Lecture notes for the spring 2021 lecture

Dynamical Systems

University of Fribourg
Livio Liechti

Contents

Introduction

Much of the theory of dynamical systems we know today originates from attempts to understand celestial mechanics.

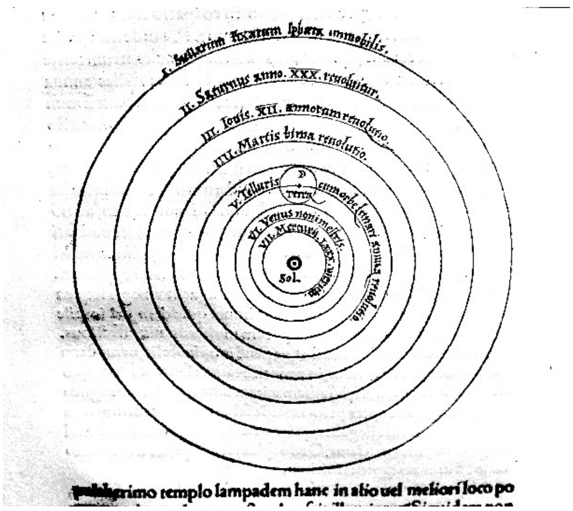

Figure 1: An illustration of the heliocentric solar system by Copernicus.

By Newton's second law of motion,

$$\ddot{x} = \sum F,$$

the study of planetary motion of our solar system reduces to solving a differential equation. Now a problem is that this differential equation is too difficult to solve explicitly. Indeed, we cannot even solve the three-body problem in general, even though certain explicit solutions are known.

Instead of solving the problem of planetary motion explicitly, one might hope to show that our solar system is stable, that is, a slight change in the initial conditions (for example due to an asteroid passing closeby) does only slightly change the long-term behaviour of the system. However, Poincaré has shown that the three-body problem is unstable: for every $\varepsilon > 0$ there exist two initial points x_1, x_2 with distance $< \varepsilon$ in the phase space $\mathbb{R}^{18}$ of the three-body problem, such that the solutions starting at x_1 and x_2 diverge with exponential speed.

This means that in order to predict long-term behaviour, our lack of precise measurement is an actual problem. However, one can still try to understand a dynamical system qualitatively, without calculating an explicit solution. For example, one can ask the following questions:

- are there periodic orbits?
- starting from a point in the phase space, which regions of the phase space will the solution visit?

We will cover a result due to Poincaré and Bendixon that provides a qualitative understanding for ordinary differential equations in the plane $\mathbb{R}^2$, such as the motion of a pendulum: every bounded orbit either converges or is asymptotically periodic. Already in $\mathbb{R}^3$, this does not hold anymore: a well-known example is given by the Lorenz flow, an orbit of which is shown in Figure 2 below.

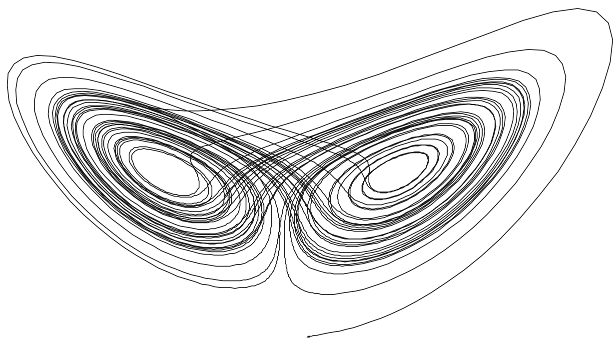

Figure 2: A bounded orbit of the Lorenz flow.

Since its beginnings, the theory of dynamical systems has developed considerably and contributed to many other areas of pure mathematics, but also real-world problems. The following is an example where the underlying problem does not describe a dynamical situation per se, but a solution can be found by adapting a dynamical view. Assume we want to rank the popularity of twitter profiles[1]. For this, we start from the assumption that everybody has popularity 1, and introduce the discrete-time dynamical system where in each step, every twitter profile distributes all its popularity evenly among the twitter profiles it is following. One can repeat this step many times, with the hope that the solution will converge. By a result of Perron and Frobenius, it actually does. The popularity in the limit gives the

[1] as of January 8 2021, it is not obvious anymore who is most popular

ranking. With undisputable success, Google uses a slightly more elaborate model of this scheme to rank the importance of webpages.

Course structure.

We will start out by covering basic notions of dynamical systems, both in the discrete-time and in the continuous-time case, and see some first examples. We then deal with linear dynamical systems and in particular cover the result by Perron–Frobenius evoked in this introduction. After covering basic notions and results from topological dynamics and ergodic theory, we will finally focus on low-dimensional dynamical systems: homeomorphisms and diffeomorphisms of the circle, vector fields in the plane (including the Poincaré–Bendixon theorem) and homeomorphisms of surfaces.

Further reading.

The books of Brin-Stuck [2] and Katok-Hasselblatt [3] contain most of what is treated in this course, and much more. These two books are suggested as basic references on the course materials and dynamical systems in general. More specific references may be given as they become relevant.

Acknowledgments.

These lecture notes are strongly influenced by Pierre Dehornoy's course on Dynamical systems 2012 in Bern. I thank Lennart Harms and Florine Pierroz for correcting some inconsistencies and typos in the script.

Chapter 1

First examples and basic notions

In the introduction, we have already met the two types of dynamical systems we will consider: discrete-time and continuous-time dynamical systems. Here, we start by giving a formal definition of these concepts.

Definition 1.1. *A discrete-time dynamical system is given by a non-empty set X and a map $f : X \to X$.*

Definition 1.2. *A continuous-time dynamical system is given by a non-empty set X and a one-parameter family of maps $\{\varphi^t : X \to X\}$ satisfying $\varphi^{s+t} = \varphi^s \circ \varphi^t$ and $\varphi^0 = id_X$, where the parameter t runs over $\mathbb{R}$ or $\mathbb{R}_{\geq 0}$.*

A continuous-time dynamical system is also called a *flow* in case the parameter t runs over $\mathbb{R}$, or a *semiflow* in case the parameter t runs over $\mathbb{R}_{\geq 0}$. In order to not give every basic definition in two flavours, we will just start by considering discrete-time dynamical systems, and come back to the continuous-time analogues later.

1.1 Discrete-time dynamical systems

Example 1.3. The following are basic examples of discrete-time dynamical systems. We will study many of the concepts we encounter with the help of these examples.

- permutations of finite sets,
- translations in $\mathbb{R}$, $t_\alpha : x \mapsto x + \alpha$,
- rotations on the circle $\mathbb{R}/\mathbb{Z}$, $r_\alpha : x \mapsto x + \alpha \bmod 1$,
- expansions of the circle $\mathbb{R}/\mathbb{Z}$, $E_m : x \mapsto mx \bmod 1$.

1.1.1 Orbits and conjugacy

Given a dynamical system, one of the key points of interest is the evolution of points over time. This is captured by the notion of *orbit*.

Definition 1.4. *Let (X, f) be a discrete-time dynamical system. The positive semiorbit $\mathcal{O}^+_f(x)$ is the set $\{x, f(x), f^2(x), \dots\}$. If f is invertible, the negative semiorbit $\mathcal{O}^-_f(x)$ is the set $\{x, f^{-1}(x), f^{-2}(x), \dots\}$.*

The *orbit* of a point x is then simply the union of the positive and the negative semiorbit.

Definition 1.5. *Let (X, f) be a discrete-time dynamical system. Some point $x \in X$ is a fixed point if $f(x) = x$ and a periodic point if $f^n(x) = x$ for some integer $n \geq 1$. The number n is called a period of x and the smallest such n is the minimal period. If $f^m(x)$ is periodic for some m, then x is called eventually periodic.*

Exercise 1.6. Study fixed, periodic and eventually periodic points for the examples at the beginning of this section (Example 1.3).

Exercise 1.7. Show that for a circle rotation with $\alpha \in \mathbb{R} \setminus \mathbb{Q}$, every orbit is a dense subset of the circle.

As in many mathematical theories, there is a notion of equivalence of objects. In the theory of dynamical systems, it goes by the name of *conjugacy.*

Definition 1.8. *Let (X, f) and (Y, g) be two discrete-time dynamical systems. A semiconjugacy from (Y, g) to (X, f) is a surjective map $\pi : Y \to X$ such that $f \circ \pi = \pi \circ g$.*

Definition 1.9. *A conjugacy is an invertible semiconjugacy.*

Conjugate dynamical systems have the same properties. As an exercise, you can show that if (X, f) and (Y, g) are conjugated by π, then x is periodic with minimal period n if and only if $\pi(x)$ is periodic with minimal period n. Being presented with some unknown dynamical system to study, it is a basic idea to find conjugacies or semiconjugacies to systems that are better understood.

Exercise 1.10. When are two permutations conjugate or semiconjugate?

We now describe an nontrivial example of a conjugacy.

Proposition 1.11. *Let $f, g : [0, 1] \to [0, 1]$ be two homeomorphisms, and suppose $f(0) = g(0) = 0$ and $f(x), g(x) > x$ for $x \in (0, 1)$. Then there exists a continuous conjugacy between f and g.*

Before we procede to the proof of Proposition 1.11, we note that in this case the dynamical system we consider has additional structure: $X = [0, 1]$ is a topological space and f, g are continuous. In this case we speak of *topological dynamical systems*. Also the conjugacy described in Proposition 1.11, being continuous, is what we call a *topological conjugacy*. One could add other structure to a dynamical system in a similar way with, for example, differentiability: one could suppose X to be a smooth manifold, and f to be of class C^1, C^2 or C^∞.

Proof of Proposition 1.11. Choose an arbitrary point x in $(0, 1)$.

We have $\lim_{n\to\infty} f^n(x) = 1$. Indeed, $\{f^n(x)\}_n$ is an increasing sequence in $[0, 1]$ and hence must converge. But the limit point l satisfies $f(l) = l$, so $l = 1$.

With the same argument, we get $\lim_{n\to\infty} g^n(x) = 1$, $\lim_{n\to-\infty} f^n(x) = 0$ and $\lim_{n\to-\infty} g^n(x) = 0$.

Define $I^n = \big[f^n(x), f^{n+1}(x)\big)$ and $J^n = \big[g^n(x), g^{n+1}(x)\big)$ for all $n \in \mathbb{Z}$. Clearly, the interval $(0, 1)$ is the disjoint union of all the I^n, and also the disjoint union of all the J^n.

Choose a homeomorphism $\pi : I^0 \to J^0$, with $\pi(x) = x$.

In order to extend π to a homeomorphism of $[0, 1]$, we note that for $y \in (0, 1)$, there exists a unique $z \in I^0$ and a unique $n \in \mathbb{Z}$ such that $y = f^n(z)$. We now define $\pi(y) = g^n(\pi(z))$. Finally, we extend to the boundary of the unit interval by $\pi(0) = 0, \pi(1) = 1$.

The map π is a topological conjugacy (checking the final details of this claim is left as an exercise). □

1.1.2 Shifts

Shifts are the principal examples of what is called *symbolic dynamics*, not being related to functions on the real numbers, a priori. Given a finite set of symbols $\mathcal{A}$, define $\Sigma_{\mathcal{A}} = \mathcal{A}^{\mathbb{Z}}$ to be the set of bi-infinite sequences $(x_i)_{i\in\mathbb{Z}}$ and $\Sigma^+_{\mathcal{A}} = \mathcal{A}^{\mathbb{N}}$ to be the set of infinite sequences $(x_i)_{i\in\mathbb{N}}$ with elements $x_i \in \mathcal{A}$.

Definition 1.12. *The two-sided full shift $\sigma : \Sigma_{\mathcal{A}} \to \Sigma_{\mathcal{A}}$ is the map*

$$(x_i)_{i\in\mathbb{Z}} \overset{\sigma}{\mapsto} (x_{i+1})_{i\in\mathbb{Z}}$$

shifting the indices of the sequences by one. Similarly, the one-sided full shift $\sigma : \Sigma^+_{\mathcal{A}} \to \Sigma^+_{\mathcal{A}}$ is the map

$$(x_i)_{i\in\mathbb{N}} \overset{\sigma}{\mapsto} (x_{i+1})_{i\in\mathbb{N}}.$$

The fixed points and periodic points of a full shift are easy to determine: they are simply the constant sequences and the periodic sequences, respectively.

We will now construct a semiconjugacy from a one-sided full shift to the expanding map $E_2, x \mapsto 2x \bmod 1$ on the circle $\mathbb{R}/\mathbb{Z}$.

We consider the alphabet $\mathcal{A} = \{0, 1\}$ and define the map

$$\begin{aligned} b_2 : \Sigma^+_{\{0,1\}} &\to \mathbb{R}/\mathbb{Z} \\ (x_0, x_1, \dots) &\mapsto \sum_{i \in \mathbb{N}} \frac{x_i}{2^{i+1}} \bmod 1. \end{aligned}$$

Proposition 1.13. *The map b_2 is a semiconjugacy $(\Sigma^+_{\{0,1\}}, \sigma) \to (\mathbb{R}/\mathbb{Z}, E_2)$.*

Proof. The map b_2 is surjective as any number in $[0, 1)$ can be written in base 2. We note that b_2 is injective outside the countable set of sequences that are eventually constant, but we will not need this observation here.

It remains to show $E_2 \circ b_2 = b_2 \circ \sigma$. For the left side, we have

$$\begin{aligned} E_2 \circ b_2(x_0, x_1, \dots) &= E_2 \left(\sum_{i \in \mathbb{N}} \frac{x_i}{2^{i+1}} \right) \bmod 1 \\ &= \sum_{i \in \mathbb{N}} \frac{x_i}{2^i} \bmod 1 \\ &= \sum_{i \in \mathbb{N}} \frac{x_{i+1}}{2^{i+1}} \bmod 1. \end{aligned}$$

The last equality follows from the fact that $\frac{x_0}{1} = 0 \bmod 1$. For the right side, we have

$$b_2 \circ \sigma(x_0, x_1, \dots) = b_2(x_1, x_2, \dots) = \sum_{i \in \mathbb{N}} \frac{x_{i+1}}{2^{i+1}},$$

which finishes the proof. □

Exercise 1.14. Show that $\sigma : \Sigma^+_{\{0,1\}} \to \Sigma^+_{\{0,1\}}$ has 2^k periodic points of period k. Use this and the semiconjugacy b_2 to deduce that $E_2 : \mathbb{R}/\mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ has $2^k - 1$ periodic points of period k.

1.2 Continuous-time dynamical systems

The notions of orbits, fixed points, periodic points, conjugacy, semiconjugacy, and so on are defined as in the discrete-time case. We repeat only one of the definitions to illustrate the similarity.

Definition 1.15. *Let (X, φ^t) be a continuous-time dynamical system. The positive semiorbit $\mathcal{O}^+_\varphi(x)$ is the set $\cup_{t \in \mathbb{R}_{\geq 0}} \varphi^t(x)$. If φ^t is a flow, the negative semiorbit $\mathcal{O}^-_\varphi(x)$ is the set $\cup_{t \in \mathbb{R}_{\leq 0}} \varphi^t(x)$.*

1.2.1 Flows and vector fields

Flows or continuous-time dynamical systems naturally arise from ordinary differential equations that do not depend on the time t. Suppose $\dot{x} = F(x)$ is a differential equation in $\mathbb{R}^n$, where $F : \mathbb{R}^n \to \mathbb{R}^n$ is continuously differentiable.

Definition 1.16. *A solution of $\dot{x} = F(x)$ with initial data $x_0 \in \mathbb{R}^n$ is a differentiable curve $\gamma : J \to \mathbb{R}^n$, where $J \subset \mathbb{R}$ is some interval containing the origin, such that $\gamma(0) = x_0$, and for all $t \in J$, we have $\dot{\gamma}(t) = F(\gamma(t))$.*

The result of Picard-Lindelöf[1] tells us that for each $x \in \mathbb{R}^n$, there exists a unique solution $\gamma(t)$ starting from x at time 0 and defined for all $t \in J$.

Example 1.17. For the differential equation $\dot{x} = x^2$ in $\mathbb{R}$, the curve

$$\gamma(t) = \frac{x_0}{1 - \frac{t}{x_0}}$$

is the solution with initial data x_0. The maximal interval of definition J is $(-\infty, x_0)$. We note that solutions need not be defined on all $\mathbb{R}$.

We are now ready to define the flow associated with a differential equation. It can be thought of as the collection of all the local solutions, for all initial data.

Definition 1.18. *Let $\dot{x} = F(x)$ be a differential equation in $\mathbb{R}^n$, where we assume F to be continuously differentiable. The associated flow is the map $\varphi : U \to \mathbb{R}^n$, where U is an open subset of $\mathbb{R} \times \mathbb{R}^n$ containing $\{0\} \times \mathbb{R}^n$, satisfying for $x \in \mathbb{R}^n$ that $\varphi(\cdot, x)$ is the maximally defined solution with initial data x.*

Proposition 1.19. *The flow associated with a differential equation is a continuous-time dynamical system.*

Proof. Clearly, φ defines a one-parameter family of maps $\mathbb{R}^n \to \mathbb{R}^n$ (in order to obtain the notation from our definition of continuous-time dynamical systems, we write the first coordinate in superscript) with $\varphi^0 = \mathrm{id}_X$. What we have to show is that $\varphi^{s+t} = \varphi^s \circ \varphi^t$. Specialised to a point $x \in \mathbb{R}^n$, we want to show $\varphi(s+t, x) = \varphi(s, \varphi(t, x))$. But this follows from the uniqueness of the solution of the differential equation $\dot{x} = F(x)$ at the point $\varphi(t, x)$. □

Example 1.20. For $F(x, y) = (-y, x)$ in the real plane $\mathbb{R}^2$, we get the flow $\varphi^t(x, 0) = (x \cos t, x \sin t)$.

[1] or Cauchy-Lipschitz if we are francophone

1.2.2 Flowboxes

The flowbox theorem is a structural result stating that in the neighbourhood of a regular point, a flow is conjugate (via a diffeomorphism) to a flow in a box that moves only the last coordinate. It is a structural theorem we will use later when studying vector fields and flows in the plane. Before stating the result, we recall how a diffeomorphism transports a vector field from its domain to its range.

Definition 1.21. *Let F be a vector field on some open set $U \subset \mathbb{R}^n$, and let $h : U \to V$ be a diffeomorphism. Recall that the pushforward $G = h_* F$ of the vector field F by h is defined to be $G(y) = \mathrm{D}_{h^{-1}(y)} h(F(h^{-1}(y)))$.*

As an exercise, one directly verifies that if $\gamma(t)$ is a solution of $\dot{x} = F(x)$, then $h \circ \gamma(t)$ is a solution of $\dot{y} = h_* F(y)$.

