Dynamical Systems

An LTCC course

Wolfram Just and Franco Vivaldi

School of Mathematical Sciences
Queen Mary University of London

Last updated: January 16, 2019

These lecture notes provide a brief introduction to some topics in dynamical systems theory, suitable for a short course.

1) Dynamical systems: continuous and discrete time.

2) Linear stability analysis and beyond: hyperbolicity, structural stability.

3) Bifurcations: what happens when a state becomes unstable.

4) Symbolic dynamics: turning orbits into strings of symbols.

5) Chaos: when the future of a deterministic system cannot be determined.

The presentation is informal and hands-on. The accompanying supplementary notes by F Vivaldi provide some background material in analysis and linear algebra, and also deal with one-dimensional flows, which are not part of this course. Suggestions for further reading will be found in the last section.

1 Dynamical systems

Informally, a dynamical system is a system whose states are represented by the points of a set; these states evolve with time according to a deterministic law, specified by a differential or difference equation. The system is not subjected to the effects of noise[1].

For instance, let us consider a single particle in one-dimension, subjected to a force. The state of this system is determined by the particle's position and velocity; the time-evolution is governed by Newton's law $F = ma$, which can be written as a pair of first-order ordinary differential equations: $\dot{x} = y, \dot{y} = F(x)/m$. This law determines the state at time t provided the state at time 0 is given. This is a dynamical system on the set $\mathbb{R}^2$; the time is a real number.

Let M be any set and $f : M \to M$ be any function. Choose $x_0 \in M$, and, using the first-order difference equation $x_{t+1} = f(x_t),\ t \geqslant 0$, define recursively an infinite sequence $(x_0, x_1, x_2, \dots)$ of elements of M. This is a dynamical system on M; the time is an integer.

Definition 1.1 *A dynamical system consists of a set M, called the phase space, and a one-parameter family of transformations $\Phi_t : M \to M$ such that, for all $x \in M$,*

i) $\Phi_0(x) = x$;

ii) $\Phi_t(\Phi_s(x)) = \Phi_{t+s}(x)$, *for all* $s, t \geqslant 0$.

If $t \in \mathbb{R}$ (respectively, $t \in \mathbb{Z}$), then we speak of a continuous time (respectively, discrete time) dynamical system.

The family $\{\Phi_t\}$ is a semi-group of transformations. If ii) holds also for negative t, s, then we have a group, and the system is said to be invertible. The set $\Gamma(x_0) = \{\Phi_t(x_0), t \geqslant 0\}$ is called the (forward) orbit of the point x_0, called the initial condition.

For discrete time dynamical systems, the set M can be quite arbitrary. For continuous time, or for systems of physical origin, M is a manifold[2]. For continuous time, the family of functions Φ_t is called a flow.

---

[1] In presence of noise, we speak of a *stochastic* dynamical system.

[2] A set that, locally, looks like $\mathbb{R}^n$ for some n, see page 16.

1.1 Flows

We begin with some examples of continuous-time dynamical systems.

Example 1.1 The harmonic oscillator. This is one of the simplest dynamical system of mathematical physics, yet a very significant one. The phase space is $M = \mathbb{R}^2$; the basic information is displayed below[3].

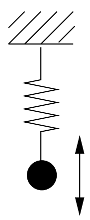

equations	initial condition	solution
$\dot{x} = y$	$x(0) = x_0$	$x(t) = x_0 \cos(\omega t) + \dfrac{y_0}{\omega}\sin(\omega t)$
$\dot{y} = -\omega^2 x$	$y(0) = y_0$	$y(t) = -\omega x_0 \sin(\omega t) + y_0 \cos(\omega t)$

with $\Phi_t(x_0, y_0) = (x(t, x_0, y_0), y(t, x_0, y_0))$.

The phase space space foliates into invariant ellipses. The figure on the right illustrates the geometrical meaning of the associative property ii) in definition 1.1.

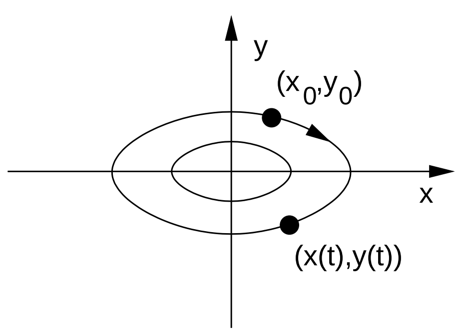

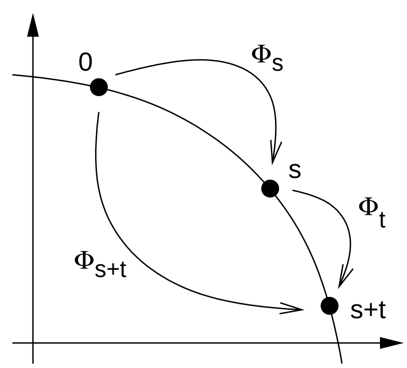

The definition of a dynamical systems requires that, for any initial condition in phase space, the orbit must be unique and defined for all future times. This is not guaranteed for the solutions of a differential equation $\dot{x} = f(x)$, and so there are differential equations which cannot be regarded as dynamical systems.

[3]The dot denotes differentiation with respect to time.

Example 1.2 The solutions are not defined for all times. The system

$$\dot{x} = x^2 \qquad x(0) = x_0 \tag{1}$$

has solution

$$x(t, x_0) = \frac{x_0}{1 - x_0 t}.$$

For $x_0 \neq 0$, this function is not defined over the whole $\mathbb{R}$. It is defined in the interval $(-\infty, 1/x_0)$ for $x_0 > 0$, and $(1/x_0, \infty)$ for $x_0 < 0$. This is because the real function $f(x) = \dot{x}$ grows too rapidly near infinity, so that it is possible to reach infinity, or come back from infinity, in a finite time.

Example 1.3 Lack of uniqueness. The differential equation

$$\dot{x} = 2\sqrt{|x|}, \quad \text{with} \quad x(0) = 0$$

has two distinct solutions for $t \geqslant 0$, namely $x(t) = 0$ and $x(t) = t^2$. This is because the function $f(x) = \dot{x}$ is not Lipschitz continuous at the equilibrium point $x = 0$ (the constant solution), and hence it does not vanish rapidly enough near it. So it is possible to reach the equilibrium point, or move away from it, in a finite time.

To ensure uniqueness of solutions, it suffices to require that f be of class C^1.

Example 1.4 Forced van der Pol oscillator.

$$\begin{aligned} \dot{x} &= x - \sigma y - x(x^2 + y^2) \\ \dot{y} &= \sigma x + y - y(x^2 + y^2) - \gamma. \end{aligned} \tag{2}$$

For any value of the real parameters σ, γ, this equation defines a vector field $\mathbf{v}(z) = \dot{z} = (\dot{x}, \dot{y})$ in $\mathbb{R}^2$. Plotting this vector field (here for $\sigma = 0.2$ and $\gamma = 0.3$) gives already an impression of the behaviour of the solutions. Dynamical systems theory is about understanding how the flow relates to the vector field.

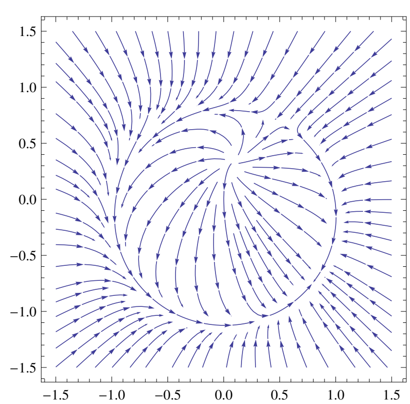

For special parameter values we are able to compute the flow analytically. If $\gamma = 0$, then using polar coordinates $x = r\cos(\varphi)$, $y = r\sin(\varphi)$, the equations of motion read

$$\dot{r} = r - r^3, \quad \dot{\varphi} = \sigma\,. \tag{3}$$

The special solution $r(t) = 0$ yields the constant solution $z(t) = (x(t), y(t)) = (0,0)$ (an equilibrium point, or fixed point), while $r(t) = 1$, $\varphi(t) = \sigma t + \varphi(0)$ results in the harmonic solution $z(t) = (\cos(\sigma t + \varphi(0)), \sin(\sigma t + \varphi(0)))$. Let $T = 2\pi/\sigma$. Then this solution satisfies $z(t) = z(t+T)$ for all t, and we call it a periodic orbit (or a cycle), with period T. A periodic orbit of a flow is a closed non-self intersecting curve in phase space. In the present example, if $\sigma \neq 0$, then this periodic orbit is a so-called limit cycle, since it is approached by nearby orbits —see below. For $\gamma = 0$, it is even possible to write down the general solution.

The long-time behaviour of the orbits of one- and two-dimensional flows is well-understood. In one dimension, the solutions are either constant, or strictly monotonic. Indeed let $\dot{x} = f(x)$ be a smooth differential equation over $\mathbb{R}$, and let $x = x(t)$ be a non-constant solution (i.e., $\dot{x}(t) = f(x(t))$ and $f(x_0) \neq 0$). Then separation of variables yields $\int_{x(0)}^{x(t)} \mathrm{d}x/f(x) = t$. Since the integral on the left-hand side is bounded (and smooth in t), we have $f(x) \neq 0$ for $x \in [x(0), x(t)]$. Because t is arbitrary, $f(x(t))$ does not change sign. Therefore the solution is monotonic, and hence it is either unbounded or tends to a finite limit $x(t) \to x_*$. It then follows by continuity that $x(t) = x_*$ is a constant solution of the differential equation.

In higher dimensions, limit sets of flows can be much more complicated than fixed points, and we need a machinery to characterise this form of convergence. Let $\Gamma(z_0) = \{\Phi_t(z_0) : t \in \mathbb{R}\}$ be the orbit of z_0.

A point w is an ω-limit point of the orbit $\Gamma(z_0)$ if there is a sequence (t_k) of times such that

$$\lim_{j\to\infty} t_k = \infty \qquad \text{and} \qquad \lim_{j\to\infty} \Phi_{t_k}(z_0) = w. \tag{4}$$

In other words, w is an ω-limit point if for any neighbourhood U of w there is a (smallest) time $t^* = t^*(U)$ such that $\Phi_{t^*}(z_0) \in U$.

The set of all ω-limit points of $\Gamma(z_0)$ is called the ω-limit set of $\Gamma(z_0)$, denoted by $\omega(z_0)$ [or $\omega(\Gamma)$, since this set is a property of the orbit]. By reversing the direction of time, and replacing ∞ by $-\infty$, we obtain the analogous concept of α-limit set.

The simplest situation occurs when $\lim_{t\to\infty}\Phi_t(z_0) = z_*$, where z_* is a fixed point. Then, by letting $t_k = k$ and $w = z^*$ in (4), we find that $z^* \in \omega(z_0)$. Since there cannot be any other ω-limit point (why?), we have $\omega(z_0) = \{z_*\}$.

Another possibility is that $\omega(z_0) = \alpha(z_0) = z_*$, with $z_0 \neq z_*$, that is, the orbit of z_0 approaches z_* in both time directions. We then say that z_0 is a homoclinic point, and that $\Gamma(z_0)$ is a homoclinic orbit. The point z_0 is heteroclinic if $\omega(z_0) = z_*$, and $\alpha(z_0) = w_*$ where w_* is an equilibrium distinct from z_*. Heteroclinic orbits are commonplace in flows on the line. By contrast, homoclinic orbits are not possible on the line, but can be found on the circle.

Let Γ be a periodic orbit. If $z_0 \in \Gamma$, then $\alpha(z_0) = \omega(z_0) = \Gamma$. However, if $z_0 \notin \Gamma$, then it is possible that one of the sets $\alpha(z_0)$ and $\omega(z_0)$ is equal to Γ and the other isn't, as in (3). Finally, if Γ is a homo/heteroclinic orbit, the possibility exists that $\omega(z_0) \supset \Gamma$ for $z_0 \notin \Gamma$ but $\omega(z_0) \neq \Gamma$ for $z_0 \in \Gamma$, as the following example illustrates.

Example 1.5 Let us consider the differential equation

$$\begin{aligned} \dot{\phi} &= v \\ \dot{v} &= -\sin(\phi) - \gamma v\left(\frac{v^2}{2} - \cos(\phi) - 1\right) \end{aligned} \tag{5}$$

where $\gamma > 0$ denotes a fixed real parameter. Since the vector field depend periodically on ϕ, the phase space is a cylinder: $z = (\phi, v) \in \mathbb{S} \times \mathbb{R}$. This equation has been tailored in such a way that an orbit with initial condition close to $z = (0,0)$ approaches the curve $\mathrm{H}(\phi, v) = 1$ where $\mathrm{H}(\phi, v) = v^2/2 - \cos(\phi)$. Such a curve is invariant, since

$$\dot{\mathrm{H}} = \frac{\mathrm{dH}}{\mathrm{d}\phi}\dot{\phi} + \frac{\mathrm{dH}}{\mathrm{d}v}\dot{v} = -\gamma v^2(\mathrm{H} - 1),$$

and consists of three orbits: the fixed point $(\pi, 0)$, and two distinct orbits which approach the fixed point in both time directions, namely two homoclinic orbits[4]. The time-dependence of ϕ and the phase portrait clearly show that the orbit approaches from opposite sides the equilibrium point $(\pi, 0)$ (which coincides with $(-\pi, 0)$, since we are on the torus), lingering longer and longer around it, without ever settling down.

[4]These orbits are also called separatrices, since they divide the phase space into regions whose orbits have different topological properties.

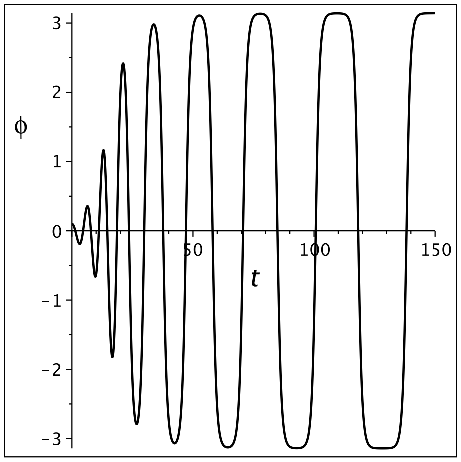

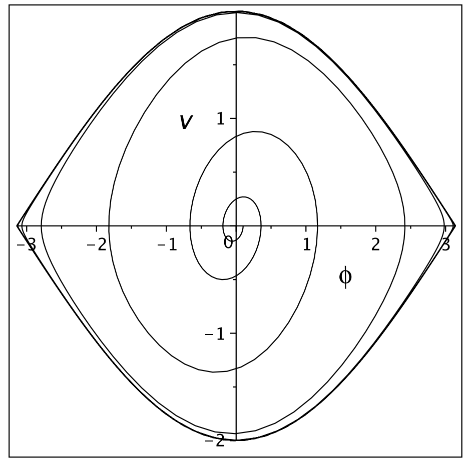

Exercise 1.1. Let $n \in \mathbb{N}$. Define *i*) a flow on the line with *n* heteroclinic orbits; *ii*) a flow on the circle with a homoclinic orbit; *iii*) a flow on the plane with *n* heteroclinic orbits [*Hint:* modify the ϕ-dependence in (3)].

The following theorem describes all possible limit sets of a planar system.

Theorem 1.1 (Poincaré-Bendixson) *Let $\dot{z} = f(z)$ be a planar system with a finite number of equilibrium points. If the orbit $\Gamma(z_0)$ of z_0 is bounded, then one of the following is true:*

i) *The ω-limit set $\omega(z_0)$ is a single equilibrium point z^* and $\Phi_t(z_0) \to z^*$ as $t \to \infty$.*

ii) *$\omega(z_0)$ is a periodic orbit Γ_* and $\Gamma(z_0)$ is either equal to Γ_*, or it spirals towards Γ_* on one side of it.*

iii) *$\omega(z_0)$ consists of equilibrium points and orbits whose α- and ω-limit sets are equilibrium points.*

In particular, any bounded ω-limit set which contains no equilibrium points is a periodic orbit [case *ii*)]. The proof of the Poincaré-Bendixson Theorem is exploits the fact that solutions in two-dimensional phase spaces cannot cross, which eventually leads to the existence of a monotonic one-dimensional Poincaré map (see next section).

