

Characterisation of local properties and prediction in chaotic dynamical systems

Von der Fakultät für Mathematik und Naturwissenschaften
der Carl von Ossietzky Universität Oldenburg
zur Erlangung des Grades und Titels eines
Doktors der Naturwissenschaften (Dr. rer. nat.)
angenommene Dissertation

von
Markus Harle
geboren am 15.02.1976 in Wertingen

Gutachterin: Prof. Dr. Ulrike Feudel
Zweitgutachter: Dr. Jan Freund
Tag der Disputation: 18.07.2007

Abstract

Many physical systems of great social and economic importance exhibit chaotic behaviour that makes it difficult to predict future states of the system, often already for intermediate prediction times τ . This thesis aims at identifying regions in state space with extraordinary properties. These properties allow us to make claims about the state space region probably visited by the trajectory in the near future or the predictability of the system in a given state. The systems studied include conservative as well as dissipative systems. The identification of special regions in state space is done using two different kinds of local exponents, namely finite time Lyapunov exponents and maximum growth exponents, as well as ensemble studies.

One of the main parts of this thesis is concerned with identifying regions in state space where the trajectory gets trapped near islands of regular motion in conservative systems. This manifests itself in the finite time Lyapunov exponents approaching zero. These measure the short-time growth of infinitesimal errors, so one can say that the motion becomes less and less chaotic the closer one is to such an island. We find multiple maxima in the distribution of finite time Lyapunov exponents that can be attributed to pieces of the trajectory that stay in the vicinity of islands of different orders. This finding explains a feature recently published by a different group.

In the next part we look for a relation between predictability and the distribution of special points in state space, points of homoclinic tangency. We establish that the predictability is enhanced around such points. This is only so for one kind of local exponents, though, the finite time Lyapunov exponents. One can say that these are closely linked to the dynamics of the system, whereas the “maximum growth exponents” are not. The latter also measure the short-time growth of infinitesimal errors, but in a different way. The size of the region of enhanced predictability was studied and found to shrink exponentially with increasing prediction time τ .

The third major part of this treatise investigates the growth of finite errors, governed by the full nonlinear system, using an ensemble approach. Comparison of the maximum nonlinear error growth with local exponents shows good

agreement for the maximum growth exponents for small initial errors and short prediction times, but not for the finite time Lyapunov exponents. We establish a systematic dependence of worst-case error growth on prediction time τ , which is exponential for small τ and a power law for large τ . Various measures are used to study how the ensemble size influences the accuracy of worst-case, average and best-case error growth. For the especially interesting worst-case we find that the error due to using a small ensemble (compared to a hypothetical infinite ensemble) grows with decreasing ensemble size according to a power law. Furthermore, the worst case is the one that can actually be handled most accurately with a small ensemble.

What is made strikingly clear by the results of the last two parts mentioned is the fact that the choice of local exponents is very important when trying to quantify predictability using growth of infinitesimal errors. Depending on one's interest, one or the other type of local exponents may be the better choice.

Zusammenfassung

Viele physikalische Systeme von großer sozialer und ökonomischer Bedeutung zeigen chaotisches Verhalten, das eine Vorhersage zukünftiger Zustände des Systems schwierig macht, und zwar oft schon für mittelfristige Vorhersagen. Diese Dissertation hat zum Ziel, Regionen im Zustandsraum mit außergewöhnlichen Eigenschaften zu ermitteln. Diese Eigenschaften ermöglichen zum einen Aussagen über diejenige Region im Zustandsraum, in der sich die Trajektorie wahrscheinlich in der näheren Zukunft aufhalten wird. Zum anderen ermöglichen sie Aussagen über die Vorhersagbarkeit des Systems in einem gegebenen Zustand. Unter den betrachteten Systemen sind sowohl konservative als auch dissipative Systeme vertreten. Spezielle Regionen im Zustandsraum werden mit Hilfe zweier verschiedener Arten von lokalen Exponenten, nämlich *finite time* Lyapunov-Exponenten und *maximum growth* Exponenten, und mit Ensemblerechnungen ermittelt.

Einer der Hauptteile dieser Arbeit befasst sich mit dem Auffinden von Regionen im Zustandsraum konservativer Systeme, in denen die Trajektorie in der Nähe von Inseln regulärer Bewegung eingefangen wird (*trapping*). Dies äußert sich darin, daß sich die Werte der *finite time* Lyapunov-Exponenten Null nähern. Diese messen das kurzfristige Anwachsen infinitesimaler Störungen, weshalb man die Bewegung als umso weniger chaotisch bezeichnen kann, je näher sie einer solchen Insel ist. Wir haben mehrere Maxima in der Verteilung der *finite time* Lyapunov-Exponenten gefunden, die Abschnitten der Trajektorie zugeordnet werden können, die in der Nähe von Inseln verschiedener Ordnung bleiben. Diese Erkenntnis erklärt ein Resultat, das kürzlich von einer anderen Gruppe veröffentlicht wurde.

Im nächsten Teil suchen wir nach einer Beziehung zwischen Vorhersagbarkeit und der Verteilung besonderer Punkte im Zustandsraum, sogenannter *homoclinic tangencies*. Wir zeigen, daß um solche Punkte herum eine verbesserte Vorhersagbarkeit besteht. Dies gilt jedoch nur für eine Art von lokalen Exponenten, die *finite time* Lyapunov-Exponenten. Diese sind eng mit der Dynamik des Systems verknüpft, die *maximum growth* Exponenten hingegen nicht. Die Größe der Region verbesserter Vorhersagbarkeit schrumpft exponentiell mit zunehmendem Vorhersagezeitraum τ .

Der dritte Hauptteil dieser Abhandlung untersucht das Anwachsen endlicher Störungen, welches durch das volle nichtlineare System beschrieben wird, mit Hilfe von Ensembles. Ein Vergleich des maximalen nichtlinearen Fehlerwachstums mit den lokalen Exponenten zeigt eine gute Übereinstimmung für die *maximum growth* Exponenten im Falle kleiner Anfangsfehler und kurzer Vorhersagezeiträume. Dies gilt nicht für die *finite time* Lyapunov-Exponenten. Das Anwachsen der Fehler im schlimmsten Fall, dem *worst case*, hängt systematisch vom Vorhersagezeitraum τ ab, und zwar exponentiell für kleine τ und über ein Potenzgesetz für große τ . Verschiedene Maße werden benutzt, um den Einfluß der Ensemblegröße auf die Genauigkeit von *worst-case*-, mittlerem und *best-case*-Fehlerwachstum zu untersuchen. Für den besonders interessanten *worst case* ergibt sich, daß der Fehler aufgrund der Benutzung eines kleinen Ensembles (im Vergleich zu einem hypothetischen unendlich großen Ensemble) von der Ensemblegröße über ein Potenzgesetz abhängt. Ferner ist der *worst case* tatsächlich derjenige Fall, der mit einem kleinen Ensemble am besten behandelt werden kann.