Theorem 1.22 (Flowbox theorem)**.** *Let F be a continuously differentiable vector field on some open set $U \subset \mathbb{R}^n$. If $F(x_0) \neq 0$ for some point $x_0 \in U$, then there exists a diffeomorphism h from a neighbourhood U_0 of x_0 to a neighbourhood V of $0 \in \mathbb{R}^n$ such that $h_* F = (0, \dots, 0, 1)$.*

Proof. Take an affine hyperplane H centered at x_0 which is transverse to $F(x_0)$ at x_0. Now take a neighbourhood H_0 of x_0 in H such that H is transverse to F in all points of H_0. Writing φ for the flow associated with F, we consider for $\varepsilon > 0$ small enough

$$\begin{aligned} \psi : H_0 \times (-\varepsilon, \varepsilon) &\to \mathbb{R}^n \\ (x, t) &\mapsto \varphi^t(x). \end{aligned}$$

One verifies directly that $\mathrm{D}_{(x_0,0)}\psi$ is invertible, recalling from the definitions that $\varphi^0(x) = x$ and $\frac{\partial}{\partial t}|_{t=0}\varphi^t(x_0) = F(x_0)$. This implies that ψ is a local diffeomorphism. Note that for $T = (0, \dots, 0, 1)$, we have $\psi_* T = F$. So we may simply take $h = \psi^{-1}$, suitably restricted. □

1.2.3 Pendula and Lyapunov functions

Studying physical systems, one often considers a phase space of the problem at hand and the energy of the system, which should stay constant in time along the solutions. We will discuss pendula.

Example 1.23 (Phase space of a pendulum)**.** Denoting by θ the angle of the pendulum away from its equilibrium state, the motion of a frictionless pendulum is gouverned by the differential equation

$$\ddot{\theta} = -\sin\theta.$$

We introduce the variable $\xi = \dot{\theta}$ and obtain $\dot{\theta} = \xi$ and $\dot{\xi} = -\sin\theta$. In particular, the solution must, in the (θ, ξ)-plane, satisfy the differential equation given by $F : \mathbb{R}^2 \to \mathbb{R}^2, F(\theta, \xi) = (\xi, -\sin\theta)$.

At this point, it is an exercise to sketch the vector field F.

We now introduce the energy

$$E(\theta, \xi) = 1 + \frac{\xi^2}{2} - \cos\theta$$

and note, as an exercise, that if $\gamma(t) = (\theta(t), \xi(t))$ is a solution to F, then

$$\frac{\mathrm{d}}{\mathrm{d}t} E(\theta(t), \xi(t)) = 0.$$

So, the solutions lie on the level sets of the function E.

As an exercise, draw the level sets of E and describe, using also your sketch of the vector field F, the orbits of the flow associated with $(\dot\theta, \dot\xi) = F(\theta, \xi)$.

In the example above, we have seen a function that is constant along the orbits of the flow. Such a function is called *first integral.*

Example 1.24 (Damped pendulum)**.** The differential equation for a damped pendulum is $\ddot\theta = -\alpha\dot\theta - \sin\theta$, where $\alpha > 0$ is a damping parameter. It can be restated in the phase plane as in the example above, to become $\dot\theta = \xi$ and $\dot\xi = -\alpha\xi - \sin\theta$. This time, it is an exercise to verify that the energy

$$E(\theta, \xi) = 1 + \frac{\xi^2}{2} - \cos\theta$$

is decreasing along non-constant solutions.

A continuous function that is decreasing along nontrivial forward orbits of a flow is called *Lyapunov function.* Strict local minima of Lyapunov functions have strong stability properties. We first give the definition and then a stability result for Lyapunov functions.

Definition 1.25. *A fixed point x_0 of a flow φ^t on $X \subset \mathbb{R}^n$ is*

- *stable if for all $\varepsilon > 0$ there exists $\delta > 0$ such that $||x - x_0|| < \delta$ implies, for all $t \geq 0$, that $||\varphi^t(x) - x_0|| < \varepsilon$.*
- *asymptotically stable if it is stable and there exists a neighbourhood U of x_0 such that for $x \in U$, $\varphi^t(x)$ converges to x_0 as $t \to \infty$.*

Theorem 1.26. *Suppose x_0 is a fixed point of the flow φ^t on $X \subset \mathbb{R}^n$. If there exists a Lyapunov function L in a neighbourhood of x_0 that has a strict local minimum at x_0, then x_0 is asymptotically stable.*

Proof. Let W be a compact neighbourhood of x_0 such that $L(x_0)$ is the strict minimum of L on W. Let B be an open ε-ball around x_0, for ε small enough such that $B \subset W$. We further define $m_0 := \min L|_{W\setminus B} > L(x_0)$ and V_{m_0} to be the connected component of $L^{-1}([L(x_0), m_0)) \cap W$ containing x_0.

Clearly, V_{m_0} contains an open neighbourhood of x_0. We also note $V_{m_0} \subset B$. Indeed, if there existed $x \in V_{m_0} \setminus B$, then $L(x)$ would have to be both $< m_0$ and $\geq m_0$, a contradiction.

For any $x \in V_{m_0}$, the forward orbit $\varphi^t(x)$, for $t \geq 0$, stays inside V_{m_0}. Indeed, by the definition of a Lyapunov function, we have $L(\varphi^t(x)) < m_0$ for all t. In order to prove stability of the fixed point x_0, one can now simply choose δ such that the δ-ball around x_0 is contained in V_{m_0}.

Let us now prove asymptotic stability of x_0. For an arbitrary $x \in V_{m_0}$, let $a := \lim_{t\to\infty} L(\varphi^t(x))$. Assume for a contradiction that the orbit $\varphi^t(x)$ does not converge to x_0. Take $x_1 \neq x_0$ to be an accumulation point of $\varphi^t(x)$. Such an accumulation point has to exist since $\varphi^t(x)$ stays inside the compact set W. Let $t_1, \ldots, t_n, \cdots \to \infty$ be a sequence with $\lim_{n\to\infty} \varphi^{t_n}(x) = x_1$.

By the definition of L, we have $L(x_1) > L(x_0)$ and $L(\varphi^t(x_1)) < L(x_1)$ for all $t > 0$. In particular, we get

$$a > L(\varphi^t(x_1)) = \lim_{n\to\infty} L(\varphi^t(\varphi^{t_n}(x))) = \lim_{n\to\infty} L(\varphi^{t+t_n}(x)) = a,$$

a contradiction. We conclude that $\varphi^t(x)$ must in fact converge to x_0, proving asymptotic stability of the fixed point x_0. □

Exercise 1.27. Fix $\varepsilon > 0$ and consider the flow φ associated with the planar differential equation

$$\begin{aligned}\dot{x} &= -y + \varepsilon x(1 - x^2 - y^2)\\ \dot{y} &= x + \varepsilon y(1 - x^2 - y^2).\end{aligned}$$

a) Does φ have fixed points? Does φ have periodic orbits?

b) Find a Lyapunov function for φ on $\mathbb{R}^2 \setminus \{(x, y) : x^2 + y^2 = 1\}$.

c) For $(x, y) \neq (0, 0)$, what is the set of accumulation points of the orbit $\varphi^t(x, y)$?

Exercise 1.28. Consider the first-order differential equation $\dot{x} = \lambda x - x^3$, where $\lambda \in \mathbb{R}$ is a real parameter. Determine the fixed points of the associated flow. Are they stable? Are they asymptotically stable? To illustrate this, trace the graph of the multi-map that associates to every λ the fixed points of the associated flow (and color the branches of the graph according to the stability of the fixed points). Describe what happens around $\lambda = 0$.

1.3 Continuous-time vs. discrete-time systems

There are ways to go from discrete-time dynamical systems to continuous-time dynamical systems and back. The easiest way to associate a discrete-time system with a flow φ^t is to simply let $f = \varphi^{t_0}$, for any fixed time t_0.

This construction amounts to look at the effects of a flow in discrete time intervals only. However, very few maps f arise in this way: assuming we consider a topological flow on a topological manifold, every such map is necessarily isotopic to the identity.

Definition 1.29. *For a discrete-time dynamical system* (X, f) *and a so-called ceiling function* $c : X \to \mathbb{R}_{\geq 0}$*, we define the suspension of* (X, f) *to be the continuous-time dynamical system given by the flow* φ^t *induced by the constant vector field* $(0, 1)$ *on* $X_c := \{(x, t) : x \in X, t \in [0, c(x)]\}/\sim$*, where* $(x, c(x)) \sim (f(x), 0)$*.*

One can picture the suspension as follows: starting from a base point $(x, 0)$, it increases the second coordinate at unit speed until it hits the ceiling $c(x)$. The point $(x, c(x)) \in X_c$ is identified with the point $(f(x), 0)$, and so again the flow increases the second coordinate at unit speed until it hits the ceiling $c(f(x))$, and so on.

Exercise 1.30. Let (X_c, φ^t) be a suspension of (X, f). Show that the orbit of $x \in X$ under f is periodic if and only if the orbit of $(x, 0) \in X_c$ under φ^t is periodic. Find conditions on X and c such that the same equivalence holds for density of orbits.

The reverse direction of the suspension construction does not always work; it is captured by the notion of *cross-section*.

Definition 1.31. *Given a semiflow* $\varphi^t : Y \to Y$*, a cross-section is a subset* $X \subset Y$ *such that for all* $y \in Y$*, the set* $T_y = \{t \in \mathbb{R}_{\geq 0} : \varphi^t(y) \in X\}$ *is nonempty and discrete.*

Definition 1.32. *Let* X *be a cross-section for* $\varphi^t : Y \to Y$*. For* $x \in X$*, the first return time is* $\tau(x) := \min T_x$*. The first return map* $f : X \to X$ *is then defined to be* $f(x) = \varphi^{\tau(x)}(x)$*.*

Example 1.33. We consider the torus $\Sigma_1 = \mathbb{R}^2/\mathbb{Z}^2$. For $a, b \in \mathbb{R}$, let φ^t be the linear flow defined by

$$\varphi^t(x, y) = (x + ta, y + tb) \bmod 1.$$

If $b \neq 0$, then the circle $\{(x, 0) \in \mathbb{R}^2/\mathbb{Z}^2 : x \in [0, 1)\} \simeq \mathbb{R}/\mathbb{Z}$ is a cross-section for φ^t. The first return time is $\frac{1}{|b|}$ and the first return map is given by $x \mapsto x + \frac{a}{b} \bmod 1$.

While a cross-section does not always exist, it does always do so for suspensions.

Exercise 1.34. Let (X, f) be a discrete-time dynamical system and let c be a ceiling function. Show that $X \times \{0\}$ is a cross-section for the suspension of (X, f) with ceiling function c, and that the discrete-time dynamical system $(X \times \{0\}, \text{first return map})$ is conjugate to (X, f).

We note that the existence of a cross-section for a continuous-time dynamical system is a strong property, since the system can then be studied via a discrete-time dynamical system of dimension one less.

Chapter 2

Linear dynamical systems

Linear dynamical systems are fairly well-understood and form an important class of examples, also with respect to applications.

A discrete-time linear dynamical system on $\mathbb{R}^n$ is defined by the map

$$f(x) = Ax,$$

where $A \in \mathrm{Mat}_n(\mathbb{R})$. A continuous-time linear dynamical system is defined by the solutions to the linear ordinary differential equation

$$\dot{x} = Ax,$$

where $A \in \mathrm{Mat}_n(\mathbb{R})$. Recall that the solution curves for this differential equation are of the type

$$\gamma(t) = e^{tA}x_0,$$

where $x_0 = \gamma(0)$.

2.1 Low-dimensional examples

Dimension 1. In dimension one, a discrete-time dynamical system is given by multiplication with a scalar λ. In particular, we have $f^n(x) = \lambda^n x$, and we distinguish three cases. If $|\lambda| < 1$, then 0 is an *attractive* fixed point. Here, *attractive* is used as a synonym for *asymptotically stable.* If $|\lambda| = 1$, then 0 is a stable fixed point, but not asymptotically stable. If $|\lambda| > 0$, then 0 is still a fixed point, but it is neither stable nor asymptotically stable. In fact, it is a *repulsive* fixed point. What we mean by this is that 0 is an asymptotically stable fixed point for f^{-1}.

For continuous-time dynamical systems, we can proceed in the same way. Remember that a solution curve is given by $\gamma(t) = e^{ct}x$. Now, we get the same behaviour for the fixed point at 0 by distinguishing between the three cases $c < 0, c = 0$ and $c > 0$.

Dimension 2. We distinguish multiple cases.
Case 1. Assume the matrix $A \in \mathrm{Mat}_2(\mathbb{R})$ is conjugate to $\left(\begin{smallmatrix} \lambda & 0 \\ 0 & \mu \end{smallmatrix}\right)$, where $\lambda, \mu \in \mathbb{R}$. In this case, we have $A^n = \left(\begin{smallmatrix} \lambda^n & 0 \\ 0 & \mu^n \end{smallmatrix}\right)$. As an exercise, draw sketches of the orbits for the following possibilities:

1) $0 < \lambda = \mu < 1$,
2) $0 < \lambda < \mu < 1$,
3) $0 < \lambda < 1 < \mu$,
4) $1 < \lambda, \mu$ (this can be obtained by reversing previous cases),
5) $\mu < 1 = \lambda$.

In the case of a continuous-time dynamical system defined by $B \sim \begin{pmatrix} p & 0 \\ 0 & q \end{pmatrix}$, we have $e^{tB} = \begin{pmatrix} e^{tp} & 0 \\ 0 & e^{tq} \end{pmatrix}$, and we again get a similar case distinction. Orbits of the continuous-time case can be thought of as interpolating between the points of the orbits of the discrete-time case.

Exercise 2.1. Show that the dynamical system defined by multiplication with the matrix $\begin{pmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{pmatrix}$ is topologically conjugate to the dynamical system defined by multiplication with the matrix $\begin{pmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{4} \end{pmatrix}$. Show that no smooth such conjugacy exists.

Case 2. Assume the matrix $A \in \mathrm{Mat}_2(\mathbb{R})$ is conjugate to to Jordan block of size two $\left(\begin{smallmatrix} \lambda & 1 \\ 0 & \lambda \end{smallmatrix}\right)$, where $\lambda \in \mathbb{R}$. In this case, we have $A^n \sim \left(\begin{smallmatrix} \lambda^n & n\lambda^{n-1} \\ 0 & \lambda^n \end{smallmatrix}\right)$, which yields the formula

$$f^n(x, y) = \lambda^n(x + \frac{n}{\lambda}y, y)$$

For the linear map f defined by multiplication with the matrix A. As an exercise, draw a sketch of the orbits of this dynamical system, for the case where $0 < \lambda < 1$.

Exercise 2.2. Show that the dynamical system defined by multiplication with the matrix $\left(\begin{smallmatrix} \lambda & 1 \\ 0 & \lambda \end{smallmatrix}\right)$, where $\lambda \in \mathbb{R}$, is topologically conjugate to the dynamical system defined by multiplication with the matrix $\left(\begin{smallmatrix} \lambda & 0 \\ 0 & \lambda \end{smallmatrix}\right)$. Note again that there can be no smooth conjugacy.

Case 3. The matrix $A \in \mathrm{Mat}_2(\mathbb{R})$ has no real eigenvalues and

$$A \sim \begin{pmatrix} r\cos\theta & r\sin\theta \\ -r\sin\theta & r\cos\theta \end{pmatrix}.$$

Again, one can make a qualitative case distinction between $|r| < 1, |r| = 1$ and $|r| > 1$, and show that there is a topological conjugacy to the the dynamical system given by multiplication with the matrix $\left(\begin{smallmatrix} r & 0 \\ 0 & r \end{smallmatrix}\right)$.

2.2 Contractions

Contractions provide examples of maps with the simplest possible asymptotic behaviour: every point converges to a unique fixed point under iterates of the map.

Definition 2.3. *Let X be a metric space. A map $f : X \to X$ is a* contraction *if there exists $\lambda < 1$ such that for all $x, y \in X$,*

$$d(f(x), f(y)) \leq \lambda d(x, y).$$

The asymptotic behaviour of points under a contraction is described by Banach's fixed point theorem.

Theorem 2.4 (Banach fixed point theorem)**.** *Let X be a complete metric space, and let $f : X \to X$ be a contraction. Then f has a unique fixed point in X, and for every $x \in X$, the sequence $f^n(x)$ converges to the fixed point of f with exponential speed.*

If you do not recall the proof of Banach's fixed point theorem, it is an exercise in Cauchy sequence calculus to find it again.

For a matrix A, we let $r(A)$ be its *spectral radius*, that is, the largest absolute value among all eigenvalues of A. We will now show that matrices with spectral radius < 1 are contractions. To show this, we need to define suitable norms on $\mathbb{R}^n$. Recall that if $||\cdot||$ is a norm on $\mathbb{R}^n$, then we define the norm of a linear map A by

$$||A|| := \sup_{||v||=1} ||Av||.$$

Clearly, the norm of A is equal to $r(A)$ in case A is diagonalisable. Otherwise, we still have $||A|| \geq r(A)$. The following proposition also includes the case of matrices that cannot be diagonalised.

Proposition 2.5. *For every $\delta > 0$, there exists a norm $||\cdot||$ on $\mathbb{R}^n$ such that $||A|| \leq r(A) + \delta$.*

Proof. Basically, what is left (after our discussion leading up to the proposition) is to deal with Jordan blocks of size greater than one. Let us consider the example of a Jordan block

$$J = \begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 1 & 0 \\ 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & \lambda \end{pmatrix}.$$

We want to find a norm $||\cdot||$ on $\mathbb{R}^4$ such that the norm of

$$\begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 1 & 0 \\ 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & \lambda \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \lambda \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} + \begin{pmatrix} x_2 \\ x_3 \\ x_4 \\ 0 \end{pmatrix}$$

is bounded from above by $(\lambda+\delta)||(x_1, x_2, x_3, x_4)^\top||$. By the triangle equality, this is satisfied in case we have $||(x_2, x_3, x_4, 0)^\top|| \leq \delta||(x_1, x_2, x_3, x_4)^\top||$. For the norm $||\cdot||$, we now simply choose

$$||(x_1, x_2, x_3, x_4)^\top|| := |x_1| + \frac{1}{\delta}|x_2| + \frac{1}{\delta^2}|x_3| + \frac{1}{\delta^3}|x_4|.$$

and leave it as an exercise to finish the details and to give a proof of the general case (see, for example, Proposition 1.2.2. in [3]). □

In conclusion, Proposition 2.5 tells us that a matrix A with spectral radius $r(A) < 1$ is a contraction and therefore, by Banach's fixed point theorem, has very particular asymptotic behaviour.

2.3 Conjugacy to the linear part

In this section, we consider the question of when a differentiable discrete-time dynamical system $(\mathbb{R}^n, f)$ is conjugate to its linearisation around a neighbourhood of a fixed point.

2.3.1 Examples in dimension one

Example 2.6. Let $f : \mathbb{R} \to \mathbb{R}$ be continuously differentiable, with $f(0) = 0$ and derivative $0 < f'(0) < 1$. Then on some open neighbourhood U of 0, we have $0 < f'(x) \leq \lambda < 1$, and by the mean value theorem f is a strictly increasing function on U such that $f(x) < x$. In particular, by the argument used in the proof of Proposition 1.11, we obtain that on U, the map f is conjugate to its linearisation $x \mapsto f'(0)x$.

What is the important feature of the above example, admitting a conjugacy to its linear part? Is it the fixed point being attractive?