The above theorem illustrates a key aspect of the theory of dynamical systems. One is not concerned with methods of solution of equations of motion, but rather in *qualitative* properties of families of solutions, here expressed in the language of limit sets. Solving equations of motion is often done by computers, whereby we surrender to the difficulty/impossibility of finding analytic solutions. Moreover, inspecting particular solutions is not necessarily informative.

1.2 From continuous to discrete time

Let Φ_t be the flow of a differential equation $\dot{z} = f(z)$ on a space M. There are several ways of constructing from it a discrete-time dynamical system.

Time-advance maps

Given Φ_t, by merely restricting the time t to the integers, we obtain a discrete-time dynamical system of the same space. The map Φ_1 is called the (unit) time-advance map of the flow. It sends a point $z(0) \in M$ to the point $z(1)$. One may restrict the time in other ways, for instance to the additive group $h\mathbb{Z}$, for some h. A numerical integration scheme is a time-advance map. Given $z(0)$, it outputs an approximate value for $z(h)$, where h is a very small positive number. By iterating this map a large number of times, one obtains an approximate numerical solution for large t. More sophisticated schemes adjust the size of the time-step h to the location in phase space.

Poincaré maps

Let Φ_t be as above, and let the space M be N-dimensional. Consider a smooth $N-1$ dimensional surface Σ with the property that *i*) the vector field is nowhere tangent to Σ; *ii*) the orbit of every point of Σ eventually returns to Σ. The Poincaré map $\mathscr{P}$ sends the point $w \in \Sigma$ to the first intersection of the orbit of w with Σ.

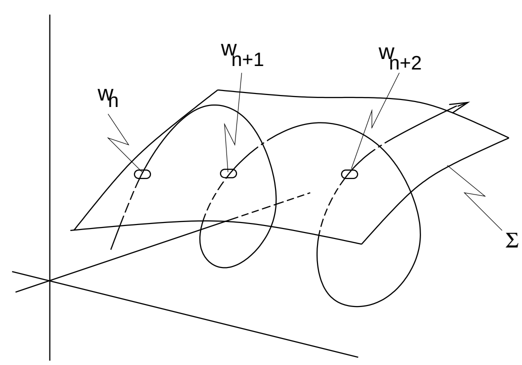

The surface Σ is called a surface of section of the flow. To define the Poincaré map, we first introduce the notion of first-return time $\tau(w)$ to the surface of section at the point

w

$$\tau : \Sigma \to \mathbb{N}_0 \qquad \tau(w) = \inf\{t \in \mathbb{R} : t > 0, \Phi_t(w) \in \Sigma\}. \tag{6}$$

(The infimum ensures that $\tau(w) = 0$ if w is an equilibrium point.) Then we define the Poincaré map as follows:

$$\mathscr{P} : \Sigma \to \Sigma \qquad \mathscr{P}(w) = \Phi_{\tau(w)}(w).$$

Choosing a surface of section requires some care. For example, consider the harmonic oscillator (p. 2). The x-axis is a surface of section, and from symmetry we see that $\mathscr{P}(x) = -x$. To avoid collecting unnecessary information, we can either record only every other crossing of the surface, or modify the surface so that there is a single crossing during each period: $\Sigma = \{(x, y) \in \mathbb{R}^2 : y = 0, x \geqslant 0\}$, which is a ray. In either case, the Poincaré map is the identity: every point is a fixed point. Because the rotations are isochronous (same angular frequency), the first-return time is constant: $\tau(x) = 2\pi/\omega$.

In general, the first-return time is not constant, and it could be unbounded. Let us consider again example 1.5, p. 5. The one-dimensional surface $\Sigma = \{(\phi, v) \in \mathbb{S} \times \mathbb{R} : v = 0, 0 \leqslant \phi \leqslant \pi\}$ is a surface of section for all orbits lying within the region $\mathrm{H}(\phi, v) \leqslant 1$. The map $\mathscr{P}$ has two fixed points, $\phi = 0$ and $\phi = \pi$, and for all other points we have $\mathscr{P}(\phi) > \phi$. Considering the nature of the motion near the separatrix, we infer that this Poincaré map has unbounded first-return time: $\lim_{\phi \to \pi} \tau(\phi) = \infty$.

1.3 Linear systems

There is a complete theory of linear dynamical systems, which leads to their solution and classification. (By contrast, nonlinear systems are solvable only in exceptional cases.)

A linear system of differential equations with constant coefficients

$$\begin{aligned}
\dot{x}_1(t) &= A_{1,1}x_1(t) + A_{1,2}x_2(t) + \cdots + A_{1,n}x_n(t) \\
\dot{x}_2(t) &= A_{2,1}x_1(t) + A_{2,2}x_2(t) + \cdots + A_{2,n}x_n(t) \\
&\vdots \\
\dot{x}_n(t) &= A_{n,1}x_1(t) + A_{n,2}x_2(t) + \cdots + A_{n,n}x_n(t)
\end{aligned}$$

can be written as

$$\dot{z}(t) = \mathbf{A}\, z(t), \tag{7}$$

where $z = (x_1, x_2, \ldots, x_n)^T$, and $\mathbf{A}$ denotes the square matrix with coefficients $A_{i,j}$. Letting

$$\mathbf{A}\,\mathbf{u}_\ell = \lambda_\ell \mathbf{u}_\ell \qquad \mathbf{u}_\ell \in \mathbb{R}^n$$

denote the eigenvalue equation for $\mathbf{A}$, particular solutions of the linear system are given by exponential functions $z(t) = \exp(\lambda_\ell t)\mathbf{u}_\ell$. If the eigenvectors yield a basis of the phase space $\mathbb{R}^n$ (that is, if the eigenvalues are distinct, or if geometric multiplicity equals algebraic multiplicity), then the general solution —the flow— can be written as a linear combination of exponentials

$$z(t) = c_1 e^{\lambda_1 t}\mathbf{u}_1 + c_2 e^{\lambda_2 t}\mathbf{u}_2 + \cdots + c_n e^{\lambda_n t}\mathbf{u}_n,$$

where the coefficients c_ℓ are determined by the initial conditions. (A similar expression exists if eigenvectors fail to form a basis.) For a complex conjugate pair of eigenvalues, the corresponding coefficients are complex conjugate, to result in a real-valued solution.

The initial condition $z(0) = (0,0)$ yields the so-called trivial solution of the linear system: $z(t) \equiv (0,0)$. If $\mathrm{Re}(\lambda_\ell) < 0$ for any $\ell = 1, \ldots, n$, then the general solution tends towards the trivial solution regardless of the initial condition: $z(t) \to (0,0)$. In this case the trivial solution is said to be stable. If we have at least one eigenvalue λ_ℓ with $\mathrm{Re}(\lambda_\ell) > 0$, then there are constants c_ℓ, i.e., there are initial conditions, such that $\|z(t)\|$ grows exponentially with time. The trivial solution is said to be unstable.

Example 1.6 Stable focus:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} -1 & -1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

Eigenvalues and eigenvectors:

$$\lambda_1 = -1 + i, \quad \mathbf{u}_1 = \begin{pmatrix} 1 \\ -i \end{pmatrix}, \qquad \lambda_2 = -1 - i, \quad \mathbf{u}_2 = \begin{pmatrix} 1 \\ i \end{pmatrix}.$$

The general solution (flow) with complex-conjugate coefficient $c, \overline{c}$ and $c = |c|\exp(i\alpha)$ is:

$$\begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = c e^{\lambda_1 t}\mathbf{u}_1 + \overline{c} e^{\lambda_2 t}\mathbf{u}_2 = \begin{pmatrix} 2|c|e^{-t}\cos(t+\alpha) \\ 2|c|e^{-t}\sin(t+\alpha) \end{pmatrix}.$$

The time-dependence displays an exponentially damped oscillation. In the x_1-x_2 phase plane the solution determines a curve. If we use polar coordinates $x_1 = r\cos(\varphi)$ and $x_2 = r\sin(\varphi)$, then the time-dependence reads $r(t) = 2|c|\exp(-t)$ and $\varphi(t) = t + \alpha$. If we eliminate the time t, the curve is given by $r(t) = 2|c|\exp(\alpha)\exp(-\varphi(t))$, a logarithmic spiral.

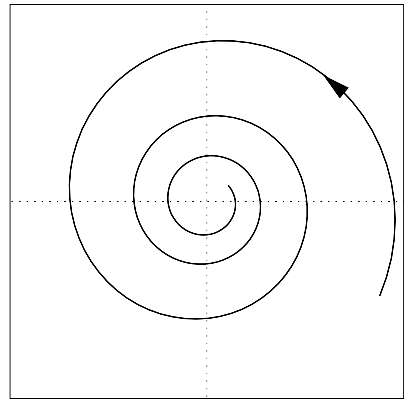

(It is a common misconception that the orientation of the spiral is linked to the imaginary part of the eigenvalues.)

Example 1.7 Saddle point:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

Eigenvalues and eigenvectors:

$$\lambda_1 = 1, \quad \mathbf{u}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad \lambda_2 = -1, \quad \mathbf{u}_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

The general solution (flow) with real-valued expansion coefficients is:

$$\begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = c_1 e^{\lambda_1 t}\mathbf{u}_1 + c_2 e^{\lambda_2 t}\mathbf{u}_2 = \begin{pmatrix} c_1 e^t + c_2 e^{-t} \\ c_1 e^t - c_2 e^{-t} \end{pmatrix}.$$

It is easy to verify that $x_1^2 - x_2^2 = 4c_1 c_2$, i.e., the curves in phase space (for different initial conditions) are hyperbolae. Only initial conditions with $c_1 = 0$ (which is the eigenspace of $\mathbf{u}_2$) yield solutions which tend to the fixed point in the origin. All other initial conditions yield diverging orbits.

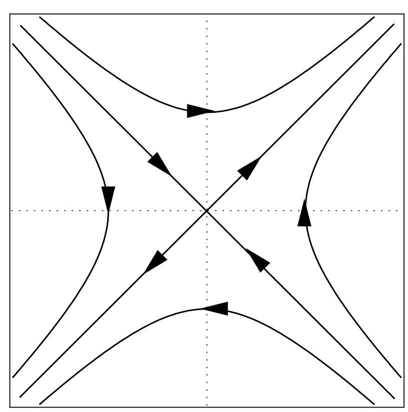

The two eigenvectors $\mathbf{u}_1$ and $\mathbf{u}_2$ determine the so-called unstable and stable direction of the saddle, respectively. More generally, eigenvectors with $\mathrm{Re}(\lambda_\ell) > 0$ ($\mathrm{Re}(\lambda_\ell) < 0$) span the so-called unstable (stable) eigenspace.

Example 1.8 Stable node:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} -1 & 0 \\ 0 & -2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

Eigenvalues and eigenvectors:

$$\lambda_1 = -1, \quad \mathbf{u}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad \lambda_2 = -2, \quad \mathbf{u}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The general solution (flow) with real-valued coefficients is

$$\begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = c_1 e^{\lambda_1 t}\mathbf{u}_1 + c_2 e^{\lambda_2 t}\mathbf{u}_2 = \begin{pmatrix} c_1 e^{-t} \\ c_2 e^{-2t} \end{pmatrix}$$

All solutions tend towards the origin, but the relaxation along the $\mathbf{u}_2$ direction is faster. It is easy to verify that the orbits in phase space are given by $x_2 = x_1^2 c_2 / c_1^2$ (if $c_1 \neq 0$), which is a foliation of parabolas.

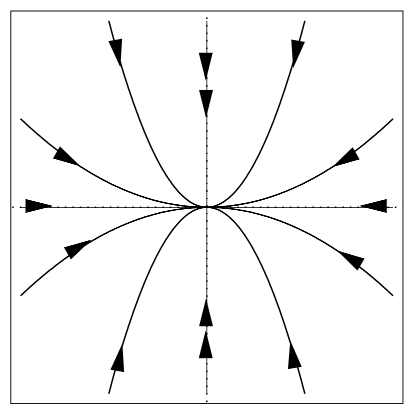

We have given examples of stable focus, saddle, and stable node. There are also unstable versions of foci and nodes, corresponding to eigenvalues with positive real part. The orbits of the unstable systems are obtained from the orbits of the stable systems by reversing the direction of travel.

These examples cover the typical cases occurring in two-dimensional linear ordinary differential equations. In section 3, we shall meet 'non generic' cases, representing transitions between these systems.

1.4 Classification of planar linear systems

Linear systems are classified using linear algebra. Let $\mathbf{A}$ be a 2×2 real matrix. Jordan's theorem states that there is an invertible 2×2 real matrix $\mathbf{P}$ such that

$$\mathbf{P}^{-1}\mathbf{A}\mathbf{P} = \mathbf{J}, \tag{8}$$

where $\mathbf{J}$ is one of the following matrices:

$$\mathbf{J}_1 = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \qquad \mathbf{J}_2 = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} \qquad \mathbf{J}_3 = \begin{pmatrix} \alpha & \beta \\ -\beta & \alpha \end{pmatrix} \tag{9}$$

and $\lambda_1, \lambda_2, \lambda, \alpha, \beta \in \mathbb{R}$, with $\beta \neq 0$. The matrices $\mathbf{A}$ and $\mathbf{J}$ have the same eigenvalues, since trace and determinant are preserved by the linear coordinate change (8).

The matrix $\mathbf{J}$ appearing in (8) is called the Jordan canonical form of the matrix $\mathbf{A}$.

Thus if we change co-ordinates in an appropriate way, then any matrix will be transformed into precisely one of the three Jordan matrices (9), and hence any real linear differential equation on the plane can be reduced to one of three canonical forms. A unifying formalism is achieved by introducing the notion of exponential of a square matrix $\mathbf{A}$:

$$e^{\mathbf{A}} = \mathbf{1} + \mathbf{A} + \frac{\mathbf{A}^2}{2!} + \cdots = \sum_{k\geqslant 0} \frac{\mathbf{A}^k}{k!},$$

where $\mathbf{1}$ denotes the identity matrix of the appropriate dimension. This expression converges with respect to the norm

$$\|\mathbf{A}\| := \sup_{z\neq 0} \frac{\|\mathbf{A}z\|}{\|z\|}.$$

It can be shown that the solution of the differential equation (7) has the neat form

$$z(t) = e^{\mathbf{A}t} z(0). \tag{10}$$

We write the matrix exponential in the case in which $\mathbf{A}$ is one of the Jordan matrices (9).

$$\begin{aligned}
\exp(\mathbf{J}_1 t) &= \begin{pmatrix} e^{\lambda_1 t} & 0 \\ 0 & e^{\lambda_2 t} \end{pmatrix} \\
\exp(\mathbf{J}_2 t) &= e^{\lambda t} \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \\
\exp(\mathbf{J}_3 t) &= e^{\alpha t} \begin{pmatrix} \cos(\beta t) & \sin(\beta t) \\ -\sin(\beta t) & \cos(\beta t) \end{pmatrix}.
\end{aligned}$$

1.5 Some remarks on stability

There are countless definitions of stability in the literature, and we must spell out how this term is interpreted in this course.

A fixed point z_* is (Lyapounov) stable if for every neighbourhood of V of z_* there is a neighbourhood U of z_* such that $\Phi_t(z_*)$ remains in V for all times. This means that all points sufficiently close to z_* remain forever near z_* (although they do not necessarily approach z_*).

Lyapounov stability is equivalent to the statement that the family of time-advance maps Φ_t (see definition 1.1) is equicontinuous at x_*, namely continuous and with the same variation for all t.

A fixed point z_* is asymptotically stable if it is stable and if there is a neighbourhood U of z_* such that, for any $x_0 \in U$, $\Phi_t(x_0)$ converge to z_* as $t \to \infty$. (Requiring stability is not redundant here —see exercises.) A point may be stable but not asymptotically stable, the trivial example being the identity: $\Phi_t = \mathbf{1}$. For nontrivial examples see exercises.

In this course, we use the term unstable as the logical negation of stable. This is not universally accepted, and often this term is interpreted in the following stronger sense: a fixed point z_* is unstable if there is a neighbourhood U of z_* such that for all $z_0 \in U$ the point $\Phi_t(z_0)$ eventually leaves U. This is not the logical negation of stability, and hence there are fixed points which are neither stable nor unstable —see exercises.