Durch die Ergebnisse der letzten beiden Teile wird eines besonders augenfällig: Die Entscheidung, welche lokalen Exponenten man benutzt, ist sehr wichtig, wenn die Vorhersagbarkeit durch das Anwachsen infinitesimaler Störungen quantifiziert werden soll. Je nach Zielsetzung kann die eine oder andere Art lokaler Exponenten die bessere Wahl sein.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Predictability of dynamical systems	2
1.3	Limitations of predictability	6
1.4	Overview of this thesis	11
2	Concepts, methods and systems under investigation	15
2.1	Exponents as a measure for divergence of trajectories	15
2.1.1	Global Lyapunov exponents	15
2.1.2	Local exponents	19
2.1.3	Numerical calculation	22
2.2	Model systems studied	27
2.2.1	Hénon map	28
2.2.2	Standard map	28
2.2.3	Quasiperiodicity map	30
3	Local exponents in chaotic layers of Hamiltonian systems	35
3.1	Island structure in chaotic layers	36
3.2	Spectrum of local exponents	38
3.3	Distributions of local exponents	43
4	Local exponents and homoclinic tangencies	49
4.1	Definition of homoclinic tangencies	49
4.2	Calculation of homoclinic tangencies	52
4.3	Connection between homoclinic tangencies and regions of enhanced predictability	53
4.3.1	Hénon map	55
4.3.2	Standard map	59
5	Predictability studies using ensembles	63
5.1	Error growth factors as a measure of divergence of trajectories	64

Contents

5.2	Distributions of error growth factors	67
5.3	Comparison of error growth factors and local exponents	69
5.4	Dependence of the worst-case error growth on prediction time	75
5.5	Minimal ensemble size needed	79
6	Summary	89
	Bibliography	97
	Acknowledgements	105
	Lebenslauf	107

Chapter 1

Introduction

1.1 Motivation

The study of predictability is of considerable economic and social importance. Its relevance is obvious from the widespread application of forecasts in fields as diverse as stock prices (Sornette and Zhou, 2006), risk assessment of epidemics (Thomson et al., 2006), local precipitation patterns (Moron et al., 2003) and flood forecasts (Koussis et al., 2003) or sediment transport in rivers (McConchie et al., 2005). In order to alleviate the consequences of severe weather conditions, for instance, the reliability of forecasts has to be taken into account. Different means have been developed to do this by investigating the predictability for given weather conditions, i. e. a given point in state space of the dynamical system describing the relevant processes.

Unfortunately, the task of predicting the future behaviour of a system becomes difficult if the system under consideration is nonlinear and exhibits chaotic behaviour. On the other hand, this has led to a fascinating area of research.

Making a forecast is one thing, judging in advance how good it is is a different matter. If one cannot easily improve the forecast itself, one would at least like to have some quantitative estimation of how reliable the forecast is. Due to the exponential divergence of trajectories in chaotic systems, this becomes even more important when the time horizon of the forecast is extended. Since it is an essential issue to judge the quality of forecasts, the main aim of this work is to identify situations when a forecast can be trusted almost blindly or situations when a forecast should be considered little more than a not-so-educated guess. In other words, we would like to identify states of a dynamical system for which predictability is good and states for which it is bad.

Before considering the various influences limiting predictability, let us first clarify what we are talking about, i. e. define some terms involved.

1.2 Predictability of dynamical systems

The term dynamical system will be used herein to denote a set of rules that describe how a state \vec{x} of this system at a given time t is carried over to a new state \vec{x}' at a future time t' . We restrict our investigations to deterministic systems without any stochastic elements. Therefore, given the state of the system at a given time, $\vec{x}(t)$, the future states are completely determined by the rules defining the dynamical system.

The time variable t can, in general, be either continuous or discrete, i. e. $t \in \mathbb{R}$ or $t \in \mathbb{Z}$, respectively. In the case of continuous time, the rules describing how a state of the system evolves with time are typically written in the form of *differential equations*. The most simple way is as a set of first-order ordinary differential equations:

$$\frac{dx_j(t)}{dt} = F_j(x_1(t), \dots, x_d(t); p_1, \dots, p_k), \quad j \in \{1, \dots, d\}, \quad (1.1)$$

where x_j are the components of the system state vector \vec{x} and d is the dimension of the state space of the system. The values p_1, \dots, p_k are parameters that do not depend on time. If we start with a given system (with given parameter values p_k) and change one or more parameter values, what do we get? It would be possible to consider the resulting system a completely different system altogether, of course. However, it is more convenient not to do so, since the behaviour of dynamical systems is often similar over a range of parameter values. With vector notation, Eq. 1.1 can be simplified to

$$\dot{\vec{x}}(t) = \vec{F}(\vec{x}(t); \vec{p}). \quad (1.2)$$

In the case of discrete time, the evolution of system states is usually described by a *map* M :

$$x_j(t+1) = M_j(x_1(t), \dots, x_d(t); p_1, \dots, p_k), \quad j \in \{1, \dots, d\}, \quad (1.3)$$

or in short

$$\vec{x}(t+1) = M(\vec{x}(t); \vec{p}). \quad (1.4)$$

A *trajectory* or *orbit* is the path in the d -dimensional state space that the system describes as it evolves in time, starting at a given initial state $\vec{x}(t_0)$. Trajectories can take very different forms. First of all, one can differentiate between

“short-term” or transient behaviour and “long-term” or asymptotic behaviour. Typically the first part of a trajectory shows transient behaviour, producing markedly different characteristics from the asymptotic part of a trajectory. Depending on the initial state $\vec{x}(t_0)$ the system may need a shorter or longer time to finally reach its asymptotic behaviour. In this work we are only interested in the asymptotic behaviour, so we discard the transient part of trajectories. A notion related to the asymptotic behaviour of a trajectory is the *omega limit set* $\omega(\vec{x})$ of \vec{x} . This is the set of all accumulation points of the trajectory starting at \vec{x} .

An *invariant set* S is a subset of state space such that for each $\vec{x}(0) \in S$ the following holds: For every time $t \in \mathbb{R}$, $\vec{x}(t) \in S$. Here $\vec{x}(t)$ is the system state at time t , evolved from the state $\vec{x}(0)$ at time $t = 0$. In other words, if the state of the system is part of an invariant set, the system will never leave this set (neither for $t \rightarrow \infty$ nor for $t \rightarrow -\infty$). Such a situation can either arise when a trajectory starts on an invariant set or if a trajectory reaches an invariant set after the transient. Simple examples for an invariant set are a fixed point or a limit cycle.

A special case of an invariant set is an *attractor*. There is not one definition that fits all purposes, because in some situations it is preferable to relax a constraint in a definition instead of excluding a few special cases. We provide a quite general definition that was suggested by Milnor (1985): A closed subset A of state space is an attractor if

- the *basin of attraction*¹ $\rho(A) = \{\vec{x} \mid \omega(\vec{x}) \subset A\}$ has non-zero measure, i. e. there is a positive, non-zero probability that a randomly chosen state will approach A in the limit of infinite time, and
- there is no smaller subset A' of A such that $\rho(A')$ is the same as $\rho(A)$ except for a set of measure zero, i. e. by reducing A to A' one would miss some important part of the attractor.

This definition has the advantage that there is always at least one attractor present in a smooth dynamical system. Attractors are always closed invariant sets, and there are statements one can make about unions of attractors being attractors again and about maximal and minimal attractors (the one containing all others and those not reducible, respectively). These concepts are not part of the work presented here, however, so we refer to Milnor (1985) for details.

The asymptotic behaviour of trajectories can be classified into four cases. The simplest is a *steady state*, also called fixed point, equilibrium point or singular point of the dynamical system. As these different names try to suggest, such a

¹ Milnor (1985) uses “basin of attraction” only for open sets and the term “realm of attraction” for the general case.

