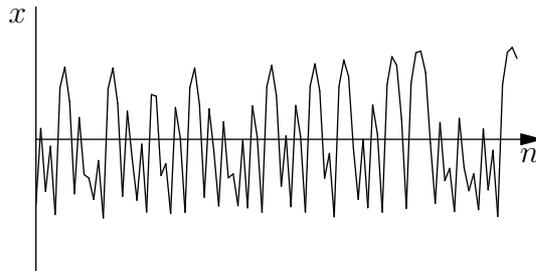


FIGURE 46. The x -coordinates of a trajectory as a chaotic signal.

There is another, terribly important, way in which the map $f: \Lambda \rightarrow \Lambda$ is chaotic. If we partition Λ into disjoint sets $\Lambda_1, \dots, \Lambda_k$, then we can code trajectories of f by recording which partition element the iterate $f^n(x)$ lands in. Thus to a trajectory $\{f^n(x)\}$ we associate the sequence

$$\omega = (\dots, i_{-2}, i_{-1}, i_0, i_1, i_2, \dots) \in \Sigma_k,$$

where i_j is such that $f^j(x) \in \Lambda_{i_j}$. This is very similar to our coding of Cantor-like sets with the symbolic space Σ_k^+ ; the primary difference here is that since the map f is invertible, we consider pre-images as well, and obtain a doubly infinite sequence $(i_j)_{j \in \mathbb{Z}}$.

What sequences in Σ_k do we obtain as codings of trajectories in Λ ? The answer depends on which partition we choose; for example, if i_j and $i_{j'}$ are such that $f(\Lambda_{i_j}) \cap \Lambda_{i_{j'}} = \emptyset$, then no sequence ω which contains the symbol i_j followed by the symbol $i_{j'}$ can correspond to a trajectory in Λ . In general, it is not possible to find a partition such that all sequences are admissible; however, for a uniformly hyperbolic system such as f , it is possible to find a partition $\{\Lambda_1, \dots, \Lambda_k\}$ such that for some $k \times k$ transition matrix A , the admissible sequences are precisely those which lie in Σ_A .

Thus while a uniformly hyperbolic map f cannot usually be described as a Bernoulli process (Σ_k), it can always be described as a Markov process (Σ_A). This fact can be used to derive various properties which are characteristic of chaos; for example, f has infinitely many periodic orbits, the set of periodic orbits is dense in Λ , almost every point $x \in \Lambda$ has a trajectory which is dense in Λ . . . the list goes on.

b. Quantifying the attractor. Having discussed some of the qualitative properties of the attractor Λ and its implications for the dynamics of f , we turn our attention to quantitative questions. In particular, having obtained a set Λ which has a fractal structure, we ask the natural question: what is the Hausdorff dimension of Λ ?

Intuitively, we expect dimension to be additive with respect to direct products; after all, the direct product of \mathbb{R}^m and \mathbb{R}^n is \mathbb{R}^{m+n} , and so it seems natural to conjecture that in general,

$$(88) \quad \dim_H(A + B) = \dim_H A + \dim_H B.$$

If (88) holds when A is the Cantor-like set $C = \Lambda \cap (D^2 \times \{\theta\})$ and B is the interval $(-\varepsilon, \varepsilon)$, then it follows that

$$(89) \quad \dim_H \Lambda = (\dim_H C) + 1,$$

since Λ can be decomposed into finitely many pieces, each of which is strongly equivalent to $C \times (-\varepsilon, \varepsilon)$.

EXERCISE 18. Using the product measures $m_H(\cdot, \alpha) \times m_H(\cdot, \beta)$, show that $\dim_H(A + B) \geq \dim_H A + \dim_H B$ in general.

Exercise 18 establishes one half of (88). However, the reverse inequality is not true in general; a counterexample to this effect was first produced by

Besicovitch, who also proved that equality *does* hold under slightly stronger assumptions.

THEOREM 65. *If $A, B \subset \mathbb{R}^n$ are such that $\dim_H A = \underline{\dim}_B A = \overline{\dim}_B A$, then equality holds in (88).*

Note that we only require coincidence of the Hausdorff and box dimensions for *one* of the two sets A and B . In particular, since the three dimensional quantities coincide for the interval $(-\varepsilon, \varepsilon)$, Theorem 65 establishes the expression (89) for the Hausdorff dimension of the Smale–Williams solenoid.

Of course, we still need to compute $\dim_H C$. In the simplest case where $\lambda_1 = \lambda_2 = \lambda$, the construction of C is exactly of the sort dealt with by Moran’s theorem, and we have

$$\dim_H C = \frac{\log 2}{-\log \lambda}.$$

In particular, we get

$$(90) \quad \dim_H \Lambda = 1 + \frac{\log 2}{-\log \lambda} = \log 2 \left(\frac{1}{\log 2} - \frac{1}{\log \lambda} \right),$$

where the reasons for the seemingly unusual form of this expression will soon be made clear. Indeed, we observe that $\log \lambda$ is nothing but the Lyapunov exponent in the contracting direction (the stable plane corresponding to the discs $D^2 \times \{\theta\}$), while $\log 2$ is the Lyapunov exponent in the expanding direction (the unstable line corresponding to $\{(x, y)\} \times (-\varepsilon, \varepsilon)$); thus the denominators of the terms in (90) are the Lyapunov exponents in various directions along the attractor. One can also define the entropy in this case, and for an appropriate measure μ , it turns out to be $\log 2$; thus (90) is reminiscent of our earlier expression (61), which related the Hausdorff dimension of a set to the entropy of a measure and the Lyapunov exponent of a map.

Although we do not give further details here, one can rigorously establish various relationships in the spirit of (90) which connect dimension, entropy, and Lyapunov exponents. The relatively simple form of this expression for the Smale–Williams solenoid is due to the linearity of the map f ; in the non-linear case, the expressions quickly become more complicated. One could also attempt to use this approach to study the case where $\lambda_1 \neq \lambda_2$, in which Moran’s theorem cannot be used to calculate $\dim_H C$; a complete analysis of this case is still elusive, and the problem remains open.

Lecture 29

a. The non-conformal case. In the specific case $\lambda_1 = \lambda_2$, the map f given in (86) is *conformal*; the amount of contraction along the stable manifold (the disc) is the same in all directions, as is the amount of expansion along the unstable manifold (which in this case is only one-dimensional, so there can only be one rate of expansion). In this case, we saw last time that the Hausdorff dimension of the attractor Λ is given by (90).

For $\lambda_1 \neq \lambda_2$, we are in the non-conformal case, which is much more difficult. This case was studied by the German mathematician Hans Bothe, who considered a more general class of maps f , in which the functions $\cos \theta$ and $\sin \theta$ in (86) are replaced by arbitrary periodic functions $z_1(\theta)$ and $z_2(\theta)$, which changes the geometry of how the image $f(P)$ wraps around the z -axis. Bothe obtained a general result for “typical” functions z_1 and z_2 , but it was not until 1997 that the Hungarian mathematician Károly Simon proved that \sin and \cos belong to this “typical” class, and established that for the Smale–Williams solenoid Λ with $\lambda_1 < \lambda_2 < 1/8$, we have¹⁷

$$\dim_H \Lambda = 1 + \frac{\log 2}{-\log \lambda_2}.$$

Somewhat surprisingly, the smaller value λ_1 , which corresponds to a direction of faster contraction, does not affect the Hausdorff dimension of the attractor!

b. The attractor for the FitzHugh–Nagumo map. Returning to the FitzHugh–Nagumo map (76), we recall that for a particular range of values of A , the map f has a trapping region R , as shown in Figure 41. This ensures the existence of an attractor $\Lambda \subset R$ as in (85), and it is natural to ask what features Λ shares with the Smale–Williams solenoid, since both are attractors.

For example, we mentioned at the end of Lecture 28(a) that the Smale–Williams solenoid contains infinitely many periodic points, and what is more, the set of periodic points is dense in the attractor. In some sense, this runs parallel to the fact that for a Morse–Smale system, the set of periodic points (which in that case is finite) attracts all orbits, and hence all long-term behaviour takes place near periodic points. Does this also hold true for the attractor for the FitzHugh–Nagumo map?

Or we may recall the hyperbolic structure found at each point of the Smale–Williams solenoid, and ask if a similar structure can be found at each point of the attractor for the FitzHugh–Nagumo map. Is it true that for every $x \in \Lambda$ there exist vectors v_1 and v_2 such that the map f is asymptotically expanding in the direction given by v_1 , and asymptotically contracting in the direction given by v_2 ?

Computer simulations suggest that both of the above questions are answered in the affirmative; however, no rigorous proofs are known.

¹⁷This was later extended to include all $\lambda_2 < 1/2$ by Jörg Schmeling.

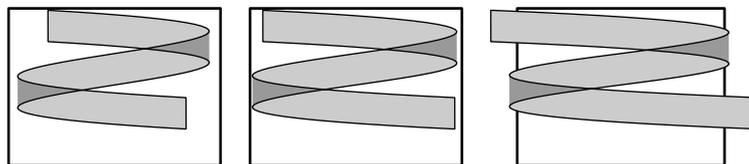


FIGURE 47. Escape from the trapping region.

c. From attractors to horseshoes. Figure 47 shows how the image $f(R)$ of the trapping region changes as A increases. The “arms” of the image become longer and longer, until when A is large enough, $f(R)$ is no longer contained in R , as in the third picture. Thus R is no longer a trapping region, and the closest we have to an analogue of the attractor Λ is the set of all points whose trajectories remain within R .

Rather than try to describe this set for the FitzHugh–Nagumo map, which leads to rather complicated geometric considerations, we begin by considering a model case in which the picture is much cleaner. Consider a map f in \mathbb{R}^2 which acts on the square $R = [0, 1] \times [0, 1]$ as shown in Figure 48; first the square is squeezed in the vertical direction by a factor of $\lambda < 1/2$ and stretched in the horizontal direction by a factor of $\mu > 2$, after which it is bent so that $f(R) \cap R$ consists of two rectangles, each of width 1 and height λ .

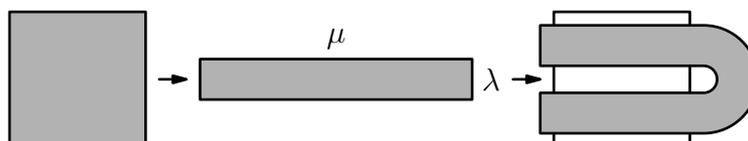


FIGURE 48. A horseshoe.

Notice that significant portions of the square R are mapped to the area outside R ; this is reminiscent of the interval maps we considered in the first few lectures, under which some points could only be iterated a few times because their trajectories were carried outside the domain of definition of the map. It is also similar to the behaviour we observed in Lecture 25(c) for the family of logistic maps with $c < -2$, for which some (indeed, almost all) trajectories escape to infinity, and where we were able to describe the set of points with bounded trajectories as a Cantor set in the interval $[-p_2, p_2]$.

Thus we ask the same question for the horseshoe map f shown in Figure 48; which points in R can be iterated infinitely often? That is, what is the set of points with trajectories which remain in R for all time?

Figure 49 duplicates Figure 48, but highlights the parts of R whose image lies in R ; in particular, we see that the set of points for which f can be iterated at least twice is the union of the two rectangles R_1 and R_2 , each of which has height 1 and width μ^{-1} .

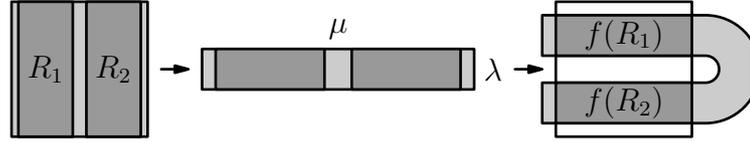


FIGURE 49. Points with two iterates.

So which points can be iterated three times? In order to have $f^2(x) \in R$, we must have $f(x) \in R_1 \cup R_2$; Figure 50 shows the set of points whose image lands in R_1 or R_2 . Observe that

$$\begin{aligned} f^{-1}(R_1) &= R_{11} \cup R_{21}, \\ f^{-1}(R_2) &= R_{12} \cup R_{22}, \end{aligned}$$

and that $R_{i_1 i_2} \subset R_{i_1}$ for all $i_1, i_2 \in \{1, 2\}$. Continuing this process, we see that the set of points which have n images in R is the union of 2^n rectangles $R_{i_1 \dots i_n}$, each of width μ^{-n} and height 1, which are characterised by

$$(91) \quad R_{i_1 \dots i_n} = R_{i_1} \cap f^{-1}(R_{i_2}) \cap \dots \cap f^{-(n-1)}(R_{i_n}).$$

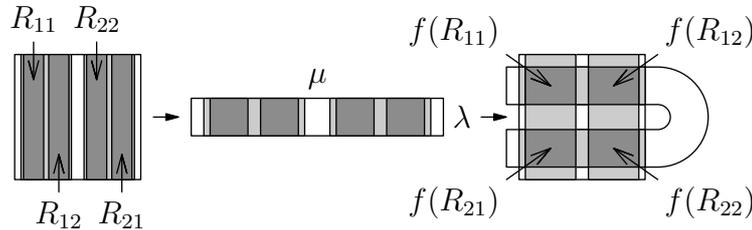


FIGURE 50. Points with three iterates.