In section 2.3 we shall introduce the term structurally stable, quite unrelated to those given above.

Exercise 1.2. Give an example of a one-dimensional smooth flow with a fixed point which is

(a) not stable, yet all orbits converge to it;

(b) stable but not asymptotically stable;

(c) neither stable nor unstable, according to the strong definition of instability given above.

2 Hyperbolicity

In the previous section we computed the flow of linear systems. This knowledge will be used to investigate local stability of equilibrium points of nonlinear systems. This is possible as long as the so-called marginally stable cases —eigenvalues with zero real part— are absent. By introducing nonlinear extensions of local linear systems, the linear analysis can also provide information about global features of the dynamics.

2.1 Fixed points and linear stability

Let $\dot{z} = f(z)$ denote a differential equation on $\mathbb{R}^n$. A fixed point (or equilibrium point) z_* is a root of the equation $f(z) = 0$. The initial condition $z(0) = z_*$ yields a constant solution $z(t) = z_*$.

Example 2.1 We compute the fixed points of the forced van der Pol oscillator (p. 3) for the parameter values $\sigma = 0$ and $\gamma > 0$. The fixed points are the solutions of the following system of equations, derived from (2):

$$\begin{aligned} 0 &= x - x(x^2+y^2) \qquad &\Rightarrow \qquad & x = 0, \text{ or } x^2+y^2 = 1 \\ 0 &= y - y(x^2+y^2) - \gamma \qquad &\Rightarrow \qquad & x^2+y^2 \neq 1 \text{ (since } \gamma > 0). \end{aligned}$$

Denoting the fixed point by $z_* = (x_*, y_*)$ we find:

$$y_* - y_*^3 = \gamma \quad \text{and} \quad x_* = 0 \qquad (\sigma = 0, \gamma > 0). \tag{11}$$

We see that if $|\gamma| > 2/(3\sqrt{3})$, there is a single solution, and three solutions if $|\gamma| < 2/(3\sqrt{3})$. The fixed point with negative ordinate $y_* < -1$ exists for all $\gamma > 0$, while for $0 < \gamma < 2/(3\sqrt{3})$ the two additional fixed points have positive ordinate, $0 < y_* < 1/\sqrt{3}$ and $1/\sqrt{3} < y_* < 1$, respectively. The critical parameter value $\gamma = 2/(3\sqrt{3})$ corresponding to $y_* = 1/\sqrt{3}$ leads to a bifurcation —see section 3.

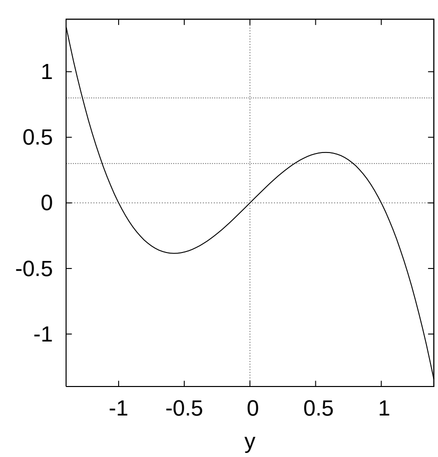

If z_* is a fixed point of the differential equation $\dot{z} = f(z)$, then approximate solutions near z_* can be obtained by linearising the equation. If $z(t) = z_* + \delta z(t)$ denotes a solution initially close to the fixed point, that is, if the norm of the vector $\delta z(t) = z(t) - z_*$ is small when $|t|$ is small, then expanding the vector field near z_* in Taylor series yields

$$\begin{aligned} \delta \dot{z} &= \dot{z} = f(z) = f(z_* + \delta z) \\ &= f(z_*) + \mathrm{D}f(z_*)\delta z + O(\delta z^2) = \mathrm{D}f(z_*)\delta z + O(\delta z^2), \end{aligned} \tag{12}$$

where $\mathrm{D}f(z_*)$ denotes the Jacobian matrix of f at z_*. The expression $O(\delta z^2)$ contains all monomials in $\delta x, \delta y$ of total degree at least 2 ($\delta x^2, \delta y^2, \delta x \delta y, \delta x^3$, etc.). Since $\|\delta z\|$ is small, the nonlinear contributions can be neglected, and the linearised equation

$$\delta \dot{z} = \mathrm{D}f(z_*)\delta z, \tag{13}$$

called the variational equation, provides precious information on the behaviour of the solutions.

We have three scenarios. If all the eigenvalues of the Jacobian have negative real part, then the increment $\delta z(t)$ decays exponentially (the fixed point solution is stable), and the linear approximation remains valid for all times (see example 1.8). If one of the eigenvalues has a positive real part, then the increment increases exponentially for almost every initial condition (see example 1.7). Eventually, the orbit will move away from z_*, where the linear approximation is no longer valid, and such an exponential behaviour can no longer be inferred. Finally, if all eigenvalues of the Jacobian have zero real part, then the linearised system provides limited information, since the dynamics will be determined by higher-order terms in the expansion (12).

Example 2.2 We return to the fixed points of the forced van der Pol oscillator, discussed in example 2.1. The Jacobian matrix is given by

$$\mathrm{D}f(x,y) = \begin{pmatrix} 1 - 3x^2 - y^2 & -\sigma - 2xy \\ \sigma - 2xy & 1 - x^2 - 3y^2 \end{pmatrix}. \tag{14}$$

Specialising this expression to the parameter ranges $\sigma = 0, \gamma > 0$, and the fixed points (11), we obtain

$$\mathrm{D}f(0,y_*) = \begin{pmatrix} 1 - y_*^2 & 0 \\ 0 & 1 - 3y_*^2 \end{pmatrix}, \qquad y_* - y_*^3 = \gamma. \tag{15}$$

There is a fixed point with $y_* < -1$ for all positive values of γ, and clearly the corresponding eigenvalues of the Jacobian matrix are both negative. Thus this fixed point is a stable node. For $0 < \gamma < 2/(3\sqrt{3})$ there are two additional fixed points with $0 < y_*^2 < 1/3$ and $1/3 < y_*^2 < 1$, respectively, the first one being an unstable node, the second one a saddle.

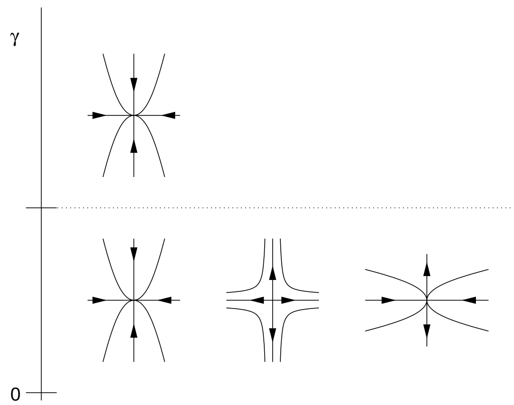

Two issues need to be resolved. First, the above heuristic arguments must be made more rigorous; we shall do so in the next session, with the Hartman-Grobman theorem. Second, we would like to piece together the various local information obtained by linearisation, to obtain a global qualitative view of the dynamics. We shall provide a glimpse of the possibilities in section 3.2.

2.2 Conjugacy and invariant manifolds

Informally, a manifold is a set that, locally, looks like Euclidean space $\mathbb{R}^n$, for some n. This property gives local coordinates, so that on a manifold we can do calculus and geometry. A sphere is a two-dimensional manifold because, locally, it looks like $\mathbb{R}^2$ (hence people believed that the Earth was flat). The circle, the torus, the cylinder are also two-dimensional manifolds. The eigenspaces of a saddle point are one-dimensional manifolds. By contrast, two lines intersecting transversally do not form a manifold, because there is no neighbourhood of the point of intersection where this set looks like $\mathbb{R}^n$, for any n.

More precisely, a set M is a differentiable manifold if it is provided with a finite or countable collection of *charts*, such that every point is represented in at least one chart. A chart is an open set $U \subset \mathbb{R}^n$ together with a one-to-one mapping $\phi : U \to \phi(U) \subset M$. If z and z' in two charts U and U' have the same image in M, then z and z' must have neighbourhoods $V \subset U$ and $V' \subset U'$ with the same image in M. In this way we get a map $(\phi')^{-1} \circ \phi : V \to V'$. Two charts U and U' are *compatible* if the n components of this map are differentiable[5]. An atlas is a union of compatible charts, and two atlases are equivalent

[5]Adjacent charts of a road atlas clearly have this property.

if their union is also an atlas. Finally, a differentiable manifold is a class of equivalent atlases.

To analyse structural similarities between dynamical systems we shall require two classes of functions.

Definition 2.1 *A bijective map f is a homeomorphism, if both f and f^{-1} are continuous. The map f is a diffeomorphism if both f and f^{-1} have continuous first derivatives.*

Let $f : X \to Y$ be a function. For a homeomorphism we only need continuity, hence it suffices that X and Y be topological spaces; for a diffeomorphism, X and Y must be smooth manifolds, because we must be able to differentiate.

Example 2.3 Consider the two differential equations (on $\mathbb{R}^+$)

$$\dot{x} = -x, \qquad \dot{y} = -y/2, \quad (x, y \geqslant 0)\,.$$

The second equation can be transformed into the first by the substitution $y = h(x) = \sqrt{x}$, which is bi-continuous in the given range. The two flows

$$\Phi_t(x_0) = e^{-t}x_0, \qquad \Psi_t(y_0) = e^{-t/2}y_0$$

are related via the function h by the equation $\Psi_t(h(x_0)) = h(\Phi_t(x_0))$.

Definition 2.2 *The flows Φ_t and Ψ_t on the spaces X and Y, respectively, are said to be topologically conjugate, if there exists a homeomorphism $h : X \to Y$ such that $\Psi_t \circ h = h \circ \Phi_t$, for all t.*

The attribute 'topological' refers to the continuity of h, and corresponding orbits of conjugate flows, namely, $\Phi_t(z_0)$ and $\Psi_t(h(z_0))$, are distorted versions of one another. Topological conjugacy can be defined locally, that, is, in a neighbourhood U of some point in phase space. In this case the relation $\Psi_t \circ h = h \circ \Phi_t$, will hold for all t such that $\Phi_t(z_0) \in U$.

Definition 2.3 *A fixed point z_* of the differential equation $\dot{z} = f(z)$ is said to be hyperbolic if all eigenvalues of the Jacobian matrix $\mathrm{D}f(z_*)$ have non-zero real part.*

The term hyperbolic is unfortunate, because it suggests that the fixed point is a saddle (where some eigenvalues have positive real part, some have negative real part, and —in some cases— the orbits are hyperbolae). The above definition does not require this. The following result is the rigorous justification of linear stability analysis.

Theorem 2.1 (Hartman-Grobman) *Let z_* be a hyperbolic fixed point of the differential equation $\dot{z} = f(z)$. Then there exists a neighbourhood U of z_* such that the flow of the differential equation is topologically conjugate in U to the flow of the variational equation $\dot{w} = \mathrm{D}f(z_*)w$. In two-dimensions, if f is smooth, then so is the conjugacy function.*

Definition 2.4 *Let z_* be a hyperbolic fixed point of the flow Φ_t, and let U be a neighbourhood of z_*. The sets*

$$\begin{aligned} W^s_U(z_*) &= \{z \in U : \Phi_t(z) \to z_* \text{ as } t \to \infty \text{ and } \Phi_t(z) \in U \text{ for all } t \geqslant 0\} \\ W^u_U(z_*) &= \{z \in U : \Phi_t(z) \to z_* \text{ as } t \to -\infty \text{ and } \Phi_t(z) \in U \text{ for all } t \leqslant 0\} \end{aligned}$$

are called (local) stable and unstable manifolds of z_* in U.

Thus the stable (unstable) manifold of a hyperbolic fixed point z_* is the set of points in U whose ω- (α-) limit set is z_*, see section 1.1, provided that the forward (backward) orbit of these points never leaves U.

The Hartman-Grobman theorem implies that the stable and unstable manifolds are tangent to the linear eigenspaces of the Jacobian at z_*. One can extend the local manifolds to global manifolds by transportation through the flow:

$$W^s(z_*) = \bigcup_{t \leqslant 0} \Phi_t(W^s_U(z_*)) \qquad \text{and} \qquad W^u(z_*) = \bigcup_{t \geqslant 0} \Phi_t(W^u_U(z_*)).$$

2.3 Structural stability

In example 2.2 we have observed an important phenomenon: the dynamical behaviour does not necessarily change when a parameter is varied, or when a perturbation is applied. This is a form of stability, called structural stability.

Definition 2.5 *The differential equation $\dot{z} = f(z)$ is structurally stable if for any function g of class C^1, and any sufficiently small $\varepsilon > 0$, the flow of $\dot{z} = f(z) + \varepsilon g(z)$ is conjugate to the flow of $\dot{z} = f(z)$.*

If the property above holds only in a neighbourhood of a point, then we call the flow locally structurally stable.

Let z_* be a hyperbolic fixed point of $\dot{z} = f(z)$s; in particular, the Jacobian $\mathrm{D}f(z_*)$ is not singular. Then the equation $f_\varepsilon(z) = f(z) + \varepsilon g(z) = 0$ has a solution z'_* close to z_*, from the implicit function theorem; that is, the perturbed differential equation $\dot{z} = f_\varepsilon(z)$ also has a fixed point. Since $\mathrm{D}f(z_*)$ has eigenvalues with non-vanishing real parts, and these eigenvalues depend continuously on matrix elements, the same holds for the eigenvalues of $\mathrm{D}f_\varepsilon(z'_*)$. Thus both fixed points z_* and z'_* are hyperbolic. The Hartman-Grobman theorem tells us that $\dot{z} = f(z)$ and $\dot{z} = f_\varepsilon(z)$ are conjugate to the corresponding variational equations, $\delta\dot{z} = \mathrm{D}f(z_*)\delta z$ and $\delta\dot{z} = \mathrm{D}f_\varepsilon(z'_*)\delta z$, respectively. Next we have the following result.

Theorem 2.2 *A necessary and sufficient condition for the topological equivalence of two linear systems, all of whose eigenvalues have non zero real part, is that the number of eigenvalues with negative (hence positive) real part be the same in both systems.*

In particular, all asymptotically stable fixed points are topologically equivalent.

The Jacobians $\mathrm{D}f(z_*)$ and $\mathrm{D}f_\varepsilon(z'_*)$ satisfy the assumption of the above theorem. It follows that for all sufficiently small ε, the systems $\dot{z} = f(z)$ and $\dot{z} = f_\varepsilon(z)$ are topologically conjugate in a neighbourhood of the respective fixed points. We have established that the system $\dot{z} = f(z)$ is structurally stable in a neighbourhood of a hyperbolic fixed point. In particular, saddle points are locally structurally stable. One can extend statements of this type to more general invariant hyperbolic structures.

2.4 Hyperbolic structures for diffeomorphisms

Hyperbolic fixed points and stable/unstable manifolds can also be introduced for discrete-time systems, more precisely for diffeomorphism (see definition 2.1). In presence of complicated dynamics, these are simpler to handle than flows.

Let $f : M \to M$ be a diffeomorphism of a manifold M. A fixed point z_* is a point of M such that $z_* = f(z_*)$. Expanding f in Taylor series near z_*, and neglecting all non-linear terms, we derive the discrete-time analogue of the variational equation (13):

$$w_{t+1} = \mathrm{D}f(z_*)w_t, \qquad w = z - z_*. \tag{16}$$

As in the continuous time case, the stability of the fixed points, and their classification are based on the study of the eigenvalues of the Jacobian $\mathrm{D}f$, and its Jordan form.

However, the superficial similarity between the variational equations (13) and (16) is misleading, and the stability criterion is quite different in the two cases. Comparing the respective solutions makes this clear [cf. (10)]:

$$\delta z(t) = e^{\mathrm{D}f(z_*)t}\delta z(0) \qquad\qquad w_t = (\mathrm{D}f(z_*))^t w_0.$$

Thus the stability condition associated to the eigenvalues λ of the Jacobian is given by $|\lambda| < 1$ in the discrete time case (16), whereas in the continuous time case (13) is $\mathrm{Re}(\lambda) < 0$.

The definition of conjugacy is simpler in the discrete-time case (cf. definition 2.2).