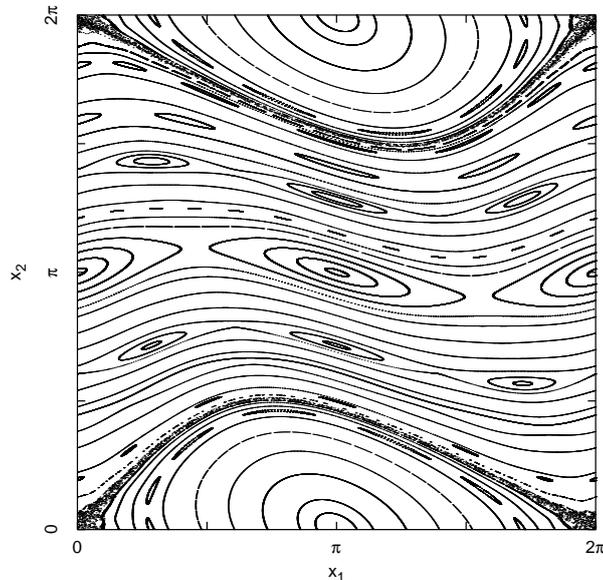


Figure 1.1: Dynamics of a two-dimensional map showing examples of different types of motion. There are fixed points at $(0,0)$ and $(\pi,0)$, a periodic orbit $(0,\pi) \leftrightarrow (\pi,\pi)$ and chaotic motion in a layer connecting $(0,0)$ with $(2\pi,0)$ and $(2\pi,0)$ with $(2\pi,2\pi)$. Since this map is 2π -periodic in both coordinates, this is actually one and the same layer. Quasiperiodic motion results in closed curves, some of which run more or less horizontally from edge to edge, while others form nested ellipses around the periodic orbits. Details about this map can be found in Sec. 2.2.2.

state does not change with time but remains the same for all times. In the case of *periodic motion*, the trajectory is no longer a single point, but a closed curve in state space. This curve is also called a limit cycle. Furthermore, there is the possibility of *quasiperiodic motion*, which can be thought of as more than one periodic motion (with incommensurate frequencies) influencing the system state. The last case is *chaotic motion*. It differs qualitatively from the types of motion mentioned so far and is the one we investigate. The reasons for this are firstly that it is of the highest practical importance and secondly that it is the most interesting from a theoretical point of view. Examples for these types of motion can be seen in Fig. 1.1.

In practice one is often faced with the problem of guessing the future state of a natural, real-world system given the current state. This guess at the future state is

called a *prediction*. The “guessing” is done by simulating a mathematical model of the real-world system on a computer for a desired amount of time. After this amount of real time has passed one can compare the prediction with the true state of the real-world system. This comparison will lead one to the conclusion that the prediction was good or bad or, to put it more scientifically, more or less accurate. The agreement between prediction and the observed true state of the system is known in meteorology as *accuracy*. The term *skill* (Gringorten, 1958; see also Ziehmann, 2001), sometimes used in the same sense as accuracy, is normally used to describe the ability of making better forecasts, i. e. predictions. Better here means better than a given reference with which the new method or model is compared. What this reference is depends on the situation at hand. It could be random chance, the persistence of the latest observation or a climatological average (Sanders, 1953), to name but a few examples.

When we speak of the *predictability* of a dynamical system we are talking about the influence of errors, e. g. measurement errors contained in the initial state of the model system. To judge the predictability of the system in a given state, one has to consider how these errors evolve over time. If they increase over time then the agreement of model system state and real-world state will deteriorate with time. On the other hand, if the initial errors decrease then the state of the model system and the state of the real-world system will actually come into better and better agreement. It is rather obvious that in the second case the prediction has a relatively high probability of agreeing with the final true state of the real-world system. In the first case, however, it could be considered a curious coincidence if the prediction is accurate, especially for a long prediction time.

In many dynamical systems making predictions is not a problem. If the system is in a steady state the prediction is trivial. For a system exhibiting periodic motion it is not difficult to predict the future state either, even for very long prediction times. The case of quasiperiodic motion is more involved, because one needs a longer stretch of time to get an idea of the typical behaviour of the system, but the prediction is not more complicated in principle. The difficulties arise when the trajectory is not so well-behaved any more. Since we do not consider stochastic systems this leads us to chaotic systems. These are of the highest importance because the natural systems that one is interested in predicting are typically chaotic.

The term *chaos* was first used in connection with dynamical systems by Li and Yorke (1975). Since then, various definitions have emerged. A broad consensus exists that the essential ingredient is sensitive dependence on initial conditions. One way of putting this more precisely is by requiring the largest global Lyapunov

exponent to be positive². We adopt this notion, which is also the prevailing one, and consider a system to be chaotic if it has a positive global Lyapunov exponent.

A widely known definition of chaotic maps given by Devaney (1992) also requires two more characteristics apart from sensitive dependence on initial conditions: The periodic points must form a dense subset and the system must be transitive. The last requirement is related to the mixing property and means that every open subset of state space will reach any other open subset given enough time. Actually the sensitive dependence on initial conditions follows from the other two requirements, as has been shown by Banks et al. (1992). On the other hand, none of the other two requirements imply the third (Assaf and Gadbois, 1992): Denseness of periodic points and sensitive dependence on initial conditions do not imply transitivity; transitivity and sensitive dependence on initial conditions do not imply that the periodic points are dense.

Auslander and Yorke (1980) define chaos as instability of all trajectories together with the existence of a dense orbit. This definition is compared to sensitivity to initial conditions in the context of topological dynamics by Akin and Kolyada (2003). Various concepts of sensitivity and chaos are compared in Kolyada (2004). There is also the notion of topological chaos with the hallmark property of a positive topological entropy (Adler et al., 1965). Furthermore, one can start with two points being mapped and study the statistical properties of distances between these points, arriving at distributional chaos (Schweizer and Smítal, 1994). This was recently compared to Devaney’s notion of chaos by Oprocha (2006).

Now that we have clarified the important terms used in this thesis let us move to the reasons why predictability is limited in chaotic systems and to the concepts used to describe this limitation.

1.3 Limitations of predictability

The problems arising in the prediction of chaotic systems can result from two very different underlying causes. The first are deficiencies in the model itself, the second are uncertainties about the true state of the real-world system at the time the simulation is started. We will deal with these two in turn.

The first cause, deficiencies in the model, is virtually always present when modelling a real-world system. It is a well-known fact that good models rely on a good trade-off between simplicity of the model and accuracy in representing all

²The concept of Lyapunov exponents will be introduced and discussed in detail in Sec. 2.1.1.

physical processes involved. Without the latter, the effect to be studied will not show up at all, the results will be unrealistic; without the first, one cannot identify the cause of the effect and the system may not even be treatable numerically, not to mention the possibility of analytical results.

The model in general has deficiencies of different kinds. Even in very simple cases like a pendulum or billiard balls colliding, one has to make simplifying assumptions in order to arrive at a model that is not too complicated. In the examples, the inclusion of friction and inelastic processes could in principle increase the accuracy of the model. However, the price for this is a complication of the model that can make it unsuitable for investigation. Fortunately, many interesting effects and typical traits of real-worlds systems can already be found in very simple models, so-called “toy models”.

Perhaps the most basic deficiency of a model is that sometimes the underlying processes of the physical system are not even understood well enough to be modelled properly. This incomplete knowledge leads to an incomplete or imprecise model which may well lack vital components and is therefore of limited use for the prediction of future states of the real-world system. This deficiency can range from wrong parameter values to wrong functional relationships to downright omission of important processes because they have not been discovered yet.

The influence of model errors can be studied by running a complicated, “perfect” model, consisting of two parts, and another model that only replicates one part of the perfect model, the other part is parametrised. One can then compare results of these two models and study the dependence on the parametrisation. A good presentation of this can be found in Smith (2000). Recently Knopf et al. (2005) investigated the difference in dynamics between forced and coupled dynamics of a simple atmosphere-ocean system. The forced system consists of an atmosphere part that is not coupled to the ocean part of this (i. e. the forced) system, but is driven by the ocean variables from an independent atmosphere-ocean system initialised in a slightly different state. The authors then determined mean locking times, that is the fraction of time the systems are synchronised (the state of the forced system is very close to the state of the forcing system). Upon varying the strength of the coupling between atmosphere and ocean, it was found that the forced system has a much more complex bifurcation diagram and phase space portrait than the forcing system alone, with the possibility of additional stable periodic orbits occurring. This causes “artificial bistability” in the forced system in some parameter ranges, making the occurrence of locking dependent on the initial state. This bistability is called “artificial” because it is not present in the fully coupled system. What is of considerable importance for forecasts is the fact that unstable periodic orbits, arising from periodic ones after a bifurcation, can lead to intermittent locking. In this case the forced system synchronises with the forcing system for a while, but then there are quite long stretches of time

during which the states seem to be unrelated. Such a situation of locking makes one assume rather good predictability, when in fact the closeness of trajectories is only caused by using a forced system instead of a fully coupled system.

