

Universitext

Jürgen Jost

Dynamical Systems

Examples of Complex Behaviour

With 50 Figures

 Springer

Jürgen Jost
Max Planck Institute
for Mathematics in the Sciences
Inselstr. 22
04103 Leipzig
Germany
e-mail: jost@mis.mpg.de

Mathematics Subject Classification (2000): 37-XX, 34Cxx, 34Dxx

Library of Congress Control Number: 2005925883

ISBN-10 3-540-22908-6 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-22908-7 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springeronline.com
© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: by the author using a Springer \LaTeX macro package
Cover design: Erich Kirchner, Heidelberg

Printed on acid-free paper SPIN 10860185 41/TechBooks - 5 4 3 2 1 0

Preface

Our aim is to introduce, explain, and discuss the fundamental problems, ideas, concepts, results, and methods of the theory of dynamical systems and to show how they can be used in specific examples. We do not intend to give a comprehensive overview of the present state of research in the theory of dynamical systems, nor a detailed historical account of its development. We try to explain the important results, often neglecting technical refinements and, usually, we do not provide proofs.¹

One of the basic questions in studying dynamical systems, i.e. systems that evolve in time, is the construction of invariants that allow us to classify qualitative types of dynamical evolution, to distinguish between qualitatively different dynamics, and to study transitions between different types. It is also important to find out when a certain dynamic behavior is stable under small perturbations, as well as to understand the various scenarios of instability. Finally, an essential aspect of a dynamic evolution is the transformation of some given initial state into some final or asymptotic state as time proceeds. The temporal evolution of a dynamical system may be continuous or discrete, but it turns out that many of the concepts to be introduced are useful in either case.

We first introduce some general notions and exemplify them for systems of ordinary differential equations. We classify some simple types of dynamical behavior, like fixed points, and discuss the stability issue. We introduce the notion of typical or generic behavior and study bifurcations, i.e. transitions between different types of behavior. Attractors represent important asymptotic dynamical invariants. Another aspect is the distinction between dynamically contracting and expanding directions and those that are neither. The latter constitute the so-called center manifold and encode the dynamically nontrivial part of the evolution.

¹ All proofs can be readily found in the references provided in the bibliography. My conscience as a mathematician does not allow me to suggest that you study dynamical systems without seeing the proofs of the difficult results. Therefore, I hope that you will consult at least some of these references.

The theory of Conley allows a detailed investigation of qualitative features of dynamical systems in terms of discrete algebraic invariants. The theory is presented in detail.

Kolmogorov introduced the fundamental asymptotic invariant for a dynamical system, the entropy. The topological entropy is an important tool for analyzing so-called chaotic behavior, and the method of symbolic dynamics transforms a continuous scenario into a discrete one.

The metric aspects of entropy allow us to discuss the issue of complexity and the absence or presence of intrinsic scales of a dynamical process. The measure theoretic entropy establishes a fundamental connection with Shannon's concept of information. Lyapunov exponents of a dynamical system are easier to compute than the entropy, but can sometimes provide an alternative to the latter for analyzing the relation between expansion and contraction of a dynamical process. For that aspect, a rather complete theory exists under a certain assumption of structural stability, called hyperbolicity.

We also discuss cellular automata and the more general Boolean networks as examples of discrete dynamical systems.

Of course, this short survey cannot treat the field of dynamical systems exhaustively. The most important omission is perhaps the theory of Hamiltonian and integrable dynamical systems and its profound connections with symplectic geometry for which a standard presentation is available in several recent textbooks.

While any individual mathematician who develops a new concept or demonstrates an important result is rightly proud of her or his achievement, in general we mathematicians are inclined to consider important mathematical theories and results to be the common property of all mathematicians, if not of all of mankind. Perhaps for that reason, we are not always very diligent in tracing the history of individual contributions, and this may serve as a faint excuse for not always carefully searching and listing all individual references in the present survey.

The present book emerged from series of lectures given at Leipzig University and the Santa Fe Institute for the Sciences of Complexity to rather diverse audiences. I thank them all for their interest, their inspiring questions, and their constructive criticism. I am grateful to Antje Vandenberg and Pengcheng Zhao for technical help.

Table of Contents

1	Introduction	1
2	Stability of dynamical systems, bifurcations, and generic properties	7
2.1	Some general notions	7
2.2	Autonomous systems of ODEs	9
2.3	Examples: Bifurcation depending on a parameter $\lambda \in \mathbb{R}$	15
2.4	Chaos in differential and difference equations. The concept of an attractor	22
2.5	Interaction, or the interplay between concentration or reaction and diffusion	36
2.6	Discrete and continuous systems. The Poincaré return map ..	41
2.7	Stability and bifurcations; generic properties	42
2.8	The Hopf bifurcation	46
2.9	Lotka-Volterra equations	48
2.10	Stable, unstable, and center manifolds	52
3	Discrete invariants of dynamical systems	61
3.1	The topology of graphs	61
3.2	Floer homology	62
3.3	Conley theory: examples and results	70
3.4	Cohomological Conley index	77
3.5	Homotopical invariants	80
3.6	Continuation properties of the Conley index	92
3.7	The discrete Conley index	93
4	Entropy and topological aspects of dynamical systems	99
4.1	The entropy of a process as an asymptotic quantity	99
4.2	Positive entropy and chaos	103
4.3	Symbolic dynamics	108
5	Entropy and metric aspects of dynamical systems	111
5.1	The metric approach to topological entropy	111
5.2	Complexity and intrinsic scales	114

6	Entropy and measure theoretic aspects of dynamical systems	119
6.1	Probability spaces and measure preserving maps	119
6.2	Ergodicity	121
6.3	Entropy and information	126
6.4	Invariant measures	138
6.5	Stochastic processes	142
6.6	Stochastic bifurcations	149
7	Smooth dynamical systems	153
7.1	Lyapunov exponents	153
7.2	Hyperbolicity	156
7.3	Information loss	165
8	Cellular automata and Boolean networks as examples of discrete dynamical systems	169
8.1	Cellular automata	169
8.2	Boolean networks	175
	References	181
	Index	185

1 Introduction

A dynamical system is a system that evolves in time through the iterated application of an underlying dynamical rule. That transition rule describes the change of the actual state in terms of itself and possibly also previous states. The dependence of the state transitions on the states of the system itself means that the dynamics is recursive. In particular, a dynamical system is not a simple input-output transformation, but the actual states depend on the system's own history. In fact, an input need not even be given to the system continuously, but rather it may be entirely sufficient if the input is only given as an initial state and the system is then allowed to evolve according only to its internal dynamical rule. This will represent the typical paradigm of a dynamical system for us.

The transition rules for dynamical system will typically depend on certain parameters. Investigating the qualitative nature of this dependence constitutes an important aspect of the theory of dynamical systems.

The application of the transition rule can happen either at discrete time steps, with the time parameter denoted by t taking values in the (positive) integers $\mathbb{Z}(\mathbb{N})$, or infinitesimally with continuous underlying time taking values in \mathbb{R} or \mathbb{R}^+ as in differential equations. If time is continuous, we assume that the transition rules lead to an evolution that is continuous w.r.t. some appropriate topology. The qualitative dynamical behavior of the system may, however, change due to phase transitions or bifurcations. Or, from a different perspective, there may be a transient and an asymptotic dynamical regime. If time is discrete, we need to select a class of permissible state transitions that preserve the identity of the system.

The underlying rule may be rather simple, but its iterated application may still create an asymptotic behavior as time goes to infinity that is not so easy to predict and analyze from the dynamical rule itself. In fact, in many, and perhaps typical, cases, there is no simpler way to obtain or predict the final result than to let the dynamical system run itself. Thus, in hindsight, the attitude of Laplace seems rather naive that, given the complete initial conditions, future states of the world could be computed. The point is that the dynamical evolution may be so complicated that, for that prediction, a computer that is essentially as powerful as the world itself would be needed, and this computation would then take about as long as the corresponding

evolution of the world itself. While this already touches on some issues from theoretical computer science, the theory of dynamical systems can offer powerful insights that help to clarify this point.

A dynamical system transforms an initial state in its state space through the dynamical iteration into an asymptotic final state, and this transformation is what one wishes to understand. Thus, the fundamental question is what one can say about the asymptotic final state from the knowledge of the initial condition and the dynamical rule. In particular, one is interested in its qualitative aspects, in a sense to be made precise. In spite of the complexity of the computation just argued, it is typically still possible to derive qualitative results that are stable under large classes of variations of the initial conditions or the dynamical rule. Often, in fact, that asymptotic state is qualitatively simpler than the initial state the dynamics started with. Thus, the underlying space of possible initial states might get transformed into a simple collection of attractors. That collection then captures the essential aspects of the corresponding initial conditions. In this sense, the attractors constitute a classification of the possible initial conditions. Knowing those attractors, and in addition the unstable invariant states the dynamical system possesses, allows us to reconstruct the qualitative features of the underlying space of initial conditions. Thus, no essential information gets lost during the dynamical process. Conversely, knowing the topology of that underlying space allows to derive constraints on the system of invariant sets of the dynamics. This interplay between information contained in the space of initial conditions and relationships between dynamically invariant states is one of the main themes of the theory of dynamical systems.

This transformation of the initial into the asymptotic final states is sometimes considered as a computation performed by the dynamical system, and our above discussion may suggest such a point of view. In this interpretation, the initial state is the input of the system, and the asymptotic final state that is achieved by letting the system run on the basis of this input then is the output. Since one might prefer to consider computation as the purposeful extraction from external input of information that is relevant for the system, and since I do not wish to enter into a discussion of purpose or meaning here, it is perhaps preferable to employ the more neutral term of translation for describing this dynamic transformation process.

There seems to be some contradiction in the preceding. On one hand, I have argued that prediction is impossible because the complexity of the dynamics cannot be reduced. On the other hand, I have claimed that the dynamical system can achieve a simple classification of a large and diverse set of inputs by assigning to each input one attractor out of some small collection, and that the essential qualitative aspects of that input space are reflected and preserved in the collection of asymptotic invariant states of the dynamical system. In fact, one should see these aspects as complementary rather than contradictory. A dynamical system may simultaneously possess complexifying

features that amplify small differences, in the sense of a so-called chaotic evolution, as well as simplifying tendencies because some degrees of freedom constrain and dominate the other ones. It may not be easy to predict to which attractor some given initial state is asymptotically drawn, and two very similar initial states may subsequently evolve in completely different directions. When we pass to a different scale of resolution, however, some of these fluctuations, unpredictabilities, and perturbations may average out and we may see a rather regular global picture. Thus, we may not be able to know everything by any method short of running the dynamical process itself, but we might still be able to capture certain qualitative and global aspects reliably. The latter is the basis of all scientific modelling. When we wish to describe some part of reality in a formal model, we need to identify the essential degrees of freedom of the system or process whose knowledge allows an understanding and perhaps even prediction of the important and relevant aspects, at least at a global scale or in the long term. The mathematical theory of dynamical systems offers important conceptual and technical tools for that as we shall also explore in this monograph. After all, the task of science is not the prediction of every tiny detail of a complicated system, but rather an understanding of its essential and dominant features.

Another clarification should be made: When I speak about evolution, this term is employed in the sense used by physicists, i.e., simply as a temporal process, but not with the meaning and connotations it carries in biology.

There are two important questions that may be posed at this rather general level:

- The **isomorphism problem**: What are the relevant criteria for comparing two dynamical systems and calling them isomorphic if they exhibit the same values of these criteria? The isomorphism problem can of course be approached more successfully for small and constrained systems than for large ones with many independent degrees of freedom. Also, the problem is better posed in situations where the system is not continuously exposed to unpredictable external inputs, but rather where it starts from some initial condition as its sole input and then is left alone to run according to its dynamical rule. We shall be interested in qualitative properties of dynamical systems, and not in quantitative ones, and we wish to describe these qualitative aspects through certain continuous or discrete invariants. Thus, isomorphic systems need to have the same values of all such invariants. One of the most important invariants for dynamical systems is the entropy.
- The **identity problem**: This refers to the evolution of a given system rather than the comparison of two different systems. In contrast to the isomorphism problem, the identity problem tends to be rather trivial for small systems and becomes really meaningful only for larger systems with

many elements or degrees of freedom. In particular, we are concerned here with a situation where the possible states are taken from a large state space and can be quite intricate and complicated. We can then ask which part of the state space can be explored from given initial conditions and what types of constraints restrict the possibilities for the system evolution. Also, here we may well assume that the system is perpetually exposed to some input from its environment that it is unable to predict, but to which it needs to react. We then want to understand how external input influences the internal dynamics. Conceptually, this input can be formalized as random, unpredictable, stochastic, that is, as noise, or it could be encoded in certain parameters that govern the dynamical evolution. In some cases, these parameters are considered as fixed while, in other cases, they may themselves evolve, perhaps on a much larger time scale. – Naturally, then, for this problem a less complete and penetrating answer can be achieved than for the first one.

How can we identify a system in state $S(t_1)$ observed at time t_1 with the same system in state $S(t_2)$ observed at time t_2 ? The problem is that $S(t_2)$ may be qualitatively different from $S(t_1)$ as it may have undergone structural changes during its evolution from t_1 to t_2 . The question then is to find out what types of structural changes are possible during the evolution of a given system and what types can be excluded, for example on the basis of certain invariants that have to remain constant during the evolution process. While in some cases the parameters governing the dynamical system are externally given, and perhaps constant, in other situations they themselves obey an intrinsic dynamical rule, but typically on a slower time scale. In that case, the combination of the evolution on two different time scales may lead to intrinsically caused bifurcations, that is, qualitative dynamical changes. Often, the point of view can also be reversed in the sense that the dynamical evolution itself dynamically separates into different time scales, with the slower one containing the essential aspects of the evolution of the system.

Intimately connected with these two problems is a third one, namely

- The **stability problem**: When is the qualitative behavior of a dynamical system insensitive to small perturbations? In terms of the question of the transformation of the initial condition into an asymptotic final state, this is the issue of **dynamical stability** and means that slight variations of the initial conditions should typically yield similar, if not identical, final states. This aspect will be made precise through the concept of an attractor and its basin of attraction. However, even the presence of such an attractor does not guarantee such a stabilization. Namely, many dynamical systems exhibit so-called chaotic behavior, that is, differences in the initial conditions are ever more amplified as time proceeds. There may even exist so-called strange or chaotic attractors that, while attracting nearby initial values, carry an internal dynamics that amplifies differences. In terms of

the isomorphism problem, this is the issue of **structural stability** and means that a system obtained from a given one through a slight perturbation of a parameter remains isomorphic to the original one. In terms of the identity problem, this means that the system does not undergo structural changes during a certain time period, but remains qualitatively invariant. However, as this is often not the case, one also needs to investigate the possible structural changes occurring either through the variation of a parameter or in the course of the evolution of a given system. Again, one may ask what types of structural changes are stable under slight perturbations or generic in the sense that they occur in typical situations. Of course, the preceding discussion involves several notions that still need to be assigned a precise mathematical meaning.

Finally, a fourth important problem concerns

- The **statistical behavior of dynamical systems**: What are the dynamic properties of some system for average initial conditions or parameter values? How much information can be extracted from the observation of a single evolution of the system for a fixed choice of initial conditions or parameters?

Also, the dynamic rules may contain deterministic as well as stochastic components, and the effect of the latter should also be investigated. Ergodic theory is concerned with the most important aspect of the statistical behavior of dynamical systems, namely the following question:

Let (X, μ) be a measure space, and let $T : X \rightarrow X$ be measure preserving and bijective (possibly up to sets of measure 0). The **ergodic problem** is to isolate conditions under which the temporal average of a measurable function $f : X \rightarrow \mathbb{R}$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{\nu=0}^{n-1} f(T^\nu x),$$

exists ($T^\nu = T \circ \dots \circ T$; ν times) and coincides with the spatial average or mean

$$\int_X f(y) \mu(dy),$$

at least for μ -almost all x . Thus, if the system satisfies this ergodic property then, for obtaining the average of a function over the state space, it suffices to look at the dynamic iterates of a single (typical) point in that state space X . Such dynamic iterates are often given as **time series**.

In fact, we now have two objects here, namely the transformation T and the measure μ . Thus, given T , we can first try to find such a measure μ that is left invariant by T , and we can then ask whether we can also find an ergodic invariant measure, perhaps with certain further properties imposed.

Conversely, if the measure is given, for example as the volume measure for some underlying Riemannian metric on X , then perhaps T will not leave it

invariant, but it still offers the perspective of considering all initial conditions simultaneously instead of a single one. This is particularly useful when we study differential equations. The traditional local approach was to take an initial condition and ask whether the differential equation possesses a unique solution for that initial condition for all times, and then whether this solution depends continuously or even smoothly on that initial condition for all finite times. The modern global approach rather considers the differential equation as a dynamical system that, for each time t (> 0), transforms the space of initial conditions (corresponding to time 0 by an arbitrary convention) into the space of the corresponding values at time t . Under appropriate conditions (essentially the ones that ensure unique existence for all times and smooth dependence on initial conditions), this yields a diffeomorphism of the space of initial conditions, that is, the space on which the dynamical system operates. It is important to realize, however, that as $t \rightarrow \infty$ these diffeomorphisms need not converge to a limit diffeomorphism. For example, if the system has a global point attractor, then the solution of the differential equation will converge to that point for any initial condition, and so, our diffeomorphisms will also converge to a point map and not to a diffeomorphism.

In other situations, for example those typically arising in statistical mechanics, we do not know the initial value precisely as a point in our state space, but rather only as a certain probability measure. Thus, again, we do not consider iterates of a single point and perhaps seek a fixed point, but rather iterates of a probability measure and look for an invariant measure to which those iterates converge.

2 Stability of dynamical systems, bifurcations, and generic properties

2.1 Some general notions

A flow (semiflow) is a family

$$F_t : X \rightarrow X$$

of maps of a set X (state or phase space) into itself, for $t \in \mathbb{R}$ ($t \geq 0$), satisfying

- (i) $F_0 = Id$
- (ii) $F_{t+s} = F_t \circ F_s$ for all $t, s \in \mathbb{R}$ ($t, s \geq 0$) ((semi) group property).

Here, t is considered as a time parameter. For a given initial state x_0 , we study the process

$$x(t) := F_t x_0$$

transforming the initial state into the state at time t .

The map

$$t \mapsto x(t)$$

is called a **trajectory**, and $\{x(t) : t \in \mathbb{R}(t \geq 0)\}$ is the **orbit** of x_0 . x_0 is called a **fixed point** or a **stationary point** if

$$x(t) = x_0 \text{ for all } t.$$

More generally, a trajectory or an orbit is called **periodic** if

$$x(t + \omega) = x(t) \text{ for some } \omega \geq 0 \text{ and all } t.$$

(ii) implies that

$$x(t + s) = F_t x(s),$$

i.e., $y(t) := x(t + s)$ is the state at time t of the process with initial state $y(0) = x(s)$.

In the discrete case, we consider instead the iterates of a given map

$$(2.1) \quad \begin{aligned} &F : X \rightarrow X, \\ &x_{n+1} = F(x_n) \text{ for } n \in \mathbb{Z} \quad (n \in \mathbb{N}). \end{aligned}$$

An important example of a discrete dynamical system is obtained by letting

$$F = F_1$$

be the time 1-map of a (semi) flow as above. In that case

$$x(n) = F^n(x_0)$$

are the values of the flow at integer times.

More generally, we could let F in (2.1) also depend on n , $F = F_n$. That would be a so-called non-autonomous situation.

The dynamical system given by (2.1) is invertible when each state $x(n)$ has a unique predecessor $x(n - 1)$, that is, if two trajectories can never merge.

Let X be a topological space, and $F : X \times \mathbb{R} \rightarrow X$ a flow with $F(x, 0) = x$, and put $x(t) := F(x, t)$. For $y \in X$, we put

$$(2.2) \quad \alpha(y) := \bigcap_{t \in \mathbb{R}} \overline{y((-\infty, t))},$$

$$(2.3) \quad \omega(y) := \bigcap_{t \in \mathbb{R}} \overline{y((t, \infty))}.$$

$\alpha(y)$ tells us where y came from in the infinite past and, likewise, $\omega(y)$ encodes where it goes in the infinite future. An important point is that, in particular in the examples that are of interest for us, many points y may have the same α or ω limit sets. Thus, looking at the collection of all such limit sets, we see to what extent the asymptotic dynamics reduces or simplifies the underlying space X on which it operates.

When X is a metric space with metric $d(., .)$, we have

$$\begin{aligned} \alpha(y) &= \{z : \text{for all } \epsilon > 0, \text{ there exist arbitrarily large } t \text{ with } d(y(-t), z) < \epsilon\} \\ \omega(y) &= \{z : \text{for all } \epsilon > 0, \text{ there exist arbitrarily large } t \text{ with } d(y(t), z) < \epsilon\}. \end{aligned}$$

Similar concepts apply also in the time discrete case. If we consider the system given by $x_n = F(x, n)$, for instance, the iterates F^n (that is, $x_n = F^n(x)$) for $n \in \mathbb{N}$ of a continuous map $F : X \rightarrow X$, we have the forward limit set

$$(2.4) \quad \omega(y) = \{z \in X : \text{for any neighborhood } U \text{ of } z, \\ \text{there exist arbitrarily large } n \text{ with } y(n) \in U\}.$$

Again, in a metric space X , this condition can be formulated in metric terms. The forward limit set $\omega(y)$ is the set of points to which the orbit of y comes arbitrarily close infinitely often in forward time. The limit set $\omega(y)$ may or may not contain points from the orbit of y . A fixed point and a periodic orbit are their own limit sets, but they can also be the limit sets for other points as well, namely for those attracted by it (see 2.4). y and its iterates $y(n)$ have the same limit set. Therefore, when $z \in \omega(y)$, then also the entire forward orbit of z , i.e. all iterates $z(n)$, $n \in \mathbb{N}$, lie in $\omega(y)$ as well.

2.2 Autonomous systems of ODEs

Let $f = (f^1, \dots, f^d) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a mapping of class C^1 . We consider the system of ODEs

$$(2.5) \quad \dot{x}^i(t) = f^i(x^1(t), \dots, x^d(t)) \text{ for } i = 1, \dots, d,$$

with $\dot{x}^i = \frac{d}{dt}x^i$. Such a system is called **autonomous** because f does not depend explicitly on t (but implicitly through the dependence of x on t). One may also consider non-autonomous systems,

$$(2.6) \quad \dot{x}^i(t) = \phi^i(t, x^1(t), \dots, x^d(t)) \text{ for } i = 1, \dots, d,$$

with an explicit dependence on t , but such systems can be converted into an autonomous form by introducing a new dependent variable x^{d+1} and the equation

$$(2.7) \quad \dot{x}^{d+1}(t) = f^{d+1}(x^1(t), \dots, x^d(t), x^{d+1}(t)) \equiv 1.$$

One should note, however, that this may turn linear (non-autonomous) equations into non-linear (autonomous) ones; for example

$$\dot{x} = \sin \omega t$$

becomes

$$\begin{aligned} \dot{x}^1 &= \sin \omega x^2 \\ \dot{x}^2 &= 1 \end{aligned}$$

which couples the dependent variables x^1, x^2 in a non-linear manner.

The important point about autonomous systems is that they are invariant under time shifts. This means that, if we consider the solution of (2.5)¹ $x_1(t)$ with initial values $x_1(t_1) = \xi$ and the solution $x_2(t)$ with the same initial values, but starting at time t_2 , that is, $x_2(t_2) = \xi$, then for all $t \geq t_2$, $x_2(t) = x_1(t+t_1-t_2)$. In other words, the behavior of the solution (obviously) depends on the initial values, that is, where or how it starts, but not on the starting time, that is, when it starts.

By the theorem of Picard-Lindelöf, (2.5) defines a local flow in the sense that for every initial state x_0 , the solution $x(t)$ of this system exists on some time interval

$$-T < t < T, \text{ for some } T > 0.$$

If $f^i(x_0) = 0$ for $i = 1, \dots, d$ then x_0 is a stationary point of our local dynamical system. In order to investigate the local behavior near such a stationary

¹ assuming that there exists a unique solution, see below

point, we linearize the problem near x_0 and first investigate the local behavior of the linearized problem. We then try to find out whether the latter behavior persists in the original problem.

We may assume w.l.o.g. that

$$x_0 = 0.$$

We thus study the linearized system

$$(2.8) \quad \dot{x}(t) = Ax,$$

with $A = \left(\frac{\partial f^i}{\partial x^j}(x_0) \right)_{i,j=1,\dots,d}$, $x = (x^1, \dots, x^d)$.

We consider the case $d = 2$.

We first look at the case where A has two real eigenvalues α_1 and α_2 and can be diagonalized. After a linear change of coordinates, our system then becomes

$$(2.9) \quad \begin{aligned} \dot{x}^1(t) &= \alpha_1 x^1(t) \\ \dot{x}^2(t) &= \alpha_2 x^2(t), \end{aligned}$$

and hence

$$(2.10) \quad \begin{aligned} x^1(t) &= e^{\alpha_1 t} x^1(0) \\ x^2(t) &= e^{\alpha_2 t} x^2(0). \end{aligned}$$

If α_1 and α_2 are both negative, then $x(t)$ converges to 0 with exponential speed while, in the case where α_1 and α_2 are both positive, $x(t)$ exponentially expands. In fact, in both cases $x(t)$ moves along the curves

$$(x^1)^{\alpha_2} = \text{const.} (x^2)^{\alpha_1},$$

because




$$\frac{(x^1(t))^{\alpha_2}}{(x^2(t))^{\alpha_1}}$$

remains constant.

In the first case, $x = 0$ is called a **sink** or **node**, and it is a stable fixed point for $t \rightarrow \infty$, while in the second case, a **source**, it is unstable for $t \rightarrow \infty$. The two cases are interchanged under a time reversal $t \rightarrow -t$.

If $\alpha_2 < 0 < \alpha_1$, then the fixed point 0 is neither stable nor unstable because any initial point on the x^2 -axis converges to 0, while all other initial points diverge under the flow. This is called a **saddle**.

Finally, we consider the case where A has two complex conjugate eigenvalues $\alpha \pm i\vartheta$. After a linear change of coordinates, we then obtain the system





$$\dot{x}(t) = \begin{pmatrix} \alpha & \vartheta \\ -\vartheta & \alpha \end{pmatrix} x(t),$$

and hence

$$x(t) = e^{\alpha t} \begin{pmatrix} \cos \vartheta t & \sin \vartheta t \\ -\sin \vartheta t & \cos \vartheta t \end{pmatrix} x(0).$$

If $\alpha < 0$, $x(t)$ moves on a spiral towards 0, if $\alpha > 0$, it expands on such a spiral, while in case $\alpha = 0$, it moves on a circle around 0.



The last case $\alpha = 0$ is principally different from all the other cases, because it is not structurally stable in the sense that an arbitrarily small variation of $\alpha = 0$ changes the qualitative behavior of the system. Even worse, while in the other cases the qualitative behavior of the original system near the fixed point 0 is the same as that of the linearized system, this is in general no longer so in case $\alpha = 0$. This example will be taken up again in 2.8.

We now revert to general dimension d and formulate

Definition 1. *The fixed point 0 is called hyperbolic if all eigenvalues of the linearized system have nonvanishing real part.*

Thus, the dynamic behavior near a hyperbolic fixed point is structurally stable. This assertion will be made more precise in the sequel.

In the time-discrete case, we may consider the time-1-maps of the above examples. We then obtain linear maps of the form

$$\begin{aligned}x^1 &\mapsto e^{\alpha_1} x^1 \\x^2 &\mapsto e^{\alpha_2} x^2\end{aligned}$$

or

$$\begin{pmatrix} x^1 \\ x^2 \end{pmatrix} \mapsto e^\alpha \begin{pmatrix} \cos \vartheta & \sin \vartheta \\ -\sin \vartheta & \cos \vartheta \end{pmatrix} \begin{pmatrix} x^1 \\ x^2 \end{pmatrix},$$

where $x = 0$ again is a fixed point.

Analogous to the definition in the time continuous case, we formulate the

Definition 2. *A linear map $A : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called hyperbolic if it has maximal rank and none of its eigenvalues has absolute value 1.*

Let $f : M \rightarrow M$ be a map. $p \in M$ is called a periodic point of f of period n if $f^n(p) = p$. A periodic point of f of period n is thus a fixed point of f^n .

As another class of examples, we consider gradient flows. Here, $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is a C^2 -function, that is, twice continuously differentiable, and we consider

$$(2.11) \quad \dot{x}(t) = -DF(x(t))$$

where D stands for the gradient; in coordinates, (2.11) is

$$(2.12) \quad \dot{x}^i(t) = -\frac{\partial F(x(t))}{\partial x^i} \text{ for } i = 1, \dots, d.$$

The minus sign is purely conventional – one could as well consider the positive instead of the negative gradient flow.

We compute

$$(2.13) \quad \frac{d}{dt} F(x(t)) = DF(x(t))\dot{x}(t) = -|\dot{x}(t)|^2.$$

Thus, F is a decreasing function along any flow line, and strictly so except at those points where $DF(x) = 0$. These latter points are called the **critical points** of F . These then are precisely the stationary points for our dynamics. We now consider a flow line $x(t)$. Such a flow line could, for example, be specified, by imposing the condition

$$x(0) = x_0$$

for some $x_0 \in \mathbb{R}^d$. We write

$$(2.14) \quad x(\pm\infty) := \lim_{t \rightarrow \pm\infty} x(t),$$

assuming that these limits exist. We now assume that the first and second derivatives of F are bounded on \mathbb{R}^d . Then

$$(2.15) \quad \ddot{x}(t) = -D^2F(x(t))DF(x(t))$$

is also bounded and, consequently, $\dot{x}(t)$ is uniformly Lipschitz continuous. From (2.13), we obtain for $t_1, t_2 \in \mathbb{R}$

$$(2.16) \quad \begin{aligned} F(x(t_1)) - F(x(t_2)) &= - \int_{t_1}^{t_2} \frac{d}{dt} F(x(t)) dt = \int_{t_1}^{t_2} |\dot{x}(t)|^2 dt \\ &= \int_{t_1}^{t_2} |DF(x(t))|^2 dt. \end{aligned}$$

Therefore, if $F(x(t))$ is bounded along our flow line for $t \rightarrow \infty$, (2.16) and the Lipschitz continuity of $\dot{x}(t)$ imply that

$$(2.17) \quad \lim_{t \rightarrow \infty} DF(x(t)) = \lim_{t \rightarrow \infty} \dot{x}(t) = 0.$$

We are thus inclined to believe that, as $t \rightarrow \infty$, $x(t)$ converges to a critical point $x(\infty)$ of F , and analogously for $t \rightarrow -\infty$. From the preceding considerations, we can conclude at least that, if F is bounded, we can find some sequence t_n tending to ∞ for which $x(t_n)$ converges to a critical point $x(\infty)$. We now linearize our system about $x(\infty)$. To make contact with the considerations in the beginning of this section, we assume w.l.o.g. that $x(\infty) = 0$. The linearized problem is then

$$(2.18) \quad \dot{x}(t) = Ax$$

with

$$A = \left(\frac{\partial^2 F}{\partial x^i \partial x^j} (x(\infty)) \right)_{i,j=1,\dots,d}.$$

Since A is symmetric because F is twice continuously differentiable, it has only real eigenvalues $\alpha_1, \dots, \alpha_d$ and, after a linear change of coordinates, the linearized dynamics is given by

$$(2.19) \quad x^i(t) = e^{\alpha_i t} x^i(0).$$

We call the critical point $x(\infty)$ **non-degenerate** if all eigenvalues of the Hessian $D^2F(x(\infty)) = \left(\frac{\partial^2 F}{\partial x^i \partial x^j} (x(\infty)) \right)_{i,j=1,\dots,d}$ are non-zero. This means

that $x(\infty)$ is hyperbolic in the sense of the above definition as a stationary point for our dynamics. In that case, (2.19) implies that, for the linearized dynamics, the convergence to $x(\infty)$ which takes place along a flow line is exponential. (Of course, we obtain convergence to $x(\infty)$ only when the initial values $x(0)$ fall into linear combinations of coordinate directions x^i that correspond to negative eigenvalues α_i . Other initial values asymptotically move away from $x(\infty)$ under the dynamics.) By the Theorem of Hartman-Grobman to be stated in 2.7 below, in the hyperbolic case, the situation is structurally stable in the sense that the linearized system already captures the qualitative dynamics of the original system in the vicinity of the stationary point. Thus, when $x(\infty)$ is a non-degenerate critical point, any flow line for which $x(\infty)$ occurs as an accumulation point in fact converges to it exponentially. Thus, when all critical points of F are non-degenerate, and if F is bounded, then any flow line $x(t)$ converges for $t \rightarrow \pm\infty$ to critical points $x(\pm\infty)$ with exponential speed. Another consequence of the fact that the linearized dynamics approximates the original dynamics locally in the vicinity of a critical point is that those critical points, assuming that they are non-degenerate, are in fact isolated.

The example of gradient flows will be taken up in 3.2 below.

For a gradient flow with non-degenerate critical points, the asymptotic situation is thus quite simple; namely, the whole space is dynamically transformed into the discrete set of critical points of the underlying function F . When F has more than one critical point, this transformation of $x \in \mathbb{R}^d$ into $x(\infty)$ for the flow line with $x(0) = x$ is then not continuous anymore as the asymptotic image, that is, the set of critical points, is not connected while \mathbb{R}^d of course is.

We conclude this section with a simple observation. Given a solution $x(t)$ of a system of ODEs (2.5), we obtain a time discrete dynamical system by putting $x_n = x(n)$, that is, by evaluating the flow at integer times. In the linear case, the transition is explicit, as in going from (2.9) to (2.10). In particular, we see that by exponentiating from a flow to a time discrete dynamical system, an eigenvalue 0 for the linearized flow corresponds to an eigenvalue (of absolute value) 1 for the linear discrete iteration. In the sequel, we shall see that an eigenvalue 0 for a flow and an eigenvalue of absolute value 1 for a discrete iteration both correspond to a special type of behavior that is not covered by the general theory of hyperbolic dynamical systems to be presented below.


2.3 Examples: Bifurcation depending on a parameter $\lambda \in \mathbb{R}$

A. Time continuous systems

The notation in the sequel will always be $x = x(t), y = y(t), \dot{x} = \frac{d}{dt}x(t)$ etc., for $t \in \mathbb{R}$; we shall write (x, y) in place of (x^1, x^2) .

1.


$$(2.20) \quad \dot{x} = -x^2 + \lambda.$$



A fixed point emerges at $\lambda = 0$ and bifurcates for $\lambda > 0$ into an attracting and a repelling one.

2.

$$(2.21) \quad \dot{x} = -x^3 + \lambda x (= x(-x^2 + \lambda)).$$




The attracting fixed point bifurcates into two attracting points and one repelling point.

3.

$$(2.22) \quad \begin{aligned} \dot{x} &= y - x(x^2 + y^2 - \lambda) \\ \dot{y} &= -x - y(x^2 + y^2 - \lambda). \end{aligned}$$

$(0, 0)$ is a fixed point for all λ .

$\lambda < 0$: $(0, 0)$ is a globally exponentially attracting fixed point, since



in that case


$$(2.23) \quad \frac{d}{dt} \log(x^2 + y^2) = 2(-x^2 - y^2 + \lambda) \leq 2\lambda < 0,$$

and hence

$$\log(x^2 + y^2) \leq 2\lambda t + \log(x(0)^2 + y(0)^2) \Rightarrow x^2 + y^2 \leq e^{2\lambda t}(x(0)^2 + y(0)^2).$$

$\lambda = 0$: $(0, 0)$ is still a globally attracting fixed point since, for $(x, y) \neq (0, 0)$, we have $\frac{d}{dt} \log(x^2 + y^2) < 0$, but no longer exponentially attracting.

$\lambda > 0$: $(0, 0)$ is a repelling fixed point, and a periodic orbit emerges for $x^2 + y^2 = \lambda$. This periodic orbit is attracting since for $x^2 + y^2 < \lambda$, $\frac{d}{dt} \log(x^2 + y^2) > 0$ by (2.23), but for $x^2 + y^2 > \lambda$, $\frac{d}{dt} \log(x^2 + y^2) < 0$.



The system linearized at $(x, y) = (0, 0)$,

$$\begin{aligned} \dot{x} &= y + \lambda x \\ \dot{y} &= -x + \lambda y \end{aligned}$$


has eigenvalues $\lambda \pm i$, and thus their real part vanishes at $\lambda = 0$. For $\lambda < 0$, the qualitative behavior of the original system is the same as for the linearized one. For $\lambda > 0$ this holds at least in some neighborhood of $(0, 0)$, which however depends on λ . For $\lambda = 0$ the two systems behave differently. This is an example of a Hopf bifurcation where a periodic orbit emerges at the transition from an attracting focus to a repelling one. While the situation at $\lambda = 0$ is not itself structurally stable, the Hopf bifurcation as such is structurally stable, in a sense to be made precise below (see 2.8).

4.

$$(2.24) \quad \begin{aligned} \dot{x} &= y \\ \dot{y} &= x - x^2 + \lambda y. \end{aligned}$$


$(0, 0)$ and $(1, 0)$ are always fixed points.

$(0, 0)$ is a saddle since for $\lambda = 0$ flow lines in the first and third quadrants in the vicinity of $(0, 0)$ are repelled while those in the second and fourth quadrants are attracted. For $\lambda = 0$, \dot{y} remains invariant under reflection across the x -axis while \dot{x} changes its sign. This yields the following diagram for $\lambda = 0$.



In particular, for the saddle $(0, 0)$, there exists a homoclinic orbit, i.e., one that starts at this point and returns to it.

For $\lambda < 0$, the corresponding vector field points more strongly towards the x -axis while for $\lambda > 0$ it points more strongly away from it than for $\lambda = 0$. Thus, for $\lambda < 0$ the orbits are attracted more strongly towards the x -axis and hence intersect this axis earlier than for $\lambda = 0$, while we have the opposite effect for $\lambda > 0$. We thus obtain the following diagrams:



(The situation for $\lambda > 0$ is obtained through the inversion

$$x \mapsto x, y \mapsto -y, t \mapsto -t$$

from the one for $\lambda < 0$.) In particular, for $\lambda \neq 0$ the homoclinic orbit no longer exists. (While an orbit that starts at one fixed point and returns to this same point is called homoclinic, an orbit between two different fixed points is called heteroclinic.)


B. Time discrete systems

We consider the dynamical iteration $x_{n+1} = f(x_n)$. We have the following **dynamical picture**.

We project alternately vertically onto the graph of f and horizontally onto the diagonal. The first step here associates the value $f(x_n)$ to the argument x_n at time n , and the second step turns this value into the argument x_{n+1} for the next step.

Fixed points correspond to the intersections of the graph of f with the diagonal. They are attracting or repelling depending on whether the absolute value of the derivative of the graph at the intersection point is smaller or larger than 1.

We now investigate the iteration of a map



$$f_\lambda : \mathbb{R} \rightarrow \mathbb{R},$$

that depends on a parameter $\lambda \in \mathbb{R}$.

1. $f_\lambda(x) = x + x^2 + \lambda$


$\lambda < 0$: fixed points at

$$x = \pm\sqrt{-\lambda}$$


$$x = \sqrt{-\lambda} \text{ repelling}$$

$$x = -\sqrt{-\lambda} \text{ attracting for } \lambda > -1$$


and repelling for $\lambda < -1$.



$\lambda = 0$: fixed point at $x = 0$, neither attracting nor repelling



$\lambda > 0$: no fixed point



2. $f_\lambda(x) = x + x^3 + \lambda x$.

$\lambda < 0$: attracting fixed point at $x = 0$, repelling ones at $\pm\sqrt{-\lambda}$

$\lambda = 0$: fixed point at $x = 0$

$\lambda > 0$: repelling fixed point at $x = 0$

3. $f_\lambda(x) = x^2 - \lambda x$

exhibits a bifurcation at $\lambda = 1$. 0 is a fixed point for all f_λ , attracting for $\lambda \leq 1$, repelling for $\lambda > 1$. There is another fixed point at $x = 1 + \lambda$, and thus at $x = 2$ for $\lambda = 1$. In particular, locally 0 is the only fixed point of f_λ . We consider now f_λ^2 :

$$f_\lambda^2(x) = x(1 + (\lambda - 1)(\lambda + 1 + \lambda x - 2x^2) + x^2(x - 2))$$

In addition to 0, f_λ^2 has other fixed points, namely the solutions of

$$(2.25) \quad (\lambda - 1)(\lambda + 1 + \lambda x - 2x^2) + x^2(x - 2) = 0.$$

The term that is independent of λ , $x^2(x - 2)$, possesses a double zero at $x = 0$, besides a simple zero at $x = 2$. For x in the vicinity of 0, the λ -dependent term (considered as a perturbation term for λ near 1) is dominated by $\lambda^2 - 1$. Thus, for $\lambda < 1$, (2.25) has no further solutions in the vicinity of 0 while, for $\lambda > 1$, it has two of them. For $\lambda \leq 1$, f_λ^2 has an attracting fixed point at $x = 0$, while for $\lambda > 1$, there is a repelling one, compensated by two attracting ones nearby. These latter ones, however, are not fixed under f_λ . Thus, we have a bifurcation with period doubling.

Thus, bifurcation theory analyzes the qualitative dependence of the solution of a dynamical system on parameters. In concrete situations, one may then ask where these parameters come from and what controls their variation. In some cases, they are simply externally given or controlled. In other cases, they are also intrinsically evolving, but more slowly than the dynamical system itself. For example, they can be control parameters in some feed-back system. This means that the dynamical evolution can tune its own parameters so that the system remains in some desired region of its state space. An important class of dynamical systems where the fast dynamical evolution is coupled with a slow parameter evolution is constituted by neural networks. In a neural network, this slow parameter evolution is supposed to represent learning on the basis of accumulated dynamical experience. Here, and in many other cases, the two different time scales, a fast dynamical one and a slower one for the parameter evolution, are part of the system design. In many dynamical systems, however, there is an intrinsically controlled separation into different time scales. We shall return to a systematic discussion of this aspect in 2.10. R.Thom[54] proposed to conceptualize biological growth and metamorphosis as an unfolding through parameter-induced bifurcations. Again, these parameters evolve on a slow time scale and shape the qualitative aspects of the fast dynamics, that is, of metabolism in the biological setting.

And so, again, the slow time scale of the parameters controls the fast time scale of the dynamics, but in turn depends also on the accumulated effects of the latter.

2.4 Chaos in differential and difference equations. The concept of an attractor

Chaos means sensitive dependence on initial conditions. One of the properties of chaos is that, if the dynamics is given by

$$(2.26) \quad \dot{x}(t) = F(x(t)) \quad \text{for } t \geq 0,$$

($x = (x^1, \dots, x^d)$ being vector valued),
then the difference ratio

$$(2.27) \quad \frac{|x_1(t) - x_2(t)|}{|x_1(0) - x_2(0)|},$$

or, more generally, when x lies in a metric space $(X, d(\cdot, \cdot))$, the ratio

$$(2.28) \quad \frac{d(x_1(t), x_2(t))}{d(x_1(0), x_2(0))},$$

$x_1(t), x_2(t)$ both being solutions of (2.26), can become arbitrarily large as we vary $x_1(0), x_2(0)$ and let t tend to infinity. This implies that, although the rule governing (2.26) is deterministic, if we can measure the initial conditions $x(0)$ only with a fixed precision, we cannot predict the long time evolution of $x(t)$. Chaotic behavior usually is characterized by at least one positive Lyapunov exponent, or by positive entropy: concepts to be defined in subsequent chapters.

Of course, one source of diverging paths for different initial conditions might be that the solutions themselves become unbounded; already for a linear equation for a scalar x ,

$$\dot{x}(t) = \alpha x(t)$$

with $\alpha > 0$, we have

$$\frac{|x_1(t) - x_2(t)|}{|x_1(0) - x_2(0)|} = e^{\alpha t}$$

which goes to ∞ as $t \rightarrow \infty$. But this type of behavior is not very interesting by itself; rather, we are asking whether (2.27) can diverge even if all solutions of (2.26) remain bounded. Still, this is not enough, as any non-attracting fixed point will generate such a behavior. Rather, one also requires the presence of a more complicated attractor structure. E.g. one of the definitions proposed for chaotic behavior requires that there exists a so-called strange attractor, one that is neither a fixed point nor a limit cycle, and perhaps not even a

higher-dimensional torus, but rather a nowhere rectifiable set. Examples are Cantor-type sets. Strange attractors are typically (but not necessarily) of non-integral Hausdorff dimension.

For understanding this aspect better, it might seem desirable to have some formal definitions here. However, it turns out that even one of the most fundamental concepts of the theory of dynamical systems, namely that of an attractor, does not possess a universally accepted definition. Nevertheless, it is insightful to discuss this concept here. The subsequent discussion applies to dynamical systems with continuous and discrete time alike. Thus, while some definitions are formulated in the time discrete case, it is obvious how to transfer them to the continuous case. In particular, some examples in the sequel will be taken from the time continuous case; they can in turn be easily transferred to the time discrete case by considering the iteration of the time 1 map as explained in 2.1.

We start with the definition of an attractor as presented in [24] and employed by most mathematicians.

Definition 3. *A compact set A in a topological space X is called an attractor for the continuous map $F : X \rightarrow X$ if there exists a neighborhood U of A with*

$$(2.29) \quad F(U) \subset U$$

and

$$(2.30) \quad A = \bigcap_{n \in \mathbb{N}} F^n(U).$$


Of course, an analogous definition is possible for a semi-flow $F : X \times \mathbb{R}^+ \rightarrow X$. Under the conditions of Definition 3, every point in the vicinity of A is ultimately attracted by A under the iteration of the map F . Moreover, when a point is close to A , that is, in one of the sets $F^n(U)$ for possibly large n , then its iterates have to stay in that set as well since that set also is mapped into itself by (2.29). This property is also encoded in the concept of stability as formulated by Lyapunov: A is stable in the sense of Lyapunov if for every neighborhood U of A , there exists a neighborhood V of A with the property that, for any $x \in V$, the orbit $F^n(x)$, $n \in \mathbb{N}$, stays in U .

Example:

We consider the differential equation

$$(2.31) \quad \dot{\theta} = \sin^2(\pi\theta)$$

as defined on the unit circle $S^1 = \{z = e^{2\pi i\theta}\}$. (As always, we identify $\theta + 1$ with θ .) $\theta = 0$ then is a fixed point with a homoclinic orbit. When we start with a small positive value of $\theta(0)$, $\theta(t)$ increases under the flow until it eventually approaches $\theta = 1$ which is identified with $\theta = 0$. Thus, every point on the circle, and, in particular, every point in the vicinity of $\theta = 0$ is eventually



attracted by $\theta = 0$. However, $\theta = 0$ is not an attractor in the sense of Definition 3 because it does not possess a neighborhood that contracts to it under the dynamical system. In other words, Definition 3 requires the attraction to be local, and this is not the case in the present example. $\theta = 0$ is not Lyapunov stable. In fact, the only attractor for this dynamical system is the whole circle itself.

Definition 3 does not imply that A consists only of the forward limit sets in the sense of (2.4) for the points in its vicinity. It could be larger as we shall see in Example 2 below.

Moreover, (2.29) and (2.30) imply that A is also invariant under the iterations of F :

$$(2.32) \quad F(A) \subset A.$$

Definition 4. *The largest open set U satisfying the conditions of Definition 3 is called the basin of attraction of A .*

Thus, both the attractor A and its basin of attraction $B(A)$ are dynamically invariant. When the map F is invertible, we can also extend the flow to negative times, that is, consider the iterates F^n for all $n \in \mathbb{Z}$ and not only for those in \mathbb{N} . Then both A and $B(A)$ remain invariant, but, of course, A will not be attracting in backward time, as the iterates F^{-n} for $n \in \mathbb{N}$ will rather move the points in the vicinity away from A . Thus, the attracting property of A here depends on the choice of a direction of time whereas the invariance is not affected by that. Thus, when we shall discuss Conley theory below, starting in 3.3, we shall encode the essential properties of a dynamical process defined for forward and backward time in collections of invariant sets. There, however, we need to impose some additional restrictions on the invariant sets considered. For example, we wish to distinguish between A and $B(A)$ which, after all, are both invariant. One difference, of course, is that A is required to be compact whereas $B(A)$ is open. However, when F is continuous (which


we usually assume) and if X itself is compact, then also $\overline{B(A)}$ is invariant and compact. Thus, in 3.3 and subsequent sections, we shall require that the invariant sets considered be isolated, that is, possess an open neighborhood whose closure does not contain any other invariant set.

In a related, but different direction, one might wish to refine the definition of an attractor by requiring that it be irreducible; see for example [13] or [41]. This means that the attractor contains a dense orbit, that is, there exists some point $x \in A$ such that, for every other point $y \in A$ and any neighborhood V of y , we can find an $n \in \mathbb{N}$ with $F^n(x) \in V$. In particular, A then is the forward limit set $\omega(x)$ of that orbit. This eliminates superfluous parts. In the time continuous case, such a requirement excludes that an attractor consists of disjoint pieces. A variant of that definition states that we can find a neighborhood U of A such that, for all $z \in U$, the forward orbit is contained in U and the limit set $\omega(z)$ is all of A .

Let us consider some

Examples:

1. (2.31) does not possess any irreducible attractor, because the circle does not contain a dense orbit. In fact, for every θ , the limit set $\omega(\theta)$ is the point $\theta = 0$. That point, however, was seen above not to be an attractor.
2. A saddle x_1 and a sink x_0 with two heteroclinic orbits from x_1 to x_0 for a system of ODEs in the plane. That system of orbits is an attractor in the sense of Definition 3. It is also isolated. It does not contain a dense orbit, however, because all the points x in its vicinity with the exception of the saddle fixed point x_1 will ultimately be attracted by the sink x_0 , that is, $\omega(x) = x_0$ except for $x = x_1$. Thus, it is not an irreducible attractor.



3. This is one of the most fundamental examples in the theory of dynamical systems, see e.g. [24]. We consider the 2-dimensional torus $T := S^1 \times S^1$, that is, the product of two circles and represent it as $\mathbb{R}^2/\mathbb{Z}^2$, that is, as the unit square $Q := \{(x_1, x_2) \in \mathbb{R}^2 : 0 \leq x_1, x_2 \leq 1\}$ with opposite sides identified: $(x_1, 0) \sim (x_1, 1), (0, x_2) \sim (1, x_2)$. For $\omega_1, \omega_2 \in \mathbb{R}$, we consider the flow

$$(2.33) \quad F^t(x_1, x_2) := (x_1 + \omega_1 t, x_2 + \omega_2 t) \pmod{1}.$$


In the square Q , the flow lines are straight lines with slope $\sigma = \frac{\omega_2}{\omega_1}$ (reappearing at the opposite side with the same coordinate value when they hit one of the sides of Q). If σ is rational, $\sigma = \frac{m_2}{m_1}$, then every orbit closes after traversing Q m_2 times in the vertical (x_2) and m_1 times in the horizontal (x_1) direction. Equivalently, the orbit winds m_1 times around the first S^1 factor and m_2 times around the second one before closing, assuming that m_1, m_2 are chosen minimal. In this situation, the torus T itself is an attractor in the sense of Definition 3, but it is not irreducible. The closed orbits, however, are not attractors in either sense, because they do not attract any open neighborhood. They are invariant, of course, but not isolated.