Letting n go to infinity, we see that the non-escaping set is given by

$$(92) \quad \Gamma^+ = \{x \in R \mid f^n(x) \in R \ \forall n \geq 0\} = C_{\mu^{-1}} \times [0, 1],$$

where $C_{\mu^{-1}} \subset [0, 1]$ is a Cantor set with both ratio coefficients equal to μ^{-1} ; in particular, this implies that

$$\dim_H \Gamma^+ = \dim_H C_{\mu^{-1}} + 1 = \frac{\log 2}{\log \mu} + 1.$$

The story so far has largely been a retelling of a familiar tale from two earlier examples, the piecewise linear interval maps and the logistic maps with escape. There is a twist in the plot, however; unlike either of those two maps, the horseshoe map f is one-to-one, and hence invertible on its image. Consequently, we are interested in points for which both the forward *and backward* trajectories remain in R . That is, we are also interested in

$$\Gamma^- = \{x \in R \mid f^n(x) \in R \ \forall n \leq 0\}.$$

Since $f(R)$ does not cover the entire square R , we see that the only points in R with any pre-images at all are those which lie in $f(R_1)$ or $f(R_2)$ (see Figure 49). Write $S_i = f(R_i)$; then the set of points in R with one backwards iterate is the union of the two rectangles S_1 and S_2 , each of which has width 1 and height λ .

The set of points with two backwards iterates is

$$f(f(R) \cap R) \cap R = f(S_1 \cup S_2) \cap R = S_{11} \cup S_{12} \cup S_{22} \cup S_{21},$$

as shown in Figure 51, where $S_{i_1 i_2} = S_{i_1} \cap f(S_{i_2})$. Continuing, we obtain rectangles $S_{i_1 \dots i_n}$ of width 1 and height λ^n characterised by

$$(93) \quad S_{i_1 \dots i_n} = S_{i_1} \cap f(S_{i_2}) \cap \dots \cap f^{(n-1)}(S_{i_n}),$$

and we see that

$$\Gamma^- = [0, 1] \times C_\lambda,$$

where $C_\lambda \subset [0, 1]$ is a Cantor set with both ratio coefficients equal to λ ; hence

$$\dim_H \Gamma^- = 1 + \frac{\log 2}{-\log \lambda}.$$

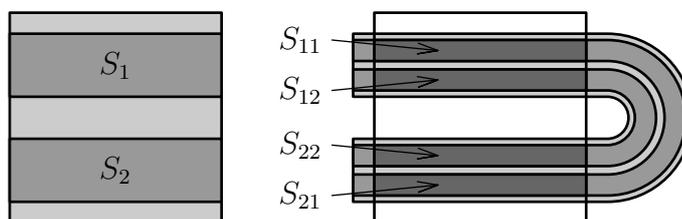


FIGURE 51. Points with two backwards iterates.

Now we can describe the set of points for which all forward and backward iterates remain in R ; this set is given by

$$\begin{aligned} \Gamma &= \{x \in R \mid f^n(x) \in R \forall n \in \mathbb{Z}\} \\ &= \Gamma^+ \cap \Gamma^- \\ &= (C_{\mu^{-1}} \times [0, 1]) \cap ([0, 1] \times C_\lambda) \\ &= C_{\mu^{-1}} \times C_\lambda, \end{aligned}$$

and we have

$$(94) \quad \begin{aligned} \dim_H \Gamma &= \dim_H C_{\mu^{-1}} + \dim_H C_\lambda \\ &= \log 2 \left(\frac{1}{\log \mu} + \frac{1}{-\log \lambda} \right). \end{aligned}$$

Observe that $\log \lambda$ is the Lyapunov exponent in the vertical (contracting) direction, and that $\log \mu$ is the Lyapunov exponent in the horizontal (expanding) direction; the notion of entropy can also be defined in this case, and is equal to $\log 2$. Thus we have another example of the relationship between dimension, entropy, and Lyapunov exponents.

The set Γ is known as the *Smale horseshoe*; it is a better-behaved relative of the non-escaping set for the FitzHugh–Nagumo map with large A . It has many important dynamical properties, among them the fact that periodic points of f are dense in Γ , and that every point of Γ is hyperbolic.

Perhaps even more important than these properties is the following: if we perturb the map f slightly, and consider a map g in the plane which is “close to” f in some appropriate sense, then the entire qualitative description we have given remains true for g as well. The rectangle R and the basic sets $R_{i_1 \dots i_n}$, $S_{i_1 \dots i_n}$ will be replaced with slightly less regular regions (“distorted” rectangles), but Γ will still be the direct product of two Cantor sets, and will lie in a neighbourhood of the original Γ . In particular, the map g is still chaotic, just as f is.

Lecture 30

a. Symbolic dynamics on the Smale horseshoe. In the previous lecture, we described the geometry of the set Γ of points which do not escape from R , by observing that Γ is the direct product of the Cantor sets C_λ and $C_{1/\mu}$. The considerations which led us to this conclusion can also be used to obtain information about the dynamics of the map $f: \Gamma \rightarrow \Gamma$.

Recall that the set of points in R whose image is also in R is $R_1 \cup R_2$, where R_1 and R_2 are the rectangles shown in Figure 49. It follows that $\Gamma \subset R_1 \cup R_2$, and so given a point $x \in \Gamma$, we have $f^n(x) \in R_1 \cup R_2$ for all n . Define a sequence $\omega^+ = (i_1, i_2, i_3, \dots) \subset \Sigma_2^+$ such that $f^{n-1}(x) \in R_{i_n}$ for all $n \geq 1$; in this manner we can associate to each point $x \in \Gamma$ a sequence in the symbolic space Σ_2^+ .

Upon the further observation that $f^{(n-1)}(f(x)) = f^n(x) \in R_{i_{n+1}}$, we see that the point $f(x)$ is coded by the sequence $\sigma(\omega^+) = (i_2, i_3, i_4, \dots)$, and it looks like we are well on our way to establishing a topological conjugacy between the map $f: \Gamma \rightarrow \Gamma$ and the shift $\sigma: \Sigma_2^+ \rightarrow \Sigma_2^+$.

Once again, however, there is a twist in the plot. If we follow the recipe from our previous encounters with symbolic dynamics, then the sequence ω^+ should determine the point x uniquely as the only point in the infinite intersection

$$E^+(\omega^+) = R_{i_1} \cap R_{i_1 i_2} \cap \dots \cap R_{i_1 \dots i_n} \cap \dots$$

However, as Figures 49 and 50 suggest, each rectangle $R_{i_1 \dots i_n}$ has width μ^{-n} and height 1; this means that the intersection $E^+(\omega^+)$ is a vertical line, rather than a single point! In fact, any point on the vertical line $E^+(x) = E^+(\omega^+)$ passing through x has a forward trajectory which is coded by the same sequence ω^+ .

We see, then, that the coding of the forward trajectory is not enough to determine x uniquely. We also need to code the backward trajectory $\{f^{-n}(x)\}_{n=0}^\infty$, which we do via the sets $S_{i_1 \dots i_n}$, as indicated in (93). Thus we obtain a sequence $\omega^- = (j_1, j_2, \dots)$ such that $f^{-(n-1)}(x) \in S_{j_n}$ for all $n \geq 1$. Once again, ω^- does not determine x uniquely; we have

$$E^-(\omega^-) = S_{j_1} \cap S_{j_1 j_2} \cap \dots \cap S_{j_1 \dots j_n} \cap \dots,$$

where each rectangle $S_{j_1 \dots j_n}$ has width 1 and height λ^n , and so the intersection $E^-(x) = E^-(\omega^-)$ is the horizontal line through x .

Although neither the forward itinerary ω^+ nor the backward itinerary ω^- of x is by itself enough to determine x uniquely, the combination of the two does suffice. Indeed, the vertical line $E^+(x)$ and the horizontal line $E^-(x)$ meet in a single point, x itself. Thus if we denote by $\Sigma_2 = \{1, 2\}^{\mathbb{Z}}$ the space of all *doubly* infinite sequences of ones and twos, we can define a coding map

$$\begin{aligned} h: \quad \Sigma_2 &\rightarrow \Gamma, \\ \omega = (\omega^-, \omega^+) &\mapsto E^-(\omega^-) \cap E^+(\omega^+), \end{aligned}$$

where (ω^-, ω^+) denotes the concatenation $(\dots, j_2, j_1, i_1, i_2, \dots)$.

Evaluating the coding map h for a given doubly infinite sequence $\omega = (\dots, i_{-2}, i_{-1}, i_0, i_1, i_2, \dots)$ requires us to decide where the “centre” of the sequence is; that is, we must decide once and for all which integer k will be such that $\omega^+ = (i_k, i_{k+1}, \dots)$ and $\omega^- = (i_k, i_{k-1}, \dots)$. If we take $k = 0$, then $\omega^+ = (i_0, i_1, \dots)$ and $\omega^- = (i_{-1}, i_{-2}, \dots)$, so x is the unique point in Γ for which $f^n(x) \in R_{i_n}$ for all $n \in \mathbb{Z}$.

Applying the shift map $\sigma: \Sigma_2 \rightarrow \Sigma_2$ to a sequence ω shifts the “centre” to the right by one; it is common to demarcate ω^+ and ω^- with a dot, so that we may write

$$\sigma(\dots i_{-2}i_{-1}.i_0i_1i_2\dots) = \dots i_{-2}i_{-2}i_{-2}i_0.i_1i_2\dots;$$

in particular, the construction of the coding map h makes it transparent that the following diagram commutes:

$$(95) \quad \begin{array}{ccc} \Sigma_2 & \xrightarrow{\sigma} & \Sigma_2 \\ \downarrow h & & \downarrow h \\ \Gamma & \xrightarrow{f} & \Gamma \end{array}$$

The horizontal lines $E^-(x)$ and the vertical lines $E^+(x)$ have a dynamical meaning. f contracts each vertical line uniformly by a factor of λ , and so given $y \in E^+(x)$, we have

$$d(f^n(x), f^n(y)) = \lambda^n d(x, y) \xrightarrow{n \rightarrow +\infty} 0.$$

Thus $E^+(x)$ is the stable direction through x for the map f . Similarly, observe that f^{-1} contracts horizontal lines uniformly by a factor of $1/\mu$, and so for $y \in E^-(x)$, we have

$$d(f^{-n}(x), f^{-n}(y)) = \mu^{-n} d(x, y) \xrightarrow{n \rightarrow +\infty} 0.$$

Thus $E^-(x)$ is the stable direction through x for the map f^{-1} , and the *unstable* direction through x for the map f .

We see from all this that the dynamics of the symbolic space encode the hyperbolic structure of the horseshoe map; the stable direction for f through a point x consists of those points whose forward itineraries eventually agree with the forward itinerary of x , as given by the sequence ω^+ , and the unstable direction is given similarly, as those points whose backwards itineraries eventually agree with ω^- .

EXERCISE 19. Using the fact that $f: \Gamma \rightarrow \Gamma$ is conjugate to $\sigma: \Sigma_2 \rightarrow \Sigma_2$, show that f has 2^n points of period n (that is, fixed points of f^n) lying in Γ , for all $n \geq 1$.

b. Variations on the horseshoe map. One can show that the coding map h is a homeomorphism, and so the symbolic space Σ_2 captures the topology of Γ as well as the dynamics of f . However, as was the case with the interval maps we discussed at the beginning of the course, the symbolic

approach does not capture all the quantitative geometric information about the horseshoe Γ . In the first place, neither Σ_2 nor the map σ depends on the value of the parameters λ, μ , while varying these quantities certainly changes the set Γ , and in particular, changes its Hausdorff dimension. Furthermore, as we mentioned at the end of the previous lecture, we can perturb the map f slightly and still carry out the whole procedure, obtaining a non-linear horseshoe which is still modeled by Σ_2 .

Nevertheless, in each of these cases, the fact that the map $f: \Gamma \rightarrow \Gamma$ is modeled by the two-sided shift $\sigma: \Sigma_2 \rightarrow \Sigma_2$ still gives us some tools with which to work. For example, the Non-Uniform Mass Distribution Principle allows us to get information about $\dim_H \Gamma$ by examining various measures on Γ ; the conjugacy with the symbolic case means that any shift-invariant measure on Σ_2 determines an f -invariant measure on Γ . Thus we have Bernoulli measures, Markov measures, etc. even on a non-linear horseshoe.