Example 2.7. Let $f : \mathbb{R} \to \mathbb{R}$ be defined by $f(x) = x - x^3$. As in the above example, 0 can be shown to be an attractive fixed point. Moreover, the point 0 is the only fixed point of f. However, we have $f'(0) = 1$, so that the linearisation of f at 0 is the identity, the map for which every point is fixed. We obtain that there is no neighbourhood of 0 on which f is conjugate to its linearisation.

A similar counterexample can be made where 0 is a repulsive fixed point: simply take $f(x) = x + x^3$. We see that the obstruction to conjugate a map to its linearisation around a fixed point in dimension one occurs when the linearisation is the identity (actually, $\pm$ the identity).

2.3.2 The Hartman-Grobman theorem

Definition 2.8. *A matrix $A \in \mathrm{Mat}_n(\mathbb{R})$ is called hyperbolic if none of its eigenvalues is contained in $\mathbb{S}^1 \subset \mathbb{C}$.*

Theorem 2.9 (Hartman-Grobman, 1959)**.** *Let $\Omega \subset \mathbb{R}^n$ be an open neighbourhood of the origin, and let $F : \Omega \to F(\Omega)$ be a C^1-diffeomorphism with $F(0) = 0$. If d_0F is hyperbolic and invertible, then there exists a neighbourhood U of the origin, and a homeomorphism $H : U \to H(U)$ fixing the origin, such that $H \circ d_0F = F \circ H$.*

Let us fix some notation for the proof of Theorem 2.9. We set $T := d_0F$ and let $F = T + f$. We will be looking for the map $H = \mathrm{Id} + h$.

Proof of a contracting case. Assume that T is contracting, that is, for every eigenvalue λ of T, we have $|\lambda| \leq r := r(T) < 1$. Choose $\varepsilon > 0$ such that $r + \varepsilon < 1$, and let $||\cdot||$ be a norm on $\mathbb{R}^n$ such that $||T|| \leq r + \varepsilon$. Such a norm exists by Proposition 2.5. For the moment, we furthermore assume $F = T$ outside some neighbourhood of 0, and that f is δ-Lipschitz for some constant $\delta < 1 - r - \varepsilon$. (We will justify these assumptions later in the general proof of Hartman-Grobman.) We are looking for $H = \mathrm{Id} + h$, with h bounded and continuous.

Fact: The space of continuous bounded maps $h : \mathbb{R}^n \to \mathbb{R}^n$ is a Banach space for the supremum norm $||h||_\infty := \sup ||h(x)||$.

We now start from the equation that we would like to solve, and see that a solution indeed exists. We want to find H such that $H \circ T = F \circ H$. This is equivalent to the equation $H = F \circ H \circ T^{-1}$. Inserting $\mathrm{Id} + h$ for H and $T + f$ for F, we obtain

$$\mathrm{Id} + h = (T + f) \circ (\mathrm{Id} + h) \circ T^{-1},$$

and singling out h on the left side, this is equivalent to

$$h = ((T + f) \circ h + f) \circ T^{-1}.$$

Now, we define a self-map of the Banach space of continuous bounded maps by $\varphi : h \mapsto ((T + f) \circ h + f) \circ T^{-1}$.

Claim. The map φ is a contraction.

Proof of the claim:

$$\begin{aligned} ||\varphi(h) - \varphi(h')||_\infty &= ||(T + f) \circ h - (T + f) \circ h'||_\infty \\ &\leq ||T \circ h - T \circ h'||_\infty + ||f \circ h - f \circ h'||_\infty \\ &\leq ||T|| \cdot ||h - h'||_\infty + \delta ||h - h'||_\infty \\ &\leq (r + \varepsilon + \delta)||h - h'||_\infty. \end{aligned}$$

For the first equality, we used that T^{-1} is bijective and hence does not change the supremum norm. By the Banach fixed point theorem, we now know that a unique fixed point h of φ exists. This unique fixed point h solves the equation $h = ((T+f)\circ h + f)\circ T^{-1}$ and hence also the initial equation $H \circ T = F \circ H$.

One final detail is left to show, namely that H is a homeomorphism. We verify this as follows. Using the exact same argument, but exchanging F and T, we obtain a unique map $G = (\mathrm{Id} + g)$ such that

$$G \circ F = T \circ G.$$

In particular, we obtain that

$$(G \circ H) \circ T = T \circ (G \circ H).$$

We know there exists a unique map $G \circ H$ that solves this equation (for example, by letting $F = T$ in our argument). However, one solution is the identity, and therefore we get $G \circ H = \mathrm{Id}$. This implies that G and H are inverses of each other and, in particular, H is a homeomorphism. □

Proof of an expanding case. The expanding case can be proved in the same way as the contracting case, replacing T with it's inverse. Assume that for every eigenvalue λ of T, we have $|\lambda| \geq R > 1$. Choose some $\varepsilon > 0$ such that $\frac{1}{R} + \varepsilon < 1$, and choose a norm $||\cdot||$ on $\mathbb{R}^n$ such that $||T^{-1}|| \leq \frac{1}{R} + \varepsilon$. We now write $F^{-1} = T^{-1} + \overline{f}$ and assume again (and justify later) that $F = T$ outside some neighbourhood of 0 and that $\overline{f}$ is δ-Lipschitz for $\delta < \frac{1}{R} - \varepsilon$.

We are done if we can find a solution h to

$$\begin{aligned} H \circ T = F \circ H &\iff F^{-1} \circ H \circ T = H \\ &\iff (T^{-1} + \overline{f}) \circ (\mathrm{Id} + h) \circ T = \mathrm{Id} + h \\ &\iff (\overline{f} + (T^{-1} + \overline{f}) \circ h) \circ T = h, \end{aligned}$$

and we can do so using the exact same Banach fixed point argument as in the contracting case. □

In order to prove Hartman-Grobman, we now decompose $\mathbb{R}^n$ according to the contracting and the expanding directions of the linear part. We need some definitions to make this precise.

Definition 2.10. *For a hyperbolic matrix $A \in \mathrm{Mat}_n(\mathbb{R})$, we define*

$$\begin{aligned} E^s &:= \text{generalised eigenspace associated to eigenvalues } |\cdot| < 1 \\ &= \bigoplus_{n\geq 1,\ |\lambda|<1} \ker(A - \lambda \mathrm{Id})^n \\ E^u &:= \text{generalised eigenspace associated to eigenvalues } |\cdot| > 1 \\ &= \bigoplus_{n\geq 1,\ |\lambda|>1} \ker(A - \lambda \mathrm{Id})^n \end{aligned}$$

We note that hyperbolicity of a matrix $A \in \mathrm{Mat}_n(\mathbb{R})$ ensures $\mathbb{R}^n = E^s \oplus E^u$.

Definition 2.11. *For a hyperbolic matrix $A \in \mathrm{Mat}_n(\mathbb{R})$, a norm $||\cdot||$ on $\mathbb{R}^n$ is called* adapted *to A if*

i) $||A|_{E^s}|| < 1$,

ii) $||(A|_{E^u})^{-1}|| < 1$,

iii) if $x = x_s + x_u$, where $x_s \in E^s$ and $x_u \in E^u$, then $||x|| = \max\{||x_s||, ||x_u||\}$.

The existence of adapted norms follows from Proposition 2.5. We are now ready for the general proof of Hartman-Grobman.

Proof of Hartman-Grobman. The rough idea is to split $\mathbb{R}^n$ into an expanding and a contracting subspace for T, and to adapt our argument to this splitting. Our first step is to justify the assumptions we left open in the contracting and the expanding cases. For this, let b_α be a bump function on $\mathbb{R}^n$ that equals 1 on $B_\alpha(0)$ and 0 outside $B_{2\alpha}(0)$, such that $||d_v b_\alpha|| \leq \frac{2}{\alpha}$ for $v \in \mathbb{R}^n$. We will soon choose the number α so that it suits our purpose.

We write $F = T + f$ on Ω, and consider $F_\alpha = T + b_\alpha f$ on $\mathbb{R}^n$. Then, on $B_\alpha(0)$ we have $F_\alpha = F$, and outside $B_{2\alpha}(0)$, we have $F_\alpha = T$.

Claim: For $\delta > 0$, there exists an α such that $b_\alpha f$ is δ-Lipschitz.

Proof of the claim: By the mean value theorem, we know

$$\mathrm{Lip}(b_\alpha f) \leq \max_{v \in \mathbb{R}^n} ||d_v b_\alpha f|| = \max_{v \in B_{2\alpha}(0)} ||d_v b_\alpha f||.$$

In order to verify the claim, we now calculate

$$d_v(b_\alpha f) = f(v) \cdot d_v b_\alpha + b_\alpha(v) d_v f.$$

Since $||d_v b_\alpha|| \leq \frac{2}{\alpha}$ and $b_\alpha(v) \leq 1$, this implies

$$||d_v(b_\alpha f)|| \leq \frac{2}{\alpha} \sup_{v \in B_{2\alpha}(0)} ||f(v)|| + \sup_{v \in B_{2\alpha}(0)} ||d_v f||,$$

which tends to 0 as $\alpha \to 0$. Indeed, recall that $F = d_0 F + f$, so $d_v f$ is continuous in v and $d_0 f$ is the zero map, and $||f(v)||$ tends to zero quicker than $||v||$ by the definition of differentiability. This proves the claim.

By assumption, the linear map T is hyperbolic. We decompose along the expanding and contracting generalised eigenspaces of T: $\mathbb{R}^n = E^s \oplus E^u$, and $T = T_{ss} \oplus T_{uu}$. Furthermore, let $||\cdot||$ be a norm on $\mathbb{R}^n$ adapted to T. We now choose a δ such that $0 < \delta < 1 - \max\{||T_{ss}||, ||T_{uu}||^{-1}\}$ and use the above claim to choose α such that $b_\alpha f$ is δ-Lipschitz.

We again consider the equation

$$\begin{aligned} H \circ T = F \circ H &\iff (\mathrm{Id} + h) \circ T = (T + f) \circ (\mathrm{Id} + h) \\ &\iff h \circ T = T \circ h + f \circ (\mathrm{Id} + \mathrm{h}) \end{aligned}$$

and split it according to $\mathbb{R}^n = E^s \oplus E^u$ (a single subscript denotes the corresponding coordinate in the splitting):

$$\begin{aligned} h_s \circ T &= T_{ss} \circ h_s + f_s \circ (\mathrm{Id} + h), \\ h_u \circ T &= T_{uu} \circ h_u + f_u \circ (\mathrm{Id} + h). \end{aligned}$$

Solving for h_s (on the left side), and h_u (on the right side), these two equations are equivalent to:

$$\begin{aligned} h_s &= (T_{ss} \circ h_s + f_s \circ (\mathrm{Id} + h)) \circ T^{-1}, \\ h_u &= T_{uu}^{-1} \circ (h_u \circ T - f_u \circ (\mathrm{Id} + h)). \end{aligned}$$

We now define the operator $\varphi = (\varphi_s, \varphi_u)$ on the Banach space of continuous bounded functions $h = (h_s, h_u) : \mathbb{R}^n \to \mathbb{R}^n = E^s \oplus E^u$ by the equations:

$$\begin{aligned} \varphi_s(h) &= (T_{ss} \circ h_s + f_s \circ (\mathrm{Id} + h)) \circ T^{-1}, \\ \varphi_u(h) &= T_{uu}^{-1} \circ (h_u \circ T - f_u \circ (\mathrm{Id} + h)). \end{aligned}$$

The norm $|| \cdot ||$ on $\mathbb{R}^n$ is adapted to the splitting $\mathbb{R}^n = E^s \oplus E^u$, so in order to show that φ is a contraction, it suffices to estimate $||\varphi_s(h) - \varphi_s(h')||_\infty$ and $||\varphi_u(h) - \varphi_u(h')||_\infty$. We have

$$\begin{aligned} ||\varphi_s(h) - \varphi_s(h')||_\infty &= ||T_{ss} \circ (h_s - h'_s) + f_s \circ (h_s - h'_s)||_\infty \\ &\le ||T_{ss}|| \cdot ||h_s - h'_s||_\infty + ||f_s \circ (h_s - h'_s)||_\infty \\ &\le (||T_{ss}|| + \delta)||h - h'||_\infty. \end{aligned}$$

Similarly, we obtain

$$\begin{aligned} ||\varphi_u(h) - \varphi_u(h')||_\infty &= ||T_{uu}^{-1}|| \cdot ||(h_u - h'_u) \circ T - f_u \circ (h - h')||_\infty \\ &\le ||T_{uu}^{-1}||(||(h_u - h'_u) \circ T||_\infty + ||f_u \circ (h - h')||_\infty) \\ &\le ||T_{uu}^{-1}||(1 + \delta)||(h - h')||_\infty \\ &\le (||T_{uu}^{-1}|| + \delta)||(h - h')||_\infty. \end{aligned}$$

Since we chose δ such that $0 < \delta < 1 - \max\{||T_{ss}||, ||T_{uu}||^{-1}\}$, the operator φ is indeed a contraction, so there exists a unique fixed point $h = (h_s, h_u)$ solving our equation $H \circ T = F \circ H$.

One can show that H is a homeomorphism exactly as in the expanding case we already dealt with. □

Remark 2.12. There is also a continuous-time version of the theorem of Hartman-Grobman: let $X : \Omega \to \mathbb{R}^n$ be a C^1 vector field with $X(0) = 0$ such that flow is defined on $(-\varepsilon, \varepsilon)$. If d_0X has no eigenvalue with real part $= 0$, then there exists a neighbourhood U of the origin, as well as a homeomorphism $H : U \to H(U)$ fixing the origin such that H conjugates the flow of X with the flow of d_0X.

2.4 Nonnegative matrices

This section closely follows the exposition of the appendix in Bogopolski's book [1]. A matrix $A \in \mathrm{Mat}_n(\mathbb{R})$ is called *nonnegative* or *positive* if all its coefficients are nonnegative or positive, respectively. We denote this by $A \geq 0$ or $A > 0$, respectively.

Definition 2.13. *A matrix $A \in \mathrm{Mat}_n(\mathbb{R})$ is called reducible if there exists a permutation matrix P such that*

$$P^{-1}AP = \begin{pmatrix} X & Y \\ 0 & Z \end{pmatrix},$$

where X and Z are square matrices. A matrix $A \in \mathrm{Mat}_n(\mathbb{R})$ is called irreducible if it is not reducible.

Exercise 2.14. Show that if $A \in \mathrm{Mat}_n(\mathbb{R})$ is irreducible, then so is $A^\top$.

With every matrix $A \in \mathrm{Mat}_n(\mathbb{R})$, $A \geq 0$, we associate an oriented graph $\Gamma(A)$ in the following way: the vertices are numbered $1, \dots, n$ and there is an oriented edge from vertex i to vertex j if and only if $A_{ij} > 0$. The following exercise reformulates irreducibility of a matrix in terms of the graph $\Gamma(A)$.

Exercise 2.15. Let $A \in \mathrm{Mat}_n(\mathbb{R})$, $A \geq 0$. Show that the following statements are equivalent:

a) A is irreducible,

b) for all $0 \leq i, j \leq n$, there exists $N(i,j) \in \mathbb{N}$ such that $(A^{N(i,j)})_{ij} > 0$,

c) for each pair of vertices i and j of $\Gamma(A)$, there exists a path of oriented edges from i to j.

Lemma 2.16. *Let $A \in \mathrm{Mat}_n(\mathbb{R})$, $0 \neq A \geq 0$ be irreducible. Then the matrix $B = \sum_{i=0}^{n-1} A^i$ is positive.*

Proof. From $A^0 = I_n$, we know $B_{ii} > 0$. Let us now show that $B_{ij} > 0$ for $i \neq j$. By Exercise 2.15, there exists an oriented path in $\Gamma(A)$ from vertex i to vertex j. We can assume that this path visits each vertex at most once, and hence is of length $l \leq n - 1$. If this path is given by the sequence of vertices $v_{i_0}, v_{i_1}, \dots, v_{i_l}$, then this implies $A_{i_0 i_1} A_{i_1 i_2} \cdots A_{i_{l-1} i_l} > 0$. This gives us $(A^l)_{ij} > 0$ and hence $B_{ij} > 0$. □

Theorem 2.17 (Perron-Frobenius)**.** *Let* $A \in \mathrm{Mat}_n(\mathbb{R})$*,* $0 \neq A \geq 0$ *be irreducible. Then* A *has a unique positive eigenvector up to scaling, and the corresponding eigenvalue is positive.*

Proof. Let $A \in \mathrm{Mat}_n(\mathbb{R})$, $0 \neq A \geq 0$ be irreducible. We denote by

$$\Delta = \{x \in \mathbb{R}^n : 0 \leq x, ||x|| = 1\}$$

the set of non-negative unit vectors in $\mathbb{R}^n$, and we let $u = (1, \ldots, 1)$ the row vector of size n with constant coefficients 1. Here, we use the one-norm: $||x|| = ||x||_1 = \sum_{i=1}^n |x_i|$, where $x \in \mathbb{R}^n$. The number

$$\lambda = \sup\{\rho : \text{ there exists } x \in \Delta \text{ such that } Ax \geq \rho x\}$$

must be finite. Indeed, a number ρ satisfying the condition in the definition of λ is bounded by the sum of the coefficients of A: the inequality $Ax \geq \rho x$ for $x \in \Delta$ implies $uAu^\top \geq uAx \geq \rho ux = \rho$. Using the compactness of Δ, one can show that the supremum is attained, that is, there exists $y \in \Delta$ such that $Ay \geq \lambda y$.

We now want to show $Ay = \lambda y$. For a contradiction, assume $Ay \neq \lambda y$. Then we have $BAy > \lambda By$, where B is the matrix of Lemma 2.16, which is positive and has diagonal entries ≥ 1. Since $AB = BA$, this allows us to deduce $Ax > \lambda x$ for $x = By/||By|| \in \Delta$. This contradicts the maximality of λ and hence we must have $Ay = \lambda y$. This proves existence of a non-negative eigenvector. From $0 < By = \sum_{k=0}^{n-1} \lambda^k y$ we further deduce $y > 0$, which proves the existence of a positive eigenvector. Furthermore, $Ay = \lambda y$ now directly implies $\lambda > 0$.

What is left to show is the uniqueness of y. We first note that the exact same argument can be used to produce a row vector $z > 0$ and a number $\mu > 0$ such that $zA = \mu z$. Then we have $\mu zy = zAy = \lambda zy$ and $zy > 0$, which implies $\mu = \lambda$. Now, let y' be a positive eigenvector of A and let λ' be the corresponding eigenvalue. As above, we argue that $\lambda' = \mu$ and hence $\lambda' = \lambda$. We now assume for a contradiction that y' is not a multiple of y. If this is the case, then the points $y'/||y'||, y/||y|| \in \Delta$ are distinct. The line containing these two points consists of eigenvectors to λ and intersects the boundary of $\{x \in \mathbb{R}^n : 0 \leq x\}$ in some point v. In particular, $v/||v|| \in \Delta$ is an eigenvector of A for the eigenvalue λ, and hence also an eigenvector of B. In particular, Bv is a multiple of v and at least one coordinate must be zero. This is a contradiction, since for all vectors $x \geq 0$ we have $Bx > 0$. We have shown that y' must be a multiple of y. □

The eigenvector and the eigenvalue obtained via Theorem 2.17 are called the *Perron-Frobenius eigenvector* and *Perron-Frobenius eigenvalue*, respectively. They enjoy the following properties.