Deficiencies of the model itself and their consequences are not the object of this study. Instead, we investigate the problem of predictability from the point of view of a “perfect model”. This means that we only deal with uncertainties in the initial conditions.

As we have noted, in chaotic systems trajectories originating from nearby points will diverge. This causes small initial deviations to grow exponentially over time, at least for some time. After a while the finite size of the system prohibits any further growth, the deviation then has the same order of magnitude as the size of the system. The averaged effective rate of the divergence of trajectories in the limit of small deviations and large time is measured by the Lyapunov exponents of the system (Oseledec, 1968; Eckmann and Ruelle, 1985). However, these are only of limited use for many applications as the time horizon of the prediction is usually much smaller than the global time scale of the system over which Lyapunov exponents are defined. Furthermore, the systems under investigation are typically very inhomogeneous. The divergence of trajectories can be fast or slow, depending on the initial state. Therefore local exponents based on the concept of Lyapunov exponents have been studied which govern the evolution of uncertainties about a particular initial condition for a finite time horizon (Grassberger et al., 1988; Abarbanel et al., 1991; Eckhardt and Yao, 1993; Smith et al., 1999). This local exponent is then associated with the initial system state and used as a measure of predictability. Large exponents represent fast divergence, hence bad predictability, whereas small values of local exponents stand for good predictability.

As long as the sign of the local exponent is positive, trajectories starting close to each other will diverge. They may diverge slowly if the value is close to zero, but they will diverge. If the local exponent is negative, however, such trajectories will actually converge, meaning that initial deviations will shrink with time. States for which all local exponents are negative are therefore very interesting for forecasting, since any initial uncertainty will become smaller, no matter which orientation the uncertainty had.

Since local exponents can be computed over times relevant for usual forecasts, they can be considered as good candidates for measures of predictability (Ziehmann et al., 2000). In view of this aim, one has to note that the method used in this work is more appropriate to describe a hindcast as opposed to a forecast. After all, the predictability is only determined at a later time, after the system has already evolved further. Thus, predictability can only be judged

in this way by examining the past. Keeping this in mind, one goal of this thesis is relating the predictability to different properties of the dynamical system. This can then provide a possibility for insights into predictability by studying the related property of a given system.

Apart from these local exponents, other measures have been introduced and studied to get a grasp of “how chaotic”, how unpredictable a given system is. Nicolis et al. (1983) define a non-uniformity factor (NUF) based on “local divergence rates,” by which they mean finite-time Lyapunov exponents (in the sense of this work) for short times. They then define the non-uniformity factor as the standard deviation of the local divergence rates along a trajectory. Perhaps not surprisingly, when studying the dependence of NUF on system parameters, they find features that do not correspond with features the Lyapunov exponents show. Extending this research, Nese (1989) studied the three-dimensional Lorenz system (Lorenz, 1963). It was found that in spite of the asymptotic divergence of trajectories, a substantial number of trajectory pairs did not diverge even after times as long as 20% of the typical time scale of the system. A coarse division of the attractor into regions according to their non-uniformity was also given.

In the meantime, regions of very good predictability have been found in various systems (Ziehmann et al., 2000). By this we mean regions where for every initial state an uncertainty will actually decrease, at least for short time. The reason for this is that there exist only some regions on the attractor where the divergence is large and others where divergence is small or where one can even find convergence of trajectories. In the long run the large divergence dominates, leading to a positive value of the largest global Lyapunov exponent.

Local exponents implicitly assume that the uncertainties are infinitesimally small. This enables one to arrive at analytical results that are valid for whole classes of systems. It also simplifies calculation, but is not necessarily a good approximation of the actual situation. Such methods cannot be expected to yield good results for finite errors of considerable size, a situation often occurring in practice. To deal with this problem further ideas were proposed that apply to uncertainties of finite size. One such measure, introduced by Aurell et al. (1997), is the so-called *finite size Lyapunov exponent* (FSLE). This is based on the idea of following an initial perturbation of size δ until it grows to size Δ , where both of these values can be chosen freely. When this happens the corresponding time $T(\delta, \Delta)$ is recorded and the perturbation is rescaled to size δ , leaving the direction unchanged. After repeating this many times, one calculates the FSLE as $\lambda(\delta, \Delta) = \langle T(\delta, \Delta)^{-1} \rangle \ln \frac{\delta}{\Delta}$, where $\langle \cdot \rangle$ denotes an average over the samples of $T(\delta, \Delta)$. This concept is useful in the study of transport processes, for example mixing processes in the ocean (d’Ovidio et al., 2004).

Operational weather forecasting models of today use methods involving the full nonlinear system. One approach, which we will also investigate in this work, is the study of ensembles: One chooses a number of states in the proximity of the current system state and follows the trajectories originating from all these states. For a given future time one can then not only give a forecast, i. e. the future state corresponding to the original “true” state of the system; one can also compare this forecast to the future states of the ensemble members. This allows one to judge how probable it is that a particular part of the forecast is correct, for example a hailstorm occurring.

When using ensembles in operational weather forecasting one has to deal with a high-dimensional system (10^7 variables is not unusual) and is therefore interested in minimising the additional computational cost caused by the ensemble. If the size of the ensemble cannot be increased at will, the choice of initial states of the ensemble members will become crucial. At the moment there are two operational weather forecast groups using ensembles.

The European Centre for Medium-Range Weather Forecasts (ECMWF) in Reading, UK, uses “optimal perturbations” that are chosen along the directions that grow fastest for short times, i. e. along singular vectors (Buizza et al., 1993; Palmer et al., 1994). These directions can be changed to arrive at even larger perturbation growth, using the fact that nonlinear interactions inhibit the growth of the fastest-growing singular vectors (Barkmeijer, 1997).

The Environmental Modeling Center (EMC) of the National Centers for Environmental Prediction (NCEP) in Camp Springs, USA, takes a different approach. There, one is interested in generating realistic perturbations that resemble the differences between the model state and the state of the real-world system. This is done by starting with an arbitrary perturbation to the state of the model system, integrating the model for a short time from both initial states (with and without perturbation), then subtracting the “true” state from the perturbed one and finally scaling this difference in states to the same size as the original perturbation. After a transient of up to a few days, the perturbations lie in the directions that grow fastest under the nonlinear system. This method is called “breeding” or “breeding of growing modes” (Toth and Kalnay, 1993).

A comparison of these bred vectors (BVs) and Lyapunov vectors (associated with global error growth) shows that there are similarities but also significant differences. In particular, bred vectors derived from different initial perturbations remain different and do not all collapse onto the direction of the leading Lyapunov vector. When starting with infinitesimal perturbations and continuing the process for a long time, one arrives at the leading Lyapunov vectors (Kalnay et al., 2002). In fact, Toth and Kalnay (1997) found the bred vectors to be superpositions of the leading Lyapunov vectors.

Using the bred vector technique Patil et al. (2001) measured the effective lo-

cal finite-time dimensionality of the Earth’s atmosphere as an example of a spatiotemporally chaotic system. The effective dimension of the subspace spanned by the local bred vectors is calculated via the singular values of the covariance matrix of the bred vectors (interpreted as columns of a matrix). It was found that spatial regions of a local dimensionality as low as 2.5 exist. This means that 2 to 3 of the local bred vectors alone account for the variability present in all bred vectors, and indicates that in such a region the local bred vectors effectively span a space of substantially lower dimension than that of the full space.