The conjugacy with the shift also lets us find many other interesting f -invariant subsets of the horseshoe Γ ; in particular, given a 2×2 transition matrix A , the subshift of finite type Σ_A is shift-invariant, and so $h(\Sigma_A) \subset \Gamma$ is f -invariant. Thus the horseshoe contains all manner of intricate fractals which are preserved by the dynamics of f .

Finally, we can consider the horseshoe-like map $g: R \rightarrow \mathbb{R}^2$ which maps the square R to the region shown in Figure 52. Following the same line of argument as for f , one finds that the set of non-escaping points Γ' is homeomorphic to $\Sigma_3 = \{1, 2, 3\}^{\mathbb{Z}}$, and so the restriction of f to Γ' is conjugate to the full shift on *three* symbols. If λ and μ are the contraction and expansion ratios in the vertical and horizontal directions, respectively, then Γ' is the direct product of Cantor sets in the interval with ratio coefficients λ and μ^{-1} , and so

$$\dim_H \Gamma' = \frac{\log 3}{-\log \lambda} + \frac{\log 3}{\log \mu} = \log 3 \left(\frac{1}{-\log \lambda} + \frac{1}{\log \mu} \right).$$

This has the same form as (94), giving Hausdorff dimension in terms of entropy and Lyapunov exponents, but the entropy term has been replaced by $\log 3$. In general, we could consider a horseshoe-like map with n branches, and would find the same formula, with the entropy replaced by $\log n$. This reflects the fact that entropy somehow measures the *complexity* of the system; for these maps, the system is entirely linear on the invariant horseshoe, and so all the complexity comes from how many times that set is folded back into itself by the map.

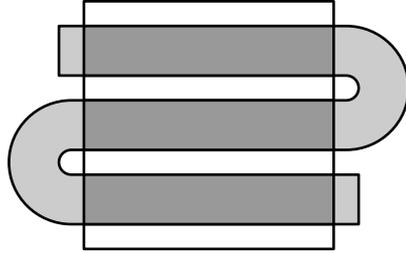


FIGURE 52. A horseshoe with three branches.

Lecture 31

a. Transient and persistent chaos. If we choose a typical point $x = (x_1, x_2)$ in the Smale horseshoe Γ and observe its trajectory $f^n(x) = (x_1^{(n)}, x_2^{(n)})$ by plotting either $x_1^{(n)}$ or $x_2^{(n)}$ as a function of n , we see a chaotic signal which persists as $n \rightarrow \infty$, without ever settling down into periodic behaviour.

However, even though the Smale horseshoe Γ is the largest f -invariant set in the square R , it has zero Lebesgue measure, and so if we simply choose a point from the square R “at random”, it will lie outside of Γ with probability 1. In particular, after some number of iterations, the trajectory will leave the square; as long as it remains in the square, it appears chaotic, but once it leaves R , all bets are off.

We have not defined the map f outside of R , and there are various ways in which it could be extended. Depending on how f is extended to \mathbb{R}^2 and where the trajectory leaves the square, there are a number of possible behaviours we might observe; for example, it might be attracted to a fixed point outside R , or approach a periodic orbit outside R , or it might follow such an orbit for a while (and hence appear quite regular) before re-entering R and going through a period of apparent chaos as it follows the horseshoe again and is eventually repelled. These all fall under the general heading of *intermittent chaos* (or *transient chaos*); such behaviour is characteristic of typical trajectories in the neighbourhood of a horseshoe.

Thus there are two general sorts of chaotic behaviour that we may observe, which more or less correspond to the last two examples we have studied. Persistent chaos, in which the trajectory never settles down to a periodic orbit, is indicative of the presence of an attractor, and hence of a trapping region containing the trajectory’s starting point. In intermittent chaos, on the other hand, the chaotic period eventually ends and the trajectory becomes more regular; this is indicative of the presence of a horseshoe, and occurs when the trajectory enters a region R which is mapped into itself in a manner resembling the horseshoe map.

b. Continuous-time systems. Up to this point, all the systems we have studied have been given in terms of a map from some domain to itself. Such systems are known as *discrete-time* systems, since “time” moves in discrete increments, corresponding to how many times the map has been iterated. For the remainder of this course, we will focus our attention on *continuous-time* systems, which are specified by ordinary differential equations (ODEs) rather than maps. Before examining the connections between the discrete and continuous-time cases, we will point out a striking difference between the two—for discrete-time systems, we were able to produce chaotic behaviour in maps of any dimension, while continuous-time systems cannot be chaotic in fewer than three dimensions.

We begin by recalling some of the basic notions regarding differential equations. The system of ODEs specifying a continuous-time system may be written as a single ODE for a vector-valued function \mathbf{x} . That is, we consider functions $\mathbf{x}: (a, b) \rightarrow \Omega \subset \mathbb{R}^n$ which solve the following equation:

$$(96) \quad \dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t)),$$

where $\mathbf{F}: \Omega \rightarrow \mathbb{R}^n$ is the function defining the ODE on the domain¹⁸ Ω . For the time being, we consider the case $n = 2$; if we write $\mathbf{x} = (x_1, x_2)$ and $\mathbf{F}(\mathbf{x}) = (F_1(\mathbf{x}), F_2(\mathbf{x}))$, then (96) may be written coordinate-wise as

$$(97) \quad \dot{x}_1(t) = F_1(x_1(t), x_2(t)),$$

$$(98) \quad \dot{x}_2(t) = F_2(x_1(t), x_2(t)).$$

If the function \mathbf{F} is continuously differentiable—that is, if all the partial derivatives $\frac{\partial F_i}{\partial x_j}$ exist and are continuous—then the standard existence and uniqueness theorem from the theory of ODEs implies that the system has a unique solution on some interval $-\varepsilon < t < \varepsilon$. By gluing together the solutions on these small intervals, the solution may be extended to some maximal interval $a < t < b$; the endpoints a and b are either infinite or the points at which $\mathbf{x}(t)$ reaches the boundary of Ω . If Ω is unbounded, then it is possible for a solution $\mathbf{x}(t)$ to reach infinity in finite time; this phenomenon can be avoided by requiring, for example, that $\|\mathbf{F}(\mathbf{x})\|$ be bounded.

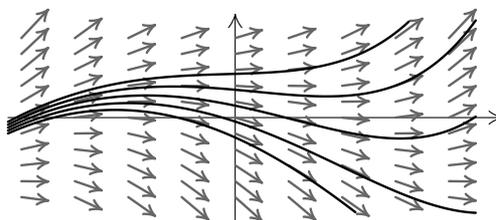


FIGURE 53. Some integral curves for a vector field.

The ODE (96) can also be given a geometric interpretation. The vector $\mathbf{F}(\mathbf{x})$ specifies a direction and length; placing this vector so that it originates at the point \mathbf{x} , we obtain a *vector field* in the plane. Solutions of the ODE (96) correspond to *integral curves* of the vector field; that is, curves whose tangent vector at each point is exactly the element of the vector field at that point, as shown in Figure 53.

Interpreting an ODE in terms of its associated vector field can be of great utility when it comes to answering certain global questions regarding fixed points, as results from index theory, etc. can be brought into play.

The integral curves $\gamma: (a, b) \rightarrow \mathbb{R}^2$ which represent solutions of the ODE admit various parametrisations, as do all curves; however, the requirement

¹⁸To be precise, we take our domain Ω to be a simply connected open subset of \mathbb{R}^n , that is, an open subset such that every loop in Ω can be contracted to a point without ever leaving Ω .

that the tangent vector $\gamma'(t)$ have the same length as the vector $\mathbf{F}(\gamma(t))$ fixes a unique natural parametrisation. Writing $\gamma_{\mathbf{x}}$ for the unique curve with $\gamma_{\mathbf{x}}(0) = \mathbf{x}$, we see that $\gamma_{\mathbf{x}}(t)$ is the point in \mathbb{R}^2 which the system reaches after time t , if the initial condition at time 0 was given by \mathbf{x} . The map which takes \mathbf{x} to $\gamma_{\mathbf{x}}(t)$ is called the *time- t map* of the system, and is denoted φ_t .

This defines a one-parameter family of maps $\varphi_t: \Omega(t) \rightarrow \Omega$, where $\Omega(t)$ consists of the points $\mathbf{x} \in \Omega$ such that $\gamma_{\mathbf{x}}(\tau)$ remains in Ω for $0 \leq \tau \leq t$. If $\Omega = \mathbb{R}^2$ and $\|\mathbf{F}(x)\|$ is bounded, then φ_t is defined on all of \mathbb{R}^2 .

Each of these maps corresponds to evolving the system forward (or backward) by the appropriate amount of time, and they are related to each other by the *group property*

$$(99) \quad \varphi_{t+s} = \varphi_t \circ \varphi_s = \varphi_s \circ \varphi_t;$$

that is, $\varphi_{t+s}(\mathbf{x}) = \varphi_t(\varphi_s(\mathbf{x})) = \varphi_s(\varphi_t(\mathbf{x}))$ for all $\mathbf{x} \in \mathbb{R}^2$ and $t, s \in \mathbb{R}$.

DEFINITION 66. A *flow* on \mathbb{R}^n is a one-parameter family of one-to-one differentiable maps $\varphi_t: \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $\varphi_0(\mathbf{x}) = \mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^n$ and (99) holds for all $s, t \in \mathbb{R}$.

Flows provide the third way of looking at continuous-time systems; the three descriptions of such a system in terms of an ODE, a vector field, and a flow are all equivalent, and which one is most suitable depends on the circumstances.¹⁹

To describe the trajectories of the solutions to an ODE, we begin (as always), by finding the fixed points. In the algebraic language of the flow φ_t , these are the points \mathbf{x} such that $\varphi_t(\mathbf{x}) = \mathbf{x}$ for all $t \in \mathbb{R}$; in the geometric language of vector fields, these are the points at which the vector fields vanish; and in the language of ODEs, these are the zeros of the function \mathbf{F} , the points at which $\mathbf{F}(\mathbf{x}) = \mathbf{0}$, or in terms of the coordinate functions,

$$(100) \quad F_1(x_1, x_2) = F_2(x_1, x_2) = 0.$$

Suppose for the present that (100) has only finitely many solutions $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$. Then near the fixed point $\mathbf{x}^{(i)}$, we have

$$(101) \quad \mathbf{F}(\mathbf{x}) = \mathbf{F}(\mathbf{x}^{(i)}) + D\mathbf{F}(\mathbf{x}^{(i)})(\mathbf{x} - \mathbf{x}^{(i)}) + \mathbf{r}(\mathbf{x}),$$

where $\mathbf{r}(\mathbf{x})$ is an error term of order $o(\|\mathbf{x} - \mathbf{x}^{(i)}\|)$. Because $\mathbf{x}^{(i)}$ is fixed, the ODE (96) becomes

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t)) = D\mathbf{F}(\mathbf{x}^{(i)})(\mathbf{x}(t) - \mathbf{x}^{(i)}) + \mathbf{r}(\mathbf{x}),$$

and in the new coordinates $\mathbf{v} = \mathbf{x} - \mathbf{x}^{(i)}$, we have

$$\dot{\mathbf{v}}(t) = D\mathbf{F}(\mathbf{x}^{(i)})\mathbf{v}(t) + o(\|\mathbf{v}\|).$$

¹⁹The relationship between discrete-time and continuous-time systems is clearest when we consider the latter in terms of flows. If we restrict our attention to integral values of s and t , then (99) reduces to (1), which reflects the fact that for a fixed value of t , say $t = 1$, the time- t map of a continuous-time system defines a discrete-time system.

Since the perturbation from the linear case is small in a neighbourhood of the fixed point, we may hope to describe the solutions of (96) near the fixed points—and in particular, determine the stability of those fixed points—by first describing the solutions of the linear system

$$\dot{\mathbf{v}}(t) = D\mathbf{F}(\mathbf{x}^{(i)})\mathbf{v}(t),$$

which depend on the real part²⁰ of the eigenvalues of $D\mathbf{F}(\mathbf{x}^{(i)})$. In two dimensions, there are only three non-degenerate possibilities:

- (1) Both eigenvalues have negative real part, in which case the trajectories $\mathbf{v}(t)$ are the curves shown in one of the three phase portraits in Figure 27. If the trajectories are those shown in Figure 27(a), the fixed point is called an *attracting focus*; if they are the curves shown in Figure 27(b) or (c), the fixed point is called a *attracting node*.
- (2) Both eigenvalues have positive real part, in which case the trajectories are the same curves as in the previous case, but move in the opposite direction, away from the fixed point instead of towards it. In this case the fixed point is either a *repelling focus* or a *repelling node*.
- (3) One eigenvalue is negative and the other positive; in this case the trajectories follow the curves shown in Figure 28, and the fixed point is called a *saddle*. The attracting and repelling directions (horizontal and vertical in the picture) correspond to the eigenvectors of $D\mathbf{F}(\mathbf{x}^{(i)})$.