Proposition 2.18. *Let $A \in \mathrm{Mat}_n(\mathbb{R})$, $0 \neq A \geq 0$ be a irreducible, and let λ be its Perron-Frobenius eigenvalue. Then we have the following:*

a) the Perron-Frobenius eigenvalue of $A^\top$ also equals λ,

b) for each eigenvalue τ of A, we have $|\tau| \leq \lambda$,

c) up to real multiples, the eigenvector for λ is unique.

Proof. The statement a) follows directly from the proof of Theorem 2.17.

As for b), assume we are given $\tau \in \mathbb{C}$, and $0 \neq x \in \mathbb{C}^n$ with $Ax = \tau x$. Let $x' = (|x_1|, \ldots, |x_n|)^\top$. From

$$|\tau| x_i' = |a_{i1}x_1 + \cdots + a_{in}x_n| \leq |a_{i1}x_1| + \cdots + |a_{in}x_n|$$

it follows that $Ax' \geq |\tau| x'$. For a row vector $z > 0$ with $zA = \lambda z$, we now get $\lambda z x' = zAx' \geq |\tau| z x'$. From $zx' > 0$ we conclude $\lambda \geq |\tau|$.

As for c), we let y be a positive and v be an arbitrary real eigenvector for the eigenvalue λ. For $r \in \mathbb{R}$ large enough, $v + ry$ is a positive eigenvector for A. By Theorem 2.17, $v + ry$ and hence v must be a real multiple of y. □

Exercise 2.19. Strengthen the assertion b) of Proposition 2.18 in case of a positive matrix $A \in \mathrm{Mat}_n(\mathbb{R})$, that is, $A > 0$. In this case, show that $|\tau| < \lambda$ holds for every eigenvalue $\tau \neq \lambda$ of A.

Exercise 2.20. Formulate a problem about the evolution of a population over time, and find an equilibrium state using Perron-Frobenius theory.

Exercise 2.21. Make the example of ranking the popularity of twitter profiles from the introduction rigorous using Perron-Frobenius theory.

2.4.1 Convergence to the positive eigendirection

The following set of exercises give a way to obtain the theorem of Perron-Frobenius from a more dynamical perspective, using the Banach fixed point theorem. As a result, we will see that the iterates of any positive vector will converge to the unique positive eigendirection of a positive matrix.

Exercise 2.22. Let $\Delta^+ = \{x \in \mathbb{R}^n : 0 < x, ||x|| = 1\}$ be the set of positive unit vectors in $\mathbb{R}^n$. Show that

$$d(x, y) := \log \left(\frac{\max_i \frac{x_i}{y_i}}{\min_i \frac{x_i}{y_i}} \right)$$

defines a metric on Δ^+, where we use the notation $x = (x_1, \ldots, x_n)$.

Exercise 2.23. For a positive matrix $A \in \mathrm{Mat}_n(\mathbb{R}_{>0})$, show that

$$f : \Delta^+ \to \Delta^+, \quad x \mapsto \frac{Ax}{||Ax||}$$

is a contraction with respect to the metric d.

Exercise 2.24. Using the Banach fixed point theorem, show that a positive matrix $A \in \mathrm{Mat}_n(\mathbb{R}_{>0})$ has a unique eigenvector $v \in \Delta^+$, and $f^n(x) \to v$ as $n \to \infty$, for every $x \in \Delta^+$.

Chapter 3

Topological dynamics

In this chapter, the underlying set X of the dynamical system is supposed to be a topological space, and the map[1] is supposed to be continuous. We will use topological properties that are invariant under topological conjugacy. Having a dense orbit is one example of such a property.

3.1 Limits and minimal sets

Definition 3.1. *Let (X, f) be a disctrete-time topological dynamical system. The ω-limit of $x \in X$ is defined to be*

$$L^f_\omega(x) := \{\lim_{i\to\infty} f^{n_i}(x) : n_i \to \infty, f^{n_i}(x) \text{ converges}\}.$$

Similarly, the α-limit of $x \in X$ is defined to be

$$L^f_\alpha(x) := \{\lim_{i\to\infty} f^{n_i}(x) : n_i \to -\infty, f^{n_i}(x) \text{ converges}\}.$$

For the example of rotations $r_\alpha : \mathbb{R}/\mathbb{Z} \to \mathbb{R}/\mathbb{Z}$, we have two cases. If α is rational, then both the ω-limit and the α-limit of any point $x \in \mathbb{R}/\mathbb{Z}$ consist of the orbit of x under r_α. If α is irrational, then both limits consist of the entire unit circle $\mathbb{R}/\mathbb{Z}$.

Exercise 3.2. Show that $L^f_\omega(x)$ and $L^f_\alpha(x)$ are closed and f-invariant sets (that is, closed sets A satisfying $f(A) \subset A$).

Definition 3.3. *A nonempty closed f-invariant subset $Y \subset X$ is called minimal set if there exists no proper nonempty closed f-invariant subset of Y. If X itself is a minimal set, then we say f is minimal.*

[1] or the flow–but we will mainly focus on discrete time-systems, the definitions being similar in the continuous-time case

A simple example of a minimal set is given by a periodic orbit. Not every map has a periodic orbit, but nevertheless, if we assume X to be compact, then it must contain a minimal set for any continuous map $f : X \to X$. We will state the result without giving the proof, which is based on a combination of Zorn's lemma and the compactness of X.

Proposition 3.4. *Let $f : X \to X$ be a topological dynamical system. If X is compact, then X contains a minimal set for f.*

Exercise 3.5. Let X be a compact topological space and let $f : X \to X$ be a continuous map. Show that

a) $Y \subset X$ is minimal if and only if $L^f_\omega(y) = Y$ for all $y \in Y$,

b) $Y \subset X$ is minimal if and only if the forward orbit of every $y \in Y$ is dense in Y.

3.2 Topological transitivity

Definition 3.6. *We say that a topological dynamical system $f : X \to X$ is topologically transitive if there exists a point $x \in X$ whose forward orbit is dense in X.*

Irrational rotations of the unit circle are examples of topological transitive dynamical systems. Density of every forward orbit is shown in Exercise 1.7.

Exercise 3.7. Show that if X has no isolated points, then some forward orbit $\mathcal{O}^+_f(x)$ is dense in X if and only if the ω-limit $L^f_\omega(x)$ is dense in X. Give a counterexample to this equivalence for a topological space X that has isolated points.

The following lemma can be used to produce plenty of non-examples of topological transitivity.

Lemma 3.8. *Let X be a normal Hausdorff space. If $f : X \to X$ has an attractive periodic orbit γ and at least one other orbit disjoint from γ, then f is not topologically transitive.*

Proof. Let $\gamma = \{f^n(x)\}$ be the periodic attractive orbit. There exists an open neighbourhood U of γ such that $f(U) \subset U$. For $y \in X$, if $f^k(y) \in U$ for some $k \in N$, then the forward orbit of y is not dense in $X \backslash U$. If $f^k(y) \notin U$ for all k, then the orbit is not dense in U. □

The following proposition gives a condition for topological transitivity.

Proposition 3.9. *Let X be locally-compact and Hausdorff, and $f : X \to X$ be continuous. If for all nonempty open subsets $U, V \subset X$ there exists an integer $n \geq 0$ such that $f^n(U) \cap V \neq \varnothing$, then f is topologically transitive.*

Proof. By hypothesis, for each nonempty open subset $V \subset X$, we have that $\bigcup_{n\in\mathbb{N}} f^{-n}(V)$ is dense in X. Let $\{V_i\}$ be a countable basis for the topology of X, and define

$$Y := \bigcap_i \bigcup_{n\in\mathbb{N}} f^{-n}(V_i).$$

As a countable intersection of open and dense subsets, Y is itself dense in X by Baire's category theorem. In particular, Y is nonempty, so now let $y \in Y$. By definition, the forward orbit of y visits every basis set V_i of the topology, and hence is dense in X. □

Example 3.10. The matrix $A = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$ defines a linear map $\mathbb{R}^2 \to \mathbb{R}^2$ that preserves the integer lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$. In particular there is a well-defined quotient map $T_A : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$. Since the determinant of the matrix A is 1, we can also define an inverse map. This implies that the map T_A is a homeomorphism of the torus $\mathbb{R}^2/\mathbb{Z}^2$.

The matrix $A = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$ has Perron-Frobenius eigenvector $v = (\varphi, 1)^\top$ to the eigenvalue φ^2, where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. The vector v has irrational slope, and one can show that $\mathbb{R}v$ is dense in $\mathbb{R}^2/\mathbb{Z}^2$. We know from Perron-Frobenius theory that an orbit will get "projectively close" to the direction v, so it seems reasonable to assume that $T_A : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$ is topologically transitive.

In the next section, we will give a proof of an even stronger property: topological mixing.

3.3 Topological mixing

Definition 3.11. *We say that a topological dynamical system $f : X \to X$ is topologically mixing if for all nonempty open subsets $U, V \subset X$, there exists an integer $n_0 \geq 0$ such that for all $n \geq n_0$, we have $f^n(U) \cap V \neq \varnothing$.*

We remark that for nice enough topological spaces (certainly if X is a manifold), topological mixing implies topological transitivity by Proposition 3.9. The converse, however, does not hold.

Exercise 3.12. Show that $r_\alpha : \mathbb{R}/\mathbb{Z} \to \mathbb{R}/\mathbb{Z}$ is not topological mixing.

More generally:

Exercise 3.13. Show that an isometry of a metric space (X, d) with at least two points is not topologically mixing.

We now get back to the homeomorphism $T_{\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)} : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$.

Proposition 3.14. *The homeomorphism $T_{\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)} : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$ is topologically mixing.*

Proof. Recall that $\mathbb{R}v$ is dense in $\mathbb{R}^2/\mathbb{Z}^2$, where $v = (\varphi, 1)^\top$ is the Perron-Frobenius eigenvector of the matrix $A = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, and where φ is the golden ratio.

For each $\varepsilon > 0$, there exists a collection of ε-balls centered on $\mathbb{R}v$ that covers the torus $\mathbb{R}^2/\mathbb{Z}^2$. By compactness, there exists a finite subcollection of such balls that also covers the torus. We deduce that there exists a bounded segment $S_0 \subset \mathbb{R}v$ whose ε-neighbourhood covers the torus $\mathbb{R}^2/\mathbb{Z}^2$.

Translations in $\mathbb{R}^2$ are isometries. We deduce that for any $w \in \mathbb{R}^2$ fixed, the ε-neighbourhood of the segment $S_0 + w \subset w + \mathbb{R}v$ also covers the torus $\mathbb{R}^2/\mathbb{Z}^2$.

In summary, we have shown that for each $\varepsilon > 0$, there exists $L(\varepsilon) > 0$ such that every segment S of length $\geq L(\varepsilon)$ that is parallel to $\mathbb{R}v$ is ε-dense in the torus $\mathbb{R}^2/\mathbb{Z}^2$, that is, $d(y, S) \leq \varepsilon$ for all $y \in \mathbb{R}^2/\mathbb{Z}^2$.

Now, let $U, V \subset \mathbb{R}^2/\mathbb{Z}^2$ be two nonempty open subsets of the torus. Take a point $y \in V$ and let $\varepsilon > 0$ such that $B_\varepsilon(y) \subset V$. The subset U contains a segment of length $\delta > 0$ in some translate of $\mathbb{R}v$.

Let $\lambda = \varphi^2$ be the Perron-Frobenius eigenvalue of A, and choose N such that $\lambda^N \delta > L(\varepsilon)$. Now, for all $n \geq N$, we have that $A^n(U)$ contains a segment of length $\lambda^N \delta > L(\varepsilon)$ in some translate of $\mathbb{R}v$. This implies that $A^n(U)$ is ε-dense in the torus, and in particular must intersect V. □

Exercise 3.15. Generalise Proposition 3.14 to any matrix $A \in \mathrm{GL}_2(\mathbb{Z})$ with $|\mathrm{trace}(A)| > 2$.

Exercise 3.16. Show that topological transitivity and topological mixing are invariants of topological conjugacy.

Chapter 4

Vector fields in the plane

The goal of this chapter is to study flows induced by vector fields in the real plane $\mathbb{R}^2$. In particular, we want to prove the structural theorem by Poincaré–Bendixson mentioned in the introduction.

Throughout the chapter, the notion of ω- and α-limits will play a prominent role. We therefore quickly repeat the definition, this time for flows.

Definition 4.1. *Let φ^t be a continuous flow on a topological space X, and let $x \in X$ The ω-limit of x is defined to be*

$$L_\omega^\varphi(x) = \{\lim_{i\to\infty} \varphi^{t_i}(x) : t_i \to \infty,\ \varphi^{t_i}(x) \text{ converges}\}.$$

Similarly, the α-limit is defined to be

$$L_\alpha^\varphi(x) = \{\lim_{i\to\infty} \varphi^{t_i}(x) : t_i \to -\infty,\ \varphi^{t_i}(x) \text{ converges}\}.$$

4.1 Vector fields on the line

Before approaching the statement of the theorem of Poincaré–Bendixson, we take a step back and consider, for one moment, vector fields and their associated flows one dimension lower, that is, on the real line $\mathbb{R}$. We leave the qualitative study of the limits in this case as an exercise.

Exercise 4.2. Let $F : \mathbb{R} \to \mathbb{R}$ be a continuously differentiable vector field on $\mathbb{R}$ and let φ^t be the associated flow. Show that the ω- and the α-limit of every point $x \in \mathbb{R}$ is determined by the following data:

- the zero set of the vector field F, which is a closed set $A \subset \mathbb{R}$,
- the sign of F on the connected component $x \in \mathbb{R} \setminus A$.

4.2 Limits in $\mathbb{R}^2$

In the case of vector fields on the line, every ω- or α-limit is either empty or a single point. For the plane, we have encountered a more intriguing limit set in Exercise 1.27: in this exercise, for all points except the origin the ω-limit is a whole periodic orbit.

There is more that can happen in the plane: instead of being a single periodic orbit, some ω-limit could be a union of orbits connecting zeros of the vector field F. To create an example for this, you can take the vector field from Exercise 1.27 and change it slightly so that it has zeros on the unit circle (which in the exercise is the period orbit), but still the unit circle is the ω-limit of all points except the origin.

It is the content of the theorem of Poincaré–Bendixson that basically nothing else can happen[1].

Theorem 4.3 (Poincaré–Bendixson 1901)**.** *Let $F : \mathbb{R}^2 \to \mathbb{R}^2$ be a continuously differentiable vector field, and let φ^t be its associated flow. Let $x \in \mathbb{R}^2$ be such that the flow $\varphi^t(x)$ is defined for all $t > 0$ and stays bounded. Then either*

a) $L^\varphi_\omega(x)$ is a periodic orbit, or

b) for all $y \in L^\varphi_\omega(x)$, the limits $L^\varphi_\omega(y)$ and $L^\varphi_\alpha(y)$ consist of zeroes of F.

In particular, it follows directly that if the ω-limit of a point x contains no fixed point of the flow, that is, a zero of F, then it consists of a single periodic orbit. This happens for example if an orbit stays within a bounded domain of the plane that does not contain any zero of the vector field.

Exercise 4.4. In the setting of Theorem 4.3, show that

a) the limit set $L^\varphi_\omega(x)$ is compact,

b) the limit sets $L^\varphi_\omega(y)$ and $L^\varphi_\alpha(y)$ are connected,

c) if F has finitely many zeros, then $L^\varphi_\omega(x)$ consists of one of the following: one fixed point, one periodic orbit, or a finite number of fixed points and a set of orbits γ_i such that $L^\varphi_\omega(\gamma_i)$ and $L^\varphi_\alpha(\gamma_i)$ consist of one fixed point each.

Exercise 4.5. Find a closed surface Σ of genus > 0 and a vector field F on Σ such that the statement of Theorem 4.3 fails.

[1] We state the theorem of Poincaré–Bendixson for continuously differentiable vector fields. However, all we need, again and again, in the proof is the continuity of the flow and the uniqueness of local solutions. It would therefore be sufficient to choose the vector field locally Lipschitz.

4.3 Transversals and first-return maps

Our proof of Theorem 4.3 relies heavily on the notion of a *transversal* to the flow. As always, we let $F : \mathbb{R}^2 \to \mathbb{R}^2$ be continuously differentiable.

Definition 4.6. *A transversal to F through $x \in \mathbb{R}^2$ is a continuously differentiable embedded segment $S \subset \mathbb{R}^2$ containing x and transverse to the vector field F.*

The existence of transversal at nonsingular point of F follows directly from Theorem 1.22. We record this fact as the following lemma.

Lemma 4.7. *If $F(x) \neq 0$, then there exists a transversal through x.*

Definition 4.8. *A periodic orbit γ of the flow φ is called attractive if there exists an open neighbourhood U of γ such that for all $y \in U$, $L_\omega^\varphi(y) = \gamma$. A periodic orbit γ is called repulsive if there exists an open neighbourhood U of γ such that for all $y \in U$, $L_\alpha^\varphi(y) = \gamma$.*

The definition of repulsiveness of a periodic orbit γ given here is only to be used like this for vector fields in the plane.

Definition 4.9 (First-return map for a transversal)**.** *Let x_0 be a periodic point of period T, which simply means that $\varphi^T(x_0) = x_0$, and let S be a transversal through x_0. Choose a neighbourhood $U \subset S$ of x_0 such that $\varphi^t(x_0) \cap U = \{x_0\}$. Let*

$$W := \{\varphi^t(x) : x \in S, -\varepsilon < t < \varepsilon\}$$

with $\varepsilon > 0$ small enough such that $\varphi^t(y) = \varphi^{t'}(y)$ for $-\varepsilon < t, t' < \varepsilon$ implies $t = t'$. Further, define

$$N := \{x \in U : \varphi^T(x) \in W\}.$$

By the continuity of the flow φ, N is an open neighbourhood of x_0 in S. We finally define the first-return map $P : N \to S$ by

$$x \mapsto P(x) = \varphi^{T+s(x)}(x),$$

where $s(x)$ is the unique number such that $|s(x)| < \varepsilon$ and $\varphi^{T+s(x)}(x) \in S$.

Lemma 4.10. *A periodic orbit γ through x_0 is attractive if and only if x_0 is attractive for the first-return map P.*

Proof. We first deal with the only if direction. For any point $x \in N$, the sequence $\{P^i(x)\}_{i\in\mathbb{N}}$ equals some sequence $\{\varphi^{t_i}(x)\}_{i\in\mathbb{N}}$. Since γ is compact, the latter sequence has an accumulation point y on γ. We have $y \in N$ which implies $y = x_0$, since x_0 is the only point of intersection of the periodic orbit γ and the transversal N. This implies $P^i(x) \to x_0, i \to \infty$.

For the if direction, we note that by hypothesis, there exsists some open neighbourhood U of x_0 in N such that for every $x \in U, P^i(x) \to x_0$ for $i \to \infty$. We want to show that $L^{\varphi}_{\omega}(x) = \gamma$. This then implies that some neighbourhood U' of γ, obtained by taking the forward image of U under the flow, satisfies the condition for attractivity.