Definition 2.6 *The maps f and g on the spaces X and Y, respectively, are said to be topologically conjugate, if there exists a homeomorphism $h: X \to Y$ such that $g \circ h = h \circ f$.*

Note that there is no explicit reference to time. Indeed if f and g are conjugate, then

$$(g \circ g) \circ h = g \circ h \circ f = h \circ (f \circ f).$$

that is, $f \circ f$ and $g \circ g$ are also conjugate via h. Then, writing f^n for Φ_n, as it is customary for maps, an easy induction shows that $h \circ f^n = g^n \circ h$ for all n.

In place of definition 2.3, we now have

Definition 2.7 *A fixed point z_* of a diffeomorphism f is said to be hyperbolic if all eigenvalues of the Jacobian matrix $\mathrm{D}f(z_*)$ have absolute value different from 1.*

Finally, there is a Hartman-Grobman theorem for diffeomorphisms.

Example 2.4 Arnold 'cat map'. We consider the following linear map of the two-dimensional torus $\mathbb{T}^2 = \mathbb{S} \times \mathbb{S}$:

$$\begin{aligned} x_{n+1} &\equiv x_n + y_n \pmod 1 \\ y_{n+1} &\equiv x_n \pmod 1. \end{aligned}$$

This map has the fixed point $z_* = (0,0)$. Because of linearity, the map and its Jacobian coincide, and the variational equation in matrix form reads:

$$\begin{pmatrix} x_{n+1} \\ y_{n+1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} x_n \\ y_n \end{pmatrix}.$$

The eigenvalues and eigenvectors are given by

$$\lambda^u = \frac{1+\sqrt{5}}{2}, \qquad \mathbf{u}^u = \begin{pmatrix} 2 \\ -1+\sqrt{5} \end{pmatrix}$$

$$\lambda^s = \frac{1-\sqrt{5}}{2}, \qquad \mathbf{u}^s = \begin{pmatrix} 2 \\ -1-\sqrt{5} \end{pmatrix}$$

where the superscripts u, s refer to stable and unstable, since $|\lambda^s| < 1$ and $|\lambda^u| > 1$, respectively (recall that for maps these are the conditions for stability/instability).

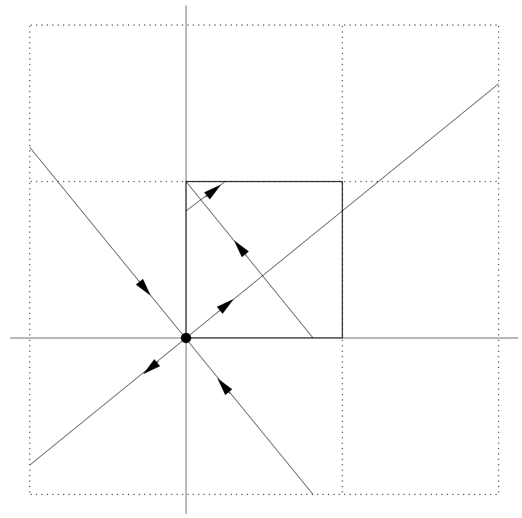

On the torus, the stable and unstable manifold of the fixed point intersect, giving rise to homoclinic points. They are defined just as in the continuum time case, namely their forward and backward orbits approach the fixed point at the origin. The number of homoclinic points is infinite, because the stable and unstable manifolds have irrational slope, and hence they never close up on the torus. It is possible to show that such homoclinic points are *dense* on the torus.

The construct of stable and unstable manifolds is not restricted to fixed points: it can also be defined for orbits of dynamical systems (and indeed to invariant sets, namely sets which are not changed by the dynamics).

Definition 2.8 *Let $z_{n+1} = f(z_n)$ be a diffeomorphism. The sets*

$$\begin{aligned} W^s(z) &= \{z' : \|f^n(z') - f^n(z)\| \to 0 \text{ as } n \to \infty\} \\ W^u(z) &= \{z' : \|f^n(z') - f^n(z)\| \to 0 \text{ as } n \to -\infty\} \end{aligned}$$

are called stable and unstable sets of the orbit $z_n = f^n(z)$.

For the cat map, these stable and unstable manifolds are determined solely by the eigenvectors of the Jacobian, and the latter does not depend on co-ordinates. For this reason, at every point in phase space there are two directions, the stable and unstable direction, along which the dynamics contracts and expands the space, respectively. Such a structure is called a *hyperbolic structure*; its presence is regarded as the signature of chaotic dynamics —see section 5.

3 Bifurcations

We investigate the topological changes that occur in dynamical systems when one or more parameters are varied. These transitions involve a violation of hyperbolicity, We shall consider three scenarios:

1. A pair of fixed points appears out of nowhere (saddle-node bifurcation).

2. A stable fixed point becomes unstable while it ejects a limit cycle (Hopf bifurcation).

3. A limit cycle disappears by colliding with an unstable fixed point (homoclinic bifurcation).

The analysis of the locus of parameter values at which these phenomena occur will provide valuable information on the structure of parameter space.

3.1 Centre manifolds

The behaviour in a neighbourhood of a hyperbolic fixed point z_* is captured by the linear part of the dynamics, characterised by stable and unstable manifolds tangent to the linear eigenspaces of the Jacobian matrix. In fact, using linear superposition, the flow can be decomposed into parts associated with the stable and unstable subspaces (formally, a direct sum of one-dimensional flows). A similar statement applies if the Jacobian has eigenvalues with vanishing real part. In this case however, in addition to stable and unstable manifolds, the so-called centre manifold appears, which contains the dynamics associated to the eigenvalues with vanishing real part. In presence of a centre manifold, the dynamics is no longer structurally stable, namely, its topological properties may change under the effect of an arbitrarily small perturbation.

Theorem 3.1 (Centre Manifold Theorem) *Let z_* be a fixed point of the differential equation $\dot{z} = f(z)$. Let E^s, E^u, and E^c denote the linear eigenspaces of the Jacobian matrix $\mathrm{D}f(z_*)$, which correspond to eigenvalues $Re(\lambda) < 0$, $Re(\lambda) > 0$, and $Re(\lambda) = 0$ respectively. Then there exist smooth invariant manifolds W^s and W^u tangent to E^s and E^u at z_*, and an invariant manifold W^c tangent to E^c. The stable and unstable manifolds W^s and W^u are unique.*

In the non-hyperbolic case, the main problem is to determine the dynamics on the centre manifold, since the dynamics on the other manifolds is known up to topological conjugacy. In particular, if the unstable manifold is empty ($\mathrm{Re}(\lambda) \leqslant 0$), then the orbits near z_* converge exponentially to the centre manifold, and the long-time dynamics near z^* is captured by a lower-dimensional dynamical system.

Example 3.1 We consider the following nonlinear system:

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= -y - x^3. \end{aligned}$$

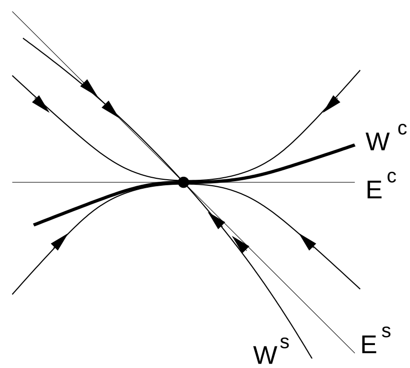

There is an equilibrium point $z_* = (0,0)$, with Jacobian

$$\mathrm{D}f(0,0) = \begin{pmatrix} 0 & 1 \\ 0 & -1 \end{pmatrix}.$$

The linear analysis gives:

$$\lambda^s = -1, \quad \mathbf{u}^s = \begin{pmatrix} -1 \\ 1 \end{pmatrix}, \qquad \lambda^c = 0, \quad \mathbf{u}^c = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

The centre manifold is tangent to $\mathbf{u}^c$ at the origin, so we let

$$y = h(x) = \sum_{k \geqslant 0} a_k x^k = 0 + 0 \cdot x + a_2 x^2 + a_3 x^3 + \cdots$$

Since the centre manifold is invariant, we write:

$$\dot{y} = -h(x) - x^3 = \frac{\mathrm{d}h}{\mathrm{d}x}\dot{x} = \frac{\mathrm{d}h}{\mathrm{d}x} h(x)$$

which gives

$$\begin{aligned} -a_2 x^2 - (a_3 + 1)x^3 + O(x^4) &= (2a_2 x + O(x^2))(a_2 x^2 + O(x^3)) \\ &= 2a_2^2 x^3 + O(x^4). \end{aligned}$$

Thus $a_2 = 0$, $a_3 = -1$ and

$$\dot{x} = y = h(x) = -x^3 + O(x^4).$$

We see that the solution $x(t)$ tends to zero as $t \to \infty$, so the nonlinear terms make the origin stable (nonlinear stability).

3.2 Bifurcations of fixed points

Let $\dot{z} = f_\alpha(z)$ be a differential equation in $\mathbb{R}^n$, where f_α depends smoothly on a real parameter α. We assume that there is a fixed point $z_* = z_*(\alpha)$, whose Jacobian matrix also depends on α. Suppose that for an isolated critical parameter value $\alpha = \alpha_*$ the fixed point z_* is non-hyperbolic. Several possibilities exist for the behaviour of the system at and near α_*.

3.2.1 Saddle-node bifurcation

Suppose that Jacobian matrix $\mathrm{D}f_{\alpha_*}(z_*)$ has a single zero eigenvalue:

$$f_{\alpha_*}(z_*) = \mathbf{0} \qquad \text{and} \qquad \det \mathrm{D}f_{\alpha_*}(z_*) = 0\,.$$

where $\mathbf{0} = (0, \ldots, 0) \in \mathbb{R}^n$. This is a system of $n+1$ equations in as many unknowns, namely the n coordinates of the fixed point z_*, and the critical parameter α_*.

Assume that there is a single eigenvalue with vanishing real part. Then the dynamics in the vicinity of z_* is captured by a differential equation $\dot{x} = h(x)$ on a one-dimensional centre manifold. Choosing the coordinates so that the origin is the fixed point, the constant term of h vanishes. By assumption, the linear term also vanishes, while the coefficient of the quadratic term —if it is not equal to zero— can be normalised to ± 1 by rescaling the time. Then the equation of motion on the centre manifold reads $\dot{x} = \pm x^2 + O(x^3)$. Let us choose the negative sign (the other case is analogous). If we consider parameter values close to α_*, then the equation of motion on the centre manifold will contain additional small contributions, $\dot{x} = a_0 + a_1 x - x^2 + O(x^3)$, a so-called *unfolding*. Introducing the coordinate change $w = x - c$, one obtains $\dot{w} = a_0 + a_1 c - c^2 + (a_1 - 2c)w - w^2 + O(w^3)$. By choosing $c = a_1/2$, the linear term can be eliminated and one ends up with the so-called normal form of the bifurcation:

$$\dot{w} = \mu - w^2\,. \tag{17}$$

Here the parameter μ is a function of the original parameter α of the system. The expression tells us that there is no fixed point for $\mu < 0$ and a pair of stable and unstable fixed points $w_* = \pm\sqrt{\mu}$ for $\mu > 0$. This means that if we cross the bifurcation point $\mu = 0$ in parameter space, a pair of fixed points is generated/destroyed: a saddle-node bifurcation.

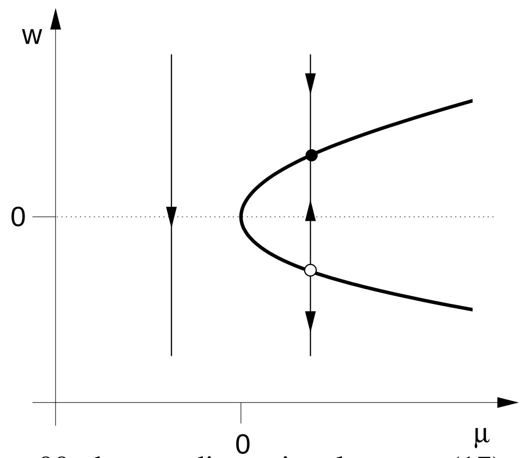

This terminology is justified as follows. For $\mu > 00$, the one-dimensional system (17) on the centre manifold has a stable and an unstable fixed point. If the motion in the transversal direction is expanding, then the former becomes a saddle, and the latter an unstable node; if it is contracting, the former becomes a stable node, and the latter a saddle. In any case, there is a saddle and a node.

The saddle-node bifurcation is an instance of a codimension-one bifurcation, meaning that the critical parameter set (a point) has one dimension fewer than the parameter space (a line). The same would apply for a critical parameter curve in a two-dimensional parameter space, etc.

The following result makes the above reasoning rigorous.

Theorem 3.2 (Saddle-node bifurcation) *Let $\dot{z} = f_\alpha(z)$ be a differential equation depending smoothly on a parameter α. At $\alpha = \alpha_*$, we assume that there exists a fixed point z_* such that*

SN1: $\mathrm{D}f_{\alpha_*}(z_*)$ *has a simple eigenvalue* 0 *with right eigenvector* $\mathbf{v}$ *and left eigenvector* $\mathbf{w}$. *All other eigenvalues have non-zero real part;*

SN2: $\mathbf{w} \cdot \left. \dfrac{\partial f_\alpha(z)}{\partial \alpha} \right|_{(z_*, \alpha_*)} \neq 0$ *(transversality condition);*

SN3: $\mathbf{w} \cdot \mathrm{D}^2 f_{\alpha_*}(z_*)(\mathbf{v}, \mathbf{v}) \neq 0$ *(quadratic non-degeneracy).*

Then there exists a smooth curve of fixed points $\xi_*(\alpha)$ with $\xi_*(\alpha_*) = z_*$. Depending on the signs of the expressions in SN2 *and* SN3*, there are no fixed points near z_* if $\alpha < \alpha_*$ (or $\alpha > \alpha_*$). The two fixed points near α_* are hyperbolic.*

The condition SN3 is defined via

$$\mathbf{w} \cdot \mathrm{D}^2 f(z)(\mathbf{u}, \mathbf{v}) = \sum_{k,l,n} w_k \frac{\partial^2 f^{(k)}(z)}{\partial x_n \partial x_l} u_n v_l \qquad f = (f^{(1)}, \ldots, f^{(n)}).$$

Example 3.2 Forced van der Pol oscillator. The fixed point equation can be written as [see eq. (2)]

$$\begin{pmatrix} 0 \\ \gamma \end{pmatrix} = \begin{pmatrix} 1 - r^2 & -\sigma \\ \sigma & 1 - r^2 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \qquad r^2 = x^2 + y^2.$$

Regarding this expression as a linear system and solving for x and y, we obtain an equation for r^2, the square distance from the origin of the fixed point $z_* = (x_*, y_*)$:

$$r^2 = \frac{\gamma^2}{(1 - r^2)^2 + \sigma^2}.$$

We now require that the Jacobian matrix (14) has a vanishing eigenvalue, namely zero determinant:

$$0 = \det(\mathrm{D}f(z_*)) = (1 - r^2)^2 - 2r^2(1 - r^2) + \sigma^2.$$

Re-arranging the last two equations, we obtain

$$\begin{aligned} \gamma^2 &= 2r^4(1 - r^2) \\ \sigma^2 &= (3r^2 - 1)(1 - r^2). \end{aligned}$$

These expressions can be viewed as a parametric representation of a curve in the γ-σ plane, with the curve parameter r^2 in the range $1/3 \leqslant r^2 \leqslant 1$. The values $r^2 = 1, 2/3$, and $1/3$ correspond, respectively, to the origin, the cusp, and the upper intersection with the γ-axis.

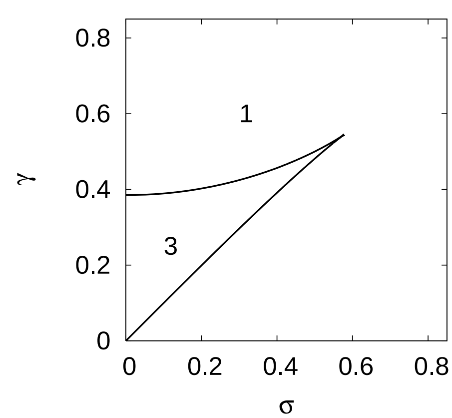

Using the analysis for $\sigma = 0$ (example 2.2, p. 15), and structural stability, we can identify the region in parameter space where the pair of fixed points exists, without additional computation. The presence of a cusp in the bifurcation diagram suggests that our analysis is still incomplete. We shall continue it in the next section.