In order to analyze the case where σ is irrational, we consider the intersections of the flow lines with the vertical boundary $I := \{x_1 = 0\}$ (this can be identified with the second S^1 factor). This is an example of a Poincaré return map to be discussed in 2.6. The x_2 coordinate changes by the amount $\sigma \pmod{1}$ between two subsequent crossings of this circle (we identify $x_2 = 0$ and $x_2 = 1$, of course). We therefore obtain an induced discrete time dynamical system R_σ on the circle S^1 given by the sequence of these crossing points. We claim that every orbit of this dynamical system is dense in the circle. If the orbit is not dense, the complement of its closure is a non-empty open invariant set in S^1 . It is then given by a collection of disjoint intervals. Let J be an interval of greatest length in this collection. Its iterates $R_\sigma^n J$ cannot overlap as otherwise their union would be a longer such interval. Furthermore, since σ is irrational, all these iterates must also be distinct because, otherwise, an endpoint p of J would return to itself, $p = p + m\sigma \pmod{1}$, and $m\sigma$ would be an integer, in contradiction to σ being irrational. Therefore, all the intervals $R_\sigma^n J$ are disjoint and of equal length. This makes it impossible to fit all of them into the circle S^1 of finite length ($=1$ in the convention adopted here). This contradiction shows that the complement of the closure of an orbit is empty, and therefore that orbit is dense in the circle. Returning to the torus T , since the images of this circle I under our flow F^t cover the whole torus, every orbit of F^t is then dense in T . We conclude that, for irrational σ , the torus T is the only attractor, in any sense, for the flow F^t (in fact, not only in forward time as is the issue here, but also in backward time).

Obviously, the example can be generalized to a flow on an n -dimensional torus T^n . Such a flow with irrational frequencies is called quasiperiodic because all orbits come back arbitrarily close to their initial positions, and in that sense are almost periodic, but never exactly.

In 7.2, we shall even find it useful to require that the periodic orbits be dense in Λ . This turns out to be a technical requirement for structural stability. So far, we have considered an attractor as an object that is dynamically stable, that is, any point on a small perturbation of it will asymptotically be moved back to the attractor under the dynamical iteration. In 7.2, however, we shall rather be concerned with structural stability, that is, whether an attractor persists without changing its qualitative features when the dynamical system itself is slightly perturbed. In that section, we shall study the important class of Axiom A diffeomorphism which include many chaotic systems. While chaos is characterized by an extreme form of dynamical instability, namely small perturbations of the initial values can be amplified to an arbitrarily large degree, those Axiom A diffeomorphisms are structurally stable in a sense to be made precise. Essentially, this means that small smooth perturbations of the parameters defining the system can be compensated by a continuous transformation of the dynamical behavior. Thus, this type of chaotic dynamics is so rich, in a sense because of its dynamical instability, that it contains all possible variations already in itself and is therefore structurally stable.

One might also wish to relax the requirement for an attractor that it attracts an open neighborhood of itself. For example, a homoclinic orbit from a saddle point enclosing a repelling fixed point in a flow in the plane is attracting from the inside, but not from the outside.



Another example arises when we put two fixed points of the type described by (2.31) on the unit circle, each of them attracting from one side, say in the clockwise direction, and repelling at the other side. If we subject the dynamics to small random perturbations, then an orbit will always approach one of the fixed points in the clockwise direction, but when close to that fixed point, the effect of a perturbation may carry it to the other side whence it may start to approach the other fixed point where then a perturbation may

cause the same effect. Thus, the perturbed dynamics will cause orbits to randomly oscillate between the two fixed points. More important examples of such a phenomenon arise for certain chaotic systems. In that sense, Milnor [30] defines an attractor A for a measurable map $F : X \rightarrow X$ of a space X equipped with a measure by requiring that the set $B(A)$ of points attracted, the basin of attraction of A , have positive measure and that there be no smaller $A' \subset A$ with $B(A')$ coinciding with $B(A)$ up to a set of measure 0. Thus, some points in the vicinity of A are allowed to escape from A . The positive measure requirement for the basin of attraction excludes saddle points and the like.

There are then also various concepts of strange and chaotic attractors in the literature. These qualifications might refer to the geometric properties of A or the dynamical properties of F . In the original reference [45], a strange attractor for a semi-flow was meant to be one that is not a fixed point or a limit cycle. Subsequently, it was often required that it should not be the union of finitely many smooth submanifolds of X (usually a differentiable manifold in applications). See also the discussion in 7.2 below. Other authors, however, then rather defined it through sensitive dependence on initial conditions, that is, a dynamical property. Before elaborating on that aspect, it is helpful to look at some examples of chaotic dynamics.

The first example of a system of ODEs with this type of behavior as expressed by a complicated attractor structure was discovered by E. Lorenz. In his system, $x = (x^1, x^2, x^3)$ takes its values in \mathbb{R}^3 , and the equations are

$$(2.34) \quad \begin{aligned} \dot{x}^1 &= 10(x^2 - x^1) \\ \dot{x}^2 &= x^1(28 - x^3) - x^2 \\ \dot{x}^3 &= x^1x^2 - \frac{8}{3}x^3. \end{aligned}$$

The chaotic behavior of this system was discovered through numerical simulations by Lorenz [28] in 1963. A rigorous mathematical demonstration that chaos occurs here has been achieved only quite recently by Tucker [55]. Below, however, we shall see an example of a one-dimensional difference equation where the chaotic behavior can be studied in a more explicit manner.

The system (2.34) is quite simple as the nonlinearities are only quadratic. Another, even simpler system that exhibits chaotic behavior was subsequently discovered by Rössler:

$$(2.35) \quad \begin{aligned} \dot{x}^1 &= -x^2 - x^3 \\ \dot{x}^2 &= x^1 + 0.2x^2 \\ \dot{x}^3 &= 0.2 + x^3(x^1 - 5.7) \end{aligned}$$

Of course, by now many other chaotic systems of ODEs are known. They all have at least three components and, in fact, one can show that chaos cannot occur for one- or two-dimensional systems of ODEs.

For difference equations, i.e.,

$$(2.36) \quad x(t + 1) = F(x(t)),$$


however, chaos can already occur for one component, as in the famous logistic equation

$$(2.37) \quad x(t + 1) = \rho x(t)(1 - x(t)).$$

The behavior here depends crucially on the parameter ρ . We require $0 \leq \rho \leq 4$ so that the unit interval be mapped to itself. This equation always has 0 as a fixed point, but this fixed point becomes unstable for $\rho > 1$. Another fixed point is at $x_1 = 1 - \frac{1}{\rho}$. Its stability depends on $F'(x_1)$, here with $F(x) = \rho x(1 - x)$. As

$$F'(x_1) = -\rho + 2,$$

$|F'(x_1)|$ becomes larger than 1 for $\rho > 3$.




As $F'(x_1) = -1$ for $\rho = 3$, and $F'(x_1) < -1$ for $\rho > 3$, small perturbations of x_1 lead to fluctuations about this fixed point that first increase in size; for larger perturbations, however, the global shape of the function F leads to a decrease of the fluctuations. For a certain size of the perturbation, it then stays constant, and we obtain a period of order 2, i.e. two additional fixed points of the system


$$(2.38) \quad \begin{aligned} x(t + 2) &= F(F(x(t))) \\ &= \rho^2 x(t) - (\rho^2 + \rho^3)x(t)^2 + 2\rho^3 x(t)^3 - \rho^3 x(t)^4. \end{aligned}$$

Remark: From the perspective of physics, it might seem strange that here different powers of ρ are added. In particular, ρ must be a dimensionless

quantity. Thus, the dynamics considered here do not involve dimensional quantities as in traditional physics.



For increasing values of ρ , these new fixed points in turn become unstable, and we obtain a period of order 4. Increasing ρ further, the period keeps doubling, and the intervals of ρ for which a certain period is stable become smaller and smaller until at $\rho_c \approx 3.5700$, chaotic behavior sets in, in the sense that all periodic solutions of period 2^n become unstable. For $\rho > \rho_c$, there exist some locally attracting cycles with periods $2^n k$, for some odd k . $\rho^* \approx 3.8284$ turns out to be the parameter value when a solution of period 3 appears.



We then also get prime periods of the form $3 \cdot 2^n$. In fact, due to the following fundamental theorem of Sharkovsky, period 3 is the ultimate event.

Theorem 1. *We order the positive integers as follows*

$$1 \triangleleft 2 \triangleleft 2^2 \triangleleft 2^3 \triangleleft \dots \triangleleft 2^m \triangleleft \dots \triangleleft 2^k (2n-1) \triangleleft \dots \triangleleft 2^k \cdot 3 \triangleleft \dots \triangleleft 2 \cdot 3 \triangleleft \dots \triangleleft n-1 \triangleleft \dots \triangleleft 9 \triangleleft 7 \triangleleft 5 \triangleleft 3.$$

Let $F : I \rightarrow I$ be a continuous map of the compact interval I into itself. If F has a periodic point of prime period² p , then it also has periodic points for any prime period $q \triangleleft p$.

A proof of Sharkovsky's theorem can be found for example in the book of Katok-Hasselblatt [24].

In order to analyze the behavior of the logistic map, we utilize

Definition 5. Let $F : \mathbb{R} \rightarrow \mathbb{R}$ be smooth. The Lyapunov exponent of the orbit $x(n) = F(x(n-1))$ ($n \in \mathbb{N}$), $x(0) = x$, is

$$(2.39) \quad \lambda(x) := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{\nu=1}^n \log |F'(x(\nu))|,$$

provided that the limit exists.

The quantity $\lambda(x)$ is undefined when the orbit contains a point with $|F'(x(\nu))| = 0$, and one needs some regularization procedure. To a fixed point, one assigns the Lyapunov exponent $-\infty$. A negative Lyapunov exponent indicates that the directions along an orbit are asymptotically contracted, while a positive one means that they are asymptotically expanded. (The numerical computation of Lyapunov exponents has to be conducted with care. In cases where the orbit is expanding one might worry about instabilities because then small errors might accumulate. We shall see in 7.1, however, that generically, the Lyapunov exponent is stable under small perturbations of an orbit.)

Therefore (cf. [2]) we introduce the following

Definition 6. The orbit $x(n) = F(x(n-1))$ is called chaotic if it possesses a positive Lyapunov exponent and if it is not asymptotically periodic, that is, it does not converge to a periodic orbit ($y(n)$ with $y(n+N) = y(n)$ for some N and all n) in the sense that $\lim_{n \rightarrow \infty} |x(n) - y(n)| = 0$.

We shall see in a moment that the Lyapunov exponent of the logistic map with $\rho = 4$ is $\log 2$ for almost all x .

Definition 7. The limit set $\omega(x)$ of a chaotic orbit $F^n(x)$ is called a chaotic set if it contains x . A Milnor attractor is a chaotic attractor if it is a chaotic set.

Thus, while a strange attractor has been defined in [45] through its irregular geometric properties, a chaotic attractor is characterized solely by dynamical properties. Thus, it may well be geometrically regular as the following **example** shows:

² Prime period means that the orbit is not a multiple of another one with a smaller period (which then has to be a divisor of the original one).

$$F : \mathbb{C} \rightarrow \mathbb{C}$$


$$(r, \theta) \mapsto (r^{1/2}, 2\theta) \text{ in polar coordinates } z = re^{2\pi i\theta}.$$

Here, the unit circle has dense orbits that are not asymptotically periodic with Lyapunov exponent $\log 2$. Thus, it is a chaotic set. Except for the origin $z = 0$ which is a fixed point, all other orbits are attracted to the unit circle and also possess the same Lyapunov exponent. Thus, the unit circle is a chaotic attractor. It is geometrically completely regular, and therefore not strange.

While the Lyapunov exponents for the general logistic map $F(x) = \rho x(1-x)$ are not readily computed explicitly, we can obtain the exponent for the case $\rho = 4$ by conjugating that logistic map to the tent map

$$Z : [0, 1] \rightarrow [0, 1],$$

$$Z(x) = \begin{cases} 2x & \text{for } 0 \leq x \leq \frac{1}{2} \\ 2(1-x) & \text{for } \frac{1}{2} \leq x \leq 1. \end{cases}$$



For the tent map, we have $|Z'(x)| = 2$ for all $x \neq 1/2$, and therefore almost every orbit has the Lyapunov exponent $\log 2$. The logistic map for $\rho = 4$ is conjugated to the tent map via $\sin^2(\frac{\pi x}{2})$ ($F(\sin^2(\frac{\pi x}{2})) = \sin^2(\frac{\pi Z(x)}{2})$). Letting ϕ be the inverse of that map, we have $\phi(F(x(\nu))) = Z(\phi(x(\nu)))$, and thus $F'(x(\nu)) = Z'(\phi(x(\nu))) \frac{\phi'(x(\nu))}{\phi'(F(x(\nu)))}$. Hence, using $F(x(\nu)) = x(\nu + 1)$, the Lyapunov exponent is for almost every orbit

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{\nu=0}^{n-1} \log |F'(x(\nu))| \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_0^{n-1} (\log |Z'(\phi(x(\nu)))| + \log |\phi'(F(x(\nu)))| - \log |\phi'(F(x(\nu+1)))|) \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_0^{n-1} (\log |Z'(\phi(x(\nu)))| + \log |\phi'(F(x(0)))| - \log |\phi'(F(x(n)))|) \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_1^n \log |Z'(\phi(x(\nu)))| \\
&= \log 2,
\end{aligned}$$

as promised.

We shall discuss general results about Lyapunov exponents in 7.1.

The preceding example of a difference equation (namely, the logistic equation) with chaotic behavior of solutions is much simpler than those of systems of ODEs leading to such a behavior. In fact, the solutions of a single ODE cannot exhibit such complex behavior. The reason is simply that for a solution of

$$(2.40) \quad \dot{x}(t) = F(x(t)),$$

x being real valued, a solution can asymptotically grow or decay to a fixed point or $\pm\infty$, but it cannot turn back. Along any non-constant trajectory, $\dot{x}(t)$ is either always positive or always negative. As soon as $\dot{x}(t)$ is 0, we are at a fixed point, and the solution stays there for ever. In fact, such a fixed point cannot even be reached in finite time from any other point provided F is smooth (at least Lipschitz continuous). In particular, (2.40) cannot even have periodic solutions. For example, when compared with the logistic difference equation, the logistic differential equation³

$$\dot{x} = \rho x(1 - x)$$

is trivial to analyze: it has two fixed points, namely 0 and 1. For $\rho > 0$, the first one is unstable as the linearization at 0 is

$$\dot{\xi} = \rho\xi,$$

whereas 1 is stable, since there the linearization is

$$\dot{\xi} = -\rho\xi.$$

³ Of course, the reader will realize that the analogy between the logistic differential and difference equations is not really valid as the latter is not the discretization of the former.

The situation becomes more interesting if we introduce delays and consider

$$(2.41) \quad \dot{x}(t) = \rho x(t)(1 - x(t - \tau)).$$

(As all logistic equations, this can serve as a model for growth with inherent self-limitation, for example because of limited resources. For a positive τ , the limiting factor only acts with some delay.) More generally, one can also consider convolution type models like

$$(2.42) \quad \dot{x}(t) = \rho x(t) \left(1 - \int_{-\infty}^t \eta(t-s)x(s)ds\right)$$

with a weighting factor η , thus taking into account all past values of x , and not only the one at the previous time τ .

We return to the simple equation (2.41); its steady states are $x = 0$ and $x = 1$. We assume that $\rho > 0$ and, by rescaling the time t , we may then assume w.l.o.g. that $\rho = 1$. $x = 0$ is unstable. The linearization about $x = 1$ is

$$(2.43) \quad \dot{\xi}(t) = -\xi(t - \tau).$$

We wish to find solutions of the form

$$(2.44) \quad \xi(t) = ce^{\lambda t}.$$

Inserting this into (2.43) gives the transcendental equation

$$(2.45) \quad \lambda = -e^{-\lambda\tau}$$

Are there solutions of (2.45) with $\operatorname{Re}\lambda > 0$? Such solutions would imply the instability of $x = 1$ as a fixed point of (2.41). In general, the answer to this question will depend on τ .

Clearly, if λ is a real solution of (2.45), it has to be negative. If however λ also has an imaginary part,

$$\lambda = \mu + i\omega,$$

from (2.45) we obtain

$$(2.46) \quad \mu = -e^{-\mu\tau} \cos \omega\tau$$

and

$$(2.47) \quad \omega = e^{-\mu\tau} \sin \omega\tau.$$

We may assume here that $\omega > 0$ because, if $\mu + i\omega$ is a solution, so is $\mu - i\omega$. (2.46) admits a solution with $\mu > 0$ precisely if $\cos \omega\tau < 0$. In other words, if $\omega\tau < \pi/2$, then we have $\mu < 0$, i.e. stability. In that case, (2.47) gives

$$e^{-\mu\tau} \sin(\omega\tau) \tau < \pi/2.$$

The bifurcation, i.e. $\mu = 0$, occurs at $\omega\tau = \pi/2$, and $e^{-\mu\tau} \sin(\omega\tau) \tau = \pi/2$. But $e^{-\mu\tau} \sin(\omega\tau) = 1$ in this case, and hence $\tau = \pi/2$ (and so also $\omega = 1$). We want to study the behavior for $\tau = \pi/2 + \varepsilon$, $0 < \varepsilon \ll 1$, more closely. We expand $\omega = 1 + \sigma$ with $|\mu|, |\sigma| \ll 1$, and obtain from the expansions of (2.46) and (2.47)

$$\begin{aligned}\sigma &\approx -\frac{\pi}{2}\mu \\ \mu &\approx \varepsilon + \frac{\pi}{2}\sigma,\end{aligned}$$

and hence

$$\begin{aligned}\mu &\approx \frac{\varepsilon}{1 + \pi^2/4}, \\ \sigma &\approx -\frac{\varepsilon\pi}{2(1 + \pi^2/4)}.\end{aligned}$$

Inserting

$$\begin{aligned}\lambda &= \mu + i(1 + \sigma) \\ &\approx \frac{\varepsilon}{1 + \pi^2/4} + i\left(1 - \frac{\varepsilon\pi}{2(1 + \pi^2/4)}\right)\end{aligned}$$

into (2.44), and recalling the expansion

$$x(t) \approx 1 + \xi(t),$$

we obtain

$$x(t) \approx 1 + \operatorname{Re}\left(c \exp\left(\frac{\varepsilon t}{1 + \pi^2/4}\right) \exp\left(it\left(1 - \frac{\varepsilon\pi}{2(1 + \pi^2/4)}\right)\right)\right).$$

The instability thus approximately grows with oscillations with period

$$\frac{2\pi}{1 - \frac{\varepsilon\pi}{2(1 + \pi^2/4)}} \approx 2\pi,$$

which is four times the critical delay $\tau = \pi/2$. The exponentially growing factor depends on the slow time scale εt . The preceding example is taken from Murray [35].

In contrast to single ODEs, single delay-differential equations can also give rise to chaotic dynamics. For example, this is the case for

$$\dot{x}(t) = -\alpha x(t) + P(x(t - \tau)),$$

where P is a so-called single hump function, i.e. $P(0) = 0$, $P(x) > 0$ for $x > 0$, $P(x) \rightarrow 0$ for $x \rightarrow \infty$, $P'(x) > 0$ for $x < x_0$, $P'(x) < 0$ for $x > x_0$, and so P has a single maximum at x_0 .

There are many textbooks available that present a detailed study of several chaotic dynamical systems, including the relevant numerical and visualization techniques, for example [2], [51]. Techniques and concepts from chaotic dynamics are also quite useful for the analysis of time series; see in particular [23]. Namely, while a stochastic system produces an erratic, entirely random outcome, a dynamical system as considered here, even if it is chaotic, obeys a deterministic evolution rule. Non-linear time series analysis then develops tools to distinguish between stochastic and chaotic processes on the basis of observations and then offers methods for reconstructing chaotic attractors or at least for evaluating some of their properties, such as certain dimensions. In particular, when those dimensions are low, that is, when only few degrees of freedom are involved, then even for a chaotic dynamical system, the underlying determinism allows short-term predictions to be made that can be far more accurate than simple expectation values available for stochastic dynamics. This is one instance where mathematical insights can lead to methods of great commercial value.

An interpretation of the chaos occurring in some of the preceding equations and systems that is useful for our purposes is that differences or fluctuations in the initial conditions are converted into differences in the temporal behavior of the solutions.

2.5 Interaction, or the interplay between concentration or reaction and diffusion

In the preceding section, we discussed single equations, or systems of finitely many equations. We found that the qualitative behavior can depend crucially on whether the time is continuous or discrete, and whether all reactions come into effect simultaneously, or whether there exist time delays. We now wish to study the case of interacting units, the state of each being modeled by one or several equations. Under certain circumstances, it may still be possible to describe the whole system by e.g. a single equation, namely if that equation describes some aggregated or averaged quantity, like a total population in an ecological model. Typically, however, the situation is not spatially homogeneous, and so the spatial distribution of such a quantity cannot be neglected. Thus, our quantities will not only depend on time t , but also on a spatial position x . In other words, we shall need to study equations for quantities $u(x, t)$. As before, we shall consider continuous state values, i.e. u takes its values in \mathbb{R} , or more generally in some \mathbb{R}^n . Of course, it is also possible to consider discrete state values, contained in \mathbb{Z} , or even in $\{0, 1\}$, but we shall return to that issue only later.

Like time, space can also be either discrete or continuous, and we shall explore the various resulting possibilities a little.

While at each individual site, the state will typically be described by some reaction equation, the relation between different sites is often one of diffusion of quantities. In other words, a state quantity aggregated at one site will have the tendency to diffuse to other sites, for examples if it is present at those other sites in smaller amounts. To which other sites it can diffuse from a given site is determined by the spatial connectivity pattern of the system, and that pattern often plays a crucial role in determining the qualitative behavior of the resulting dynamic process.

In the continuous case, the simplest diffusion process is described by the heat equation:

$$(2.48) \quad u_t(x, t) := \frac{\partial u}{\partial t}(x, t) = D\Delta u(x, t) := D \sum_{i=1}^d \frac{\partial^2 u}{(\partial x^i)^2},$$

with $u : \Omega \times [0, \infty) \rightarrow \mathbb{R}$, Ω some domain in \mathbb{R}^d , and D being a positive constant, the so-called diffusion coefficient.

If we discretize space and time, with step sizes h and k , resp., we obtain the difference equation

$$(2.49) \quad \begin{aligned} & \frac{1}{k}(u^{h,k}(x, t+k) - u^{h,k}(x, t)) \\ &= \frac{D}{h^2} \sum_i \{u^{h,k}(x^1, \dots, x^{i-1}, x^i + h, \dots, x^d, t) \\ & \quad + u^{h,k}(x^1, \dots, x^{i-1}, x^i - h, \dots, x^d, t) - 2u^{h,k}(x^1, \dots, x^i, \dots, x^d, t)\}. \end{aligned}$$

The underlying intuition is that the quantity $u(x, t)$ diffuses from x uniformly to all directions at a rate D . Of course, the value of u at x at the next time step then results from the balance between the amount diffusing out of x to its neighbors and the amount coming to x from its neighbors. In particular, if the value of u at x at time t exceeds the values at the neighboring sites, then it will be decreased. More generally, the maximum principle is valid, saying that the maximum of u w.r.t. the spatial position x is non-increasing in time t , and similarly the minimum is nondecreasing. This, together with the linear structure of the equation, entails a smoothing of the solution u as t increases.

If $h^2 = 2dDk$, (2.49) becomes even simpler:

$$(2.50) \quad u(x, t+k) = \frac{1}{2d} \sum_{i=1}^d (u(x^1, \dots, x^i + h, \dots, x^d, t) + u(x^1, \dots, x^i - h, \dots, x^d, t)).$$

Thus, $u(x, t + k)$ is the arithmetic mean of the values of u at time t at the $2d$ spatial neighbors of x . (2.50) simply is an iterated local averaging process. More generally, one may consider some kernel $h(x, y)$ with

$$(2.51) \quad \sum_y h(x, y) = 1 \text{ for all } x$$

and put

$$(2.52) \quad u(x, t + k) = \sum_y h(x, y)u(y, t).$$

Of course, in the situation of a continuous space, the sums should be replaced by integrals. In any case, this formulation also naturally applies to nonuniform spatial structures, e.g. where the number of neighbors varies with x , or where the neighbors may carry different weights. Also, h need not be nonnegative here. For example, it might be positive for close neighbors y of x , and negative for more distant ones, thereby modeling short-range excitation and long-range inhibition (this point will be taken up only at some later stage). If h is always nonnegative, then (2.52) describes some iterated weighted local averaging. This clearly cannot lead by itself to any complex patterns, as it smoothes out local deviations and thereby destroys information. Thus, in a certain sense, it has the opposite effect to that of the chaotic dynamics described in the preceding section that destroy structures by bringing ever more information to the fore. One might then speculate that some interplay of diffusion and reaction-concentration processes might generate complex structures. Of course, there exist various possibilities for combining these two types of processes. Let us discuss three of them:

1) Reaction-diffusion equations:

Here, the two processes are coupled linearly:

$$(2.53) \quad u_t(x, t) = D\Delta u(x, t) + f(u(x, t)), \text{ with the diffusion coefficient } D > 0,$$

for some function f . Thus, in the absence of diffusion, we have the ODE (ordinary differential equation)

$$(2.54) \quad \dot{u} = f(u).$$

This also holds in the case where u is spatially constant, i.e. independent of x . If that happens initially, i.e. at time $t = 0$, then this persists, at least in the absence of fluctuations. Of course, it remains to study the stability of such a spatially homogeneous state. An important aspect of reaction-diffusion equations, however, is that for a sufficiently large diffusion rate D (depending on the size of the domain, the boundary conditions that might be supposed, and on f), the solution asymptotically (i.e. as $t \rightarrow \infty$) tends to become spatially homogeneous. Of course, if, conversely, D is very

small, then the evolutions of the states of the individual sites become more independent of each other. The preceding model directly extends to the case where u is vector instead of scalar valued. Since, as we have explained, a single ODE like (2.54) cannot exhibit a qualitatively rich behavior, this is an important aspect. For example, for the reaction part, one might take the Lorenz system, or the Hodgkin-Huxley equations for the activity in space-clamped neurons, or any chemical reaction model.

2) Coupled map lattices

Here, the diffusion is superimposed onto the reaction. The original model is discrete in space and time: if the space variable runs through the integers \mathbb{Z} , we have

$$u(n, t + 1) = \frac{\alpha}{2}(f(u(n - 1, t)) + f(u(n + 1, t))) + (-\alpha + 1)f(u(n, t)).$$

Here $\alpha > 0$ is the diffusion constant, and in many applications it is assumed to be small. For a more general spatial structure, the appropriate generalization is

$$(2.55) \quad u(x_i, t + 1) = \frac{\alpha}{\# \text{ neighbors of } x_i} \sum_{x_j \text{ neighbor of } x_i} (f(u(x_j, t)) - f(u(x_i, t)) + f(u(x_i, t))).$$

Of course, as before, one may also employ a general neighborhood interaction function $h(x_i, x_j)$.

An analogous model for a continuous space-time is

$$u_t(x, t) = \alpha \Delta f(u(x, t)) + f(u(x, t)).$$

Since

$$\Delta f \circ u = f'(u) \Delta u + f''(u) \sum_{i=1}^d \left(\frac{\partial u}{\partial x_i} \right)^2,$$

we first of all obtain a nonlinearity that is quadratic in the first derivatives of the solution u , and secondly, the equation changes its type from a forward to a backward heat equation when $f'(u)$ becomes negative. In that case, the smoothing properties of the Laplace operator (see [22]) do not apply anymore to the solution u . While the quadratic nonlinearity can be handled in the scalar case, it can lead to the formation of singularities if u is vector valued.

3) In our last model, we change the order of reaction and diffusion in the coupled map lattice and superimpose a reaction dynamics onto a diffusion process. In the discrete case, the resulting model is equivalent to the coupled map lattice in 2) and simply follows from putting $z = f(u)$ in the latter to obtain

$$z(n, t + 1) = f\left(\frac{\alpha}{2}(z(n - 1, t) + z(n + 1, t)) + (-\alpha + 1)z(n, t)\right)$$

or, for a more general spatial structure again,

$$z(x_i, t + 1) = f\left(\frac{\alpha}{\# \text{ neighbors of } x_i} \sum_{x_j \text{ neighbor of } x_i} (z(x_j, t) - z(x_i, t)) + z(x_i, t)\right).$$

For a continuous space-time, however, the models are no longer equivalent. Here, we get the equation

$$z_t(x, t) = f(\alpha \Delta z(x, t) + z(x, t)).$$

Even in the discrete case, however, we should note that as f in general is not invertible, the relation $z = f(u)$ does not imply a bijection between the solutions of the two models.

The space and time discrete version of this model is also well suited for the situation where the state values are discrete as well. One simply needs to take a function f that only assumes discrete (e.g. integer) values.

We return to (2.55), that is,

$$(2.56) \quad u(x_i, t + 1) = \frac{\alpha}{\# \text{ neighbors of } x_i} \sum_{x_j \text{ neighbor of } x_i} (f(u(x_j, t)) - f(u(x_i, t)) + f(u(x_i, t))).$$

We may consider this as a system of difference equations where the individual equations can be labeled by the index i , that is, we could write it in the form

$$(2.57) \quad u^i(t + 1) = F^i(u^1(t), \dots, u^n(t)).$$

We may also consider a continuous time analogue,

$$(2.58) \quad \dot{u}^i(t) = \Phi^i(u^1(t), \dots, u^n(t)).$$

This is a standard system of coupled ordinary differential equations, and its solution is given by a flow in n -dimensional space. This means that we consider $u = (u^1, \dots, u^n)$ as a map from our time interval, typically \mathbb{R}^+ , into n -dimensional (Euclidean) space. This is the standard setting for the theory of differential equations. For certain purposes, or in certain situations, however, it is useful to develop a different interpretation. Namely, the mappings $F = (F^1, \dots, F^n)$ or $\Phi = (\Phi^1, \dots, \Phi^n)$ may have only sparse entries, that is, n might be quite large while for example F^{i_0} , for any particular index i_0 , might only depend on a small number of other components u^{j_1}, \dots, u^{j_m} where the indices j_1, \dots, j_m , including perhaps their number m , depend on i_0 . In other words, for the dynamic evolution of u^{i_0} , only certain specific other components of u directly occur in the dynamical rule. The other indices only play an indirect role for u^{i_0} because they influence – either directly or again indirectly – the evolution of those components of u on which u^{i_0} directly depends. In order to be able to treat that situation, it is convenient to change the conceptual setting somewhat and not consider (2.57) or (2.58) as a single vector valued dynamical system, but rather as a dynamical network with scalar valued individual dynamics. This means the following. We consider a graph Γ whose nodes correspond to the indices i and where we insert an edge from j to i

and call j a neighbor of i when u^j occurs (non-trivially) as an argument of F^i (or Φ^i in the time-continuous case). At each node i , we then have a scalar dynamics given by F^i or Φ^i , with a direct input from those other nodes j that have edges to i . In that situation, the evolution at node i is coupled to those at its neighbors j . We may then analyze how the coupling structure, that is, the topology of the underlying graph Γ , influences the dynamics. While this can be treated in quite some generality we discuss only some simple aspects here:

- Typically, the evolution of u^i depends in particular on the state of u^i itself; thus, the graph should carry an edge from i to itself, but that is usually omitted.
- In the extreme case when all other components u^j enter into the evolution equation for every u^i , the underlying graph is fully connected, that is, every node is connected with every other one.
- We may consider limits of small coupling, that is, for example systems of the form

$$(2.59) \quad u^i(t+1) = f^i(u^i(t)) + \epsilon \sum_{j \text{ neighbor of } i} f^{ji}(u^j(t))$$

and let the coupling parameter ϵ be small and perhaps tend to 0.

2.6 Discrete and continuous systems. The Poincaré return map

We have already observed that, from a time continuous dynamical system, a time discrete one can be obtained trivially by taking the time-1-map. There exists, however, a different and often more useful such construction, namely the Poincaré return map. We consider a hypersurface S in the state space M of the flow $(x, t) \mapsto f^t(x)$. Each time a trajectory crosses S , we record the position of that intersection, and we thus obtain a sequence of points $\dots, y(n-1), y(n), y(n+1), \dots$ in S . We thus obtain a discrete dynamical system on a state space S whose dimension is one less than the dimension of the state space M for the original flow. This map preserves many properties of that flow but, in general, there may also arise artefacts coming from the choice of S . The situation becomes quite transparent, however, for a periodic orbit Γ of a differentiable flow $(x, t) \mapsto f^t(x)$ on a differentiable manifold M . For an orbit Γ , we consider a transversal hypersurface S in M through a point $y_0 \in \Gamma$. Let y_0 have period T , i.e.,

$$y_0 = f^T(y_0).$$

It can then be shown that for each $x \in S$, sufficiently close to y_0 , there exists $T(x)$ near T with

$$f^{T(x)}(x) \in S.$$

The map

$$P : x \mapsto f^{T(x)}(x)$$

yields then a map of a neighborhood U of y_0 in S to S , the Poincaré return map. P has a fixed point at y_0 and the same regularity properties as the flow itself. P forgets or ignores the behavior of the flow in the direction of the periodic orbit which is trivial anyway and on top of that typically makes the formulation of the hypotheses more complicated. Namely, at least in the autonomous case, the linearization $Df^T(y_0)$ has the eigenvalue 1 in the direction of the periodic orbit and thus does not satisfy a hyperbolicity assumption. P rather encodes the behavior in the transversal directions.

Definition 8. Let $f : M \rightarrow M$ be a continuously differentiable map of a differentiable manifold with a fixed point $y_0 \in M$. y_0 is called hyperbolic if the derivative of f at y_0 , that is, the linear self-map

$$Df(y_0) : T_{y_0}M \rightarrow T_{y_0}M$$

of the tangent space to M at y_0 , is a hyperbolic linear map.

Similarly, if $y_0 \in M$ a periodic point of period n , i.e.

$$f^n(y_0) = y_0,$$

then y_0 is called a hyperbolic periodic point for f if $Df^n(y_0) : T_{y_0}M \rightarrow T_{y_0}M$ is a hyperbolic linear map. A (periodic) orbit then is called a hyperbolic (periodic) orbit if all its points are hyperbolic.

Thus, y_0 is a hyperbolic periodic point precisely if it is a hyperbolic fixed point of f^n .

Definition 9. Let $f^t : M \rightarrow M$ be a differentiable flow with a periodic orbit Γ , and let $y_0 \in \Gamma$. If y_0 is a fixed point, it is called hyperbolic if $Df^t(y_0) : T_{y_0}M \rightarrow T_{y_0}M$ is hyperbolic for all $t \neq 0$. If y_0 is not a fixed point, but periodic with period $T > 0$, it is called a hyperbolic periodic point if $Df^T(y_0) : T_{y_0}M \rightarrow T_{y_0}M$ has 1 as a simple eigenvalue and no further eigenvalue of absolute value 1. The orbit Γ is called hyperbolic if all points on Γ are hyperbolic.

The eigenvalue 1 in this definition of course corresponds to the tangent direction of the orbit. Thus, y_0 is a hyperbolic point for the flow f^t precisely if it is a hyperbolic fixed point of the corresponding Poincaré return map.

2.7 Stability and bifurcations; generic properties

This section develops some general concepts that will help us to understand some examples in subsequent sections. The insights emerging from these concepts are also fundamental for the abstract theory of dynamical systems.

Theorem 2. (Hartman-Grobman)

Let $p \in M$ be a hyperbolic fixed point of the map $f : M \rightarrow M$ as in Definition 8. Then there exist neighborhoods U_1, U_2 of p in M as well as neighborhoods V_1, V_2 of 0 in T_pM , and a homeomorphism

$$h : U_1 \cup U_2 \rightarrow V_1 \cup V_2$$

with the property that the diagram

$$\begin{array}{ccc} U_1 & \xrightarrow{f} & U_2 \\ h \downarrow & & \downarrow h \\ V_1 & \xrightarrow{Df(p)} & V_2 \end{array}$$

commutes.

Thus, f is locally topologically conjugate to its linearization $Df(p)$ at a hyperbolic fixed point, in the sense of the following

Definition 10. Two maps $f : M \rightarrow M$ and $g : N \rightarrow N$ are called topologically conjugate if there exists a homeomorphism $h : M \rightarrow N$ for which the diagram

$$\begin{array}{ccc} M & \xrightarrow{f} & M \\ h \downarrow & & \downarrow h \\ N & \xrightarrow{g} & N \end{array}$$

commutes, that is, if $f = h^{-1} \circ g \circ h$.

Definition 11. Let $1 \leq \ell \leq k \leq \infty$. Then a C^k -map⁴ $f : M \rightarrow M$ is called C^ℓ structurally stable if f possesses a neighborhood \mathcal{U} in the C^ℓ -topology with the property that each $g \in \mathcal{U}$ is topologically conjugate to f .

One should note that two different categories are occurring here, namely smooth and continuous maps, and we are neither requesting that each map in some C^0 -neighborhood of f is topologically conjugate to f , nor that the conjugation can be accomplished by a diffeomorphism. Either of these requirements would be too strong for the results to follow. With the first of these requirements, one would have to admit that isolated fixed points get conjugated into larger pointwise invariant sets. While with the second requirement, the flows corresponding to the systems of differential equations

$$\begin{aligned} \dot{x}^1 &= \alpha x^1 \\ \dot{x}^2 &= \alpha x^2 \end{aligned}$$

⁴ C^k means k times continuously differentiable; C^∞ -maps are infinitely often differentiable. Similarly, C^0 means continuous.

and

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

would no longer be conjugate for $\varepsilon \neq 0$ even in the case $\alpha \neq 0$.

Theorem 3. *Let p be a fixed point of the C^1 -map $f : M \rightarrow M$, and assume that $Df(p) : T_p M \rightarrow T_p M$ has maximal rank. Then f is structurally stable in some neighborhood of p if and only if p is hyperbolic.*

The positive direction of this theorem follows from the Hartman-Grobman theorem which says that, in the hyperbolic case, f is conjugate to its linearization, and two hyperbolic linear maps are conjugate if and only if they have the same number of eigenvalues with absolute value < 1 (and hence also the same number with absolute value > 1) and the determinants have the same sign on the corresponding spaces. These are obviously structurally stable conditions.

The negative direction is not difficult either; one considers non-hyperbolic linear maps that can be perturbed into hyperbolic ones while exhibiting a qualitatively different behavior.

Analogous constructions and results are possible for the time continuous case. Two flows are topologically conjugate if their orbits are transformed into each other by a homeomorphism, allowing a time reparametrization on each orbit. In particular, the period lengths of periodic orbits need not be invariant under topological conjugation. The corresponding Hartman-Grobman theorem for flows states that a flow with a hyperbolic fixed point is conjugate to its linearization at this fixed point.

The essence of the concept of structural stability is that small variations of the parameters defining the dynamical system can be compensated by a continuous transformation of the variables of the system. Thus, a variation of those parameters does not produce a different qualitative behavior, or, to express it somewhat differently, the system is able to reproduce in itself already all the behavior that can be achieved by small perturbations. In particular, the dynamics should not be qualitatively affected by small round-off errors or noise, as in computer simulations or measurements, although this is of course a subtle issue in practice. We shall see in 7.2 that structural stability needs a somewhat particular type of dynamical behavior, namely some form of global hyperbolicity. This has implications for the significance of the concept of structural stability, which, however, are not yet fully explored.

We now introduce concepts for situations that are not structurally stable. There, a small change of the dynamical systems, for example caused by the variation of some parameter, leads to a qualitative change of the behavior, a so-called bifurcation. In turn, however, these qualitative changes themselves might be structurally stable in the sense that, if we consider the dynamical system and the parameter variation together, the situation is stable under small perturbations.

Definition 12. *Let $(f_\lambda)_{\lambda \in \Lambda}$ (Λ open in \mathbb{R}^n) be a family of maps $f_\lambda : M \rightarrow M$ of class C^∞ , depending on some parameter λ , and twice continuously differentiable with respect to λ . $\lambda_0 \in \Lambda$ is called a bifurcation point if f_{λ_0} is not C^2 -structurally stable. We say that we have a structurally stable n -parameter bifurcation at the bifurcation point λ_0 if, for each family $(g_\lambda)_{\lambda \in \Lambda}$ contained in some sufficiently small C^2 -neighborhood of $(f_\lambda)_{\lambda \in \Lambda}$, there exist $\varepsilon > 0$, a reparametrization $\lambda \mapsto \varphi(\lambda)$, and a family $(h_\lambda)_{\lambda \in \Lambda}$ of local homeomorphisms such that*

$$g_{\varphi(\lambda)} = h_\lambda^{-1} \circ f_\lambda \circ h_\lambda$$

for $|\lambda - \lambda_0| < \varepsilon$.

We have discussed several examples of bifurcations already in 2.3 above. The structural stability depends typically on the number n of parameters. R.Thom has classified all structurally stable bifurcations for small values of n , and he has applied this theory to questions of developmental biology, see [54]. People have also discussed the qualitative explanation of phenomena in the social sciences by such a bifurcation analysis. Often, the exact determination of the bifurcation parameters is difficult here. The advantage of this theory, also called catastrophe theory in this context, however rests on the fact that it is not necessary to know the precise functional dependency of the observed quantities on the causal ones for obtaining a qualitative description of possible types of transitions, pattern changes, etc. At least, the theory offers model situations that can be analyzed exactly and are stable under perturbations. These can then serve as analogies for gaining some understanding of phenomena in the realm of the social sciences that cannot be captured themselves with complete detail and precision. Sometimes, however, the theory is applied too naively, and qualitative analogies get confounded with causal explanations.

One would like to have a notion of a property P being typical for a certain parameter-dependent class of dynamical systems. The starting point would be the requirement that i) for every λ_0 in the parameter space Λ , there exists an arbitrary small perturbation λ of λ_0 in Λ such that for the parameter value λ , P holds, and that ii) if P holds for $\lambda_0 \in \Lambda$, it so does for every sufficiently small perturbation λ of λ_0 .

In other words, we require that the subset $A(P)$ of Λ of those parameters for which P holds be dense and open in Λ .

Actually, it turns out to be practical to weaken this requirement a little:

Definition 13. *A subset Λ_0 of a complete metric space Λ is called generic if it contains a countable intersection of open and dense subsets of Λ , and elements of Λ_0 in this case are called generic parameter values if Λ is considered as a space of parameters.*

Thus, instead of asking for an open and dense set, we only ask for a countable intersection of such sets. The reason behind this is Baire's theorem stating that, in a complete metric space, any countable intersection of open and dense sets is dense.

2.8 The Hopf bifurcation

Recalling an example already discussed in 2.3, we wish to investigate an important bifurcation, the (Andronov-)Hopf bifurcation.

Here, when varying some real parameter a periodic orbit emerges at the transition from an attracting to a repelling fixed point. As an example, consider the following system of ODEs in \mathbb{R}^2 ,

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

depending on a real parameter $\alpha \in \mathbb{R}$.

For all α , $(0, 0)$ is a fixed point. For $\alpha < 0$, this fixed point is globally exponentially attracting. This is seen by considering the Lyapunov function $\log(x^2 + y^2)$ (a Lyapunov function by definition is a function that is strictly decreasing along every flow line):

$$\frac{d}{dt} \log(x^2 + y^2) = 2(-x^2 - y^2 + \alpha) \leq 2\alpha < 0.$$

Thus, $\log(x^2 + y^2)$ decreases along every flow line, and then so does $x^2 + y^2$, and therefore each flow line has to lead to $(0, 0)$. This is a structurally stable situation that is invariant under small perturbations of α .

For $\alpha = 0$, $(0, 0)$ is still globally attracting, but no longer exponentially so. We still have

$$\frac{d}{dt} \log(x^2 + y^2) < 0 \text{ for } (x, y) \neq (0, 0),$$

but this expression is no longer bounded away from 0. The situation at $\alpha = 0$ is no longer structurally stable.

For $\alpha > 0$, $(0, 0)$ is repelling, and there exists a periodic orbit $x^2 + y^2 = \alpha$ that is attracting. Namely, we have

$$\frac{d}{dt} \log(x^2 + y^2) \begin{cases} > 0 & \text{for } x^2 + y^2 < \alpha \\ = 0 & \text{for } x^2 + y^2 = \alpha \\ < 0 & \text{for } x^2 + y^2 > \alpha. \end{cases}$$

Thus, when we are on the circle $x^2 + y^2 = \alpha$, we stay there and since, \dot{x} and \dot{y} do not vanish there, it is a nontrivial periodic orbit. When we are outside or inside that circle, we move towards it.

The situation near $(0, 0)$ is structurally stable.

We thus obtain a family, depending on α , of periodic orbits that emerge from the fixed point at the transition from $\alpha = 0$ to $\alpha > 0$. This family of periodic orbits represents a structurally stable bifurcation, that is, such a family remains under perturbations of the above system.

In order to understand this better, let us consider once more the linearized system

$$\begin{aligned} \dot{x} &= y + \alpha x \\ \dot{y} &= -x + \alpha y. \end{aligned}$$

The eigenvalues are $\alpha \pm i$, with

$$\begin{aligned} \text{imaginary part} &\neq 0, \\ \text{but real part} &= 0 \text{ for } \alpha = 0. \end{aligned}$$

Thus, at the bifurcation value $\alpha = 0$, a pair of complex conjugate nonzero eigenvalues crosses the imaginary axis. This is the characteristic criterion for the Hopf bifurcation.

Here, in the linear system, at $\alpha = 0$ all orbits are periodic, namely circles, about $(0, 0)$, while for $\alpha \neq 0$ there is no periodic orbit at all. Thus, here the whole family of periodic orbits is concentrated at a single parameter value, while when the linear system is perturbed by a higher order term, that family gets distributed among different parameter values. The situation at $\alpha = 0$ itself is not structurally stable while the behavior of the whole family is, namely the emergence of a family of periodic orbits at the transition from an attracting to a repelling fixed point.

In fact, the preceding bifurcation where a stable fixed point continuously changed into a stable periodic orbit was a so-called supercritical Hopf bifurcation. In contrast to this, in a subcritical Hopf bifurcation, an unstable periodic orbit coalesces into a stable fixed point so that the latter becomes repelling and no stable orbit is present anymore in its vicinity when the relevant parameter passes the bifurcation value. Thus, the dynamic behavior undergoes a discontinuous transition. Let us study this again in a concrete example. We consider

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

depending on a real parameter $\alpha \in \mathbb{R}$ as before. We now have

$$(2.60) \quad \frac{d}{dt} \log(x^2 + y^2) = 2(-(x^2 + y^2)^2 + 2(x^2 + y^2) + \alpha).$$

This becomes 0 when

$$x^2 + y^2 = 1 \pm \sqrt{1 + \alpha}.$$

Thus, whenever this value is real and nonnegative, we obtain that $x^2 + y^2$ remains constant along a solution, that is, the orbit is a circle. When α is smaller than -1 , no such solution exists. For $\alpha = -1$, we find precisely one solution whereas, for $-1 < \alpha < 0$, we obtain two solutions, of radii $0 < \rho_1 < \rho_2$, say. The right-hand side of (2.60) is negative for $0 < \rho := \sqrt{x^2 + y^2} < \rho_1$, but positive for $\rho_1 < \rho < \rho_2$ and negative again beyond ρ_2 . Thus, the orbit at ρ_1 is repelling whereas the one at ρ_2 is attracting. When α increases to 0, the repelling periodic orbit at ρ_1 moves into the attracting fixed point at 0. When α then becomes positive, both the repelling periodic orbit and the attracting fixed point disappear, or, more precisely, the latter turns into a repelling fixed point. Only the attracting periodic orbit at ρ_2 remains. The solution of our system of ODEs then has no option but to move away from the no longer attracting fixed point at 0 to the periodic orbit at ρ_2 . This is a subcritical Hopf bifurcation, as already mentioned.

We point out that the linearization at 0 is the same for both examples, the supercritical and the subcritical Hopf bifurcation. The crucial fact is that the linearization possesses a pair of complex conjugate eigenvalues whose real parts vanish at the bifurcation point. In fact, according to the theorem of E.Hopf, this is precisely the criterion for such a bifurcation where a stable fixed point bifurcates into a family of periodic orbits which may be either stable or unstable.

The reader may wonder why no other cancellations occur, or why a fixed point cannot simply change from attracting to repelling without a periodic orbit moving in, or why there is a repelling orbit between the attracting fixed point and the attracting orbit in the situation described for the subcritical Hopf bifurcation, and so on. This will be clarified in the next chapter through so-called Conley theory which states that certain invariants must remain constant even across bifurcations.

2.9 Lotka-Volterra equations

The purpose of this section is to investigate a certain class of systems of ODEs which has found applications in biology and ecology and which exhibits some interesting phenomena.

The general Lotka-Volterra equation for n populations is

$$(2.61) \quad \dot{x}_i = x_i \left(a_i + \sum_{j=1}^n b_{ij} x_j \right) \text{ for } i = 1, \dots, n.$$

Here x_i is the size of the i^{th} population, and a_i is its intrinsic growth or decay rate in the absence of the other populations, while b_{ij} is the strength of the effect that the j^{th} population has on the i^{th} one.

Of course, a_1 is positive (negative) iff x_1 has an inherent tendency to grow (decay), while b_{ij} is positive (negative) iff x_j enhances (inhibits) the growth of x_i , e.g. if population i feeds on (is preyed upon by) population j ; both b_{ij} and b_{ji} are negative if the two corresponding populations compete.

Let us examine one of the simplest cases, the two-dimensional predator-prey model without interspecific competition, given by

$$(2.62) \quad \begin{aligned} \dot{x}_1 &= x_1(a_1 + b_{12}x_2) && (x_1 \text{ is the prey}) \\ \dot{x}_2 &= x_2(a_2 + b_{21}x_1) && (x_2 \text{ is the predator}), \end{aligned}$$

with

- $a_1 > 0$ (the prey-population grows in the absence of predators)
- $a_2 < 0$ (the predator population decays in the absence of prey)
- $b_{12} < 0$ (the prey is fed upon by the predators)
- $b_{21} > 0$ (the presence of prey leads to growth of the predator population).

Of course, we are only interested in solutions satisfying

$$x_i(t) \geq 0 \quad \text{for } i = 1, 2 \text{ and all } t \geq 0.$$

We start with some trivial observations:

$$(x_1, x_2) = (0, 0) \text{ is a fixed point.}$$

This fixed point is a saddle as one sees by linearization. In fact, the x_1 -axis is an orbit where the solution expands according to $x_1(t) = x_1(0)e^{a_1 t}$, $x_2(t) = 0$, whereas the x_2 -axis is a contracting orbit with $x_1(t) = 0$, $x_2(t) = x_2(0)e^{a_2 t}$ as $a_2 < 0$.