In each of these cases, the behaviour of the non-linear system is qualitatively the same as the behaviour of the linear system, provided we consider a small enough neighbourhood of the fixed point. However, there are also various degenerate cases which may occur; for example, if both eigenvalues are purely imaginary, then the trajectories of the linear system are concentric circles around the fixed point, as shown in Figure 54. In this case, the non-linear effects may (or may not) qualitatively change the behaviour of the system in an arbitrarily small neighbourhood of the origin, and so we gain no information about the stability of the fixed point.

²⁰This is in contrast to the criteria for linear maps, where it was the absolute value of the eigenvalues which determined the stability. In fact, the eigenvalues of the time-1 map φ_1 are of the form e^λ , where λ is an eigenvalue of $D\mathbf{F}$, and we observe that e^λ lies on the unit circle if and only if λ lies on the imaginary axis.

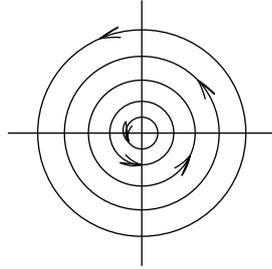


FIGURE 54. Trajectories around a centre.

Lecture 32

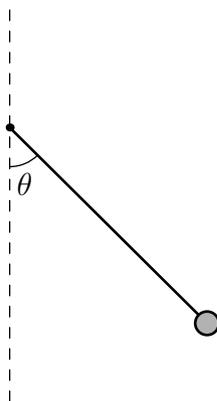


FIGURE 55. A pendulum.

a. The pendulum. One of the standard examples of a non-linear differential equation is the pendulum, a massless rod with one end fixed and a point mass at the other end. We let θ denote the angle by which the rod is displaced from the vertical, as shown in Figure 55. Neglecting effects of air resistance, the equation governing the motion of the pendulum is

$$(102) \quad \ddot{\theta} + a \sin \theta = 0,$$

where $a > 0$ is a parameter which depends on the length of the rod, the mass on the end, and the force of gravity.

We require *two* initial conditions, θ and $\dot{\theta}$, in order to specify a particular solution of (102). Thus the *phase space* of the pendulum is two-dimensional; if we write $x_1 = \theta$ and $x_2 = \dot{\theta}$, then (102) may be rewritten as

$$(103) \quad \begin{aligned} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= -a \sin x_1, \end{aligned}$$

or even more succinctly as $\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x})$, where $\mathbf{F}(x_1, x_2) = (x_2, -a \sin x_1)$, which allows us to use the language of the previous lecture.

The fixed points of (103) occur at $(k\pi, 0)$ for $k \in \mathbb{Z}$; however, since adding 2π to x_1 does not change the physical state of the system, there are actually only two fixed points. One of these corresponds to a pendulum hanging motionless, pointing straight down (when k is even); the other corresponds to a pendulum balancing on its pivot, pointing straight up (when k is odd).

Intuitively, we feel that the first of these is “stable”, while the second is “unstable”; we may attempt to confirm and clarify this intuition by linearising around the fixed points, as in the previous lecture. Upon doing so, we discover that at $\mathbf{x} = (0, 0)$, both eigenvalues of $D\mathbf{F}(\mathbf{x})$ are purely imaginary,

and so the linearised system is a rotation, while at $\mathbf{x} = (\pi, 0)$, one eigenvalue is positive and one is negative, and so $(\pi, 0)$ is a saddle.

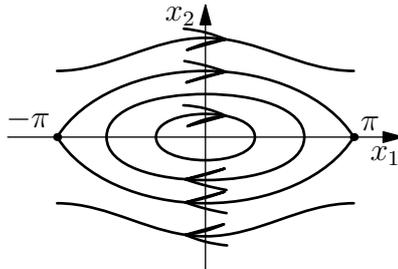


FIGURE 56. The phase portrait for a pendulum.

These results on the stability of the fixed points are reflected in Figure 56, which shows the phase portrait for (103). Observe that near $(0, 0)$, the phase portrait is as shown in Figure 54, while near $(\pi, 0)$ (equivalently, near $(-\pi, 0)$), it looks like Figure 28, although now the stable and unstable directions are no longer horizontal and vertical.

What do the various trajectories in the phase portrait correspond to in the physical system? Note that every trajectory intersects the vertical axis, which corresponds to the configuration where the pendulum is pointed straight down. The height at which the trajectory intersects the axis corresponds to the speed with which the pendulum is moving when it reaches this configuration.

If this speed is relatively small, then the pendulum will rise a little ways, eventually reach a maximum height (equivalently, a maximum angle of displacement), and then fall again, passing back through the line $x_1 = 0$ with the same speed it initially had, but in the other direction, and so on; this corresponds to the closed nearly-elliptical orbits around the origin in Figure 56. After some finite time T , the system is back where it started, and so for these trajectories we have $\mathbf{x}(t + T) = \mathbf{x}(t)$ for all $t \in \mathbb{R}$; these are *periodic orbits*.

If we start the pendulum off with a greater initial speed, then it will reach a greater height before reversing direction; it will also take longer to reach its maximum angle, and so the period T increases. At some critical initial speed, the maximum angle will be equal to π ; that is, the pendulum has enough energy that gravity will not pull it back before it reaches the top.²¹

In fact, if the pendulum reaches the top in finite time, then it will have some momentum left over (even if only a very small amount), which will be enough to carry it over the top, and into another complete rotation, so that it eventually reaches the bottom again, at which point it has exactly the same speed it began with. Thus there is some T such that $\mathbf{x}(t + T) = \mathbf{x}(t) + (2\pi, 0)$

²¹Many of us have tried to accomplish this on a swing set as a child.

for all $t \in \mathbb{R}$; these orbits are not periodic from the point of view of the system in the plane, but since changing x_1 by 2π does not change the physical system, they are periodic in terms of the pendulum itself.²²

What happens, though, if the pendulum has *exactly* enough energy to reach the top; enough that gravity will not stop it short, but not enough that it will have any momentum left over? Then it will move more and more slowly as time goes on, but will never stop (in which case it would reverse direction) or reach the top (in which case it would have some momentum left and would keep going). In a manner of speaking, it reaches the top, but in infinite time.

There are two trajectories in Figure 56 which correspond to this situation; one runs from $(-\pi, 0)$ to $(\pi, 0)$, passing through a point on the positive x_2 -axis, while the other runs from $(\pi, 0)$ to $(-\pi, 0)$, passing through a point on the negative x_2 -axis. Each of these curves is the unstable manifold for one fixed point, and the stable manifold for the other. Although both of these curves have finite length as curves in \mathbb{R}^2 , as trajectories it takes a solution of (103) an infinite amount of time to move the entire length of either one.

Trajectories such as these, which both originate and terminate in a fixed point, are known as *homoclinic* (if they begin and end in the same fixed point) or *heteroclinic* (if they begin and end in different fixed points), and often act as separatrices between regions of different qualitative behaviour. In this case, the two trajectories just described separate the orbits which “oscillate” (as we usually expect a pendulum to do) from the orbits which “spin”.

b. Two-dimensional systems. We now turn our attention to general two-dimensional continuous-time systems; as with discrete-time systems, a useful first step in analysing any given system is to find its fixed points and periodic orbits, and then classify them by stability. Fixed points were discussed in the previous lecture, so we now consider periodic orbits.

In a discrete-time system, a periodic orbit is just a finite collection of points, and it could have any of the stabilities available to a fixed point; stable, unstable, saddle, etc. In a continuous-time system, on the other hand, a periodic orbit is a closed curve, and so is of higher dimension than a periodic orbit for a map. This seemingly innocuous distinction is largely responsible for the absence of chaos in two-dimensional continuous-time systems, as we shall now see.

PROPOSITION 67. *Let $\gamma: \mathbb{R} \rightarrow \mathbb{R}^2$ be a periodic solution of the ODE (96), and suppose that γ is isolated; that is, that there exists some open neighbourhood $U \subset \mathbb{R}^2$ containing the curve γ such that U contains no other periodic orbits. Then γ is either stable and attracts every nearby trajectory, or it is unstable and repels every nearby trajectory.*

²²A more satisfactory model takes the phase space to be not \mathbb{R}^2 , but the cylinder; that is, \mathbb{R}^2 wrapped up so that the x_1 -axis becomes a circle.

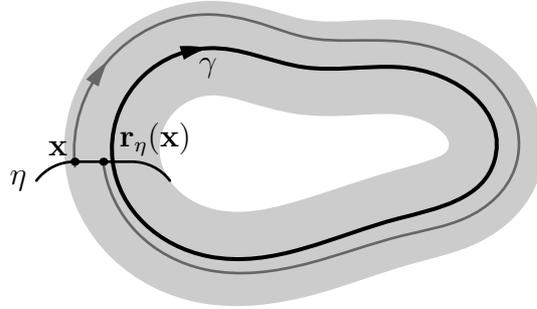


FIGURE 57. The Poincaré section for a transversal curve η .

PROOF. Fix a curve η which intersects the curve γ transversally, as shown in Figure 57. Then we define a map \mathbf{r}_η which acts on η , called the *Poincaré section for η* , as follows: each point $\mathbf{x} \in \eta$ defines a unique solution of (96), which remains near γ , and so eventually intersects η again; $\mathbf{r}_\eta(\mathbf{x})$ is defined to be the first point at which this intersection occurs.

Now suppose that $\mathbf{r}_\eta(\mathbf{x})$ is closer to γ than \mathbf{x} is. We iterate the Poincaré section by continuing the solution curve through $\mathbf{r}_\eta(\mathbf{x})$; this curve lies closer to γ than the solution curve through \mathbf{x} does, and in particular, $\mathbf{r}_\eta^2(\mathbf{x})$ lies closer to γ than $\mathbf{r}_\eta(\mathbf{x})$ does. This follows because solutions of (96) are unique, and so no two solution curves can cross each other; thus the second curve (starting at $\mathbf{r}_\eta(\mathbf{x})$) cannot cross the first curve (starting at \mathbf{x}) and escape. This is the piece of the argument which makes explicit use of the fact that we are working in two dimensions.

Observe that fixed points of the Poincaré section correspond to periodic solutions of the ODE; thus because γ is an isolated periodic orbit, the map \mathbf{r}_η has no fixed points except the intersection of γ and η . Since the trajectory of \mathbf{x} under the Poincaré section moves monotonically along η , it must converge to a fixed point of \mathbf{r}_η , which shows that the solution curve beginning \mathbf{x} approaches the periodic orbit γ . Similarly, any trajectory beginning close enough to γ approaches γ , and so the periodic orbit is stable.

A similar argument applies if $\mathbf{r}_\eta(\mathbf{x})$ is further away from γ than \mathbf{x} is, in which case γ is unstable, and these are the only two options. \square

One consequence of Proposition 67 is that a periodic orbit for a two-dimensional continuous-time system cannot be a “saddle”, and so the menagerie of possible local behaviours is tamer than it was for discrete-time systems.

There are a number of other results which hold in the two-dimensional continuous-time case, which are of a more global character.

PROPOSITION 68. *If γ is a periodic solution of (96), then the region enclosed by γ contains a fixed point of the system.*

SKETCH OF PROOF. One shows that this region is homeomorphic to a disc (this is the Schoenflies Theorem, a stronger version of the Jordan Curve

Theorem), and then applies the Brouwer Fixed Point Theorem to the time-1 map. \square

PROPOSITION 69. *Suppose that $\mathbf{F}: \Omega \rightarrow \mathbb{R}^2$ is divergence-free; that is,*

$$(104) \quad \frac{\partial F_1}{\partial x_1} + \frac{\partial F_2}{\partial x_2} \neq 0$$

for every $\mathbf{x} = (x_1, x_2) \in \Omega$. Then Ω contains no periodic orbits.