3.2.2 Hopf bifurcation

The second possibility by which a hyperbolic fixed point can change stability is via a complex conjugate pair of eigenvalues (vanishing real part, but non-vanishing imaginary part). Suppose α_* is a parameter value where such an instability takes place. The (non-hyperbolic) fixed point and the parameter obey

$$f_{\alpha_*}(z_*) = \mathbf{0} \qquad \text{and} \qquad \det(\mathrm{D}f_{\alpha_*}(z_*) - i\omega\mathbf{1}) = 0,$$

where $\omega \neq 0$ denotes the imaginary part of one of the eigenvalues. These $n+2$ real equations for the unknowns (z_*, α_*, ω) give again one algebraic constraint for the parameter value, that is, the equations determine a codimension-one manifold in parameter space where the instability occurs.

A complex conjugate pair implies a two-dimensional centre manifold (see also example 1.6). The corresponding equation of motion on the centre manifold and the unfolding in a neighbourhood of the bifurcation point requires some substantial —if straightforward— computations. The resulting two-dimensional differential equation is most conveniently written in polar coordinates ($x = r\cos(\varphi)$, $y = r\sin(\varphi)$) as

$$\dot{r} = \mu r \pm r^3 + O(r^4), \qquad \dot{\varphi} = \omega + O(r).$$

The quantity μ is the unfolding parameter, while the sign in front of the cubic term characterises two variants of the same bifurcation.

Let us first consider the negative sign (the so-called *supercritical* case). The fixed point $r = 0$ [that is, $(x,y) = (0,0)$] is stable for $\mu < 0$ and unstable for $\mu > 0$. If $\mu > 0$ then there is a stable periodic solution $r = \sqrt{\mu}$, $\varphi = \omega t$, given by $x(t) = \sqrt{\mu}\cos(\omega t)$, $y(t) = \sqrt{\mu}\sin(\omega t)$, namely a limit cycle.

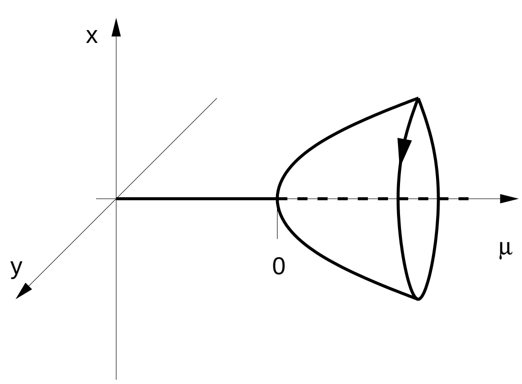

If the cubic coefficient is $+1$ (the so-called *sub-critical* case), then the stability properties of the fixed points are the same, but the system develops an unstable periodic solution $r = \sqrt{-\mu}$, $\varphi = \omega t$, i.e., $Z(t) = \sqrt{-\mu}(\cos(\omega t), \sin(\omega t))$ for $\mu < 0$.

This type of bifurcation, whereby the loss of stability of a fixed point gives rise to a limit cycle, is called a Hopf bifurcation. The analysis presented here be made rigorous with a few additional technical assumptions.

Example 3.3 Let us return to the analysis of the bifurcations of the forced van der Pol oscillator. The condition for a Hopf bifurcation in a two-dimensional system, that is, a complex conjugate imaginary pair of eigenvalues, results in $\mathrm{Tr}(\mathrm{D}f(z_*)) = 0$ and $\det(\mathrm{D}f(\underline{z}_*)) > 0$. With the Jacobian matrix (see example 3.2) we thus obtain

$$0 = 2 - 4r^2, \qquad (1-r^2)^2 - 2r^2(1-r^2) + \sigma^2 > 0\,.$$

With the condition for the fixed point (see example 3.2), we arrive at

$$r^2 = \frac{1}{2}, \qquad \sigma^2 > \frac{1}{4}, \qquad \gamma^2 = \frac{1}{2}\left(\frac{1}{4} + \sigma^2\right),$$

which is a ray in the σ^2-γ^2 plane, originating from the point $r^2 = 1/2$ on the saddle-node curve, located on the upper branch, near (but not at) the cusp.

We can now complete the analysis of the bifurcation diagram of the forced van der Pol oscillator, by re-visiting examples 1.4, 2.2, and 3.2.

Within the domain bounded by the saddle-node bifurcation line the flow contains a stable node/focus (SN), an unstable node/focus (UN), and a saddle (SD). Furthermore, example 1.4 tells us that the flow possesses a stable limit cycle (SLC) at $\gamma = 0$. That suggests that the limit cycle is destroyed at the Hopf bifurcation line in a subcritical Hopf bifurcation leaving a stable focus (SF) behind.

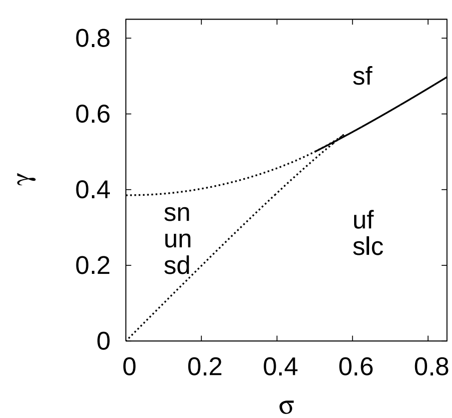

We have been therefore able to gain substantial insight into the phase portrait (up to topological equivalence) without actually integrating the underlying differential equations. There are still exceptional points in the bifurcation diagram, such as the point where the Hopf bifurcation line terminates at the saddle-node bifurcation, or the cusp formed by the merging of two saddle-node bifurcation lines. These degenerate cases are related to codimension-two bifurcations.

3.3 A homoclinic bifurcation

So far we have considered changes in the topology of a phase portrait related to changes of stability of fixed points. Such local changes may have global repercussions; for instance, the system (17) has no bounded orbits for negative parameter values, but if the parameter is positive, 'half' of the orbits are bounded. There are other mechanisms that can trigger global changes in the dynamics when a parameter is changed. In this section we illustrate the so-called homoclinic bifurcation for two-dimensional flows.

We consider a planar system with a saddle point, namely a hyperbolic fixed point with a stable and an unstable manifold. This system has the property that these manifolds have a non-empty intersection, which is a homoclinic orbit. Specifically, we consider the following nonlinear oscillator ('the fish'):

$$\dot{x} = y, \qquad \dot{y} = 2x - 3x^2. \tag{18}$$

There are two fixed points, $(0,0)$ and $(2/3,0)$. The Jacobian matrix reads

$$\mathrm{D}f(x,y) = \begin{pmatrix} 0 & 1 \\ 2-6x & 0 \end{pmatrix}$$

hence the former is a saddle (eigenvalues $\pm\sqrt{2}$ with stable and unstable eigenvectors $(1,\mp\sqrt{2})$), and the latter is a centre (complex conjugate eigenvalues on the unit circle). The Hamiltonian function[6]

$$\mathrm{H}(x,y) = \frac{1}{2}y^2 - x^2 + x^3 \tag{19}$$

is a conserved quantity; indeed

$$\dot{\mathrm{H}} = y\dot{y} + (-2x+3x^2)\dot{x} = y(2x-3x^2) + (-2+3x^2)y = 0.$$

[6]Physically, it represents the energy.

It follows that the level sets of the function H, namely the curves $\mathrm{H}(x,y) = \text{const.}$ are union of orbits. In particular, the level set $\mathrm{H}(x,y) = 0$ comprises four orbits, one of which is homoclinic. All other level sets consist of one or two orbits, depending on the value of H.

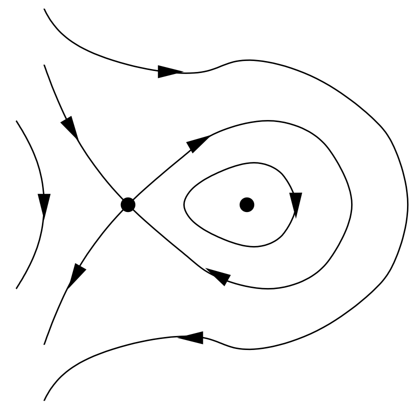

Next we modify equations (18) as follows (the 'dissipative fish'):

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= 2x - 3x^2 - \gamma y(\mathrm{H}(x,y) - \mu) \end{aligned}$$

where $\mathrm{H}(x,y)$ is given in (19), γ and μ are parameters, and $\gamma \geqslant 0$. The additional term plays the role of a 'damping force'. The function H is no longer invariant, and one finds:

$$\dot{\mathrm{H}} = -\gamma y^2(\mathrm{H}(x,y) - \mu),$$

giving

$$\mathrm{H}(x(t),y(t)) - \mu = \big(\mathrm{H}(x(0),y(0)) - \mu\big) \exp\Big(-\gamma \int_0^t y^2(t')\mathrm{d}t'\Big).$$

Therefore, as $t \to \infty$ either $y(t) \to 0$ or $\mathrm{H}(x(t),y(t)) \to \mu$. The first case corresponds to one of the fixed point solutions $(0,0)$ or $(2/3,0)$, or to the stable manifold of the saddle point $(0,0)$. This set of points is exceptional (zero Lebesgue measure).

The second possibility $\mathrm{H}(x(t),y(t)) \to \mu$, which affects almost all initial conditions, depends on the value of the parameter μ. If $\mu = \mu_0 = 0$, then the level curve $\mathrm{H}(x,y) = 0$ comprises a homoclinic orbit, approached by all points inside the resonant domain, as well as the unbounded orbit emanating from the saddle-point, approached by all points outside. In the parameter range $-4/27 < \mu < \mu_0 = 0$ ($-4/27$ is the value of H at the stable equilibrium) the equaiton $\mathrm{H}(x,y) = \mu$ determines a closed orbit, which is stable (a limit cycle), together with an unbounded orbit. By contrast, no closed orbit exists if $\mathrm{H}(x,y) < -4/27$ or $\mathrm{H}(x,y) > 0$.

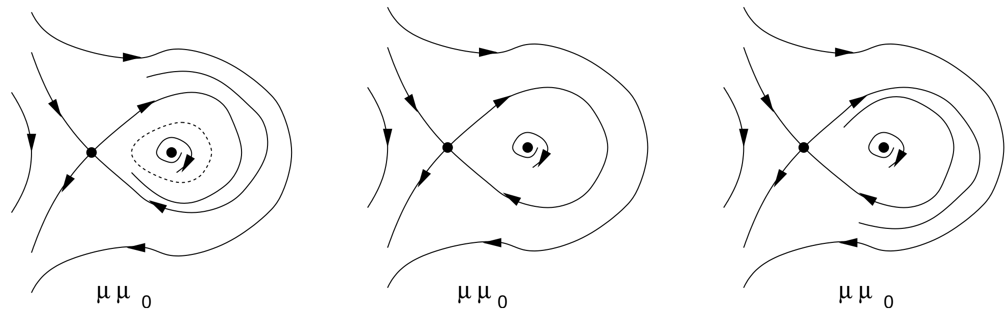

The change in the topology of the flow which occurs at $\mu_0 = 0$ does not involve a change of stability of fixed points, but rather the collision of a limit cycle with a saddle point; the former briefly morphs into a homoclinic orbit, and then disappears. When the limit cycle exists ($\mu < 0$), the bounded orbits form a two-dimensional domain in phase space, namely the strip bounded by the two branches of the stable manifold of the saddle point. Such a strip extends to infinity. By contrast, if $\mu > 0$, the bounded orbits consist of two fixed points, and the one-dimensional stable manifold of the saddle, with one branch coming from infinity, and the other spiralling away from the unstable focus.

The invariant sets of maps and higher-dimensional flows become more varied and exotic, and so do their bifurcations. As parameters change, one may observe infinite cascades of bifurcations —the mechanism for the creation of chaos— or collisions of strange invariant sets.

4 Symbolic dynamics

Symbolic dynamics is the study of sequences of symbols associated with orbits of dynamical systems. This is a tool of wide applicability, needed to handle the complex phenomena —much richer that those we have examined so far— found in the study of dynamical systems.

These phenomena first appear in three-dimensional flows, and their analysis present substantial difficulties. By contrast, complicated dynamics is already present in one-dimensional maps. For this reason, we now shift our attention to the latter, where there is a chance to develop the basic theory, and to illustrate several constructs which are relevant to more general systems.

Let $f : M \to M$, $x_{n+1} = f(x_n)$ be a discrete dynamical system on a set M. We recall some terminology and notation. The transformations Φ_n of definition 1.1 will now be denoted by f^n, with $f^0 = \mathbf{1}$, and $f^{n+1} = f \circ f^n$, that is

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

A point $x_* \in M$ is a fixed point if $x_* = f(x_*)$. If x_0 is a fixed point of f^n, for some $n \in \mathbb{N}$, then we say that the orbit of x_0 is periodic. Such an orbit is a sequence consisting of indefinite repetitions of a finite sub-sequence. If $(x_0, x_1, \ldots x_{n-1})$ are distinct, and $x_n = x_0$, then we say that n is the (minimal) period of the orbit.

4.1 Bernoulli shift map

We introduce a well-known 'toy-model' of irregular dynamics, a map of disarming simplicity, with orbits of seemingly unlimited complexity: the Bernulli shift map, or doubling map.

Let $I = [0, 1)$, and let $\mathrm{B} : I \to I$ be given by

$$\begin{aligned} x_{n+1} &= \mathrm{B}(x_n) \equiv 2x_n (\bmod 1) \\ &= \begin{cases} 2x_n & \text{if} \quad 0 \leqslant x_n < 1/2 \\ 2x_n - 1 & \text{if} \quad 1/2 \leqslant x_n < 1. \end{cases} \end{aligned}$$

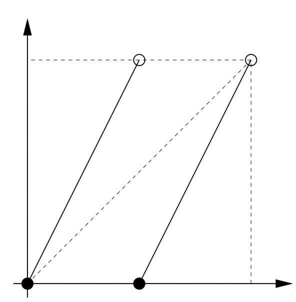

By identifying $[0,1)$ with the unit circle

$$\mathbb{S}^1 = \{z \in \mathbb{C} : |z| = 1\} = \{e^{2\pi i x} : x \in [0,1)\}$$

the map B becomes a continuous map of a compact domain. Indeed B is conjugate —via the exponential function— to the restriction to $\mathbb{S}^1$ of the analytic map $g(z) = z^2$:

$$g(z) = z^2 = e^{2\pi i 2x} = e^{2\pi i 2x (\text{mod } 1)} = e^{2\pi i \mathrm{B}(x)}.$$

On $\mathbb{S}^1$ we have the distance (metric):

$$|x - x'| = \min_{m \in \mathbb{Z}} \{|x - x' + m|\}.$$

To find the periodic points of B we fix $n \in \mathbb{N}$ and solve the congruence

$$x = \mathrm{B}^n(x) \equiv 2^n x (\text{mod } 1).$$

This gives $x = 2^n x - k$, for some $k \in \mathbb{Z}$, that is, $x = k/(2^n - 1)$, $k = 0, 1, \ldots, 2^n - 2$. Each value of k corresponds to a periodic point whose minimal period is a divisor of n. The set of all such points, denoted by $\mathrm{Per}_n(\mathrm{B})$, is given by

$$\mathrm{Per}_n(\mathrm{B}) = \{x \in [0,1) : \mathrm{B}^n(x) = x\} = \Big\{\frac{k}{2^n - 1} : k \in \mathbb{Z}, 0 \leqslant k \leqslant 2^n - 2\Big\}.$$

For $n = 1$, we have $k = 0$, and $x_* = 0$. For $n = 2$, we have again the fixed point ($k = 0$), while the values $k = 1, 2$ give one periodic orbit of minimal period 2, namely $\{1/3, 2/3\}$.