Thus, in particular, the positive quadrant $x_1(t) > 0, x_2(t) > 0$ is invariant.

A more interesting fixed point is given by

$$(2.63) \quad \bar{x}_1 = -\frac{a_2}{b_{21}}, \quad \bar{x}_2 = -\frac{a_1}{b_{12}}$$


All the other orbits in the positive quadrant are periodic, circling this fixed point counterclockwise. This easily follows from the observation that, for

$$(2.64) \quad V(x_1, x_2) := b_{21}(\bar{x}_1 \log x_1 - x_1) - b_{12}(\bar{x}_2 \log x_2 - x_2),$$

we have

$$\begin{aligned} \frac{d}{dt}V(x_1(t), x_2(t)) &= -a_2 \frac{\dot{x}_1}{x_1} - b_{21}\dot{x}_1 + a_2 \frac{\dot{x}_2}{x_2} + b_{12}\dot{x}_2 \text{ from (2.63)} \\ &= 0 \text{ from (2.64)} \end{aligned}$$

so that $V(x_1, x_2)$ is a constant of motion. V attains its unique maximum at (\bar{x}_1, \bar{x}_2) , and so the curves $V(x_1, x_2) \equiv \text{constant}$ are circles, that is, closed curves, around this point. The motion on such a circle is counterclockwise because in the case $x_1(t) > \bar{x}_1, x_2(t) > \bar{x}_2$ for example, we have $\dot{x}_1(t) < 0, \dot{x}_2(t) > 0$.



On the line $x_1 = \bar{x}_1$ (isocline), we have $\dot{x}_2(t) = 0$ and, on $x_2 = \bar{x}_2, \dot{x}_1(t) = 0$. Thus, the system yields a periodic oscillation of the prey and predator populations.

If T is the period of an oscillation, we have

$$0 = \log x_1(T) - \log x_1(0) = \int_0^T \frac{d}{dt} \log x_1(t) dt = \int_0^T (a_1 + b_{12}x_2(t)) dt,$$

and hence

$$(2.65) \quad \frac{1}{T} \int_0^T x_2(t) dt = -\frac{a_1}{b_{12}} = \bar{x}_2,$$

and similarly

$$(2.66) \quad \frac{1}{T} \int_0^T x_1(t) dt = -\frac{a_2}{b_{21}} = \bar{x}_1,$$

so that the time averages of the oscillating populations are given by their values at the equilibrium point. In particular, this observation allows us to read off the effect of changing one of the coefficients in (2.61) through external influence on the time averages of the two populations.

The behavior of the preceding system with its family of periodic orbits is not stable under small perturbation. Namely if we also include interspecific competition to formulate the system

$$(2.67) \quad \begin{aligned} \dot{x}_1 &= x_1(a_1 + b_{11}x_1 + b_{12}x_2) \\ \dot{x}_2 &= x_2(a_2 + b_{21}x_1 + b_{22}x_2) \end{aligned}$$

with

$$\begin{aligned} b_{11} &< 0 \text{ (the members of population 1 compete for food)} \\ b_{22} &\leq 0, \end{aligned}$$

the behavior of the system changes.

Besides $(0, 0)$, we now get a second fixed point on the positive x_1 -axis, namely $(-\frac{a_1}{b_{11}}, 0)$. This fixed point is always attractive for x_1 , because in case $x_2(t) = 0$, we are looking at the logistic equation

$$\dot{x}_1(t) = x_1(a_1 + b_{11}x_1) \text{ with } a_1 > 0, b_{11} < 0.$$

Whether it is also attractive for x_2 depends on the sign of $\dot{x}_2(t)$ for $x_2(t)$ small and $x_1(t) = -\frac{a_1}{b_{11}}$, i.e. on whether

$$a_2b_{11} - a_1b_{21} > 0.$$

In that case, there is no other fixed point in the positive quadrant, and in fact for any solution

$$\lim_{t \rightarrow \infty} x_2(t) = 0,$$

so that the predator becomes extinct. If, however,

$$a_2b_{11} - a_1b_{21} < 0,$$

then

$$\begin{aligned} \bar{x}_1 &= \frac{a_2b_{12} - a_1b_{22}}{b_{11}b_{22} - b_{12}b_{21}} > 0 \\ \bar{x}_2 &= \frac{a_1b_{21} - a_2b_{11}}{b_{11}b_{22} - b_{12}b_{21}} > 0 \end{aligned}$$

is a fixed point in the positive quadrant.

With $V(x_1, x_2)$ as in (2.64), we now compute

$$\frac{d}{dt}V(x_1(t), x_2(t)) = -b_{11}b_{21}(\bar{x}_1 - x_1(t))^2 + b_{12}b_{22}(\bar{x}_2 - x_2(t))^2 \geq 0,$$

and in fact this derivative is positive unless $(x_1, x_2) = (\bar{x}_1, \bar{x}_2)$. This means that $V(x_1(t), x_2(t))$ is increasing on every orbit, and equilibrium is possible only at its maximum, namely at the fixed point (\bar{x}_1, \bar{x}_2) . The orbits in the positive quadrant then all spiral counterclockwise towards this fixed point.

2.10 Stable, unstable, and center manifolds

In this section, we want to develop some more precise instruments for uncovering and analyzing non-generic phenomena in dynamical systems. Unfortunately, the established terminology is such that the term *stable* will be employed with a meaning which is different from the one in the preceding. It does not refer to structural stability in the sense of stability under perturbations of the system itself, but rather identifies those directions that are dynamically stable in the sense that they move initial states closer together.

We first consider a C^1 -map

$$f : U \rightarrow \mathbb{R}^d,$$

where U is a neighborhood⁵ of $y_0 \in \mathbb{R}^d$ with a hyperbolic fixed point at y_0 ; thus

$$f(y_0) = y_0,$$

and the derivative

$$Df(y_0)$$

does not have eigenvalues of absolute value 1. As usual, we consider the iterates f^n . The stable manifold $M^s(y_0)$ then consists of all points $x \in U$ that satisfy $f^n x \in U$ for all n and

$$f^n x \rightarrow y_0$$

for $n \rightarrow \infty$. Since we are assuming that we have a hyperbolic fixed point, we may even omit the second condition, provided that U is chosen small enough. Thus, we only have to require that the orbit $f^n x$ does not leave U . Those points that are not attracted by the fixed point y_0 are repelled. The unstable manifold, however, does not consist of all the points that get repelled, but only of those that are asymptotically coming from y_0 in the following sense: $y \in M^u(y_0)$ if there exists a sequence $(x_n)_{n \in \mathbb{N}} \subset U$ with $f x_n = x_{n-1}$ for $n \in \mathbb{N}$ and

$$x_n \rightarrow y_0 \text{ for } n \rightarrow \infty.$$

In the hyperbolic case, the second condition can again be omitted if U is chosen small enough.

In the case when f is a diffeomorphism, the unstable manifold of f is nothing other than the stable manifold of f^{-1} as can be seen from the definitions. In order to be also able to handle the case of non-hyperbolic fixed points later on, it turns out to be expedient to formulate the subsequent statements under the more general condition that, for some $\rho \geq 0$, the linearization of

⁵ Neighborhoods are, of course, always assumed to be open.

the map at the fixed point y_0 does not have an eigenvalue of absolute value ρ . This property is called ρ -pseudo-hyperbolicity.

Theorem 4. (Theorem of Hadamard-Perron on the stable and unstable manifolds):

Let U be a neighborhood of $0 \in \mathbb{R}^d$,

$$f \in C^1(U, \mathbb{R}^d)$$

with

$$f(0) = 0.$$

Let $\rho > 0$.

Assume that the derivative $Df(0)$ does not have an eigenvalue of absolute value ρ . Let V_ρ^s, V_ρ^u be the subspaces of $T_0\mathbb{R}^d \cong \mathbb{R}^d$ corresponding to eigenvalues of absolute value $< \rho$ and $> \rho$, respectively, and let $A^{s,u} = Df(y)|_{V_\rho^{s,u}}$. By elementary linear algebra, we find some norm $\|\cdot\|$ on \mathbb{R}^d satisfying

$$\|A^s\| < \rho, \|(A^u)^{-1}\| < \rho^{-1}.$$

And for $R > 0$, we define

$$U(R) := \{x \in \mathbb{R}^d : \|x\| < R\},$$

$$U^{s,u}(R) := \{x \in V_\rho^{s,u} : \|x\| < R\}.$$

If $\rho \leq 1$ then for sufficiently small $R > 0$, the stable manifold

$$M^s(0) := \bigcap_{n \geq 0} f^{-n}U(\rho^n R) = \{x_0 \in U(R) : f^n x_0 \in U(\rho^n R) \text{ for all } n\}$$

is the graph of a C^1 -map

$$\varphi^s : U^s(R) \rightarrow U^u(R)$$

with

$$\varphi^s(0) = 0, D\varphi^s(0) = 0.$$

If $\rho \geq 1$, then for sufficiently small $R > 0$ the unstable manifold

$$M^u(0) := \bigcap_{n \geq 0} f^n \bigcap_{\nu=0}^n f^{-\nu}U(\rho^{\nu-n} R)$$

$$= \{x_0 \in \mathbb{R}^d : \text{there exists a sequence } (x_n)_{n \in \mathbb{N}} \in \mathbb{R}^d$$

$$\text{with } f x_n = x_{n-1}$$

$$\text{and } x_n \in U(\rho^{-n} R) \text{ for all } n \in \mathbb{N}\}$$

is the graph of a C^1 -map

$$\begin{aligned} \varphi^u : U^u(R) &\rightarrow U^s(R) \\ \text{with } \varphi^u(0) &= 0, D\varphi^u(0) = 0. \end{aligned}$$

If f is of class C^k , then so are the maps $\varphi^{s,u}$.

The preceding results easily generalize to (finite-dimensional) differentiable manifolds, and even to those Banach spaces that carry differentiable functions with compact support.

One method for proving the preceding result is the graph transformation method of Hadamard. For the following sketch, we consider the case $\rho = 1$. For sufficiently small $R > 0$, we consider

$$\mathcal{L}_1 := \{\psi : U^u(R) \rightarrow U^s(R),$$

where ψ is Lipschitz continuous with Lipschitz constant ≤ 1 and $\psi(0) = 0\}$.

For $\psi \in \mathcal{L}_1$, we consider the graph

$$\Gamma(\psi) := \{(x, \psi(x)), x \in U^u(R)\},$$

which is contained in $U(R)$ because of the condition on the Lipschitz constant of ψ . One shows then that f maps \mathcal{L}_1 into itself, in the sense that, for $\psi \in \mathcal{L}_1$, $f\Gamma(\psi) \cap U(R)$ is again the graph of some map from \mathcal{L}_1 which we denote by $f_{\#}\psi$. As a contraction of \mathcal{L}_1 w.r.t. the C^0 -norm, $f_{\#}$ then has a fixed point in \mathcal{L}_1 , and this is the desired map φ^u .

Since f is approximated by its linearization $Df(0)$, $f_{\#}$ compresses the graph of ψ in the stable direction V^s and stretches it in the unstable direction V^u . Therefore, derivatives of ψ are made smaller by $f_{\#}$, and φ^u turns out to be differentiable with $D\varphi^u(0) = 0$. Furthermore, for $(x^1, x^2) \in U^u(R) \times U^s(R)$, the distance from $(x^1, \varphi^u(x^1)) \in U^u(R) \times U^s(R)$ is decreased by f by some factor < 1 . As consequence, $f^n(x^1, x^2)$, when it remains in $U(R)$, converges to some point on the graph of φ^u , and the graph of φ^u can be identified with $\bigcap_{n \geq 0} f^n U(R)$ (or, more precisely, with the expression given above for $M^u(0)$).

The stable and the unstable manifold, and the maps φ^s and φ^u of the preceding theorem as well, are uniquely determined. This is no longer the case in the

Theorem 5. (Theorem on the center-stable and center-unstable manifolds):

Let the assumptions of the preceding theorem continue to hold. In the case $\rho \geq 1$, for sufficiently small $R > 0$ there exists a differentiable map

$$\varphi^{0s} : U^s(R) \rightarrow U^u(R)$$

with $\varphi^{0s}(0) = 0$, $D\varphi^{0s}(0) = 0$, whose graph $M^{0s}(0)$, the center-stable manifold, is locally invariant under f in the sense that

$$\begin{aligned} fM^{0s}(0) \cap U(R) &\subset M^{0s}(0) \\ f^{-1}M^{0s}(0) \cap U(R) &\subset M^{0s}(0). \end{aligned}$$

Moreover, $M^{0s}(0)$ is locally attracting for f^{-1} in the sense that, for $x_n \in \bigcap_{\nu=0}^n f^{-\nu}U(R)$, the distance between x_n and $M^{0s}(0)$ converges to 0 as $n \rightarrow \infty$.

In the case $\rho \leq 1$, for sufficiently small $R > 0$, there exists a differentiable map

$$\varphi^{0u} : U^u(R) \rightarrow U^s(R)$$

with $\varphi^{0u}(0) = 0$, $D\varphi^{0u}(0) = 0$, whose graph $M^{0u}(0)$, the center-unstable manifold, is locally invariant under f in the sense that

$$fM^{0u}(0) \cap U(R) \subset M^{0u}(0).$$

Moreover, $M^{0u}(0)$ is locally attracting for f in the sense that for

$$x_n \in \bigcap_{\nu=0}^n f^{-\nu}U(R)$$

the distance between $f^n x_n$ and $M^{0u}(0)$ converges to 0 for $n \rightarrow \infty$.

$M^{0s}(0)$ contains the stable manifold $M^s(0)$, and $M^{0u}(0)$ the unstable manifold $M^u(0)$.

The stable manifold corresponds to those eigenvalues of the linearization whose absolute value is < 1 , the center-stable one to those of absolute value ≤ 1 . Analogously statements are valid for the unstable and the center-unstable manifolds (eigenvalues > 1 or ≥ 1 , respectively). The dynamics inside the stable manifold is exponentially contracting, that in the unstable manifold exponentially expanding. Both dynamics can be described thus in a clear and simple manner. Both of them correspond to a hyperbolic behavior and, at least locally, their analysis is entirely trivial. Thus, in the present perspective, the preceding considerations about hyperbolic aspects only served to pave the way for concentrating on the non-trivial ones as encoded in the intersection

$$M^0(0) = M^{0s}(0) \cap M^{0u}(0),$$

called the center manifold. It is tangential to that subspace of $T_0\mathbb{R}^d \cong \mathbb{R}^d$ which belongs to eigenvalues of $Df(0)$ of absolute value 1. In contrast to the

stable and unstable directions, the dynamics on the center manifold cannot be described by some simple general statements. It rather reflects the individual aspects of the dynamical system under consideration.

If V^0, V' are subspaces of $T_0\mathbb{R}^d$ which correspond to the eigenvalues of absolute value 1 or $\neq 1$, respectively, and if $U^0(R) := \{x \in V^0 : \|x\| < R\}, U' := \{x \in V' : \|x\| < R\}$, then, in the situation of the preceding theorem, the center manifold is given by a C^1 -map

$$\varphi^0 : U^0(R) \rightarrow U'(R)$$

with $\varphi^0(0) = 0, D\varphi^0(0) = 0$ whose graph is locally invariant for the dynamics. A more detailed explanation of the preceding concepts and ideas can be found in the booklet of Ruelle [16], for example.

As indicated, the preceding results also hold in (certain) Banach spaces, i.e., in the infinite-dimensional case. It often happens, however, even in such infinite-dimensional situations that the dimension of the center manifold is finite. This means that the dynamics can be characterized by only finitely many degrees of freedom. This corresponds to the enslaving principle of synergetics, a theory of H. Haken [14, 15]. In any case, it is of fundamental importance that typically there are only very few directions corresponding to the special value 1 of the parameter ρ measuring the local expansion or contraction rates. On a sufficiently long time scale, these few directions then capture all the important and non-trivial aspects of the dynamical process as the contracting directions relax to equilibrium already on a shorter time scale, while the expanding ones move out of sight on that time scale. Therefore, the essential dynamics is dominated or governed by few degrees of freedom and, consequently, a simplified model can still represent the main features of the underlying process. This insight from the theory of dynamical systems thus has implications far beyond the mathematical theory and goes to the heart of scientific modeling.

Analogous results hold for time continuous systems

$$(x, t) \mapsto f^t(x) \quad (x \in U \subset \mathbb{R}^d, t \in \mathbb{R}).$$

This is obtained by first applying the preceding results to the time-1-map

$$f(x) := f^1(x)$$

and then noting that the stated invariance properties hold not only for $n \in \mathbb{N}$, but also for $t \geq 0$.

We now wish to explain these results for systems of ordinary differential equations:

$$(2.68) \quad \dot{w} = g(w, \lambda)$$

with $w \in \mathbb{R}^d$ and a parameter λ taken from a neighborhood of $0 \in \mathbb{R}$. Let $g(0, \lambda) = 0$ for all λ . We consider the linearization

$$(2.69) \quad \dot{w} = Aw \text{ with } A := \frac{\partial}{\partial w}g(0, 0).$$

Let d_s, d_u, d_z be the sums of the multiplicities of the eigenvalues of A as solutions of the characteristic equation $\det(A - \mu Id) = 0$, with real part smaller than, larger than, equal to 0, resp., and we identify the corresponding subspaces of $T_0\mathbb{R}^d \cong \mathbb{R}^d$ with $\mathbb{R}^{d_s}, \mathbb{R}^{d_u}, \mathbb{R}^{d_z}$. The dynamics is then conjugate to that of

$$(2.70) \quad \begin{aligned} \dot{x} &= \gamma(x, \lambda) && \text{for } x \in \mathbb{R}^{d_z} \\ \dot{u} &= -u && \text{for } u \in \mathbb{R}^{d_s} \\ \dot{v} &= v && \text{for } v \in \mathbb{R}^{d_u} \end{aligned}$$

with a function γ to be determined from g . In other words, we have chosen local coordinates so that the stable, unstable, and center manifolds become linear subspaces. The dynamics of u and v is totally simple,

$$\begin{aligned} u(t) &= e^{-t}u(0) \\ v(t) &= e^t v(0). \end{aligned}$$

The interesting aspects occur on the center manifold, which has become \mathbb{R}^{d_z} here. Since, however, typically the problem consists in conjugating system (2.68) to system (2.70), we return to (2.68) and, after a linear coordinate transformation, write it as

$$(2.71) \quad \begin{aligned} \dot{x} &= A^0x + h(x, y, \lambda) \\ \dot{y} &= A^1y + k(x, y, \lambda). \end{aligned}$$

We assume first that the parameter λ does not occur (it will be simple to reduce the general case to this one later on), i.e.

$$(2.72) \quad \begin{aligned} \dot{x} &= A^0x + h(x, y) \\ \dot{y} &= A^1y + k(x, y). \end{aligned}$$

Let the decomposition be such that $x \in \mathbb{R}^{d_z}$ and all eigenvalues of A^0 have real part 0, while $y \in \mathbb{R}^{d'}$ ($d' = d_s + d_u$) and all eigenvalues of A^1 have a real part different from 0.

The center manifold is now described by an equation

$$y = \varphi(x).$$

Since this equation has to be flow-invariant, differentiating w.r.t. t , we obtain the equation

$$\begin{aligned}
 (2.73) \quad 0 &= \dot{x}\varphi'(x) - \dot{y} \\
 &= \varphi'(x)(A^0x + h(x, \varphi(x))) - (A^1\varphi(x) + k(x, \varphi(x))) \\
 &=: \Lambda(\varphi)(x) \quad (\text{with } \varphi'(x) = \frac{d}{dx}\varphi(x)).
 \end{aligned}$$

From this equation, we may now derive a Taylor expansion for the center manifold. The dynamics on the center manifold is then given by the equation

$$(2.74) \quad \dot{x} = A^0x + h(x, \varphi(x)).$$

The equation

$$(2.75) \quad y = \varphi(x)$$

is called the slaving equation. Equation (2.74) describes the essential degrees of freedom of the system and is called the reduced equation.

For treating examples, it will be useful to make the remark after (2.73) more precise by the following lemma:

Lemma 1. *Let the functions h and k in (2.72) be of class C^m ($m \in \mathbb{N}$). Then φ in (2.75) is also of that class. Further, if for $\psi \in C^m$*

$$\Lambda(\psi)(x) = O(|x|^m) \text{ as } |x| \rightarrow 0,$$

then

$$|\varphi(x) - \psi(x)| = O(|x|^m) \text{ as } |x| \rightarrow 0$$

The lemma says that an approximate solution $\psi(x)$, that is, one that satisfies (2.75) approximately, also is an approximation of a solution $\varphi(x)$ of the same order. And we only need an approximation of a solution because the qualitative behavior is locally determined by the lowest nontrivial term in an expansion of the reduced equation.

Examples:

1.

$$\begin{aligned}
 \dot{x} &= -xy, \\
 \dot{y} &= -y + x^2 \quad \text{for } x, y \in \mathbb{R}.
 \end{aligned}$$

The linearization at $(0, 0)$ is

$$\begin{aligned}
 \dot{x} &= 0 \\
 \dot{y} &= -y.
 \end{aligned}$$

Thus, the y -direction is stable while the x -direction is central and thus requires a more refined analysis. For the linearized equation $y = \varphi(x)$,

the x -axis is pointwise invariant, but we shall see that for the original nonlinear problem this, of course, no longer holds. By (2.73), the enslaving equation $y = \varphi(x)$ leads to

$$-x\varphi(x)\varphi'(x) + \varphi(x) - x^2 = 0.$$

With $\psi(x) = x^2$, we have

$$A(\psi)(x) = -2x^4$$

and hence by the lemma,

$$y = \varphi(x) = x^2 + O(|x|^4)$$

and, inserting this, we obtain the dynamics on the center manifold

$$\dot{x} = -x^3 + O(|x|^5).$$

Thus, in contrast to the linearized problem, 0 is a global attractor for the original problem.

2.

$$\begin{aligned} \dot{x} &= -x^3 \\ \dot{y} &= -y \quad \text{for } x, y \in \mathbb{R}. \end{aligned}$$

Here, we see directly that 0 is a global attractor. This example demonstrates that the center manifold need not be unique. Namely, for any constants α, β the curve

$$y = \varphi(x; \alpha, \beta) := \begin{cases} -\alpha e^{-\frac{1}{2x^2}} & \text{for } x < 0 \\ 0 & \text{for } x = 0 \\ \beta e^{-\frac{1}{2x^2}} & \text{for } x > 0 \end{cases}$$

is flow-invariant and therefore yields a center manifold. This example also shows that, even for real analytic equations, the center manifolds need not be real analytic. Refinements of this example show that they need not even be of class C^∞ .

We now return to the situation with the parameter λ and simply add the equation

$$\dot{\lambda} = 0.$$

This means that we consider λ on an equal footing with x . In particular, it becomes a central direction. The slaving equation now becomes

$$y(t) = \varphi(x(t), \lambda),$$

and the reduced equation is

$$\dot{x} = A^0 x + h(x, \varphi(x, \lambda), \lambda).$$

3.

$$\begin{aligned}\dot{x} &= \lambda x - x^3 + xy \\ \dot{y} &= -y + y^2 - x^2 \quad (\lambda, x, y \in \mathbb{R})\end{aligned}$$

in the vicinity of $x = y = \lambda = 0$.

The equation for the center manifold, $y = \varphi(x, \lambda)$, now leads to

$$\begin{aligned}0 &= \dot{x} \frac{\partial \varphi}{\partial x} + \dot{\lambda} \frac{\partial \varphi}{\partial \lambda} - \dot{y} \\ &= (\lambda x - x^3 + x\varphi) \frac{\partial \varphi}{\partial x} - (-\varphi + \varphi^2 - x^2) = \Lambda(\varphi)(x, \lambda).\end{aligned}$$

Of course, the fact that $\dot{\lambda} = 0$ made the derivative $\frac{\partial \varphi}{\partial \lambda}$ drop out. With $\psi(x, \lambda) = -x^2$, we have

$$\Lambda(\psi)(x, \lambda) = O((|x| + |\lambda|)^3)$$

and thus by the lemma

$$\varphi(x, \lambda) = -x^2 + O((|x| + |\lambda|)^3) \text{ as } \lambda, x \rightarrow 0.$$

Inserting this into the reduced equation yields

$$\dot{x} = \lambda x - 2x^3 + O((|x| + |\lambda|)^4) \text{ as } \lambda, x \rightarrow 0.$$

Thus, near $\lambda = 0, x = 0$ the structure of the flow is determined by

$$\dot{x} = \lambda x - 2x^3.$$

The bifurcation at $\lambda = 0$, namely the transition from a stable fixed point for $\lambda < 0$ to two stable and one unstable ones for $\lambda > 0$, now takes place on the center manifold, as follows also from the general theory.

It is also useful to make the following observation: As we see from examples 1 and 3, the strategy for finding an approximate solution ψ of the slaving equation is simply to choose ψ to cancel the lowest order term in the implicit slaving equation $\Lambda(\phi) = 0$ and ignore higher order terms.

Conceptually, in the present section, we have separated the general and universal aspects of a dynamical system, as encoded in the stable and unstable directions, from the individual ones that are represented by the center manifold. We have also emphasized the fundamental aspect that these individual, central degrees of freedom capture the dominant features of the dynamics to which the other ones get subordinated. Of course, this is a local picture. On the global scale, hyperbolic dynamical systems, that is, essentially those for which the center manifold is trivial, still present a rich behavior as will be analyzed in more detail in 7.2. The theory of hyperbolic systems is thus that part of the theory of dynamical systems which can capture the universal aspects of the global behavior. By their very nature, the central dynamics do not succumb to such a general theory. Therefore, they need to be modeled and analyzed more individually in specific examples.

3 Discrete invariants of dynamical systems


The theories of Conley and Floer

The guiding principle of this chapter is to associate discrete topological invariants with a dynamical system, so that qualitative differences between systems get expressed in different values of these invariants, as an approach to classifying dynamical systems. Although the ideal that these invariants can always distinguish non-isomorphic systems cannot be achieved here, nevertheless the invariants constructed in this chapter allow us to detect quite subtle distinctions between dynamical systems.

Of course, the theory depends on certain assumptions, and it works best for dynamical systems on compact state spaces. The basic idea was introduced by Morse theory which is concerned with gradient flows on Riemannian manifolds. The theories of Conley and Floer represent far-reaching extensions and refinements of this theory.

3.1 The topology of graphs

In this section, we wish to explain some important ideas in an elementary example. Let Γ be a graph consisting of a set V of vertices and an edge set E . The edges are directed. Therefore, at each vertex, we can distinguish between inward and outward pointing edges.



The **Euler characteristic** of Γ is defined as

$$\chi(\Gamma) := \#V - \#E,$$

i.e. as the number of vertices minus the number of edges. A vertex v is called **critical** if the number of outward edges is different from 1. The **index** of a critical v is defined as

$$i(v) := \begin{cases} 0 & \text{if there is no outward edge} \\ 1 & \text{if there is more than 1 outward edge} \end{cases}$$

Finally, we define the multiplicity of a critical v with $i(v) = 1$ as

$$m(v) := \text{number of outward edges minus 1.}$$

We then have

Theorem 6.

$$\chi(\Gamma) = m_0 - m_1,$$

with

$$m_0 := \#\{v \text{ critical, } i(v) = 0\}$$

$$m_1 := \sum_{\substack{v \text{ crit.} \\ i(v)=1}} m(v)$$

(if the multiplicity happens to be $m(v) = 1$ for all critical v with $i(v) = 1$, then $m_1 = \#\{v \text{ critical, } i(v) = 1\}$).

The content of this easily demonstrated result is that the computation of a global topological invariant of Γ , namely its Euler characteristic, can be reduced to the evaluation of local quantities at certain particular vertices, namely the critical ones. This is also the basic principle of the theories of Morse and Floer.

3.2 Floer homology

We consider a situation that we have already introduced in 2.2, namely gradient flows. Instead of working on \mathbb{R}^d , however, here we work in a more general setting, namely on a differentiable manifold. This differentiable manifold could be viewed as a submanifold of some Euclidean space, and we would then consider a dynamics that is constrained to that submanifold. Even though every (compact) differentiable manifold can indeed be realized as a submanifold of some Euclidean space by the embedding theorem of Whitney, it is usually preferable to adopt an intrinsic point of view, that is, not to make use of such an embedding. In any case, the analysis of 2.2 easily extends to the case of a differentiable manifold as the underlying space.

Let X be a compact differentiable manifold, and let $F : X \rightarrow \mathbb{R}$ be a smooth function. $x_0 \in X$ is called a **critical point** of F if the first derivatives of F vanish at x_0 :

$$DF(x_0) = 0.$$

At a critical point x_0 , the matrix of the second derivatives, the Hessian $D^2F(x_0)$, can be defined independently of the choice of local coordinates,

and x_0 is called **nongenerate** if none of the eigenvalues of $D^2F(x_0)$ vanishes. In that situation, the **index** of x_0 is defined as

$$i(x_0) := \text{number of negative eigenvalues of } D^2F(x_0).$$

Now $DF(x_0)$ is a 1-form, and we introduce a Riemannian metric on X in order to obtain a tangent vector that is dual to $DF(x)$ at every x , and which we likewise denote by $DF(x)$. This will be needed in order to define the gradient flow of F . (In fact, somewhat less is actually needed, namely a bounded smooth vector field $V_F(x)$ on X that vanishes precisely at the critical points of F and for which $DF(x)V_F(x)$, the 1-form DF applied to that vector field, is positive elsewhere.)

We observed already in 2.2 that the dynamical system defined by our gradient flow transforms the space X as the space of initial conditions for the flow lines ($x(0) = x \in X$) asymptotically into the finite (because they are discrete and X is now assumed compact) set of critical points of F . We might thus expect that this set of critical points somehow captures the qualitative aspects of X , that is, its topological invariants. We are now going to see that this is indeed so but, for a deeper understanding, we shall also need to take relations between those critical points into account.

We now assume:

1. All critical points of F are nongenerate.
2. The space of flow lines or orbits for the negative gradient flow, i.e. of the solutions of

$$\dot{x}(t) = -DF(x(t)) \text{ for } t \in \mathbb{R},$$

$$\lim_{t \rightarrow -\infty} x(t) = p, \quad \lim_{t \rightarrow \infty} x(t) = q,$$

between the two critical points p and q , is either empty or of dimension

$$i(p) - i(q).$$

(These conditions are satisfied for generic smooth functions F .)

In 2.2, we have already explored the consequences of the first assumption, and those, of course, continue to hold in our manifold setting. In particular, we always find limits of a flow line $x(t)$ as $t \rightarrow \pm\infty$, and these limits are critical points of F .

The second assumption then implies that there are no flow lines between critical points of the same index while in case

$$i(p) = i(q) + 1$$

the number of flow lines is finite (as each such flow line is 1-dimensional, the space of lines then is 1-dimensional as well). In the latter case, we put

$$\rho(p, q) := (\text{number of flow lines from } p \text{ to } q) \bmod 2$$

Thus, $\rho(p, q)$ is 1 when the number of flow lines from p to q is odd, and it is 0 otherwise.

We now define the boundary operator ∂ : For a critical point p of index $i(p) = i$, let

$$\partial_i p := \sum_{\substack{q \text{ critical} \\ i(q)=i(p)-1}} \rho(p, q)q$$

(the sum here is taken in the vector space generated by the critical points of F over \mathbb{Z}_2 ; for example

$$\begin{aligned} & (a_1q_1 + a_2q_2) + (b_1q_1 + b_2q_2) && (q_1, q_2 \text{ critical points,} \\ & && a_1, a_2, b_1, b_2 \in \mathbb{Z}_2) \\ & = (a_1 + b_1)q_1 + (a_2 + b_2)q_2, \end{aligned}$$

where the sum in brackets takes place in the field \mathbb{Z}_2 (i.e. we are adding mod 2) while the sum between brackets is a formal sum.)

Thus, we are counting flow lines from p to critical points q with $i(q) = i(p) - 1 \pmod{2}$. We then have

Theorem 7. (Floer) $\partial_{i-1} \circ \partial_i = 0$ for all i .


Explanation:

$$\begin{aligned} (\partial_{i-1} \circ \partial_i)p &= \partial_{i-1}(\partial_i p) \\ &= \partial_{i-1} \left(\sum_{i(q)=i(p)-1} \rho(p, q)q \right) \\ &= \sum_{i(q)=i(p)-1} \rho(p, q) \left(\sum_{i(r)=i(q)-1} \rho(q, r)r \right). \end{aligned}$$

The content of this theorem is that the flow lines from p to r with $i(r) = i(p) - 2$ that are broken, i.e. go through another critical point q , always occur in pairs, i.e.

$$\sum_q \rho(p, q)\rho(q, r) = 0 \bmod 2 \text{ for all } p \text{ and } r.$$

On the basis of this theorem, we can define **homology groups**: We have the kernel of ∂_i



$$\ker \partial_i = \{ \Sigma a_\nu p_\nu : i(p_\nu) = i, a_\nu \in \mathbb{Z}_2, \text{ with } \partial(\Sigma a_\nu p_\nu) = 0 \text{ (} = \Sigma a_\nu \partial p_\nu, \partial \text{ has been linearly extended) } \},$$

and the image of ∂_i

$$\text{im } \partial_i = \{ q \text{ with } i(q) = i - 1, \text{ that can be represented as } q = \partial(\Sigma a_\nu p_\nu) \}.$$

(Note that these spaces are vector spaces over \mathbb{Z}_2 and therefore in particular additive groups.) Then $\partial_i \circ \partial_{i+1} = 0$ implies that

$$\text{im } \partial_{i+1} \subset \ker \partial_i$$

(if p is a boundary, then its own boundary is 0).

Definition 14. The i^{th} homology group of X is

$$H_i(X, \mathbb{Z}_2) := \ker \partial_i / \text{im } \partial_{i+1}$$

(all boundaries are set to 0).

Theorem 8. The groups $H_i(X, \mathbb{Z}_2)$ do not depend on the choice of the function $F : X \rightarrow \mathbb{R}$ (which, however, has to satisfy the assumptions 1. and 2. above), and conversely, we can always find such a function F . These groups therefore define topological invariants of the differentiable manifold X .

The i^{th} **Betti number** of X is defined as

$$b_i(X) := \dim H_i(X, \mathbb{Z}_2),$$

and the **Euler characteristic** is

$$\chi(X) := \sum_i (-1)^i b_i(X),$$

i.e. the alternating sum of the Betti numbers.

Of course, the Euler characteristic can also be defined topologically, for example on the basis of a triangulation, and the result then coincides with the above value. As for the Euler characteristic of a graph, we see that it can be computed from local expressions, namely, in the present situation, relations between critical points of index difference 1. Actually, these expressions are not strictly local as they depend on connections between two different points, but we also have

Theorem 9. *Let m_i be the number of critical points of F of index i . Then*

$$m_j \geq b_j \text{ for all } j,$$

and


$$\sum_j (-1)^j m_j = \chi(X).$$

A proof of this result will be given in 3.5 below.

In particular, we can compute a topological invariant of X , its Euler characteristic, from purely local information about the critical points of the function F . In fact, the only such information needed are the indices of the critical points. These indices encode the number of directions in which F decreases, and so the gradient flow moves away, near the critical points. This observation will be the starting point for Conley theory as presented in the next sections. There, instead of gradient flows with isolated (non-degenerate) critical points, one studies arbitrary flows and their local invariant sets and again counts (in a suitable topological sense) the outflowing directions near each such invariant set.

Let us now consider some examples for the computation of the Betti numbers $b_i = \dim H_i(X, \mathbb{Z}_2)$. In these examples, we take X as a submanifold of \mathbb{R}^3 and F as the height function, that is, $F(x)$ is simply the value of the third coordinate x^3 of the point $x = (x^1, x^2, x^3)$.

In the first example, X is diffeomorphic to S^2 :




$$\begin{aligned} \partial p_1 &= q \text{ and } \partial p_2 = q, \\ \partial(p_1 + p_2) &= 2q = 0, \\ \text{and so} \\ \ker \partial_2 &\text{ is generated by } p_1 + p_2, \\ \text{im } \partial_2 &\text{ generated by } q. \end{aligned}$$

$$\begin{aligned} \text{Next, } \partial q &= 2r = 0, \\ \ker \partial_1 &\text{ is generated by } q, \\ \text{im } \partial_1 &= \{0\}, \text{ and} \\ \partial r &= 0, \\ \ker \partial_0 &\text{ is generated by } r. \end{aligned}$$

Thus,

$$\begin{aligned} H_2(X, \mathbb{Z}_2) &= \ker \partial_2 \text{ is 1-dimensional: } b_2 = 1 \\ H_1(X, \mathbb{Z}_2) &= \ker \partial_1 / \text{im } \partial_2 \text{ is 0-dimensional: } b_1 = 0 \\ H_0(X, \mathbb{Z}_2) &= \ker \partial_0 / \text{im } \partial_1 \text{ is 1-dimensional: } b_0 = 1. \end{aligned}$$

In the next example, X is a torus:



$$\begin{aligned} \partial p &= 2q_1 + 2q_2 = 0 \Rightarrow \ker \partial_2 \text{ is generated by } p, \text{ im } \partial_2 = \{0\} \\ \partial q_1 &= \partial q_2 = 2r = 0 \Rightarrow \ker \partial_1 \text{ is generated by } q_1 \text{ and } q_2, \text{ im } \partial_1 = \{0\} \\ \partial r &= 0 \Rightarrow \ker \partial_0 \text{ is generated by } r. \end{aligned}$$

Thus, for the torus,

$$\begin{aligned} H_2(X, \mathbb{Z}_2) &= \ker \partial_2 && \text{is 1-dimensional: } b_2 = 1 \\ H_1(X, \mathbb{Z}_2) &= \ker \partial_1 / \text{im } \partial_2 && \text{is 2-dimensional: } b_1 = 2 \\ H_0(X, \mathbb{Z}_2) &= \ker \partial_0 / \text{im } \partial_1 && \text{is 1-dimensional: } b_0 = 1. \end{aligned}$$

We have used the boundary operator ∂ that relates the critical points of a generic smooth function $F : X \rightarrow \mathbb{R}$ in order to define the homology groups $H_i(X, \mathbb{Z}_2)$. For the subsequent discussion of the theory of Conley, we shall also need relative homology groups. For that purpose, let A be a compact subset of the compact differentiable manifold X , and let $F : X \rightarrow \mathbb{R}$ be a smooth function with the property that flow lines may enter, but not leave, A . This means that, if

$$\begin{aligned} \dot{x}(t) &= -DF(x(t)) \text{ for all } t \in \mathbb{R} \\ \text{and } x(t_0) &\in A \text{ for some } t_0 \in \mathbb{R} \cup \{-\infty\}, \\ \text{then also } x(t) &\in A \text{ for all } t \geq t_0. \end{aligned}$$

We then obtain a boundary operator ∂^A if, in the preceding constructions, we only take those critical points into account that lie in $X \setminus A$. Thus, for a critical point $p \in X \setminus A$, we put

$$\partial^A p := \sum_{\substack{q \text{ crit. pt. in } X \setminus A \\ i(q)=i(p)-1}} \rho(p, q) q.$$

Because of the above condition, all flow lines between the critical points p and q are contained in $X \setminus A$ themselves. With the help of this boundary operator, we then obtain the relative homology groups

$$H_i(X, A; \mathbb{Z}_2) := \ker \partial_i^A / \text{im } \partial_{i+1}^A.$$

For example, when $X = S^2$ and A is a point which we can then take as the minimum point of our function f , we obtain

$$\begin{aligned} b_2(S^2, \text{point}) &= 1 \\ b_1(S^2, \text{point}) &= 0 \\ b_0(S^2, \text{point}) &= 0. \end{aligned}$$

When X is the torus T^2 and A is a circle which we then take to consist of the critical points q_2 and r and the two flow lines between them in the above example, we obtain

$$\begin{aligned} b_2(T^2, S^1) &= 1 \\ b_1(T^2, S^1) &= 1 \\ b_0(T^2, S^1) &= 0, \end{aligned}$$

simply because the circle S^1 that we take out of the torus accounts for one of the two generators of H_1 and the generator of H_0 .

The preceding construction can be generalized as follows. Let $A \subset Y \subset X$ be compact and let $F : X \rightarrow \mathbb{R}$ satisfy:

(i) If for some flow line, i.e.

$$\begin{aligned} \dot{x}(t) &= -DF(x(t)) \quad \text{for all } t \in \mathbb{R}, \\ x(t_0) &\in A \text{ for some } t_0 \in \mathbb{R} \cup \{-\infty\}, \end{aligned}$$

then there is no $t > t_0$ with $x(t) \in Y \setminus A$.

(ii) If for some flow line

$$x(t_1) \in Y, x(t_2) \in X \setminus \overset{\circ}{Y}, \text{ with } -\infty \leq t_1 < t_2 \leq \infty,$$

then there exists $t_1 \leq t_0 \leq t_2$ with

$$x(t_0) \in A.$$

Thus, by (i), flow lines cannot re-enter $Y \setminus A$ from within A , while, by (ii), they can exit Y only through A .

In this situation, for a critical point $p \in Y \setminus A$, we consider the boundary operator

$$\partial^{Y,A} p := \sum_{\substack{q \in Y \setminus A \text{ critical} \\ i(q) = i(p) - 1}} \rho(p, q) q$$

and define

$$H_i(Y, A; \mathbb{Z}_2) := \ker \partial_i^{Y,A} / \text{im } \partial_{i+1}^{Y,A}.$$

Actually, for the discussion of Conley theory, we shall need cohomology in place of homology groups. These are dual to homology groups in the following sense. If $(C_k, \partial_k)_{k \in \mathbb{N}}$ is a chain complex, i.e. the maps

$$\partial_k : C_k \rightarrow C_{k-1}$$

satisfy $\partial_{k-1} \circ \partial_k = 0$ for all k , we put

$$C^k := \text{Hom}(C_k, \mathbb{Z})$$

and define the boundary operators δ^k by

$$\langle \delta^k c^k, \gamma_{k+1} \rangle = \langle c^k, \partial_{k+1} \gamma_{k+1} \rangle$$

for $c^k \in C^k, \gamma_{k+1} \in C_{k+1}$. Then again

$$\delta^{k+1} \circ \delta^k = 0,$$

and the groups

$$H^k := \ker \delta^k / \operatorname{im} \delta^{k+1}$$

are called the cohomology groups of the complex (C^k, δ^k) .

Remark 1. Floer homology can also be defined with \mathbb{Z} in place of \mathbb{Z}_2 coefficients. Then, one needs to introduce orientations in order to assign a sign ± 1 to flow lines connecting critical points of index difference 1. The terms $\rho(p, q)$ in the definition of ∂ can then take positive or negative values if the signed flow lines are counted. Again, one shows that $\partial \circ \partial = 0$, the basis for the introduction of homology groups. Treating the orientations is somewhat subtle, however, and so at this point we refer the reader to Schwarz [47], Jost [21].

For applications in the calculus of variations, the following observation is useful: for the construction of the boundary operator ∂ and therefore for the homology groups, one does not really need to know the indices of critical points. It suffices to know the index difference between critical points, in order to check which flow lines need to be counted. In other words, one only needs a relative index, not an absolute one.

3.3 Conley theory: examples and results

Let X be a metric space, $F : X \times \mathbb{R} \rightarrow X$ a flow (and so, in particular, $F(x, 0) = x$), and put $x(t) := F(x, t)$.

For $N \subset X$, we let

$$I(N) := \{y \in N : y(\mathbb{R}) \subset N\}$$

be the set of points of N that remain in N for all positive and negative times under the flow. $N \subset X$ is called invariant if

$$I(N) = N.$$

The basic and guiding idea of this chapter is that both the topology of the underlying space X and the qualitative properties of the dynamical system defined by F are encoded in its invariant sets. (Of course, such a statement needs some qualifications; for example, it is best if X is compact; otherwise, for example if $X = \mathbb{R}^n$, then one needs some control near infinity, but this will not be made precise here.) Clearly, X itself is invariant, but that does not yield any simplification. However, except under very special circumstances,

there also exist smaller invariant sets, and those contain the essential information. More precisely, in order to have the complete information, we need a collection of such invariant sets satisfying the following properties.

A **Morse decomposition** of a compact invariant set N is a finite collection of disjoint compact invariant subsets N_i of N , $i = 1, \dots, n$, the so-called Morse sets, that permit a so-called admissible ordering (N_1, N_2, \dots, N_n) such that, whenever

$$y \in N \setminus \bigcup_{i=1}^n N_i,$$

there exist indices $i < j$ with

$$\alpha(y) \subset N_i, \omega(y) \subset N_j.$$

Here, $\alpha(y)$ and $\omega(y)$ are the asymptotic limit sets defined at the end of 2.1. For $i \leq j$, $N_{ij} := \{y \in N : \alpha(y), \omega(y) \subset N_i \cup N_{i+1} \cup \dots \cup N_j\}$.

In an admissible ordering of a Morse decomposition, N_{ij} may be exchanged with $\bigcup_{\nu=i}^j N_\nu$, and vice versa, to obtain another Morse decomposition with an admissible ordering. An important point in the sequel will be that this procedure does not affect the basic topological information contained in a Morse decomposition. For a gradient flow as considered in 3.2, with isolated critical points, the finest Morse decomposition, that is, the one with the smallest Morse sets, is simply given by the collection of its critical points, with an ordering such that the initial point of any flow line has a smaller label than the final point. We obtain coarser Morse decompositions by lumping together some critical points with all the orbits between them. This possibility of shifting scales, that is, comparing the topological information at a finer, more localized scale with the one at a coarser, perhaps even global, scale, is one of the key aspects of Conley theory.

We now treat an example in some detail. In \mathbb{R}^2 , we consider the system of differential equations

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= \lambda y + x(x - 1)(x - \mu) \end{aligned}$$

with parameters $\lambda, \mu \in \mathbb{R}$, $\mu > 1$.

For each λ , the fixed points of this system are $(0, 0)$, $(1, 0)$ and $(\mu, 0)$. More detailed properties depend on λ as we shall now analyze.

The linearization at $(0, 0)$ is

$$\begin{aligned} \dot{\xi} &= \eta \\ \dot{\eta} &= \lambda \eta + \mu \xi. \end{aligned}$$

The corresponding matrix

$$A_0 = \begin{pmatrix} 0 & 1 \\ \mu & \lambda \end{pmatrix}$$

has eigenvalues $\alpha_{1,2} = \frac{\lambda}{2} \pm \sqrt{\frac{\lambda^2}{4} + \mu}$, i.e. one positive and one negative real eigenvalue. Therefore, for every λ , $(0, 0)$ is a saddle point. The linearization at $(1, 0)$ is

$$\begin{aligned} \dot{\xi} &= \eta \\ \dot{\eta} &= \lambda\eta - (\mu - 1)\xi. \end{aligned}$$

The corresponding matrix

$$A_1 = \begin{pmatrix} 0 & 1 \\ 1 - \mu & \lambda \end{pmatrix}$$

has eigenvalues $\alpha_{1,2} = \frac{\lambda}{2} \pm \sqrt{\frac{\lambda^2}{4} + 1 - \mu}$. For $\lambda = 0$, these are purely imaginary while for $\lambda^2 \geq 4(\mu - 1)$ they are real and of the same sign. Thus, for $\lambda^2 \geq 4(\mu - 1)$, we find a node (a source for $\lambda \geq 2\sqrt{\mu - 1}$, and a sink for $\lambda \leq -2\sqrt{\mu - 1}$) while, for $0 < \lambda^2 < 4(\mu - 1)$, we have a spiral. The behavior at $\lambda = 0$ requires a more detailed analysis because, as we have already seen, in the case of two purely imaginary eigenvalues, the fixed point is not hyperbolic and it may be possible that the linearization exhibits a behavior that is qualitatively different from that of the original system.

Finally, the linearization at $(\mu, 0)$ is

$$\begin{aligned} \dot{\xi} &= \eta \\ \dot{\eta} &= \lambda\eta + \mu(\mu - 1)\xi \end{aligned}$$

and the situation is qualitatively the same as at $(0, 0)$. Again, we have a saddle point for all λ . We now start with the case $\lambda = 0, \mu = 2$. In this case, we have two symmetries that leave our system invariant:


1.

$$\begin{aligned} t &\rightarrow -t \\ x &\rightarrow x \\ y &\rightarrow -y \end{aligned}$$

2.

$$\begin{aligned} t &\rightarrow -t \\ x &\rightarrow 2 - x \\ y &\rightarrow y. \end{aligned}$$

This implies that, except at the fixed points, the orbits intersect the x -axis as well as the line $x = 1$ orthogonally. Between $x = 0$, and $x = 1$ as well as for $x > 2$, the x -axis is intersected by the orbits from below and, on each such orbit, the intersection with the x -axis occurs at its smallest x -value. In the other regions, the opposite behavior occurs. For $y > 0$, the line $x = 1$ is intersected by the orbits from the left, and that intersection point corresponds to a local maximum of the y -coordinate, again with the opposite behavior for $y < 0$.




These facts and symmetry arguments imply that there exists an orbit above the x -axis from the fixed point at $(0,0)$ to the one at $(2,0)$, and an orbit below the x -axis from $(2,0)$ to $(0,0)$. (Such orbits that connect two different fixed points are called heteroclinic.) Between these two orbits, we find closed orbits about the fixed point at $(1,0)$. The orbits that lie outside this region are unbounded.

We now consider the case $\lambda = 0, \mu > 2$. (For $\lambda = 0, 1 < \mu < 2$, we find the same qualitative behavior after exchanging the fixed points at $(0,0)$ and $(\mu,0)$.) Symmetry 1 persists, but symmetry 2 disappears.


Instead of the two heteroclinic orbits found for $\mu = 2$, we obtain for $\mu > 2$ a homoclinic orbit that emanates from the fixed point at $(0,0)$, intersects the x -axis between $x = 1$ and $x = \mu$ and then returns to $(0,0)$ as follows from symmetry arguments. Inside that homoclinic orbit, we again find a family of periodic orbits around the fixed point at $(1,0)$. All other orbits are unbounded.

It is remarkable that, although for $\lambda = 0$ the fixed point at $(1,0)$ is non-hyperbolic, the behavior of the orbits in the vicinity of this point is the same as for the linearized system.


We next consider the case $\mu > 2, \lambda > 0$; the case $\lambda < 0$ is symmetric to that one. As λ grows, the point where the orbit that starts at $(0,0)$ and enters the first quadrant intersects the x -axis moves to the right. Also, the symmetry 1 disappears. The fixed point at $(1,0)$ becomes hyperbolic and, for



$0 < \lambda < 2\sqrt{\mu-1}$, we obtain a spiral. One of the spiral curves starting at $(1, 0)$ goes into the fixed point at $(0, 0)$. All other orbits are unbounded.