$\underline{L^{\varphi}_{\omega}(x) \subset \gamma}$: suppose $y \notin \gamma$ and $\{\varphi^{t_i}(x)\}$ accumulates on y. Pick a subsequence such that $\varphi^{t_i} \to y$. There exist $s_i \to s$ such that $\varphi^{t_i+s_i} \in N$. Then $\varphi^{t_i+s_i}$ must converge to x_0 by our assumption on P. But it also converges, by definition, to $\varphi^s(y)$. Hence $\varphi^s(y) = x_0$, so $y \in \gamma$.

$\underline{\gamma \subset L^{\varphi}_{\omega}(x)}$: If $y \in \gamma$, then $y = \varphi^t(x_0)$ for some t. Now if $\varphi^{t_i}(x) \to x_0$, then $\varphi^{t+t_i}(x) \to y$, so $y \in L^{\varphi}_{\omega}(x)$. □

Exercise 4.11. For the following planar vector fields, determine their limit cycles, and discuss whether they are attracting, repelling, or neither of the two:

a) $F(x, y) = \alpha(x^2 + y^2 - 1)(x, y) + (y, -x)$, for $\alpha \in \mathbb{R}$,

b) $F(x, y) = (x^2 + y^2 - 1)^2(x, y) + (y, -x)$,

c) $F(x, y) = \frac{\sin(x^2+y^2)}{x^2+y^2}(x, y) + (y, -x)$.

4.4 The proof of Poincaré–Bendixson

As always, we take $F : \mathbb{R}^2 \to R^2$ be continuously differentiable and we let φ^t be the associated flow. We furthermore place ourselves in the assumptions for Theorem 4.3. We need two more lemmas for the proof of Theorem 4.3.

Lemma 4.12. *Let γ be the orbit of $y \in \mathbb{R}^2$. If $y \in L^{\varphi}_{\omega}(x)$, then $\gamma \subset L^{\varphi}_{\omega}(x)$, and further $L^{\varphi}_{\omega}(y) \cup L^{\varphi}_{\alpha}(y) \subset L^{\varphi}_{\omega}(x)$.*

Proof. By assumption, there exists a sequence $t_i \to \infty$ with $\varphi^{t_i}(x) \to y$. It follows that $\varphi^{t_i+T} \to \varphi^T(y)$ for all T. In particular, $\varphi^T(y) \in L^{\varphi}_{\omega}(x)$ for all T, and in particular $\gamma \subset L^{\varphi}_{\omega}(x)$. This proves the first statement. For the second statement, we note that $L^{\varphi}_{\omega}(x)$ is a closed subset of the plane containing γ. This means that any set contained in the topological closure of γ is also a subset of $L^{\varphi}_{\omega}(x)$. In particular, $L^{\varphi}_{\omega}(y)$ and $L^{\varphi}_{\alpha}(y)$, being contained in the closure of γ, are subsets of $L^{\varphi}_{\omega}(x)$. □

Lemma 4.13 (Key lemma)**.** *Let I be a transversal to F. If there exist real numbers $t_1 < t_2 < t_3$ such that $\varphi^{t_1}(x), \varphi^{t_2}(x), \varphi^{t_3}(x) \in I$, then $\varphi^{t_2}(x)$ lies between $\varphi^{t_1}(x)$ and $\varphi^{t_3}(x)$ on I.*

It is recommended to accompany the following proof with at least one sketch.

Proof. Consider the simple closed curve c defined by $\varphi^{[t_1,t_2]}(x)$ and the portion of I between $\varphi^{t_1}(x)$ and $\varphi^{t_2}(x)$. By the Jordan curve theorem, the complement of c in the plane has two connected components and the orbit $\varphi^t(x)$ must stay inside one of them for $t > t_2$, namely the connected component reached by following the flow a little bit at $\varphi^{t_2}(x)$. In particular, $\varphi^{t_2}(x)$ must lie between $\varphi^{t_1}(x)$ and $\varphi^{t_3}(x)$ on I. $\square$

We are finally ready to proof Theorem 4.3.

Proof of Theorem 4.3. Assume we have $x \in \mathbb{R}^2$ with bounded positive semi-orbit. Let $y \in L^{\varphi}_{\omega}(x)$, and let $z \in L^{\varphi}_{\omega}(y) \cup L^{\varphi}_{\alpha}(y)$. (As an exercise, prove that such a z must exist).

We assume that $F(z) \neq 0$. We want to show that in this case y is periodic with orbit γ, and $L^{\varphi}_{\omega}(x) = \gamma$.

Let I_z be a transversal through z, and let V be a product neighbourhood of I_z such that

$$\varphi^t : (-\varepsilon, \varepsilon) \times I_z \to V$$

is an embedding.

Claim 1: $\varphi^{t>0}(y) \cap I_z = \{z\}$.

In order to prove this claim, we suppose there exist $y_1 \neq y_2 \in I_z$ and there exist sequences $s_i \to \infty$ and $t_i \to \infty$ such that $\varphi^{s_i}(x) \to y_1$ and $\varphi^{t_i}(x) \to y_2$. By changing s_i and t_i by at most ε, we may and do assume that $\varphi^{s_i}(x)$ and $\varphi^{t_i}(x)$ are in I_z.

For i, j large enough, we now have that $d(\varphi^{s_i}(x), y_1), d(\varphi^{t_j}(x), y_2) < \eta$, where $\eta = \frac{1}{2}d(y_1, y_2)$. In particular, we can find real numbers $a < b < c$ such that $\varphi^a(x), \varphi^b(x), \varphi^c(x) \in I_z$ contradict Lemma 4.13. This means that the positive semi-orbit $\varphi^{t>0}(x)$ accumulates on at most one point in I_z, which must be the point z. (Recall that the point z is an accumulation point for $\varphi^{t>0}(x)$ by Lemma 4.12).

Now, since $\varphi^{t>0}(y) \subset L^{\varphi}_{\omega}(x)$, this implies that $\varphi^{t>0}(y) \cap I_z = \{z\}$ and proves the claim.

Claim 2: z is periodic. To prove this claim, take a sequence $t_i \to \infty$ such that $\varphi^{t_i}(x) \to z$. Without loss of generality, we may assume that $\varphi^{t_i}(x) \in I_z$. Every point $\varphi^{t_i}(x)$ comes back to the product neighbourhood V within time $(T - \varepsilon, T + \varepsilon)$ for some time $T \in \mathbb{R}$. By continuity, there exists some time T_z such that $\varphi^{T_z}(z) \in I_z$. Furthermore, the point $\varphi^{T_z}(z)$ is an accumulation point of $\varphi^{t>0}(x)$. Since the latter orbit only has one accumulation point on I_z, namely, z, this implies $\varphi^{T_z}(z) = z$. So z is periodic, which proves the claim.

Note that since z is periodic, also y must be periodic by Claim 1. We denote by T_y its minimal period. The third and last claim will finish the proof.

Claim 3: $\varphi^{[0,T_y]}(y) = L^{\varphi}_{\omega}(x)$.

We first note that the inclusion $\varphi^{[0,T_y]}(y) \subset L_\omega^\varphi(x)$ is given by Lemma 4.12. For the other inclusion, let I_y be a transversal through y. There exists a sequence $t_i \to \infty$ such that $\varphi^{t_i}(x) \to y$ and $\varphi^{t_i}(x) \in I_y$.

Let A_i be the annulus bounded by $\varphi^{[0,T_y]}(y)$ and $\varphi^{[t_i,T_i]} \cup$ some part of I_y, where T_i is the first-return time for $\varphi^{t_i}(x)$ to come back to the transversal I_y. Since $\varphi^{t_i}(x)$ is monotone on the transversal I_y by Lemma 4.13, we have the infinite chain of inclusions

$$\cdots \subset A_{i+1} \subset A_i \subset \cdots .$$

For i_0 large enough, the annulus A_{i_0} contains no zero of the vector field F. By continuity of the flow and compactness of A_{i_0}, there exists some $T > 0$ such that for all $p \in A_{i_0}$, the forward orbit of p intersects I_y before time T. In particular, for all $j \geq i_0$ and for all $p \in A_j$ there exists a time T_p such that $\varphi^{T_p}(p) \in A_{j+1}$. This implies that

$$L_\omega^\varphi(x) \subset \bigcap_{i \geq i_0} A_i = \varphi^{[0,T_y]}(y),$$

which yields $\varphi^{[0,T_y]}(y) = L_\omega^\varphi(x)$ and thus finishes the proof. $\square$

Exercise 4.14. Let $F : \mathbb{R}^2 \to \mathbb{R}^2$ be a continuously differentiable vector field in the plane. Prove that a simply connected region of the plane that is closed and contains a positive semi-orbit also contains a fixed point of the flow.

Give a counterexample after replacing the plane $\mathbb{R}^2$ by the space $\mathbb{R}^3$.

Exercise 4.15. Construct a continuously differentiable vector field on the torus Σ_1 that contains exactly one periodic orbit γ, and such that for every point $x \in \Sigma_1$, the ω- and α-limit equals γ.

Is such a construction possible in the plane $\mathbb{R}^2$?

Chapter 5

Homeomorphisms of the circle

The goal of this chapter is to understand conjugacy classes of homeomorphisms and diffeomorphisms of the circle S^1.

We first set up some notation and recall some basics from topology. We let $S^1 := \mathbb{R}/\mathbb{Z}$, and we let $\pi : \mathbb{R} \to S^1$ be the natural projection (in terms of covering theory, this is the universal covering map). For each continuous map $f : S^1 \to S^1$ and each $x \in \pi^{-1}(f(0))$, there exists a unique $F : \mathbb{R} \to \mathbb{R}$ with $\pi \circ F = f \circ \pi$, such that $F(0) = x$. Such a map F is called a *lift* of f.

Definition 5.1. *We let* $\text{Homeo}^+(S^1)$ *be the group of orientation-preserving homeomorphisms* $f : S^1 \to S^1$.

Note that f preserving the orientation of S^1 is equivalent to any lift F of f being an increasing function. We will use the following lemma from covering theory using the fact that any two lifts F must differ by an orientation-preserving deck transformation of the universal covering $\pi : \mathbb{R} \to S^1$. These deck transformations are given by integer translations.

Lemma 5.2. *Let* $f \in \text{Homeo}^+(S^1)$*, and let* $F_1, F_2 : \mathbb{R} \to \mathbb{R}$ *be two lifts of* f*. Then there exists* $k \in \mathbb{Z}$ *such that* $F_1(x) = F_2(x) + k$ *for all* $x \in \mathbb{R}$.

In particular, there must exist a unique lift F of f such that $F(0) \in [0, 1)$. We call this lift the *canonical lift* of f.

Remark 5.3. A lift F of a homeomorphism $f : S^1 \to S^1$ commutes with integer translations: $F(x + k) = F(x) + k$.

Finally, we define the set of all homeomoprhisms of $\mathbb{R}$ that are obtained by lifting orientation-preserving homeomorphisms of the circle.

Definition 5.4. *We define* $\widetilde{\text{Homeo}}^+(S^1)$ *to be the group of homeomorphisms* $F : \mathbb{R} \to \mathbb{R}$ *such that* $F(x + 1) = F(x) + 1$ *for all* $x \in \mathbb{R}$.

5.1 The rotation number

We now look at the slope of the lines obtained by connecting x with $F^n(x)$. It turns out that these slopes stabilise for $n \to \infty$. Using this, we will obtain a conjugacy invariant for homeomorphisms of the circle.

Proposition 5.5. *Let $F \in \widetilde{\text{Homeo}}^+(S^1)$. Then for all $x \in \mathbb{R}$, the limit*

$$\lim_{n\to\infty} \frac{F^n(x)}{n}$$

exists, and it does not depend on $x \in \mathbb{R}$.

This limit for a lift F of f does, however, depend on choice of lift. However, the following exercise shows that the residue modulo 1 is independent of the choice of lift.

Exercise 5.6. Let F_1 and F_2 be two lifts of $f \in \text{Homeo}^+(S^1)$ that differ by k, that is, $F_1(x) = F_2(x) + k$. Show that also the corresponding limits differ by k, that is,

$$\lim_{n\to\infty} \frac{F_1^n(x)}{n} = \lim_{n\to\infty} \frac{F_2^n(x)}{n} + k.$$

Proof of Proposition 5.5. Let $F \in \widetilde{\text{Homeo}}^+(S^1)$. We define the auxiliary function $u_n(x) = F^n(x) - x$.

Caim: for all $x, y \in \mathbb{R} : |u^n(x) - u^n(y)| < 1$.

Proof of Claim: There exists $p \in \mathbb{Z}$ such that $y \in [x + p, x + p + 1)$. Since F commutes with integer translations, so does F^n. In particular, we obtain $F^n(x + p) = F^n(x) + p$. This gives us

$$F^n(y) \in [F^n(x + p), F^n(x + p + 1)) = [F^n(x) + p, F^n(x) + p + 1),$$

and hence

$$|u^n(x) - u^n(y)| = |F^n(x) - F^n(y) + y - x| < 1,$$

which proves the claim.

We note that the claim implies that if the limit exists, then it does not depend on $x \in \mathbb{R}$. Indeed, we have

$$|F^n(x) - F^n(y)| = |u_n(x) - u_n(y) + x - y| < 1 + |x - y|,$$

a difference which converges to zero when dividing by n, as $n \to \infty$.

It remains to show that the limit exists. For this, we calculate

$$\begin{aligned} u_{n+m}(x) - u_n(x) - u_m(x) = \\ &= (F^{nm}(x) - F^n(x) + F^n(x) - x) - (F^n(x) - x) - (F^m(x) - x) \\ &= (F^{nm}(x) - F^n(x)) - (F^m(x) - x) \\ &= u_m(F^n(x)) - u_m(x) < 1. \end{aligned}$$

This implies that the sequence $\{u_n(x) + 1\}_{n\in\mathbb{N}}$ is subadditive, that is

$$u_{n+m}(x) + 1 \leq u_n(x) + 1 + u_m(x) + 1.$$

Fekete's lemma states that for every subadditive sequence r_n in the real numbers, the sequence $\frac{r_n}{n}$ has a limit in $\mathbb{R} \cup \{-\infty\}$, as $n \to \infty$. The proof of Fekete's lemma is an exercise in Analysis. In particular, we obtain that the sequence $\frac{u_n(x)+1}{n}$ has a limit in $\mathbb{R} \cup \{-\infty\}$, as $n \to \infty$.

Similarly, we have that $\{u_n(x) - 1\}_{n\in\mathbb{N}}$ is a superadditive sequence and therefore the sequence $\frac{u_n(x)-1}{n}$ has a limit in $\mathbb{R} \cup \{+\infty\}$, as $n \to \infty$.

This implies that the sequence $\frac{u_n(x)}{n}$ has a limit in $\mathbb{R}$ as $n \to \infty$, and hence the sequence $\frac{F^n(x)}{n} = \frac{u_n(x)+x}{n}$ has a limit in $\mathbb{R}$ as $n \to \infty$. □

Definition 5.7. *For $F \in \widetilde{\text{Homeo}}^+(S^1)$, the translation number $\tau(F)$ is defined to be $\lim_{n\to\infty} \frac{F^n}{n}$.*

Definition 5.8. *For $f \in \text{Homeo}^+(S^1)$, the rotation number $\rho(f)$ is defined to be $\tau(F)$ mod 1 for any lift F of f.*

The translation number and the rotation number are well-defined thanks to Proposition 5.5 and Exercise 5.6.

Example 5.9. We can calculate the following translation and rotation numbers.

1. Let $F : \mathbb{R} \to \mathbb{R}, x \mapsto x + \alpha$ for $\alpha \in \mathbb{R}$. Then $\tau(F) = \alpha$.
2. Let $r_\alpha : S^1 \to S^1, x \mapsto x + \alpha \bmod 1$ for $\alpha \in \mathbb{R}$. Then $\rho(f) = \alpha \bmod 1$.
3. If $f : S^1 \to S^1$ has a fixed point, then $\rho(f) = 0$.

For the rotation number, we often omit writing mod 1 and simply give a representative $\in [0, 1)$.

Exercise 5.10. Show the equality $\rho(f^p) = p \cdot \rho(f)$. Use this equality to prove that if f has a periodic point of period q, then $\rho(f) = \frac{p}{q}$ for some $p \in \mathbb{Z}$.

Exercise 5.11. Show that the rotation number is not additive in general. More precisely, find two homeomorphisms f_1 and f_2 of S^1 such that

$$\rho(f_1 \circ f_2) \neq \rho(f_1) + \rho(f_2).$$

Proposition 5.12. *If $f, g \in \mathrm{Homeo}^+(S^1)$ are semiconjugate, then they have the same rotation number: $\rho(f) = \rho(g)$.*

Proof. Suppose there exist a continuous and surjective map $h : S^1 \to S^1$ such that $f \circ h = h \circ g$. Choose lifts $F, G, H : \mathbb{R} \to \mathbb{R}$. Replacing F by an integer translate, that is $F(x)$ by $F(x) + k$, we can assume that

$$F \circ H = H \circ G.$$

By induction, we extend this equality to powers of F and G, yielding

$$F^n \circ H = H \circ G^n.$$

In particular, we have

$$\tau(F) = \lim_{n\to\infty} \frac{F^n(H(x))}{n} = \lim_{n\to\infty} \frac{H(G^n(x))}{n}.$$

We note that $H(y) - y$ is bounded for $y \in \mathbb{R}$, so

$$\lim_{n\to\infty} \frac{H(G^n(x))}{n} = \lim_{n\to\infty} \frac{G^n(x)}{n} = \tau(G).$$

This finally yields $\rho(f) = \rho(g)$. □

Corollary 5.13. *A rotation r_α is conjugate to a rotation r_β if and only if $\alpha = \beta \bmod 1$.*

Exercise 5.14. Find $f, g \in \mathrm{Homeo}^+(S^1)$ with $\rho(f) = \rho(g)$ but f and g not conjugate.

We will divide the circle homeomorphisms into two classes: those with rational rotation number and those with irrational rotation number, and study them separately. We start with rational rotation numbers.

5.2 Rational rotation numbers

In the whole section, we assume $f \in \mathrm{Homeo}^+(S^1)$.

Lemma 5.15. *If $\rho(f) = 0$, then f has a fixed point.*

Proof. Suppose f has no fixed point. We can write $f(x) = x + g(x) \bmod 1$ where g is continuous and $g(x) \in (0, 1)$ for all $x \in S^1$. Compactness of the unit circle implies that $g(x)$ has a minimum and a maximum, that is, there exists some large $N \in \mathbb{N}$ such that

$$\frac{1}{N} \leq g(x) \leq 1 - \frac{1}{N}$$

for all $x \in S^1$. Take the canonical lift $F : \mathbb{R} \to \mathbb{R}$ of f. It is defined by $F(x) = x + g(\pi(x))$ for $x \in \mathbb{R}$. Taking $x = 0$, this implies

$$x + 1 \leq F^N(x) \leq x + (N - 1).$$

Taking limits, we obtain $\tau(F) \in (\frac{1}{N}, 1 - \frac{1}{N})$ and hence $\rho(f) \neq 0$, a contradiction. □

Lemma 5.16. *If $\rho(f) = \frac{p}{q}$ with $\gcd(p, q) = 1$, then f has a periodic point of period q. Furthermore, for every lift x of a periodic point, and for the lift $F : \mathbb{R} \to \mathbb{R}$ of f with $\tau(F) = \rho(f)$, we have $F^q(x) = x + p$.*

Proof. By Exercise 5.10, we have $\rho(f^q) = q\rho(f) = 0 \bmod 1$. By the previous lemma, f^q has a fixed point, implying that f has a point of period q.