Many studies, including the work presented here, involve computer experiments with long trajectories. This is necessary because otherwise the statistics would not be good enough to draw conclusions. It should be noted that problems with the validity of long numerical trajectories due to roundoff errors are known to exist for some systems. The shadowing theorem (Anosov, 1967; Bowen, 1975) states that for a certain class of dynamical systems, there exists a true trajectory that stays arbitrarily close to a given numerical trajectory for an arbitrarily long time. This theorem is only proven to hold for hyperbolic systems, but is believed to apply to a wider class of systems. However, situations can arise in which the shadowability of the system under consideration can be drastically limited, meaning that a true trajectory closely follows a numerical one for only a rather short time. Such situations can occur when a Lyapunov exponent fluctuates around zero or the system exhibits an unstable dimension variability—effects expected to be common in high-dimensional systems (Dawson et al., 1994; Sauer et al., 1997). Thus the connection between predictability and shadowing in nonhyperbolic systems is still an intriguing problem which requires further studies to shed light on the predictability of high-dimensional systems.

1.4 Overview of this thesis

The goal of this work is to find suitable measures to characterise local properties and predictability of dynamical systems. In particular, we are interested in the local dynamics of uncertainties or errors, of initial deviations from the “true” state of the system. Since we assume a perfect model, the “true” state of the model system is in fact the true state, not an approximation of a real-world state.

Measures suitable for this task come in two fundamentally different kinds: measures for infinitesimal deviations from the true state and measures applicable for finite deviations. Both of these approaches were used to study predictability in chaotic dynamical systems. Furthermore, the measures used to study infinitesimal deviations proved useful for the characterisation of the dynamics in state space as well.

In order to study the growth of infinitesimal initial deviations we investigated two different kinds of local exponents, one of them being local Lyapunov exponents (which we will call “finite time Lyapunov exponents”), the other is based on singular values (“maximum growth exponents”). A detailed discussion of these measures is presented in Sec. 2.1.2.

A connection between smaller values of finite time exponents and the geometry of the stable direction was found by Tang and Boozer (1996) for conservative systems in the context of diffusive transport. Our approach yields a connection between predictability and homoclinic tangencies, which involve stable as well as unstable directions.

For the study of finite error growth we use ensembles initialised around the true state of the system. In contrast to the complicated, high-dimensional models used in operational weather forecasting, our models are rather simple and low-dimensional. This allows us to systematically study error growth in dependence of various parameters, including ensemble size.

The remaining part of this thesis is organised as follows:

Chapter 2 lays the foundation by presenting the theoretical concepts and methods that are common to all studies presented thereafter. Those methods or concepts that are specific to a particular topic will be introduced later, when it is appropriate. Also, the model systems studied are introduced in this chapter.

Chapter 3 presents the results published in Harle and Feudel (2007). This work is not directly related to predictability and could be viewed as an interesting by-product of developing and investigating the local measures for infinitesimal initial deviations. In particular, we investigate local exponents for trajectories inside the chaotic layer of a Hamiltonian system. Such trajectories are related to transport of passive tracers in a fluid. One very important aspect is the mixing of these tracers. In nature, it can account for better nutrient supply for organisms floating on the surface of the sea (Abraham et al., 2000), for example, or for transitions in reaction-diffusion systems (Neufeld et al., 2002). Motion in the chaotic layer of a Hamiltonian system is known to show stickiness when the trajectory comes close to an island of regular motion (Zaslavsky, 1998), spending a long time in the vicinity of these islands. We set out to characterise the motion in the chaotic layer by means of local exponents describing infinitesimal error growth. That is, we want to characterise the “closeness” of the trajectory to islands of regular motion with certain values of local exponents. By doing so we can explain the recent findings by Szezech et al. (2005) of a bimodal distribution of local exponents.

Chapter 4 deals with the question of relating regions of good (or bad) predictability, in the sense of slow (or fast) growth of infinitesimal errors, to other

properties of the dynamical system. This should prove advantageous, since one can hope to gain further insights by studying these related properties. One possibility is to identify particular regions in state space that stand out by providing starting positions for forecasts of either very good or very bad predictability. This information can then be used by checking whether the current system state lies within such a region. As a candidate for this kind of property of dynamical systems we investigated homoclinic tangencies, which will be introduced in this chapter. It can be argued that around a point of homoclinic tangency the predictability should be enhanced. The results of this work, which builds on Ziehmann et al. (2000) and Jaeger and Kantz (1997), are published in Harle and Feudel (2005).

Chapter 5 moves from infinitesimal errors to errors of finite size, whose growth is studied using ensembles. This is necessary when the initial error size is too large or the prediction time horizon is too long for the local exponents to yield an accurate description of error growth. In practice this is the common case, because even with a perfect model the initial state cannot be determined to infinite accuracy by measurements. For this reason operational weather forecasting of today, for example, uses methods involving the full nonlinear system (e. g. Buizza, 1996; Toth and Kalnay, 1993). When studying error growth using ensembles we want to get an idea of the worst-case scenario as well as the most favourable and the average situation with respect to predictability. The dependence of finite error growth on the parameters describing the initial ensemble states and on the prediction time horizon is studied. The intent is to establish systematic relationships, for example between the worst-case error growth and the prediction time. Such relationships may then be used to estimate times when predictability breaks down or sensible ensemble sizes. The results of chap. 5 are published in Harle et al. (2006).

Finally, in chap. 6 we summarise the results presented in chaps. 3–5 and comment on possible future extensions of our work.

1.4. *Overview of this thesis*

Chapter 2

Concepts, methods and systems under investigation

This chapter introduces the theoretical concepts as well as the methods that are common to the following chapters. Things specific to a particular topic will be introduced where appropriate.

Throughout this treatise, vectors are written with an arrow on top, like this: \vec{x} . The time dependence is expressed with parentheses: $\vec{x}(t)$, $\vec{x}(n)$. The case of discrete time (maps) is treated the same in order to avoid possible confusion with state variables: x_1 is the first component of vector \vec{x} , not the first iterate of some state x_0 .

2.1 Exponents as a measure for divergence of trajectories

2.1.1 Global Lyapunov exponents

The Lyapunov exponents are a measure of divergence of two trajectories that are originally infinitesimally close to each other (s. Fig. 2.1). One of the hallmarks of chaos is the fact that such a pair of trajectories diverges exponentially fast, i. e. the distance d between the trajectories typically grows with time t as

$$d(t) = d(0) \exp(\Lambda t) , \quad (2.1)$$

where Λ is called a Lyapunov exponent. In fact, this is the very property that is most commonly used to determine if a system exhibits chaos or not.

To be more precise, we consider a trajectory $\vec{x}(t)$ and a neighbouring trajectory $\vec{y}(t) = \vec{x}(t) + \vec{u}(t)$. The deviation of $\vec{y}(t)$ from $\vec{x}(t)$ is given by the vector $\vec{u}(t)$, the norm $\|\vec{u}(t)\|$ takes the role of $d(t)$ above. In general, the exponent in

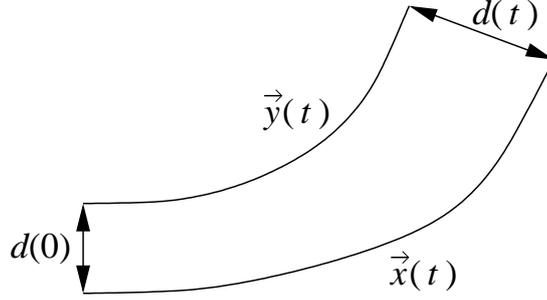


Figure 2.1: Exponential growth of initial deviation $d(0)$ between trajectories $\vec{x}(t)$ and $\vec{y}(t)$.