If λ reaches a critical value λ_μ , in addition to the heteroclinic orbit from $(1, 0)$ to $(0, 0)$, another heteroclinic orbit from $(0, 0)$ to $(\mu, 0)$ appears. Again, all other orbits are unbounded.



If λ grows further, we finally obtain a heteroclinic orbit from $(1, 0)$ to $(0, 0)$ and another from $(1, 0)$ to $(\mu, 0)$. Again, all other orbits are unbounded. In case $\mu = 2$, this is so for all $\lambda > 0$.

Thus, for $\lambda = 0$ and $\mu = 2$, the maximal compact invariant set N consists of the two heteroclinic orbits and the region they enclose, which in turn consists



of periodic orbits and the fixed point in the middle. There is no nontrivial Morse decomposition of N .

For $\lambda = 0$ and $\mu > 2$, such an N consists of two components N' and N'' . N' contains the homoclinic orbit corresponding to $(0, 0)$ and as before the enclosed set of periodic orbits including the fixed point $(1, 0)$, while N'' consists of the isolated fixed point $(\mu, 0)$ only. Nontrivial Morse decompositions are given by $N_1 = N', N_2 = N''$ and $N_1 = N'', N_2 = N'$.

For $\mu > 2$ and $0 < \lambda < \lambda_\mu$, N consists of the heteroclinic orbit C from $(1, 0)$ to $(0, 0)$ as well as the point $(\mu, 0)$. The nontrivial Morse decompositions are

$$\begin{aligned} N_1 &= \{C\}, N_2 = \{(\mu, 0)\} \\ N_1 &= \{(\mu, 0)\}, N_2 = \{C\} \\ N_1 &= \{(\mu, 0)\}, N_2 = \{(1, 0)\}, N_3 = \{(0, 0)\} \\ N_1 &= \{(1, 0)\}, N_2 = \{(0, 0)\}, N_3 = \{(\mu, 0)\} \\ N_1 &= \{(1, 0)\}, N_2 = \{(\mu, 0)\}, N_3 = \{(0, 0)\} \\ N_1 &= \{(\mu, 0), (1, 0)\}, N_2 = \{(0, 0)\} \\ N_1 &= \{(1, 0)\}, N_2 = \{(\mu, 0), (0, 0)\}. \end{aligned}$$


For $\mu > 2$ and $\lambda = \lambda_\mu$, N consists of the heteroclinic orbits C_1 from $(1, 0)$ to $(0, 0)$ and C_2 from $(0, 0)$ to $(\mu, 0)$. The nontrivial Morse decompositions are:

$$\begin{aligned} N_1 &= \{C_1\}, N_2 = \{(\mu, 0)\} \\ N_1 &= \{(1, 0)\}, N_2 = \{C_2\} \\ N_1 &= \{(1, 0)\}, N_2 = \{(0, 0)\}, N_3 = \{(\mu, 0)\}. \end{aligned}$$

Finally, for $\mu > 2$ and $\lambda > \lambda_\mu$, N consists of the two heteroclinic orbits C_1, C_2 , from $(1, 0)$ to $(0, 0)$ and $(\mu, 0)$, and the nontrivial Morse decompositions are:

$$\begin{aligned} N_1 &= \{(1, 0)\}, N_2 = \{(0, 0)\}, N_3 = \{(\mu, 0)\} \\ N_1 &= \{(1, 0)\}, N_2 = \{(\mu, 0)\}, N_3 = \{(0, 0)\} \\ N_1 &= \{(1, 0)\}, N_2 = \{(\mu, 0), (0, 0)\} \\ N_1 &= \{C_1\}, N_2 = \{(\mu, 0)\} \\ N_1 &= \{(0, 0)\}, N_2 = \{C_2\}. \end{aligned}$$

We now show an illustration of the Morse decompositions for $\mu = 3$:



In order to extract the desired qualitative information, namely our invariants, from a Morse decomposition, we now need to introduce the fundamental concept of an index pair. A compact $N \subset X$ is called an **isolating neighborhood**, and $I \subset N$ is called an **isolated invariant set** if

$$I(N) \subset \overset{\circ}{N} \quad (= \text{the interior of } N) .$$

In other words, $S \subset X$ is an isolated invariant set if it is invariant and possesses an open neighborhood whose closure contains no other invariant sets besides S .

A pair (M_0, M_1) of compact sets $M_1 \subset M_0$ is called an **index pair** for the isolated invariant set S if

- (i) $\overline{M_0 \setminus M_1}$ is an isolating neighborhood of S with $M_1 \cap S = \emptyset$.
- (ii) If $y \in M_1, y([0, t]) \subset M_0$, then $y([0, t]) \subset M_1$
- (iii) If $y \in M_0$ and $y(t) \notin M_0$ for some $t > 0$, then there exists $0 \leq t_0 < t$ with $y([0, t_0]) \subset M_0, y(t_0) \in M_1$.

M_1 is the **exit set** from M_0 for the flow. (ii) means that the flow cannot return to M_0 from M_1 , and (iii) means that a flow line exiting M_0 has to pass through M_1 .

Theorem 10. *Let S be an isolated invariant set, (N_1, \dots, N_n) an admissible ordering of a Morse decomposition of S . Then there exists a **Morse filtration**, i.e. compact sets*

$$M_0 \supset M_1 \supset \dots \supset M_n$$

with the property that, whenever $i \leq j$,

$$(M_{i-1}, M_j)$$

is an index pair for N_{ij} .

In particular, (M_0, M_n) is an index pair for S , and (M_{i-1}, M_i) is an index pair for N_i .

3.4 Cohomological Conley index

Before introducing homotopy invariants which require additional concepts from algebraic topology, we discuss cohomological invariants which only need the concepts introduced in 3.2.

Let I be an isolated invariant set and let (M_0, M_1) be an index pair for I . We put $CH^*(I) := H^*(M_0, M_1)$ (coefficients in \mathbb{Z} or \mathbb{Z}_2).

Theorem 11. *The cohomological Conley index $CH^*(I)$ does not depend on the choice of the index pair for I , that is, any two index pairs yield the same cohomology groups.*

Examples:


1. $I = \emptyset \Rightarrow CH^*(\emptyset) = 0$.

2. I a hyperbolic fixed point with unstable manifold of dimension n


$$\Rightarrow CH^k(I) = \begin{cases} \mathbb{Z} & \text{for } k = n \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 12. (McCord) *If $CH^*(I(N)) = \mathbb{Z}$, then $I(N)$ contains a fixed point.*

2a) $n = 1$:




2b) $n = 0$



$$M_1 = \emptyset$$

Morse index 0

However: 3.



$$CH^k(I) = \begin{cases} \mathbb{Z} & \text{for } k = 0 \\ 0 & \text{otherwise} \end{cases}$$


Attracting periodic orbit and hyperbolic fixed point of index 2

In this case, a more refined analysis is possible:

$$I_1 = \text{periodic orbit} : CH^*(I_1) = H^*(S^1) = \begin{cases} \mathbb{Z} & \text{for } k = 0, 1 \\ 0 & \text{otherwise.} \end{cases}$$

$$I_2 = \text{fixed point} : CH^*(I_2) = H^*(S^2, \text{point}) = \begin{cases} \mathbb{Z} & \text{for } k = 2 \\ 0 & \text{otherwise.} \end{cases}$$

4. Unstable periodic orbit



$$\text{red: } M_1$$

$$CH^k(I) = \begin{cases} \mathbb{Z} & \text{for } k = 1, 2 \\ 0 & \text{otherwise} \end{cases}$$

5. Let I be an invariant set which is normal hyperbolic for a differentiable flow f ; this means that the linearization does not have any eigenvector transversal to I with an eigenvalue with vanishing real part. The local unstable manifold is then homotopic to a vector bundle. Let this vector bundle be of rank n and orientable. By the isomorphism theorem of Thom, a deep result in algebraic topology that we shall not discuss here in detail, we then have


$$CH^k(I) = H^{k-n}(I).$$

Corollary 1. *Let I be a hyperbolic, periodic orbit with an orientable unstable manifold of dimension $n + 1$. Then*

$$CH^k(I) = \begin{cases} \mathbb{Z} & \text{for } k = n, n + 1 \\ 0 & \text{otherwise.} \end{cases}$$

This result includes the periodic orbits in the examples 3 and 4.

6. Conversely, however, a Conley index as in the corollary does not imply the existence of a periodic orbit as the next, graphically displayed, example shows.



red: M_1
 2 fixed points (index 2
 and 1) with 2 hetero-
 clinic orbits

7. A hyperbolic periodic orbit with a two-dimensional local unstable manifold that is homeomorphic to a Möbius strip, and hence non-orientable. Contracting the boundary of the Möbius strip to a point, we obtain $(\mathbb{R}P^2, \text{point})$; therefore

$$CH^k(I, \mathbb{Z}_2) = \begin{cases} \mathbb{Z}_2 & \text{for } k = 1, 2 \\ 0 & \text{otherwise.} \end{cases}$$

This and more general examples can be investigated by the Thom isomorphism with \mathbb{Z}_2 coefficients.

8. Of course, it is also instructive to analyze the examples from 3.2 with the help of the present concepts. We urge the reader to carry this out as an exercise before we proceed to the more general homotopical invariants.

3.5 Homotopical invariants

We shall need to summarize some concepts from algebraic topology. The reader who happens to be familiar with them already can proceed directly to the statement of the next theorem.

We consider pairs (X, A) , $A \subset X$ closed, of topological spaces, and we abbreviate $X = (X, \emptyset)$. All maps will be assumed to be continuous, that is, we work in the category of topological spaces and continuous maps between them.

Let

$f : (X, A) \rightarrow (Y, B)$ be continuous with $f(A) \subset B$.

Two such maps are homotopic, that is, continuously deformable into each other, $f_0 \sim f_1 : (X, A) \rightarrow (Y, B)$, when, with $I = [0, 1]$, there exists an

$$F : (X, A) \times I \rightarrow (Y, B) \quad \text{with } F(x, t) = f_t(x) \text{ for } t = 0, 1.$$

The crucial point for the concept of homotopy is the continuity of F (as contained in our general assumption).

A homotopy equivalence between two such pairs of spaces, $(X, A) \sim (Y, B)$, is defined by requiring that there exist

$$f : (X, A) \rightarrow (Y, B), \quad g : (Y, B) \rightarrow (X, A)$$

with $f \circ g \sim id_Y, g \circ f \sim id_X$.

Let us consider some examples, first for pairs of the form $(X, \emptyset), (Y, \emptyset)$. We let $D := \{x = (x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 \leq 1\}$ be the closed unit disk and 0 the origin in \mathbb{R}^2 . Then the identity map of D , $f_1 := id_D$, and the constant map f_0 that maps all of D to 0 are homotopic: the homotopy is provided by $F(x, t) = tx$. This then also implies that the space D is homotopically equivalent to the space consisting only of the origin, that is, of a single point. If we want to have nontrivial pairs here, we take $(D, \{0\})$ and $(\{0\}, \{0\})$. Another example of homotopically equivalent spaces is given by the cylinder $Z := \{(x_1, x_2, x_3) \in \mathbb{R}^3 : x_1^2 + x_2^2 = 1, -1 \leq x_3 \leq 1\}$ and the circle $S := \{x = (x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 = 1\}$, with the required map from the cylinder to the circle simply collapsing the third coordinate x_3 to 0 , and the map from the circle to the cylinder embedding the former into the latter as the circle $x_3 = 0$. Again, this also shows the homotopy equivalence of (Z, A) and (S, A) , where A is any subset of S . The disk D and the circle S , however, are not homotopically equivalent. If one does not know how to prove that, one might wish to consult any elementary introduction to algebraic topology, for example [52]. Even without that, however, the reader should see that the preceding statement is equivalent to the fact that the identity map of the circle is not homotopic to a constant map.

Of course, the reader can also do the standard exercises here, like verifying that homotopy equivalence is an equivalence relation on the set of pairs of topological spaces.

(X, x_0) ($x_0 \in X$) is called a punctured space. (X, A) yields the punctured space X/A by identifying all $x \in A$ as a single point. More formally: $x \approx y$ if $x = y$ or $x, y \in A$ is an equivalence relation, and we let $[A]$ be the equivalence class of $x \in A$. We then put

$$X/A := (X/\approx, [A]).$$

Thus, we collapse the whole subset A to a single point. Moreover, we have the important special case

$$X/\emptyset = (X \amalg p, p), p \notin X.$$

(Here, \amalg means the disjoint union, that is, here we are adding a point p to X that is not contained in X .)

A continuous map $f : (X, A) \rightarrow (Y, B)$ induces $[f] : X/A \rightarrow Y/B$ via $[f][x] := [f(x)]$.

If $(X, A) \sim (Y, B)$ then also $X/A \sim Y/B$.

In the sequel, we simply write X/A for the homotopy class $[X/A]$ of X/A .

An example occurring in our applications arises when we consider Z/A where A consists of the two boundary circles $x_3 = \pm 1$ of the cylinder Z . The reader should convince her/himself that this space is homotopically equivalent to $S^2/\{p_1, p_2\}$, the sphere modulo two points, for example the north and the south pole. This in turn is homotopically equivalent to T^2/S' where the torus T^2 is obtained by rotating the circle $S' := \{(x_1, 0, x_3) : x_1 = 2 + \sin \theta, x_3 = \cos \theta\}$ about the x_3 -axis.

The join of two punctured spaces is defined as follows:

$$(X, x_0) \vee (Y, y_0) := X \amalg Y/\{x_0, y_0\}$$

(glueing at the distinguished points x_0, y_0).

$\bar{0} := x/x, x$ point.

$A \subset X$ is called a strong deformation retract of X if there exists a continuous map $r : X \times [0, 1] \rightarrow X$ satisfying

$$\begin{aligned} r(x, t) &= x \quad \forall x \in A \\ r(x, 0) &= x \quad \forall x \in X \\ r(x, 1) &\in A \quad \forall x \in X. \end{aligned}$$

Thus, we shrink the whole space X into its subset A while leaving all points in A itself unaffected.

From the above examples, we see that the origin (or, likewise any other point in D) is a strong deformation retract of the unit disk D , and so is the circle S for the cylinder Z . This circle S , however, is not a strong deformation retract of the unit disk D .

With these topological concepts, we can now state the fundamental results of Conley theory.

Theorem 13. *If $(M_0, M_1), (M'_0, M'_1)$ are index pairs for the same isolated invariant set S , then M_0/M_1 and M'_0/M'_1 are homotopically equivalent.*

This result justifies the following fundamental

Definition 15. The **Conley index** $h(S)$ of the isolated invariant set S is defined as the homotopy type $[M_0/M_1]$ where (M_0, M_1) is any index pair for S .

Theorem 14. Let $M_0 \supset M_1 \supset \dots \supset M_n$ be a Morse filtration of the isolated invariant set S . Then the relative Betti numbers $b_\nu(M_{j-1}, M_j)$ satisfy

$$\sum_{j=1}^n \sum_{\nu} b_\nu(M_{j-1}, M_j)t^\nu = \sum_{\nu} b_\nu(M_0, M_n)t^\nu + (1+t)Q(t)$$

as a power series for the indeterminate t where $Q(t)$ is a polynomial with nonnegative integer coefficients.

We note that the relative Betti numbers occurring in the preceding theorem in turn have been defined through a dynamical system, namely a gradient flow for an appropriate Morse function, in 3.2. Here, in contrast, they appear as invariants of the underlying spaces that are used to detect information about a dynamical system. Thus, there really is an intimate connection between the topology of the space and the qualitative properties of a dynamical system operating on it, and one can be used to elucidate the other. The topological and the dynamical information appear as two sides of the same coin.

Theorem 15. Let I_1, I_2 be disjoint isolated invariant sets. Then the disjoint union of I_1 and I_2 , written as $I_1 \amalg I_2$, is an isolated invariant set, and

$$h(I_1 \amalg I_2) = h(I_1) \vee h(I_2)$$

when \vee is the join of topological spaces as explained above.

Theorem 16. Let N be an isolating neighborhood of the isolated invariant set I . If $h(I) \neq \bar{0}$, then $I \neq \emptyset$, i.e. N contains an entire orbit of the flow.

Proof. $h(\emptyset) = \bar{0}$.

q.e.d.

We now discuss some algebraic tools for deriving such relationships as in Theorem 14. Let $Z \subset Y \subset X$ be topological spaces. We then have the long exact cohomology sequence

$$\begin{aligned} 0 &\rightarrow H^0(X, Y) \rightarrow H^0(X, Z) \rightarrow H^0(Y, Z) \\ &\xrightarrow{\delta^0} H^1(X, Y) \rightarrow H^1(X, Z) \rightarrow H^1(Y, Z) \\ &\xrightarrow{\delta^1} H^2(X, Y) \rightarrow \dots \end{aligned}$$

Exactness here means that the composition of any two subsequent maps in this sequence is 0. For details about the definition of these maps, we need to

refer to any introduction to algebraic topology, such as [52].

Assume that all cohomology groups are finitely generated. We put

$$b_j(X, Y) := \dim H^j(X, Y)$$

$$\nu_j(X, Y, Z) := \dim (\operatorname{im} \delta^j).$$

Lemma 2. *Let $0 \rightarrow A_1 \xrightarrow{a_1} A_2 \xrightarrow{a_2} A_3 \rightarrow \dots$ be an exact sequence of linear maps between vector space. Then for all $k \in \mathbb{N}$ $\dim A_1 - \dim A_2 + \dim A_3 - \dots - (-1)^k \dim A_k + (-1)^k \dim (\operatorname{im} \delta^k) = 0$.*

Proof. If $\ell : V \rightarrow W$ is linear, then $\dim V = \dim (\ker \ell) + \dim (\operatorname{im} \ell)$. The exactness implies that

$$\dim (\ker a_j) = \dim (\operatorname{im} a_{j-1}).$$

Hence

$$\dim A_j = \dim (\operatorname{im} a_{j-1}) + \dim (\operatorname{im} a_j)$$

and with

$$\dim A_1 = \dim (\operatorname{im} a_1)$$

the claim follows.

q.e.d.

The Lemma implies that

$$\sum_{j=0}^m (-1)^j (b_j(X, Y) - b_j(X, Z) + b_j(Y, Z)) - (-1)^m \nu_m(X, Y, Z) = 0.$$

Hence

$$(3.1) \quad (-1)^{m-1} \nu_{m-1}(X, Y, Z) =$$

$$(-1)^m \nu_m(X, Y, Z) - (-1)^m b_m(X, Y) + (-1)^m b_m(X, Z) - (-1)^m b_m(Y, Z).$$

We put

$$P(t, X, Y) := \sum_{m \geq 0} b_m(X, Y) t^m$$

$$Q(t, X, Y, Z) := \sum_{m \geq 0} \nu_m(X, Y, Z) t^m$$

Multiplying (3.1) by $(-1)^m t^m$ and summing yields

$$(3.2) \quad Q(t, X, Y, Z) = -tQ(t, X, Y, Z) + P(t, X, Y) - P(t, Y, Z) + P(t, Y, Z).$$

We can now prove Theorem 14.

Theorem 17. *Let $M_n \subset M_{n-1} \subset \dots \subset M_1 \subset M_0$ be a filtration of M_0 , for example a Morse filtration of the Morse decomposition (N_1, \dots, N_n) of the isolated invariant set I . Then*

$$\sum_{j=1}^n P(t, M_{j-1}, M_j) = P(t, M_0, M_n) + (1+t)Q(t)$$

with $Q(t) = \sum_{j=1}^{n-1} Q(t, M_{j-1}, M_j, M_n)$

(a polynomial in t with nonnegative integer coefficients).

Proof. (3.2) for $M_{j-1} \supset M_j \supset M_n, j = 1, \dots, n - 1$ implies that $P(t, M_{j-1}, M_j) + P(t, M_j, M_n) = P(t, M_{j-1}, M_n) + (1+t)Q(t, M_{j-1}, M_j, M_n)$. Summation w.r.t. j yields the claim.

q.e.d.

We now put

$$P(t, h(I)) := P(t, M_0, M_1)$$

for an index pair (M_0, M_1) for the isolated invariant set I .


Corollary 2. *Let I be an isolated invariant set, (N_1, \dots, N_n) an admissible ordering of a Morse decomposition of I . Then*

$\sum_{j=1}^n P(t, h(N_j)) = P(t, h(I)) + (1+t)Q(t)$, where $Q(t)$ is a polynomial with nonnegative integer coefficients.

Let us return once more to the example of 3.2, in order to discuss also global aspects of the Conley index. The example X in question is diffeomorphic to S^2 , the two-dimensional sphere, with a function f with 4 critical points, namely two local maxima p_1, p_2 , one minimum r , and one saddle point q . For each critical point, an index pair (M_0, M_1) is indicated in the figure. For each such pair, M_1 is marked in red, $\partial M_0 \setminus M_1$ in green. Thus, each fixed point is considered as an isolated invariant set. For each of the two maxima, the Conley index is the homotopy type of $D/\partial D$, ($D = \text{disk}$), i.e., of (S^2, point) , and thus

$$CH^k(p_i) = \begin{cases} \mathbb{Z} & \text{for } k = 2 \\ 0 & \text{otherwise.} \end{cases}$$

For the saddle point q , we obtain the homotopy type of square/(top and bottom boundary lines), i.e., of interval/end points, i.e. of (S^1, point) , as in the figure for example 2a). Thus



$$CH^k(q) = \begin{cases} \mathbb{Z} & \text{for } k = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Finally, the Conley index of the minimum r is given by D/\emptyset , as in example 2b), i.e., by a point. Thus

$$CH^k(r) = \begin{cases} \mathbb{Z} & \text{for } k = 0 \\ 0 & \text{otherwise.} \end{cases}$$


In order to connect this example with the theorems formulated in the preceding §§, we have to construct a Morse filtration. For that purpose, we shall deform our index pairs while preserving the topological information obtained. Thus, for the index pairs for the maxima p_1, p_2 , we enlarge the corresponding sets M_0 to $M_0(p_1)$ and $M_0(p_2)$ until they touch the set $M_0(q)$. All the rest will be assigned as $M_0(r)$ to the minimum r . Obviously, this will not affect the cohomological indices.

We consider also the entire space X as an isolated invariant set, and we employ now the letters M_0, M_1, \dots , for a Morse filtration, namely the Morse decomposition (p_1, p_2, q, r) . We put

$$\begin{aligned} M_0 &= X, \quad M_1 = \overline{X \setminus M_0(p_1)}, \quad M_2 = \overline{M_1 \setminus M_0(p_2)}, \\ M_3 &= \overline{M_2 \setminus M_0(q)}, \quad M_4 = \overline{M_3 \setminus M_0(r)} = \emptyset. \end{aligned}$$

Then $(M_0, M_1), (M_1, M_2), (M_2, M_3), (M_3, M_4)$ are index pairs for p_1, p_2, q, r , homotopically equivalent to the index pairs constructed above.

We now wish to evaluate the general formula



$$\sum_{j=1}^4 \sum_{\nu} b_{\nu}(M_{j-1}, M_j)t^{\nu} = \sum_{\nu} b_{\nu}(M_0, M_4)t^{\nu} + (1+t) Q(t)$$

for our example (thus, $n = 4$). According to the above computations, we obtain on the left-hand side

$$2t^2 + t + 1.$$

The term on the right-hand side is

$$\sum_{\nu} b_{\nu}(S^2, \emptyset) t^{\nu} = t^2 + 1,$$

because $b_2(S^2) = 1$, $b_1(S^2) = 0$, $b_0(S^2) = 1$, as we have deduced above in 3.2 with the help of Floer homology.

Thus, the remaining term is


$$(1+t) Q(t) = t^2 + t = (1+t)t.$$

This term contains the excess of the local information gathered on the left-hand side over the global quantity $\sum_{\nu} b_{\nu}(M_0, M_n)t^{\nu}$ on the right-hand side. This term results from the fact that on the left-hand side, in contrast to the right-hand side, we do not include the cancellations implied by the boundary operator, as should become clear when comparing with the computation of Floer homology.

We can localize this in an even more precise manner (and the flexibility to interpolate between local and global information constitutes one of the

decisive strengths of Conley theory, as we shall see in more detail in 3.6). For this purpose, we consider two invariant sets I_1 and I_2 .

I_1 consists of the critical points p_1, p_2 , and q , and the connecting orbits, i.e. the two flow-lines from p_1 and p_2 , resp., to q . I_2 contains only the critical point r .



An index pair for I_1 consists now of the red curve as $M_1(I_1)$ and the set above it as $M_0(I_1)$. This is extended to a Morse filtration by adding to each of them the entire lower part, i.e.,

$$M_0 = X, M_1 = \overline{M_0 \setminus M_0(I_1)}, M_2 = \emptyset.$$

The homotopy type of M_0/M_1 then is (S^2, point) as before, while that of M_1/M_2 is (point) . Thus, the general formula

$$\sum_{j=1}^2 \sum_{\nu} b_{\nu}(M_{j-1}, M_j) t^j = \sum_{\nu} b_{\nu}(M_0, M_2) t^{\nu} + (1+t) Q(t)$$


becomes

$$t^2 + 1 = t^2 + 1,$$

i.e. , $Q(t) = 0$. The additional topological information is now entirely contained within $M_0(I_1)$, and the decomposition is too coarse for representing the local information.

This example exhibits another interesting aspect that will lead to the next §. Namely, inside $M_0(I_1)$, we can deform the flow such that the saddle point q moves into one of the two maxima, and the saddle point and the maximum

then cancel each other. While the invariant set I_1 changes, the index pair $(M_0(I_1), M_1(I_1))$ remains the same. Our coarse Morse decomposition thus remains invariant under deformations that take place within an index pair. The preceding, finer, Morse decomposition needs however to be adapted, and readers should try for themselves to discuss Morse decompositions for the middle diagram in the next figure.



We now consider a general gradient flow as in 3.2:

$$\dot{x} = -\nabla f(x),$$

with f having only isolated critical points p_1, \dots, p_n on the compact manifold M . $N_i = \{p_i\}$ then constitutes a Morse decomposition (not yet necessarily admissibly ordered), as follows.

For $x \in M$ either $x(t) \equiv x$ or $f(x(t)) < f(x)$ for $t > 0$. In the first case, f is constant on $\alpha(x), \omega(x)$, and $\alpha(x), \omega(x)$ then satisfy $\nabla f(y) = 0$.

Thus, either $\nabla f(x) = 0$ or $f(\omega(x)) < f(\alpha(x))$.

If the ordering is such that $f(p_i) \leq f(p_j)$ for $i > j$, it is admissible. As discussed in 2.2, in a neighborhood of a nondegenerate critical point, the flow is given in local coordinates $x \in \mathbb{R}^d$ by

$$\begin{aligned} \dot{x}_- &= A_- x_- + g_-(x) \\ \dot{x}_+ &= A_+ x_+ + g_+(x) \end{aligned}$$

with

$$\begin{aligned} g_{\pm}(0) &= g'_{\pm}(0) = 0 \\ |g_{\pm}|, |g'_{\pm}| &\leq \delta \end{aligned}$$

for $|x| \leq 2$, where δ is arbitrarily small for a suitable choice of coordinates and, for

$$x = (x_-, x_+) \in \mathbb{R}^{d^-} \times \mathbb{R}^{d^+} = \mathbb{R}^d,$$

we have

$$\begin{aligned} \langle x_-, A_- x_- \rangle &\leq -\lambda |x_-|^2 \\ \langle x_+, A_+ x_+ \rangle &\geq \lambda |x_+|^2 \text{ for some } \lambda > 0. \end{aligned}$$

This is the so-called Morse lemma.

Let $Q := \{x \in \mathbb{R}^d : |x_{\pm}| \leq 1\}$. For $x \in \partial Q$, $|x_-| = 1$ or $|x_+| = 1$.

If for $x \in Q$, $|x_-| \leq |x_+|$, then $\frac{d}{dt}|x_-|^2 = 2\langle x_-, A_- x_- + g_-(x) \rangle \leq -\lambda |x_-|^2$ for $4\delta \leq \lambda$.

If for $x \in Q$, $|x_+| \leq |x_-|$, then $\frac{d}{dt}|x_+|^2 \geq \lambda |x_+|^2$.

$M_0 := Q$, $M_1 := \{x \in Q, |x_+| = 1\}$ then constitute an index pair, and

$P(t, h(p_i)) = t^{d_i^-}$ as well as $(M_0, M_1) \sim (D^{d_i^-}, \partial D^{d_i^-})$.

(M, \emptyset) is an index pair for M and, with

$P(t, h(M)) = \sum_{j=0}^d b_j(M)t^j$, the general theory implies that

$$\sum_{i=1}^n t^{d_i^-} = \sum_{j=0}^d b_j(M)t^j + (1+t)Q(t),$$

and so

$$(3.3) \quad m_j := \#\{\text{critical points } p_i \text{ with } d_i^- = j\} \geq b_j(M).$$

The last result is the content of Theorem 9 above.

Remark 2. Let X be compact, $\varphi : X \rightarrow \mathbb{R}$ continuous. We put $X^\mu := \{x \in X : \varphi(x) \leq \mu\}$ for $\mu \in \mathbb{R}$, $\mu_0 := \max_{x \in X} \varphi(x) > \mu_1 > \dots > \mu_n$ for real numbers μ_0, \dots, μ_n , $M_i := X^{\mu_i}$ ($M_i = \emptyset$ for $\mu_i < \min_{x \in X} \varphi(x)$).

Then $X = M_0 \supset M_1 \supset \dots \supset M_n$ is a filtration of X by sublevel sets of φ . Therefore

$$\sum_{j=1}^n \sum_{m \geq 0} b^m(M_{j-1}, M_j)t^m = \sum_{m \geq 0} b^m(M_0, M_n)t^m + (1+t)Q(t).$$

This is a localization in the image. In order to also obtain a localization in the domain, as in Morse theory, one uses a flow. One may define even without using a flow, however, what a critical point of a continuous function $\varphi : X \rightarrow \mathbb{R}$ is, and even some kind of an index of such a critical point. Namely, $x_0 \in X$ is noncritical for φ precisely if there exists a neighborhood U of x_0 in X with a continuous map

$$\Phi : U \times [0, 1] \rightarrow X$$

such that

$$\begin{aligned} \varphi(\Phi(x, t)) &\text{ is monotonically decreasing in } t \in [0, 1] \text{ for all } x \in U, \\ \varphi(\Phi(x, 1)) &< \varphi(x_0) \text{ for all } x \in U. \end{aligned}$$

For defining the index of a critical point x_0 of φ , we put

$$X^-(x_0) := \{x \in X : \varphi(x) < \varphi(x_0)\}.$$

The index then is the homotopy type of

$$(X^-(x_0) \cup \{x_0\}) / X^-(x_0).$$

3.6 Continuation properties of the Conley index

Let $f^\Lambda : X \times \mathbb{R} \rightarrow X$ be a family of flows, continuous with respect to $\lambda \in \Lambda$, a compact, locally contractible, connected metric space with metric $d(\cdot, \cdot)$.

Lemma 3. *Let N be an isolating neighborhood for f^{λ_0} , $\lambda_0 \in \Lambda$. Then there exists $\varepsilon > 0$ such that, for all $\lambda \in \Lambda$ with $d(\lambda, \lambda_0) < \varepsilon$, N is an isolating neighborhood for f^λ .*

This is no longer true for index pairs.

Theorem 18. *Under the assumptions of the lemma,*

$$h(\text{Inv}^\lambda(N)) = h(\text{Inv}^{\lambda_0}(N))$$

($\text{Inv}^\lambda(N)$ is the invariant set of N with respect to the flow f^λ).

To proceed, we consider

$$\begin{aligned} F : X \times \Lambda \times \mathbb{R} &\rightarrow X \times \Lambda \\ (x, \lambda, t) &\mapsto (f^\lambda(x, t), \lambda). \end{aligned}$$

F is a flow on $X \times \Lambda$ and is called a parameterized flow. We also put

$$N \subset X \times \Lambda, N^\lambda := N \cap (X \times \{\lambda\}).$$

Definition 16. *For $\lambda_i \in \Lambda, i = 0, 1$, let I^i be an isolated invariant set for f^{λ_i} . I^0 and I^1 are related by continuation if there exists an isolating neighborhood N of the parameterized flow with $N \subset X \times \Lambda$, $\text{Inv}^{\lambda_i}(N^{\lambda_i}) = I^i$ ($i = 0, 1$) (invariant set of N^{λ_i} with respect to f^{λ_i}).*

Theorem 19. *If I^0 and I^1 are related by continuation, then $h(I^0) \sim h(I^1)$.*

General formulation: Let X be a compact metric space. Let the compact-open topology be given on the space Φ of flows $\varphi : X \times \mathbb{R} \rightarrow X$, i.e., $\varphi_n \rightarrow \varphi$ if φ_n converges to φ uniformly on compact sets in $X \times \mathbb{R}$. Φ is complete w.r.t. this topology.

Let $N \subset X$ be closed with $\overset{\circ}{N} \neq \emptyset$.

Let $\mathcal{U}(N) := \{f \in \Phi : N \text{ is an isolating neighborhood for } f\}$.

$\mathcal{U}(N)$ is open in Φ by the above lemma.

Let $A(X) := \{V \subset X, V \text{ closed}\}$.

Put $\mathcal{S} := \{(f, I) : I \text{ is an isolated invariant set for } f \in \Phi\} \subset \Phi \times A(X)$.

For $N \subset X$ closed, $\overset{\circ}{N} \neq \emptyset$, we consider both

$$\begin{aligned} \sigma_N : \mathcal{U}(N) &\rightarrow \mathcal{S} \\ f &\mapsto (f, I_f) \end{aligned}$$

where I_f is a maximal invariant subset of N with respect to f , and the projection

$$\pi_1 : \mathcal{S} \rightarrow \Phi$$

with

$$\pi_1 \circ \sigma_N = id : \mathcal{U}(N) \rightarrow \mathcal{U}(N).$$

Let now \mathcal{S} carry the topology generated by the sets $\sigma_N(U), U \subset \mathcal{U}(N)$ open.

Theorem 20. $\pi_1 : \mathcal{S} \rightarrow \Phi$ is a local homeomorphism. In other words, (\mathcal{S}, Φ, π) is a sheaf.¹

Theorem 21. Let (f^0, I^0) and (f^1, I^1) be contained in the same component of \mathcal{S} . Then the corresponding Conley indices coincide.

It is, of course, instructive to analyze our examples of bifurcations discussed in Chapter 2, for example the Hopf bifurcation, in the light of the preceding results.

So far, starting with 3.3, we have been essentially following the original reference [7] and also using the exposition in [50]. Another good presentation of the theory can be found in [46].

3.7 The discrete Conley index

We consider a continuous map $f : X \rightarrow X$ of a topological space X . Let $J \subset \mathbb{Z}$ be an interval, i.e., a set of the form $J = \{m, m + 1, \dots, M - 1, M\}, m, M \in \mathbb{Z}$ then

$$\gamma : J \rightarrow X$$

is called a solution of f if

$$\gamma(n) = f(\gamma(n - 1)),$$

whenever $n - 1, n \in J$. If $0 \in J$ and $\gamma(0) = x$, then γ is called a solution through the point x .

For $N \subset X$, we put

$$I^+(N) := \{x \in X : f^n(x) \in N \text{ for all } n \geq 0\} \text{ (forward invariant set)}$$

$$I^-(N) := \{x \in X : \exists \text{ solution } \gamma : \{n \leq 0\} \rightarrow N \text{ with } \gamma(0) = x\} \\ \text{(backward invariant set)}$$

$$I(N) := I^+(N) \cap I^-(N) \text{ (invariant set) .}$$

¹ This is a terminology that was first introduced in complex analysis for studying continuation properties of holomorphic and meromorphic functions.

If $f : X \rightarrow X$ is a homeomorphism, then

$$I(N) = \{y \in X : f^n(y) \in N \ \forall n \in \mathbb{Z}\}.$$

A compact $N \subset X$ is called an isolating neighborhood, and $I(N)$ is called an isolated invariant set, if

$$I(N) \subset \overset{\circ}{N}.$$


$M \subset N$ is called positively invariant with respect to N if

$$f(M) \cap N \subset M.$$

Examples:

$$f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$$


$$(x, y) \mapsto (2x, \frac{1}{2}y).$$



0 is an isolated invariant point

$$H^k(M_0, M_1) = \begin{cases} \mathbb{Z} & \text{for } k = 1 \\ 0 & \text{otherwise} \end{cases}$$

$M_0 = \text{box}$
 M_1 shaded



$$H^k(M'_0, M'_1) = \begin{cases} \mathbb{Z} & \text{for } k = 0, 1 \\ 0 & \text{otherwise} \end{cases}$$

$M'_0 = \text{both boxes}$
 M'_1 shaded

We conclude that in the discrete case the cohomology of an index pair for an isolated invariant set depends on the choice of this index pair. Therefore, we need to refine the construction.

A pair (M_0, M_1) of compact subsets $M_1 \subset M_0 \subset N$ of the isolating neighborhood of the isolated invariant set $I(= I(N))$ is called an index pair for I in N (with respect to f) if

- (i) I is contained in the interior of $M_0 \setminus M_1$,
- (ii) M_0 and M_1 are positively invariant with respect to N ,
- (iii) $\overline{M_0 \setminus M_1}, f(\overline{M_0 \setminus M_1}) \subset \overset{\circ}{N}$.

Thus, if $x \in M_0$, either $x \in \overline{M_0 \setminus M_1}$ or $x \in M_1 \setminus \overline{M_0 \setminus M_1}$.

In the first case, $f(x) \in \overset{\circ}{N}$ by (iii) and hence $f(x) \in M_0$ by (ii). Thus, $f(x)$ can lie outside of M_0 only in the second case. In particular, we have that, if $x \in M_0, f(x) \notin M_0$, then $x \in M_1$ (in other words, $f(\overline{M_0 \setminus M_1}) \subset M_0$). If now $x \in M_1$, then by (ii) either $f(x) \in M_1$ or $f(x)$ lies outside of N and hence also outside of M_0 . In this sense, M_1 can be considered as the exit set.

For these constructions, the isolating neighborhood of the isolated invariant set I is considered as given. One may then show the existence of index pairs for I in N with respect to f .

We have the inclusion

$$f(M_0) \setminus (M_0 \cap f(M_0)) \subset f(M_1),$$

because, for an index pair, $f(M_0 \setminus M_1) \subset M_0$. Therefore, $M_0 \cup f(M_0) = M_0 \cup f(M_1)$. Thus, by excision of $f(M_1)$, we obtain an isomorphism

$$r : H^*(M_0 \cup f(M_0), M_1 \cup f(M_1)) \rightarrow H^*(M_0, M_1).$$

for the (Alexander-Spanier) cohomology.²

Furthermore,

$$f : (M_0, M_1) \rightarrow (M_0 \cup f(M_0), M_1 \cup f(M_1))$$

induces a map

$$f^* : H^*(M_0 \cup f(M_0), M_1 \cup f(M_1)) \rightarrow H^*(M_0, M_1).$$

We define then the index map for the index pair (M_0, M_1) as

$$f_{M_0, M_1}^* := f^* \circ r^{-1} : H^*(M_0, M_1) \rightarrow H^*(M_0, M_1).$$

Let V be a graded vector space with an endomorphism $f : V \rightarrow V$ which preserves the grading ($\deg f = 0$).

We put

² We again need to refer to an introduction to algebraic topology here.

$$\begin{aligned} \text{g-ker}(f) &:= \bigcup \{f^{-n}(0) : n \in \mathbb{N}\} \\ \text{g-im}(f) &:= \bigcap \{f^n(V) : n \in \mathbb{N}\}. \end{aligned}$$

Then f induces a monomorphism

$$f' : V/\text{g-ker}(f) \rightarrow V/\text{g-ker}(f).$$

A monomorphism $f : V \rightarrow V$ induces an isomorphism

$$f'' : \text{g-im}(f) \rightarrow \text{g-im}(f).$$

Putting $LV := \text{g-im}(f')$, the Leray reduction of an endomorphism $f : V \rightarrow V$ is then

$$Lf := (f'')'' : LV \rightarrow LV.$$

Let I be an isolated invariant set with index pair (M_0, M_1) . The corresponding cohomological Conley index of I then is defined as

$$CH^*(I) := CH^*(M_0, M_1) := (LH^*(M_0, M_1), Lf_{M_0, M_1}^*).$$

That the index of an isolated invariant set I is well-defined is a consequence of the next result.

Theorem 22. *Let I be an isolated invariant set for the continuous map $f : X \rightarrow X$. Then*

- (i) *There exists an index pair (M_0, M_1) for I .*
- (ii) *The cohomological Conley index does not depend on the choice of this index pair.*


If f happens to be the time-1-map of a flow then, for each isolated invariant set of the flow, an index pair (M_0, M_1) can be constructed that is also an index pair for f and for which the above maps

$$r, f^* : H^*(M_0 \cup f(M_0), M_1 \cup f(M_1)) \rightarrow H^0(M_0, M_1)$$


are homotopic to each other, and hence f_{M_0, M_1}^* is the identity. Therefore, the cohomological Conley index for the flow can be identified with the Conley index with respect to the time-1-map.

Examples:

1) The horse-shoe map of S. Smale:



$M_0 =$ square $ABCD$
 $M_1 =$ shaded region of M_0
 $f(M_0) =$ horse-shoe
 $f(M_1) =$ shaded in red



$$f_{M_0, M_1}^* : H^*(M_0, M_1) \rightarrow H^*(M_0, M_1)$$


$$\alpha \mapsto \alpha + \beta$$

$$\beta \mapsto -\alpha - \beta$$

$$\Rightarrow (f_{M_0, M_1}^*)^2 = 0$$

$$\Rightarrow CH^*(M_0, M_1) = (0, 0).$$

2) The G-horse-shoe map



$M_0, M_1, f(M_0), f(M_1)$ as in 1),
 but this time


$$f_{M_0, M_1}^1 : H^1(M_0, M_1) \rightarrow H^1(M_0, M_1)$$

$$\alpha \mapsto \alpha + \beta$$

$$\beta \mapsto \alpha + \beta,$$

and therefore $(f_{M_0, M_1}^1)^2 = 2f_{M_0, M_1}^1$ and $CH^*(M_0, M_1; \mathbb{Q}) = (\mathbb{Q}, 2id)$ (with \mathbb{Z} -coefficients, the Conley index would be 0 again, because $2id : \mathbb{Z} \rightarrow \mathbb{Z}$ is not an isomorphism).

3)



$M_0, M_1, f(M_0), f(M_1)$ as in 1)

$$f_{M_0, M_1}^1 : H^1(M_0, M_1) \rightarrow H^1(M_0, M_1)$$

$$\alpha \mapsto \alpha + \beta + \gamma$$

$$\beta \mapsto -\alpha - \beta - \gamma$$

$$\gamma \mapsto \alpha + \beta + \gamma$$

$$CH^*(M_0, M_1, \mathbb{Q}) = (\mathbb{Q}, id).$$

4 Entropy and topological aspects of dynamical systems

4.1 The entropy of a process as an asymptotic quantity

We consider the continuous map

$$T : X \rightarrow X$$

of a compact state space, and we wish to measure how much information this map generates. The basic idea is here to determine how much additional information we obtain when, in addition to the position of a point $x \in X$, we also know the position of Tx . This basic idea will be made precise in several steps. First of all, we need to quantify the amount of information contained in the position of x . In order to distinguish between different positions, we need to partition X into subsets A_1, \dots, A_m , with the aim that the information about x comprises in which of these subsets it lies. This does not seem to look very canonical, since the choice of m , the number of our subsets, is arbitrary, as are the possible criteria for selecting these subsets. Later on, however, this arbitrariness will be eliminated by a limit process. In any case, the sets A_1, \dots, A_m should not be chosen in a completely arbitrary manner, but they should respect whatever structure X carries. Since so far X is only assumed to be a topological space, the only meaningful requirement at this point is that the sets A_1, \dots, A_m be open. Moreover, each conceivable position of x in X should carry some information. Therefore, the sets A_1, \dots, A_m should cover X , i.e.,

$$X = \bigcup_{i=1}^m A_i,$$

even though this may have the consequence that some x are contained in more than one of the A_i , i.e., that the A_i do not partition X . This deficit will again be eliminated by our subsequent limit process.

We denote the covering $\{A_1, \dots, A_m\}$ by \mathcal{A} . In order to avoid redundancies as much as possible, we first select a subcovering of \mathcal{A} of smallest possible cardinality, i.e., we seek to cover X by as few as possible of the sets A_1, \dots, A_m from \mathcal{A} . In the sequel, we shall occasionally assume without loss of generality that \mathcal{A} has been chosen minimal in the sense that it does not contain any

nontrivial subcovering, i.e. that none of the sets in \mathcal{A} can be omitted without violating the covering property.

In order to quantify the information, we assign the same measure, namely $\frac{1}{m}$, to each A_i . The entropy of the covering \mathcal{A} , again assumed to be minimal, is then simply

$$H(\mathcal{A}) := - \sum_{i=1}^m \frac{1}{m} \log_2 \frac{1}{m} = \log_2 m$$

or, more generally, if \mathcal{A} need not be minimal, the logarithm of the number of sets of its smallest subcovering. Since we did not assume any further structure on X , we have no possibility to measure the size of the sets A_1, \dots, A_m . This would be different if we had a probability measure on X . In that situation, we could require that all the subsets A_1, \dots, A_m be measurable and not only cover, but also partition, X , i.e.,

$$A_i \cap A_j = \emptyset \quad \text{for } i \neq j.$$

Of course, then we ignore the topological structure of X and abandon the requirement that the A_i be open, and for the sequel, in place of the continuity of T , we should require that T be measure preserving, i.e.,

$$\mu(T^{-1}(A)) = \mu(A) \quad \text{for all measurable subsets } A \text{ of } X.$$

The entropy of the partition $\mathcal{A} = \{A_1, \dots, A_m\}$ of X w.r.t. the measure μ would then be

$$H_\mu(\mathcal{A}) := - \sum_{i=1}^m \mu(A_i) \log_2 \mu(A_i)$$

Of course, this raises the question of how the concepts and quantities to be defined depend on the choice of the measure μ . In order to obtain canonical quantities, one constructs the entropy of the transformation T with respect to every measure μ that is preserved by T and then forms the supremum with respect to all such μ . It will turn out that the resulting quantity coincides with the one obtained from the above topological approach. This, however, is not so surprising because, for a partition of X by a fixed number m of subsets, the entropy is largest if all sets in the partition are of equal measure $\frac{1}{m}$. We shall return to these measure theoretic aspects in Chapter 6.

Finally, we can also introduce a notion of entropy when there exists a distance measure, i.e., a metric d on X . This will again lead to the same result as the other approaches, as will be explained in Chapter 5. We return to the topological situation. We need one more technical definition, namely the join

$$\mathcal{A} \vee \mathcal{B}$$

of two coverings $\mathcal{A} = \{A_1, \dots, A_m\}, \mathcal{B} = \{B_1, \dots, B_n\}$. This join consists of all sets of the form $A_i \cap B_j, A_i \in \mathcal{A}, B_j \in \mathcal{B}$, and so is a covering itself.

We now let the transformation T enter. Together with \mathcal{A} , we also consider the covering $T^{-1}\mathcal{A}$ that consists of all the sets $T^{-1}A_i, A_i \in \mathcal{A}$. Since T is assumed continuous, all these sets are open. We consider the covering

$$\mathcal{A} \vee T^{-1}\mathcal{A}.$$

The number of sets in this refined covering, when compared with the number of sets in the original covering \mathcal{A} , then yields information about how much additional knowledge we gain from the position of Tx when compared to the position of x . To be precise, the knowledge of both x and Tx selects one of the sets in the covering $\mathcal{A} \vee T^{-1}\mathcal{A}$, namely a set $A_i \cap T^{-1}A_j$ with $x \in A_i, Tx \in A_j$, because Tx is contained in A_j precisely when x is in $T^{-1}A_j$. (If, instead of coverings, we worked with partitions, this set $A_i \cap T^{-1}A_j$ would even be uniquely determined. The non-uniqueness in the present context will play no role later on and will be eliminated by a limit process.) Thus, we should consider now $H(\mathcal{A} \vee T^{-1}\mathcal{A})$. This, however, is not yet all the information that we can gain from T . For example, we could also evaluate the position of $T^2x (= T(Tx))$, i.e., consider the covering

$$\mathcal{A} \vee T^{-1}\mathcal{A} \vee T^{-2}\mathcal{A}.$$

Of course, this can be iterated. In order to normalize our quantities, we divide by the number of iterations and define the entropy of T with respect to the covering \mathcal{A} as

$$h(T, \mathcal{A}) := \lim_{N \rightarrow \infty} \frac{1}{N} H(\mathcal{A} \vee T^{-1}\mathcal{A} \vee \dots \vee T^{-(N-1)}\mathcal{A}).$$

When each iteration of T and the determination of the corresponding image of x is considered as a measurement, then at each step we divide by the number of measurements taken in order to obtain the average amount of information per measurement from T .

As already indicated, the construction becomes independent of the choice of the covering \mathcal{A} when we define the entropy of T as

$$h_{top}(T) := h(T) := \sup\{h(T, \mathcal{A}) : \mathcal{A} \text{ covering } X\}.$$

Of course, we must verify that the above limit and the supremum both exist. This is not difficult; essentially, it rests on the monotonicity property

$$H(\mathcal{A}) \leq H(\mathcal{B})$$

which holds whenever the covering \mathcal{B} refines \mathcal{A} in the sense that every element of \mathcal{B} is contained in some element of \mathcal{A} . For instance, $\mathcal{A} \vee \mathcal{B}$ refines both

coverings \mathcal{A} and \mathcal{B} as follows directly from the definitions.

The detailed proof of the existence of the limit introduced for the definition of $h(T, \mathcal{A})$ goes as follows. Let $m_{\mathcal{C}}$ be the cardinality of a minimal subcovering of a covering \mathcal{C} of X , i.e., the number of subsets of X contained in \mathcal{C} that are at least required for covering X . In particular, $m_{\mathcal{A} \vee \mathcal{B}} \leq m_{\mathcal{A}} m_{\mathcal{B}}$.

Then

$$\begin{aligned} H(\mathcal{A} \vee \mathcal{B}) &= \log_2 m_{\mathcal{A} \vee \mathcal{B}} \\ &\leq \log_2(m_{\mathcal{A}} m_{\mathcal{B}}) \\ &= H(\mathcal{A}) + H(\mathcal{B}). \end{aligned}$$

Hence also

$$\begin{aligned} H(\mathcal{A} \vee T^{-1}\mathcal{A} \vee \dots \vee T^{-(N+N') + 1}\mathcal{A}) \\ \leq H(\mathcal{A} \vee \dots \vee T^{-N+1}\mathcal{A}) + H(T^{-N}\mathcal{A} \vee \dots \vee T^{-(N+N') + 1}\mathcal{A}) \\ \leq H(\mathcal{A} \vee \dots \vee T^{-N+1}\mathcal{A}) + H(\mathcal{A} \vee \dots \vee T^{-N'+1}\mathcal{A}), \end{aligned}$$

because $H(T^{-1}\mathcal{A}) \leq H(\mathcal{A})$ by definition of $T^{-1}\mathcal{A}$, with equality if T is surjective.