PROOF. Suppose γ is a periodic orbit in Ω , and let $\Omega' \subset \Omega$ be the region enclosed by γ , so that $\gamma = \partial\Omega'$. Then applying Stokes' Theorem, we have

$$\begin{aligned} \iint_{\Omega'} \frac{\partial F_1}{\partial x_1} + \frac{\partial F_2}{\partial x_2} dx_1 dx_2 &= \int_{\gamma} F_1 dx_2 - F_2 dx_1 \\ &= \int_0^T F_1(\mathbf{x}(t))\dot{x}_2(t) - F_2(\mathbf{x}(t))\dot{x}_1(t) dt \\ &= \int_0^T F_1(\mathbf{x}(t))F_2(\mathbf{x}(t)) - F_2(\mathbf{x}(t))F_1(\mathbf{x}(t)) dt \\ &= 0. \end{aligned}$$

However, since the original integrand is continuous and non-vanishing on Ω' , it must be either positive everywhere or negative everywhere, and hence the integral cannot be zero. This contradiction implies that no periodic orbits exist in Ω . \square

In order to state the general result regarding the absence of chaos in two dimensions, we adapt the notion of an ω -limit set to the continuous-time case.

DEFINITION 70. If $\gamma: \mathbb{R} \rightarrow \mathbb{R}^2$ is a trajectory in \mathbb{R}^2 with $\gamma(0) = \mathbf{x}$, then the ω -limit set of \mathbf{x} is

$$(105) \quad \omega(\mathbf{x}) = \{ \mathbf{p} \in \mathbb{R}^2 \mid \mathbf{p} = \lim_{n \rightarrow \infty} \mathbf{x}(t_n) \text{ for some sequence } t_n \rightarrow \infty \}.$$

The α -limit set $\alpha(\mathbf{x})$ is defined similarly, but with the requirement that $t_n \rightarrow -\infty$ instead.

EXERCISE 20. Show that each ω -limit set $\omega(\mathbf{x})$ is closed and φ_t -invariant for all t , and that the same is true of the set $\Omega = \bigcup_{\mathbf{x} \in \mathbb{R}^2} \omega(\mathbf{x})$ of all points in \mathbb{R}^2 which lie in some ω -limit set.

A complete description of the possible ω -limit sets for flows in \mathbb{R}^2 is given by the following theorem, whose proof we omit.

THEOREM 71 (Poincaré–Bendixson). *If γ is a bounded trajectory of a flow φ_t with initial condition $\gamma(0) = \mathbf{x}$, then one of the following occurs:*

- (1) $\omega(\mathbf{x})$ is a union of fixed points and heteroclinic (or homoclinic) orbits.
- (2) γ is a periodic orbit.
- (3) $\omega(\mathbf{x})$ is a periodic orbit to which γ converges.

The Poincaré–Bendixson theorem rules out the existence of chaos for two-dimensional flows by completely describing all the possible asymptotic behaviours, all of which are quite regular. Thus in order to observe chaos in a continuous-time system, we must look to higher dimensions, which we now proceed to do.

c. The Lorenz equations. For flows on \mathbb{R}^3 , we do not have the proper topological context to make the theorems in the previous section work; in particular, a periodic orbit γ does not need to be the boundary of a φ_t -invariant region homeomorphic to a disc. Thus life can be much more interesting in higher dimensions; to illustrate this fact, we study a particular system of ODEs in \mathbb{R}^3 , the *Lorenz equations*

$$(106) \quad \begin{aligned} \dot{x} &= -\sigma x + \sigma y, \\ \dot{y} &= rx - y - xz, \\ \dot{z} &= xy - bz. \end{aligned}$$

Here r (called the *Reynolds number*) is the leading parameter, while σ and b will be fixed at $\sigma = 10$ and $b = 8/3$. In the next lecture, we will examine the behaviour of solutions of (106) as r ranges from 0 to some number $R > 25$.

But first, a little history. The equations (106) were first considered in 1963 by Edward Lorenz, a meteorologist at M.I.T. who was studying the motion of the atmosphere. To a good approximation, the atmosphere is governed by the Navier–Stokes equations of fluid motion; however, these are enormously difficult to solve, and even an approximate numeric solution requires a powerful computer.

Four years earlier, Lorenz had begun using one of the more powerful computers available at that time²³ to simulate the motion of the atmosphere. More precisely, he considered a layer of fluid between two horizontal plates, where the lower plate is heated and the upper one is cooled. If the difference in temperature ΔT between the plates is small, then heat will flow by conduction from the bottom plate to the top plate, and the fluid will remain motionless, and this is a stable equilibrium configuration; as ΔT increases, this equilibrium configuration becomes unstable, and a small disturbance is enough to cause convection cells to form in the fluid, rotating vortices of fluid carrying warm fluid from the bottom plate to the top plate, and cooler fluid back down. As ΔT increases even further, these convection cells become unstable, and the fluid flow eventually becomes turbulent—that is, chaotic.

Initially, Lorenz considered a system of twelve equations in twelve variables, which were obtained as Fourier coefficients of the functions in the Navier–Stokes equations; upon being given the initial conditions, the computer would calculate the (approximate) trajectory of the system. One day, Lorenz wanted to take a closer look at a part of the previous day’s simulation, and so he entered as the initial condition the output from midway

²³Sixty computations a second!

through the calculated trajectory. To his surprise, the results of this simulation, which should have matched perfectly with the previous results, instead diverged quite quickly!

Lorenz soon realised what the problem was; the computer stored all initial data and intermediate calculations to an accuracy of 6 digits, but only printed out the first 3 digits. Thus when Lorenz entered the previous day's results, he unknowingly introduced a small error term, on the order of 10^{-4} . Rather than dying away or at least remaining small as the system evolved in time, this error term grew exponentially, until it became large enough to make the two trajectories appear completely unrelated.

After further investigation, Lorenz was able to replace the system of twelve ODEs with the system (106) of three ODEs which now bears his name; although this system does not capture all the details of the original system, it displays the same behaviour. The leading parameter r plays the role of ΔT ; for small values of r , as we will see, trajectories of the system are quite simple. However, when r becomes sufficiently large, (106) displays one of the hallmarks of chaos, *sensitive dependence on initial conditions* at every point; that is, any two nearby trajectories eventually diverge exponentially, as described above.

Lecture 33

a. Beyond the linear mindset. One of the simplest ODEs is the one-dimensional equation

$$(107) \quad \dot{x} = ax,$$

where $a \in \mathbb{R}$ is a constant; one quickly learns in any introductory course in differential equations that solutions of (107) have the form $x(t) = e^{at}x_0$, where $x_0 \in \mathbb{R}$ is the initial condition at $t = 0$. One later learns that this solution generalises to higher dimensions, and that for an $n \times n$ matrix A , the vector ODE

$$(108) \quad \dot{\mathbf{x}} = A\mathbf{x}$$

is solved by trajectories of the form

$$(109) \quad \mathbf{x}(t) = e^{At}\mathbf{x}_0,$$

where the exponential of a matrix X is defined by

$$e^X = \sum_{n=0}^{\infty} \frac{A^n}{n!} = I + X + \frac{X^2}{2} + \frac{X^3}{6} + \cdots .$$

This allows us to solve an arbitrary linear ODE; however, non-linear ODEs are much more difficult to deal with, and explicit closed-form solutions such as (109) are not generally available. Certain special cases do admit such solutions—for example, the equations of celestial mechanics derived from Newton’s law of universal gravitation can be solved explicitly (one often says they can be *integrated*)—and until the mid-twentieth century, the primary goal in the study of non-linear ODEs was to find clever solutions to particular equations, or classes of equations.

All systems for which explicit solutions could be found turned out to be Morse–Smale, and the prevailing philosophy held that this was the only possible sort of behaviour. Thanks to the work of Jacques Hadamard, Heinz Hopf, and Gustav Hedlund, geometers knew that the geodesic flow on a negatively curved manifold gave rise to a chaotic system; however, among physicists and mathematicians working in the field of differential equations, it was believed that a chaotic signal, such as that observed by Lorenz in the solution of (106), could only be the result of a random noise from the external environment of the system.

Lorenz’s examination of the system (106) demonstrated that a chaotic signal could be produced by a completely deterministic process; because he published his results in the *Journal of Atmospheric Science*, it took some time for the mainstream of the mathematics and physics communities to become aware of them. However, the word eventually spread, and the phenomenon of *deterministic chaos* is now an integral part of our understanding of the natural world.

b. Examining the Lorenz system. To begin our analysis of the system (106), we first find the fixed points at which $\dot{x} = \dot{y} = \dot{z} = 0$ and determine their stability via the Jacobian determinant $D\mathbf{F}$. It is helpful to observe that the system is invariant under the reflection in the z -axis given by

$$x \mapsto -x, \quad y \mapsto -y, \quad z \mapsto z,$$

and so solutions away from this axis come in symmetric pairs. In particular, any fixed point with $x \neq 0$ or $y \neq 0$ has a twin on the other side of the z -axis.

For $\sigma, b \neq 0$, the fixed point conditions $\dot{x} = 0$ and $\dot{z} = 0$ imply that $y = x$ and $z = x^2/b$. Thus the condition $\dot{y} = 0$ may be written as

$$rx - x - \frac{x^3}{b} = 0.$$

Thus the system always has a fixed point at $(0, 0, 0)$, and any other fixed point must satisfy

$$\frac{x^2}{b} = r - 1.$$

Hence for $0 < r < 1$, the only fixed point is the origin; at $r = 1$, a pitchfork bifurcation occurs, and two more fixed points appear:

$$p_1 = \left(\sqrt{b(r-1)}, \sqrt{b(r-1)}, r-1 \right),$$

$$p_2 = \left(-\sqrt{b(r-1)}, -\sqrt{b(r-1)}, r-1 \right).$$

To check the stability of these fixed points, we find the eigenvalues of the Jacobian matrix

$$D\mathbf{F}(\mathbf{x}) = \begin{pmatrix} -\sigma & \sigma & 0 \\ r - z & -1 & -x \\ y & x & -b \end{pmatrix}.$$

At the origin, we have $x = y = z = 0$, and so

$$D\mathbf{F}(\mathbf{0}) = \begin{pmatrix} -\sigma & \sigma & 0 \\ r & -1 & 0 \\ 0 & 0 & -b \end{pmatrix}$$

has $-b$ as one eigenvalue, and the other two eigenvalues are the roots of

$$\lambda^2 + (\sigma + 1)\lambda + \sigma(1 - r) = 0,$$

which is the characteristic polynomial of the 2×2 matrix in the upper-left corner of $D\mathbf{F}(\mathbf{0})$. The constant negative eigenvalue $-b$ with eigenvector $(0, 0, 1)$ indicates that the vertical direction is always contracting at the origin. For $0 < r < 1$, the other two eigenvalues are also negative, and the origin is an attracting fixed point; for $r > 1$, one of these eigenvalues is negative and the other is positive, and so the origin is a hyperbolic fixed point, with two contracting directions and one expanding direction. Trajectories moving along the latter direction are attracted to one of the two fixed points

p_1 and p_2 , which exist precisely when $r > 1$; this is the situation shown in Figure 58.

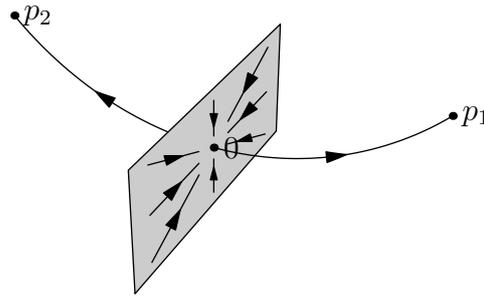


FIGURE 58. The fixed points for the Lorenz system when $r - 1$ is small.

Initially, all the eigenvalues of $D\mathbf{F}(p_1)$ and $D\mathbf{F}(p_2)$ are real and negative; as r increases, two of the eigenvalues become complex. p_1 and p_2 remain stable, and in particular, still attract the trajectory along the unstable manifold from the origin, but now trajectories approach these fixed points along the spirals shown in Figure 59.

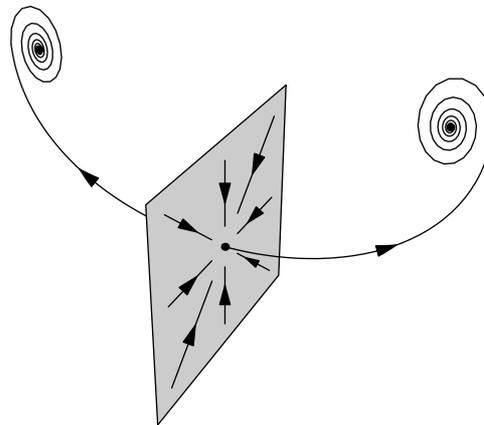


FIGURE 59. Changes in the behaviour of orbits as r increases.

As r increases still further, the spirals enlarge, and the trajectory along the unstable manifold from the origin takes longer and longer to approach p_1 or p_2 , as shown in Figure 60.

Through all of this, the system defined by (106) is Morse–Smale; all trajectories which begin on the stable separatrix approach the origin, and all other trajectories approach either p_1 or p_2 , depending on which side of the stable separatrix they begin on.