Now let $x \in \mathbb{R}$ be a lift of a periodic point for f. There exist natural numbers $n, m \in \mathbb{N}$ such that $F^n(x) = x + m$. This gives

$$\lim_{k\to\infty} \frac{F^{nk}(x)}{nk} = \lim_{k\to\infty} \frac{(F^n)^k(x)}{kn} = \lim_{k\to\infty} \frac{x + km}{kn} = \frac{m}{n}.$$

This implies that $\frac{m}{n} = \frac{p}{q}$ or, in other words, there exists some $l > 0$ such that $(m, n) = (lp, lq)$. Assume now that $F^q(x) < x + p$. This implies that $F^{2q}(x) < x + 2p$, and so on: we obtain $F^{lq}(x) < x + lp$, which contradicts $F^n(x) = x + m$. We obtain a similar contradiction if we assume $F^q(x) > x + p$. Hence, we must have $F^q(x) = x + p$. □

We note that the previous proof in particular shows that if $\rho(f) = \frac{p}{q}$ with $\gcd(p, q) = 1$, then every periodic point for f has minimal period q. We summarise and finalise our study by the following result.

Theorem 5.17. *If $\rho(f) = \frac{p}{q}$ with $\gcd(p, q) = 1$, then f has a periodic point and all its periodic points are of minimal period q. Furthermore, the order of the points $(x, f(x), f^2(x), \ldots, f^{q-1}(x))$ on S^1 is the same as the order of $(0, \frac{p}{q}, \frac{2p}{q}, \ldots, \frac{(q-1)p}{q})$.*

We note that the points $(0, \frac{p}{q}, \frac{2p}{q}, \ldots, \frac{(q-1)p}{q})$ equal the orbit of 0 under the rational rotation r_α with angle $\alpha = \frac{p}{q}$.

Proof. Let $i \in \mathbb{N}$ be the minimal natural number such that $(x, f^i(x))$ contains no point of the orbit of x under f. If we define $I := [x, f^i(x))$, we obtain that S^1 is the disjoint union of the intervals $I, f(I), \ldots, f^{q-1}(I)$. In particular, the order of the images of x on S^1 is the following:

$$x < f^i(x) < f^{2i}(x) < \cdots < f^{(q-1)i}(x) < x.$$

We claim that we have $ip = 1 \bmod q$. Indeed, consider a lift $\tilde{x} \in \mathbb{R}$ of the periodic point x. Let $\bar{F}$ be the lift of f^i such that $\bar{F}^q(\tilde{x}) = \tilde{x} + 1$, and let F be the lift of f such that $F^q(\tilde{x}) = \tilde{x} + p$.

There exists a number $k \in \mathbb{Z}$ such that $\bar{F} = F^i + k$. In particular, we have both $F^{qi}(\tilde{x}) = \tilde{x} + ip$ and $F^{qi}(\tilde{x}) = (\bar{F} + k)^q(\tilde{x}) = \bar{F}^q(\tilde{x}) + qk = \tilde{x} + 1 + qk$. This implies that $ip = 1 + qk$, proving the claim.

To finish the proof, we note that $ip = 1 \mod q$ implies that the order of $\{0, \frac{p}{q}, \dots\} \mod 1$, which is the order of $\{0, p, 2p, \dots\} \mod q$, is the following: $0 < ip \mod q < 2ip \mod q < \dots$. □

5.3 Irrational rotation numbers

In the whole section, we assume $f \in \mathrm{Homeo}^+(S^1)$. The goal of this section is to show the following classification result for surface homeomorphisms with irrational rotation number.

Theorem 5.18 (Poincaré classification)**.** *If $\rho(f) \notin \mathbb{Q}$, then there exists a map $h : S^1 \to S^1$ that is monotone and continuous and $h \circ f = r_{\rho(f)} \circ h$. Moreover, if f is topologically transitive, then h is invertible. If f is not topologically transitive, then h is not injective.*

We note that Theorem 5.18 shows that a circle homeomorphism f with irrational rotation number is semiconjugate to the rotation with the same (irrational) rotation number. Furthermore, if f has a dense orbit, then the semiconjugacy is in fact a conjugacy.

5.3.1 Denjoy examples

The first thing we do is to verify that there exist f with irrational rotation number that are not topologically transitive. This shows that the second part of Theorem 5.18 is indeed necessary, as such a map cannot be conjugate to an irrational rotation. Such f go by the name of *Denjoy examples.*

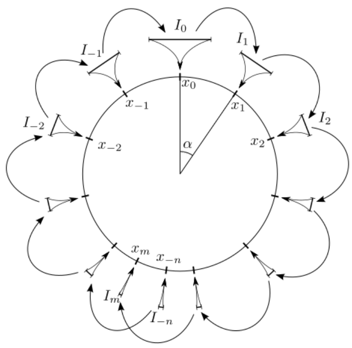

Figure 5.1: An illustration of the Denjoy example.

Example 5.19 (Denjoy examples)**.** Let $(l_n)_{n\in\mathbb{Z}}$ be real numbers $l_n > 0$ such that $\sum_{n\in\mathbb{Z}} l_n = 1$. We start with the irrational rotation r_α but modify it in the following way. Choose $x_0 \in S^1$. For each $n \in \mathbb{Z}$, insert an interval I_n of length l_n at $x_n = f^n(x_0)$ and choose any homeomorphism $h_n : I_n \to I_{n+1}$. Define $f_\alpha(x)$ to be $h_n(x)$ on I_n, for each $n \in \mathbb{Z}$, and $r_\alpha(x)$ on all other points, see Figure 5.1. We note that since we leave many orbits of r_α untouched, the rotation number of f_α must be the same as the rotation number r_α, which equals α. Furthermore, no orbit of f_α is dense, as every orbit visits each interval I_n at most once.

Exercise 5.20. Read in [2] how continuously differentiable Denjoy examples can be constructed. We will show later that no such examples exist that are twice continuously differentiable.

5.3.2 Proof of the Poincaré classification

As always, let $f \in \text{Homeo}^+(S^1)$. Before we start with the proof of the Poincaré classification, we ponder the density of orbits.

Lemma 5.21. *If $\rho(f) \notin \mathbb{Q}$, then for all $x \in S^1$ and for all integers $n \neq m$, every forward orbit of x under f intersects the interval $I := [f^m(x), f^n(x)]$.*

Proof. We assume that $m > n$, the other case being treated similarly. We want to show that

$$S^1 = \bigcup_{k=1}^{\infty} f^{-k}(I).$$

If this was not the case, then we would have

$$S^1 \not\subset \bigcup_{k=1}^{\infty} f^{-k(m-n)}(I) = \bigcup_{k=1}^{\infty} [f^{-(k-1)m+kn}(x), f^{-km+(k+1)n}(x)].$$

The endpoints of the intervals $f^{-k(m-n)}$ match, which implies that the sequence of endpoints $f^{-k(m-n)}(f^n(x))$ converges monotonically to a $p \in S^1$, as $k \to \infty$. This point p must therefore be a fixed point of f^{m-n} and hence a periodic point for f, contradicting $\rho(f) \notin \mathbb{Q}$. □

Proposition 5.22. *If $\rho(f) \notin \mathbb{Q}$, then for all $x, y \in S^1$: $L^f_\omega(x) = L^f_\omega(y)$.*

Proof. Suppose $x_0 \in L^f_\omega(x)$. There exists a sequence a_n with $f^{a_n}(x) \to x_0$ for $n \to \infty$. By Lemma 5.21, there exists a sequence b_n such that

$$f^{b_n}(y) \in [f^{a_n}(x), f^{a_{n+1}}(x)].$$

But this implies that $f^{b_n}(y) \to x_0$, as $n \to \infty$. This proves the inclusion $L^f_\omega(x) \subset L^f_\omega(y)$, and the reverse inclusion is obtained simply by exchanging x and y. □

This implies the following result.

Corollary 5.23. *If one orbit of f with $\rho(f) \notin \mathbb{Q}$ is dense, then every orbit of f is dense.*

The following Lemma is a key ingredient to the proof of the Poincaré classification.

Lemma 5.24. *Let $F \in \widetilde{\mathrm{Homeo}}^+(S^1)$ with $\tau(F) \notin \mathbb{Q}$. Then for all $x \in \mathbb{R}$:*

$$F^{n_1}(x) + m_1 < F^{n_2}(x) + m_2 \iff n_1\tau(F) + m_1 < n_2\tau(F) + m_2.$$

Proof. We first note that there exists no $x \in \mathbb{R}$ such that

$$F^{n_1}(x) + m_1 = F^{n_2}(x) + m_2$$

for $n_1 \neq n_2$. Indeed, we would have

$$F^{n_1-n_2}(x) = F^{-n_2}(F^{n_2}(x) + m_2 - m_1) = x + m_2 - m_1$$

and in particular, $\pi(x)$ is a periodic point for f, a contradiction.
We have shown that either $F^{n_1}(x) + m_1 < F^{n_2}(x) + m_2$ for all $x \in \mathbb{R}$ or $F^{n_1}(x) + m_1 > F^{n_2}(x) + m_2$ for all $x \in \mathbb{R}$.
We now prove " $\implies$ ". Note that $F^{n_1}(x)+m_1 < F^{n_2}(x)+m_2$ is equivalent to $F^{n_1-n_2}(x) < x + m_2 - m_1$. By the above remark, it suffices to consider the inequality for $x = 0$, in which case it reads $F^{n_1-n_2}(0) < m_2 - m_1$. By an inductive argument, we get $F^{k(n_1-n_2)}(0) < k(m_2 - m_1)$.
This implies $\tau(F) \leq \frac{m_2-m_1}{n_1-n_2}$. The inequality is strict since $\tau(F) \notin \mathbb{Q}$, which proves the direction " $\implies$ ".
The reverse implication is obtained by logical contraposition, repeating the same proof with switched inequality signs. □

We are now ready to prove the Poincaré classification.

Proof of Theorem 5.18. Our goal is to construct a semiconjugacy $H : \mathbb{R} \to \mathbb{R}$ for the lifts that descends to a semiconjugacy $h : S^1 \to S^1$.
Let F be a lift of f. We first construct H on one orbit. Let $x \in \mathbb{R}$ and define $B := \{F^n(x) + m : n, m \in \mathbb{Z}\} \subset \mathbb{R}$. We now define H on B as follows:

$$H : B \to \mathbb{R}, \quad F^n(x) + m \mapsto n\tau(F) + m$$

By Lemma 5.24, H is monotone. Furthermore, we have

$$H \circ F = T_{\tau(F)} \circ H.$$

We note that $H(B)$ is dense in $\mathbb{R}$.

We now prove that H extends continuously to $H : \overline{B} \to \mathbb{R}$. Let $y \in \overline{B}$. Then at least one of the two limits

$$\lim_{x_n \nearrow y,\ x_n \in B} H(x_n)$$

and

$$\lim_{x_n \searrow y,\ x_n \in B} H(x_n)$$

exists, since H is monotone. If both limits exist, then they must agree. Indeed, if not, then $R \setminus H(B)$ contains an open interval, which contradicts the density of $H(B)$ in $\mathbb{R}$. We have therefore have $H : \overline{B} \to \mathbb{R}$ continuous and monotone.

The compement $R \setminus \overline{B}$ is an open set; it is the union of open intervals. On such an open interval I, the limit values $H(I^-)$ and $H(I^+)$ on the two endpoints $\in \overline{B}$ must agree. Otherwise we would get that $H(\overline{B})$ is not dense in $\mathbb{R}$. So for each such interval I, we define

$$H|_I = H(I^+) = H(I^-) = \text{const.}$$

This finally yields a continuous and monotone function $H : \mathbb{R} \to \mathbb{R}$ with

$$H \circ F = T_{\tau(F)} \circ H.$$

Furthermore, the function H is invertible if and only if $\overline{B} = \mathbb{R}$, which is the case if and only if the orbit of every point is dense under $f : S^1 \to S^1$.

The induced map $h : S^1 \to S^1$ is a semiconjugacy satisfying all propoerties of the theorem. □

5.4 Arnold tongues

In this section, we present a series of exercises exploring how the rotation number $\rho(f)$ depends on the homeomorphism f. For inspiration on the exercises, you can consult the second half of the Section 11.1 in the book of Katok and Hasselblatt [3].

Exercise 5.25. Show that the rotation number $\rho(f)$ depends continuously on the function f, in the C^0-topology.

Exercise 5.26. Assume that $\rho(f) = \frac{p}{q}$. Show that if for a periodic point x of f, we have $(f^q)'(x) \neq 1$, then all sufficiently close perturbations of f have rotation number $\frac{p}{q}$. Give a counterexample for the statement if all periodic points x satisfy $(f^q)'(x) = 1$.

Exercise 5.27. Let $F_1 \in \widetilde{\text{Homeo}}^+(S^1)$ such that $\tau(F) \notin \mathbb{Q}$. Furthermore, let $F_2 \in \widetilde{\text{Homeo}}^+(S^1)$ be a sufficiently small perturbation of F_1 such that $F_2(x) > F_1(x)$ for all $x \in \mathbb{R}$. Show that $\tau(F_2) > \tau(F_1)$.

Exercise 5.28 (Arnold tongues)**.** For $(a, b) \in [0, 1] \times [0, \frac{1}{2\pi}]$, define the map $f_{a,b} : S^1 \to S^1$ by

$$f_{a,b}(x) = x + a + b \sin(2\pi x) \bmod 1.$$

a) Show that $f_{a,b}$ is a homeomorphism of S^1.

b) For $p/q \in \mathbb{Q} \cap [0, 1]$, define the *Arnold tongue* $A_{p/q}$ to be the set

$$A_{p/q} = \{(a, b) \in [0, 1] \times [0, \frac{1}{2\pi}] : \rho(f_{a,b}) = p/q\}.$$

Show that each tongue $A_{p/q}$ is closed and intersects the line $b = 0$ exactly at the point $a = p/q$.

c) Show that each tongue $A_{p/q}$ intersects every line $b =$ const. in a closed interval of positive length, except for $b = 0$.

d) Show that the union of the tongues $A_{p/q}$ is dense in $[0, 1] \times [0, \frac{1}{2\pi}]$.

e) Show that for $b \in [0, \frac{1}{2\pi})$ fixed, the function $a \mapsto \rho(f_{a,b})$ is a devil's staircase.

5.5 Diffeomorphisms of the circle

Definition 5.29. *For a map $f : S^1 \to \mathbb{R}$ we define its variation to be*

$$\mathrm{Var}(f) := \sup_{0 \le x_1 \le \cdots \le x_n \le 1} \sum_{k=1}^{n} |f(x_k) - f(x_{k+1})|,$$

where the supremum is taken over all $n \in \mathbb{N}$. We say such a map f has bounded variation if $\mathrm{Var}(f) < +\infty$.

We will prove the following theorem due to Denjoy.

Theorem 5.30 (Denjoy)**.** *Let $f \in \mathrm{Homeo}^+(S^1)$ with $\rho(f) \notin \mathbb{Q}$. If f' has bounded variation, then f is topologically transitive.*

We note that every C^1 map $S^1 \to \mathbb{R}$ is Lipschitz and hence has bounded variation. As a consequence of Denjoy's theorem and the Poincaré classification of surface homeomorphisms with irrational rotation number, we get that Denjoy examples are not possible in the class of twice continuously differentiable homeomorphisms of S^1.

Corollary 5.31. *If $f \in \mathrm{Homeo}^+(S^1)$ with $\rho(f) \notin \mathbb{Q}$ is of class C^2, then f is conjugate to the rotation $r_{\rho(f)}$.*

Proof of Theorem 5.30. We assume f is not topologically transitive and aim to arrive at a contradiction. In this case, the ω-limit $L_\omega^f(0)$ is a closed and nowhere dense subset of the unit circle. In particular, $S^1 \setminus L_\omega^f(0)$ is a union of open intervals. Let $I = (a, b)$ be one of them.

We note that the intervals $f^n(I)$, $n \in \mathbb{Z}$, must be pairwise disjoint, as otherwise we could produce a periodic point of the map f, which would contradict $\rho(f) \notin \mathbb{Q}$. This implies

$$\sum_{n\in\mathbb{Z}} l(f^n(I)) \leq 1,$$

where $l(f^n(I)) = \int_a^b (f^n)'(t)dt$ is the length of the interval $f^n(I)$.

Before stating a lemma we will need in the proof, we note that for an orientation-preserving diffeomorphism $f : S^1 \to S^1$, we must have $f'(x) > 0$ for all $x \in S^1$.

Lemma 5.32. *Let $J \subset S^1$ be such that the interiors of $J, f(J), \ldots, f^{n-1}(J)$ are pairwise disjoint. Let $g = \log(f')$, and fix $x, y \in J$. Then for all $n \in \mathbb{Z}$, we have*

$$\mathrm{Var}(g) \geq |\log((f^n)'(x)) - \log((f^n)'(y))|.$$

Proof. Using that all interiors of the intervals $J, f(J), \ldots, f^{n-1}(J)$ are pairwise disjoint, we get

$$\begin{aligned}\mathrm{Var}(g) &\geq \sum_{k=0}^{n-1} \left|g(f^k(y)) - g(f^k(x))\right| \geq \left|\sum_{k=0}^{n-1} g(f^k(y)) - g(f^k(x))\right| = \\ &= \left|\log \prod_{k=0}^{n-1} f'(f^k(y)) - \log \prod_{k=0}^{n-1} f'(f^k(x))\right| \\ &= \left|\log((f^n)'(y)) - \log((f^n)'(x))\right|,\end{aligned}$$

which is what we wanted to show. □

Take $x \in S^1$ and n such that $f^k(x) \notin [x, f^n(x)]$ for $0 < |k| < n$. Such an n, and in fact infinitely many of them, exists. Indeed, the irrational rotation number of f implies that the orbit of any point x is ordered in the same way as the orbit of x under the irrational rotation $r_{\rho(f)}$, by Lemma 5.24, and the orbit of x under the irrational rotation is dense.