For $\eta_N := H(\mathcal{A} \vee \dots \vee T^{-N+1}\mathcal{A})$, we then have

$$\eta_{N+N'} \leq \eta_N + \eta_{N'} \quad \text{for } N, N' \in \mathbb{N}.$$

Thus, for fixed $i \in \mathbb{N}$ and $N = \nu i + j$ with $0 \leq j < i$,

$$\begin{aligned} \frac{\eta_N}{N} &\leq \frac{\eta_{\nu i} + \eta_j}{N} \leq \frac{\eta_{\nu i}}{\nu i} + \frac{\eta_j}{\nu i} \leq \frac{\nu \eta_i}{\nu i} + \frac{\eta_j}{\nu i} \\ &= \frac{\eta_i}{i} + \frac{\eta_j}{\nu i} \end{aligned}$$

and hence

$$\overline{\lim}_{N \rightarrow \infty} \frac{\eta_N}{N} \leq \frac{\eta_i}{i} \quad \text{for all } i.$$

Since also

$$\inf_i \frac{\eta_i}{i} \leq \lim_{N \rightarrow \infty} \frac{\eta_N}{N},$$

the existence of the limit follows.

q.e.d.

It is not hard to convince oneself that the topological entropy remains invariant under topological conjugation, as follows.

Theorem 23. *Let $T_i : X_i \rightarrow X_i$ be continuous maps of the compact Hausdorff spaces X_i , $i = 1, 2$, $\varphi : X_1 \rightarrow X_2$ a homeomorphism with $\varphi T_1 = T_2 \varphi$. Then*

$$h_{top}(T_1) = h_{top}(T_2).$$

Under the assumptions of the theorem, the two maps T_1 and T_2 are called topologically conjugate.

References for this section are [38, 57].

4.2 Positive entropy and chaos


Since the definition of the entropy depends on limit processes, one may expect that the entropy of a map T cannot be computed directly, but that for computing it or, if that should not be feasible, for estimating it, one will need general theorems. First of all, however, there exist maps whose entropy obviously vanishes, like constant maps or the identity map of X .

Next, one can prove that every homeomorphism of the unit circle S^1 has vanishing entropy, and therefore so does every homeomorphism of a compact interval in \mathbb{R} . The proof (for which we refer to [24]) depends on the classification of homeomorphisms of S^1 ; this classification is subtle, but many homeomorphisms of S^1 are conjugate to a rotation which obviously has vanishing entropy.

The perhaps simplest map of non-zero entropy is the tent map already introduced in 2.4,

$$Z : [0, 1] \rightarrow [0, 1],$$

$$Z(x) = \begin{cases} 2x & \text{for } 0 \leq x \leq \frac{1}{2} \\ 2(1-x) & \text{for } \frac{1}{2} \leq x \leq 1 \end{cases}$$



This map has the entropy $\log_2 2 = 1$ because, at every step, the precision of the knowledge about the position is doubled. (It is not difficult to carry out the computation for this example explicitly.) This phenomenon also occurs explicitly if each real number in the unit interval is represented as a binary sequence, i.e.,

$$x = 0, b_1 b_2 \dots b_m \dots,$$

with $b_i \in \{0, 1\}$, and one puts

$$Tx = 0, b_2 b_3 \dots$$

i.e., multiplies x by 2 and then omits or forgets the first digit. (It does not matter here that the representation of x as a binary sequence is not unique.) If any number x can only be measured with a finite precision, for example up to the k^{th} digit, then the application of T generates one additional digit and thus doubles the precision. This map is called the one-sided 2-shift. In this interpretation, T operates on the space of all binary sequences $\{(b_1, b_2, b_3, \dots)\}$ by

$$T(b_1, b_2, b_3, \dots) = (b_2, b_3, b_4, \dots).$$

Similarly, on the space Y_2 of two-sided binary sequences,

$$(\dots, b_{-2}, b_{-1}, b_0, b_1, b_2, \dots),$$

we define the Bernoulli shift

$$(4.1) \quad T(\dots, b_{-1}, b_0, b_1, \dots) = (\dots, b'_{-1}, b'_0, b'_1, \dots)$$

by $b'_m := b_{m+1}$. The Bernoulli shift again has the entropy $\log_2 2 = 1$, because at each step at any position one additional bit of information is generated.

If we consider sequences whose elements can assume p different states, i.e.,

$$Y_p := \{(\dots, a_{-1}, a_0, a_1, \dots) \text{ with } a_i \in \{0, 1, \dots, p-1\} \text{ for all } i \in \mathbb{Z}\}$$

or

$$Y_p^0 := \{(a_1, a_2, \dots) \text{ with } a_i \in \{0, 1, \dots, p-1\} \text{ for all } i \in \mathbb{N}\},$$

then the corresponding shifts have the entropy $\log_2 p$.

Of course, we use the product topology on these sequence spaces, as they are infinite products of the spaces $\mathbb{Z}_p = \{0, 1, \dots, p-1\}$.

The entropy of the Bernoulli shift on the space Y_2 of two-sided binary sequences can also be simply computed as follows. We cover the space Y_2 by the two open sets

$$A_0 := \{(\dots, b_{-1}, b_0, b_1, \dots) : b_0 = 0\}$$

and

$$A_1 := \{(\dots, b_{-1}, b_0, b_1, \dots) : b_0 = 1\}.$$

Thus

$$\mathcal{A} = \{A_0, A_1\},$$

and so

$$T^{-1}\mathcal{A} = \{\{b_1 = 0\}, \{b_1 = 1\}\},$$

and

$$\mathcal{A} \vee T^{-1}\mathcal{A}$$

is the covering consisting of the sets

$$\begin{aligned} A_{00} &= \{b_0 = 0, b_1 = 0\}, A_{01} = \{b_0 = 0, b_1 = 1\}, \\ A_{10} &= \{b_0 = 1, b_1 = 0\}, A_{11} = \{b_0 = 1, b_1 = 1\}. \end{aligned}$$

The process can be iterated in this manner, and we obtain

$$H(\mathcal{A} \vee T^{-1}\mathcal{A} \vee \dots \vee T^{-(N-1)}\mathcal{A}) = - \sum_{N>1}^{2^N} \frac{1}{2^N} \log_2 2^{-N} = N,$$

and thus

$$h(T) \geq h(T, \mathcal{A}) = \lim_{N \rightarrow \infty} \frac{1}{N} N = 1.$$

Since on the other hand it is easily checked that, for an arbitrary covering \mathcal{A} of Y_2 at each step, the number of sets in the covering asymptotically gets at most doubled, we have

$$H(\mathcal{A} \vee T^{-1}\mathcal{A} \vee \dots \vee T^{-(N-1)}\mathcal{A}) \leq \log_2(2^N |\mathcal{A}|) + o(N) = N + o(N)$$

($|\mathcal{A}| :=$ number of elements of \mathcal{A}), and hence

$$h(T, \mathcal{A}) \leq 1 \quad \text{for all coverings } \mathcal{A}.$$

Thus also

$$h(T) \leq 1,$$

and altogether the equality $h(T) = 1$ results.

In an analogous and elementary manner, one also verifies that the Bernoulli shift on the space of \mathbb{Z}_p -valued sequences has the entropy $\log_2 p$.

For later purposes, we also observe that the periodic sequences, that is, those with $T^N(\dots, b_{-1}, b_0, b_1, \dots) = (\dots, b_{-1}, b_0, b_1, \dots)$ for some $N \in \mathbb{N}$, or equivalently, those with $b_{i+N} = b_i$ for all i , are dense in Y_p . Namely, the topology of

Y_p is generated by the cylinder sets (that is, these sets are open and form a basis for the topology) $A_{\beta_0, \dots, \beta_{N-1}}$ consisting of those sequences with $b_i = \beta_i$ for $i = 0, \dots, N-1$. These, however, contain the periodic sequences with $b_{i+kN} = \beta_i$ for $i = 0, \dots, N-1$ and $k \in \mathbb{Z}$.

We return to the interval maps and say that a continuous map $T : I \rightarrow \mathbb{R}$ of a closed interval I in \mathbb{R} possesses a fold map¹ $[a, b]$ if $a < c < b \in I$ and

$$(4.2) \quad [a, b] \subset T[a, c] \cap T[c, b].$$

Thus, the entire interval $[a, b]$ is covered by the image of either one of two subintervals. For the above tent map, this holds for $a = 0, c = \frac{1}{2}, b = 1$. We then have the theorem that every map $T : I \rightarrow I$ with a fold map has entropy

$$h(T) \geq \log_2 2 = 1.$$

Conversely, for every continuous $T : I \rightarrow I$ ($I \subset \mathbb{R}$ compact interval) with

$$h(T) > 0,$$

there exists $k \in \mathbb{N}$ for which T^k has a fold map. Among other things, this implies that such a T must have a periodic point whose period is not a power of 2. For more details, see [24].

In summary, the situation for interval maps is quite clear-cut: homeomorphisms have vanishing entropy while positive entropy is generated by foldings.

There also exist more general and abstract results. For instance, as explained above, two continuous maps $T : X \rightarrow X, S : Y \rightarrow Y$ are called topologically conjugate if there exists a homeomorphism $\varphi : X \rightarrow Y$ with

$$\varphi \circ T = S \circ \varphi$$

and, more generally, S is called a factor of T if there exists a surjective continuous map $\psi : X \rightarrow Y$ with

$$\psi \circ T = S \circ \psi$$

If S is a factor of T , then

¹ Often, this is also called a horse-shoe, but as its properties are somewhat different from the horse-shoe introduced in 3.7 and further studied in 7.2, we have to give it a different name here. The essential differences to that horse-shoe are that the fold map in one dimension is not injective and that the horse-shoe of 3.7 does not fit into the original square so that, during the dynamic iteration, more and more of it gets mapped to the outside of the original square and so disappears from the picture.


$$h(T) \geq h(S),$$

and so, in particular, topologically conjugate maps share the same entropy. For example, the tent map is topologically conjugated to the quadratic map


$$x \mapsto 4x(1 - x)$$

of the unit interval. Therefore, the latter also has entropy $\log_2 2 = 1$.

In higher dimensions, of course, the situation is no longer as simple. In a certain sense, nevertheless, horse-shoes² (see also the discussion at the end of 3.7) as indicated in the figure



i.e.



are typical building blocks for constructing maps of positive entropy. See also the discussion in 7.2. Horse-shoes are also called baker's maps³ since the kneading of dough visualizes the iteration of such maps.

The entropy of a continuous map of a compact state space can only be positive if at each step new information is generated. One possibility to visualize this generation of information is to consider, for a given covering or partition (depending on whether we study the topological or the measure theoretic

² This is somewhat different from the horse-shoe introduced in 3.7, but the purpose of this section is to stress the similarities. The horse-shoe can easily be made into a bijective self-map of the square whereas Smale's horse-shoe as considered in 3.7 and 7.2 maps part of the square to its outside so that it leaves the picture, and the chaotic phenomena can become transient.

³ Again, the terminology is somewhat ambiguous in the literature; often a somewhat different map, but with similar properties, is called baker's map.

entropy), points $x \neq y$ that are not contained in different sets and so cannot be distinguished by any observation based on the given covering or partition. In order that T generates information, there must exist some $n \in \mathbb{N}$ for which $T^n x$ and $T^n y$ lie in different such sets so that the dynamical iterates of x and y eventually can be distinguished by an observation. This can also be interpreted as sensitive dependence on the initial conditions in the sense of chaotic dynamics. An essential point here, however, is that the distance between $T^n x$ and $T^n y$ cannot be arbitrarily increased for larger and larger n , because the state space is compact. Thus, in addition to the diverging orbits, there must also exist converging ones. This interplay makes chaotic dynamics interesting and subtle.

If one only considers the orbit of a single point $x \in X$ under the iterates of T , the generation of information can be interpreted in the sense that the orbit $T^n x$ never traverses a regular periodic sequence of the sets A_1, \dots, A_m . Thus, knowing the sets containing $x, Tx, \dots, T^n x$ never suffices for predicting the one containing $T^{n+1}x$. Otherwise, asymptotically, no new information would be generated, and the entropy of T would vanish.

4.3 Symbolic dynamics

We have already noted an analogy between the tent map

$$Z : [0, 1] \rightarrow [0, 1]$$

$$Z(x) = \begin{cases} 2x & \text{for } 0 \leq x \leq \frac{1}{2} \\ 2(1-x) & \text{for } \frac{1}{2} \leq x \leq 1 \end{cases}$$

and the one-sided 2-shift. We want to examine this in more detail and for that purpose we establish a direct relationship between these two maps. As before, every $x \in [0, 1]$ is represented as a binary sequence $b = (b_0, b_1, b_2, \dots)$, but this time in a somewhat unusual manner:

$$b_k := b_k(x) := \begin{cases} 0 & \text{if } Z^k(x) < \frac{1}{2} \\ 1 & \text{if } Z^k(x) > \frac{1}{2}. \end{cases}$$

We neglect the case where the orbit of x under the iterates of Z contains the point $\frac{1}{2}$.

In this coding, Z operates as the one-sided 2-shift

$$T(b_0, b_1, b_2, \dots) = (b_1, b_2, \dots),$$

because the binary representation of Zx is given by

$$b_k(Zx) = \begin{cases} 0 & \text{for } Z^k(Zx) < \frac{1}{2} \\ 1 & \text{for } Z^k(Zx) > \frac{1}{2}, \end{cases}$$

i.e.,

$$b_k(Zx) = b_{k+1}(x).$$


One may show that, when neglecting the preimages of $\frac{1}{2}$ which constitute a countable and hence null set, the relation between x and the binary sequence b yields a continuous bijection between $[0, 1]$ and the space of all binary sequences. Therefore, the tent map Z is conjugate to the one-sided 2-shift T . Thus, for instance, the logistic map $F : [0, 1] \rightarrow [0, 1]$,

$$F(x) = 4x(1 - x),$$

which is conjugated to Z via $\sin^2(\frac{\pi x}{2})$ ($F(\sin^2(\frac{\pi x}{2})) = \sin^2(\frac{\pi Z(x)}{2})$), is conjugate to this shift as well.

With the help of this conjugation, i.e., by analyzing the shift, we can also study the chaos aspects of the iterates of F or Z . We employ the theorem of Hardy saying that almost every $x \in [0, 1]$ has the property that its binary representation contains every finite sequence of 0s and 1s. For such an x and any $y \in [0, 1]$ (whose orbit should not contain $\frac{1}{2}$) and $N \in \mathbb{N}$, we choose the block $(b_0, b_1, \dots, b_{N-1})$ of the first N digits of our binary representation of y and are then assured that this block occurs also somewhere in the binary representation $(\beta_0, \beta_1, \beta_2, \dots)$ of x , e.g. $(b_0, \dots, b_{N-1}) = (\beta_k, \beta_{k+1}, \dots, \beta_{k+N-1})$. This means that y and $Z^k x$ have distance at most 2^{-N} since the first N digits of their binary representations coincide. Since y and N are arbitrary, this means that the orbit of x approaches any point in $[0, 1]$ arbitrarily closely and, in fact, this even happens infinitely often. Thus, the orbit of any such x is dense in $[0, 1]$, and this holds for all x in a subset of $[0, 1]$ of full measure. On the other hand, we have already seen that orbits with initial points which are arbitrarily close together eventually separate and diverge. This example demonstrates how symbolic dynamics, i.e. the dynamics on a discrete space which is conjugate to the original dynamics defined on a continuum permits a simple analysis of the qualitative properties of that original dynamic process.

We can also use the two-sided Bernoulli shift $T : Y_2 \rightarrow Y_2$ to analyze the higher dimensional horse-shoe introduced in the previous section.



We put

$$(4.3) \quad \Lambda := \bigcap_{n \in \mathbb{Z}} F^n(S).$$

$S \cap F(S)$ consists of S_0 and S_1 . $S \cap F(S) \cap F^2(S)$ consists of 4 vertical rectangles $S_{ij} = S_i \cap F(S_j) = F(S^i) \cap F^2(S^j)$, $i, j \in \{0, 1\}$, and iteratively $\bigcap_{\nu=0}^n F^\nu(S)$ consists of 2^n exponentially thin vertical rectangles $S_{i_1, \dots, i_n} = \bigcap_{\nu=1}^n F^\nu(S^{i_\nu})$ ($i_\nu \in \{0, 1\}$). In the limit, the intersection $\bigcap_{n=1}^\infty F^n(S)$ is a family of vertical segments $\bigcap_{n=1}^\infty F^n(S^{i_n})$, more precisely, the product of a vertical segment with a Cantor set in the horizontal direction. By the same pattern, $\bigcap_{n=0}^\infty F^{-n}(S)$ is the product of a horizontal segment with a Cantor set in the vertical direction. Consequently, $\Lambda = \bigcap_{n \in \mathbb{Z}} F^n(S)$ is the product of two Cantor sets, and thus a Cantor set itself. We then consider

$$(4.4) \quad \begin{aligned} \varphi : Y_2 &\rightarrow \Lambda \\ \varphi(\dots, b_{-1}, b_0, b_1, \dots) &:= \bigcap_{n=-\infty}^\infty F^n(S^{b_n}). \end{aligned}$$

φ is a homeomorphism that conjugates the two-sided Bernoulli shift and the restriction of the diffeomorphism F to Λ . This will be utilized below in 7.2. In particular, since conjugacy maps periodic points to periodic points, and since we have observed in 4.2 that the periodic points for the Bernoulli shift are dense in Y_2 , we conclude that the periodic points for the horse-shoe F are dense in Λ .

Further details about the issues raised here can be found in [24].

5 Entropy and metric aspects of dynamical systems

5.1 The metric approach to topological entropy

We recall the definition of topological entropy from 4.1 with the help of coverings. For a covering \mathcal{A} of X , $H(\mathcal{A})$ was defined as the logarithm of the minimal number of open sets in \mathcal{A} that are needed to cover X . If \mathcal{B} is a refinement of \mathcal{A} , i.e., if \mathcal{B} is a covering of X for which every set in \mathcal{B} is contained in some set from \mathcal{A} , then

$$H(\mathcal{B}) \geq H(\mathcal{A}),$$

and thus also

$$h(\mathcal{B}, T) \geq h(\mathcal{A}, T).$$

Thus, if $(\mathcal{U}_n)_{n \in \mathbb{N}}$ is a family of coverings with the property that \mathcal{U}_{n+1} is always a refinement of \mathcal{U}_n and any covering \mathcal{A} is refined by some \mathcal{U}_n , then

$$h_{top}(T) = \lim_{n \rightarrow \infty} h(\mathcal{U}_n, T).$$

If X now carries a metric d and if $(\mathcal{U}_n)_{n \in \mathbb{N}}$ is a family of such refining coverings (i.e., \mathcal{U}_{n+1} refines \mathcal{U}_n for all n as above), with

$$\text{diam}(\mathcal{U}_n) \rightarrow 0$$

(where the diameter $\text{diam}(\mathcal{U}_n)$ is the supremum of the diameters of the elements of \mathcal{U}_n , computed via the metric d), then every covering \mathcal{A} gets refined by some \mathcal{U}_n , and we therefore have

$$h_{top}(T) = \lim_{n \rightarrow \infty} h(\mathcal{U}_n, T).$$

This indicates that, following Bowen (see [38, 57]), it is useful to define the entropy with the help of a metric d on X . For that purpose, for $x, y \in X$ and a homeomorphism $T : X \rightarrow X$, one defines a dynamical sequence of metrics

$$(5.1) \quad d_n(x, y) := \sup_{0 \leq k \leq n-1} d(T^k x, T^k y)$$

and concludes from the preceding remark that

$$h_{top}(T) = \lim_{\varepsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log_2 (\text{minimal number of} \\ \text{distance balls of radius } \varepsilon \\ \text{with respect to } d_n \text{ that cover } X).$$

In particular, the quantity on the right-hand side does not depend on the metric d (assuming, as always, that X is compact). This can also be seen by observing that this quantity is the same for uniformly equivalent metrics and, on a compact space, any two metrics are uniformly equivalent.

A distance ball is a set of the form

$$B_n(x, \varepsilon) := \{y \in X : d(T^k x, T^k y) \leq \varepsilon \text{ for } 0 \leq k \leq n-1\}.$$

We might as well use open balls but, for the sequel, closed balls will be a little more convenient.

We say that a subset F of X is (n, ε) -spanning if, for every $y \in X$, there exists some $x \in F$ with

$$d_n(x, y) \leq \varepsilon,$$

i.e.,

$$d(T^k x, T^k y) \leq \varepsilon \text{ for } 0 \leq k \leq n-1.$$

For understanding this definition, it is useful to consider orbits $\{x, Tx, \dots, T^{n-1}x\}$ of length n of points of X under the map T . If distances can be measured at most with precision ε , then the orbits of two points x, y with

$$d_n(x, y) \leq \varepsilon$$

cannot be distinguished. Thus, an (n, ε) -spanning set F for every point $y \in X$ contains a point $x \in F$, whose orbit of length n cannot be distinguished from the one of y through measurements with precision at most ε .

The centers of a covering of X by balls of radius ε with respect to d_n thus constitute such an (n, ε) -spanning set, and the above minimal number of such balls that cover X is then

$$r(n, \varepsilon) := \text{minimal cardinality of} \\ (n, \varepsilon)\text{-spanning sets.}$$

According to our above interpretation, $r(n, \varepsilon)$ is thus a measure for the number of orbits $\{x, Tx, \dots, T^{n-1}x\}$ of length n with precision ε .

Therefore

$$(5.2) \quad h_{top}(T) = \lim_{\varepsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log_2 r(n, \varepsilon).$$

We now call a set $E \subset X$ (n, ε) -separated if, for any two different $x, y \in E$, there exists some $k \in \{0, \dots, n-1\}$ with

$$d(T^k x, T^k y) > \varepsilon.$$

We then write

$$s(n, \varepsilon) := \text{maximal cardinality of} \\ (n, \varepsilon)\text{-separated sets.}$$

We have

$$s(n, \varepsilon) \geq r(n, \varepsilon),$$

because an (n, ε) -separated set of maximal cardinality must be (n, ε) -spanning since otherwise we could find another element of X that is (n, ε) -separated from those of the given set.

On the other hand,

$$s(n, \varepsilon) \leq r(n, \frac{\varepsilon}{2}),$$

because we can choose any $(n, \frac{\varepsilon}{2})$ -spanning set F and then, for every (n, ε) -separated set E and $x \in E$, find precisely one $y \in F$ with

$$d_n(x, y) \leq \frac{\varepsilon}{2},$$

since F is $(n, \frac{\varepsilon}{2})$ -spanning, and in such a manner that different elements in E correspond to different elements in F , as E is (n, ε) -separated.

Of course, $s(n, \varepsilon)$ can also be considered as a measure for the number of orbits $\{x, \dots, T^{n-1}x\}$ of length n with precision ε .

Altogether, these inequalities between r and s imply that

$$(5.3) \quad h_{top}(T) = \lim_{\varepsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log_2 s(n, \varepsilon).$$

The various different characterizations of $h_{top}(T)$ turn out to be useful for its computation.

The preceding constructions can also be carried through for maps that are somewhat more general than homeomorphisms. For example, we can consider the m -fold covering

$$\phi_m : S^1 \rightarrow S^1$$

(we may represent S^1 as the unit circle in \mathbb{C} and restrict the complex map $z \mapsto z^m$ to S^1). For this map, $s(n, \varepsilon)$ and $r(n, \varepsilon)$ behave like εm^n , and thus we have

$$h_{top}(\phi_m) = \log_2 m > 0 \quad \text{for } m > 1.$$

It is not too difficult to show that, for every Lipschitz map

$$T : X \rightarrow X$$

of a compact metric space X with Lipschitz constant

$$L(T) := \sup_{x \neq y \in X} \frac{d(Tx, Ty)}{d(x, y)},$$

the entropy is bounded by

$$\dim X \max(0, \log_2 L(T))$$

where the dimension $\dim X$ has to be defined appropriately, in such a manner of course that it coincides with the usual dimension if X happens to be a differentiable manifold. For that purpose, one defines $b(\varepsilon)$ as the minimal cardinality of a covering of X by balls of radius ε and then

$$\dim X := \limsup_{\varepsilon \rightarrow 0} \frac{\log b(\varepsilon)}{|\log \varepsilon|}.$$

More information on the metric approach to topological entropy and technical details can be found in [24, 38, 57].

5.2 Complexity and intrinsic scales

For maps of positive entropy on a compact space, there cannot be any intrinsic length scale because the entropy can only be positive if infinitesimal information gets enlarged further and further. Therefore, for instance all gradient flows on compact state spaces have vanishing entropy.¹

For this reason, we wish to introduce a notion from [20] that can evaluate the dynamic behavior of maps of vanishing entropy and determine an intrinsic length scale for them. Let (X, d) be a metric space, typically compact, and let

$$T : X \rightarrow X$$

be a continuous map. Let ε be positive; ε will assume the role of a scale. As before, $E \subset X$ is called ε -separated if

$$d(x, y) > \varepsilon$$

¹ There is also no intrinsic length scale for self-similar structures, i.e. structures that remain invariant under rescalings. Some fractals have this property, and this is the reason for certain similarities between the theories of chaotic systems and fractals.

for all $x, y \in E$ with $x \neq y$, and we let

$$s(\varepsilon) := \max\{ \text{cardinality}(E) : E \subset X \text{ } \varepsilon\text{-separated} \}$$

be the maximal number of ε -separated elements in X . $E \subset X$ is called ε -divergent if, for any two $x \neq y \in E$ with $d(x, y) \leq \varepsilon$ and every $n_0 \in \mathbb{N}$, there exists some $n \geq n_0$ with

$$d(T^n x, T^n y) > \varepsilon.$$

In other words, two different elements of E have to be either initially separated by an amount ε or there have to exist arbitrarily large iterates of our transformation T that ε -separates them. In particular, this definition is thus set up in such a manner that any ε -separated set is automatically ε -divergent. The decisive point, however, is that elements that are initially closer than ε have to be separated by a dynamic iteration, in order to contribute. (A variant of our definition would require that for all $x \neq y \in E$ with $d(x, y) \leq \varepsilon$, we have

$$d(T^n x, T^n y) > \varepsilon$$

for all sufficiently large $n \in \mathbb{N}$, i.e., elements stay separated after sufficiently many iterations. This appears somewhat restrictive, however.)

Let then

$$\sigma(\varepsilon) := \max\{ \text{cardinality}(E) : E \subset X \text{ } \varepsilon\text{-divergent} \}.$$

As noted above, any ε -separated set is also ε -divergent, and hence

$$\sigma(\varepsilon) \geq s(\varepsilon).$$

We then define the complexity of $T : X \rightarrow X$ as

$$c(T) := \sup_{\varepsilon > 0} \log \frac{\sigma(\varepsilon)}{s(\varepsilon)}.$$

In particular, $c(T)$ is always non-negative.

In contrast to the entropy, the supremum here need not be realized for $\varepsilon \rightarrow 0$, but can be attained at some positive value of ε . That value then yields the scale of largest complexity for our transformation T . The complexity $c(T)$ vanishes if T is contracting in the sense that

$$d(Tx, Ty) \leq d(x, y) \text{ for all } x, y \in X.$$

For gradient flows, complexity is generated at critical points of positive index. Such a critical p remains invariant itself under the flow, but some of its neighboring points are moved into different directions, and in general there

are the more such directions the higher the index (this is only a rough intuitive correspondence as one may construct particular situations where it does not hold.) An ε of largest complexity will then be of the order of magnitude of the average distance between such critical points. ε should not be too large because otherwise points would not get dynamically separated, but not too small either because otherwise there would exist too many ε -separated points already so that the denominator $s(\varepsilon)$ would be too large. In any case, critical points of higher index constitute those places where complexity in the sense of the definition is produced, or, expressed in different words, in whose neighborhood one obtains the most information about the qualitative behavior of the dynamical system.

$c(T)$ is invariant under isometries, but not necessarily under homeomorphisms of X , for the reason that the choice of the metric entails a uniform scale ε on all of X . This might not always be desirable, and so we now formulate a notion of complexity for maps of topological spaces that permits varying scales.

Thus, let X now be a compact topological space, and let

$$T : X \rightarrow X$$

be continuous again.

For a covering of X , let

$$\nu(\mathcal{U}) := \text{minimal cardinality of all subcoverings of } \mathcal{U};$$

this should be considered as an analogue of $s(\varepsilon)$ above. For a covering \mathcal{U} , we define a dynamic covering

$$\mathcal{U}_T := \bigcup_{\substack{U_n \in \mathcal{U}, n_0 \in \mathbb{N} \\ (n=0, n_0, n_0+1, \dots)}} (U_0 \cap T^{-n_0} U_{n_0} \cap T^{-(n_0+1)} U_{(n_0+1)} \cap \dots).$$

If $x, y \in X$ happen to be contained in the same intersection, i.e.

$$T^n x, T^n y \in U_n \text{ for } n = 0 \text{ and } n = n_0, n_0 + 1, \dots,$$

and if this is so at every step, i.e., for every choice of U_n , then x and y cannot be separated by the dynamical iteration. The different elements of \mathcal{U}_T thus correspond to the possibilities of dynamical separation of the points in X . We now define the (topological) complexity of T as

$$\gamma(T) := \sup_{\substack{\mathcal{U} \text{ open} \\ \text{covering of } X}} \log \frac{\nu(\mathcal{U}_T)}{\nu(\mathcal{U})}.$$

A covering for which this supremum is achieved is then a covering of maximal complexity. The size² of the individual elements of such a covering then

² provided that we can measure this quantity, for example if we happen to have a metric or a measure

yields the local scale of highest complexity of T and, in contrast to the preceding metric concept, it can now vary across X . In contrast to $c(T)$, $\gamma(T)$ is invariant under homeomorphisms of X .

In contrast to the entropy, the complexity as defined here is a relative quantity, because we form the quotient of two quantities one of which measures the static separation property of the space and the other the dynamic separation property of the process. Of course, dividing by $s(\varepsilon)$ or $\nu(U)$ can also be considered as a normalization, in analogy to the fact that, for the definition of the entropy, one divides by a factor that measures the number of dynamic iterations. Finally, the complexity as defined here is an intrinsic quantity in the sense that the disjoint union of several identical systems has the same complexity as each individual one.

6 Entropy and measure theoretic aspects of dynamical systems

6.1 Probability spaces and measure preserving maps

This section contains some technical foundations for the sequel.

Let X be a set. A σ -algebra of subsets of X is a set \mathcal{B} of subsets of X satisfying:

- (i) $X \in \mathcal{B}$.
- (ii) If $B \in \mathcal{B}$, then so is $X \setminus B$.
- (iii) If $B_n \in \mathcal{B}$ for all $n \in \mathbb{N}$, then so is $\bigcup_{n \in \mathbb{N}} B_n$.

These properties imply:

- (iv) $\emptyset \in \mathcal{B}$.
- (v) If $B_1, \dots, B_m \in \mathcal{B}$, then so is $\bigcap_{j=1}^m B_j$.

(X, \mathcal{B}) is then called a measurable space.

If X is a topological space, we shall always consider the σ -algebra of Borel sets, i.e. the smallest σ -algebra containing all open subsets of X .

A probability measure on (X, \mathcal{B}) is a function

$$\mu : \mathcal{B} \rightarrow [0, 1]$$

satisfying:

- (i) $\mu\left(\bigcup_{n \in \mathbb{N}} B_n\right) = \sum_{n \in \mathbb{N}} \mu(B_n)$, if $B_i \cap B_j = \emptyset$ for all $i \neq j$, i.e., if the sets B_n are pairwise disjoint.
- (ii) $\mu(X) = 1$.
- (i) implies
- (iii) $\mu(\emptyset) = 0$.

A triple (X, \mathcal{B}, μ) with the preceding properties is called a probability space. We let $M(X)$ be the set of all probability measures on (X, \mathcal{B}) . $M(X)$ contains the Dirac measures supported at the points of X . $M(X)$ is convex because the convex combination of two probability measures is again a probability measure.

A transformation $T : X \rightarrow X$ of a probability space (X, \mathcal{B}, μ) is called measurable if for all $A \in \mathcal{B}$ also $T^{-1}(A) \in \mathcal{B}$ (i.e. $T^{-1}\mathcal{B} \subset \mathcal{B}$). A continuous transformation of a topological space is always measurable for any (Borel) measure because the preimages of open sets are open, and so the preimages of Borel sets are Borel.

T is called measure preserving if

$$(6.1) \quad \mu(T^{-1}(A)) = \mu(A)$$

for all $A \in \mathcal{B}$. The reason that we take T^{-1} here instead of T itself is that T need not be injective. For example, we can consider multiple coverings of the unit circle $S^1 = \{z \in \mathbb{C} : |z| = 1\}$,

$$T : S^1 \rightarrow S^1 \\ z \mapsto z^n \text{ for } n \in \mathbb{Z}.$$

All these transformations are measure preserving for the Lebesgue (=Hausdorff=Haar) measure on S^1 in the sense of (6.1), even though they are not injective for $n \neq \pm 1$.

We can now define isomorphy and conjugation in the measure theoretic sense. To do this, let $(X_i, \mathcal{B}_i, \mu_i)$ be probability spaces, $T_i : X_i \rightarrow X_i$, $i = 1, 2$, measure preserving.

Definition 17. (i) T_1 and T_2 are called isomorphic if there exist $M_i \subset \mathcal{B}_i$ with

$$\mu_i(M_i) = 1 \text{ and } T_i(M_i) \subset M_i \text{ for } i = 1, 2$$

and an invertible, measure preserving

$$\varphi : M_1 \rightarrow M_2$$

with

$$\varphi T_1 x = T_2 \varphi x \text{ for all } x \in M_1$$

(ii) If the above φ is not necessarily invertible, then T_2 is called a factor of T_1 .

(Example: If S and T are measure preserving, then S is a factor of $S \times T$).

$A, B \in \mathcal{B}$ are said to be measure equivalent if $\mu(A \Delta B) = 0$.¹ The space $\tilde{\mathcal{B}}$ of equivalence classes constitutes a Borel σ -algebra with induced measure $\tilde{\mu}$ ($\tilde{\mu}(\tilde{B}) = \mu(B)$). $(\tilde{\mathcal{B}}, \tilde{\mu})$ is called a measure algebra.

¹ Here, $A \Delta B$ denotes the symmetric difference $(A \setminus B) \cup (B \setminus A)$.

A measure preserving map $\psi : X_1 \rightarrow X_2$ induces a homomorphism of measure algebras

$$\begin{aligned} \tilde{\psi}^{-1} : (\tilde{\mathcal{B}}_2, \tilde{\mu}_2) &\rightarrow (\tilde{\mathcal{B}}_1, \tilde{\mu}_1) \\ \tilde{B} &\mapsto \widetilde{\psi^{-1}B}. \end{aligned}$$

Definition 18. T_1 and T_2 are called *conjugate* if there exists an isomorphism

$$\Phi : (\tilde{\mathcal{B}}_2, \tilde{\mu}_2) \rightarrow (\tilde{\mathcal{B}}_1, \tilde{\mu}_1)$$

of measure algebras with

$$\Phi \tilde{T}_2^{-1} = \tilde{T}_1^{-1} \Phi.$$

If Φ is only assumed to be a homomorphism, then T_2 is called the *semi-conjugate image* of T_1 .

6.2 Ergodicity

General assumption:

Let $T : X \rightarrow X$ be a measure preserving transformation of the probability space (X, \mathcal{B}, μ) , i.e., for all $E \in \mathcal{B}$, we have $T^{-1}E \in \mathcal{B}$ and $\mu(T^{-1}E) = \mu(E)$.

We have the

Theorem 24. (Poincaré Recurrence Theorem)

Let $E \in \mathcal{B}$ with $\mu(E) > 0$. The iterates $(T^n)_{n \in \mathbb{N}}$ then map almost every $x \in E$ infinitely often back into E .

Proof. $F := E \cap \bigcap_{N=0}^{\infty} \bigcup_{n=N}^{\infty} T^{-n}E \subset E$.

For each $x \in F$, there exist arbitrarily large $m \in \mathbb{N}$ with $T^m x \in E$, and the iterates of T therefore map any such x arbitrarily often to E .

Since $T^{-1}(\bigcup_{n=N}^{\infty} T^{-n}E) = \bigcup_{n=N+1}^{\infty} T^{-n}E \subset \bigcup_{n=N}^{\infty} T^{-n}E$ and T is measure preserving, it follows that

$$\begin{aligned} \mu \left(\bigcup_{n=0}^{\infty} T^{-n}E \right) &= \mu \left(\bigcup_{n=N}^{\infty} T^{-n}E \right) \text{ for all } N \\ &= \mu \left(\bigcap_{N=0}^{\infty} \bigcup_{n=N}^{\infty} T^{-n}E \right) \end{aligned}$$

and, because of $E \subset \bigcup_{n=0}^{\infty} T^{-n}E$, therefore

$$\mu(F) = \mu(E).$$

q.e.d.

If $T^{-1}E = E$ for some $E \in \mathcal{B}$, then also $T^{-1}(X \setminus E) = X \setminus E$, and the dynamic behavior of T on X can therefore be reduced to the one on the two subsets E and $X \setminus E$. If $0 < \mu(E) < \mu(X) (= 1)$, this simplifies the situation, while in the cases $\mu(E) = 0$ or $\mu(E) = 1$, one may omit the null sets E or $X \setminus E$, respectively, without changing the measure theoretic behavior of T .

Definition 19. T is called ergodic if for every $E \in \mathcal{B}$ with $T^{-1}E = E$ either $\mu(E) = 0$ or $\mu(E) = 1$.

Theorem 25. T is ergodic if and only if any $f \in L^2(\mu)$ satisfying $f(Tx) = f(x)$ for almost all $x \in X$ is constant (almost everywhere).

Proof. \Rightarrow : We consider the sets

$$E(k, n) := \{x \in X \mid 2^{-n}k \leq f(x) < 2^{-n}(k + 1)\} \quad (k \in \mathbb{Z}, n \in \mathbb{N}).$$

These sets remain invariant under T up to null sets if $f(Tx) = f(x)$ for almost all x , and therefore their measure is either 0 or 1. Since for fixed $n \in \mathbb{N}$, X is the disjoint union of the sets $E(k, n), k \in \mathbb{Z}$, precisely one of them has measure 1, with index $k = k(n)$. Then $E := \bigcap_{n \in \mathbb{N}} E(k(n), n)$ is likewise of measure 1, and f is constant on E , and hence almost everywhere on X .

\Leftarrow : Let $T^{-1}E = E$ for some $E \in \mathcal{B}$. The characteristic function χ_E is then of class $L^2(\mu)$, and $\chi_E(Tx) = \chi_E(x)$ for all $x \in X$. Therefore $\chi_E = \text{const}$, and so $\chi_E = 0$ or 1. Thus $\mu(E) = \int \chi_E d\mu = 0$ or 1.

q.e.d.

Examples:

1. $S^1 = \{|z| = 1\} \subset \mathbb{C}$ with the Lebesgue (=Hausdorff=Haar) measure,

$$\begin{aligned} T : S^1 &\rightarrow S^1 \\ z &\mapsto a \cdot z \text{ for } a \in S^1 \end{aligned}$$

is ergodic precisely if a is not a root of unity, that is, if T is not periodic.

2. Bernoulli shift

Let $Y := \{0, 1, \dots, \ell - 1\}$ be our state space,

$(p_0, \dots, p_{\ell-1})$ with $p_i > 0$ for all i and $\sum_{i=0}^{\ell-1} p_i = 1$ yields a probability measure p on Y .

Let $X := Y^{\mathbb{Z}}$ (the space of two-sided sequences with values in Y), and let \mathcal{B} be generated by sets of the form

$$\begin{aligned} A = \{ (x_n)_{n \in \mathbb{Z}} : & \quad x_m \in Y \text{ for } m < n_1 \text{ or } m > n_2, \\ & \quad x_{n_1} \in A_{n_1}, x_{n_1+1} \in A_{n_1+1}, \dots, x_{n_2} \in A_{n_2} \}, \\ n_1 \leq n_2, \quad & A_i \subset Y, \end{aligned}$$

the so-called cylinder sets, with the product measure μ determined by p , i.e.,

$$(6.2) \quad \mu\{x : x_{i_1} = j_1, \dots, x_{i_k} = j_k\} = p_{j_1} p_{j_2} \dots \cdot p_{j_k} \quad (j_i \in Y \text{ for all } i).$$

The shift $\sigma : X \rightarrow X$

$$(x_n)_{n \in \mathbb{Z}} \mapsto (y_n)_{n \in \mathbb{Z}} \text{ with } y_n = x_{n+1}$$

is then measure preserving on X .

Claim: σ is ergodic.

Proof. Let $E \in \mathcal{B}$ with $\sigma^{-1}E = E$. Intuitively, the argument is that, if at the position i , the value of all the sequences contained in E is restricted, for example $x_i = j$, then the shift invariance of E implies that this is then also so at any other position. Any such set that restricts the values of all the sequence elements, however, has vanishing measure by (6.2). On the other hand, if there are no such restrictions, then E has to be the whole space. It is also instructive to look at the formal argument: For $\varepsilon > 0$ we find an A of the form described above with $\mu(E \Delta A) < \varepsilon$, because $E \in \mathcal{B}$. Therefore

$$|\mu(E) - \mu(A)| > \varepsilon.$$

Let A be as above, $n_0 > n_2 - n_1$, $B = \sigma^{-n_0}A$. Then $\mu(B) = \mu(A)$ and

$$\mu(B \cap A) = \mu(B)\mu(A) = \mu(A)^2.$$

Since $\sigma^{-1}E = E$, we have

$$\mu(E \Delta B) = \mu(\sigma^{-n_0}E \Delta \sigma^{-n_0}A) = \mu(E \Delta A).$$

Therefore, we also have $\mu(E \Delta (A \cap B)) < 2\varepsilon$ and

$$|\mu(E) - \mu(A \cap B)| < 2\varepsilon$$

and

$$\begin{aligned} |\mu(E) - \mu(E)^2| &\leq |\mu(E) - \mu(A \cap B)| + |\mu(A \cap B) - \mu(E)^2| \\ &\leq 2\varepsilon + |\mu(A)^2 - \mu(E)^2| \\ &\leq 2\varepsilon + \mu(A)|\mu(A) - \mu(E)| \\ &\quad + \mu(E)|\mu(A) - \mu(E)| \\ &< 4\varepsilon \\ \Rightarrow \mu(E) &= \mu(E)^2 \Rightarrow \mu(E) = 0 \text{ or } 1. \end{aligned}$$

q.e.d.

3. Markov shift:

$X = Y^{\mathbb{Z}}$ and σ as in Example 2, but this time we use a different measure:

Let $A = (a_{ij})_{i,j=0,\dots,n-1}$ with $a_{ij} \geq 0$ for all i, j , $\sum_{j=0}^{n-1} a_{ij} = 1$ for all i .

Then there exists

$$p = (p_0, \dots, p_{n-1}) \text{ with } p_i > 0 \text{ for all } i, \sum_{i=0}^{n-1} p_i = 1,$$

$$pA = p \text{ (i.e. } \sum_{i=0}^{n-1} p_i a_{ij} = p_j \text{ for all } j).$$

$\mu_A\{x : x_i = j_0, x_{i+1} = j_1, \dots, x_{i+k} = j_k\} := p_{j_0} a_{j_0 j_1} \dots a_{j_{k-1} j_k}$ (p_{j_0} tells us the probability of the occurrence of the symbol j_0 , and $a_{j_\ell j_{\ell+1}}$ is the transition probability from symbol j_ℓ to symbol $j_{\ell+1}$.) μ_A can be extended as a probability measure to (X, \mathcal{B}) . σ leaves μ_A invariant and is called (p, A) -Markov shift. (Example 2 is a special case of Example 3 with $a_{ij} = p_j$ for all i, j) σ is ergodic on (X, \mathcal{B}, μ_A) if and only if A is irreducible, i.e., for all i, j , there exists $m \in \mathbb{N}$ with $a_{ij}^{(m)} > 0$ ($A^m = (a_{ij}^{(m)})$ for $m \in \mathbb{N}$).

Theorem 26. (Birkhoff ergodic theorem) *Let $f \in L^1(\mu)$. Then $\frac{1}{n} \sum_{i=0}^{n-1} f(T^i x)$ converges for almost all x to some $f^* \in L^1(\mu)$, with $f^* \circ T = f^*$ almost everywhere, and $\int f^* d\mu = \int f d\mu$.*

If T is ergodic, then f^ is constant almost everywhere, and thus*

$$f^* = \int f d\mu \text{ almost everywhere.}$$

In other words: If T is ergodic, then for all $f \in L^1(\mu)$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i x) = \int f d\mu \text{ for almost all } x.$$

Analogously, for an ergodic continuous (semi) flow $(T_t)_{t \geq 0}$

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_t x) dt = \int_X f d\mu \text{ for almost all } x.$$

Thus, the temporal mean coincides with the spatial mean. The Birkhoff ergodic theorem implies that, when we consider a time series, that is, the values which some observable f takes on the orbit of some point x that happens to be the starting point for a dynamical iteration, then in the ergodic case the average of this time series converges to the expectation value of that observable. Thus, following a single dynamical orbit already samples the whole space.

Conceptually similar, but only obtaining a weaker type of convergence that does not allow the above strong implications for time series is

Theorem 27. (Von Neumann ergodic theorem)

If $f \in L^p(\mu), 1 \leq p < \infty$, then there exists $f^* \in L^p(\mu)$, with $f^* \circ T = f^*$, such that $\frac{1}{n} \sum_{i=0}^{n-1} f(T^i x) \rightarrow f^*$ in $L^p(\mu)$.

Corollary 3. T is ergodic if and only if for all $A, B \in \mathcal{B}$:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \mu(T^{-i} A \cap B) = \mu(A)\mu(B).$$

Proof. \Rightarrow : Putting $f = \chi_A$ in Birkhoff’s ergodic theorem implies that

$$\begin{aligned} \frac{1}{n} \sum_{i=0}^{n-1} \chi_A(T^i x) &\rightarrow \mu(A) \text{ almost everywhere, and hence} \\ \frac{1}{n} \sum_{i=0}^{n-1} \chi_A(T^i x)\chi_B &\rightarrow \mu(A)\chi_B \text{ almost everywhere, and hence again} \\ \frac{1}{n} \sum_{i=0}^{n-1} \mu(T^{-i} A \cap B) &\rightarrow \mu(A)\mu(B) \text{ (theorem on dominated convergence).} \end{aligned}$$

\Leftarrow : Let $T^{-1}E = E$ for some $E \in \mathcal{B}$. Choosing $A = B = E$ gives $\frac{1}{n} \sum_{i=0}^{n-1} \mu(E) \rightarrow \mu(E)^2$, and thus $\mu(E) = \mu(E)^2$ and then $\mu(E) = 0$ or 1 . q.e.d.

Definition 20. (i) T is called weakly mixing if

$$\text{for all } A, B \in \mathcal{B} : \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} |\mu(T^{-i} A \cap B) - \mu(A)\mu(B)| = 0$$

(ii) T is called strongly mixing if

$$\text{for all } A, B \in \mathcal{B} : \lim_{n \rightarrow \infty} \mu(T^{-n} A \cap B) = \mu(A)\mu(B).$$

This condition captures a decay of correlations under the iteration of T . Strong mixing is equivalently expressed by requiring that, for all L^2 -functions f, g on (X, μ) , the correlation

$$(6.3) \quad C_n(f, g) := \int f(T^n x)g(x)d\mu(x) - \int f(x)d\mu(x) \int g(x)d\mu(x)$$

converges to 0 as $n \rightarrow \infty$. One may take $f = g$ here to understand the meaning of this condition, which is that the values of the function f evaluated after n iterations of the process become independent of those of g . Note that $\int f(x)d\mu(x) = \int f(T^n x)d\mu(x)$ because T preserves the measure μ . One also speaks about exponentially mixing when that correlation decays exponentially, that is,

$$(6.4) \quad |C_n(f, g)| \leq c(f, g)\rho^n$$

for some $\rho < 1$ and some constant $c(f, g)$.

Strong mixing implies weak mixing implies ergodicity, as follows directly from Corollary 3. The converses are not true in general. Strong mixing means that, for any set A , $T^{-n}A$ becomes asymptotically independent of any other set B , weak mixing that this is so up to certain exceptional n , and ergodicity means that this property holds on average. See [57] for more details.

Example: Rotating S^1 by a factor a is ergodic if a is not a root of unity, as explained above, but not weakly mixing since, for sufficiently small intervals A, B , for most i , we have $T^{-i}A \cap B = \emptyset$, and hence $\mu(T^{-i}A \cap B) = 0$, $\mu(A)\mu(B) > 0$.

Examples that are weakly, but not strongly mixing, while generic, are harder to present explicitly. The Bernoulli shift is strongly mixing.

Good references about ergodic theory are [24, 38, 48, 57].

6.3 Entropy and information

General assumptions: (X, \mathcal{B}, μ) is a probability space, $T : X \rightarrow X$ is measure preserving, $\alpha = \{A_1, \dots, A_n\}$ is a partition of X , i.e. $\mu(A_i) > 0$ for all i , $A_i \cap A_j = \emptyset$ for $i \neq j$, and $\mu(X \setminus \bigcup_{i=1}^n A_i) = 0$.

The entropy of α is then defined as

$$H(\alpha) := - \sum_{i=1}^n \mu(A_i) \log_2 \mu(A_i).$$

Now let $\alpha = \{A_1, \dots, A_n\}$ and $\beta = \{B_1, \dots, B_m\}$ be partitions of X . We then also have the partition $\alpha \vee \beta := \{A_i \cap B_j : i = 1, \dots, n, j = 1, \dots, m\}$.

The entropy of T with respect to α is

$$(6.5) \quad h(\alpha, T) := h_\mu(\alpha, T) := \lim_{N \rightarrow \infty} \frac{1}{N} H(\alpha \vee T^{-1}\alpha \vee \dots \vee T^{-N+1}\alpha).$$

The existence of this limit is a consequence of the following observations.