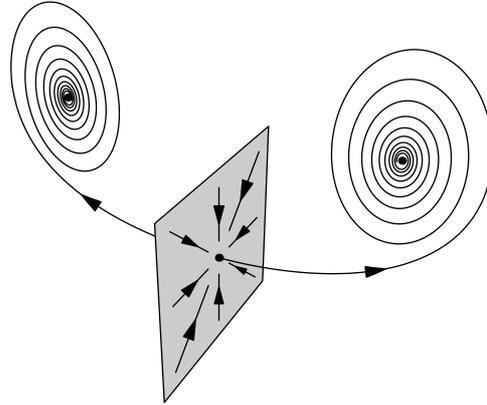


FIGURE 60. Weaker attraction and larger spirals as r increases still further.

Finally, at some critical value $r = r_0 \approx 13.926$, things change.²⁴ For $r < r_0$, the two halves of the unstable manifold from the origin are *heteroclinic* orbits which approach $\mathbf{0}$ as $t \rightarrow -\infty$ and p_1 or p_2 as $t \rightarrow \infty$; as r approaches r_0 , the spirals widen and these trajectories come closer and closer to the *stable* manifold from the origin, which is shown as a vertical plane in Figure 60.

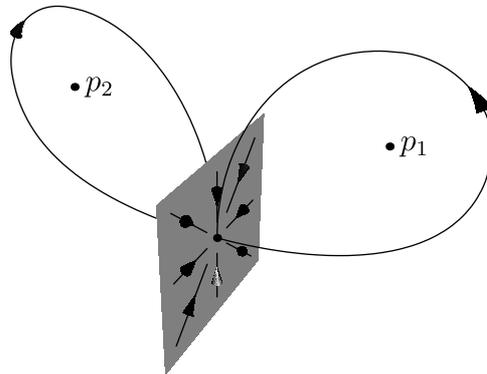


FIGURE 61. Appearance of homoclinic orbits at $r = r_0 \approx 13.926$.

For $r < r_0$, the trajectories along the unstable manifold come close to the stable manifold, but are eventually attracted to the fixed point p_1 or p_2 . When r reaches r_0 , this changes, and the trajectories become homoclinic; rather than spiraling in to p_1 or p_2 , they approach the origin from the vertical direction, as shown in Figure 61. For $r > r_0$, a completely new picture emerges; after circling one fixed point, the trajectories along the

²⁴Because r_0 is found as a result of numerical computations, we can only give an approximate value.

unstable manifold from the origin return past the stable separatrix²⁵ and approach the other fixed point, as shown in Figure 62. The fixed points p_1 and p_2 are still stable, but now an unstable periodic orbit appears around each one, separating trajectories which spiral in from trajectories which pass the separatrix and approach the other fixed point.

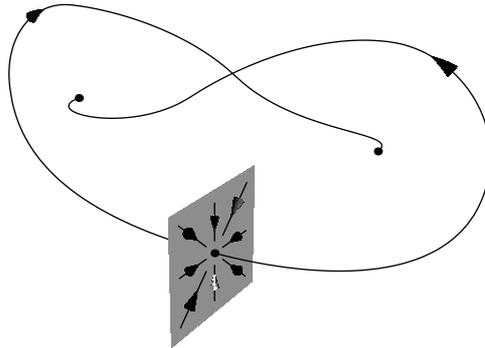


FIGURE 62. A change in behaviour when $r > r_0$.

Although this simplified picture suggests that the system may be Morse–Smale at this point, things are actually rather more complicated than that. We will see in the next lecture that there is a horseshoe hidden in this picture, and that as a result, the system is no longer Morse–Smale (since it has infinitely many periodic points) and in fact exhibits intermittent chaos.

²⁵Without intersecting it, which suggests that the geometry of the stable manifold is quite complicated, as indeed it is.

Lecture 34

a. Homoclinic orbits and horseshoes. Setting aside the Lorenz system for the moment, we examine some of the concepts introduced in the last few lectures in the discrete-time setting. In particular, we can define homoclinic and heteroclinic orbits for a discrete-time system just as we did for a flow; we will focus our attention on the former, for it turns out that the existence of a homoclinic orbit has striking consequences for the dynamics near the corresponding fixed point. In particular, such a fixed point is contained in a horseshoe, a closed invariant set on which the dynamics are topologically conjugate to the Smale horseshoe, or to the full shift on a two-sided symbolic space.

To see why this is so, let us first consider the Smale horseshoe itself, the maximal invariant set Γ for the map $f: R \rightarrow \mathbb{R}^2$ shown in Figure 48. As we saw in (95), the map $f: \Gamma \rightarrow \Gamma$ is topologically conjugate to the full shift $\sigma: \Sigma_2 \rightarrow \Sigma_2$. This implies that it has two fixed points, corresponding to the sequences $\omega_1 = (\dots 1.11 \dots)$ and $\omega_2 = (\dots 2.22 \dots)$.

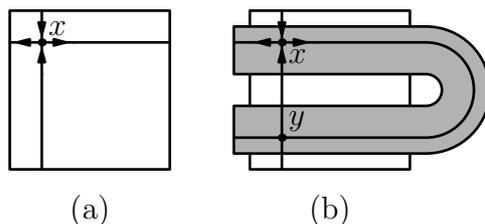


FIGURE 63. Finding a homoclinic orbit in a horseshoe.

Write $x = h(\omega_1)$, and recall that the (local) stable and unstable manifolds at x are the vertical and horizontal lines, respectively, passing through x , shown in Figure 63(a). Because these are invariant under the action of f (being defined in terms of asymptotic behaviour of trajectories), the unstable manifold contains the entire image of the horizontal line through x , which is the sideways “U”-shaped curve shown in Figure 63(b). In particular, it contains the point y , which therefore lies on both the stable and unstable manifolds of x ; it follows that the trajectory of y approaches x in both the forward and backward directions, and hence is a homoclinic orbit.²⁶

The situation here differs from the continuous-time case discussed in the previous lecture; here the intersection between the stable and unstable manifolds is transversal and happens in finite time, away from the fixed point, rather than being tangent and happening asymptotically as $t \rightarrow \infty$, as we saw for the Lorenz system.

²⁶This establishes the existence of a homoclinic orbit for both fixed points of f ; in fact, every *periodic* point also has stable and unstable manifolds, whose points of intersection lie on homoclinic orbits. Since periodic points are dense in the horseshoe, we see that homoclinic orbits are really quite ubiquitous.

This relationship between horseshoes and homoclinic orbits might remain a mere curiosity, were it not for the fact that the implication actually runs both ways. Setting aside the particular form of the map defining the Smale horseshoe, let us consider an arbitrary map f with a transverse homoclinic point y for a hyperbolic fixed point x , as shown in Figure 64(a).

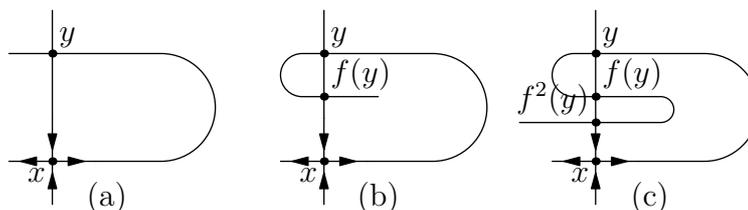


FIGURE 64. Consequences of a homoclinic point.

The unstable manifold is invariant, and so it also passes through $f(y)$, which lies on the stable manifold between y and x ; as Figure 64(b) shows, this forces the unstable manifold to fold back on itself. A similar argument applies to $f^2(y)$, as Figure 64(c) shows, and indeed to any $f^n(y)$; thus the unstable manifold folds back on itself infinitely often, and is stretched further and further between successive intersections with the unstable manifold. Not only that, but the stable manifold is left invariant by the action of f^{-1} , and must pass through all the points $f^n(y)$ for $n < 0$, so it folds back on itself infinitely often as well. The resulting picture is known as a *homoclinic tangle*, part of which is shown in Figure 65 (with slightly different notation – a better picture will follow eventually).

Observe that the picture in Figure 64 is somewhat idealised, and that the intersections between the stable and unstable manifolds may not be orthogonal initially; however, as $n \rightarrow \pm\infty$, the angle of intersection at $f^n(y)$ goes to $\pi/2$.

One sees immediately from Figure 65 that the geometric structure of the stable and unstable manifolds, and hence of the orbits of the system, is fantastically complicated; Henri Poincaré, who first discovered this picture in 1889 in conjunction with his work on the three-body problem, remarked,

One must be struck by the complexity of this shape, which I do not even attempt to illustrate. Nothing can give us a better idea of the complication of the three-body problem, and in general of all problems of dynamics for which there is no uniform integral.

The fundamental result which ties all this together is due to Smale himself, who showed that given a transverse homoclinic intersection for a fixed point x , one can find a rectangle R containing x and an integer n such

we choose a section of the plane $z = r - 1$ as shown in Figure 66; having fixed

$$E \subset \{ (x, y, z) \in \mathbb{R}^3 \mid z = r - 1 \},$$

we define a map $T: E \rightarrow E$ by

$$T(\mathbf{x}) = \varphi_{\tau(\mathbf{x})}(\mathbf{x}),$$

where $\tau(\mathbf{x}) > 0$ is minimal such that $\varphi_{\tau(\mathbf{x})}(\mathbf{x}) \in E$.²⁸ If \mathbf{x} is a fixed point of the flow, then set $T(\mathbf{x}) = \mathbf{x}$.

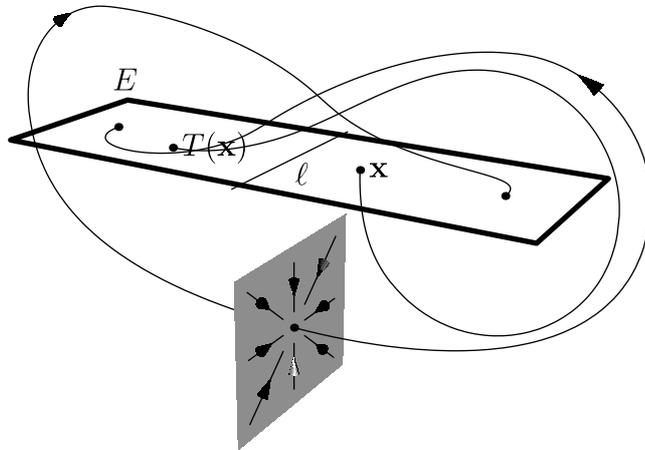


FIGURE 66. Defining the Poincaré section on E .

Since the domain of T is a two-dimensional space, the map T is in some ways simpler than the flow φ_t ; however, because not all trajectories which begin in E necessarily return to E , the return map T may not be defined on all of E . In particular, it may not be a continuous map on all of E , and so we pay a price for the simplification. We immediately encounter this difficulty when we consider points in the intersection of E with the stable manifold through the origin. These are both two-dimensional surfaces, which intersect in a curve ℓ , shown schematically in Figure 66 as a line; if \mathbf{x} is any point on ℓ , then the trajectory $\varphi_t(\mathbf{x})$ approaches $\mathbf{0}$, and never crosses E again, so $T(\mathbf{x})$ is undefined.

Consider then a point $\mathbf{x} \in E$ which lies just to the left of ℓ ; the trajectory $\varphi_t(\mathbf{x})$ will follow the stable manifold towards the origin for some ways before diverging and following the unstable manifold out towards near the edge of its range, eventually passing outside the edge of E , and then intersecting E again on the other side of ℓ , as shown in Figure 66.

As \mathbf{x} approaches ℓ , $T(\mathbf{x})$ approaches p_2 , and so any continuous extension of T to the line ℓ must have $T(\mathbf{x}) = p_2$ for all $\mathbf{x} \in \ell$; however, an identical

²⁸We remark in passing that the same construction may be used for discrete time maps; given a map $f: X \rightarrow X$ and a subset $A \subset X$, we can define the Poincaré first return map $f_A: A \rightarrow A$ in an analogous fashion.

argument requires $T(\mathbf{x}) = p_1$, and so the return map has no continuous extension to the line ℓ .

Despite this complication, James Kaplan and James Yorke were able to demonstrate the existence of a horseshoe-like structure for the map T . They showed that for carefully chosen regions $A, B, C, D \subset E$, the map acts as shown in Figure 67, and so one may once again carry out the procedure in the construction of the Smale horseshoe and obtain a Cantor-like set as the maximal invariant set for T . As trajectories of the Lorenz system come near this horseshoe, they exhibit chaotic behaviour for a finite period of time before being repelled, and so the system displays transient chaos.