Now let $J := f^{-n}[x, f^n(x)] = [f^{-n}(x), x]$. Then $f^k(J) = [f^{-n+k}, f^k(x)]$ and the intervals $J, f(J), \ldots f^{n-1}(J)$ are pairwise disjoint. We now apply the lemma to the interval J with $y = f^{-n}(x)$. This yields

$$\begin{aligned}\mathrm{Var}(g) &\geq \left|\log\left((f^n)'(x)\right) - \log((f^n)'(y))\right| \\ &= \left|\log\left(\frac{(f^n)'(x)}{(f^n)'(f^{-n}(x))}\right)\right| \\ &= \left|\log\left((f^n)'(x)(f^{-n})'(x)\right)\right|.\end{aligned}$$

So for infinitely many $n \in \mathbb{N}$, we obtain

$$
\begin{aligned}
l(f^n(I)) + l(f^{-n}(I)) &= \int_I (f^n)'(x)dx + \int_I (f^{-n})'(x)dx \\
&= \int_I (f^n)'(x) + (f^{-n})'(x)dx \\
&\geq 2 \int_I \sqrt{(f^n)'(x)(f^{-n})'(x)}dx \\
&\geq 2 \int_I \sqrt{\exp(-\mathrm{Var}(g))}dx \\
&= 2 \cdot l(I) \exp\left(-\frac{\mathrm{Var}(g)}{2}\right) = \text{const.} > 0,
\end{aligned}
$$

since we assume that the variation of f' is bounded and hence so is the variation of $g = \log(f')$. This contradicts the fact that $\sum_{n\in\mathbb{Z}} l(f^n(I)) \leq 1$, and finishes the proof. $\square$

Chapter 6

Ergodic theory

Ergodic theory is the study of statistical properties of dynamical systems. The name originates from statistical mechanics, where the "ergodic hypothesis" asserts that observables satisfy an asymptotic equality of their average over time and of their average over the space, in short:

$$\text{time average} = \text{space average}.$$

Before proving such an equality via Birkhoff's ergodic theorem, we consider basic definitions and a first recurrence theorem, the Poincaré recurrence.

6.1 Poincaré recurrence

Let μ be a measure on the space X. We do not recall the basic definitions from measure theory. For a minimal refresher on the notions we encounter, read the first pages of the chapter on ergodic theory in [2].

Definition 6.1. *The triple (X, μ, f) is a measure-preserving discrete-time dynamical system if f is measurable and $f^*\mu = \mu$, that is, for all measurable sets $A \subset X$, we have $\mu(f^{-1}(A)) = \mu(A)$.*

While we will only deal with discrete-time dynamical systems, an analogue definition can be given for continuous-time ones.

Example 6.2. The rotation r_α of the circle S^1 is measure-preserving with respect to the Lebesgue measure $\mathcal{L}^1$.

Exercise 6.3. Show that $(S^1, \mathcal{L}^1, x \mapsto 2x \bmod 1)$ is measure-preserving.

Example 6.4. The rotation $(S^1, r_{\frac{p}{q}})$ admits many invariant measures: take any measure on the interval $[0, \frac{p}{q})$ and transport it to a measure on S^1 by the rotation.

Theorem 6.5 (Poincaré recurrence)**.** *Let (X, μ, f) be a measure-preserving dynamical system with $\mu(X) < \infty$. If $A \subset X$ is measurable, then for μ-almost every $x \in A$, the orbit $\mathcal{O}_f^+(x)$ visits A infinitely many times.*

Proof. Define the measurable sets

$$E_n = \bigcup_{k=n}^{\infty} f^{-k}(A),$$
$$E = \bigcap_{n=0}^{\infty} E_n.$$

Let $A^* = A \cap E$ be the set of points $x \in A$ whose orbit returns to A infinitely many times. That is, the points $x \in A$ for which there exists an infinite sequence

$$0 < k_1 < k_2 < \dots$$

with $f^{k_i}(x) \in A$ for all i. In fact, by definition, we obtain $f^{k_i}(x) \in A^*$ for all i.

We now want to show that $\mu(A^*) = \mu(A)$. For all $m \geq n$ and by the definition of the sets E_i, we have $E_m = f^{n-m}(E_n)$. As f is measure-preserving, this implies that all sets E_n have the same measure: $\mu(E_n) = \mu(E_0)$ for all $n \geq 0$. Since $E_n \subset E_{n-1}$ for all n, we obtain that $\mu(E) = \mu(E_0)$ as well. This finally implies

$$\mu(A^*) = \mu(A \cap E) = \mu(A \cap E_0) = \mu(A),$$

where the last inequality follows from the fact that $A \subset E_0$. □

Exercise 6.6. Find a counterexample to the Poincaré recurrence theorem with $\mu(X) = \infty$.

Example 6.7 (Zermelo's paradox)**.** Assume there is a box with a wall in its middle, and the half to the left of the box is filled with some perfect gas, while the right half is in a vacuum. Then remove the wall so the gas can move freely in the box. Will all the gas once again be contained in the left half of the box?

Let X is the space of all positions and velocities of all particles, and let A is the subspace of points that position all gas particles on the left half of the box. There is an invariant measure μ for this dynamical system, the Maxwell-Boltzmann measure. So by the Poincaré recurrence theorem, after some time, all particles of the gas will actually be back in the left half of the box.

This is called a paradox because is seemingly contradicts the second law of thermodynamics, which asserts that natural processes run in one sense and are not reversible.

Exercise 6.8. Discuss a solution of Zermelo's paradox.

6.2 Birkhoff's Ergodic Theorem

Given a measure-preserving dynamical system (X, μ, f), where does the orbit $\mathcal{O}_f^+(x)$ dwell? More precisely, for any subset $A \subset X$, how much time does the orbit spend in A? More formally, we look at the proportion of time spent in A after n steps,

$$\frac{1}{n}\left(\chi_A + \chi_A \circ f + \cdots + \chi_A \circ f^{n-1}\right)(x),$$

and wonder what happens if we let $n \to \infty$. If this limit exists, it tells us the asymptotic proportion of time the orbits spends in A, the "time average".

Example 6.9. For $(S^1, \mathcal{L}^1, x \mapsto 2x \bmod 1)$ and $A = [0, \frac{1}{2})$, the limit equals the proportion of 0s in the binary expansion of $x \in [0, 1)$.

Example 6.10. Consider the circle rotation $(S^1, \mathcal{L}^1, r_\alpha)$ together with the subset $A = [0, \lambda) \subset S^1$.
If $\alpha \notin \mathbb{Q}$, then we might guess that the orbit equidistributes in some sense, so that the limit equals λ, which also equals $\mathcal{L}^1(A)$.
If $\alpha = \frac{p}{q} \in \mathbb{Q}$, then the limit equals

$$\frac{\#\{k \in \mathbb{N} : k < q, x + \frac{k}{q} \in [0, \lambda)\}}{q}.$$

In order to state the Birkhoff Ergodic Theorem, we need the following definition generalising the indicator sums we consideret up to now.

Definition 6.11. *Let (X, f) be a dynamical system and let $\varphi : X \to \mathbb{C}$. We define*

$$S_n(\varphi, f) := \frac{1}{n}\sum_{k=0}^{n-1} \varphi \circ f^k,$$

the Birkhoff sums of φ with respect to f.

Theorem 6.12 (Birkhoff Ergodic Theorem, 1931)**.** *Let (X, μ, f) be a measure-preserving dynamical system, and let $\varphi \in L^1(X, \mu)$. Then*

i) the limit $\widetilde{\varphi}(x) := \lim_{n\to\infty} S_n(\varphi, f)(x)$ exists for μ-almost every $x \in X$,

ii) $\widetilde{\varphi}(x) = \widetilde{\varphi} \circ f(x)$ for μ-almost every $x \in X$,

iii) $||\widetilde{\varphi}||_{L^1} \leq ||\varphi||_{L^1}$,

iv) if $\mu(X) < \infty$, then the convergence is in $L^1(X, \mu)$:

$$||S_n(\varphi, f) - \widetilde{\varphi}||_{L^1} \to 0,$$

v) for every f-invariant measurable subset $A \subset X$ with $\mu(A) < \infty$, we have

$$\int_A \varphi d\mu = \int_A \widetilde{\varphi} d\mu.$$

6.3 Ergodic dynamical systems

In the context of measure-preserving dynamical systems, a subset $A \subset X$ is called invariant if $f^{-1}(A) = A$.

Definition 6.13. *Let (X, μ, f) be a measure-preserving dynamical system with $\mu(X) = 1$. The system is called ergodic if for every invariant measurable subset $A \subset X$, either $\mu(A) = 0$ or $\mu(A) = 1$.*

We will deduce from the Birkhoff Ergodic Theorem 6.12 that the time average and the space average agree for ergodic dynamical systems.

Corollary 6.14 (Time average = space average)**.** *If (X, μ, f) is ergodic and $\varphi \in L^1(X, \mu)$, then for μ-almost every $x \in X$, the time average equals the space average:*

$$\widetilde{\varphi}(x) = \int_X \varphi d\mu.$$

In particular, for $\varphi = \chi_A$, we get $S_n(\chi_A, f) \to \mu(A)$.

Example 6.15. The dynamical system $(S^1, \mu_x^q, r_{\frac{p}{q}})$ is ergodic, where

$$\mu_x^q = \frac{1}{q}\left(\delta_x + \delta_{x+\frac{1}{q}} + \cdots + \delta_{x+\frac{q-1}{q}}\right),$$

and δ_y is the Dirac measure with atom y.

Example 6.16. More generally, given (X, f) and a periodic point $x \in X$ of period n, the measure

$$\mu = \frac{1}{n}\left(\delta_x + \delta_{f(x)} + \cdots + \delta_{f^{n-1}(x)}\right)$$

gives an ergodic dynamical system (X, μ, f).

Example 6.17. $(S^1, \mathcal{L}^1, r_\alpha)$ is not ergodic if $\alpha \in \mathbb{Q}$. This can be checked by finding an r_α-invariant set A of measure > 0 but < 1.

The following proposition together with Birkhoff's Ergodic Theorem 6.12 will allow us to prove Corollary 6.14. It characterises ergodicity via functions that are invariant under f μ-almost everywhere.

Proposition 6.18. *Let (X, μ, f) with $\mu(X) = 1$ be a measure-preserving dynamical system. Then the following are equivalent.*

i) (X, μ, f) is ergodic,

ii) for all $\varphi : X \to \mathbb{C}$ measurable with $\varphi \circ f = \varphi$ μ-almost everywhere, it holds that φ is constant μ-almost everywhere, *iii) for all* $\varphi \in L^1(X,\mu)$ *with with* $\varphi \circ f = \varphi$ μ*-almost everywhere, it holds that* φ *is constant* μ*-almost everywhere,*

iv) for all $\varphi \in L^2(X,\mu)$ *with with* $\varphi \circ f = \varphi$ μ*-almost everywhere, it holds that* φ *is constant* μ*-almost everywhere.*

Proof. We first note that *ii)* $\implies$ *iii)* $\implies$ *iv)* follows from the inclusion of spaces of functions.

To prove *iv)* $\implies$ *i)*, we take $A \subset X$ measurable and f-invariant. The indicator function χ_A in in $L^2(X,\mu)$ and $\chi_A \circ f = \chi_A$, so by *iv)*, the indicator function χ_A ist constant except on a set of measure zero. As it can only take the values 0 or 1, we have $\mu(A) = 1$ or $\mu(A) = 0$, proving ergodicity.

For the implication *i)* $\implies$ *ii)*, suppose there exists $\varphi : X \to \mathbb{C}$ measurable such that $\varphi \circ f = \varphi$ μ-almost everywhere, but with φ not constant μ-almost everywhere. Splitting into real and imaginary parts, we may actually assume $\varphi : X \to \mathbb{R}$.

There exists $x \in \mathbb{R}$ such that $\mu(\varphi^{-1}([x,+\infty))) \neq 0$ or 1. Define

$$A' := \varphi^{-1}([x,+\infty)).$$

We have $\mu(A' \setminus f^{-1}(A') \cup f^{-1}(A') \setminus A') = 0$, since any point x in the set does not satisfy $\varphi \circ f(x) = \varphi(x)$.

As f is measure-preserving, we also get

$$\mu(f^{-k}(A') \setminus f^{-k-1}(A') \cup f^{-k-1}(A') \setminus f^{-k}(A')) = 0 \tag{6.1}$$

for all $k \geq 0$. We now define

$$A := \bigcap_{p\in\mathbb{N}} \bigcup_{n\geq p} f^{-n}(A'),$$

which is measurable and f-invariant. Furthermore, by 6.1 we obtain

$$\mu\left(\bigcup_{n\geq p} f^{-n}(A')\right) = \mu(A')$$

for all $p \in \mathbb{N}$, and finally $\mu(A) = \mu(A')$, which does not equal 0 or 1, contradicting ergodicity. □

We are now ready to prove Corollary 6.14.

Proof of Corollary 6.14. By *ii)* of Birkhoff's ergodic theorem 6.12, we have

$$\widetilde{\varphi}(x) = \widetilde{\varphi} \circ f(x)$$

for μ-almost every $x \in X$. By the above proposition, ergodicity implies that $\widetilde{\varphi}$ is constant on a full measure set. Applying $v)$ of Birkhoff's Ergodic Theorem 6.12 to $A = X$ implies

$$\int_X \varphi d\mu = \int_X \widetilde{\varphi} d\mu = \widetilde{\varphi}(x)$$

for μ-almost every point $x \in X$. □

Let us now turn to more examples.

Example 6.19. We want to show that $(S^1, \mathcal{L}^1, r_\alpha)$ is ergodic assuming the rotation parameter is irrational: $\alpha \notin \mathbb{Q}$.
Take $\varphi \in L^2(X, \mu)$ and suppose $\varphi \circ r_\alpha = \varphi$ almost everywhere. Fourier analysis tells us that the Fourier series

$$\sum_{n=-\infty}^{+\infty} c_n(\varphi) e^{2\pi i n x}$$

converges to φ in L^2, where

$$c_n(\varphi) = \int_{S^1} e^{-2\pi i n t} \varphi(t) dt.$$

Since $\varphi \circ r_\alpha = \varphi$ almost everywhere, the Fourier coefficients of $\varphi \circ r_\alpha$ and φ must agree. We calculate

$$\begin{aligned} c_n(\varphi) = c_n(\varphi \circ r_\alpha) &= \int_{S^1} e^{-2\pi i n t} \varphi(t + \alpha) dt \\ &= \int_{S^1} e^{-2\pi i n (t-\alpha)} \varphi(t) dt \\ &= e^{2\pi i n \alpha} c_n(\varphi). \end{aligned}$$

As α is irrational, $e^{2\pi i n \alpha} \neq 1$ for $n \neq 0$. This implies $c_n = 0$ for $n \neq 0$. So the function φ has only one nontrivial Fourier coefficient and so is constant on a full measure set. The function φ was chosen arbitrarily, so Proposition 6.18 implies ergodicity.

Exercise 6.20. Show that $(S^1, \mathcal{L}^1, x \mapsto 2x \bmod 1)$ is ergodic.

Example 6.21. We want to show that $(\mathbb{R}^2/\mathbb{Z}^2, \mathcal{L}^2, T_{\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)})$ is ergodic.
Let $\varphi : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{C}$ measurable with $\varphi \circ T_{\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)} = \varphi$ on a full measure subset. By Fourier analysis, the Fourier series

$$\sum_{m,n=-\infty}^{+\infty} c_{mn} e^{2\pi i (mx + ny)}$$

converges to φ in L^2. Similarly, the Fourier series

$$\sum_{m,n=-\infty}^{+\infty} c_{mn} e^{2\pi i(m(2x+y)+n(x+y))}$$

converges to $\varphi \circ T_{\left(\begin{smallmatrix}2&1\\1&1\end{smallmatrix}\right)}$ in L^2. As $\varphi \circ T_{\left(\begin{smallmatrix}2&1\\1&1\end{smallmatrix}\right)} = \varphi$ on a full measure subset, we obtain that $c_{mn} = c_{(2m+n)(m+n)}$ for all m and n. In particular, if $c_{mn} \neq 0$ for some $(m,n) \neq (0,0)$, then there exist i, j with $|i| + |j|$ arbitrarily large and $c_{ij} = c_{mn}$. But this would mean that the Fourier series diverges, a contradiction. So we have that $c_{mn} = 0$ for all $(m,n) \neq (0,0)$. In particular, φ is constant on a full measure subset of the torus. Since φ was chosen arbitrarily, Proposition 6.18 implies ergodicity.

Exercise 6.22. Show that $(\mathbb{R}^2/\mathbb{Z}^2, \mathcal{L}^2, T_A)$ is ergodic for any Anosov map T_A.

6.4 Proof of Birkhoff's Ergodic Theorem

We will need the following lemma for our proof.

Lemma 6.23 (Maximal Ergodic Theorem)**.** *Let (Y, ν, g) be a measure-preserving dynamical system, and let $\psi \in L^1(Y, \nu)$ be real-valued. Furthermore, define*

$$\psi^*(y) = \sup_{n \geq 0} S_n(\psi, g)(y).$$

Then we have

$$\int_{\{\psi^*(y)>0\}} \psi d\nu \geq 0.$$

Proof. Define

$$\begin{aligned} \psi_n(y) &:= \sup\{0, \psi(y), \psi(y) + \psi \circ g(y), \ldots, \psi(y) + \cdots + \psi \circ g^{n-1}(y)\}, \\ Y_n &:= \{y \in Y : \psi_n(y) > 0\}. \end{aligned}$$

We directly observe that $\{\psi^*(y) > 0\} = \cup_{n\in\mathbb{N}} Y_n$ and that for all $n \in \mathbb{N}$, we have $Y_n \subset Y_{n+1}$. Furthermore, we have the following equalities:

$$\begin{aligned} \psi_n &= \psi + \psi_{n-1} \circ g && \text{on } Y_n, \\ \psi_n &= 0 && \text{on } Y \setminus Y_n. \end{aligned}$$

We now obtain

$$\begin{aligned} \int_{Y_n} \psi d\nu &= \int_{Y_n} \psi_n d\nu - \int_{Y_n} \psi_{n-1} \circ g d\nu \\ &\geq \int_Y \psi d\nu - \int_Y \psi_{n-1} d\nu \\ &= \int_Y (\psi_n - \psi_{n-1}) d\nu \geq 0, \end{aligned}$$

where we use $\psi_{n-1} \geq 0$ for the inequality in the middle. Finally, this implies

$$\int_{\{\psi^*(y)>0\}} \psi d\nu = \int_{\cup_{n\in\mathbb{N}} Y_n} \psi d\nu \geq 0$$

by the dominated convergence theorem. □

We are now ready to prove Birkhoff's Ergodic Theorem.