Eq. 2.1 can depend on both the initial system state, $\vec{x}(0)$, and on the orientation of the initial deviation, $\vec{u}(0)$. The Lyapunov exponent for a dynamical system given by a map as in Eq. 1.3 is then defined by

$$\Lambda(\vec{x}(0), \vec{u}(0)) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \frac{\|\vec{u}(n)\|}{\|\vec{u}(0)\|}. \quad (2.2)$$

Since the deviations \vec{u} are assumed to be infinitesimal, they evolve according to the linear system

$$\vec{u}(n+1) = \text{D} M(\vec{x}(n)) \cdot \vec{u}(n), \quad (2.3)$$

where $\text{D} M$ is the Jacobian of the map M . Using this, Eq. 2.2 becomes

$$\Lambda(\vec{x}(0), \vec{u}(0)) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \left| \text{D} M^n(\vec{x}(0)) \cdot \frac{\vec{u}(0)}{\|\vec{u}(0)\|} \right|, \quad (2.4)$$

where

$$\text{D} M^n(\vec{x}(0)) = \text{D} M(\vec{x}(n-1)) \cdot \text{D} M(\vec{x}(n-2)) \cdot \dots \cdot \text{D} M(\vec{x}(0)) \quad (2.5)$$

is the product of the Jacobians along the trajectory.

In the case of a dynamical system with continuous time (Eq. 1.2, p. 2), the definitions are analogous, with a matrix $P(\vec{x}(0), t)$ replacing $\text{D} M^n(\vec{x}(0))$. Thus, Eq. 2.4 becomes

$$\Lambda(\vec{x}(0), \vec{u}(0)) = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \left| P(\vec{x}(0), t) \cdot \frac{\vec{u}(0)}{\|\vec{u}(0)\|} \right|, \quad (2.6)$$

with $P(\vec{x}(0), t)$ defined as the solution of the initial value problem

$$\begin{aligned} \frac{dP(\vec{x}(0), t)}{dt} &= \text{D} \vec{F}(\vec{x}(t)) \cdot P(\vec{x}(0), t), \\ P(\vec{x}(0), t=0) &= 1. \end{aligned} \quad (2.7)$$

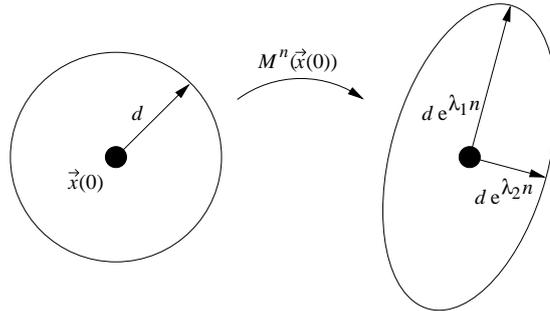


Figure 2.2: Dependence of exponential growth of infinitesimal deviations on direction.

For each $\vec{x}(0)$, the number of Lyapunov exponents is equal to the dimension d of the system, but their values do not have to be distinct. The actual value found by applying Eq. 2.4 depends on the orientation of $\vec{u}(0)$. A random orientation will result in the largest Lyapunov exponent $\Lambda_1(\vec{x}(0))$ ¹, since a randomly chosen orientation will have a component in the direction of maximum growth with probability 1. One can only encounter a smaller value than $\Lambda_1(\vec{x}(0))$ if the orientation happens to be perpendicular to the right singular vector of $DM^n(\vec{x}(0))$. A schematic picture of the situation for a two-dimensional system is shown in Fig. 2.2. This fact has to be taken into account when calculating Lyapunov exponents in practice. When one is interested in finding not only the leading Lyapunov exponent, one has to normalise the deviations $\vec{u}(t)$. This ensures that only one deviation $\vec{u}_1(t)$ grows with exponent $\Lambda_1(\vec{x}(0))$, the next with $\Lambda_2(\vec{x}(0))$ and so on.

It is known that under very general circumstances, the limits in Eq. 2.4 exist (Oseledec, 1968). Furthermore, if the system is ergodic, the set of values $\Lambda_i(\vec{x}(0), \vec{u}(0))$ is the same for almost all initial states $\vec{x}(0)$ with respect to the ergodic measure (see e. g. Eckmann and Ruelle, 1985). Therefore, we will drop the dependence on $\vec{x}(0)$ and $\vec{u}(0)$ from now on. The independence of $\vec{x}(0)$, or the limit of infinite time in Eq. 2.2, is the reason why the Λ_i are also called *global* Lyapunov exponents.

The value of Λ_1 describes the typical growth of a line segment at $\vec{x}(0)$ (except if the line happens to be perpendicular to the direction of maximum growth). When starting with two random vectors $\vec{u}^{(1)}(0)$ and $\vec{u}^{(2)}(0)$ defining a parallelogram, the probability that neither vector has a component in the directions associated with Λ_1 and Λ_2 is zero. Therefore, the area of this parallelogram grows as $\exp((\Lambda_1 + \Lambda_2)n)$. In general, the sum of the first k Lyapunov exponents describes the

¹By convention, $\Lambda_1 \geq \Lambda_2 \geq \dots \geq \Lambda_d$.

typical growth of a k -dimensional volume. When we set $k = d$, the dimension of state space, then the sum of all Lyapunov exponents describes the growth of a d -dimensional volume element. In the case of conservative systems there is no growth, so the sum of all Lyapunov exponents must be zero². In the case of conservative systems and dissipative systems with uniform contraction³ of volume elements, the values of the Lyapunov exponents have to satisfy constraints involving the right hand side of the dynamical system. For a time-continuous system \vec{F} (Eq. 1.2) this is given by

$$\sum_{i=1}^d \Lambda_i = \operatorname{div} \vec{F}. \quad (2.8)$$

For a map M (Eq. 1.4) the volume expands by $|\det D M|$ (so for conservative maps this is always 1). The determinant is the product of the eigenvalues, which are $\exp(\Lambda_i)$, so one arrives at

$$\sum_{i=1}^d \Lambda_i = \ln |\det D M|. \quad (2.9)$$

These equations have two practical applications. Firstly, one does not need to calculate all Lyapunov exponents, because one can be determined from the appropriate equation after $d - 1$ exponents have been calculated. Secondly, after calculating all d Lyapunov exponents one can use the above relationship to check the numerical values for consistency.

The global Lyapunov exponents have properties which make them a convenient tool from a theoretical point of view. Apart from the properties already mentioned, they are invariant under a smooth change of coordinates, for example. However, they are of limited value for practical purposes. The reason for this is that one is often not interested in the long-term behaviour of the system, but in short-term predictions.

One can generalise the idea behind global Lyapunov exponents in various ways to arrive at a measure of divergence of trajectories for finite times. Two particular such measures are introduced now and will be used later to quantify predictability.

²In fact, for Hamiltonian systems the Lyapunov exponents come in pairs with opposite sign (Benettin et al., 1980a).

³For systems with non-uniform contraction the sum of all Lyapunov exponents gives the average volume contraction rate along a trajectory.

2.1.2 Local exponents

Various local exponents can be defined to measure the divergence of nearby trajectories over a finite time τ . This enables one to associate a simple measure of predictability with every system state \vec{x} , since a large exponent means fast divergence of trajectories and, correspondingly, low predictability.

Other possible methods include local uncertainty q -pling times, for example, which measure the time it takes for neighbouring trajectories to increase their separation by a factor of q . Common choices are $q = 2$ and $q = e$. These times give information about when a chosen threshold is reached, something that cannot be extracted from local exponents. One should note that the inverse of such uncertainty q -pling times need not be related to local exponents, which measure the effective rate of error growth. A comparison of such times and effective growth rates was done by Smith et al. (1999). The authors prefer estimated times to growth rates, in particular since the latter require an *a priori* selection of a time scale. In the case of maps, which is the one we are studying here, time is discrete, leading to results for q -pling times that are difficult to interpret.