1. Subadditivity:

$$\begin{aligned}
 H(\alpha \vee \beta) &= - \sum_{i,j} \mu(A_i \cap B_j) \log_2 \mu(A_i \cap B_j) \\
 &= - \sum_i \mu(A_i) \sum_j \frac{\mu(A_i \cap B_j)}{\mu(A_i)} \log_2 \frac{\mu(A_i \cap B_j)}{\mu(A_i)} \\
 &\quad - \sum_{i,j} \mu(A_i \cap B_j) \log_2 \mu(A_i) \\
 (6.6) \quad &= - \sum_j \sum_i \mu(A_i) \frac{\mu(A_i \cap B_j)}{\mu(A_i)} \log_2 \frac{\mu(A_i \cap B_j)}{\mu(A_i)} \\
 &\quad - \sum_i \mu(A_i) \log_2 \mu(A_i), \text{ since } \sum_j \mu(A_i \cap B_j) = \mu(A_i) \\
 &\leq - \sum_j (\sum_i \mu(A_i \cap B_j)) \log_2 (\sum_i \mu(A_i \cap B_j)) \\
 (6.7) \quad &\quad \quad \quad + H(\alpha),
 \end{aligned}$$

since the function² $-x \log x$ is concave (its second derivative being $-1/x$), and hence

$$- \sum_i p_i x_i \log x_i \leq - (\sum_i p_i x_i) \log (\sum_i p_i x_i)$$

for non-negative p_i with $\sum p_i = 1$. Thus, since $\sum_i \mu(A_i \cap B_j) = \mu(B_j)$, the first term in (6.7) equals $H(\beta)$, and we obtain

$$(6.8) \quad H(\alpha \vee \beta) \leq H(\beta) + H(\alpha).$$

2. Time invariance (stationarity):

$$(6.9) \quad H(T^{-1}\alpha) = H(\alpha)$$

for a measure preserving T .

3. From 1. and 2.,

$$(6.10) \quad H(\alpha \vee T^{-1}\alpha \vee \dots \vee T^{-(N+N'+1)}\alpha) \leq H(\alpha \vee \dots \vee T^{-N+1}\alpha) + H(\alpha \vee \dots \vee T^{-N'+1}\alpha).$$

For $\eta_N := H(\alpha \vee \dots \vee T^{-N+1}\alpha)$, we thus have

$$(6.11) \quad \eta_N \leq \eta_{N+1}$$

because finer partitions have higher entropy. Also, by 3.,

$$(6.12) \quad \eta_{N+N'} \leq \eta_N + \eta_{N'} \quad \text{for } N, N' \in \mathbb{N},$$

and so we obtain the existence of the limit as in 4.1.

In order to arrive at another useful interpretation of $h(\alpha, T)$, we need to introduce conditional entropies. For two measurable sets A and B , with $\mu(B) > 0$, we put

² We use the convention $0 \log 0 = 0$.

$$\mu(A|B) := \frac{\mu(A \cap B)}{\mu(B)}.$$

The so-called conditional measure $\mu(\cdot|B)$ is also a probability measure on X . Now let $\alpha = (A_1, \dots, A_n)$ and $\beta = (B_1, \dots, B_m)$ be measurable partitions of X . We then define the conditional entropy of α w.r.t. β as

$$(6.13) \quad H(\alpha|\beta) := - \sum_{j=1}^m \mu(B_j) \sum_{i=1}^n \mu(A_i|B_j) \log \mu(A_i|B_j) = \sum_{i,j} \mu(A_i \cap B_j) \log \mu(A_i|B_j).$$

We then have the following

Lemma 4. *For measurable partitions $\alpha = (A_1, \dots, A_n)$ and $\beta = (B_1, \dots, B_m)$ of X , we have*

$$(6.14) \quad H(\alpha \vee \beta) = H(\beta) + H(\alpha|\beta).$$

Moreover, if the measurable partition $\gamma = (C_1, \dots, C_l)$ refines β , then

$$(6.15) \quad H(\alpha|\gamma) \leq H(\alpha|\beta).$$

Proof. By (6.13), we have, repeating the steps leading to (6.6),

$$\begin{aligned} H(\alpha \vee \beta) &= - \sum_{i,j} \mu(A_i \cap B_j) \log \mu(A_i \cap B_j) \\ &= - \sum_{i,j} \mu(A_i \cap B_j) \log \mu(B_j) - \sum_{i,j} \mu(A_i \cap B_j) \log \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \\ &= - \sum_j \mu(B_j) \log \mu(B_j) - \sum_{i,j} \mu(A_i \cap B_j) \log \mu(A_i|B_j) \end{aligned}$$

which proves the first result. For the second one, we use once more $-\sum_k p_k y_k \log y_k \leq -\sum_k p_k y_k \log(\sum_j p_j y_j)$ for non-negative coefficients p_α with $\sum p_\alpha = 1$ (concavity of $-x \log x$). Thus,

$$\begin{aligned} H(\alpha|\gamma) &= - \sum_k \mu(C_k) \mu(A_i|C_k) \log \mu(A_i|C_k) \\ &\leq - \sum_i \sum_k \mu(A_i|C_k) \mu(C_k) \log \left(\sum_j \mu(A_i|C_j) \mu(C_j) \right) \\ &= - \sum_k \sum_i \mu(A_i \cap C_k) \log \mu(A_i) \\ &= - \sum_i \mu(A_i) \log \mu(A_i) \\ &= H(\alpha). \end{aligned}$$

When we replace the measure μ in the preceding computation by the conditional measure μ_j defined by $\mu_j(A) := \frac{\mu(A \cap B_j)}{\mu(B_j)}$ for $B_j \in \beta$, we deduce the inequality (6.15). q.e.d.

From the Lemma, we deduce that

$$\begin{aligned}
 & H(\alpha \vee \dots \vee T^{-N+1}\alpha) - H(\alpha \vee \dots \vee T^{-N+2}\alpha) \\
 &= H(T^{-N+1}\alpha | \alpha \vee \dots \vee T^{-N+2}\alpha) \\
 &\leq H(T^{-N+1}\alpha | T^{-1}\alpha \vee \dots \vee T^{-N+2}\alpha) \\
 &= H(T^{-N+2}\alpha | \alpha \vee \dots \vee T^{-N+3}\alpha) \\
 &= H(\alpha \vee \dots \vee T^{-N+2}\alpha) - H(\alpha \vee \dots \vee T^{-N+3}\alpha)
 \end{aligned}$$

where in the second-to-last step we have used the shift invariance (6.9) above. This implies that the η_N defined in 3. above satisfy

$$(6.16) \quad \eta_N - \eta_{N-1} \leq \eta_{N-1} - \eta_{N-2}$$

for all $N \in \mathbb{N}$. Therefore, we have

$$\begin{aligned}
 \lim_{N \rightarrow \infty} (\eta_N - \eta_{N-1}) &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{\nu=1}^N (\eta_{\nu+1} - \eta_{\nu}) \\
 &= \lim_{N \rightarrow \infty} \frac{1}{N} (\eta_N - \eta_1) \\
 &= \lim_{N \rightarrow \infty} \frac{\eta_N}{N}.
 \end{aligned}$$

This implies the following

Theorem 28.

$$(6.17) \quad h(\alpha, T) = \lim_{N \rightarrow \infty} (H(\alpha \vee \dots \vee T^{-N+1}\alpha) - H(\alpha \vee \dots \vee T^{-N+2}\alpha)).$$

The connection between the concept of entropy and information theory is the following: Let $p(dx)$ be a probability density for finding a particle at $x \in X$. Let \mathcal{B} be the σ -algebra of those sets that are measurable with respect to $p(dx)$. For $A \in \mathcal{B}$ we then put

$$\mu(A) := \int_A p(dx).$$

The entropy is

$$(6.18) \quad H(X, p) = - \int p(dx) \log_2 p(x).$$

In the discrete case, integrals are of course replaced by sums.

For example, if $X = \{0, \dots, 7\}$ and $p(i) = \frac{1}{8}$ for all $i \in X$, then we have

$$H(X, p) = - \sum_{i=1}^8 p(i) \log_2 p(i) = \log_2 8 = 3,$$

and this expresses the fact that, for representing the elements of X as binary numbers, we need 3 digits, i.e., 3 bits.

If, however,

$$p(0) = \frac{1}{2}, p(1) = \frac{1}{4}, p(2) = \frac{1}{8}, p(3) = \frac{1}{16}, p(4) = \dots p(7) = \frac{1}{64},$$

then

$$H(X, p) = -\frac{1}{2} \log_2 \frac{1}{2} - \frac{1}{4} \log_2 \frac{1}{4} - \frac{1}{8} \log_2 \frac{1}{8} - \frac{1}{16} \log_2 \frac{1}{16} - 4 \frac{1}{64} \log_2 \frac{1}{64} = 2.$$

We may, for example, represent the elements of X as the binary sequences 0, 10, 110, 1110, 111100, 111101, 111110, 111111 and the average (with respect to p) length of such a sequence is 2. When compared with the uniform distribution, the present one has a smaller entropy.

Thus, $H(X, p)$ represents the average information that we gain from measuring x .

In the same manner, $H(\alpha)$ represents the average information gained by measuring in which of the sets A_i the particle x lies. Thus, we ignore the precise information about the position of x and retain only the information about that set A_i that contains x . Obviously, the information, i.e. the entropy, is then the smaller the coarser the partition is.

If we now have a measure preserving transformation $T : X \rightarrow X$ and measure not only in which of the sets $A_i \in \alpha$ x lies, but also in which $A_i \in \alpha$ Tx lies, i.e., in which $T^{-1}A_j \in T^{-1}\alpha$ the point x is contained, altogether we then obtain more precise information about x . However, we wish to count only that information which is obtained additionally if we already know in which $A_i \in \alpha$ Tx lies, and this is

$$(6.19) \quad H(T^{-1}\alpha|\alpha) = H(\alpha \vee T^{-1}\alpha) - H(\alpha).$$

Asymptotically, the additional information gained when we already know in which of the sets from α $x, Tx, \dots, T^{N-1}x$ lie, and then learn in which $A_i \in \alpha$ $T^N x$ lies, is then precisely the entropy of T with respect to α . This is the content of (6.17). Similarly,

$$\frac{1}{2} H(\alpha \vee T^{-1}\alpha)$$

is the average amount of information per measurement when we perform the two measurements of x and Tx in the context of the partition α , and (6.5) then describes the entropy as the asymptotic average amount of information per measurement on the orbit of x .

Remarks:

1. The measure theoretic entropy of T defined below then comes out by choosing the partition α such that, asymptotically, these measurements yield the maximal possible amount of information.
2. Choosing the logarithm with base 2 in the definition of the entropy yields the normalization that the entropy becomes 1 when asymptotically at each step the precision of the observation is doubled.

The entropy measures the new information obtained by observations, in other words, that part of the information content of the dynamical iteration, that is, the observation of orbits, that cannot be predicted from the preceding observations. Thus, the entropy encodes the random aspects of the dynamical system. For many purposes, however, it is more important to know the predictive part, that is, how much information about future events is contained in past observations. We shall now present the analysis of Grassberger[11] of this issue. We put $b_\nu = i$ when $T^{-\nu}x \in A_i$. We observe a time series b_0, b_1, b_2, \dots , and we ask how much the uncertainty about future values can be reduced through the knowledge of past values. We let $\Lambda(N)$ be the set of all such strings $b_0, b_1, b_2, \dots, b_{N-1}$ of length N . Since the cardinality of the partition α is n , there are n^N possible such strings. We let $p(\lambda)$ be the probability with which each $\lambda \in \Lambda(N)$ occurs. Since T is assumed stationary, the probabilities do not change when we shift the time index by ν , that is, when we consider the string $b_\nu, \dots, b_{\nu+N-1}$. For $N = 1$, we simply have the possible values i of b_ν which occur with the probabilities $\mu(A_i)$. This is an unconditional probability, that is, it may change when we already know the previous value $b_{\nu-1}$, or even more preceding values. This point will now be generalized and made quantitative by considering how entropies depend on the string length N .

We have the entropy

$$(6.20) \quad H(T, N) := - \sum_{\alpha=1}^{n^N} p(\lambda) \log p(\lambda) = H(\alpha \vee T^{-1}\alpha \vee \dots \vee T^{-N+1}\alpha).$$

This is the information contained in strings $b_\nu, b_{\nu+1}, \dots, b_{\nu+N-1}$ of length N . Given such a string of length N , the additional information needed to predict the value $b_{\nu+N}$ is then

$$(6.21) \quad h(T, N) := H(T, N + 1) - H(T, N).$$

These differences cannot increase with N , because the prediction of the next value becomes easier the more previous values we already know (see the argument establishing (6.16)). In fact, by (6.17),

$$(6.22) \quad h(T) := \lim_{N \rightarrow \infty} h(T, N)$$

is the entropy $h(\alpha, T)$ of T w.r.t. the covering α . (Since we wish to consider α as fixed and given, encoding the precision of our observations, we drop it from the notation for the moment.)

Then the difference of these differences,

$$(6.23) \quad \delta h(T, N) := h(T, N - 1) - h(T, N),$$

measures the average amount by which the uncertainty of $b_{\nu+N}$ decreases due to the knowledge of b_ν , that is, one more observation value back in the past. We thus need to store at least an amount of $\delta h(T, N)$ for at least N time steps for optimal prediction. Grassberger[11] therefore defines the effective measure complexity of T as

$$(6.24) \quad \begin{aligned} h_{EMC}(T) &:= \sum_N N \delta h(T, N) = \sum_N (h(T, N - 1) - h(T, N)) \\ &= \sum_{N=0}^{\infty} (h(T, N) - h(T)). \end{aligned}$$

This expression may be infinite, that is, the series need not converge. The rate of divergence then captures an important aspect of the dynamical system. Such a divergence must be sublinear. Namely, when we write

$$(6.25) \quad H(T, N) = Nh(T) + \sigma(N)$$

we see that

$$(6.26) \quad \lim_{N \rightarrow \infty} \frac{\sigma(N)}{N} = 0$$

since $h = \lim_{N \rightarrow \infty} \frac{H(T, N)}{N}$, see (6.5). If $\sigma(N)$ converges to some value σ for $N \rightarrow \infty$, then by (6.24), since $h(T, N) = h + \sigma(N + 1) - \sigma(N)$,

$$(6.27) \quad h_{EMC} = \sum_N (\sigma(N + 1) - \sigma(N)) = \sigma - \sigma(1).$$

By (6.10), the mutual information is

$$(6.28) \quad I(T, N + N') := H(T, N) + H(T, N') - H(T, N + N') \geq 0.$$

This quantity expresses the extent to which the entropy of a string of length $N + N'$ is smaller than the sum of the entropies of strings of lengths N and N' . By the stationarity of T , we can consider subsequent strings $b_\nu, \dots, b_{\nu+N-1}$ and $b_{\nu+N}, \dots, b_{\nu+N+N'-1}$ and ask by what amount the uncertainty about the second one is reduced by knowing the first one; again, this quantity is $I(T, N + N')$, cf. (6.14). By the definition of $\sigma(N)$,

$$(6.29) \quad I(T, N + N') = \sigma(N) + \sigma(N') - \sigma(N + N'),$$

and therefore,

$$(6.30) \quad \lim_{N, N' \rightarrow \infty} I(T, N + N') = \sigma$$

when that limit exists. Thus, σ captures the mutual information between strings of infinite length. When one of them is considered as the recordings from an infinite past (that is, we push the starting point ν of the dynamical iteration back to $-\infty$; when 0 is considered as the present, we can take $\nu = -N + 1$), and the other one the observations in an infinite future, then we can say that σ measures the mutual information between the past and the future of the dynamical system. Thus, σ encodes the non-random part of the dynamical iteration, that is, the predictive power of past about future observations. In general, however, $\sigma(N)$ need not remain bounded and converge to some finite value σ . Thus, more generally, the rate of divergence of $\sigma(N)$ captures that predictive power.³ A systematic presentation of these ideas can be found in [10].

We have defined the entropy $H(\alpha)$ of a partition α as the average information obtained by determining in which of the sets in this partition a point lies. We can also express this as follows: For a point $x \in X$, we have the information w.r.t. α given by

$$(6.31) \quad I_\alpha(x) = -\log \mu(A_i(x))$$

where $A_i(x)$ is the set from the partition α that contains x . We then have

$$(6.32) \quad H(\alpha) = \int I_\alpha(x) d\mu(x).$$

Likewise, given a transformation T as above, we have

$$(6.33) \quad H(\alpha \vee T^{-1}\alpha \vee \dots \vee T^{-N+1}\alpha) = \int I_{\alpha \vee \dots \vee T^{-N+1}\alpha}(x) d\mu(x).$$

The entropy was the limit for N going to infinity:

$$(6.34) \quad h(\alpha, T) := h_\mu(\alpha, T) := \lim_{N \rightarrow \infty} \frac{1}{N} H(\alpha \vee T^{-1}\alpha \vee \dots \vee T^{-N+1}\alpha).$$

It is a remarkable and useful fact that for *ergodic* T , this entropy can already be obtained as the limit of the integrands for μ -almost all $x \in X$:

³ There is one subtle point here; namely, instead of asking for the predictive value of past observations, one can also study how much the information contained in past observations can be compressed without reducing its predictive power. That is, we ask for the minimal amount of information that needs to be stored from the past for the best possible prediction of the future. Grassberger[11] calls this the true measure complexity and points out that, in general, this can be larger than the effective measure complexity.

Theorem 29. (Shannon-McMillan-Breiman) *Let the measure preserving transformation $T : X \rightarrow X$ of the probability space (X, \mathcal{B}, μ) be ergodic, and let α be a finite (or countable) partition with finite entropy $H(\alpha)$. Then*

$$(6.35) \quad \lim_{N \rightarrow \infty} \frac{1}{N} I_{\alpha \vee \dots \vee T^{-N+1}\alpha}(x) = h(\alpha, T)$$

for μ -almost all x .

A *proof* of this result can be found in [38]. This theorem is similar in spirit to the Birkhoff ergodic theorem because it also tells us that we can obtain some global information about the whole space, as encoded here in the entropy, by making measurements along a single orbit.

We now return to the general discussion of the transformation T on the probability space (X, \mathcal{B}, μ) and define the measure theoretic (Kolmogorov-Sinai) entropy of T :

$$(6.36) \quad h(T) := h_\mu(T) := \sup_{\alpha} h(\alpha, T).$$

This eliminates the auxiliary object employed in the preceding discussion, namely the partition α of X .

The entropy of T is invariant under isomorphisms and conjugations of T . If T_2 is a factor of T_1 , then $h(T_2) \leq h(T_1)$, since from any partition of X_2 one may then construct a partition of equal entropy in X_1 .

After the preceding, the following remark should be obvious and perhaps even superfluous. A crucial point for the definition of $h(T)$ is that we first take the limit $N \rightarrow \infty$ and the supremum with respect to all partitions α only after that. If, however, we first took the supremum of $H(\alpha)$ for all α , then, in the continuous case, we would always get the answer ∞ because the information gained can be increased at will by refining the partition. A fundamental point of the notion of entropy and the content of the theorem of Kolmogorov-Sinai below is that, by applying the dynamical system, i.e. by iterating the transformation T , a given partition gets refined such that asymptotically all the available information is obtained. The result remains finite in the limit (at least in the interesting cases) since we divide by the number N of observations or iterations.

The actual computation of the entropy becomes possible by the Theorem of Kolmogorov-Sinai:

Theorem 30. *Let T be an invertible measure preserving transformation of (X, \mathcal{B}, μ) . Let \mathcal{A} be a finite subalgebra of \mathcal{B} , generated by some partition α , with*

$$\bigvee_{n=-\infty}^{\infty} T^n \mathcal{A} = \mathcal{B} \text{ up to null sets .}$$

Then

$$h(T) = h(\alpha, T).$$

If even

$$\bigvee_{n=0}^{\infty} T^{-n} \mathcal{A} = \mathcal{B} \text{ up to null sets}$$

then

$$h(T) = 0.$$

An α as in the preceding theorem is called a generator of \mathcal{B} with respect to T .

Building upon the work of Rohlin, Krieger showed that an ergodic measure preserving transformation T on a Lebesgue space always possesses some finite generator. The theorem of Kolmogorov-Sinai then allows the computation of the entropy of T .

Examples:

1. For $\text{id} : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ we have

$$h(\text{id}) = 0,$$

$$\text{since } h(\text{id}, \alpha) = \lim_{n \rightarrow \infty} \frac{1}{n} H(\alpha) = 0 \text{ for all } \alpha.$$

2. If $T : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is periodic, i.e.

$$T^k = \text{id} \text{ for some } k \in \mathbb{N},$$

then likewise

$$h(T) = 0, \text{ since by Example 1, } 0 = h(T^k) = kh(T).$$

In particular, every measure preserving transformation of a finite state space has vanishing entropy.

- 3.

$$T : S^1 \rightarrow S^1$$

$$z \mapsto az \text{ for } a \in S^1.$$

If a is a root of unity, then $h(T) = 0$ by 2.

If a is not a root of unity, then $(a^{-n})_{n \in \mathbb{N}}$ is dense in S^1 . Let \mathcal{A} consist of \emptyset , S^1 , and the upper and the lower semicircle. Since $(a^{-n})_{n \in \mathbb{N}}$ is dense in S^1 , every semicircle, and therefore also every interval in S^1 , belongs to $\bigvee_{n=0}^{\infty} T^{-n} \mathcal{A}$. Thus, $\bigvee_{n=0}^{\infty} T^{-n} \mathcal{A} = \mathcal{B}$, and hence $h(T) = 0$ by the second part of the theorem of Kolmogorov-Sinai.

4. The Bernoulli shift T belonging to $Y = \{0, 1, \dots, \ell - 1\}$, $(p_0, \dots, p_{\ell-1})$ has entropy

$$-\sum_{j=0}^{\ell-1} p_j \log_2 p_j$$

Proof. Let

$$A_j := \{(x_n)_{n \in \mathbb{Z}} : x_0 = j\}.$$

$A_0, \dots, A_{\ell-1}$ then constitute a partition of $X = Y^{\mathbb{Z}}$, and we denote by \mathcal{A} the subalgebra generated by this partition. We have

$$\bigvee_{n=-\infty}^{\infty} T^n \mathcal{A} = \mathcal{B} \quad (= \text{product-}\sigma\text{-algebra}).$$

By the theorem of Kolmogorov-Sinai,

$$h(T) = h(T, \mathcal{A}).$$

For computing $h(T, \mathcal{A})$, we consider sets of the form

$$\begin{aligned} & A_{i_0} \cap T^{-1} A_{i_1} \cap \dots \cap T^{-(N-1)} A_{i_{N-1}} \\ &= \{(x_n)_{n \in \mathbb{Z}} : x_0 = i_0, \dots, x_{N-1} = i_{N-1}\}. \end{aligned}$$

The measure of such a set is $p_{i_0} \cdot p_{i_1} \cdot \dots \cdot p_{i_{N-1}}$.

Therefore

$$\begin{aligned} & \frac{1}{N} H(\mathcal{A} \vee T^{-1} \mathcal{A} \vee \dots \vee T^{-N+1} \mathcal{A}) \\ &= -\frac{1}{N} \sum_{\substack{i_0, \dots, i_{N-1} \\ =0}}^{\ell-1} p_{i_0} \cdot \dots \cdot p_{i_{N-1}} \log(p_{i_0} \cdot \dots \cdot p_{i_{N-1}}) \\ &= -\frac{1}{N} \sum_{\substack{i_0, \dots, \\ i_{N-1}}} p_{i_0} \cdot \dots \cdot p_{i_{N-1}} (\log p_{i_0} + \dots + \log p_{i_{N-1}}) \\ &= -\sum_{i=0}^{\ell-1} p_i \log p_i \end{aligned}$$

and, because of $h(T, \mathcal{A}) = \lim_{N \rightarrow \infty} \frac{1}{N} H(\mathcal{A} \vee \dots \vee T^{-N+1} \mathcal{A})$, the claim follows.

q.e.d.

5. For a Markov shift, the entropy is computed to be

$$-\sum_{i,j} p_i a_{ij} \log_2 a_{ij}.$$

This is the average information about the next symbol when reading a two-sided infinite sequence of symbols with the given transition probabilities from left to right. The summation with respect to i expresses the averaging over the symbols that can possibly occur at for instance time $t = 0$, and the summation with respect to j then contains the induced probability for the next symbol (at time $t = 1$, when time is expressed in steps of unit length). Thus, knowing the symbol i at time $t = 0$, the uncertainty about the next symbol is

$$-\sum_j a_{ij} \log_2 a_{ij}.$$

This does not depend on the symbols which occurred at preceding times, in particular, but only on the symbol i at time $t = 0$. This is the so-called Markov property, namely, that the future states depend only on the present, but not directly on the past anymore. If $a_{ij} = p_j$ for all i and j , then we have a Bernoulli shift because, in that case, the occurrence probability of the next symbol is independent of the present one, and the above formula for the entropy reduces to the one for the Bernoulli shift, i.e.

$$-\sum_j p_j \log_2 p_j, \quad \text{since } \sum_i p_i = 1.$$

Conjugate transformations share the same entropy. Thus, entropy can be used to distinguish between transformations that are not conjugate. It is a quite remarkable fact that, for Bernoulli shifts, the entropy is already a complete such invariant:

Theorem 31. (Ornstein) *If two Bernoulli shifts have the same entropy, they are conjugate to each other (and hence isomorphic).*

Here, the state space can be a Lebesgue space, i.e., isomorphic to a probability space which is the disjoint union of at most countably many points y_1, y_2, \dots with positive measure p_n of y_n , and an interval $[0, s]$, $s = 1 - \sum_n p_n$ with the standard Lebesgue measure.

The connection between the (topological) entropy introduced in Chapter 4 and the (measure theoretic) entropy introduced in the present chapter is established by

Theorem 32. *Let $T : X \rightarrow X$ be a homeomorphism of the compact metric space X . Then*

$$h_{top}(T) = \sup\{h_\mu(T) : \mu \text{ } T\text{-invariant Borel probability measure on } X\}.$$

This implies, for instance, by our computation of the measure theoretic entropy above, directly that the Bernoulli shift on the set $Y = \{0, 1, \dots, \ell - 1\}$ has topological entropy $\log_2 \ell$, since $-\sum_{j=0}^{\ell-1} p_j \log p_j$ under our conditions $0 \leq p_j \leq 1$ and $\sum_{j=0}^{\ell-1} p_j = 1$ is maximal precisely if $p_j = \frac{1}{\ell}$ for all j . (Cf. 4.2)

Finally, the following result yields a connection between the measure theoretic entropy and the constructions of the previous chapter.

Theorem 33. (Katok) *Let $T : X \rightarrow X$ be a homeomorphism of the compact metric space X , and let μ be a T -invariant ergodic Borel probability measure on X . For $n \in \mathbb{N}$, $\varepsilon > 0$, $\delta > 0$ let $r_\mu(n, \varepsilon, \delta)$ be the minimal number of balls of radius ε with respect to the metric d_n , whose union is of μ -measure at least $1 - \delta$, i.e., that cover up to a set of measure at most δ . Then*

$$(6.37) \quad h_\mu(T) = \lim_{\delta \rightarrow 0} \lim_{\varepsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log_2 r_\mu(n, \varepsilon, \delta).$$

Proofs of those results about the measure theoretic entropy stated here without proof can be found in the standard references [24, 38, 48, 57].

6.4 Invariant measures

In the preceding sections, we have looked at transformations T of a probability space that preserve the underlying measure μ and, in Theorem 32, we have taken a supremum over all the probability measures on the underlying metric space X preserved by the transformation T in order to compute the topological entropy. This naturally leads to the following questions:

- Given a transformation T , say a homeomorphism of a compact metric space X , does there exist a probability measure on X that is preserved by T ?
- While the way Theorem 32 has been formulated seems to imply that the answer to the preceding question is “yes”, one will then wish to know how many T -invariant probability measures exist, and whether, for example, the supremum in that Theorem is achieved, and whether, if the answer is positive, the resulting invariant measures can be distinguished by special properties, and thus recognized as the ones from which the topological entropy can be computed.
- How do properties like ergodicity depend on the choice of an invariant measure?
- Do such invariant measures reveal interesting structural properties of the transformation T ?

The first question is answered by

Theorem 34. (Krylov-Bogolubov) *Let $T : X \rightarrow X$ be a continuous transformation of the compact metric space X . Then there exists a probability measure on X that is invariant under T .*

It is not very hard to see that (under the assumptions of this theorem), a measure μ is invariant precisely if for all continuous functions f on X

$$(6.38) \quad \int f \circ T d\mu = \int f d\mu.$$

This provides a useful criterion for verifying that a probability measure is invariant under T . In fact, the construction of such a measure, i.e. a *proof* of Theorem 34, is not very difficult:

The space $M(X)$ of probability measures on the compact metric space X is compact w.r.t. the weak*-topology.⁴ T operates on $M(X)$ via $T_*\mu(B) := \mu(T^{-1}B)$. This means that for all continuous functions f

$$(6.39) \quad \int f d(T_*\mu) = \int f \circ T d\mu.$$

One can then find an invariant measure by taking any probability measure ν and putting

$$(6.40) \quad \mu := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} T_*^i \nu$$

(the limit, or at least the limit of some subsequence, exists by the compactness of $M(X)$) and apply the above criterion to verify that μ is indeed invariant. Of course, such an averaging under the iterates of T is a natural method for producing invariant quantities, as we have already seen at other places.

The space $M(X)$ of probability measures is convex because a convex combination of probability measures is again one. The space $M(X, T)$ of T -invariant probability measures is again convex. An important result is

Theorem 35. *A T -invariant measure μ is an extreme point of the convex set $M(X, T)$ of invariant measures precisely if it is ergodic (that is, if T is ergodic w.r.t. μ).*

In particular, if there is more than one invariant probability measure for T , then T is not ergodic w.r.t. all of them, but only w.r.t. the extreme ones. This motivates

Definition 21. *The continuous transformation T of the compact metric space X is called uniquely ergodic if there exists precisely one T -invariant probability measure.*

⁴ By definition, a sequence μ_n of measures weak*-converges to μ if for any continuous function f , $\int f d\mu_n$ converges to $\int f d\mu$.

Let us now consider some examples in order to understand the meaning and significance of invariant measures. It will be useful to make the following general observation.

Theorem 36. *Let $T : X \rightarrow X$ be a continuous transformation of the compact metric space X . Then x is a periodic point of order N , i.e., $T^N(x) = x$, precisely if the Dirac measure averaged over the orbit of x , namely $\frac{1}{N} \sum_{j=0}^{N-1} \delta_{T^j(x)}$, is an invariant probability measure for T .*

In particular, if the transformation T has a globally attracting fixed point x_0 , i.e., if $\lim_{n \rightarrow \infty} T^n(x) = x_0$ for all $x \in X$, then the Dirac measure δ_{x_0} is the unique invariant measure. Thus, such a transformation which asymptotically moves everything into the global attractor is uniquely ergodic, and ergodic w.r.t. the rather trivial Dirac measure supported at that attractor. If T has several fixed points, then all convex combinations of the corresponding Dirac measures are invariant. T is not ergodic, however, w.r.t. non-trivial such convex combinations.

More generally, suppose that, for some probability measure ν , the iterates $T_*^n \nu$ converge to some limit measure μ as $n \rightarrow \infty$. This limit measure μ is then invariant. Whether this μ is distinguished depends on whether it is independent of the choice of ν , or, if that it is not possible, at least for those ν selected from some natural class. Such a class exists if T is a smooth transformation of a differentiable manifold M . In that case, we can look at the absolutely continuous measures, i.e., those that in any local coordinates (x^1, \dots, x^m) can be written as $\rho(x^1, \dots, x^m) dx^1 \dots dx^m$ for some function ρ , called a density. We may then expect that, for any such ν , the limit μ of the iterates exists and is concentrated on the attractors of the dynamics. These attractors may have vanishing measure w.r.t. ν , but may still have an intricate internal structure, as in the examples of the so-called strange attractors. In any case, iterating a density then in a certain sense means that we are considering the dynamical evolution of all the points of M simultaneously, and this clearly is a very valuable conceptual alternative to following individual trajectories.

A somewhat different, but very useful, perspective emerges when we look at time averages of functions. We recall that the Birkhoff Ergodic Theorem 26 says that when μ is an invariant measure for T , then for $f \in L^1(\mu)$, $\frac{1}{n} \sum_{i=0}^{n-1} f(T^i x)$ converges for almost all x and, if T is ergodic, then the limit, called $E_x(f)$, does not depend on x , for μ -almost all x . In different words: If T is ergodic, then for all $f \in L^1(\mu)$

$$(6.41) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i x) = \int f d\mu \text{ for almost all } x.$$

These results hold for μ -almost all x , but we may wish to ask whether they also hold for almost all x w.r.t. to some natural starting measure like the

Lebesgue or a Riemannian volume measure. If this limit $E_x(f)$ exists and is independent of x in some set B of positive (say) Lebesgue measure, then

$$E(f) := E_x(f)$$

defines a non-negative linear operator on the space $C^0(M)$ of continuous functions on M and, by the Riesz representation theorem, it defines a Borel measure μ on M that satisfies (6.41) for all $x \in B$. In other words, by computing time averages of continuous functions on the orbits of the points in B , we obtain an invariant measure.

Definition 22. *A T -invariant probability measure on M is called an SRB (Sinai-Ruelle-Bowen) measure if (6.41) holds for all $f \in C^0(M)$ and all x in some subset B of M of positive Lebesgue (Riemannian volume) measure. The maximal such set B is called the ergodic basin $B(\mu)$ of μ .*

The question then is under which conditions such an SRB measure exists and, if it does, whether it is unique and ergodic. We shall return to that question in 7.2.

In all these cases, the invariant measures constructed through dynamical iterations capture the asymptotic behavior of the dynamical system. It is still not completely clear, however, in what generality this principle holds. In any case, we conclude that invariant measures can represent important invariants for the dynamical system.

Examples:

1. As in 6.2, we consider the rotations of S^1 ,

$$\begin{aligned} T : S^1 &\rightarrow S^1 \\ z &\mapsto a \cdot z \text{ for } a \in S^1. \end{aligned}$$

Clearly, the Lebesgue measure is invariant under T . In order to check whether there are more, we simply apply Theorem 36. Thus, if a is an N th root of unity, we find periodic points and, consequently, T is not uniquely ergodic, and in fact not even ergodic as we observed already in 6.2. If a is not a root of unity, however, then T does not have any periodic points. In fact, in this case, for any x , the orbit $\{T^n(x)\}$ is dense in S^1 . A homeomorphism of a compact metric space with this property is called *minimal*. By a result of Furstenberg, any homeomorphism of S^1 without periodic points is uniquely ergodic.

2. For the same space as in the previous example, we consider

$$\begin{aligned} T : S^1 &\rightarrow S^1 \\ z &\mapsto z^n \text{ for } n \in \mathbb{Z}. \end{aligned}$$

If $n \neq \pm 1$, this is no longer a homeomorphism. However, the Lebesgue measure μ is still invariant since we have $\mu(T^{-1}A) = \mu(A)$ for any measurable set A (but note that since T is not invertible, we do not have $\mu(T(A)) = \mu(A)$ in general).

3. For the Bernoulli shifts on the state space $Y = \{0, 1, \dots, l-1\}$ with a measure given by $p = (p_0, \dots, p_{l-1})$, we have already seen in 6.2 that the product measure on $X = Y^{\mathbb{Z}}$ determined by p , and more generally the irreducible Markov measures, are invariant.
4. For the tent map introduced in 4.2,

$$Z : [0, 1] \rightarrow [0, 1],$$

$$Z(x) = \begin{cases} 2x & \text{for } 0 \leq x \leq \frac{1}{2} \\ 2(1-x) & \text{for } \frac{1}{2} \leq x \leq 1, \end{cases}$$

the Lebesgue measure is again invariant. Since, as observed in 4.3, the logistic map

$$F(x) = 4x(1-x)$$

is conjugate to the tent map via $\sin^2(\frac{\pi x}{2})$, one computes that the logistic map leaves the measure $\frac{dx}{\pi\sqrt{x(1-x)}}$ invariant.

Proofs of the results stated in this section can be found in [24]. We also refer to the survey in [48].

6.5 Stochastic processes

We have already examined the Bernoulli shift T on elements $0, \dots, l-1$ occurring with probabilities p_0, \dots, p_{l-1} . The iterated application of this shift, which is our dynamical system, can be interpreted naturally as the stochastic process where at each step we randomly chooses one of those elements according to its probability and then records the results as a sequence, i.e., an element of $\{0, \dots, l-1\}^{\mathbb{Z}}$. This can obviously be generalized to other (stationary) stochastic processes as we are now going to explain, and this, in fact, was the original source of inspiration for measurable dynamics. We start with an abstract object which we have not yet encountered in this role, but which is conceptually useful, namely a probability space (Ω, \mathcal{B}, p) . We let $\dots, \xi_{-1}, \xi_0, \xi_1, \xi_2, \dots$ be measurable maps from Ω with values in some measurable space (Y, \mathcal{E}) , perhaps $Y = \{0, \dots, l-1\}$ as above, or, perhaps even more importantly, $Y = \mathbb{R}$. (The ξ_i are called random variables in this context, because they depend on $\omega \in \Omega$ which is considered to be randomly selected according to the probability measure p . Y is sometimes called the state space of the process, and its elements are called the states.)

The relevant space for us now is the sequence space $Y^{\mathbb{Z}}$, with its σ -algebra generated by the cylinder sets, i.e. those of the form

$$A^{n_1, \dots, n_k} = \{(\dots, y_{-1}, y_0, y_1, y_2, \dots) : y_{n_j} \in A_j, j = 1, \dots, k\}$$

for some $k \in \mathbb{N}$, some distinct $n_1, \dots, n_k \in \mathbb{Z}$, and some measurable subsets A_j of Y . We consider the map

$$(6.42) \quad \begin{aligned} \Omega &\rightarrow Y^{\mathbb{Z}} \\ \omega &\mapsto \xi(\omega) = (\dots, \xi_{-1}(\omega), \xi_0(\omega), \xi_1(\omega), \dots) \end{aligned}$$

and use this to define a measure on $Y^{\mathbb{Z}}$ via

$$(6.43) \quad \mu(A^{n_1, \dots, n_k}) := p(\omega | \xi_{n_j}(\omega) \in A_{n_j} \text{ for } j = 1, \dots, k)$$

and extension to the above σ -algebra. As before, we have the shift

$$(6.44) \quad \begin{aligned} T : Y^{\mathbb{Z}} &\rightarrow Y^{\mathbb{Z}} \\ T(\xi_n(\omega)) &= \xi_{n+1}(\omega). \end{aligned}$$

This shift leaves our measure invariant if and only if the stochastic process is stationary in the sense that

$$(6.45) \quad p(\omega | \xi_{n_j}(\omega) \in A_{n_j} \text{ for } j = 1, \dots, k) = p(\omega | \xi_{n_j+m}(\omega) \in A_{n_j} \text{ for } j = 1, \dots, k)$$

for every $m \in \mathbb{Z}$.

Obviously, an analogous construction is possible for continuous time stochastic processes, i.e., families $\xi_t : \Omega \rightarrow Y, t \in \mathbb{R}$ or \mathbb{R}^+ . In fact, whenever we have a measure preserving dynamical system $T_t : \Omega \rightarrow \Omega$, with $t \in \mathbb{R}$ (or $\mathbb{R}^+, \mathbb{Z}, \mathbb{N}$), and a measurable map $g : \Omega \rightarrow Y$, we put

$$(6.46) \quad \xi_t(\omega) := g(T_t(\omega)).$$

We then obtain a new space M , analogous to the space $Y^{\mathbb{Z}}$ above, namely the space consisting of the functions

$$(6.47) \quad t \mapsto \xi(t, \omega) := \xi_t(\omega),$$

for all the $\omega \in \Omega$. Our dynamical system T_t operating on Ω then naturally induces one on M :

$$(6.48) \quad T_t(x)(s) := x(t+s)$$

for a function $x \in M$, i.e., the time shift. This provides a natural realization of the random process $T_t : \Omega \rightarrow \Omega$, and it is often natural to identify Ω with M , as well as the corresponding measures. We have constructed the measure on $Y^{\mathbb{Z}}$, and we can analogously construct the measure on M by looking at the measures on cylinder sets, the so-called finite-dimensional distributions. It is a fundamental result of Kolmogorov that conversely, given such a family of probability measures satisfying some obvious compatibility conditions, there exists an underlying stochastic process whose finite-dimensional distributions are these measures:

Theorem 37. (Kolmogorov extension theorem) For all $t_1 < t_2 < \dots < t_k, k \in \mathbb{N}$, let p_{t_1, \dots, t_k} be probability measures on a complete separable metric space Y (for example \mathbb{R}^d) (equipped with its σ -algebra) that satisfy

$$(6.49) \quad \begin{aligned} & p_{t_1, \dots, t_{j-1}, t_{j+1}, \dots, t_k}(B_1 \times \dots \times B_{j-1} \times B_{j+1} \times \dots \times B_k) \\ &= p_{t_1, \dots, t_{j-1}, t_j, t_{j+1}, \dots, t_k}(B_1 \times \dots \times B_{j-1} \times Y \times B_{j+1} \times \dots \times B_k) \end{aligned}$$

for all times $t_i, k \in \mathbb{N}$, and all Borel subsets B_1, \dots, B_k of Y . Then there exist a probability space (Ω, \mathcal{B}, p) and a stochastic process $\xi_t : \Omega \rightarrow Y$ with

$$(6.50) \quad p_{t_1, \dots, t_k}(B_1 \times \dots \times B_k) = p(\xi_{t_1} \in B_1, \dots, \xi_{t_k} \in B_k)$$

for all t_i, B_i, k .

We now wish to explain briefly the connections with stochastic differential equations, without giving detailed proofs however. We shall develop the general framework first, and only later come to the applications to stochastic differential equations. While this carries the risk of the unmotivated appearance of general constructions long before their purpose becomes clear, it offers the advantage of a quick and logical development. So, we consider the space Ω of all continuous functions $x : [0, 1] \rightarrow \mathbb{R}$ (thus we have already performed the identification of our abstract space Ω introduced above with its natural realization, as explained) with the normalization $x(0) = 0$ (which, however, is only notationally convenient, but in no way logically essential as all our constructions will be naturally invariant under translations of \mathbb{R} ; nor is it essential that our paths are defined on the unit interval – we could as well define them for example on the positive reals). We wish to define the so-called **Wiener measure** on Ω . For that purpose, according to the general scheme described above, we first consider the cylinder sets of the paths x passing through the set $B_i \subset \mathbb{R}$ at time t_i :

$$(6.51) \quad C(t_1, \dots, t_k; B_1, \dots, B_k) := \{x \in \Omega : x(t_i) \in B_i \text{ for } i = 1, \dots, k\}$$

where $0 < t_1 < \dots < t_k$ and B_1, \dots, B_k are Borel subsets of \mathbb{R} ($k \in \mathbb{N}$). Next, we define the measure of a cylinder set as the probability that a path of normalized **Brownian motion** passes through it: We put

$$(6.52) \quad g(t, y, z) := \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{(y-z)^2}{2t}\right),$$

the probability density for a Brownian motion starting at time 0 at y to pass through z at time t .

Remark: The normalization implicit in this definition is that the stochastic process $(w(t))_{t \geq 0}$ with values in \mathbb{R} for which the random variable $w(t) - w(s)$ has the density $g(t-s, 0, z)$ has variance $E((w(t) - w(s))^2) = t - s$. If, in place of g , we use the density $\frac{1}{\sqrt{2\pi\sigma^2 t}} \exp(-\frac{(y-z)^2}{2\sigma^2 t})$, the variance will be $\sigma^2(t-s)$ instead.

Putting $t_0 = 0$ and $x_0 = 0$ for notational convenience, the measure of a cylinder set then becomes

$$\begin{aligned}
 (6.53) \quad & \mu(C(t_1, \dots, t_k; B_1, \dots, B_k)) \\
 &= \int_{B_1} \dots \int_{B_k} g(t_1 - t_0, x_1, x_0) \dots g(t_k - t_{k-1}, x_k, x_{k-1}) dx_1 \dots dx_k \\
 &= \frac{1}{\sqrt{(2\pi)^k (t_1 - t_0) \dots (t_k - t_{k-1})}} \int_{B_1} \dots \int_{B_k} \exp\left(-\frac{1}{2} \sum_{i=1}^k \frac{(x_i - x_{i-1})^2}{t_i - t_{i-1}}\right) dx_1 \dots dx_k.
 \end{aligned}$$

Wiener then showed that this measure can be extended to the smallest σ -algebra of our path space Ω that contains all the above cylinder sets for all $k \in \mathbb{N}$, and the result is called the **Wiener measure**.

To repeat, Brownian motion or, equivalently, the one-dimensional (normalized) Wiener process is defined as the stochastic process $(w(t))_{t \geq 0}$ on Ω with the above σ -algebra and equipped with the Wiener measure, with values in \mathbb{R} with $w(0) = 0$ and for which the random variable $w(t) - w(s)$ has the density $g(t - s, 0, z)$. A sample path, i.e., an element of Ω , then is continuous by definition, but with probability 1 is almost everywhere not differentiable. It may then seem strange that one defines **white noise** $\eta(t)$ as the derivative $\frac{dw(t)}{dt}$ of Brownian motion. Of course, this requires some explanation, but the result is that $\eta(t)$ is distributed according to $\frac{1}{\sqrt{2\pi}} \exp(-\frac{x^2}{2})$ and that therefore its mean or expectation value is 0 for every t , and the covariance, i.e. the expectation of the product $\eta(t_1)\eta(t_2)$ vanishes for $t_1 \neq t_2$. In this sense, the random variables $\eta(t)$ for $t \in [0, 1]$ are independent and identically distributed. Of course, one can also, and more simply, define white noise directly as a family of random variables that is independent and identically distributed, with a Gaussian distribution with mean 0. The preceding will, however, help us in analyzing stochastic differential equations later, i.e., in equations where the derivative of the unknown function is given in terms of white noise. Since white noise itself is defined as a derivative, we can then set up a scheme for integrating such equations.

The preceding can easily be generalized to produce d -dimensional Brownian motion or, equivalently, the d -dimensional (normalized) Wiener process, as the vector-valued process $w(t) = (w_1(t), \dots, w_d(t))$ whose components are one-dimensional independent Wiener processes as just defined. Correspondingly, we then also obtain d -dimensional white noise, again denoted by $\eta(t)$.

With these preparations, we can now discuss the system of stochastic differential equations

$$(6.54) \quad \frac{dx}{dt} = f(x) + s(x)\eta$$

$$(6.55) \quad x(0) = x_0$$

where x and f take their values in \mathbb{R}^d , η is d -dimensional white noise, and s stands for a $d \times d$ -matrix (s_{ij}) . Since we have defined white noise as the

derivative of the Wiener process, it is natural to attempt to represent a solution of (6.54) as

$$(6.56) \quad x(t) = \int_0^t f(x(\tau))d\tau + \int_0^t s(x(\tau))dw(\tau) + x_0.$$

A solution of our stochastic differential equation has to satisfy this relation with probability 1 for every $t \geq 0$. (One also requires the technical condition that the stochastic process $x(t)$, besides being continuous, is measurable w.r.t. the smallest σ -algebra \mathcal{B}_t for which x_0 and $w(\tau)$, $0 \leq \tau \leq t$, are measurable.) In order to make (6.56) meaningful, one has to define the so-called **stochastic integral**

$$(6.57) \quad \int_a^b \xi(\tau)dw(\tau)$$

for a stochastic process $\xi(t)$. The idea is to define this integral as a limit as $n \rightarrow \infty$ of

$$(6.58) \quad \sum_{i=1}^n \xi(\bar{t}_i)(w(t_i) - w(t_{i-1}))$$

for $a = t_0 < t_1 < \dots < t_n = b$ and intermediate points $\bar{t}_i \in [t_{i-1}, t_i]$, under the natural condition that $\max_i(t_i - t_{i-1}) \rightarrow 0$ as $n \rightarrow \infty$. Since $w(t)$ is almost surely not differentiable, however, in contrast to the situation for the Riemann integral, the result will depend strongly on the choices of the intermediate points \bar{t}_i . In fact, two possibilities for systematic choices turn out to be useful: Itô chooses $\bar{t}_i = t_{i-1}$ while Stratonovich takes $\bar{t}_i = \frac{1}{2}(t_i + t_{i-1})$. We shall use here the **Itô integral**.

Returning to our system of stochastic differential equations (6.54), it can be shown that the Itô integral in (6.56) exists under the assumption that f and s satisfy uniform Lipschitz conditions, as in the Picard-Lindelöf theorem for (non-stochastic) ordinary differential equations.

If $f = 0$ and $s(x)$ is a constant matrix, then the solution becomes

$$(6.59) \quad x(t) = s \int_0^t dw(\tau) + x_0 = sw(t) + x_0,$$

that is, a shifted version of Brownian motion, because the Itô integral satisfies $\int_0^t dw(\tau) = w(t)$ (however, we have $\int_0^t w(\tau)dw(\tau) = \frac{1}{2}w(t)^2 - \frac{1}{2}t^2$ so that the Itô integral does not always behave as the standard Riemann integral).

More generally, for a constant vector b and a constant matrix s , the Langevin equation

$$(6.60) \quad \frac{dx}{dt} = -bx + s\eta, \quad x(0) = x_0$$

has the solution

$$(6.61) \quad x(t) = x_0 e^{-bt} + s \int_0^t e^{-b(t-\tau)} dw(\tau)$$

and, since for each τ the expectation value of the white noise $\eta(\tau)$ vanishes, we obtain for the expectation value of our solution

$$(6.62) \quad E(x(t)) = e^{-bt} E(x_0).$$

It is also illuminating to consider the probability density $h(t, x)$ for $x(t)$. This means that for a measurable subset B of \mathbb{R}^d , the probability that $x(t)$ is contained in B is given by

$$(6.63) \quad \int_B h(t, y) dy.$$

It turns out that under certain regularity assumptions on f and s (one might require – although this can be considerably weakened – that the components s_{ij} of s be twice and the components f_i of f be once continuously differentiable), this density h satisfies the so-called **Fokker-Planck** (or **Kolmogorov**) equation

$$(6.64) \quad \frac{\partial h}{\partial t} = \frac{1}{2} \sum_{i,j=1}^d \frac{\partial^2}{\partial x^i \partial x^j} (a_{ij} h) - \sum_{i=1}^d \frac{\partial}{\partial x^i} (f_i h),$$

with

$$a_{ij}(x) = \sum_{l=1}^d s_{il}(x) s_{jl}(x),$$

for $t > 0$ and $x \in \mathbb{R}^d$. A proof of this result can be found in [22], for example, where also the relation between Brownian motion and harmonic functions, i.e., the situation for $f = 0$ and s being the identity matrix, is explained from the point of view of partial differential equations. (6.64) is a partial differential equation, and so, by considering the density of the stochastic process defined by our stochastic differential equation, we can in fact avoid stochastic integrals.

The preceding is presented in more detail in [26]. The probabilistic aspects are developed in [12]. A textbook reference for stochastic differential equations is [36]. As the purpose of this section was only a heuristic overview, and not a precise formulation of technical assumptions or detailed proofs, the reader is advised to consult these or other references for such details.