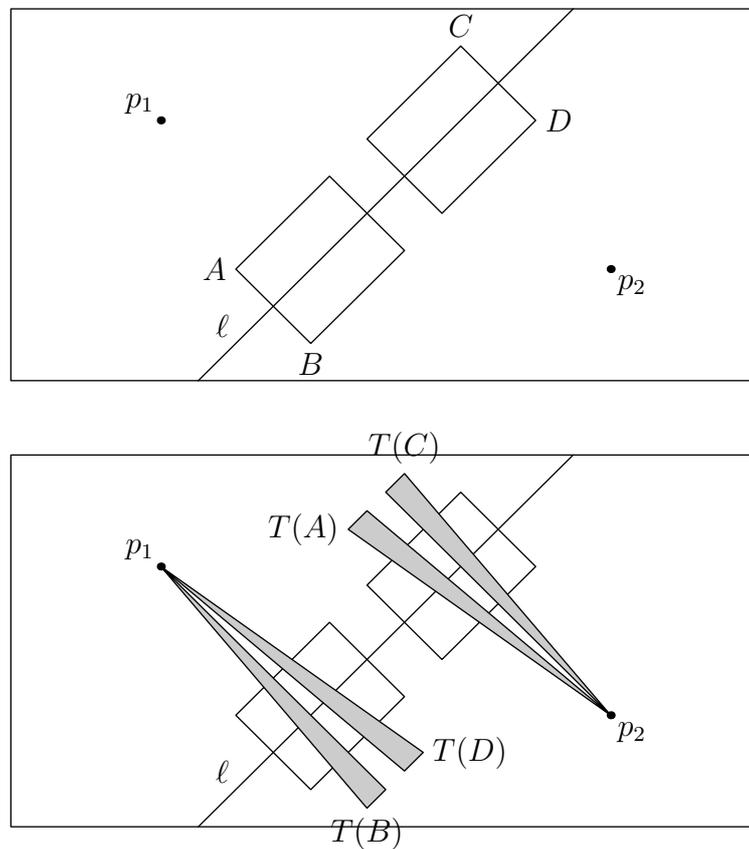


FIGURE 67. A horseshoe-like structure in the Poincaré section.

Lecture 35

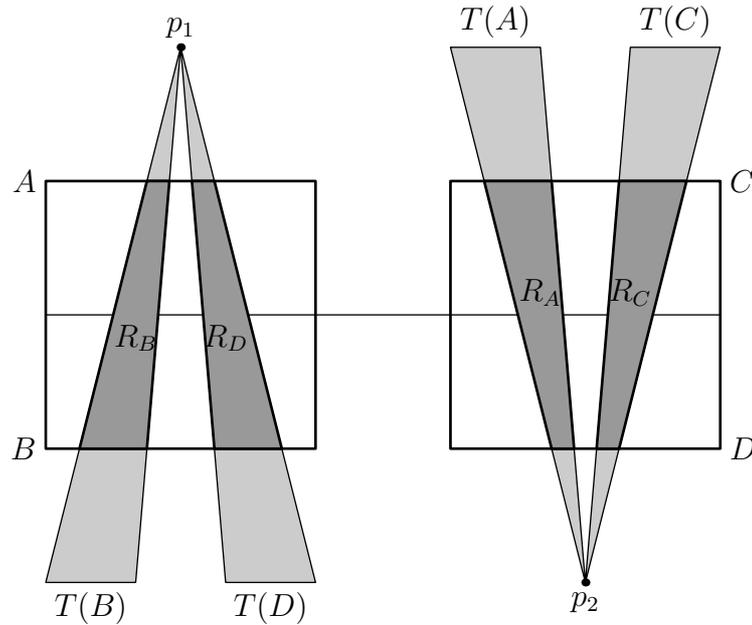


FIGURE 68. Points with one pre-image.

a. More about horseshoes in the Lorenz system. Let us take a closer look at Figure 67 and the associated Poincaré section $T: E \rightarrow E$. We are interested in the set of points whose trajectories remain in $R = A \cup B \cup C \cup D$; Figure 68 shows the topological structure of the part of the map which is significant for our purposes. The four darker trapezoids R_A, R_B, R_C , and R_D make up $f(R) \cap R$, the set of points in R with one pre-image in R ; these pre-images are shown in Figure 69, where the union of the four rectangles is the set of points in R with one *forward* image in R .

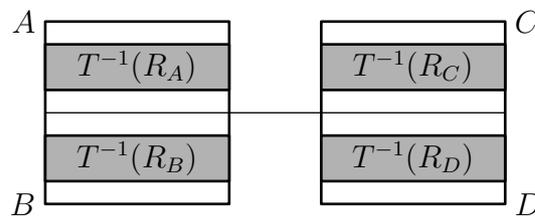


FIGURE 69. Points with one forward image.

This figure is schematic rather than quantitatively correct; it captures the topological behaviour which is observed in numerical simulations, even though the true regions do not have linear edges.

Taking the intersection of the four sets in Figure 69 with the four sets in Figure 68, we obtain eight trapezoids, whose union is the set of points in R with one forward and one backward image still in R . If we write, for example,

$$R_{A.CD} = T(A) \cap C \cap T^{-1}(D)$$

for the set of all points lying in C whose first iterate is in D and which are the image of some point in A , then the set of points with one forward and backward image in R is

$$\bigcup_{(i_{-1}.i_0i_1)} R_{i_{-1}.i_0i_1},$$

where the i_j range over the alphabet $\{A, B, C, D\}$. Similarly, the set of points with two forward and backward iterates is

$$\bigcup_{(i_{-2}i_{-1}.i_0i_1i_2)} R_{i_{-2}i_{-1}.i_0i_1i_2},$$

and we may once again write the set of all points whose entire trajectories remain in R as

$$\Lambda = \bigcap_{n \geq 1} \bigcup_{(i_{-n} \dots i_{-1}.i_0i_1 \dots i_n)} R_{i_{-n} \dots i_{-1}.i_0i_1 \dots i_n}.$$

The invariant set Λ has many of the features of the Smale horseshoe. Topologically, both are totally disconnected maximal invariant set for the relevant map.²⁹ Dynamically, in both cases we have a stable and an unstable direction through each point $\mathbf{x} \in \Lambda$ (in fact, through each point in R), which for the Lorenz horseshoe may be seen as follows. Each of the regions A, B, C, D is contracted horizontally and expanded vertically, and so there exist two curves through \mathbf{x} , one stable and one unstable, with the following properties:

- (1) If \mathbf{y} and \mathbf{z} lie on the stable curve of some point \mathbf{x} , then their orbits are asymptotic in positive time:

$$\lim_{n \rightarrow +\infty} d(T^n(\mathbf{y}), T^n(\mathbf{z})).$$

- (2) If \mathbf{y} and \mathbf{z} lie on the unstable curve of some point \mathbf{x} , then their orbits are asymptotic in negative time:

$$\lim_{n \rightarrow -\infty} d(T^n(\mathbf{y}), T^n(\mathbf{z})).$$

- (3) The tangent vectors to the stable and unstable curves at \mathbf{x} lie close to the horizontal and vertical directions, respectively.

²⁹The careful reader will observe, however, that the set Λ we consider here is not closed, due to the presence of the line of discontinuity ℓ . We may carry out the geometric construction described above to obtain a closed Cantor-like set, but this line, along with all its (countably many) pre-images under T , must be removed from that set to obtain Λ .

As always, the hyperbolic structure given by the existence of stable and unstable directions leads to chaos in one of its incarnations. This can be seen more explicitly by considering the symbolic dynamics associated to the map T ; encoding a trajectory by its itinerary through the regions A, B, C , and D , we have a correspondence between points in Λ and sequences in the symbolic space

$$\Sigma_4 = \{A, B, C, D\}^{\mathbb{Z}}.$$

As usual, the dynamics of T are modeled by the dynamics of the shift $\sigma: \Sigma_4 \rightarrow \Sigma_4$. However, not all sequences in Σ_4 correspond to points in Λ ; for example, we see from Figure 68 that $T(A) \subset C \cup D$, and so every time a point in Λ has an itinerary which includes the symbol A , it must be followed by either C or D . Thus what we have here is actually a *Markov* shift, a subshift of finite type with the following transition matrix:

$$\begin{pmatrix} 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{pmatrix}$$

The map $T: \Lambda \rightarrow \Lambda$ is topologically conjugate to this subshift of finite type,³⁰ and so all admissible itineraries for the subshift are encodings of trajectories of T .

This implies that T has a great many trajectories which appear random. However, because Λ is a horseshoe and has zero Lebesgue measure, the set of such trajectories is invisible from the point of view of the original system, since with probability 1, an arbitrarily chosen point will lie outside the horseshoe, and hence will eventually leave the region R in which chaotic behaviour is observed. Thus the chaos implied by the hyperbolic structure in this case is *transient* chaos, which *is* visible and occurs for a set of initial conditions with positive measure. With non-zero probability, a randomly chosen trajectory will appear chaotic for some period of time before leaving R and becoming more regular.

Finally, we may return to the Lorenz system itself by drawing the trajectories in \mathbb{R}^3 which connect each $\mathbf{x} \in \Lambda$ to its image $T(\mathbf{x})$; the resulting set is a *filled-in horseshoe*, which locally is homeomorphic to the direct product of \mathbb{R} and Λ . This filled-in horseshoe plays the same role for the flow of the Lorenz system as Λ did for the Poincaré section; trajectories which venture near the horseshoe follow apparently chaotic trajectories for some finite period of time before wandering away and settling down.

³⁰Or rather, to this subshift with certain trajectories removed, corresponding to the line ℓ and its preimages.

Lecture 36

a. The Lorenz attractor. The scenario illustrated in Figures 67 and 68 occurs for a broad range of parameter values; the precise locations of the regions A, B, C, D depend on the value of the leading parameter r , but the same topological outcome is observed, and the system has a horseshoe. At some value $r_1 \approx 24.05$, everything changes; as r increases beyond this value, the region $R = A \cup B \cup C \cup D$ becomes large enough to contain the fixed points p_1 and p_2 , and indeed, to contain its own image $T(R)$; it becomes a *trapping region*. Taking the intersection of all the images of R , we obtain an attractor Λ for the map T ; the dynamics of $T: \Lambda \rightarrow \Lambda$ are chaotic, and because Λ attracts nearby trajectories, this chaotic behaviour is observed for a set of initial conditions of positive measure, in contrast to the case in the previous lecture.

In fact, we must be slightly more careful, since T is not defined on the line ℓ , and so strictly speaking, R is not a trapping region in the original sense. However, we do have

$$T(R \setminus \ell) \subset R,$$

and so we may define a nested sequence of sets R_n by $R_0 = R$, $R_{n+1} = T(R_n \setminus \ell)$; we have

$$R_0 \supset R_1 \supset R_2 \supset \cdots,$$

and the attractor is given by

$$(110) \quad \Lambda = \bigcap_{n \geq 1} R_n.$$

Connecting each point $\mathbf{x} \in \Lambda$ to its image $T(\mathbf{x})$ by the corresponding trajectory in \mathbb{R}^3 gives an attractor for the Lorenz system itself; this was the object discovered by Lorenz, who originally studied the parameter value $r = 28$.

b. The geometric Lorenz attractor. Although the Poincaré map T is, conceptually speaking, a relatively simple object, any attempts to do actual calculations with it are quickly stymied by the fact that we do not have a convenient formula for T , but must rather integrate the original system for some variable period of time before obtaining $T(\mathbf{x})$. In order to bypass this difficulty, one approach is to study the *geometric Lorenz map*, which shares (or appears to share) many topological properties with the original Poincaré map, but which is given by an explicit set of formulae and thus is more amenable to concrete analysis.

The geometric Lorenz map T takes the square $R = [-1, 1] \times [-1, 1]$ into itself as follows:

$$(111) \quad T(x, y) = ((-B|y|^{\nu_0} + Bx \operatorname{sgn} y|y|^\nu + 1) \operatorname{sgn} y, ((1+A)|y|^{\nu_0} - A) \operatorname{sgn} y),$$

where $\text{sgn } y = y/|y|$ denotes the sign of y , and the parameters lie in the following ranges:

$$0 < A \leq 1, \quad 0 < B \leq \frac{1}{2}, \quad \nu > 1, \quad \frac{1}{1+A} < \nu_0 < 1.$$

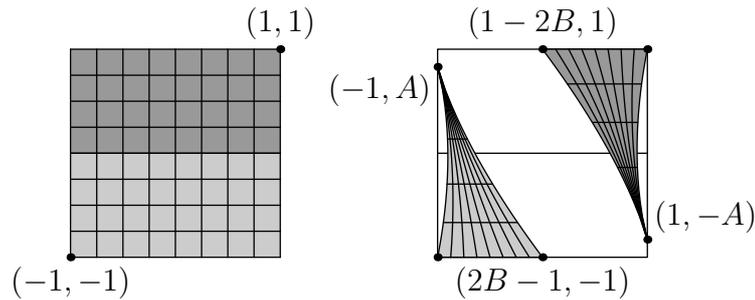


FIGURE 70. The geometric Lorenz map.