To see the bigger picture, we represent $x \in [0,1)$ with binary digits (base 2):

$$x = \frac{1}{2}\sigma_0 + \frac{1}{2^2}\sigma_1 + \frac{1}{2^3}\sigma_2 + \cdots \quad = \; 0.\sigma_0\sigma_1\sigma_2\ldots, \quad \sigma_k \in \{0,1\}.$$

For instance:

$$\frac{1}{3} = 0.010101\ldots = 0 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2^2} + 0 \cdot \frac{1}{2^3} + 1 \cdot \frac{1}{2^4} + \cdots = \sum_{k \geqslant 1} \frac{1}{4^k}.$$

How does $\mathrm{B}(x) \equiv 2x (\text{mod } 1)$ look like in binary?

$$\begin{aligned}
2x \quad &= \quad \sigma_0 + \frac{1}{2}\sigma_1 + \frac{1}{2^2}\sigma_2 + \cdots \quad &&= \; \sigma_0.\sigma_1\sigma_2\ldots \\
2x (\text{mod } 1) \quad &= \quad \frac{1}{2}\sigma_1 + \frac{1}{2^2}\sigma_2 + \frac{1}{2^3}\sigma_3 + \cdots &&= \; 0.\sigma_1\sigma_2\sigma_3\ldots.
\end{aligned}$$

The action of B on x amounts to a left shift of the binary digits of x, discarding the digit to the left of the radix point as a result of the mod operator.

We introduce the space of all semi-infinite sequences binary digits (which we write by juxtaposing digits):

$$\Omega = \{\sigma_0\sigma_1 \ldots : \sigma_k \in \{0,1\}\} \cong \{0,1\}^{\mathbb{N}}, \tag{20}$$

and we define the (left) shift map S on Ω:

$$\mathrm{S}(\sigma_0\sigma_1\sigma_2\ldots) = (\sigma_1\sigma_2\sigma_3\ldots).$$

Let $h : \Omega \to I$ be the map which sends a binary sequence to the corresponding real number. We have the commutative diagram

$$\begin{array}{ccc} \Omega & \xrightarrow{\mathrm{S}} & \Omega \\ \downarrow h & & \downarrow h \\ I & \xrightarrow{\mathrm{B}} & I \end{array}$$

namely, $\mathrm{B} \circ h = h \circ \mathrm{S}$. Moreover, the map h is continuous (with respect to the natural topology[7] on Ω). However, this relation does not provide a conjugacy, because h is not invertible, due to the identity

$$0.111111\ldots = \sum_{k\geqslant 1} \frac{1}{2^k} = 1 = 1.000000\ldots.$$

In other words, the integer 1 has two distinct binary representations, and hence so does any rational number whose denominator is a power of 2. (Think about it.)

This is not a big problem though; it suffice to remove from Ω all sequences which end up with indefinite repetitions of the digit 1. From this correspondence, we shall deduce with surprising ease many properties of the dynamics of B.

Every periodic symbol sequence yields a periodic orbit. An example will suffice to elucidate the general argument:

$$\begin{aligned} 0.\underbrace{00101}_{5}\underbrace{00101}\ldots &= \left(\frac{1}{2^3} + \frac{1}{2^8} + \frac{1}{2^{13}} + \cdots\right) + \left(\frac{1}{2^5} + \frac{1}{2^{10}} + \frac{1}{2^{15}} + \cdots\right) \\ &= \left(\frac{1}{2^3} + \frac{1}{2^5}\right) \times \sum_{k\geqslant 0} \left(\frac{1}{2^5}\right)^k = \frac{5}{31}. \end{aligned}$$

[7] A neighbourhood of $\sigma = \sigma_0\sigma_1\ldots$ is the set of all sequences which share the first n symbols with σ, for some $n \in \mathbb{N}$.

The orbit, of period 5, is obtained by applying repeatedly the map B to this rational number:

$$\frac{5}{31} \mapsto \frac{10}{31} \mapsto \frac{20}{31} \mapsto \frac{9}{31} \mapsto \frac{18}{31} \mapsto \frac{5}{31} \mapsto \cdots.$$

One can show that the periodic orbits are precisely the rational numbers in I with odd denominator. It follows that the periodic points are *dense* in I.

By the same device, the rational numbers with even denominator give rise to eventually periodic orbits, namely orbits which contain a periodic point (after which, all points must be periodic). Moreover, the exponent of the prime 2 at denominator of the initial point is equal to the amount of time it takes for the orbit to reach the first periodic point.

Now take any two points $x = 0.\sigma_0\sigma_1\sigma_2\dots$ and $x' = 0.\sigma_0'\sigma_1'\sigma_2'\dots$. The point y given by $0.\sigma_0\sigma_1\dots\sigma_{N-1}\sigma_0'\sigma_1'\sigma_2'\dots$ is close to x:

$$\begin{aligned} |x-y| &= \left|\frac{1}{2^{N+1}}(\sigma_N - \sigma_0') + \frac{1}{2^{N+2}}(\sigma_{N+1} - \sigma_1') + \cdots\right| \\ &\leqslant \frac{1}{2^{N+1}} + \frac{1}{2^{N+2}} + \cdots = \frac{1}{2^N}. \end{aligned}$$

By choosing N large enough, we can make $|x-y|$ as small as we please. But after N iterations, i.e., after N symbol shifts, y is mapped to x'. Thus we have an orbit which starts arbitrarily close x and eventually reaches x'. (We say that the Bernoulli map is topologically transitive, see p. 42.)

Exercise 4.3. Characterise all points in I whose ω-limit set under B is the origin.

Exercise 4.4. Let $n \geqslant 3$ be an odd integer, and consider the periodic orbit of B with initial condition $1/n$. Show that the maximum period such an orbit can have is $n-1$, and this can only occur if n is a prime number. Characterise these orbits geometrically. (The existence of infinitely many orbits of this type is an open problem.)

4.2 Expansive Markov maps

We want to generalise the construction developed in the previous section. Let $x_0 = 0.\sigma_0\sigma_1\sigma_2\dots$ be the initial condition of an orbit of the Bernoulli map B. The first symbol σ_0 tells us if $x_0 < \frac{1}{2}$ or $x_0 \geqslant \frac{1}{2}$. The second symbol σ_1 provides the same information for the point $\mathrm{B}(x_0)$, and so on. Thus we consider the following partition of the interval I:

$$I_0 = \left[0, \frac{1}{2}\right) \qquad I_1 = \left[\frac{1}{2}, 1\right).$$

The binary digits of the initial condition x_0 are related to the points in the orbit as follows:

$$\sigma_t = k \qquad \Leftrightarrow \qquad \mathrm{B}^t(x_0) \in I_k.$$

Let $I = [a,b]$ be a closed interval. We denote the interior (a,b) of I by $\mathrm{int}(I)$ and the length $|b-a|$ of I by $|I|$. A collection of closed intervals $\{I_0, I_1, \ldots, I_{N-1}\}$ is called a partition of I if

$$I = \bigcup_{k=0}^{N-1} I_k, \qquad \text{and} \qquad \mathrm{int}(I_k) \bigcap \mathrm{int}(I_\ell) = \emptyset, \quad k \neq \ell.$$

Definition 4.1 *A map $f: I \to I$ is called Markov map if there exists a partition $\{I_0, I_1, \ldots, I_{N-1}\}$ of I, with the property that, for all $k, \ell = 0, \ldots, N-1$, one of the following statements is true:*

$$i) \quad \mathrm{int}(I_\ell) \cap f(\mathrm{int}(I_k)) = \emptyset, \qquad ii) \quad \mathrm{int}(I_\ell) \subseteq f(\mathrm{int}(I_k)).$$

A partition with the above properties is called a Markov partition. A Markov map sends boundary points of the Markov partition to boundary points.

Example 4.1 We consider the map

$$f: [0,1] \to [0,1] \qquad f(x) = \begin{cases} 1-2x & \text{if} \quad 0 \leqslant x \leqslant \frac{1}{2} \\ x - \frac{1}{2} & \text{if} \quad \frac{1}{2} < x \leqslant 1. \end{cases}$$

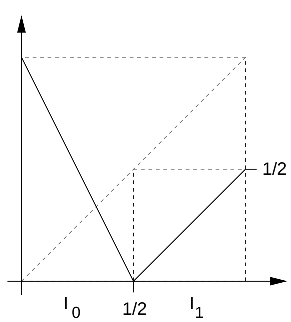

We define the partition $\{I_0, I_1\}$ as follows:

$$\begin{aligned} I_0 &= [0, \tfrac{1}{2}], \qquad \mathrm{int}(I_0) = (0, \tfrac{1}{2}) \\ I_1 &= [\tfrac{1}{2}, 1], \qquad \mathrm{int}(I_1) = (\tfrac{1}{2}, 1). \end{aligned}$$

One verifies that

$$f(\mathrm{int}(I_0)) = (0,1) \qquad \text{and} \qquad f(\mathrm{int}(I_1)) = (0, \frac{1}{2}),$$

from which we have

$$\begin{array}{ll} f(\mathrm{int}(I_0)) \supseteq \mathrm{int}(I_0) & f(\mathrm{int}(I_0)) \supseteq \mathrm{int}(I_1) \\ f(\mathrm{int}(I_1)) \supseteq \mathrm{int}(I_0) & f(\mathrm{int}(I_1)) \cap \mathrm{int}(I_1) = \emptyset. \end{array}$$

So it is not possible to reach I_1 from I_1 with a single iteration of the map, although it is possible to return to I_1 in two iterations.

Example 4.2 Let us modify the previous example as follows:

$$f : [0,1] \to [0,1] \qquad f(x) = \begin{cases} 1-2x & \text{if} \quad 0 \leqslant x \leqslant \frac{1}{2} \\ \frac{3}{2}x - \frac{3}{4} & \text{if} \quad \frac{1}{2} < x \leqslant 1. \end{cases}$$

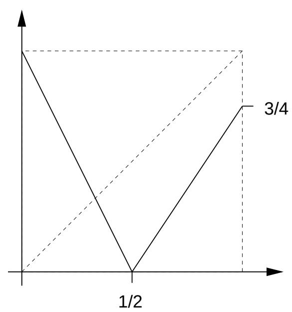

The partition $I_0 = [0, \frac{1}{2}]$, $I_1 = [\frac{1}{2}, 1]$ is not Markov. Indeed $f(\mathrm{int}(I_1)) = (0, \frac{3}{4})$, so that $f(\mathrm{int}(I_1)) \cap \mathrm{int}(I_1) = (\frac{1}{2}, \frac{3}{4}) \neq \emptyset$, and $\mathrm{int}(I_1) \not\subseteq f(\mathrm{int}(I_1))$.

(Is there another partition which is a Markov partition?)

Definition 4.2 *Let f be a Markov map. The $N \times N$ matrix defined by*

$$A_{k\ell} = \begin{cases} 1 & \textit{if} \quad f(\mathrm{int}(I_k)) \supseteq \mathrm{int}(I_\ell) \\ 0 & \textit{if} \quad f(\mathrm{int}(I_k)) \cap \mathrm{int}(I_\ell) = \emptyset \end{cases}$$

is called <u>topological transition matrix</u> of f.

Thus we have $A_{k,\ell} = 1$ if the transition $I_k \to I_\ell$ is permitted, and $A_{k,\ell} = 0$ if the transition is forbidden.

Example 4.3 Let us re-visit example 4.1

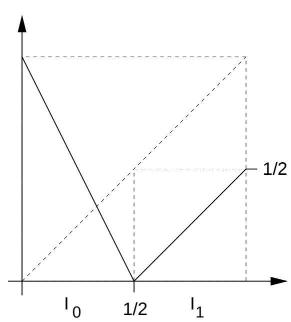

$$\begin{array}{l} I_0 \to I_0 \\ I_0 \to I_1 \\ I_1 \to I_0 \\ I_1 \not\to I_1 \end{array} \qquad \mathbf{A} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$

Definition 4.3 *A Markov map f is said to be expansive (or expanding) if f is smooth on the interior of each element of the Markov partition, and if there is $\lambda > 1$ such that $|f'(x)| \geqslant \lambda$ for all x lying in the interior of any partition element.*

It may happen that f is not expansive, but some iterate of f, say f^n, is expansive, that is, $|(f^n)'(x)| \geqslant \lambda > 1$. The term expansive is often used in this weaker sense.

Example 4.4 Let us consider again example 4.3.

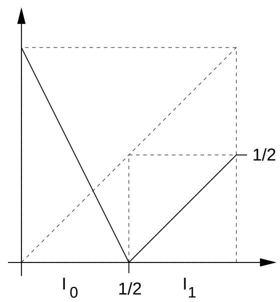

We have $|f'(x)| = 2$ if $x \in \mathrm{int}(I_0)$, but $|f'(x)| = 1$ if $x \in \mathrm{int}(I_1)$. Nevertheless, $|(f^2)'(x)| = |f'(f(x))||f'(x)| = 2 \cdot 1$ if $x \in \mathrm{int}(I_1)$, as $x \in I_1$ implies $f(x) \in I_0$, i.e., $f'(x) = 1$ and $f'(f(x)) = -2$. In addition $|(f^2)'(x)| \geqslant 2$ if $x \in \mathrm{int}(I_0)$. Thus the map is still expansive.

The idea of symbolic dynamics is as follows. We consider a map f with a partition $\{I_0, I_1, \ldots, I_{N-1}\}$ and symbols $\sigma \in \{0, 1, \ldots, N-1\}$. We assign to $x_0 \in I$ the symbol sequence $\sigma_0\sigma_1\sigma_2\cdots$ according to the rule $f^n(x_0) \in I_{\sigma_n}$. Then x_n has symbol sequence $\sigma_n\sigma_{n+1}\cdots$. In other words, the symbolic dynamics is just the left shift map.

For a Markov map with a Markov partition, we have a precise characterisation of the symbolic sequences produced by the dynamics.

Definition 4.4 *Let f be a Markov map with transition matrix* $\mathbf{A}$. *A symbol sequence $\sigma_0\sigma_1\sigma_2\ldots$ is called admissible if $A_{\sigma_n,\sigma_{n+1}} = 1$, for all $n \geqslant 0$. The shift space Ω of f is the set of all admissible symbolic sequences.*

For the Bernoulli shift map all binary sequences are admissible [equation (20), p. 34]. I this case we speak of a full shift.

Example 4.5 Let us consider the following transition matrix:

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

Then, any symbol sequences which contains the pair '11' is not admissible, since $A_{1,1} = 0$. For instance, the symbol sequence $0110010\ldots$ is not admissible. By contrast, any sequence without repeated 1s is admissible, for instance, the periodic sequence $\underbrace{001}\,001\ldots$.

Theorem 4.1 *Let $f : I \to I$ be a continuous expansive Markov map, and let Ω be the corresponding shift space. Then there is a surjective map $h : \Omega \to I$ such that, if $x = h(\sigma_0\sigma_1\ldots)$, then $f^n(x) \in I_{\sigma_n}$, for all $n \geqslant 0$.*

We illustrate the idea of proof with the tent map:

$$\mathrm{T} : [0,1] \to [0,1] \qquad \mathrm{T}(x) = 1 - |1-2x|. \tag{21}$$

The partition $I_0 = [0, \frac{1}{2}]$ and $I_1 = [\frac{1}{2}, 1]$ is Markov, with transition matrix

$$\mathbf{A} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

We see that any binary symbol sequence is admissible, and we have a full shift.

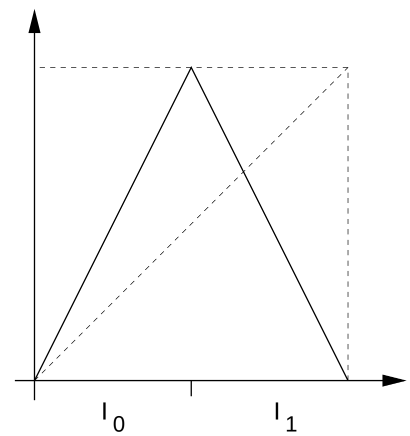

Now fix a symbol sequence, say $101\cdots$. Using this sequence, we construct a recursive sequence of intervals $I_1, I_{10}, I_{101}, \ldots$ as follows. The interval I_1 is an element of the Markov partition, with $|I_1| = \frac{1}{2}$.

Next we let $I_{10} = I_1 \cap \mathrm{T}^{-1}(I_0)$. This is the set of $x \in I$ such that $x \in I_1$ and $\mathrm{T}(x) \in I_0$. The interval I_{10} is closed, we have $|I_{10}| = \frac{1}{2} \cdot \frac{1}{2}$, $I_{10} \subseteq I_1$, and $\mathrm{T} : I_{10} \to I_0$ is bijective.