While in previous sections we have emphasized the analogies between topological and measure preserving dynamical systems, we here see a combination of them where each class has its own distinct role. Namely, we have a topological or, in the present case, an even smooth, dynamical system, namely the

flow induced by a system of ordinary differential equations on which some stochastic process is superimposed as a perturbation, for example caused by external noise. In this spirit, L. Arnold [4] develops a more general framework of random dynamical systems. There, one has a family of transformations of some space X depending on some variable ω in a probability space (Ω, \mathcal{B}, P) as before, and each time $n \in \mathbb{Z}$ or \mathbb{N} (or $t \in \mathbb{R}$ or \mathbb{R}^+ , but for the sake of variety, we here discuss the time discrete case), one of those transformations is chosen in such a manner that the natural composition property is satisfied. This means that we have transformations $T_n : \Omega \rightarrow \Omega$ (for example time shifts for our stochastic processes in case Ω is a path space as in the preceding discussion) with $T_0 = \text{id}_\Omega$ and the group property

$$(6.65) \quad T_{m+n} = T_m \circ T_n$$

and transformations

$$(6.66) \quad \begin{aligned} \phi_n : \Omega \times X &\rightarrow X \\ (\omega, x) &\mapsto \phi_n(\omega, x). \end{aligned}$$

The ϕ_n are measurable in the natural sense (as in the constructions above) if X is a measurable space, and, if X also carries a topological structure, continuous w.r.t. $x \in X$. They satisfy the cocycle property

$$(6.67) \quad \begin{aligned} \phi_0(\omega, \cdot) &= \text{id}_X \text{ for all } \omega \in \Omega \\ \phi_{m+n}(\omega, \cdot) &= \phi_m(T_n\omega, \cdot) \circ \phi_n(\omega, \cdot). \end{aligned}$$

Thus, at time n , ω is shifted to $T_n(\omega)$, and x in the fiber $\{\omega\} \times X$ in the product $\Omega \times X$ is moved to $\phi_n(\omega, x)$ in the fiber $\{T_n(\omega)\} \times X$.

Within this framework, one can then study random difference equations of the form

$$(6.68) \quad x_{n+1} = f(T_n(\omega), x_n),$$

i.e., difference equations driven by noise as input that itself evolves according to some dynamical rule.

Obviously, an analogous theory can be developed for continuous time $t \in \mathbb{R}$ or \mathbb{R}^+ in place of $n \in \mathbb{Z}$. For $t \in \mathbb{R}$, say, one considers measure preserving maps

$$(6.69) \quad \begin{aligned} \theta(t) : \Omega &\rightarrow \Omega \\ (t, \omega) &\mapsto \theta(t)\omega \quad \text{measurable in } t \\ \text{with } \theta(t+s) &= \theta(t) \circ \theta(s) \text{ for all } t, s \end{aligned}$$

and transformations

$$\begin{aligned}
\phi &: \mathbb{R} \times \Omega \times X \rightarrow X \\
&(t, \omega, x) \mapsto \phi(t, \omega)x \\
&\text{with } (t, \omega) \mapsto \phi(t, \omega)x \quad \text{measurable for all } x \\
&\quad x \mapsto \phi(t, \omega)x \quad \text{continuous for all } t, \omega \\
&\phi(0, \omega) = \text{id}_X \\
(6.70) \quad \phi(t+s, \omega) &= \phi(t, \theta(s)\omega) \circ \phi(s, \omega) \quad P\text{-almost surely for all } s, t.
\end{aligned}$$

One can then treat random differential equations

$$(6.71) \quad \dot{x} = f(\theta(t)\omega, x),$$

where $\theta(t)$ is absolutely continuous w.r.t. t , or, more generally, stochastic differential equations, where this property does not hold (recall that, in the discussion above, the stochastic term was white noise, and so $\theta(t)$ would correspond to the time shift for white noise, or the derivative of the Wiener process). The flows generated by such equations then are random dynamical systems.

Moreover, as systematically explored in [4], to a random dynamical system, one can associate various random objects. For example, a random point would be a random variable $x_0 : \Omega \rightarrow X$ which (in the case of two-sided time $t \in \mathbb{R}$) is invariant if

$$(6.72) \quad \phi(t, \omega)x_0(\omega) = x_0(\theta(t)\omega) \quad P\text{-almost surely for all } t.$$

Likewise, a random measure is a map

$$(6.73) \quad \begin{aligned} \mu &: \Omega \rightarrow \{\text{Borel probability measures on } X\} \\ \omega &\mapsto \mu_\omega \end{aligned}$$

which is invariant for the random dynamical system ϕ if

$$(6.74) \quad \phi(t, \omega)\mu_\omega = \mu_{\theta(t)\omega} \quad P\text{-almost surely for all } t.$$

6.6 Stochastic bifurcations

In order to understand the introduced abstract concepts and their scope better, we now study an example of a stochastic bifurcation – or its absence. We consider

$$(6.75) \quad \dot{x} = \lambda x - x^3 + \epsilon^{1/2}\eta$$

where η is Gaussian white noise of mean 0 and variance 1, that is, $E(\eta(t), \eta'(t)) = \delta(t-t')$, as introduced in 6.5. Thus, we are perturbing (2.21) by white

noise, and we want to understand how this affects the bifurcation pattern described in 2.3. For abbreviation, we shall put

$$f(x) := \lambda x - x^3.$$

We can ask two different questions about the influence of the noise:

1. What happens to the time course of a starting point $x_0 = x(0)$ under different realizations of the noise? A subquestion is: How does this depend on the noise level, i.e., what are the asymptotics for $\epsilon \rightarrow 0$?
2. What happens to different starting points under the same realization of the noise?

For 1: We obtain the stationary probability density

$$(6.76) \quad p(x) = c_N \exp\left(\frac{2}{\epsilon}\left(\lambda \frac{x^2}{2} - \frac{x^4}{4}\right)\right) = c_N \exp\left(\int_0^x \frac{2f(\xi)}{\epsilon} d\xi\right)$$

(where c_N is a generic normalization constant depending on λ and ϵ whose precise value is not important here) as solution of the Fokker-Planck (Kolmogorov) equation (6.64)

$$(6.77) \quad -\frac{d}{dx}((\lambda x - x^3)p) + \frac{\epsilon}{2} \frac{d^2}{dx^2} p = 0.$$

We note that the shape of $p(x)$ from (6.76) changes at $\lambda = 0$ from a one-hump to a two-hump function as evidence of a bifurcation. We have

$$(6.78) \quad \lim_{\epsilon \rightarrow 0} p(x) = \delta(0)$$

for $\lambda \leq 0$. For $\lambda < 0$, we rescale $y = \epsilon^{-1/2}x$ and obtain with $\tilde{p}(y) := p(x)$

$$(6.79) \quad \lim_{\epsilon \rightarrow 0} \tilde{p}(y) = c_N (-\lambda)^{1/2} \exp(\lambda y^2).$$

For $\lambda = 0$, we rescale instead $y = \epsilon^{-1/4}x$ and obtain

$$(6.80) \quad \lim_{\epsilon \rightarrow 0} \tilde{p}(y) = c_N \exp\left(-\frac{y^4}{2}\right).$$

Comparing the different scaling factors for $\lambda < 0$ and $\lambda = 0$, we observe that the marginal stability at $\lambda = 0$ leads to an amplification of fluctuations.

For 2: Here, we can use the treatment of [8, 4]. We shall need the concept of Lyapunov exponents as briefly introduced in 2.4 and more fully discussed in 7.1 below. In fact, we need to average the Lyapunov exponents of the individual trajectories w.r.t. the stationary density p . For determining the Lyapunov exponent, we need to linearize our equation (6.75):

$$(6.81) \quad \dot{v} = f'(x)v$$

to obtain

$$(6.82) \quad v(t) = v(0) \exp\left(\int_0^t f'(x(s)) ds\right).$$

This yields the Lyapunov exponent

$$(6.83) \quad \alpha = \int_{\mathbb{R}} f'(x)p(x)dx = -2 \int_{\mathbb{R}} \frac{f(x)^2}{\epsilon} p(x)dx$$

from (6.76). Since this expression is always negative, we obtain stability in the sense that trajectories move closer together as time t proceeds. Thus, here, we no longer see a bifurcation, the result of [8]. The reason is that all trajectories are exposed to the same noise. The noise correlates the different trajectories.

Following [8, 4], we can also understand this at a more technical level through the relationship between invariant densities as solutions of the Fokker-Planck equation (6.77) and the invariant random measure in the sense of (6.73) associated to our process.⁵ (Here, of course, for our random dynamical system $X = \mathbb{R}$.) Namely, given an invariant density $\rho(x) = p(x)dx$ (arising as a solution of (6.77)),

$$(6.84) \quad \mu_\omega = \lim_{t \rightarrow \infty} \phi(t, \theta(-t)\omega)\rho$$

is an invariant random measure in the sense of (6.73) for ϕ , and conversely, given such an invariant measure,

$$(6.85) \quad \rho = E(\mu), \text{ that is, } \rho(B) = \int_{\Omega} \mu_\omega(B) dP(\omega) \text{ for measurable sets } B,$$

is an invariant density.

In the present example, by ellipticity of the Fokker-Planck equation (6.77), $\rho(x) = p(x)dx$ is unique⁶ and therefore ergodic for our dynamical system (see the discussion in 6.4). This implies that μ_ω is also ergodic and unique. From this, one infers that μ_ω is a random Dirac measure, that is,

$$(6.86) \quad \mu_\omega = \delta_{a(\omega)}$$

⁵ Note that two different concepts of invariance are employed here. To avoid confusion, Arnold[4] then speaks of stationary instead of invariant densities.

⁶ This is a general result in the theory of elliptic partial differential equations, see e.g. [22].

for some random point $a(\omega)$. This means that almost all trajectories subjected to the noise ω asymptotically end up at the point $a(\omega)$, and again, this means that no bifurcation occurs. For the details, we need to refer to §§1.8, 9.2 in [4].⁷

⁷ The idea of the proof that μ_ω has to be a Dirac measure is the following: If not, one would have an invariant interval $[a_1(\omega), a_2(\omega)]$; this follows because the differential equation (6.75) preserves monotonicity, that is, $x_1(0) < x_2(0)$ implies $x_1(t) < x_2(t)$ for all $t \geq 0$ for solutions x_1, x_2 . But then with $x_0(\omega)$ defined as the smallest x with $\mu_\omega((-\infty, x]) \leq 1/2$ and $\mu_\omega([x, \infty)) \leq 1/2$, the sets $A^\pm(\omega) := \{\pm x_0(\omega) \leq \pm x\}$ would be invariant, and by ergodicity, each of them would have full measure, and so, $\mu_\omega(\{x_0(\omega)\}) = \mu(A^+(\omega) \cap A^-(\omega)) = 1$ so that we get a Dirac measure after all.

7 Smooth dynamical systems

7.1 Lyapunov exponents

The entropy of a transformation $F : X \rightarrow X$ measures the asymptotic generation of information by F . The quantity $h(F)$, however, is difficult to determine in general, and not only so because it is an asymptotic quantity, but in particular because a supremum over coverings or partitions of X needs to be evaluated. Therefore it makes sense to look for other quantities that measure local expansion or contraction properties of F , but that are easier to compute than the entropy. Such quantities, among others, are the Lyapunov exponents of F that which, although also defined by a limit process, can be approximated well numerically in concrete examples. Their definition, however, requires an additional structure, namely a differentiable one which is to be preserved by the dynamical system. We thus enter the realm of smooth dynamical systems.

The idea is the following. Let X carry a differentiable structure, and let F be differentiable. For each tangent vector V at some $x \in X$, we consider the iterated application of the derivative dF of F on V :

$$dF^n(x)(V),$$

determine the norm of this expression with the help of a (Riemannian) metric on X and let then n tend to ∞ . It turns out that up to sets of measure 0 (where it remains to discuss what conditions have to be required for the measure here) under the assumption that F is ergodic, i.e. that any F -invariant set is either of full or of vanishing measure, there are at most $d = \dim X$ different possibilities for the corresponding limits

$$(7.1) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \log \|dF^n(x)(V)\|$$

(where $\|\cdot\|$ is the norm coming from the Riemannian metric employed). These limits are called the Lyapunov exponents of F , and they do not depend on the measure nor on the metric employed. A Lyapunov exponent $-\infty$ corresponds essentially to a zero direction, a negative one to an exponentially contracting direction and a positive one to an exponentially expanding direction, while

a vanishing one would indicate the direction of an asymptotically conserved quantity.

We now wish to discuss some precise results, namely various variants of the **multiplicative ergodic theorem of Oseledec**.

Theorem 38. *Let $F : M \rightarrow M$ be a differentiable self-map of the compact differentiable manifold M . On M , we choose some Riemannian metric, i.e. a Euclidean norm $\|\cdot\|_x$ on each tangent space $T_x M$ that depends smoothly on x . Then there exists a Borel set B in M with $FB = B$, and*

$$\mu(B) = 1$$

for all probability measures μ on M that are preserved by F . Moreover, there exists a measurable function

$$k : B \rightarrow \mathbb{N}$$

with $k \circ F = k$ and $k(x) \leq \dim M$ for all $x \in B$, as well as measurable functions $\lambda^i : B \cap \{x : i \leq k(x)\} \rightarrow \mathbb{R} \cup \{-\infty\}$ ($1 \leq i \leq k(x)$), $-\infty \leq \lambda^1 < \lambda^2 < \dots < \lambda^{k(x)} < \infty$, with $\lambda^i \circ F = \lambda^i$. For $x \in B$ there exist linear subspaces $\{0\} = W^0(x) \subset W^1(x) \subset \dots \subset W^{k(x)}(x) = T_x M$ of the tangent space at x with $dF(x)W^i(x) \subset W^i F(x)$.

For $x \in B$, $1 \leq i \leq k(x)$, and $V \in W^i(x) \setminus W^{i-1}(x)$, we have

$$(7.2) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \log \|dF^n(x)(V)\| = \lambda^i(x).$$

The quantities λ^i are called the *Lyapunov exponents* of F . These, as well as B, k and the spaces W^i , do not depend on the choice of the Riemannian metric $\|\cdot\|$. Finally, $m^i(x) := \dim W^i(x) - \dim W^{i-1}(x)$ is called the *multiplicity* of the Lyapunov exponent $\lambda^i(x)$.

If F is a diffeomorphism, then $\lambda^i(x) > -\infty$ and $dF(x)W^i(x) = W^i(Fx)$ for all i and all $x \in B$.

The preceding result covers the topological case, with an additional differentiability assumption. In the measure theoretic case, we have:

Theorem 39. *Let F be a measure preserving transformation of the probability space (X, \mathcal{B}, μ) . Let $A : X \rightarrow \text{Hom}(\mathbb{R}^m, \mathbb{R}^m)$ (the space of linear maps $\mathbb{R}^m \rightarrow \mathbb{R}^m$) be measurable with*

$$\max(0, \log \|A(x)\|) \in L^1(X, \mu),$$

i.e. the positive part of $\log \|A(x)\|$ is integrable. Here $\|\cdot\|$ is an arbitrary norm on \mathbb{R}^m , for example the Euclidean one. Then, as before, there exist a Borel set $B \subset X$ with $\mu(B) = 1, FB \subset B$, a measurable function

$$k : B \rightarrow \mathbb{N}$$

with $k \circ F = k$, measurable real-valued functions

$$\lambda^i : B \cap \{x : i \leq k(x)\} \rightarrow \mathbb{R} \cup \{-\infty\}$$

($1 \leq i \leq k(x)$) with $\lambda^i(Fx) = \lambda^i(x)$ and

$$-\infty \leq \lambda^1(x) < \lambda^2(x) < \dots < \lambda^{k(x)}(x) < \infty,$$

and for $x \in B$ we have linear subspaces

$$\{0\} = W^0(x) \subset W^1(x) \subset \dots \subset W^{k(x)}(x) = \mathbb{R}^m \text{ of } \mathbb{R}^m$$

with $A(x)W^i(x) \subset W^i(Fx)$ for $1 \leq i \leq k(x)$.

For all $x \in B$, $1 \leq i \leq k(x)$, $V \in W^i(x) \setminus W^{i-1}(x)$, they satisfy

$$(7.3) \quad \lim_{n \rightarrow \infty} \frac{1}{2} \log \|A(F^{n-1}x) \cdot \dots \cdot A(Fx) \cdot A(x)V\| = \lambda^i(x).$$

If F is ergodic, i.e., if for every Borel set U , $FU = U$ implies that either $\mu(U) = 1$ or $\mu(U) = 0$, then $k(x)$ and the Lyapunov exponents $\lambda^i(x)$ are constant almost everywhere on X . Stated somewhat imprecisely, the theorem says that, in this case, decompositions into the flags generated by eigenspaces and averaged eigenvalues of A , $e^{\lambda^1} < e^{\lambda^2} < \dots < e^{\lambda^k}$, remain invariant under F , where $\|A^n V\|^{\frac{1}{n}} \rightarrow e^{\lambda^i}$ for $V \in W^i \setminus W^{i-1}$. Here, these flags are generated from the eigenspaces corresponding to the eigenvalues in increasing order. When F is invertible, we also have the corresponding objects coming from F^{-1} , its Lyapunov exponents being the negatives of those of F . Intersecting the corresponding flags, we then also get a decomposition into eigenspaces. In fact, under appropriate smoothness assumptions, these eigenspaces form the tangent spaces to invariant submanifolds of X . For example, the eigenspaces corresponding to negative, positive, or vanishing Lyapunov exponents span the tangent spaces of the stable, unstable, and center manifolds discussed in 2.10.

While the above results show that Lyapunov exponents are generically stable against perturbations of orbits, in general, they depend sensitively on the choice of the underlying measure.

The multiplicative ergodic theorem is a fundamental tool for the theory of random dynamical systems [4]. From that perspective, in fact, it belongs more naturally in Chapter 6 than the present one.

The following result of Ruelle establishes a connection between entropy¹ and Lyapunov exponents.

Theorem 40. *Let $F : M \rightarrow M$ be a diffeomorphism of the compact Riemannian manifold M , and let μ be a F -invariant Borel probability measure on M . Then the measure theoretic entropy of F satisfies*

$$(7.4) \quad h_\mu(F) \leq \int \sum_{i:\lambda^i(x)>0} \lambda^i(x) d\mu(x).$$

Thus, the entropy is controlled by the positive Lyapunov exponents.

Under stronger assumptions (more precisely, let F be of class C^2 and ergodic, and assume that the F -invariant measure is equivalent to the Riemannian one of M), we even have Pesin's identity

$$(7.5) \quad h_\mu(F) = \sum_{i:\lambda^i>0} \lambda^i.$$

(By the preceding, in the ergodic case, the Lyapunov exponents $\lambda^i(x)$ are constant almost everywhere, and these constants are called λ^i .) Thus, the entropy equals the sum of the positive Lyapunov exponents. In more general cases, in place of Pesin's identity, we have a formula of Ledrappier-Young [27].

7.2 Hyperbolicity

Hyperbolicity is a condition of ergodicity that is stable under small perturbations.

Examples:

- Rotations of the unit circle: irrational rotations are ergodic w.r.t. the Lebesgue measure, but one can find arbitrary small perturbations, namely rational ones, that are not. Therefore these transformations will not be classified as hyperbolic.

¹ In the preceding, we have used binary logarithms for defining the entropy, in order to bring out the connections with information theory, but natural logarithms for defining the Lyapunov exponents, in order to conform with the conventions in the literature. Either choice is arbitrary, but they are not compatible with each other. Therefore, for the results to follow, we need to change the logarithm in the definition of the entropy from binary to natural. I hope that this will not be too confusing.

$$SL(2, \mathbb{Z}) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, b, c, d \in \mathbb{Z}; ad - bc = 1 \right\}$$

operates on the Euclidean plane. It preserves the Lebesgue measure because of the determinant condition $ad - bc = 1$. Since its elements have integer entries, it leaves the lattice of points in the plane with integer coordinates invariant. We then also obtain an induced action on the two-dimensional torus T^2 , the quotient of the plane obtained by identifying points whose coordinates differ by integers. Elements of $SL(2, \mathbb{Z})$ are called elliptic if their eigenvalues are complex conjugate (and then of absolute value 1), parabolic if they are ± 1 , and hyperbolic otherwise. A hyperbolic element A of $SL(2, \mathbb{Z})$ thus has two real eigenvalues, with absolute values respectively greater and less than 1. As a transformation of the two-dimensional torus, it thus stretches one direction and shrinks the other. Such a transformation then is ergodic w.r.t. the μ -invariant μ -Lebesgue measure on T^2 . This behavior also is stable under small perturbations within the class of differentiable transformations of the torus (since $SL(2, \mathbb{Z})$ itself is discrete, it is not meaningful to speak of transformations inside this class, and we have to consider the bigger class of differentiable transformations). In this sense, it will be classified as hyperbolic. Thus, transformations of the torus that are hyperbolic in the above sense, will also be hyperbolic in the sense of the theory of dynamical systems.²

We have already seen and analyzed various examples of a behavior of hyperbolic type in sections 2.2, 2.7 and 2.10. The key point was always the assumption that the linearization does not possess any eigenvalues of absolute value 1. If that assumption holds we only need to deal with strictly expanding and strictly contracting directions so that the local behavior is easy to understand. In order to achieve a general definition, we first extend Definition 1 to sequences of linear maps.

Definition 23. *A family of linear maps $A_n : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called hyperbolic if there exists a $\rho < 1$, and decompositions $\mathbb{R}^d = W_n^+ \oplus W_n^-$ with $A_n W_n^\pm = W_{n+1}^\pm$ and*

² The above terminology comes from two-dimensional geometry where one distinguishes between Euclidean (flat), elliptic, and hyperbolic geometries. $SL(2, \mathbb{Z})$ operates as the group of isometries of the hyperbolic plane. The hyperbolic elements of $SL(2, \mathbb{Z})$ are those that operate without – proper or asymptotic – fixed points on the hyperbolic plane. One should note, however, that here we are not considering the action on the hyperbolic plane, but on the two-dimensional torus which is a quotient of the Euclidean plane. Nevertheless, the connection between the concepts of hyperbolicity in the various domains of mathematics also exhibits itself through the fact that the geodesic flow on the hyperbolic plane is hyperbolic in the sense of the theory of dynamical systems. On the other hand, there is no direct connection between hyperbolicity for dynamical systems and hyperbolic partial differential equations.

$$(7.6) \quad \|A_n|_{W_n^-}\| \leq \rho \text{ and } \|A_n|_{W_n^+}\| \leq \rho.$$

Here, $\|\cdot\|$ is the Euclidean norm of \mathbb{R}^d .

The idea now is to say that the behavior of a diffeomorphism $F : M \rightarrow M$ of a smooth manifold is hyperbolic when its operation on the collection of tangent spaces of M always satisfies this hyperbolicity condition. Of course, in order to have norms at the tangent spaces of M , we need to equip M with some Riemannian metric. In fact, a Riemannian metric on a smooth manifold M is nothing but a family of Euclidean norms on the tangent spaces $T_x M$ that depend smoothly on $x \in M$. This metric, however, only plays an auxiliary role, as in the discussion of the Floer homology of gradient flows in 3.2. In the present case, the Riemannian metric allows us to identify all tangent spaces with \mathbb{R}^d . Although this identification is only unique up to an orthogonal transformation, this property is irrelevant for our purposes as it will not affect the norms of linear maps.

Definition 24. *Let Λ be a compact set that is invariant under the diffeomorphism $F : M \rightarrow M$. Λ is called a hyperbolic set for F if, with respect to some Riemannian metric on a neighborhood U of Λ , for every $x \in \Lambda$, the family of differentials*

$$DF(F^n x) : T_{F^n x} M \rightarrow T_{F^{n+1} x} M \quad (n \in \mathbb{Z})$$

is hyperbolic.

It is then natural to formulate

Definition 25. *A diffeomorphism $F : M \rightarrow M$ of class C^k ($k \geq 1$) of the compact differentiable manifold M is called an Anosov diffeomorphism if M is a hyperbolic set for F .*

The essential point of this definition is that so-called homoclinic tangencies are excluded, that is, points where the stable and the unstable manifold of some other point become tangent to each other (the stable manifold of a point p contains those orbits that asymptotically converge to p and therefore represents those directions that are contracted, the unstable one contains those orbits that originate from p and so represents the directions that are expanded, see below for the formal definition). As should become clear below, such homoclinic tangencies generate very complicated dynamical behavior. These diffeomorphisms are named after Anosov because he proved

Theorem 41. *Anosov diffeomorphisms are structurally stable in the class of C^k -diffeomorphisms in the sense of Definition 11.*

The stability proof uses the shadowing results to be described shortly. The concept, however, is a little narrow because the existence of an Anosov diffeomorphism imposes severe topological restrictions on the underlying manifold

M . For example, when $\dim M = 2$, it has to be a torus to carry an Anosov diffeomorphism. Therefore, Smale introduced the more general concept of Axiom A diffeomorphisms. We need the preparatory

Definition 26. $p \in M$ is a non-wandering point for F if for every neighborhood U of p there exist arbitrary large integers n with $F^n(U) \cap U \neq \emptyset$.

We denote the set of non-wandering points of F by $\Omega(F)$. Of course, the periodic points of F are contained in $\Omega(F)$.

Definition 27. The diffeomorphism $F : M \rightarrow M$ satisfies Axiom A if $\Omega(F)$ is hyperbolic for F and the periodic points of F are dense in $\Omega(F)$.

While in two dimensions the periodic points are automatically dense in $\Omega(F)$ when that set is hyperbolic, this is no longer true in higher dimensions.

Axiom A diffeomorphisms obviously include Anosov diffeomorphisms. They also include generic gradient flows. Generic here means that the underlying function f has only isolated and non-degenerate critical points; see 3.2. More precisely, we have a smooth function f on our manifold M which defines a gradient flow


$$(7.7) \quad \begin{aligned} \dot{x}(t) &= -Df(x(t)) \\ x(0) &= x \end{aligned}$$

and we define the diffeomorphism $F : M \rightarrow M$ by

$$(7.8) \quad F(x) := x(1).$$

The fixed points of F then are precisely the critical points of f and, when those are non-degenerate, that is, if the Hessian D^2f does not have 0 as an eigenvalue, then at such a fixed point, all eigendirections of D^2f are either expanding (when the corresponding eigenvalue is negative) or contracting (when the eigenvalue is positive).

Moreover, Axiom A diffeomorphisms include also the horse-shoe map introduced in 3.7 and further studied in 4.3:



- S = square $ABCD$
- S^0 = lower horizontal white rectangle
- S^1 = upper horizontal white rectangle
- S_0 = left vertical white rectangle
= image of S^0
- S_1 = right vertical white rectangle
= image of S^1

$$(7.9) \quad \Lambda := \bigcap_{n \in \mathbb{Z}} F^n(S)$$

is then a hyperbolic set for F . In fact, Λ is contained in the white region in S in our figure, and F stretches the vertical and shrinks the horizontal direction there which then yields our hyperbolic splitting of the tangent space. We have already verified in 4.3 that the periodic points are dense in Λ . Thus, Axiom A is satisfied for the horse-shoe map.

A similar example is given by the solenoid: we consider the solid torus $T := S^1 \times D^2$ where S^1 is the unit circle parametrized by the angle $\theta \in [0, 2\pi]$ and $D^2 = \{z \in \mathbb{C} : |z| \leq 1\}$ is the closed unit disk. For $0 < \mu < \rho < \frac{1}{2\pi}$, we define $f : T \rightarrow T$ by

$$(7.10) \quad f(\theta, z) := (2\theta, \mu z + \rho e^{2\pi i \theta}).$$

Thus, the S^1 -direction is stretched, the D^2 -direction contracted, and the image wrapped twice around the torus T . The restrictions on ρ and μ ensure that F is injective. As for the horse-shoe, the set

$$(7.11) \quad \Lambda := \bigcap_n F^n(T)$$

is hyperbolic for F .

While we have formulated the Hadamard-Perron theorem on the stable and unstable manifolds in 2.10 only for hyperbolic fixed points, it holds for general hyperbolic sets.³ For example, for a hyperbolic element A of $SL(2, \mathbb{Z})$, the eigenvalue smaller than 1 corresponds to the contracting, stable directions, the eigenvalue larger than 1 to the expanding, unstable one. For a general formulation, we let $d(\cdot, \cdot)$ be the distance function induced by the Riemannian metric of M . For $x \in M$, we define

$$(7.12) \quad \begin{aligned} W^s(x) &:= \{y \in M : d(F^n x, F^n y) \rightarrow 0 \text{ as } n \rightarrow \infty\} \\ W^u(x) &:= \{y \in M : d(F^n x, F^n y) \rightarrow 0 \text{ as } n \rightarrow -\infty\}. \end{aligned}$$

When $\Omega(F)$ is hyperbolic, then for all $x \in \Omega(F)$, $W^s(x)$ and $W^u(x)$ are embedded submanifolds of M of complementary dimension, called the stable and the unstable manifold of x .


³ This fact can also be viewed from the following perspective. A hyperbolic set Λ for a diffeomorphism $F : M \rightarrow M$ corresponds to a hyperbolic fixed point for an induced action on some Banach manifold, namely the space $C^0(\Lambda, M)$ of continuous maps from Λ to M . The induced action $\Phi_F : C^0(\Lambda, M) \rightarrow C^0(\Lambda, M)$ is simply given by $\Phi_F(g) := FgF^{-1}$. The inclusion map $i_\Lambda : \Lambda \rightarrow M$ is then a fixed point of Φ_F , and it is hyperbolic precisely if Λ is a hyperbolic set for F .

Definition 28. *An Axiom A diffeomorphism F satisfies Smale's transversality condition if, for all $x, y \in \Omega(F)$, $W^s(x)$ and $W^u(y)$ meet transversally, that is, at a non-zero angle (equivalently, the tangent spaces $T_z W^s(x)$ and $T_z W^u(y)$ span the tangent space $T_z M$ at every point $z \in W^s(x) \cap W^u(y)$).*

We mention in passing that condition 2) in the definition of Floer homology in 3.2 is also a transversality condition in the sense of Smale's definition.

In this context, we also have Smale's homoclinic theorem and its variants ([32], [13], [40]).

Theorem 42. *Let p be a hyperbolic fixed point for the diffeomorphism F , and suppose that the stable manifold $W^s(p)$ and the unstable one $W^u(p)$ intersect transversally at some other point q . (Such a q is called a transverse homoclinic point.) Then for some $n \in \mathbb{N}$, the iterate F^n has a hyperbolic*



invariant set Λ containing p and q on which it is conjugate to the two-sided Bernoulli shift on two symbols.

Essentially, one finds a horse-shoe type structure for some iterate F^n ; we have already seen in 4.3 how the dynamics on the hyperbolic invariant set for such a horse-shoe is conjugated to a Bernoulli shift. This result again demonstrates the structural presence of horse-shoes. Also, this often suggests that we define a strange attractor by the property that it contains a transverse homoclinic orbit (see e.g. [13]). This should be compared with the discussion in 2.4. Here, we can amplify that discussion. Since both the stable and the unstable manifold of p are invariant under the dynamical iterates, all the iterates of a transverse homoclinic point q are again such points. Being on $W^s(p)$ its iterates will exponentially converge to p , but being also on $W^u(p)$, the unstable manifold must make a fold between successive such intersections. Since the unstable manifold cannot cross itself, it has to become infinitely folded and complicated when approaching p on the orbit of the transverse homoclinic point q .

A fundamental result of Mañé settles the issue of structural stability.

Theorem 43. *A diffeomorphism $F : M \rightarrow M$ is C^1 -structurally stable iff it satisfies Smale's Axiom A and transversality condition.*

In this context, we should also mention the important Kupka-Smale theorem which essentially states that those diffeomorphisms whose periodic points are all hyperbolic, and with the stable and unstable manifolds of the periodic points always intersecting transversally, are dense and generic in the class of all diffeomorphisms.

The fundamental tool of Anosov's approach to structural stability is the approximation of pseudo-orbits by actual orbits, called shadowing. Here, for a map $F : M \rightarrow M$ and some $\epsilon > 0$, a finite or infinite sequence of points x_n in M is called an ϵ -orbit or an ϵ -pseudo-orbit if

$$(7.13) \quad d(x_{n+1}, F(x_n)) < \epsilon$$

for all successive points x_n, x_{n+1} in the sequence. Such a pseudo-orbit is δ -shadowed by the orbit of a point x if

$$(7.14) \quad d(x_n, F^n(x)) < \delta$$

for all n .

Theorem 44. *Let Λ be a compact hyperbolic set for the diffeomorphism $F : M \rightarrow M$. Then there exists an open neighborhood U of Λ with the property that for all $\delta > 0$ we can find $\epsilon > 0$ such that every ϵ -orbit in U is δ -shadowed by an orbit of F .*

In particular, every ϵ -pseudo-orbit that closes up with error at most ϵ (that is, we have a sequence $x_0, x_1, \dots, x_m = x_0$ with $d(x_{n+1}, F(x_n)) < \epsilon$ for $n = 0, \dots, m-1$) is δ -shadowed by a periodic orbit of F , that is, we find $x \in U$ with $F^m(x) = x$ and $d(x_n, F^n(x)) < \delta$ for $n = 0, \dots, m$. Thus, we have a tool for obtaining periodic orbits of F based on the local linear condition of hyperbolicity together with some kind of global non-linear recurrence that assures the existence of pseudo-orbits that almost return to their starting points.

One may then ask under what conditions the shadowing orbit of a pseudo-orbit in Λ is also contained in Λ . This is the case if Λ is locally maximal in the sense that there exists a neighborhood U of Λ with

$$(7.15) \quad \Lambda = \bigcap_{n \in \mathbb{Z}} F^n(U).$$

This in turn is equivalent to Λ having a local product structure, in the sense that for every $x, y \in \Lambda$ the unique local intersection point of $W^u(x)$ and

$W^s(y)$ is contained in Λ .⁴ Also, the hyperbolic set Λ is an attractor for F precisely if, for every $x \in \Lambda$, its unstable manifold $W^u(x)$ is also entirely contained in Λ . The dynamics inside such a hyperbolic attractor may then still be chaotic because of the presence of those unstable directions.

Another useful tool for deriving important properties of hyperbolic sets is the Perron-Frobenius or transfer operator on functions $\phi : M \rightarrow \mathbb{R}$, defined by

$$(7.16) \quad L\phi(x) := \sum_{y: F(y)=x} \frac{\phi(y)}{\det DF(y)}.$$

The formula for the change of variables in integrals yields

$$(7.17) \quad \int L\phi(x)\psi(x)dm(x) = \int \phi(y)\psi(f(y))dm(y)$$

for all appropriate functions ϕ, ψ , where dm is the volume measure on M w.r.t. some Riemannian metric (or the Lebesgue measure in some local coordinates). This measure is the background measure for the SRB-measures that we hope to find. The Perron-Frobenius operator is particularly nice in the case of an expanding map F ; in that case, the determinant in the denominator is always bigger than 1, and so the Perron-Frobenius operator achieves a smoothing of functions. In the situation of an expanding map F , there is a unique SRB measure, and that measure is absolutely continuous w.r.t. the background measure m . The Perron-Frobenius operator is very useful for finding it because such measures correspond to the fixed points of L . Namely, if ϕ_0 is a non-negative fixed point of L , then $\frac{\phi_0 m}{\int \phi_0 m}$ is an F -invariant probability measure that is obviously absolutely continuous w.r.t. m , and conversely, for such a measure m_0 , the function $\frac{dm_0}{dm}$ is a fixed point of L . Thus, the strategy for finding an SRB measure is to show the existence of a fixed point of L in $L^1(M)$.

For a general hyperbolic attractor Λ for a diffeomorphism as above, we can still find a unique SRB measure m_0 supported on Λ . m_0 is ergodic, but in general no longer absolutely continuous w.r.t. the background measure m . This comes from the presence of the stable manifolds along which the measure m will be expanded by the operation of the diffeomorphism F . m_0 can be represented by disintegration along the leaves of the unstable foliation. That is, let some local neighborhood U be covered by $M^u := \bigcup_{\alpha \in A} M^\alpha$ where the M^α are the intersections of the unstable manifolds with U . Then

$$(7.18) \quad m_0(\cdot) = \int_A \mu_\alpha(\cdot) d\rho(\alpha)$$

⁴ There may be further intersection points, but they may occur only at points z for which either $F^n(z)$ does not stay close to x for all $n \leq 0$ or $F^n(z)$ does not stay close to y for all $n \geq 0$. This is meant to be excluded by the qualification “local”.

where ρ is some measure on A and the conditional probability measures μ_α are supported on S_α and absolutely continuous w.r.t. the Lebesgue measure on S_α . While A can be considered as a piece of a stable manifold, the measure ρ is singular w.r.t. to the Lebesgue measure there.

A result of Ledrappier-Young [27] says that an ergodic measure μ with compact support for a C^2 diffeomorphism is an SRB measure iff it satisfies Pesin's identity (7.5).

Similar constructions and results also hold for flows instead of iterates of diffeomorphisms. Of course, for flows, there is always one exceptional direction, namely the one in the direction of the flow lines, and we then assume that the directions transverse to the flow lines are stable or unstable.

An excellent reference for the theory of hyperbolic dynamical systems is [24]. The geometric aspects are well explained in [37].

For those readers who are familiar with Riemannian geometry (see e.g. [21]), geodesic flows on negatively curved Riemannian manifolds are a prime example.⁵ The following metric behavior is characteristic for such manifolds: when we have two pairs (p_1, q_1) and (p_2, q_2) of points, and homotopic geodesics $\gamma_i(t)$ from $p_i (= \gamma_i(0))$ to $q_i (= \gamma_i(1))$ ($t \in [0, 1]$, $i = 1, 2$), then the distance $d(\gamma_1(t), \gamma_2(t))$ is a strictly convex function of t , and the more negative the curvature is, the closer the interior points, say $\gamma_1(\frac{1}{2})$ and $\gamma_2(\frac{1}{2})$, come together. The explanation for this behavior is that, in the negative curvature case, geodesics starting from the same point in different directions are moving apart exponentially (instead of linearly as in the Euclidean case), and so, when we have two geodesics with initial and final points at some bounded distance, then in between, these geodesics have to come together more closely to compensate for that exponential divergence. This phenomenon has also been called the turnpike theorem, because the fastest route from city A_1 to city B_1 may utilize the same highway as the one from A_2 to B_2 . Thus, the optimal connections not only come together more closely than their endpoints, but even partially coincide.

While we have seen that hyperbolicity is both a general and a structurally stable condition, it is nevertheless one of the key challenges for modern research on dynamical systems to develop a general theory under weaker assumptions. For some important progress, see [56].

⁵ In fact, manifolds of constant negative curvature are called hyperbolic, and so this matches with the use of the term "hyperbolic" in the theory of dynamical systems.

7.3 Information loss

We have defined the (measure theoretic) entropy of a dynamical system given by a transformation $T : X \rightarrow X$ as the asymptotic information generated by observations of the orbit of a point. This was done with respect to some measure μ left invariant by T . On the other hand, in 6.4 and 7.2, we have also seen that, and studied how, T operates on the space of measures on X . In general, measures will get asymptotically simplified under the operation of T – in the extreme case of a global point attractor, the iterates of some measure under T will converge to a Dirac measure, the most trivial measure conceivable – and in this sense potential information contained in X , that is, in the distribution of initial values, will get destroyed. We shall now quantify that information loss, following the work of Ruelle [43, 44] (although providing a somewhat different interpretation⁶). Thus, we do not investigate here a Kolmogorov-Sinai type entropy which evaluates the information content of differences of states caused by the dynamics, but rather the difference of information content of the state distribution when subjected to the dynamical iteration. There is no direct relationship between these two strategies.

We consider a diffeomorphism $F : M \rightarrow M$ of a compact manifold. Let $\rho(x)dx$ be some probability density on M , for example the normalized density coming from the volume form of some Riemannian metric on M . With ρ , we associate the entropy

$$(7.19) \quad H(\rho) = - \int \log \rho(x) \rho(x) dx.$$

The image of ρ under F , the push-forward measure $F_*\rho$, is characterized by the property that

$$(7.20) \quad \int \phi(y) F_*\rho(y) dy = \int \phi(F(x)) \rho(x) dx$$

for all integrable functions ϕ and has the density

$$(7.21) \quad F_*\rho(y) = \frac{\rho(F^{-1}y)}{J(F^{-1}y)}$$

where $J(x) = |\det(Fx)|$ is the absolute value of the Jacobian of F . Now

⁶ In fact, we shall here call an entropy loss what Ruelle calls entropy production. That difference results because we are interested in the intrinsic aspects of the dynamical iteration whereas Ruelle is concerned about the exchange with the environment in the context of the second law of thermodynamics. Thus, the entropy reduced within the system must be exported to the environment and increase the latter's entropy.

$$\begin{aligned}
H(F_\star\rho) &= - \int \log F_\star\rho(y)F_\star\rho(y)dy \\
&= - \int (\log \rho(F^{-1}y) - \log J(F^{-1}y)) \frac{\rho(F^{-1}y)}{J(F^{-1}y)} dy \\
(7.22) \quad &= - \int (\log \rho(x) - \log J(x))\rho(x)dx.
\end{aligned}$$

Therefore, the entropy or information difference resulting from an application of F to the distribution of initial values according to the density ρ is

$$(7.23) \quad H(\rho) - H(F_\star\rho) = - \int \log J(x)\rho(x)dx.$$

This expression need not be 0. This simply expresses the fact that the entropy (7.19) is not invariant under coordinate transformations; namely, any diffeomorphism can be considered as a coordinate transformation. Of course, when we consider the inverse diffeomorphism F^{-1} in place of F , (7.23) changes its sign because the Jacobian of F^{-1} is the inverse of that of F .

Similarly, that difference at the m th iteration of F is

$$(7.24) \quad - \int \log J(x)F_\star^{m-1}\rho(x)dx$$

and, if the iterates $F_\star^m\rho$ weak \star converge to some invariant probability measure μ , then that information difference converges to

$$(7.25) \quad h_F(\mu) := - \int \log J(x)\mu(dx).$$

We assume that μ is ergodic. Then the Lyapunov exponents λ^i of F w.r.t μ are constant μ -almost everywhere on M , and the Oseledec Multiplicative Ergodic Theorem 38 implies

Lemma 5. $h_F(\mu) = - \sum \lambda^i$.

Moreover, we have the following result of Ruelle.

Theorem 45. *If μ is an SRB measure, then $h_F(\mu) \geq 0$. If F is of class C^2 and μ is singular w.r.t. the volume form dx and all the Lyapunov exponents $\lambda^i \neq 0$, then $h_F(\mu) > 0$.*

Thus, asymptotically, an information loss, but no information gain, can occur.

Proof.

$$\begin{aligned}
h_F(\mu) &= - \sum \lambda^i \\
&= (h_\mu(F) - \sum_{\lambda^j > 0} \lambda^j) - (h_\mu(F) + \sum_{\lambda^k < 0} \lambda^k) \\
&= (h_\mu(F) - \sum_{\lambda^j > 0} \lambda^j) - (h_\mu(F) - \sum_{\sigma^j > 0} \sigma^j)
\end{aligned}$$

where σ^i are the Lyapunov exponents of the inverse map F^{-1} . Now by Pesin's identity (7.5), the first term in the last line vanishes whereas, by Theorem 40, the second one is non-negative. This implies that $h_F(\mu) \geq 0$. Now if $h_F(\mu) = 0$, the preceding implies that

$$(7.26) \quad h_\mu(F) = \sum_{\lambda^j > 0} \lambda^j = - \sum_{\lambda^k < 0} \lambda^k$$

Thus, by the results mentioned at the end of 7.1, μ is an SRB measure for both F and F^{-1} , and in fact, under the conditions of the theorem, it is absolutely continuous, that is, non-singular w.r.t. dx . This completes the proof.

q.e.d.

Corollary 4. *If μ is an SRB measure for F , then the Lyapunov exponents satisfy $\sum \lambda^i \leq 0$.*

The entropy loss describes the asymptotic amount of information destroyed by the operation of the dynamical system at each step; of course when that amount is positive, the entropy of the asymptotic SRB measure μ must be

$$(7.27) \quad H(\mu) = - \int \log \mu(x) \mu(x) dx = -\infty.$$

This expresses that μ is singular w.r.t. the background measure dx when the dynamical system asymptotically at each step loses a positive amount of information about the distribution of initial values as described by dx . As explained in 7.2, the singularity of μ comes from the stable manifolds of the dynamical iteration, that is, those that are asymptotically contracted to single points. This represents an evident source of information loss. Ruelle's result says that this will not be compensated by the expansion of the unstable directions.

There is another source of information loss. To see this, we first consider an iteration T on some finite state space X containing elements x with probabilities $p(x)$, with $x(n) = Tx(n-1) \in X$. Then, with $p_n(y) := \sum_{x \in T^{n-1}X: Tx=y} p_{n-1}(x)$, $p_0(x) := p(x)$, the entropy

$$H(n) := - \sum_{x(n)} p_n(x(n)) \log p_n(x(n)) \leq H(n-1),$$

with strict inequality when T is not injective on the collection of states $x(n-1) \in X$. Thus, T destroys entropy or information contained in the initial distribution $p(x)$ by merging different initial states into single subsequent states.

More generally, we can also quantify the asymptotic information loss when the compact manifold M is covered up to a set of measure 0 by finitely many

disjoint open sets U_α and $F : M \rightarrow M$ maps each such U_α diffeomorphically onto its image, such that two such images FU_α and FU_β either coincide or are disjoint. We obtain an additional term in (7.25) coming from the loss of information caused by the non-injectivity of F . Namely, in the present situation, a measure ν can be expressed in terms of its push-forward $F_\star\nu$ by

$$(7.28) \quad \nu = \int \sigma_y F_\star\nu(dy)$$

where σ_y is a weighted sum of Dirac measures on the (finitely many) preimages of y , of total mass 1. If those weights are p_α , we obtain the folding entropy

$$(7.29) \quad f_F(\mu) := - \int \sum_\alpha p_\alpha \log p_\alpha F_\star\nu(dy).$$

(7.23) now generalizes as

$$(7.30) \quad H(\rho) - H(F_\star\rho) = - \int \log J(x)\rho(x)dx + f_F(\rho),$$

and by showing the upper semicontinuity of such expressions under weak \star convergence, Ruelle [43] concludes that the limiting entropy loss for a measure μ obtained as the weak \star limit of the averages $\frac{1}{m} \sum_{j=0}^{m-1} F_\star^j \rho$ of the iterates under F of a density ρ ,

$$(7.31) \quad - \int \log J(x)\mu(x)dx + f_F(\mu) \geq 0.$$

More generally, one can also show that the entropy loss is non-negative in the context of Axiom A diffeomorphisms.

For a systematic treatment of the topic of entropy production in dynamical systems, we refer to [19].

8 Cellular automata and Boolean networks as examples of discrete dynamical systems

8.1 Cellular automata

We have seen above how to investigate the qualitative properties of a dynamical system defined on a continuum by passing to discrete dynamics. In the case studied, this was a simple shift. We now want to introduce a more general class of discrete dynamical systems that exhibit a richer dynamical behavior, some aspects of which can in turn be studied by passing to dynamics on a continuum. The class in question is that of one-dimensional cellular automata and suitable generalizations thereof. In a certain sense, however, this chapter constitutes an anticlimax as the mathematical results that have been discovered so far about cellular automata and their generalizations are less profound than those described in the previous chapters. Thus, cellular automata represent an embarrassment to the presently known theory of dynamical systems and an important challenge for its future development.

We first consider dynamics on the space of two-sided binary sequences. We thus assign to an index $m \in \mathbb{Z}$ a value $\sigma_m \in \mathbb{Z}_2 = \{0, 1\}$. The dynamics is determined again by a map T that associates each such two-sided binary sequence $\sigma = (\dots, \sigma_{-1}, \sigma_0, \sigma_1, \dots)$ with a new sequence $T\sigma$. For a cellular automaton, we impose the restriction upon T that $(T\sigma)_m$ is determined only by the value of σ_m itself and its r neighbors to the left, namely $\sigma_{m-r}, \sigma_{m-r+1}, \dots, \sigma_{m-1}$, and to the right, $\sigma_{m+1}, \sigma_{m+2}, \dots, \sigma_{m+r}$, for some given $r \in \mathbb{N}$. Mostly, we consider the simplest case $r = 1$ only. This includes the Bernoulli shift, for which

$$(T\sigma)_m = \sigma_{m+1}$$

or, to write it as a cellular automaton rule,

$$(T\sigma)_m = 1 \quad \text{if and only if } \sigma_{m+1} = 1.$$

More generally, a cellular automaton with $r = 1$ is determined by which combinations of values of $(\sigma_{m-1}, \sigma_m, \sigma_{m+1})$ result in $(T\sigma)_m = 1$. For the Bernoulli shift, these combinations are (001), (011), (101), (111), while the other possible combinations, namely (000), (010), (100), (110), lead to $(T\sigma)_m = 0$. It is easily seen that it is no serious restriction to always require that (000) leads

to $(T\sigma)_m = 0$. This reduces the number of possible automaton rules for $r = 1$ from $2^{2^3} = 256$ to 128.

Following Wolfram, the following coding scheme for the rules is employed. One writes the combinations for $(\sigma_{m-1}, \sigma_m, \sigma_{m+1})$ in decreasing order, i.e., (111), (110), (101), (100), (011), (010), (001), (000) and, for each of them, one writes down whether it leads to 0 or to 1. This yields an eight-digit binary number. Since we require that (000) always leads to 0, the last digit is always 0, and so only 7 digits remain. The above Bernoulli shift then corresponds to the rule

$$10101010 = 170 \text{ in the decimal system.}$$

This rule generates a very simple behavior. Other rules lead to richer dynamics, sometimes even on the space of periodic sequences with (sufficiently large) period $N \in \mathbb{N}$ (i.e. $\sigma_{m+N} = \sigma_m$ for all $m \in \mathbb{Z}$). In the periodic case one can, of course, identify the index m with $m+N$ and restrict the consideration to the indices $1, 2, \dots, N$.

Wolfram divided the dynamical behavior of one-dimensional cellular automata with finite N or for initial states with infinitely many 1s into four qualitative classes:

1. After finitely many steps, the evolution leads to a homogeneous state. Thus, the initial pattern disappears.
2. The evolution leads to a simple stable or periodic state. The initial pattern develops in a finite region only.
3. The evolution leads to chaotic behavior. The initial pattern grows into ever increasing regions. Many initial states develop into self-similar structures.
4. The evolution leads to complex, spatially localized patterns that grow and shrink in a regular manner and can propagate in a soliton-like manner.

Some cellular automata of the 4th class are even universal computers in the following sense. A cellular automaton can be considered as a computer that, by its temporal evolution, applies an algorithm to the input data given as its initial configuration. A computer is called universal if by suitable initial data every possible algorithm can be produced. Since a universal computer can simulate a Turing machine, it is capable of arbitrarily complex behavior, and conversely, its calculations can only be found by explicit simulation, but cannot be predicted by a simpler scheme. In particular, as for the halting problem for Turing machines, there does not exist any finite algorithm that can decide for any initial configuration whether it will lead to the homogeneous zero state or produce structures that for arbitrarily large times possess states different from 0. It is thus impossible in principle to predict the behavior of such a cellular automaton. Even if the initial state is completely known, the long time behavior of the system can only be found by explicit

application of the dynamics.