Figure 70 shows the image of R under the action of T ; several features are immediately apparent. The two corners $(-1, -1)$ and $(1, 1)$ are fixed by T , and since the y -coordinate of $T(x, y)$ does not depend on x , T maps horizontal lines into horizontal lines. In particular, the lines $y = 1$ and $y = -1$ are mapped into themselves as follows:

$$\begin{aligned} T(x, 1) &= (1 - B + Bx, 1), \\ T(x, -1) &= (B - 1 + Bx, -1). \end{aligned}$$

The map is continuous everywhere except along the x -axis; the continuation of the map from the lower half of the square would take the x -axis to the point $(-1, A)$, while the continuation from the upper half would take it to $(1, -A)$.

Since $0 < B \leq 1/2$, the map is contracting in the horizontal direction; as long as $A > 1/2$, it is expanding in the vertical direction, and so exhibits the same sort of hyperbolic structure at every point which we have already seen in the Smale–Williams solenoid, the Smale horseshoe, and so on. In both those cases, we found an maximal invariant set for the map on which the dynamics appear chaotic, and the same is true here. This set is referred to as the *geometric Lorenz attractor*, and because $T(R \setminus \ell) \subset R$, where ℓ is the x -axis, we may construct Λ explicitly as in (110).

Figure 71 shows the first two steps in the construction of Λ . $R_1 = T(R \setminus \ell)$ is the union of two triangles with curved sides; the image of $R_1 \setminus \ell$ comprises one triangle (lighter in the picture) and one biangle (skinnier and darker) inside each of these, for a total of four regions, whose union is R_2 .

One might expect, then, that R_3 would be the union of eight regions, with one triangle and three biangles in each half of R_1 , that R_4 will have one triangle and seven biangles in each half of R_1 , and so on, always forming two “fans” with hinges at $(1, -A)$ and $(-1, A)$. Indeed, this is the general

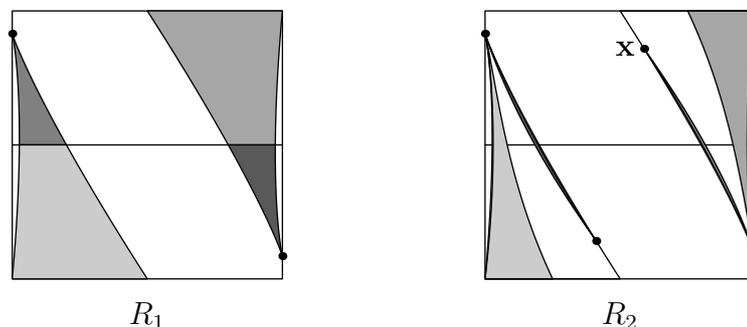


FIGURE 71. Constructing the geometric Lorenz attractor.

structure of the attractor, but things are not quite so cut and dry. Consider the biangle with one vertex at the point \mathbf{x} in Figure 71. The bottom half of this biangle will be mapped to a biangle in the left half of R_1 with one vertex at $(-1, A)$, while the top half will be mapped to a biangle in the right half with one vertex at $(1, -A)$ and the other at $T(\mathbf{x})$, and so R_3 will have one triangle and seven biangles, as expected.

However, because the map is expanding in the vertical direction, the point $T(\mathbf{x})$ will lie somewhere below \mathbf{x} , and may actually lie below ℓ ; if this is the case, then we have a biangle which does not split into two upon passing from R_3 to R_4 , and so in general, R_n may not have $2^n - 2$ biangles, as naïve reasoning would suggest.

Another way of analysing the structure of Λ is to consider its cross-section on a horizontal line $\ell_a = \{(x, y) \mid y = a\}$. We see from Figures 70 and 71 that $R_0 \cap \ell_a$ is the entire interval $[-1, 1]$, while $R_1 \cap \ell_a$ is the union of two disjoint closed subintervals, and $R_2 \cap \ell_a$ is the union of one, three, or four disjoint closed subintervals, depending on the value of a . Thus for each fixed a , the cross-section $\Lambda \cap \ell_a$ is the result of a Cantor-like construction from which certain basic intervals have been deleted, corresponding to those steps in which one of the “fingers” of the fan does not cross ℓ_a (at the present step) or did not cross ℓ (at the previous step).

Ideally, we would like to have a rule which determines which basic intervals are deleted and which remain; for example, such a construction could conceivably be the result of a Markov rule with a particular transition matrix. However, it turns out that there is no simple rule in the present case. We may gain some perspective on this fact by considering a one-dimensional factor of the geometric Lorenz map; since T maps horizontal lines to horizontal lines, we can factor out the x -coordinate and consider the one-dimensional map

$$f: [-1, 1] \rightarrow [-1, 1],$$

$$y \mapsto ((1 + A)|y|^{\nu_0} - A) \operatorname{sgn} y,$$

whose graph is shown in Figure 72. This map is reminiscent of the piecewise continuous interval maps we have already discussed at length. However, this is *not* a Markov map, because the image of an interval of continuity is not a union of such intervals; $f([-1, 0))$ contains part, but not all, of $(0, 1]$. Thus while we can pass to symbolic dynamics via the partition $\{[-1, 0), (0, 1]\}$ and obtain a topological conjugacy between $f: [-1, 1] \rightarrow [-1, 1]$ and the shift σ on some invariant subset $A \subset \Sigma_2^+$, the invariant subset A will have a quite complicated topological structure which is not given by any Markov rule.

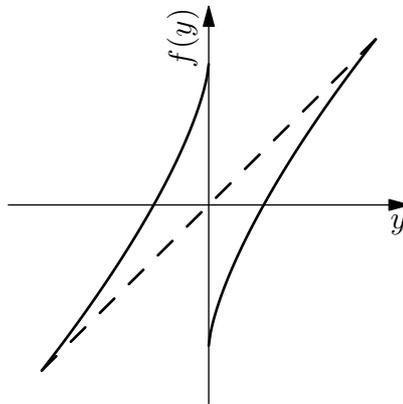


FIGURE 72. A one-dimensional factor of the geometric Lorenz map.

Lecture 37

a. Random fractals. Let us return momentarily to the penultimate idea in the previous lecture, in which we considered the cross-sections $\Lambda \cap \ell_a$ of the geometric Lorenz attractor. We found these to be Cantor-like sets with certain basic intervals deleted; however, because the intervals to be deleted are not specified by any sort of Markov rule, many of the tools we have used thus far to study the geometry of Cantor-like no longer hold any utility.

We encounter a similar situation if we consider *random fractals*; for example, we may consider a Cantor-like construction in which the basic intervals to be deleted at each step are chosen at random, or in which either the lengths or placements (or both) are chosen at random, or any number of other possibilities (we prorate for the time being any discussion of just how these things are to be “randomly” chosen). How are we to determine the Hausdorff dimension of the resulting Cantor-like set? There is no Moran-type argument available in these cases, and the direct approach appears to be impotent.

An effective way of analysing such sets is to return to the approach we took with the Non-uniform Mass Distribution Principle, and consider the class of probability measures supported on the fractal set of interest. In many cases, certain geometric properties of this set can be deduced from a variational principle involving dynamical and measure-theoretic properties of the various measures; the latter has connections to what is known as *ergodic theory*, and the whole enterprise is a part of the *thermodynamic formalism*.

b. Back to the Lorenz attractor, and beyond. As was mentioned at the end of Lecture 36(a), connecting the points of the attractor for the Poincaré section of the Lorenz equations with the corresponding trajectories yields an attractor for the Lorenz system itself, which is pictured in Figure 73. This image has become iconic for chaos theory both because of its historic significance and because of the mathematical concepts it embodies.

Despite the seminal role that it played in the study of chaos, interest in the Lorenz system began to wane in the 1980’s, for a number of reasons. For example, it turns out that interesting behaviour is observed only for a narrow range of values of the parameter r ; furthermore, given the origins of the model, the question of its relevance naturally arises. How well do these equations approximate what actually occurs at the onset of turbulence? The Lorenz system (106) was obtained by restricting our attention to three particular terms in a more complicated system; if we add some of the other terms back into the model, do we still see the same sort of behaviour?

As it happens, the particular choice of x , y , and z in (106) was somewhat serendipitous; studies of various other approximations to the Navier–Stokes

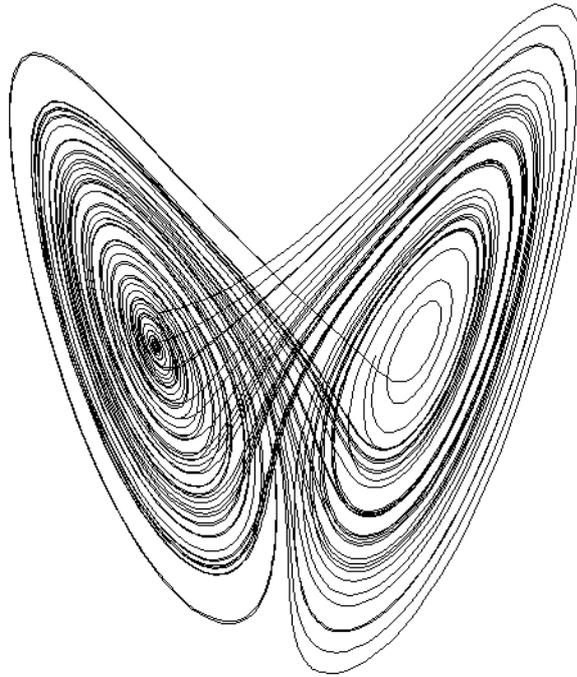


FIGURE 73. The Lorenz attractor in \mathbb{R}^3 .

equations failed to observe chaotic behaviour, which cast doubt on the relationship between (106) and the actual physical phenomenon. Nevertheless, the Lorenz system was the first example of “deterministic chaos”—the term coined by Jim Yorke, which is now commonly used to describe such behaviour—and as such was tremendously important in its own right, whatever its relationship to a physical system.

Since Lorenz’s original work, many other examples of deterministic chaos have been studied, and the mechanisms which produce it are now better understood. A typical chaotic system contains either a horseshoe (resulting from a homoclinic point) or an attractor (resulting from a trapping region); each of these is a fractal which is invariant under the action of the system, which is internally unstable (hence the chaotic behaviour), and which has zero volume (Lebesgue measure). Despite this last point, both horseshoes and attractors are observable via their effect on nearby trajectories; in the former case, this effect lasts for a finite period of time, producing intermittent chaos, while in the latter, the effect persists for all time, as nearby trajectories approach the attractor.

Depending on the context, one may see the Lorenz attractor and its many relatives referred to as “strange attractors”, emphasising their fractal

geometry, or as “chaotic attractors”, emphasising the unpredictable nature of the observed dynamics, or as “hyperbolic attractors”, emphasising the underlying dynamical instability of the attractor at each point. This final moniker has the advantage of giving pride of place to the force which drives both the fractal structure and the apparent randomness displayed by the attractor; namely, the presence at every point on the attractor of both stable and unstable directions, so that a saddle-like structure is ubiquitous. This is necessary because a trajectory with only stable directions will attract nearby trajectories, and hence cannot display chaotic behaviour, while a trajectory with only unstable directions will repel all nearby trajectories, and hence such trajectories cannot be tightly intertwined, and cannot have the complicated geometric structure we observe in fractal sets.

Furthermore, the understanding of hyperbolicity (which may be uniform or non-uniform) as the driving impulse behind chaos makes it clear that in many cases, topological and geometrical considerations alone are enough to describe a chaotic system, without the need for an explicit formula, which greatly extends the theory’s generality. For example, the geometric Lorenz map (111) can be generalised in a number of ways, such as choosing a line of discontinuity which is no longer horizontal; the resulting attractor is called a *Belych attractor*, and comes up in the study of certain non-linear electrical circuits, a far cry from the study of the atmosphere which motivated the original system!

Indeed, the real significance of Lorenz’s work is not in its contribution to atmospheric science, but in the fact that it helped fling open the doors to whole new areas of mathematics, which have since found applications across the entire spectrum of scientific research. In the process, many strands of topology, geometry, and dynamics have been woven into a single fabric. The fractal sets first introduced by Cantor and studied by Besicovitch to answer questions in classical set and function theory were conscripted by Smale to serve a key role in a particular type of dynamical system; eventually, through the work of Mandelbrot, Lorenz and others, it became apparent that far from being pathologies with limited interest, and that only for mathematicians, these “fractals” in fact lie at the heart of many important phenomena in the natural world, and still hold many deep mysteries which we have yet to understand.