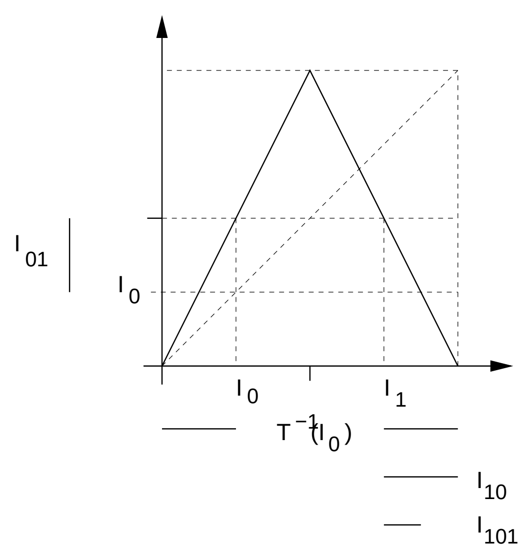

Likewise, let $I_{101} = I_{10} \cap \mathrm{T}^{-2}(I_1) = I_1 \cap \mathrm{T}^{-1}(I_{01})$: this is the set of $x \in I$ such that $x \in I_1$, $\mathrm{T}(x) \in I_0$, and $\mathrm{T}^2(x) \in I_1$.

The interval I_{101} is closed, we have $|I_{101}| = \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2}$, $I_{101} \subseteq I_{10}$, and $\mathrm{T} : I_{101} \to I_{01}$ is bijective. Continuing this process, we obtain an infinite nested sequence of closed intervals:

$$I_1 \supseteq I_{10} \supseteq I_{101} \supseteq \cdots$$

of length $1/2, 1/2^2, 1/2^3, \ldots$. Since the length of these intervals tends to zero, their intersection is the unique point $x \in I$ such that $\mathrm{T}^k(x) \in I_{\sigma_k}$ for all k. Now choose $x' \in I$. Since $\{I_0, I_1\}$ is a partition, there exists a symbol sequence such that $\mathrm{T}^k(x') \in I_{\sigma_k}$. Then $|\mathrm{T}^n(x') - \mathrm{T}^n(x)| = 2^n |x' - x|$ for any n, that is, $x' = x$. Thus, the symbolic dynamics is onto.

Now we see how the proof of theorem 4.1 will develop for a general Markov map f. Let $\sigma_0\sigma_1\sigma_2\ldots$ be an admissible symbol sequence. Then

- Admissible implies $\mathrm{int}(I_{\sigma_1}) \subseteq f(\mathrm{int}(I_{\sigma_0}))$.
- Continuity of f implies $I_{\sigma_1} \subseteq f(I_{\sigma_0})$. Thus $I_{\sigma_0\sigma_1} = I_{\sigma_0} \cap f^{-1}(I_{\sigma_1}) \neq \emptyset$, the interval $I_{\sigma_0\sigma_1}$ is closed, and $f : I_{\sigma_0\sigma_1} \to I_{\sigma_1}$ is one to one. Clearly, $I_{\sigma_0\sigma_1} \subseteq I_{\sigma_0}$.
- f expansive means $|I_{\sigma_1}| = |f(I_{\sigma_0\sigma_1})| \geqslant \lambda |I_{\sigma_0\sigma_1}|$. Thus $|I_{\sigma_0\sigma_1}| \leqslant \dfrac{1}{\lambda}|I_{\sigma_1}|$ $(\leqslant \dfrac{1}{\lambda^2}|I|)$.
- By induction ($\{I_{\sigma_0\sigma_1}\}$ is again a Markov partition, etc.), we obtain a nested sequence of closed intervals (the so-called <u>cylinder sets</u>) $I_{\sigma_0} \supseteq I_{\sigma_0\sigma_1} \supseteq I_{\sigma_0\sigma_1\sigma_2} \supseteq \ldots$, where $|I_{\sigma_0}| \leqslant \dfrac{1}{\lambda}|I|$, $|I_{\sigma_0\sigma 1}| \leqslant \dfrac{1}{\lambda^2}|I|$, ... Thus $\cap_{n\geqslant 0} I_{\sigma_0\sigma_1\ldots\sigma_n}$ contains one point x and $f^k(x) \in I_{\sigma_k}$ for $k \geqslant 0$.
- Uniqueness follows from expansivity: if both x and x' obey $f^k(x) \in I_{\sigma_k}$ and $f^k(x') \in I_{\sigma_k}$ for $k \geqslant 0$, then $|I| \geqslant |f^k(x) - f^k(x')| \geqslant \lambda^k |x - x'|$. This implies that $|x - x'| \to 0$, since $\lambda^k \to \infty$.
- Surjectivity follows from expansivity and the sets being a partition.

(Fill in all details.)

Example 4.6 We consider again example 4.1.

$$f(x) = \begin{cases} 1-2x, & x \in I_0 \\ x-\frac{1}{2}, & x \in I_1 \end{cases} \qquad I_0 = \left[0, \frac{1}{2}\right], \quad I_1 = \left[\frac{1}{2}, 1\right].$$

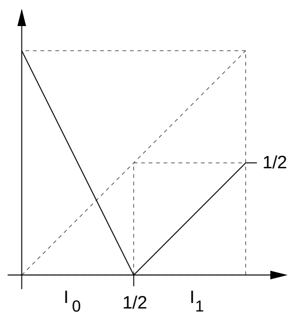

This is an (eventually) expansive Markov map, with transition matrix

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

The periodic orbits are the admissible periodic symbol sequences:

$0000\ldots$			fixed point
$010101\ldots$	$101010\ldots$		minimal period two
$001001\ldots$	$010010\ldots$	$100100\ldots$	minimal period three.

Let us compute the period-two orbit $\{x_0, x_1\}$ with

$$x_0 \cong 010101\ldots \in I_0 \qquad x_1 \cong 101010\ldots \in I_1.$$

Imposing periodicity, we obtain the system of equations

$$\begin{aligned} x_1 &= f(x_0) = 1-2x_0 \\ x_0 &= f(x_1) = x_1 - \tfrac{1}{2} \end{aligned}$$

with solution

$$x_0 = \frac{1}{6} \qquad x_1 = \frac{2}{3}.$$

In general, a periodic code will produce an affine equation with rational coefficients. So all periodic points are rational numbers.

Since the boundaries of the Markov partition are mapped onto boundary points, it follows that the boundary is contained in the stable set of a periodic orbit. This feature persists in higher-dimensional diffeomorphism, like the cat map. Segments of the stable and unstable manifold determine a Markov partition, and since the map is invertible, the symbolic dynamics results in doubly infinite symbol sequences, acted upon by a two-sided symbol shift. Many of the statements we have made in the one-dimensional case are valid for general hyperbolic dynamical systems.

5 Chaos

It is counter-intuitive that a dynamical system with a deterministic rule of evolution can display features that make it resemble a random process. The notion of sensitivity on initial conditions will resolve this apparent paradox. We now introduce the main ideas of the so-called chaotic dynamics, using simple one-dimensional maps for illustration. As the detailed description of individual orbits becomes impractical —or plainly impossible— our attention shifts to the statistical properties of large collections of orbits.

5.1 Topological chaos

Consider the map $f:[-1,1]\to[-1,1]$

$$y=f(x)=\frac{3\sqrt{3}}{2}x(1-x^2).$$

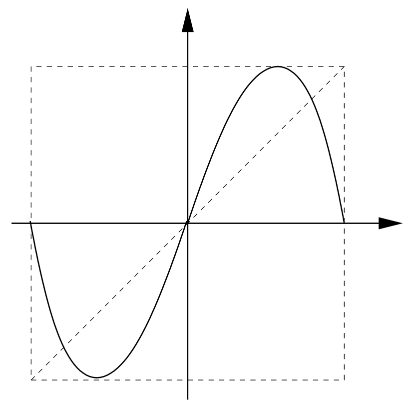

Clearly, positive (negative) initial conditions $x_0>0$ ($x_0<0$) generate non-negative (non-positive) orbits $x_n \geqslant 0$ ($x_n \leqslant 0$), since $f([0,1])=[0,1]$ and $f([-1,0])=[-1,0]$. We can 'decompose' f into two maps $f_1:[0,1]\to[0,1]$ and $f_2:[-1,0]\to[-1,0]$.

Definition 5.1 *A map $f:M\to M$ is said to be topologically transitive if for any open subsets $A,B\subseteq M$ there exists some $n>0$ such that $f^n(A)\cap B\neq\emptyset$.*

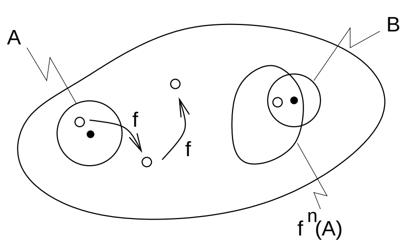

Given any pair of points, there exists a finite orbit which connects these points with arbitrary precision.

Example 5.1 The map $f(x)=\frac{3\sqrt{3}}{2}x(1-x^2)$ is not topologically transitive. Choose any $A\subseteq[-1,0]$, say $A=(-\frac{1}{2},0)$ and $B\subseteq[0,1]$, say $B=(\frac{1}{2},1)$. Then $f^n(A)\subseteq[-1,0]$ as $f([-1,0])=[-1,0]$ and $f^n(A)\cap B=\emptyset$ for all $n>0$.

Example 5.2 The map $f(x) \equiv x + \alpha \,(\text{mod } 1)$, where $\alpha \in \mathbb{R}$, is transitive if and only if $\alpha \notin \mathbb{Q}$.

Example 5.3 The Bernoulli shift map is topologically transitive, in a stronger sense than definition 5.1 —see end of section 4.1.

Example 5.4 The tent map T given in (21) is topologically transitive on $[0,1]$. We recall that (see proof of proposition 4.1) using the Markov partition $I_0 = [0, \frac{1}{2}]$, $I_1 = [\frac{1}{2}, 1]$, the cylinder sets $I_{\sigma_0\sigma_1\ldots\sigma_{n-1}}$ yield a partition as well, with each part having length 2^{-n}. Thus, if $A \subseteq [0,1]$ is an open set, then A contains at least one cylinder set $I_{\sigma_0\sigma_1\ldots\sigma_{n-1}}$ with suitable symbol string $\sigma_0\sigma_1\ldots\sigma_{n-1}$, for sufficiently large value of n. Therefore $\mathrm{T}^n(A) \supseteq \mathrm{T}^n(I_{\sigma_0\sigma_1\ldots\sigma_{n-1}}) = [0,1]$ and $\mathrm{T}^n(A) \cap B = B \neq \emptyset$ for any open set B.

Consider two nearby initial conditions x_0 and x_0', with $|x_0 - x_0'| < \varepsilon$ for some small ε. Then the distance between the first iterates is $|x_1 - x_1'| = |f(x_0) - f(x_0')| = |f'(y)||x_0 - x_0'|$, for some y lying between x_0 and x_0'. The distance shrinks if $|f'(y)| < 1$, but grows if $|f'(y)| > 1$. In the latter case $|x_n - x_n'|$ could become quite large, say, larger than some positive quantity r, even if the initial distance ε is very small. This is the so-called sensitive dependence on initial conditions (popularly named the 'butterfly effect').

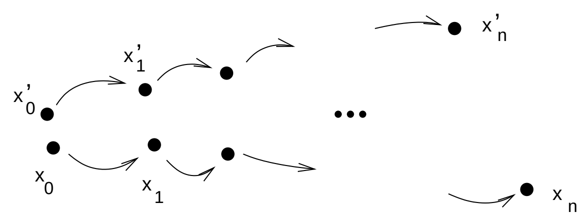

Definition 5.2 *A map $f : M \to M$ is said to have* <u>*sensitive dependence on initial conditions*</u> *if there exists an $r > 0$ such that for every $x \in M$ and $\varepsilon > 0$, there exists a point y in a open ε-neighbourhood of x, and an $n > 0$, such that $|f^n(x) - f^n(y)| \geqslant r$.*

Example 5.5 The Bernoulli shift map B has sensitive dependence on initial conditions. Let $x \in [0,1)$ and $\varepsilon > 0$ be given. Choose n such that $2^{-n} < \varepsilon$ and y such that the first n binary digits of y agree with those of x, while all remaining digits are 0 if the $(n+1)$th digit of x is 1, and 1 otherwise. Such a y is in an ε-neighbourhood of x, while $\mathrm{B}^n(x)$ and $\mathrm{B}^n(y)$ have the property that $|\mathrm{B}^n(x) - \mathrm{B}^n(y)| \geqslant 1/2$, as easily verified. So sensitivity holds with $r = 1/2$.

Example 5.6 The tent map $\mathrm{T}(x) = 1 - |2x-1|$ has sensitive dependence on initial conditions on $[0,1]$. Indeed, let $x \in [0,1]$, $\varepsilon > 0$, and let $U_\varepsilon(x)$ be an open ε-neighbourhood of x. Denote by $\sigma_0\sigma_1\sigma_2\cdots$ the symbol sequence of x with respect to the (canonical) Markov partition $I_0 = [0,\frac{1}{2}]$, $I_1 = [\frac{1}{2},1]$.

If we choose $n \in \mathbb{N}$ such that $2^{-n} < \varepsilon/2$, we have that $I_{\sigma_0\sigma_1\ldots\sigma_{n-1}} \subseteq U_\varepsilon(x)$. By definition of the symbol sequence $\mathrm{T}^n(x) \in I_{\sigma_n\sigma_{n+1}}$. Denote by $\overline{\sigma}_n$ the complementary symbol of σ_n ($\overline{\sigma}_n = 1$ if $\sigma_n = 0$, and vice-versa). Then $I_{\sigma_0\sigma_1\ldots\sigma_{n-1}\overline{\sigma}_n 0} \subseteq I_{\sigma_0\sigma_1\ldots\sigma_{n-1}} \subseteq U_\varepsilon(x)$ and for $y \in I_{\sigma_0\sigma_1\ldots\sigma_{n-1}\overline{\sigma}_n 0} \subseteq U_\varepsilon(x)$ we have $\mathrm{T}^n(y) \in I_{\overline{\sigma}_n 0}$, i.e., $|\mathrm{T}^n(x) - \mathrm{T}^n(y)| \geqslant 1/4$. Thus sensitivity holds with the choice $r = 1/4$.

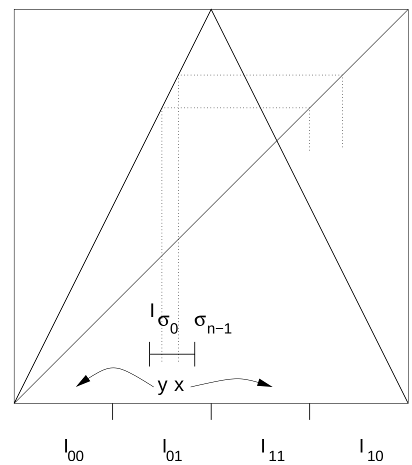

Definition 5.3 *A map $f : M \to M$ is said to be* <u>chaotic</u> *(in the topological sense) if f*

i) is topologically transitive;

ii) has sensitive dependence on initial conditions.

For an expanding Markov map, topological chaos follow automatically if the topological transition matrix is transitive, namely if there exists $n \in \mathbb{N}$ such that $\mathbf{A}^n$ has positive entries. In this case all transitions are possible for a suitable iterate of the map, and expansivity guarantees sensitivity on initial conditions.

The definitions of sensitivity, transitivity and topological chaos may also be applied to restricted domains in phase space, such as an invariant set or a recurrent set (surface of section). This often avoids trivial constraints.