Systematic computer studies, including in particular the determination of the basins of attraction of the various dynamical attractors through tracking the cellular automaton dynamics back in time, can be found in [59].

In order to assess the complexity of cellular automata, one may try various entropy concepts:

1. As explained, a cellular automaton can be seen as a self-map T of the space Y_2 of two-sided binary sequences. Therefore, we can compute the topological entropy of $T : Y_2 \rightarrow Y_2$. In practice, however, this may not be feasible.
2. If T is considered as a self-map of $(\mathbb{Z}_2)^N$, the space of binary sequences of period N , then we have a finite state space Y . Obviously, the topological entropy of every self-map of a finite state space vanishes since the cardinality of any covering is bounded by $2^{|Y|}$, the number of subsets of Y , and so it cannot grow asymptotically.

One may however utilize the following entropy. For every element i of Y , at the n^{th} iteration step, we consider the relative frequency $p_n(i)$ with which i can be reached from initial configurations. For example, if there are k elements i_1, \dots, i_k in Y with $T^n(i_\kappa) = i$, then $p_n(i) = \frac{k}{|Y|}$. Thus, an ansatz for the asymptotic entropy of T could be

$$H_0(T) := \lim_{n \rightarrow \infty} \left(- \sum_{i \in Y} p_n(i) \log_2 p_n(i) \right).$$

(We shall say more about this below when discussing Boolean networks.) If T is invertible (reversible), then the expression

$$- \sum_{i \in Y} p_n(i) \log_2 p_n(i)$$

is independent of n , and is simply $\log_2 |Y|$. If T is irreversible, i.e. if several distinct initial patterns lead to the same subsequent pattern, then the entropy

$$- \sum_{i \in Y} p_n(i) \log_2 p_n(i)$$

decreases as a function of the time n . Thus, decreasing entropy becomes a sign of irreversibility.

The preceding quantity $H_0(T)$ (which also depends on N), while easy to compute, does not yield too many insights into the dynamical behavior of T . It becomes minimal, i.e. 0, if all initial states lead to the same final state after some iterations. When we obtain a periodic cycle of length k instead of the constant final state, again attracting all initial conditions,

then the entropy is $\log_2 k$. Thus, a simple repetitive behavior does not lead to vanishing entropy here, and this fact already casts serious doubts about its utility. The entropy is maximal, i.e. $\log_2 N$, if the dynamics is reversible, that is, if every state can occur at every iteration step. This however fails to distinguish between the simple dynamics of the identity map, where $(Tb)_m = b_m$ for all m , and another dynamics that permutes the states in some complicated manner.

In any case, we are here considering a situation that is analogous to the one in 7.3. Namely, we are again not evaluating the information content of differences of state distributions, but rather comparing the information content of different state distributions. In the first case, the identity transformation yields zero entropy, because it does not lead to any state differences and thus does not bring to the fore any new information that had been hitherto hidden by a limited resolution. In the second case, for the identity or any other injective transformation, the entropy of the dynamical iterates of the initial state distribution remains maximal because the information content encoded in that distribution does not decrease. Here, we decrease the entropy of the state distribution only when the dynamics is non-injective, that is, when different states merge. This can be interpreted as forgetting some initial differences in the course of the dynamical iteration. This is of course an obvious point, namely that forgetting causes an information loss, but it is quite relevant for understanding the relationship between information theory and thermodynamics in the context of the famous second law of thermodynamics. – In any case, the main deficiency of the present approach is that no comparison between individual states and their dynamic successors is made. We rather only compare aggregate quantities, namely entropies of distributions. This disregards basically all information about the specific dynamics.

3. We now return to a Kolmogorov-Sinai type approach. We obtain a more useful quantity than in the preceding approach when we let the length N of the patterns that we evaluate go to infinity (cf. [18]). More precisely, we now consider the cellular automaton operating on infinite binary sequences, and we consider the distribution of patterns of length at most B , so-called blocks. Actually, these could be either spatial patterns as before, that is, the values at a given time of a consecutive string of sites, or alternatively sequences of length $\leq B$ of consecutive values at a fixed site, that is, temporal patterns. In the spatial case, one considers the relative frequencies $p_n(\alpha)$ with which the blocks α of values at consecutive sites appear at time n . Of course, in practice, for automata operating on infinite sequences, one usually does not have an explicit formula for these expressions, but rather needs to estimate them on the basis of observed CA computer simulations. In any case, the relevant (normalized) entropy is then

$$(8.1) \quad H_{spatial}^{nor}(T) := \lim_{B \rightarrow \infty} \lim_{n \rightarrow \infty} -\frac{1}{B} \sum_{\alpha=1}^{2^B} p_n(\alpha) \log p_n(\alpha).$$

For temporal patterns, one can work with strings of finite length, but then needs to estimate the relative frequency $q(\alpha)$ with which a block α appears as the consecutive values at some given site for an infinite running time of the CA. We then obtain an expression that is analogous to (8.1) for the (normalized) temporal entropy of the CA T ,

$$(8.2) \quad H_{temporal}^{nor}(T) := \lim_{B \rightarrow \infty} -\frac{1}{B} \sum_{\alpha=1}^{2^B} q(\alpha) \log q(\alpha).$$

One should note that the factor $1/B$ normalizes these entropies to take values between 0 and 1. The value is 0 when asymptotically only finitely many different patterns occur. It is 1 when all patterns occur with equal probability. The interesting CAs yield intermediate values for these normalized entropies.

We can obtain deeper insights by recalling the discussion of 6.3 of Grassberger's concept of effective measure complexity [11]. We consider the entropy for temporal strings of length B now without a normalization factor,

$$(8.3) \quad H(T, B) := - \sum_{\alpha=1}^{2^B} q(\alpha) \log q(\alpha).$$

This is the information contained in strings $b_i, b_{i+2}, \dots, b_{i+B-1}$ of length B of values at one site at times $i, \dots, i+B$.¹ Given such a string of length B , the additional information needed to predict the value b_{i+B} is then

$$(8.4) \quad h(T, B) := H(T, B+1) - H(T, B).$$

These differences cannot increase with B , because the prediction of the next value becomes easier the more previous values we already know. In fact, by (6.17)

$$(8.5) \quad h(T) := \lim_{B \rightarrow \infty} h(T, B)$$

is the entropy of T .

Then the difference of these differences,

$$(8.6) \quad \delta h(T, B) := h(T, B-1) - h(T, B),$$

measures the average amount by which the uncertainty of b_{i+B} decreases due to the knowledge of b_i , that is, one more site value back in the past.

¹ By definition, this does not depend on i .

We thus need to store at least an amount of $\delta h(T, B)$ for at least B time steps for optimal prediction. We thus have the effective measure complexity of T

$$\begin{aligned}
 h_{EMC}(T) &= \sum_B B \delta h(T, B) = \sum_B (h(T, B - 1) - h(T, B)) \\
 (8.7) \qquad &= \sum_{B=0}^{\infty} (h(T, B) - h(T)).
 \end{aligned}$$

4. We now present an alternative to the preceding approaches which avoids their deficiencies. We embed the discrete dynamics of a cellular automaton into a dynamics on some continuous state space. Once more, we consider a cellular automaton with N indices and two possible states for every index. Let

$$W = [0, 1]^N$$

be the unit cube in \mathbb{R}^N . Each vertex of W codes a binary sequence (b_1, \dots, b_N) of N elements, by letting b_k stand for the k^{th} coordinate of the corresponding corner. Thus, the cellular automaton defines a dynamics on the set of vertices of W . We consider the m^{th} -coordinate x^m of $x = (x^1, \dots, x^N) \in W$. Let $b := (b_{m-r}, \dots, b_m, \dots, b_{m+r})$ be a configuration that leads to

$$(Tb)_m = 1.$$

We consider

$$p_b^m(x) := \prod_{j=-r}^r p_j^m(x)$$

with

$$p_j^m(x) := \begin{cases} x^j & \text{for } b_{m+j} = 1 \\ (1 - x^j) & \text{for } b_{m+j} = 0 \end{cases}$$

and put

$$p^m(x) := \sum p_b^m(x),$$

where the sum is extended over all configurations $b = (b_{m-r}, \dots, b_{m+r})$ that lead to $(Tb)_m = 1$. If, for example, $(b_{m-1}, b_m, b_{m+1}) = (101)$ and (010) are the only states leading to $(Tb)_m = 1$, then

$$p^m(x) = x^{m-1}(1 - x^m)x^{m+1} + (1 - x^{m-1})x^m(1 - x^{m+1}).$$

Thus, $p^m(x)$ is an (inhomogeneous) polynomial of order $2r + 1$ of the components of $x = (x^1, \dots, x^N) \in \mathbb{R}^N$. We have

$$0 \leq p^m(x) \leq 1, \text{ if } 0 \leq x^j \leq 1 \text{ for all } j.$$

Thus, $p = (p^1, \dots, p^N)$ maps the unit cube W into itself. Moreover, by construction, on the vertices of W , p induces the original dynamics of the

cellular automaton T . For p , operating on W , however, we can determine the topological entropy or the Lyapunov exponents. Perhaps this is a tool to evaluate the complexity of the cellular automaton T . Numerical simulations, however, have not so far yielded conclusive results in this direction.

8.2 Boolean networks

Boolean networks are a more general class of discrete dynamical systems than cellular automata. Their potential applications have been emphasized by St.Kauffman[25].

For a Boolean network, each of N elements is updated according to a rule which, in contrast to cellular automata, can vary from element to element. Moreover, each element receives input from K other elements which no longer need to be its neighbors, but can be any among the N elements of the network. Again, each element can assume two states, 0 or 1 and, for this reason, the networks are called Boolean.

Since every element receives input from K others, there are K input combinations for it and, since every such combination can generate either 0 or 1 as output of the element under consideration, there are

$$2^{2^K}$$

Boolean functions, that is, update rules for every element. For $K=2$, we have 16 Boolean functions, namely

0 0 0	0 0 1	0 0 0	0 0 1
1 0 0	1 0 0	1 0 1	1 0 1
0 1 0	0 1 0	0 1 0	0 1 0
1 1 0	1 1 0	1 1 0	1 1 0
0	1	2	3
0 0 0	0 0 1	0 0 0	0 0 1
1 0 0	1 0 0	1 0 1	1 0 1
0 1 1	0 1 1	0 1 1	0 1 1
1 1 0	1 1 0	1 1 0	1 1 0
4	5	6	7
0 0 0	0 0 1	0 0 0	0 0 1
1 0 0	1 0 0	1 0 1	1 0 1
0 1 0	0 1 0	0 1 0	0 1 0
1 1 1	1 1 1	1 1 1	1 1 1
8	9	10	11

$\begin{array}{c} 0\ 0\ 0 \\ 1\ 0\ 0 \\ 0\ 1\ 1 \\ 1\ 1\ 1 \\ \mathbf{12} \end{array}$	$\begin{array}{c} 0\ 0\ 1 \\ 1\ 0\ 0 \\ 0\ 1\ 1 \\ 1\ 1\ 1 \\ \mathbf{13} \end{array}$	$\begin{array}{c} 0\ 0\ 0 \\ 1\ 0\ 1 \\ 0\ 1\ 1 \\ 1\ 1\ 1 \\ \mathbf{14} \end{array}$	$\begin{array}{c} 0\ 0\ 1 \\ 1\ 0\ 1 \\ 0\ 1\ 1 \\ 1\ 1\ 1 \\ \mathbf{15} \end{array}$
--	--	--	--

The rules 0 and 15 generate a constant output regardless of the input, and are hence of little interest. The output of the rules 3, 5, 10, 12 depends only on one of the two inputs, and hence they are not so interesting either.


Here is an example for $N = 3, K = 2$. We first exhibit the Boolean function for every element, that is, the input-output table

$\begin{array}{c} 2\ 3\ 1 \\ 0\ 0\ 1 \\ 1\ 0\ 0 \\ 0\ 1\ 0 \\ 1\ 1\ 0 \end{array}$	$\begin{array}{c} 1\ 3\ 2 \\ 0\ 0\ 1 \\ 1\ 0\ 1 \\ 0\ 1\ 1 \\ 1\ 1\ 0 \end{array}$	$\begin{array}{c} 1\ 2\ 3 \\ 0\ 0\ 1 \\ 1\ 0\ 0 \\ 0\ 1\ 0 \\ 1\ 1\ 0 \end{array}$
--	--	--

This yields the transition map $(\sigma_1, \sigma_2, \sigma_3) \mapsto T(\sigma_1, \sigma_2, \sigma_3)$ according to the following table:

$\begin{array}{c} 1\ 2\ 3 \\ 0\ 0\ 0 \\ 0\ 0\ 1 \\ 0\ 1\ 0 \\ 0\ 1\ 1 \\ 1\ 0\ 0 \\ 1\ 0\ 1 \\ 1\ 1\ 0 \\ 1\ 1\ 1 \end{array}$	\rightarrow	$\begin{array}{c} 1\ 2\ 3 \\ 1\ 1\ 1 \\ 0\ 1\ 1 \\ 0\ 1\ 0 \\ 0\ 1\ 0 \\ 1\ 1\ 0 \\ 0\ 0\ 0 \\ 0\ 1\ 0 \\ 0\ 0\ 0 \end{array}$
--	---------------	--

We can now represent the dynamics in the state space:




Thus, there are two chains, the smaller of them containing a cycle of period 2, the other one a fixed point. All the other 5 states are transient, that is, the dynamics never returns to them.

A small change, here of the Boolean function of the second element

$\begin{array}{c c} 2 & 3 \\ \hline 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{array}$	$\begin{array}{c c} 1 & 3 \\ \hline 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{array}$	$\begin{array}{c c} 1 & 2 \\ \hline 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{array}$
---	---	---

$\begin{array}{c} 1 & 2 & 3 \\ \hline 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{array}$	→	$\begin{array}{c} 1 & 2 & 3 \\ \hline 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{array}$
---	---	---

can induce a qualitatively different dynamics: Namely, there is only one chain left which contains a cycle of length 3:




Next follows an example for $N = 4, K = 2$. First, we show again the input-output table for the 4 elements:

$\begin{array}{c c} 2 & 3 \\ \hline 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{array}$	$\begin{array}{c c} 1 & 3 \\ \hline 0 & 0 \\ 1 & 1 \\ 0 & 1 \\ 1 & 1 \end{array}$	$\begin{array}{c c} 2 & 4 \\ \hline 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{array}$	$\begin{array}{c c} 1 & 2 \\ \hline 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{array}$
---	---	---	---

Here is the table of the dynamic transition: $(\sigma_1, \sigma_2, \sigma_3, \sigma_4) \mapsto T(\sigma_1, \sigma_2, \sigma_3, \sigma_4)$

<u>1 2 3 4</u>	→	<u>1 2 3 4</u>
0 0 0 0		0 0 1 1
0 0 0 1		0 0 0 1
0 0 1 0		0 1 1 1
0 0 1 1		0 1 0 1
0 1 0 0		0 0 0 1
0 1 0 1		0 0 1 1
0 1 1 0		1 1 0 1
0 1 1 1		1 1 1 1
1 0 0 0		0 1 1 1
1 0 0 1		0 1 0 1
1 0 1 0		0 1 1 1
1 0 1 1		0 1 0 1
1 1 0 0		0 1 0 0
1 1 0 1		0 1 1 0
1 1 1 0		1 1 0 0
1 1 1 1		1 1 1 0

This transition generates three disjoint chains:



In the first chain, each state is eventually transformed into one of the states (0011) and (0101), and these two then alternate. In the second chain, everything ends up in the invariant state (0001). The third chain contains only a cycle of length 2.

Small changes of the update rules can affect these structures strongly. For instance, when we replace the last row of the rule for element 2 by

$$\begin{array}{c|c} 1 & 3 & 2 \\ \hline 1 & 1 & 0 \end{array},$$

then (1111) will lead into (1010), that in turn to (0011), and thence as before to (0101). Thus, the second chain is broken up, and its left part now enters into the first chain. (1011), however, leads to (0001) which remains invariant.

(1110) leads to (1000), thus again eventually entering into the previous first chain. The rest remains as before.

These examples show that there are two types of states in each case. The first type are the transient states that are never revisited by the dynamics. Among them, we even find some that can only occur as initial states, but cannot be reached from any other. Such states that are out of reach from elsewhere are called Garden-of-Eden states. The second type are those states that lie on attractors. These states are traversed again and again once having been reached for the first time. Each such state is contained in a periodic cycle. In the special case of a fixed state, the period length is 1.

In the same manner as cellular automata, Boolean networks can be embedded into a dynamics on the unit cube $[0, 1]^N$, so that again one can compute Lyapunov exponents and entropies. Of course, one may also compute discrete entropies

$$-\sum_{i \in Y} p_n(i) \log_2 p_n(i)$$

as before, where $p_n(i)$ is the relative frequency of state i in the n^{th} iteration step. As before, this is a nonincreasing function of n .

The examples now permit a somewhat more detailed discussion of these quantities. In particular, the expressions $p_n(i)$ are not necessarily monotonic. We look at the first chain in the first example. For $i = (101)$ we have $p_n(i) = 0$ for $n \geq 1$, because this state can occur at most as an initial state, i.e., for $n = 0$, but not as the successor state of any other. Now $p_1(000) = \frac{1}{4}$, since (000) has two different initial states as predecessors, namely (101) and (111). However $p_1(111) = \frac{1}{8}$, since (000) is the only state leading to (111). For $n = 2$, however, the roles are reversed, since now $p_2(000) = \frac{1}{8}$, because only from the initial state (000) itself, can one return to (000) for $n=2$, while we obtain $p_2(111) = \frac{1}{4}$. More generally

$$\begin{aligned} p_{2m-1}(000) &= \frac{1}{4}, & p_{2m-1}(111) &= \frac{1}{8}, \\ p_{2m}(000) &= \frac{1}{8}, & p_{2m}(111) &= \frac{1}{4} \quad \text{for } m = 1, 2, 3, \dots \end{aligned}$$

The entropy $-\sum_i p_n(i) \log_2 p_n(i)$ stabilizes for $n \geq 2$ and becomes $-\frac{1}{8} \log_2 \frac{1}{8} - \frac{1}{4} \log_2 \frac{1}{4} - \frac{5}{8} \log_2 \frac{5}{8}$.

It would make more sense, however, to assign each of the states (000) and (111) the asymptotic probability $\frac{3}{16}$ instead of letting it oscillate between $\frac{1}{8}$ and $\frac{1}{4}$. For instance, we might put

$$p_\infty(i) := \frac{1}{n} \sum_{k=0}^{n-1} p_k(i),$$

and then define the entropy as

$$H_\infty(T) := - \sum p_\infty(i) \log_2 p_\infty(i).$$

This would make the definition independent of the transient aspects of the dynamics and capture only the types of attractors, more precisely their periods, as well as the sizes of their basins of attraction, and it would thus reflect only the asymptotic or permanent aspects of the dynamics.

Of course, Grassberger's concept of effective measure complexity as discussed for cellular automata in 8.1 can also be usefully applied here.

One is interested in statistical properties depending on K and N when the Boolean functions for each element are randomly chosen. $K = N$ of course is the most general case. Here, the average cycle length grows like $2^{N/2}$ while the number of different cycles is proportional to N only. The typical behavior of such a network is chaotic in the sense that changing the states of some element typically will lead into a different chain. Attractors thus typically have small basins of attraction.

Similar things happen for $K \geq 3$. For $K = 2$, however, we see a more stable behavior. The average cycle length is of order $N^{1/2}$ only, as is the number of cycles. Small perturbations are often corrected. If the state of some element is altered permanently, this perturbation will spread only to some limited part of the network, and some elements with periodically oscillating states can turn into ones with fixed states, and conversely. There are fewer attractors now than for $K > 2$, and these attractors have much larger basins of attraction. This is the reason for the greater stability. For $K = 1$, finally, invariant states dominate the dynamics.

The preceding results stem mainly from computer simulations (see [59] for sophisticated tools), or have been obtained with the help of methods of statistical mechanics (mean field approximation). For a survey of what is known about the properties of random Boolean networks, we refer to [1].

References

1. M.Aldana, S.Coppersmith, L.Kadanoff, Boolean dynamics with random couplings, arXiv:nlin.AO/0204062; also in: E.Kaplan, J.Marsden, K.Sreenivasan (eds.), *Perspectives and problems in nonlinear science*, Springer, 2003
2. K.Alligood, T.Sauer, J.Yorke, *Chaos*, Springer, 1996
3. D.V.Anosov, V.I.Arnold (eds.), *Dynamical systems I*, *Encycl.Math.Sciences* 1, Springer, 1988
4. L.Arnold, *Random dynamical systems*, Springer, 1998
5. C.Beck, F.Schlögl, *Thermodynamics of chaotic systems*, Cambridge Univ.Press, 1993
6. S.N.Chow, J.Hale, *Methods of bifurcation theory*, Springer, ²1996
7. C.Conley, E.Zehnder, Morse-type index theory for flows and periodic solutions for Hamiltonian equations, *Comm.Pure Appl.Math.* 37, 207 - 253 (1984)
8. H.Crauel, F.Flandoli, Additive noise destroys a pitchfork bifurcation, *J.Dyn.Diff.Eqs.* 10, 259–274, 1998
9. R.Easton, *Geometric methods for discrete dynamical systems*, Oxford Univ. Press, 1998
10. W.Ebeling, J.Freund, F.Schweitzer, *Komplexe Strukturen: Entropie und Information*, Teubner, 1998
11. P.Grassberger, Toward a quantitative theory of self-generated complexity, *Int.J.Theo.Phys.*25, 907–938 (1986)
12. G.Grimmett, D.Stirzaker, *Probability and random processes*, Oxford Univ.Press, ³2001
13. J.Guckenheimer, P.Holmes, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, Springer, 1983
14. H.Haken, *Synergetics*, Springer, ³1983
15. H.Haken, *Advanced synergetics*, Springer, 1983
16. J.Hofbauer, K.Sigmund, *The theory of evolution and dynamical systems*, Cambridge Univ. Press, 1988
17. J.Hofbauer, K.Sigmund, *Evolutionary games and population dynamics*, Cambridge Univ. Press, 1998
18. A.Ilachinski, *Cellular automata*, World Scientific, 2001
19. D.-Q.Jiang, M.Qian, M.-P.Qian, *Mathematical theory of nonequilibrium steady states*, Springer Lecture Notes Math.1833, 2004
20. J.Jost, On the notion of complexity, *Theory in Bioscienc.* 117, 161 – 171 (1998)
21. J.Jost, *Riemannian geometry and geometric analysis*, Springer, ⁴2005
22. J.Jost, *Partial differential equations*, Springer, 2002
23. H.Kantz, T.Schreiber, *Nonlinear time series analysis*, Cambridge Univ.Press, 1997

24. A.Katok, B.Hasselblatt, *Introduction to the modern theory of dynamical systems*, Cambridge Univ. Press, 1995
25. St.Kauffman, *The origins of order*, Oxford Univ. Press, 1993
26. A.Lasota, M.Mackey, *Chaos, fractals, and noise*, Springer, ²1994
27. F.Ledrappier, L.-S.Young, The metric entropy of diffeomorphisms. Part I: Characterization of measures satisfying Pesin's formula. Part II: Relations between entropy, exponents and dimension, *Ann. Math.* 122, 509–539, 540–574 (1985)
28. E.Lorenz, Deterministic non-periodic flow, *J.Atmos.Sci.*20, 130–141 (1963)
29. R.Mañé, A proof of the C^1 stability conjecture, *Publ.Math.IHES* 66, 161–210 (1988)
30. J.Milnor, On the concept of attractor, *Comm.Math.Phys.*99, 177–195 (1985)
31. K.Mischaikow, *Conley Index theory*, Springer Lecture Notes Math. 1609, 119 – 207 (1995)
32. J.Moser, *Stable and random motions in dynamical systems*, Princeton Univ.Press, 1973
33. M.Mrozek, Leray functor and cohomology index for discrete dynamical systems, *Trans. AMS* 318, 149 – 178 (1990)
34. M.Mrozek, K.Rybakowski, A cohomological Conley index for maps on metric spaces, *J. Diff. Eq.* 90, 143 – 171 (1991)
35. J.Murray, *Mathematical biology*, Springer, ³2002
36. B. Øksendal, *Stochastic differential equations*, Springer, ⁴1995
37. J.Palis, W.de Melo, *Geometric theory of dynamical systems*, Springer, 1982
38. K.Petersen, *Ergodic theory*, Cambridge Univ. Press, 1983
39. A.Pikovsky, M.Rosenblum, J.Kurths, *Synchronization*, Cambridge Univ.Press, 2001
40. C.Robinson, *Dynamical systems*, CRC Press, 1999
41. D.Ruelle, *Chaotic evolution and strange attractors*, Cambridge Univ.Press, 1989
42. D.Ruelle, *Elementary differentiable dynamics and bifurcation theory*, Academic Press, 1989
43. D.Ruelle, Positivity of entropy production in nonequilibrium statistical mechanics, *J.Stat.Phys.*85, 1–23 (1996)
44. D.Ruelle, Smooth dynamics and new theoretical ideas in nonequilibrium statistical mechanics, *J.Stat.Phys.*95, 393–468 (1999)
45. D.Ruelle, F.Takens, On the nature of turbulence, *Comm. Math. Phys.* 20, 167–192 (1971), Note concerning our paper "On the nature of turbulence", *Comm. Math. Phys.* 23, 343–344 (1971)
46. D.Salamon, Morse theory, the Conley index and Floer homology, *Bull.LMS* 22, 113–140 (1990)
47. M.Schwarz, *Morse homology*, Birkhäuser, 1993
48. Ya.G.Sinai (ed.), *Dynamical systems II*, Encycl.Math.Sciences 2, Springer, 1989
49. S.Smale, Differentiable dynamical systems, *Bull.AMS* 73, 747-817 (1967)
50. J.Smoller, *Shock waves and reaction diffusion equations*, Springer, 1983
51. J.C.Sprott, *Chaos and time-series analysis*, Oxford Univ.Press, 2003
52. R.Stoecker, H.Zieschang, *Algebraische Topologie*, Teubner, ²1994
53. St.Strogatz, *Nonlinear dynamics and chaos*, Perseus Books, 1994
54. R.Thom, *Structural stability and morphogenesis*, Benjamin, Reading, Mass., 1975
55. W.Tucker, The Lorenz attractor exists, *C.R.Acad.Sci.Paris* 328, Série I, 1197–1202 (1999)

56. M.Viana, *Stochastic dynamics of deterministic systems*, Preprint, IMPA, 1997
57. P.Walters, *An introduction to ergodic theory*, Springer, GTM 79, 1982
58. S.Wolfram, *Cellular automata and complexity*, Collected papers, Addison-Wesley, Reading, Mass., 1994
59. A.Wuensche, M.Lesser, *The global dynamics of cellular automata: An atlas of basin of attraction fields of one-dimensional cellular automata*, Addison-Wesley, 1992; see also the website www.ddlab.com
60. E.Zeidler, *Teubner-Taschenbuch der Mathematik*, Teil II, Kapitel 13, Teubner, 1995

Index

- admissible ordering, 71
- $\alpha(y)$, 8
- Andronov-Hopf bifurcation, 46
- Anosov diffeomorphism, 158, 159
- asymptotic, 1
- asymptotic dynamics, 8
- asymptotic final state, 2
- attracting fixed point, 15
- attractor, 2, 23, 140, 163, 179
- autonomous, 9
- average information, 130
- Axiom A diffeomorphism, 159, 162

- backward invariant set, 93
- Baire, theorem of, 46
- baker's map, 107
- basin of attraction, 24
- Bernoulli shift, 104, 109, 122, 126, 136–138, 142, 161, 169
- Betti number, 66
- bifurcation, 45, 93
- bifurcation point, 45
- binary sequence, 130
- Birkhoff ergodic theorem, 124
- block (in cellular automaton), 172
- Boolean function, 175
- Boolean network, 175
- Borel set, 119
- boundary operator, 64, 68, 69
- Bowen's definition of entropy, 111
- Brownian motion, 144, 145, 147

- Cantor set, 110
- catastrophe theory, 45
- cellular automaton, 169
- center manifold, 55, 58
- center-stable manifold, 54
- center-unstable manifold, 54

- chain complex, 69
- chaos, 22
- chaotic attractor, 31
- chaotic dynamics, 4, 108, 163
- chaotic orbit, 31
- chaotic set, 31
- cocycle property, 148
- cohomological Conley index, 77, 96
- cohomology group, 69
- cohomology sequence, 83
- complexity of cellular automata, 171
- complexity of transformation, 115
- conditional entropy, 128
- conditional measure, 128
- conjugate, 43
- conjugate measure preserving transformations, 121
- conjugation, 102
- Conley index, 77, 83, 85, 96
- Conley theory, 71, 82, 88
- continuation, 92
- continuous time, 1
- correlation, 125
- coupled map lattice, 39
- covering of circle, 120, 142
- critical point, 12, 62, 71, 115, 159
- critical point of a continuous function, 91
- critical vertex, 61
- cylinder set, 123, 142, 144

- deformation retract, 82
- delay, 34
- delay-differential equation, 35
- dense orbit, 25
- density, 140
- diffeomorphism, 154
- difference equation, 29

- differentiable manifold, 62
- diffusion process, 37
- Dirac measure, 119
- discrete dynamical system, 8
- discrete time, 1
- dynamical network, 40
- dynamical stability, 4
- dynamical system, 1

- effective measure complexity, 132, 173, 174
- enslaving principle, 56
- entropy, 101, 103, 129, 130, 132–134, 137, 165, 173
- entropy of a partition, 133
- entropy of Boolean network, 179
- entropy of a covering, 100
- entropy of a partition, 100, 126
- entropy of a transformation, 101
- entropy of cellular automaton, 171
- ε -divergent, 115
- ε -orbit, 162
- ε -separated, 114
- ergodic, 122, 126, 139, 153, 155, 156
- ergodic basin, 141
- ergodic problem, 5
- ergodic theorem, 124, 125
- ergodic theory, 5
- ergodic, 133
- Euler characteristic, 61, 66
- evolution, 3
- exit set, 76
- exponential correlation decay, 126
- exponentially attracting, 16
- exponentially mixing, 126

- factor, 106, 120
- final state, 2
- fixed point, 7
- Floer homology, 64, 87, 161
- Floer homology over \mathbb{Z} , 70
- Floer, theorem of, 64
- flow, 7
- flow lines, 63
- Fokker-Planck equation, 147, 150
- fold map, 106
- folding entropy, 168
- forward limit set, 8
- fractal, 114

- Garden-of-Eden state, 179
- generator, 135
- generic, 46, 63, 159
- geodesic flow, 164
- gradient flow, 12, 62, 71, 90, 114, 115, 159
- graph transformation method, 54
- group property, 148

- Hadamard-Perron, theorem of, 53, 160
- Haken, 56
- halting problem, 170
- harmonic function, 147
- Hartman-Grobman, theorem of, 43
- heat equation, 37
- heteroclinic orbit, 18, 73
- homoclinic orbit, 17, 23, 73
- homoclinic tangency, 158
- homoclinic theorem, 161
- homology group, 64, 65, 68
- homotopic, 81
- homotopy equivalence, 81
- homotopy type, 83
- Hopf bifurcation, 17, 46, 47
- horse-shoe, 97, 106, 107, 109, 159, 161
- hyperbolic, 14, 156–158
- hyperbolic linear map, 42
- hyperbolic attractor, 163
- hyperbolic family of linear maps, 157
- hyperbolic fixed point, 42
- hyperbolic fixed point, 12, 42, 77
- hyperbolic linear map, 12
- hyperbolic periodic orbit, 42, 80
- hyperbolic periodic point, 42
- hyperbolic set, 158, 160–162

- identity problem, 3
- independent and identically distributed random variables, 145
- index of a critical point, 91
- index of critical vertex, 61
- index pair, 76, 77, 83, 95
- information, 2, 129, 130
- initial state, 1, 2, 7
- input-output table, 176
- interval map, 106
- intrinsic length scale, 114
- invariant, 3, 4
- invariant measure, 5, 138, 139

- invariant set, 70, 93, 161
- invariant state, 2
- invertible, 8
- irreducible attractor, 25
- irreversibility, 171
- isolated invariant set, 76, 94
- isolating neighborhood, 76, 94
- isomorphic measure preserving transformations, 120
- isomorphism problem, 3
- Itô integral, 146

- join of coverings, 100
- join of punctured spaces, 82

- Katok, theorem of, 138
- Kauffman's Boolean networks, 175
- Kolmogorov equation, 147, 150
- Kolmogorov extension theorem, 144
- Kolmogorov-Sinai entropy, 134
- Kolmogorov-Sinai, theorem of, 134
- Krylov-Bogolubov, theorem of, 139
- Kupka-Smale, theorem of, 162

- Langevin equation, 146
- Ledrappier-Young, theorem of, 164
- length scale, 114
- Leray reduction, 96
- limit sets, 8
- linearize, 10
- local flow, 9
- logistic differential equation, 33
- logistic equation, 29
- logistic map, 31, 32, 107, 109, 142
- Lorenz, 28
- Lotka-Volterra equation, 49
- Lyapunov exponent, 31, 153–156, 167
- Lyapunov function, 46
- Lyapunov stable, 23

- Möbius strip, 80
- Mañé, theorem of, 162
- Markov property, 137
- Markov shift, 124, 136
- maximum principle, 37
- measurable space, 119
- measurable transformation, 120
- measure theoretic entropy, 156
- measure algebra, 120
- measure equivalent, 120
- measure preserving, 100, 120
- measure theoretic entropy, 134, 137, 138
- metric entropy, 111, 113
- Milnor attractor, 28
- mixing, 125
- Morse decomposition, 71, 75, 90
- Morse filtration, 76, 83, 86
- Morse lemma, 90
- Morse set, 71
- multiplicative ergodic theorem, 154
- multiplicity of critical vertex, 62
- multiplicity of Lyapunov exponent, 154
- mutual information, 132

- (n, ε) -separated, 113
- (n, ε) -spanning, 112
- node, 10
- noise, 4
- non-autonomous, 8, 9
- non-degenerate critical point, 13
- non-wandering point, 159
- nondegenerate critical point, 63
- normal hyperbolic set, 79

- $\omega(y)$, 8
- one-dimensional cellular automata, 170
- one-sided 2-shift, 104, 108
- orbit, 7
- Ornstein, theorem of, 137
- Oseledec, theorem of, 154

- parameter, 1, 4
- parameterized flow, 92
- partition, 126
- period 3, 30
- period doubling bifurcation, 21
- periodic, 7
- periodic orbit, 73, 78, 79
- periodic point, 12, 159
- Perron-Frobenius operator, 163
- Pesin's identity, 156
- phase space, 7
- Picard-Lindelöf, theorem of, 9
- Poincaré Recurrence Theorem, 121
- Poincaré return map, 26, 41
- positive entropy, 106, 107
- positively invariant set, 94

- predator-prey model, 49
- prediction, 1
- probability measure, 119
- probability space, 119, 142
- pseudo-hyperbolicity, 53
- pseudo-orbit, 162
- punctured space, 81

- quasiperiodic, 27

- Rössler, 28
- random difference equation, 148
- random differential equation, 149
- random dynamical system, 148
- random measure, 149
- random objects, 149
- random variable, 142
- reaction-diffusion equation, 38
- recursive, 1
- reduced equation, 58
- related by continuation, 92
- relative Betti number, 83
- relative homology group, 68
- relative index, 70
- repelling fixed point, 15
- Riemannian metric, 63, 154, 158
- rotation of circle, 122, 126, 135, 141, 156
- Ruelle, theorem of, 156, 166

- saddle, 10
- sample path, 145
- scale, 71
- scale of largest complexity, 115
- scientific modelling, 3
- self-similar structure, 114
- semi-conjugate, 121
- semiflow, 7
- semigroup, 7
- separated subset, 113
- separation, 117
- shadowing, 162
- Shannon-McMillan-Breiman, theorem of, 134
- Sharkovsky, theorem of, 30
- shift, 104
- σ -algebra, 119
- sink, 10
- slaving equation, 58

- Smale's horse-shoe, 97
- Smale's transversality condition, 161
- Smale, theorem of, 161
- solenoid, 160
- source, 10
- spanning subset, 112
- spatial entropy of cellular automaton, 172
- sphere, 67, 85
- spiral, 11
- SRB measure, 141, 167
- SRB-measure, 163
- stability problem, 4
- stable, 10, 52
- stable manifold, 52, 160
- state space, 7, 142
- state transition, 1
- stationary point, 7
- stationary stochastic process, 143
- statistical behavior, 5
- stochastic bifurcation, 149
- stochastic differential equation, 144, 145, 149
- stochastic integral, 146
- stochastic process, 142
- strange attractor, 28
- strange attractor, 140, 161
- Stratonovich, 146
- strong deformation retract, 82
- strongly mixing, 125
- structural change, 4
- structural stability, 5, 44
- structurally stable, 11, 17, 43, 158, 162
- structurally stable bifurcation, 45
- subcritical Hopf bifurcation, 47
- supercritical Hopf bifurcation, 47
- symbolic dynamics, 109
- synergetics, 56

- temporal entropy of cellular automaton, 173
- tent map, 32, 103, 107, 108, 142
- Thom, 21, 45
- time, 7
- time 1-map, 8
- time reversal, 10
- time series, 5
- topological complexity of transformation, 116

- topological conjugation, 102
- topological entropy, 101, 111, 137
- topological factor, 106
- topological invariant, 65, 66
- topological space, 99
- topologically conjugate, 43, 103, 106
- torus, 26, 67, 68, 157, 160
- trajectory, 7
- transfer operator, 163
- transient state, 176, 179
- transition map for Boolean network, 176
- transition probability, 124
- transition rule, 1
- transversality, 161, 162
- transverse homoclinic point, 161
- Turing machine, 170
- turnpike theorem, 164
- 2-shift, 104
- two-sided 2-shift, 104
- uncertainty, 137
- uniquely ergodic, 139
- universal computer, 170
- unstable, 10
- unstable manifold, 52, 160
- unstable periodic orbit, 78
- Von Neumann ergodic theorem, 125
- weakly mixing, 125
- white noise, 145
- Wiener measure, 144, 145
- Wiener process, 145

Universitext

- Aguilar, M.; Gitler, S.; Prieto, C.:* Algebraic Topology from a Homotopical Viewpoint
- Aksoy, A.; Khamsi, M. A.:* Methods in Fixed Point Theory
- Alevras, D.; Padberg M. W.:* Linear Optimization and Extensions
- Andersson, M.:* Topics in Complex Analysis
- Aoki, M.:* State Space Modeling of Time Series
- Arnold, V. I.:* Lectures on Partial Differential Equations
- Audin, M.:* Geometry
- Aupetit, B.:* A Primer on Spectral Theory
- Bachem, A.; Kern, W.:* Linear Programming Duality
- Bachmann, G.; Narici, L.; Beckenstein, E.:* Fourier and Wavelet Analysis
- Badescu, L.:* Algebraic Surfaces
- Balakrishnan, R.; Ranganathan, K.:* A Textbook of Graph Theory
- Balser, W.:* Formal Power Series and Linear Systems of Meromorphic Ordinary Differential Equations
- Bapat, R.B.:* Linear Algebra and Linear Models
- Benedetti, R.; Petronio, C.:* Lectures on Hyperbolic Geometry
- Benth, F.E.:* Option Theory with Stochastic Analysis
- Berberian, S. K.:* Fundamentals of Real Analysis
- Berger, M.:* Geometry I, and II
- Bliedtner, J.; Hansen, W.:* Potential Theory
- Blowey, J.F.; Coleman, J.P.; Craig, A. W. (Eds.):* Theory and Numerics of Differential Equations
- Blyth, T. S.:* Lattices and Ordered Algebraic Structures
- Börger, E.; Grädel, E.; Gurevich, Y.:* The Classical Decision Problem
- Böttcher, A; Silbermann, B.:* Introduction to Large Truncated Toeplitz Matrices
- Boltyanski, V.; Martini, H.; Soltan, P. S.:* Excursions into Combinatorial Geometry
- Boltyanskii, V. G.; Efremovich, V. A.:* Intuitive Combinatorial Topology
- Bonnans, J. F.; Gilbert, J. C.; Lemaréchal, C.; Sagastizábal, C. A.:* Numerical Optimization
- Booss, B.; Bleecker, D. D.:* Topology and Analysis
- Borkar, V. S.:* Probability Theory
- Brunt B. van:* The Calculus of Variations
- Carleson, L.; Gamelin, T. W.:* Complex Dynamics
- Cecil, T. E.:* Lie Sphere Geometry: With Applications of Submanifolds
- Chae, S. B.:* Lebesgue Integration
- Chandrasekharan, K.:* Classical Fourier Transform
- Charlap, L. S.:* Bieberbach Groups and Flat Manifolds
- Chern, S.:* Complex Manifolds without Potential Theory
- Chorin, A. J.; Marsden, J. E.:* Mathematical Introduction to Fluid Mechanics
- Cohn, H.:* A Classical Invitation to Algebraic Numbers and Class Fields
- Curtis, M. L.:* Abstract Linear Algebra
- Curtis, M. L.:* Matrix Groups
- Cyganowski, S.; Kloeden, P.; Ombach, J.:* From Elementary Probability to Stochastic Differential Equations with MAPLE
- Dalen, D. van:* Logic and Structure
- Das, A.:* The Special Theory of Relativity: A Mathematical Exposition
- Debarre, O.:* Higher-Dimensional Algebraic Geometry
- Deitmar, A.:* A First Course in Harmonic Analysis
- Demazure, M.:* Bifurcations and Catastrophes

- Devlin, K. J.:* Fundamentals of Contemporary Set Theory
- DiBenedetto, E.:* Degenerate Parabolic Equations
- Diener, F.; Diener, M. (Eds.):* Nonstandard Analysis in Practice
- Dimca, A.:* Sheaves in Topology
- Dimca, A.:* Singularities and Topology of Hypersurfaces
- DoCarmo, M. P.:* Differential Forms and Applications
- Duistermaat, J. J.; Kolk, J. A. C.:* Lie Groups
- Edwards, R. E.:* A Formal Background to Higher Mathematics Ia, and Ib
- Edwards, R. E.:* A Formal Background to Higher Mathematics IIa, and IIb
- Emery, M.:* Stochastic Calculus in Manifolds
- Endler, O.:* Valuation Theory
- Erez, B.:* Galois Modules in Arithmetic
- Everest, G.; Ward, T.:* Heights of Polynomials and Entropy in Algebraic Dynamics
- Farenick, D. R.:* Algebras of Linear Transformations
- Foulds, L. R.:* Graph Theory Applications
- Franke, J.; Härdle, W.; Hafner, C. M.:* Statistics of Financial Markets: An Introduction
- Frauenthal, J. C.:* Mathematical Modeling in Epidemiology
- Friedman, R.:* Algebraic Surfaces and Holomorphic Vector Bundles
- Fuks, D. B.; Rokhlin, V. A.:* Beginner's Course in Topology
- Fuhrmann, P. A.:* A Polynomial Approach to Linear Algebra
- Gallot, S.; Hulin, D.; Lafontaine, J.:* Riemannian Geometry
- Gardiner, C. F.:* A First Course in Group Theory
- Gårding, L.; Tambour, T.:* Algebra for Computer Science
- Godbillon, C.:* Dynamical Systems on Surfaces
- Godement, R.:* Analysis I, and II
- Goldblatt, R.:* Orthogonality and Spacetime Geometry
- Gouvêa, F. Q.:* p -Adic Numbers
- Gross, M. et al.:* Calabi-Yau Manifolds and Related Geometries
- Gustafson, K. E.; Rao, D. K. M.:* Numerical Range. The Field of Values of Linear Operators and Matrices
- Gustafson, S. J.; Sigal, I. M.:* Mathematical Concepts of Quantum Mechanics
- Hahn, A. J.:* Quadratic Algebras, Clifford Algebras, and Arithmetic Witt Groups
- Hájek, P.; Havránek, T.:* Mechanizing Hypothesis Formation
- Heinonen, J.:* Lectures on Analysis on Metric Spaces
- Hlawka, E.; Schoißengeier, J.; Taschner, R.:* Geometric and Analytic Number Theory
- Holmgren, R. A.:* A First Course in Discrete Dynamical Systems
- Howe, R., Tan, E. Ch.:* Non-Abelian Harmonic Analysis
- Howes, N. R.:* Modern Analysis and Topology
- Hsieh, P.-F.; Sibuya, Y. (Eds.):* Basic Theory of Ordinary Differential Equations
- Humi, M., Miller, W.:* Second Course in Ordinary Differential Equations for Scientists and Engineers
- Hurwitz, A.; Kritikos, N.:* Lectures on Number Theory
- Huybrechts, D.:* Complex Geometry: An Introduction
- Isaev, A.:* Introduction to Mathematical Methods in Bioinformatics
- Istas, J.:* Mathematical Modeling for the Life Sciences
- Iversen, B.:* Cohomology of Sheaves
- Jacod, J.; Protter, P.:* Probability Essentials
- Jennings, G. A.:* Modern Geometry with Applications
- Jones, A.; Morris, S. A.; Pearson, K. R.:* Abstract Algebra and Famous Impossibilities
- Jost, J.:* Compact Riemann Surfaces

- Jost, J.*: Dynamical Systems. Examples of Complex Behaviour
- Jost, J.*: Postmodern Analysis
- Jost, J.*: Riemannian Geometry and Geometric Analysis
- Kac, V.; Cheung, P.*: Quantum Calculus
- Kannan, R.; Krueger, C. K.*: Advanced Analysis on the Real Line
- Kelly, P.; Matthews, G.*: The Non-Euclidean Hyperbolic Plane
- Kempf, G.*: Complex Abelian Varieties and Theta Functions
- Kitchens, B. P.*: Symbolic Dynamics
- Kloeden, P.; Ombach, J.; Cyganowski, S.*: From Elementary Probability to Stochastic Differential Equations with MAPLE
- Kloeden, P. E.; Platen, E.; Schurz, H.*: Numerical Solution of SDE Through Computer Experiments
- Kostrikin, A. I.*: Introduction to Algebra
- Krasnoselskii, M. A.; Pokrovskii, A. V.*: Systems with Hysteresis
- Kurzweil, H.; Stellmacher, B.*: The Theory of Finite Groups. An Introduction
- Lang, S.*: Introduction to Differentiable Manifolds
- Luecking, D. H., Rubel, L. A.*: Complex Analysis. A Functional Analysis Approach
- Ma, Zhi-Ming; Roekner, M.*: Introduction to the Theory of (non-symmetric) Dirichlet Forms
- Mac Lane, S.; Moerdijk, I.*: Sheaves in Geometry and Logic
- Marcus, D. A.*: Number Fields
- Martinez, A.*: An Introduction to Semiclassical and Microlocal Analysis
- Matoušek, J.*: Using the Borsuk-Ulam Theorem
- Matsuki, K.*: Introduction to the Mori Program
- Mazzola, G.; Milmeister G.; Weissman J.*: Comprehensive Mathematics for Computer Scientists 1
- Mazzola, G.; Milmeister G.; Weissman J.*: Comprehensive Mathematics for Computer Scientists 2
- McCarthy, P. J.*: Introduction to Arithmetical Functions
- McCrimmon, K.*: A Taste of Jordan Algebras
- Meyer, R. M.*: Essential Mathematics for Applied Field
- Meyer-Nieberg, P.*: Banach Lattices
- Mikosch, T.*: Non-Life Insurance Mathematics
- Mines, R.; Richman, F.; Ruitenburg, W.*: A Course in Constructive Algebra
- Moise, E. E.*: Introductory Problem Courses in Analysis and Topology
- Montesinos-Amilibia, J. M.*: Classical Tesselations and Three Manifolds
- Morris, P.*: Introduction to Game Theory
- Nikulin, V. V.; Shafarevich, I. R.*: Geometries and Groups
- Oden, J. J.; Reddy, J. N.*: Variational Methods in Theoretical Mechanics
- Øksendal, B.*: Stochastic Differential Equations
- Øksendal, B.; Sulem, A.*: Applied Stochastic Control of Jump Diffusions
- Poizat, B.*: A Course in Model Theory
- Polster, B.*: A Geometrical Picture Book
- Porter, J. R.; Woods, R. G.*: Extensions and Absolutes of Hausdorff Spaces
- Radjavi, H.; Rosenthal, P.*: Simultaneous Triangularization
- Ramsay, A.; Richtmeyer, R. D.*: Introduction to Hyperbolic Geometry
- Rees, E. G.*: Notes on Geometry
- Reisel, R. B.*: Elementary Theory of Metric Spaces
- Rey, W. J. J.*: Introduction to Robust and Quasi-Robust Statistical Methods
- Ribenboim, P.*: Classical Theory of Algebraic Numbers
- Rickart, C. E.*: Natural Function Algebras
- Rotman, J. J.*: Galois Theory

Rubel, L. A.: Entire and Meromorphic Functions
Ruiz-Tolosa, J. R.; Castillo E.: From Vectors to Tensors
Runde, V.: A Taste of Topology
Rybakowski, K. P.: The Homotopy Index and Partial Differential Equations
Sagan, H.: Space-Filling Curves
Samelson, H.: Notes on Lie Algebras
Schiff, J. L.: Normal Families
Sengupta, J. K.: Optimal Decisions under Uncertainty
Sérroul, R.: Programming for Mathematicians
Seydel, R.: Tools for Computational Finance
Shafarevich, I. R.: Discourses on Algebra
Shapiro, J. H.: Composition Operators and Classical Function Theory
Simonnet, M.: Measures and Probabilities
Smith, K. E.; Kahanpää, L.; Kekäläinen, P.; Traves, W.: An Invitation to Algebraic Geometry
Smith, K. T.: Power Series from a Computational Point of View
Smoryński, C.: Logical Number Theory I. An Introduction
Stichtenoth, H.: Algebraic Function Fields and Codes
Stillwell, J.: Geometry of Surfaces
Stroock, D. W.: An Introduction to the Theory of Large Deviations
Sunder, V. S.: An Invitation to von Neumann Algebras
Tamme, G.: Introduction to Étale Cohomology
Tondeur, P.: Foliations on Riemannian Manifolds
Toth, G.: Finite Möbius Groups, Minimal Immersions of Spheres, and Moduli
Verhulst, F.: Nonlinear Differential Equations and Dynamical Systems
Wong, M. W.: Weyl Transforms
Xambó-Descamps, S.: Block Error-Correcting Codes
Zaanen, A. C.: Continuity, Integration and Fourier Theory
Zhang, F.: Matrix Theory
Zong, C.: Sphere Packings
Zong, C.: Strange Phenomena in Convex and Discrete Geometry
Zorich, V. A.: Mathematical Analysis I
Zorich, V. A.: Mathematical Analysis II