A different approach was suggested by Aurell et al. (1997). They focus on systems with many characteristic time scales as they are present in hydrodynamic turbulence. For each initial state, initial errors of (finite) size δ are followed until they reach a size of Δ after some time $T(\delta, \Delta)$. The *finite-size Lyapunov exponent* (FSLE) is then defined by $\lambda(\delta, \Delta) = \langle \frac{1}{T(\delta, \Delta)} \ln(\Delta/\delta) \rangle$, where $\langle \cdot \rangle$ denotes an average over the samples of $T(\delta, \Delta)$ for different initial error orientations. Through its dependence on the two parameters δ and Δ the FSLE allows one to choose the length scales one is interested in. For example, one can study the mixing up of fluid elements between these length scales (d'Ovidio et al., 2004).

The first kind of local exponents used herein is a straightforward generalization of the definition of the (global) Lyapunov exponents Λ_i in Eq. 2.4. It leads to the definition of *finite time Lyapunov exponents* (Grassberger et al., 1988):

$$\lambda_i^{(\tau)}(\vec{x}) = \frac{1}{\tau} \ln \left\| D M^\tau(\vec{x}) \vec{l}_i(\vec{x}) \right\|, \quad (2.10)$$

where $D M^\tau(\vec{x})$ is the linear propagator over a time τ along the piece of trajectory from $\vec{x} \equiv \vec{x}(0)$ to $\vec{x}(\tau - 1)$ (s. Eq. 2.5). Again, this is the definition for discrete time, to which we will restrict ourselves from now on (s. Sec. 2.2). The vectors $\vec{l}_i(\vec{x})$ (often called Lyapunov vectors) denote the orientations corresponding to the global values Λ_i at the point \vec{x} , i. e., almost every uncertainty $\vec{\epsilon}$ in the sufficiently distant past would evolve to lie in the direction $\vec{l}_1(\vec{x})$ once the trajectory reaches \vec{x} .

2.1. Exponents as a measure for divergence of trajectories

It should be noted that the ordering of the global Lyapunov exponents, $\Lambda_1 \geq \Lambda_2 \geq \dots \geq \Lambda_d$, need not carry over to the finite time exponents λ_i . This is so because $\vec{l}_1(\vec{x})$, for example, is the direction of maximum growth only in the limit of infinite time.

The other kind of local exponents we investigate is linked to the *singular values* $\sigma_i^{(\tau)}(\vec{x})$ of the product of Jacobians along the trajectory:

$$\rho_i^{(\tau)}(\vec{x}) = \frac{1}{\tau} \ln \left\| \mathrm{D} M^\tau(\vec{x}) \vec{v}_i^{(\tau)}(\vec{x}) \right\| = \frac{1}{\tau} \ln \sigma_i^{(\tau)}(\vec{x}). \quad (2.11)$$

The $\vec{v}_i^{(\tau)}(\vec{x})$ are the right singular vectors of $\mathrm{D} M^\tau(\vec{x})$ and by convention $\sigma_1^{(\tau)}(\vec{x}) \geq \sigma_2^{(\tau)}(\vec{x}) \geq \dots \geq \sigma_d^{(\tau)}(\vec{x})$, i. e., $\vec{v}_1^{(\tau)}$ is the direction which will have grown the most under the linearised dynamics after τ steps. Therefore, we refer to $\rho_i^{(\tau)}$ as *maximum growth exponents*.

The main difference between these two kinds of exponents can be formulated as follows: The $\lambda_i^{(\tau)}(\vec{x})$ in Eq. 2.10 depend on the history of the system through the vectors \vec{l}_i , while the $\rho_i^{(\tau)}(\vec{x})$ are defined by a finite piece of trajectory alone.

The terms used when referring to the various local exponents are not generally agreed upon yet. Depending on the scientific community or even the research group, the same name (e. g. finite time exponent) can stand for differently defined quantities, and the same quantities are sometimes named differently. This unfortunate state of affairs has to be kept in mind when skimming through the publications. For example, Lorenz (1965), Abarbanel et al. (1991) and Ziehmann et al. (1999) used the term “local Lyapunov exponents” for what is here called “maximum growth exponents”. On the other hand, “local Lyapunov exponents” was used for our “finite time exponents” by Nese (1989), Eckhardt and Yao (1993) and Toth and Kalnay (1993).

The finite time exponents are more common than the maximum growth exponents in theoretical studies of dynamical systems, as compared to studies dealing with practical applications like weather forecasting. As was already mentioned, their properties allow of analytical treatment leading to mathematical theorems. Besides, they are connected to the dynamics of the system due to their definition, which implicitly involves the history of the system in the form of the vectors $\vec{l}_i(\vec{x})$.

The maximum growth exponents are widely used to quantify predictability in atmospheric science (Palmer et al., 1994) and to determine low-dimensional approximations (Farrell and Ioannou, 2001). This practice is justified by the fact that observational data is repeatedly assimilated into the state of the system, i. e. the state is changed in order to better agree with the new measurements of the actual state of the physical system. Therefore, there is not enough time for perturbations to point along the globally most expanding direction $l_1(\vec{x})$, because

data assimilation can make the state of the system jump from one trajectory to another. Furthermore, the main interest lies in the estimation of the worst-case error, which is encountered when the initial error pointed in the direction of $\vec{v}_1^{(\tau)}$.

While the global Lyapunov exponents Λ_i are the same for almost all system states \vec{x} , local exponents typically depend strongly on \vec{x} . Different \vec{x} will give rise to different values of $\lambda_i^{(\tau)}(\vec{x})$ and $\rho_i^{(\tau)}(\vec{x})$, which enables one to link extreme values of local exponents to particular regions in state space. This leads to characteristic distributions of $\lambda_i^{(\tau)}(\vec{x})$ and $\rho_i^{(\tau)}(\vec{x})$ depending on the time horizon τ used. In the limit $\tau \rightarrow \infty$ both kinds of local exponents converge to Λ_i , the global Lyapunov exponents (Oseledec, 1968).

Due to the definitions of $\lambda_i^{(\tau)}(\vec{x})$ and $\rho_i^{(\tau)}(\vec{x})$, there exists a relation for the largest local exponents that holds for each state \vec{x} :

$$\rho_1^{(\tau)}(\vec{x}) \geq \lambda_1^{(\tau)}(\vec{x}), \quad (2.12)$$

because $\rho_1^{(\tau)}(\vec{x})$ involves the largest singular value, thus describing the maximum growth possible over time τ .

The mean of $\lambda_1^{(\tau)}(\vec{x})$ with respect to the natural measure of the system is always equal to the corresponding global Lyapunov exponent Λ_1 , irrespective of the value of τ . If this is not true for a numerical calculation, it is probably because the number of points used is too small. In the case of the maximum growth exponents one can show that

$$\langle \rho_1^{(2\tau)} \rangle \leq \langle \rho_1^{(\tau)} \rangle. \quad (2.13)$$

To see why this is so one can think of the product of 2τ Jacobians involved in calculating $\rho_1^{(2\tau)}(\vec{x}(0))$ (s. Eq. 2.11) as composed of two parts, one for the first τ iterations, and one for the second τ iterations:

$$D M^\tau(\vec{x}(0)) = D M^\tau(\vec{x}(\tau)) \cdot D M^\tau(\vec{x}(0)). \quad (2.14)$$

The maximum growth over the first τ iterations is given by the largest singular value of $D M^\tau(\vec{x}(0))$, $\sigma_1^{(\tau)}(\vec{x}(0))$. For the second τ iterations the maximum growth is $\sigma_1^{(\tau)}(\vec{x}(\tau))$. The combined maximum growth over 2τ iterations, $\sigma_1^{(2\tau)}(\vec{x}(0))$, can therefore be at most $\sigma_1^{(\tau)}(\vec{x}(\tau)) \cdot \sigma_1^{(\tau)}(\vec{x}(0))$. This is only the case, however, if the first left singular vector of $D M^\tau(\vec{x}(0))$ (giving the orientation of the error after τ iterations) is aligned with the first right singular vector of $D M^\tau(\vec{x}(\tau))$. In general these vectors will not be aligned, leading to $\sigma_1^{(2\tau)}(\vec{x}(0)) < \sigma_1^{(\tau)}(\vec{x}(\tau)) \cdot \sigma_1^{(\tau)}(\vec{x}(0))$.