5.2 Invariant measures

Consider the orbit $(x_0, x_1, x_2, \ldots)$ of a map $f : I \to I$. How are the points of the orbit distributed? Let $\{B_k\}$ be a partition of I into small intervals. We construct a histogram, by determining the fraction of the points that lie in each B_k:

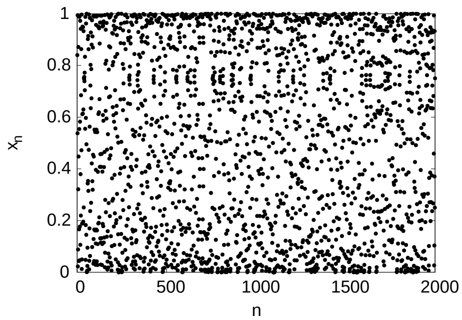

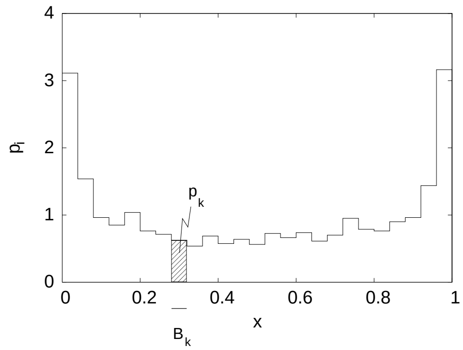

$$p_k(x_0,N) = \frac{\#\{n\,:\, n<N,\, f^n(x_0)\in B_k\}}{N} = \frac{1}{N}\sum_{n=0}^{N-1}\chi_{B_k}(x_n),$$

where χ_{B_k} is the characteristic function of the set B_k. Suppose that the expression above converges for large N, and that its limit is 'independent' from x_0 (in a sense to be made precise later). Under such conditions, the limit can be characterised by a <u>density function</u> $\rho : I \to \mathbb{R}$:

$$p_k = \lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}\chi_{B_k}(x_n) = \int_I \chi_{B_k}(x)\rho(x)\,\mathrm{d}x. \tag{22}$$

Any integrable function $h : I \to \mathbb{R}$ can be approximated arbitrarily well by a step-function, namely, by a linear combination of characteristic functions. Hence, for any integrable function h, we have

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}h(x_n) = \int_I h(x)\rho(x)\,\mathrm{d}x\,.$$

Assuming (22), we derive some properties of ρ. By construction, for all x_0 and N we have

$$\sum_k p_k(x_0,N) = 1 \qquad \text{hence} \qquad \sum_k p_k = \sum_k \lim_{N\to\infty} p(x_0,N) = \lim_{N\to\infty}\sum_k p(x_0,N) = 1.$$

It follows that

$$1 = \sum_k \int_I \chi_{B_k}(x)\rho(x)\mathrm{d}x = \int_I \sum_k \chi_{B_k}(x)\rho(x)\mathrm{d}x = \int_I \rho(x)\,\mathrm{d}x.$$

Furthermore, if the limit exists, then its value cannot depend on whether we start with x_0 or $x_1 = f(x_0)$, namely:

$$\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}h(x_n) = \lim_{N\to\infty}\left(\frac{1}{N}\sum_{n=0}^{N-1}h(x_{n+1}) + \frac{h(x_0)-h(x_N)}{N}\right) = \lim_{N=\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}h(f(x_n))\,.$$

Thus

$$\int_I h(x)\rho(x)\,\mathrm{d}x = \int_I h(f(x))\rho(x)\,\mathrm{d}x\,. \tag{23}$$

Let us derive a functional equation for ρ.

For simplicity, assume that f has two invertible branches, $f_0 : I_0 \to I,\ f_1 : I_1 \to I$. Then

$$\begin{aligned} f_0^{-1} &: I \to I_0 \\ f_1^{-1} &: I \to I_1. \end{aligned}$$

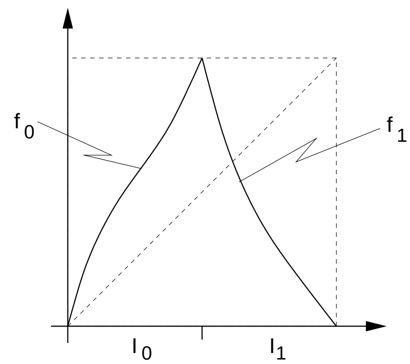

We rewrite the right-hand side of (23) as

$$\int_I h(f(x))\rho(x)\,\mathrm{d}x = \int_{I_0} h(f_0(x))\rho(x)\,\mathrm{d}x + \int_{I_1} h(f_1(x))\rho(x)\,\mathrm{d}x.$$

We now use the substitutions $x = f_0^{-1}(y)$, $\mathrm{d}x = |(f_0^{-1})'(y)|\mathrm{d}y$ in the first integral, and $x = f_1^{-1}(y)$, $\mathrm{d}x = |(f_1^{-1})'(y)|\mathrm{d}y$ in the second, to obtain

$$\begin{aligned} \int_I h(f(x))\rho(x)\,\mathrm{d}x &= \int_I h(y)\rho(f_0^{-1}(y))\frac{\mathrm{d}y}{|f_0'(f_0^{-1}(y))|} + \int_I h(y)\rho(f_1^{-1}(y))\frac{\mathrm{d}y}{|f_1'(f_1^{-1}(y))|} \\ &= \int_I h(y)\rho(y)\,\mathrm{d}y\,, \end{aligned}$$

where the last equality follows from (23). Thus

$$\rho(x) = \sum_{i=0}^{1} \frac{1}{|f'(f_i^{-1}(x))|}\rho(f_i^{-1}(x)) = \sum_{y \in f^{-1}(x)} \frac{1}{|f'(y)|}\rho(y)$$

This equation is called the Perron-Frobenius equation.

Example 5.7 Let un consider the tent map (21). We have

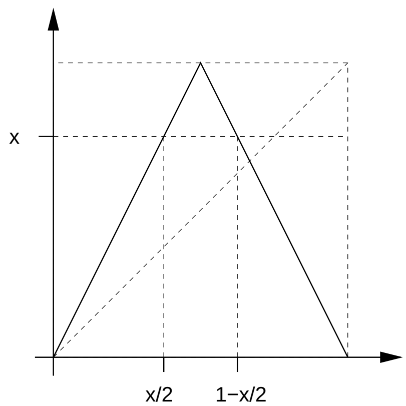

$$\begin{aligned} \mathrm{T}(x) &= 1-|1-2x| \\ \mathrm{T}^{-1}(x) &= \left\{\frac{x}{2}, 1-\frac{x}{2}\right\} \\ |\mathrm{T}'(x)| &= 2. \end{aligned}$$

The Perron-Frobenius equation reads:

$$\rho(x) = \frac{1}{2}\rho\left(\frac{x}{2}\right) + \frac{1}{2}\rho\left(1-\frac{x}{2}\right).$$

The constant function $\rho(x) = 1$ is an invariant density of the tent map since $\int_0^1 1\,\mathrm{d}x = 1$ and

$$\frac{1}{2}\rho\left(\frac{x}{2}\right) + \frac{1}{2}\rho\left(1-\frac{x}{2}\right) = \frac{1}{2} + \frac{1}{2} = \rho(x).$$

This means that the points of almost all orbits are uniformly distributed.

The density ρ may be considered as a dynamically invariant probability distribution. Indeed:

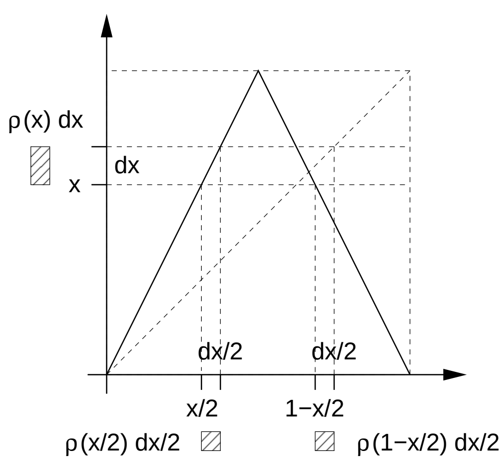

- $\rho(x)\,\mathrm{d}x$: probability to find a point in $(x, x+\mathrm{d}x)$.
- $\rho(\frac{x}{2})\,\frac{\mathrm{d}x}{2} + \rho(1-\frac{x}{2})\,\frac{\mathrm{d}x}{2}$: probability to find a point to be mapped into $(x, x+\mathrm{d}x)$.
- invariance condition

$$\rho(x)\,\mathrm{d}x = \rho\left(\frac{x}{2}\right)\frac{\mathrm{d}x}{2} + \rho\left(1-\frac{x}{2}\right)\frac{\mathrm{d}x}{2}$$

To make this kind of arguments rigorous we need the notion of measure. Let us briefly recall the main ideas (here we assume some familiarity with this concept). A measure μ is a function which assigns a non-negative weight $\mu(A)$ to any element of a collection $\mathscr{B}$ of subsets A of the phase space M. Such a collection is a σ-algebra on M, namely it contains M, and it is closed under taking of complement and countable unions. The function μ is countably additive (meaning that $\mu(\cup_i A_i) = \sum_i \mu(A_i)$ if the sets A_i are pairwise disjoint), and $\mu(\emptyset) = 0$. If $\mu(M) = 1$, then we speak of a probability measure. A set $B \in \mathscr{B}$ is said to have zero measure (full measure) if $\mu(B) = 0$ ($\mu(B) = 1$). The Lebesgue measure λ is derived by assigning to an interval I with end-points a and b, its length $\lambda(I) = |b-a|$, and then extending the definition of λ to a larger family of sets by taking limits of measures of unions of intervals. For any non negative Lebesgue integrable function ρ, the expression

$$\mu(A) = \int_A \rho(x)\mathrm{d}x \tag{24}$$

defines a measure (which is said to be absolutely continuous with respect to the Lebesgue measure). The function ρ is called the density of μ. In what follows the measures are probability measures and all the maps are considered to be continuous.

Definition 5.4 *A measure μ on M is an invariant measure for a map $f : M \to M$ if for any $B \in \mathscr{B}$, we have $\mu(f^{-1}(B)) = \mu(B)$.*

If the invariant measure has a density (i.e., if the measure is absolutely continuous with respect to the Lebesgue measure) and if the map is sufficiently smooth (e.g., piecewise smooth and expanding), then the density obeys the Perron-Frobenius equation. Thus, example 5.7 shows that the Lebesgue measure is in fact an invariant measure of the tent map.

Definition 5.5 *A map $f : M \to M$ is ergodic with respect to an invariant measure μ if for any set $B \in \mathscr{B}$, such that $B = f^{-1}(B)$, we have either $\mu(B) = 0$ or $\mu(B) = 1$*

Thus the phase space of an ergodic map cannot be decomposed into two invariant sets, each having positive measure.

Example 5.8 The Lebesgue measure is an ergodic measure of the tent map $\mathrm{T} : [0,1] \to [0,1]$, $\mathrm{T}(x) = 1 - |2x-1|$, although this is rather difficult to show. There are other ergodic measures. Consider the Dirac measure at the unstable fixed point $x_* = 0$, that is, the measure μ_0 defined by

$$\mu_0(A) = \begin{cases} 1 & \text{if } 0 \in A \\ 0 & \text{if } 0 \notin A. \end{cases}$$

This measure is an invariant measure of the tent map[8]. Indeed, if $0 \in B$, then $0 \in \mathrm{T}^{-1}(B)$ and $\mu_0(B) = \mu_0(\mathrm{T}^{-1}(B)) = 1$. If $0 \notin B$, then $0 \notin \mathrm{T}^{-1}(B)$, and $\mu_0(B) = \mu_0(\mathrm{T}^{-1}(B)) = 0$. Since the Dirac measure assigns only the values 0 and 1, ergodicity follows.

The following theorem clarifies the importance of ergodic maps.

Theorem 5.1 (Birkhoff ergodic theorem) *Let $f : M \to M$ be a measure-preserving ergodic map with invariant measure μ. Then for any (measurable) function $h : M \to \mathbb{R}$ and for μ-almost every initial conditions, we have*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} h(f^k(x)) = \int_M h(x)\mathrm{d}\mu\,. \tag{25}$$

(The expression 'μ-almost every initial conditions x' means for all x in a subset of M having μ-measure 1.)

Thus for an ergodic system, the arithmetical mean of the values of a function along an orbit —called a time-average— is almost surely equal to the integral of such a function with respect to the invariant measure —called a phase (space) average.

Example 5.9 For the tent map with the Dirac measure μ_0, the meaning of the ergodic theorem 5.1 is rather trivial. Clearly, for $x = 0$ we have $\mathrm{T}(0) = 0$, and therefore

$$h(0) = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} h(\mathrm{T}^k(0)) = \int_0^1 h(x)\mathrm{d}\mu_0\,.$$

Any subset which does not contain $x = 0$ has zero measure (with respect to μ_0). By contrast, the theorem is highly nontrivial if we consider the Lebesgue measure (taking ergodicity for granted). Then the Birkhoff ergodic theorem implies that for Lebesgue almost every initial condition the time average over the orbit exists, and it can be computed by an integral. That explains as well why one is primarily interested in special types of ergodic measures, particularly measures which are absolutely continuous with respect to the Lebesgue measure.

[8] The measure μ_0 is not absolutely continuous, that is, it cannot be written in the form (24) for some density ρ.

If $(x_0, x_1, x_2, \ldots)$ and $(x'_0, x'_1, x'_2, \ldots)$ are two orbits of a (piecewise) smooth map $f : I \to I$ on an interval then, as long as the difference $x'_k - x_k$ is sufficiently small, we have $x'_{k+1} - x_{k+1} \approx f'(x_k)(x'_k - x_k)$, and applying the chain rule of differentiation we obtain

$$\begin{aligned} |x'_n - x_n| &\approx |(f^n)'(x_0)||x'_0 - x_0| = \prod_{k=0}^{n-1} |f'(x_k)||x'_0 - x_0| \\ &= \exp\left(\sum_{k=0}^{n-1} \ln|f'(x_k)|\right) |x'_0 - x_0| \\ &= \exp\left(n \times \frac{1}{n} \sum_{k=0}^{n-1} \ln|f'(x_k)|\right) |x'_0 - x_0| . \end{aligned}$$

Therefore the so-called Lyapunov exponent

$$\Lambda(x) = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \ln|f'(f^k(x))| \tag{26}$$

quantifies the exponential rate at which nearby orbits separate. The above expression is of the type (25) with $h(x) = |\log(f'(x))|$. Hence, if f is ergodic with respect to an invariant measure μ, and if $h(x) = \ln|f'(x)|$ is μ-integrable over M, then Birkhoff ergodic theorem ensures that the limit (26) exists almost everywhere, and is independent from x. Under this circumstance, $\Lambda(x) = \Lambda$ becomes a property of the system, and we can take the condition $\Lambda > 0$ as an alternative criterion for sensitivity. This result can be extended to higher-dimensional systems, but the generalisation is nontrivial (see e.g., the so-called multiplicative ergodic theorem).

Example 5.9 illustrates the limitations of the Birkhoff ergodic theorem as a set of full μ measure could be actually quite ‘small’. One would prefer having the property of the ergodic theorem to be valid for a large set of initial conditions, say for a set of full Lebesgue measure (even if the measure μ has no density). That motivates

Definition 5.6 *An invariant measure μ of a map $f : M \to M$ is called a physical measure if*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} h(f^k(x)) = \int h(x) \mathrm{d}\mu .$$

holds for Lebesgue almost every initial condition $x \in M$.

For hyperbolic dynamical systems one can not only show the existence of such measures but one can construct such measures using the hyperbolic structure.

6 Further reading

S H Strogatz, *Nonlinear Dynamics and Chaos*, Perseus Books Publishing, Cambridge MA (1994).

One of the friendliest introductions to dynamical systems. An ideal place for a leisurely start, but lacking in mathematical sophistication.

J Guckenheimer and P Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Springer-Verlag, New York (1983).

One of the classic textbooks, which includes most of the material of this course.

C Robinson, *Dynamical Systems: Stability, Symbolic Dynamics and Chaos*, CRC Press, Boca Raton (1995).

A detailed and self-contained presentation. As the title suggest, it is very pertinent to the present course.

V I Arnold, *Ordinary Differential Equation*, The MIT Press, Cambridge (1980).

An original and lively introduction to the modern qualitative theory of ordinary differential equations, with a strong geometrical flavour. Linear systems are treated exhaustively, and there is an introduction to differential equations on manifolds. The overlap with the present course is limited.

A Katok and B Hasselblatt, *Introduction to the Modern Theory of Dynamcal Systems*, Cambridge University Press, Cambridge (1995).

A voluminous comprehensive text, with a strong emphasis on topology. An ideal reference for clarifying issues of rigour, but not suitable for a general audience.