Proof of Theorem 6.12. Without loss of generality, we assume $\varphi : X \to \mathbb{R}$. Otherwise, we can split into the real and the imaginary part.
Proof of i): For $a < b \in \mathbb{R}$, we define the sets

$$X(a,b) := \{x \in X : \liminf_{n\to\infty} S_n(\varphi, f)(x) < a < b < \limsup_{n\to\infty} S_n(\varphi, f)(x)\},$$

which are f-invariant. We can therefore apply Lemma 6.23 to the measure-preserving dynamical system $(X(a,b), \mu|_{X(a,b)}, f|_{X(a,b)})$ and to the function $\varphi(x) - b$. We note that $(\varphi - b)^* > 0$ on $X(a,b)$, so Lemma 6.23 implies

$$\int_{X(a,b)} (\varphi(x) - b) d\mu \geq 0.$$

Similarly, we apply Lemma 6.23 to the function $a - \varphi(x)$ and get

$$\int_{X(a,b)} (a - \varphi(x)) d\mu \geq 0.$$

This implies

$$\int_{X(a,b)} (a - b) d\mu \geq 0.$$

Since $a - b < 0$, this can only hold if $\mu(X(a,b)) = 0$. We have shown that there can be no set of positive measure on which the lim sup and the lim inf of the Birkhoff sums disagree.
Proof of ii): If the Birkhoff sums converge, then the limit does not depend on the first term:

$$\begin{aligned}
\lim_{n\to\infty} S_n(\varphi, f)(x) &= \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \varphi \circ f^k(x) \\
&= \lim_{n\to\infty} \frac{1}{n} \left(\varphi(x) + \sum_{k=1}^{n} \varphi \circ f^k(x) \right) \\
&= \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} (\varphi \circ f) \circ f^k(x) \\
&= \lim_{n\to\infty} S_n(\varphi \circ f, f)(x).
\end{aligned}$$

Proof of iii): We assume $\varphi \geq 0$. Otherwise we split φ into the positive and the negative part: $\varphi = \varphi_+ - \varphi_-$, and do the argument separately for φ_+ and φ_-. By Fatou's Lemma, we have

$$\int_X \varphi d\mu = \int_X S_n(\varphi, f) d\mu = \liminf_{n\to\infty} \int_X S_n(\varphi, f) d\mu \geq$$
$$\geq \int_X \liminf_{n\to\infty} S_n(\varphi, f) d\mu = \int_X \widetilde{\varphi} d\mu.$$

Proof of iv): If the function φ is bounded, note that all the terms of the sequence $S_n(\varphi, f)(x)$ are bounded by $||\varphi||_\infty$ and use the dominated convergence theorem. Here, we need that $\mu(X) < \infty$ so that the constant function $||\varphi||_\infty$ is integrable.

If the function φ is not bounded, one can approximate it by bounded ones.

Proof of v): We apply *iv*) to the measure-preserving dynamical system obtained by restriction: $(A, \mu|_A, f|_A)$ and the map $\varphi|_A$. As in *iv*), we again assume that $\varphi|_A \geq 0$.

We obtain

$$\int_A \varphi d\mu = \int_A S_n(\varphi|_A, f|_A) d\mu \longrightarrow \int_A \widetilde{\varphi} d\mu,$$

as $n \to \infty$. Since the left side is constant, this implies

$$\int_A \varphi d\mu = \int_A \widetilde{\varphi} d\mu,$$

and finishes the proof. □

Exercise 6.24. Deduce topological transitivity of Anosov maps of the torus using their ergodicity.

Chapter 7

Surface homeomorphisms

The goal of this chapter is to generalise the Anosov maps of the torus to so-called pseudo-Anosov maps on surfaces of higher genus. Strikingly, Anosov maps are topologically mixing and hence have a dense orbit, yet their periodic points are dense in the torus as well.

Exercise 7.1. Show that the periodic points for an Anosov map

$$T_A : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$$

form a dense subset of the torus.

It is our goal to construct maps of surfaces of higher genus in such a way that they have, except in finitely many points, a similar stretching and contracting behaviour as Anosov maps of the torus, and to see that the properties of topological transitivity and density of periodic points is shared by these maps.

7.1 Basic surface topology

Definition 7.2. *A surface Σ is a two-dimensional topological manifold, possibly with boundary. A surface Σ is closed if it is compact and does not have boundary.*

Example 7.3. Surfaces can be defined in many different ways. Here are some examples.

a) The preimage $f^{-1}(y)$ of a regular value y for $f : \mathbb{R}^n \to \mathbb{R}^{n-2}$ is a surface without boundary, by the implicit function theorem. An example is the 2-sphere, defined as the preimage of 1 of the function $x^2 + y^2 + z^2$.

b) One can parametrise surfaces as surfaces of revolution in $\mathbb{R}^3$: rotating some parametrised one-manifold in the xz-plane around the z-axis.

c) Take a polygon in the plane $\mathbb{R}^2$ and identify its edges pairwise. To make this precise, one has the use the notion of the quotient topology. The most common example of such a kind is the unit square with opposite sides identified: the torus.

Definition 7.4. *Let Σ_0 be the 2-sphere. For $g \geq 1$, let Σ_g be the surface obtained by taking a regular $4g$-gon in the plane and identifying opposite sides via translations. The surface Σ_g is called the surface of genus g.*

Exercise 7.5. Visualise the surface Σ_g and show that it is homeomorphic to a sphere with g handles attached.

Example 7.6. The surface with boundary obtained by taking a unit square and identifying its top and bottom edge via a reflection in the midpoint of the square is called a Moebius strip.

Definition 7.7. *A surface is orientable if it does not contain a Moebius strip as a subsurface.*

The following theorem provides a classification of closed orientable surfaces up to homeomorphisms.

Theorem 7.8 (Classification of surfaces)**.** *Every closed orientable surface Σ is homeomorphic to a surface Σ_g.*

For a slick proof of this classification result, we recommend Putman's note [6] based on a proof of Zeeman.

Exercise 7.9. Describe a closed orientable surface in a way that is complicated enough so that it is not immediately obvious to which of the surfaces Σ_g it is homeomorphic.

By the classification of surfaces, one only has to identify the genus of a closed orientable surface in order to determine its homeomorphism type. A fairly practical way to do this is by triangulations. We give a slightly informal definition.

Definition 7.10. *A triangulation of a surface Σ is a subdivision of Σ into triangles such that each pair of triangles either has exactly one edge in common, or exaclty one vertex in common, or is disjoint.*

Given a closed surface Σ that is triangulated, the Euler characteristic $\chi(\Sigma)$ is equal to $v - e + f$, where v, e and f are equal to the number of vertices, edges and faces, respectively, of the triangulation. It is a theorem that the Euler characteristic does not depend on the triangulation, and an exercise to show that it is invariant under homeomorphisms.

Exercise 7.11. Verify that the Euler characteristic of Σ_g equals $2 - 2g$.

In particular, via the equality $g = \frac{2-\chi}{2}$, we can directly read off the genus, and hence the homeomorphism type, of a closed orientable surface as soon as we have any triangulation.

7.2 Thurston's construction of pseudo-Anosov maps

This section closely follows parts of [5].

7.2.1 An example

Consider the closed surface Σ depicted in Figure 7.1.

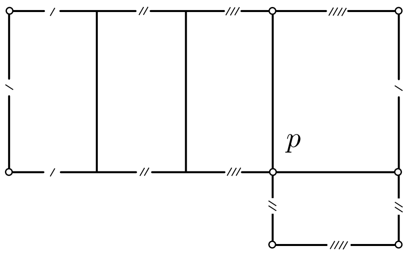

Figure 7.1: A closed surface Σ built from rectangles. The vertex p and all the other vertices which get identified with p are highlighted.

It consists of rectangles with parallel identifications of horizontal and vertical sides. We consider the horizontal and the vertical annuli H_i and V_i into which our decomposition into rectangles divides the surface. In our example, there are four vertical and two horizontal annuli. We claim (and show later in Proposition 7.13) that there exists a choice of the side lengths of the rectangles such that the ratio between the length and the width of every annulus H_i and V_i is equal to $r = \sqrt{\frac{5+\sqrt{13}}{2}} \approx 2.074$. Furthermore, we will see that this choice is unique up to a global scaling factor, so that the ratio r is unique.

Figure 7.2 shows two homeomorphisms f and g of Σ.

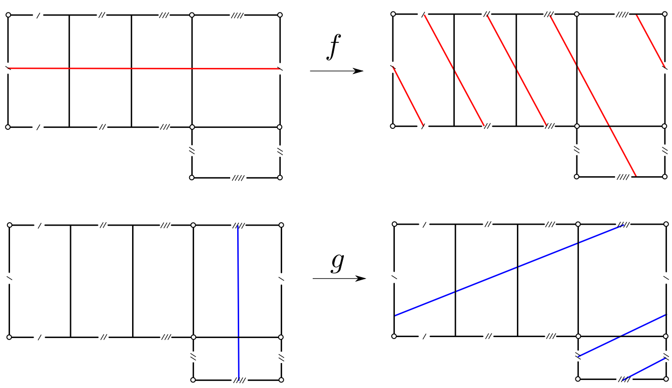

Figure 7.2: The action of two affine homeomorphisms f and g on Σ.

The homeomorphism f pointwise fixes vertical sides of the rectangles, preserves all other vertical lines and sends horizontal lines to lines with slope $-r$.

In particular, away from the point p we can write f as an affine map with linear part $\left(\begin{smallmatrix}1 & 0\\ -r & 1\end{smallmatrix}\right)$. Analogously, g pointwise fixes the horizontal sides of the rectangles, preserves horizontal lines and sends vertical lines to lines with slope $1/r$. In particular, away from the point p, also g is an affine map with linear part $\left(\begin{smallmatrix}1 & r\\ 0 & 1\end{smallmatrix}\right)$.

The composition $g \circ f$ must, away from the point p, also be an affine map with linear part

$$D = \begin{pmatrix}1 & r\\ 0 & 1\end{pmatrix}\begin{pmatrix}1 & 0\\ -r & 1\end{pmatrix} = \begin{pmatrix}1-r^2 & r\\ -r & 1\end{pmatrix}.$$

The linear part D of $g \circ f$ has two real eigenvalues $\lambda \approx -1.722$ and $1/\lambda$, and corresponding eigendirections (at every point except p). See Figure 7.3 for an illustration of how these eigendirections cover the surface Σ.

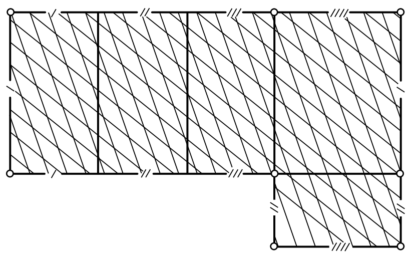

Figure 7.3: Eigendirections for $(g \circ f)$.

The map $(g \circ f)$ stretches by $|\lambda|$ and $|1/\lambda|$, respectively, along these directions. The number $|\lambda|$ is called the *dilatation* of $g \circ f$ and is an algebraic integer. It's minimal polynomial can be shown to be $t^4 - t^3 - t^2 - t + 1$.

So what about the point p? Arount the point p, the total angle in our representation of the surface Σ as a polygon in the real plane is 6π. This means that around p, the eigendirections "wind" similar to a 6-saddle, as depicted in Figure 7.4. The presence of such "singularities" for the eigendirections

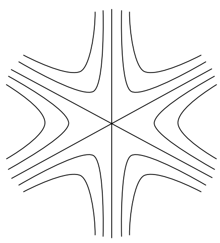

Figure 7.4: The eigendirections for λ in a neighbourhood of the point p.

of the map $(g \circ f)$ is necessary for surfaces of genus $g \geq 2$, which is why we call such maps *pseudo-Anosov* instead of *Anosov*.

Exercise 7.12. What is the genus of the surface Σ treated in this section?

7.2.2 Thurston's construction

Let Σ be a surface that is represented by a union rectangles in the real plane whose sides are parallel to the vertical and the horizontal direction. Furthermore, fix some identification of the sides via translations in the plane, so that no side remains unidentified. In particular, horizontal sides get identified with horizontal sides and vertical sides get identified with vertical sides, and identified sides must be of the same length. As in Section 7.2.1, we decompose Σ into its horizontal and vertical annuli H_i and V_i, respectively.

Proposition 7.13. *It is possible to change the width of each annulus H_i and V_j so that the ratio between the length and the width is a fixed number $r \geq 1$ for each horizontal and each vertical annulus. Furthermore, this ratio r is unique.*

Proof. Let $\Omega = \left(\begin{smallmatrix} 0 & X \\ X^\top & 0 \end{smallmatrix}\right) \geq 0$ be the geometric intersection matrix of the annuli H_i and V_j, that is, each entry of Ω is given by the number of rectangles in which the corresponding annuli overlap. Furthermore, let r be its Perron-Frobenius eigenvalue and let $v > 0$ be the corresponding eigenvector, so we have $\Omega v = rv$. We resize each horizontal and vertical annulus to have width equal to the corresponding entry of the vector v, thus changing also the length of the annuli overlapping it.

Now, the length of a horizontal annulus equals the sum of all the widths of the vertical annuli it overlaps, counted with multiplicity. This number equals $(\Omega v)_i$, where the i-th entry corresponds to the horizontal annulus under consideration. Furthermore, the width of a horizontal annulus is simply v_i. Then, the claim follows directly from $\Omega v = rv$, which implies the coordinate-wise equality $(\Omega v)_i = rv_i$. The same argument applies to vertical annuli.

Finally, v is the unique eigenvector of Ω with strictly positive entries by the Perron-Frobenius theorem 2.17. From this it follows that $r \geq 1$ is uniquely determined by the matrix Ω. □

The following way of constructing maps on surfaces is called *Thurston's construction of pseudo-Anosov maps*:

One starts with a surface Σ as above, and applies Proposition 7.13. Then one defines the maps f and g, analogously to what we did in Section 7.2.1 as affine maps outside of the points p around which the angle does not add to 2π. If one takes any product of the maps f and g such that the linear part is a hyperbolic matrix, then there exist eigenvalues λ with $|\lambda| > 1$ and $1/\lambda$ with corresponding eigendirections along which the map stretches by $|\lambda|$ or $|1/\lambda|$, respectively.

Exercise 7.14. Show that Thurston's construction is general enough to yield infinitely many examples of maps with hyperbolic linear part on every surface of genus g, for $g \geq 1$.

Hint: Start from a staircase of squares in the plane, then apply Thurston's construction.

7.3 Pseudo-Anosov maps

We finally give a definition of pseudo-Anosov maps. While the maps obtained via Thurston's construction are pseudo-Anosov by definition, we note without proof that many more pseudo-Anosov maps exist than just the ones obtained in this way.

Definition 7.15. *A homeomorphism f of a closed surface Σ is pseudo-Anosov if the map $f : \Sigma \to \Sigma$ is topologically conjugate to $\widehat{f} : \widehat{\Sigma} \to \widehat{\Sigma}$, where $\widehat{\Sigma}$ is a surface represented by rectangles in the plane, as defined at the beginning of Section 7.2.2, and the map $\widehat{f}$ is affine with hyperbolic linear part outside the finite set of points around which the total angle in our planar representation of $\widehat{\Sigma}$ does not equal 2π. The larger absolute value $|\lambda| > 1$ among the eigenvalues of the linear part is called the dilatation.*

The following is a landmark result in the study of isotopy classes of homeomorphisms of surfaces. For much more on pseudo-Anosov homeomorphisms and their ramifications, we recommend the book of Farb and Margalit [4].

Theorem 7.16 (Thurston classification of surface homeomorphisms, 1970s)**.** *Let $f : \Sigma \to \Sigma$ be a homeomorphism of a closed surface that is not isotopic to a map which preserves a system of simple closed curves on Σ. Then f is either isotopic to a periodic map or to a pseudo-Anosov map.*

To finish, we content ourselves with the following result on the orbits of pseudo-Anosov maps, showing that their dynamical properties are very similar to those of Anosov maps.

Theorem 7.17. *Let $f : \Sigma \to \Sigma$ be a pseudo-Anosov map. Then*

a) the map f is topologically mixing,

b) the periodic points for f form a dense subset of Σ.

Proof. We work with the topologically conjugate map $\hat{f} : \hat{\Sigma} \to \hat{\Sigma}$.

For a), we only give a sketch of the proof. We note that a pseudo-Anosov map must permute the members of the finite set of points around which the total angle in our planar representation of $\widehat{\Sigma}$ does not equal 2π. We assume that one of these points p is fixed (the proof of a permutation moving every point can be reduced to this case). Now consider the curve γ on $\widehat{\Sigma}$ defined by following an eigendirection for λ starting at p. One can show that this curve is not closed, otherwise we get a contradiction: mapping γ by f must give back γ, but it should also stretch by a factor of $|\lambda|$. In fact, we claim that γ is dense in $\widehat{\Sigma}$, and refer to [4] for a proof. To finish the proof, we may now copy the argument we used to show topological mixing for Anosov maps of the torus in Proposition 3.14.

For b), let S be a square in $\widehat{\Sigma}$, parametrised in such a way that

i) S does not contain any of the points around which the total angle in our planar representation of $\widehat{\Sigma}$ does not equal 2π,

ii) the vertical lines of the rectangle are parallel to the eigendirection for $1/\lambda$, and

iii) the horizontal lines are parallel to the eigendirection for λ.

Our goal is to show that S contains a periodic point, which suffices to finish the proof.

We first note that the map $\widehat{f}$, as it is affine with hyperbolic linear part of determinant ± 1, preserves the Lebesgue measure $\mathcal{L}^2$ on $\widehat{\Sigma}$. We can therefore use Poincaré recurrence 6.5 on the set S and obtain that for any N there exists $n > N$ such that $\widehat{f}^n(S) \cap S$ is nonempty.

Let $x_1 \in S$ such that $\widehat{f}^n(x_1) \in S$. Let J be be the vertical segment of S containing x_1. Since $\widehat{f}$ contracts arcs parallel to the eigendirection for $1/\lambda$, we can assume to have chosen N large enough such that $\widehat{f}^n(J) \subset S$.

There is a natural map from $\widehat{f}^n(J)$ to J, obtained by moving the arc $\widehat{f}^n(J)$ horizontally in S until it hits the arc J. Composing this map with $\widehat{f}^n$ yields a map $J \to J$. This map is a contraction with respect to the metric of the segment $J \subset S$ induced by the standard metric on the square S, and so it has a fixed point $x_2 \in J$, by the Banach fixed point theorem 2.4.

Let L be the horizontal segment of S containing x_2. Assuming to have chosen N large enough, we have $L \subset \widehat{f}^n(L)$. Here, we need that the vertical coordinate of x_2 equals the vertical coordinate of $\widehat{f}^n(x_2)$, and the fact that $\widehat{f}$ stretches arcs parallel to the eigendirection for λ by a factor > 1.

We now look at the map $(\widehat{f}^n)^{-1} : \widehat{f}^n(L) \cap S \to \widehat{f}^n(L) \cap S$. Again, this map is a contraction with respect to the metric on the segment $\widehat{f}^n(L) \cap S$, and so it has a fixed point x by the Banach fixed point theorem 2.4. We have shown that $x \in S$ is a periodic point for $\widehat{f}$ with period n. □

Bibliography

[1] O. Bogopolski: *Introduction to group theory*, European Mathematical Society, Zürich, 2008.

[2] M. Brin, G. Stuck: *Introduction to dynamical systems*, Cambridge University Press, Cambridge, 2002.

[3] A. Katok, B. Hasselblatt: *Introduction to the modern theory of dynamical systems*, Cambridge University Press, Cambridge, 1995.

[4] B. Farb, D. Margalit: *A primer on mapping class groups*, Princeton University Press, Princeton, NJ, 2012.

[5] L. Liechti: *On the spectra of mapping classes and the 4-genera of positive knots*, Inauguraldissertation, Universität Bern, 2017, `https://biblio.unibe.ch/download/eldiss/17liechti_l.pdf`

[6] A. Putman: *A quick proof of the classification of surfaces*, unpublished note, `https://www3.nd.edu/~andyp/notes/ClassificationSurfaces.pdf`