This translates to the language of maximum growth exponents ρ as

$$\begin{aligned}\rho_1^{(2\tau)}(\vec{x}(0)) &= \frac{1}{2\tau} \ln \sigma_1^{(2\tau)}(\vec{x}(0)) \\ &\leq \frac{1}{2\tau} \ln \left(\sigma_1^{(\tau)}(\vec{x}(\tau)) \sigma_1^{(\tau)}(\vec{x}(0)) \right) \\ &= \frac{1}{2} \left(\rho_1^{(\tau)}(\vec{x}(\tau)) + \rho_1^{(\tau)}(\vec{x}(0)) \right),\end{aligned}\tag{2.15}$$

the largest maximum growth exponent over 2τ iterations being less than or equal to the average of the largest exponents of the τ -iteration parts. Since this is true for all initial states $\vec{x}(0)$, the mean of the largest maximum growth exponent will not increase when τ is increased by a factor of two and inequality 2.13 follows. An analogous argument can be made for the smallest maximum growth exponent showing

$$\langle \rho_d^{(2\tau)} \rangle \geq \langle \rho_d^{(\tau)} \rangle.\tag{2.16}$$

These properties of local exponents we have just discussed can be used to validate numerical results, as has been done for the results presented in the following chapters.

An implication of the inequality 2.13 is that the mean of the distribution of $\rho_1^{(\tau)}$ is not a good estimate of the global Lyapunov exponent Λ_1 , because $\langle \rho_1^{(\tau)} \rangle \geq \Lambda_1$ for all τ .

Henceforth, we will simplify notation by omitting the explicit dependence on \vec{x} of the local exponents when it is clear from the context which system state the local exponent pertains to. It should also be kept in mind that the value of τ determines the time horizon over which the evolution of a deviation is studied.

2.1.3 Numerical calculation

The analytic calculation of (global) Lyapunov exponents is only possible in very simple cases like the generalised baker's map (Farmer et al., 1983; see also Ott, 1993). In practice, one has to determine the Lyapunov exponents – as well as the local exponents – numerically.

To this end, one starts at a given system state $\vec{x}(t = 0)$ with some random deviation vector $\vec{u}(t = 0)$ of unit length. While the system evolves according to the map M , the deviation is mapped with the linear propagator as described by Eq. 2.3. After some time, $\vec{u}(t)$ will point in the direction $l_1(\vec{x}(t))$, the direction of globally fastest growing errors. To avoid numerical overflow, the vector $\vec{u}(t)$ is renormalised to unit length after some time interval T , storing the amount of growth up to that point as u_1 . After the next interval T we get u_2 and so forth.

In the end, the logarithms of the amounts of growth can be summed and divided by the evolved time, to arrive at an estimate of Λ_1 (cf. Eq. 2.4):

$$\Lambda_1 \approx \frac{1}{nT} \sum_{j=1}^n \ln u_j. \quad (2.17)$$

The procedure can be stopped when the result has converged to within a given tolerance.

In order to calculate more than one Lyapunov exponent, one has to use more deviation vectors \vec{u}_1, \vec{u}_2 and so on, possibly up to \vec{u}_d . Because of the domination of the largest Lyapunov exponent, these vectors have to be orthogonalised along with the normalization to unit length. Instead of storing the length u_j of the vectors \vec{u} before normalization after time jT , one now stores the k -dimensional volumes $V_j^{(k)}$ of the parallelepiped spanned by the first k vectors $\vec{u}_1, \dots, \vec{u}_k$. Since this volume grows with an exponent that is the sum of the k largest Lyapunov exponents (Shimada and Nagashima, 1979), one has

$$\sum_{j=1}^k \Lambda_j \approx \frac{1}{nT} \sum_{j=1}^n \ln V_j^{(k)}. \quad (2.18)$$

The k -th Lyapunov exponent can now be calculated from the volumes in k and $k - 1$ dimensions:

$$\begin{aligned} \Lambda_k &= \sum_{j=1}^k \Lambda_j - \sum_{j=1}^{k-1} \Lambda_j \\ &\approx \frac{1}{nT} \sum_{j=1}^n \ln V_j^{(k)} - \frac{1}{nT} \sum_{j=1}^n \ln V_j^{(k-1)} \\ &= \frac{1}{nT} \sum_{j=1}^n \left(\ln V_j^{(k)} - \ln V_j^{(k-1)} \right) \\ &= \frac{1}{nT} \sum_{j=1}^n \ln \frac{V_j^{(k)}}{V_j^{(k-1)}}. \end{aligned} \quad (2.19)$$

More on this method can be found in Benettin et al. (1980b), others are described in Eckmann and Ruelle (1985) and Greene and Kim (1987).

The calculation of finite time Lyapunov exponents is along the same lines as that of the global Lyapunov exponents. One uses Eq. 2.19 for the chosen

2.1. Exponents as a measure for divergence of trajectories

time horizon τ , i. e. one sets n and T such that $\tau = nT$ and calculates

$$\lambda_k^{(\tau)}(\vec{x}) = \frac{1}{\tau} \sum_{j=1}^n \ln \frac{V_j^{(k)}}{V_j^{(k-1)}} \quad (2.20)$$

for $k \geq 2$. The first finite time exponent $\lambda_1^{(\tau)}$ is determined by the amounts of growth of perturbations in direction \vec{l}_1 , like in Eq. 2.17.

The calculation of the maximum growth exponents $\rho_i^{(\tau)}$ becomes increasingly difficult when the time horizon τ is large. The reason for this lies in the determination of the singular values of the product of τ Jacobians. The eigenvalues of each Jacobian are of the order $\exp(\Lambda_i)$, resulting in the product of the matrices being ill-conditioned. A way to overcome this problem even for very large τ is an algorithm presented in Eckmann et al. (1986). It uses a recursive QR -decomposition of the product of Jacobians and works as follows.

The starting point for the calculation of the maximum growth exponents is the matrix

$$O(\vec{x}, \tau) = D M^\tau(\vec{x}) (D M^\tau(\vec{x}))^\top, \quad (2.21)$$

that consists of a product of the τ Jacobians along the trajectory followed by the transpose of this product, so 2τ matrices altogether. It is known from the multiplicative ergodic theorem (Oseledec, 1968) that

$$\lim_{\tau \rightarrow \infty} \left(D M^\tau(\vec{x}) (D M^\tau(\vec{x}))^\top \right)^{\frac{1}{2\tau}} \quad (2.22)$$

exists and has eigenvalues $e^{\Lambda_1}, e^{\Lambda_2}, \dots, e^{\Lambda_d}$ for almost all \vec{x} . In the case of a finite piece of trajectory, the eigenvalues of $D M^\tau(\vec{x}) (D M^\tau(\vec{x}))^\top$ are interpreted as $\exp(2\rho_1^{(\tau)}(\vec{x})\tau), \exp(2\rho_2^{(\tau)}(\vec{x})\tau), \dots, \exp(2\rho_d^{(\tau)}(\vec{x})\tau)$.

The actual recursive QR -decomposition for a trajectory⁴ from $\vec{x}(t = 1)$ to $\vec{x}(\tau + 1)$ defines $Q(j), R(j)$ as

$$D M(\vec{x}(j)) \cdot Q(j - 1) = Q(j) \cdot R(j), \quad j \in \{1, \dots, \tau\}, \quad (2.23)$$

with the matrix $Q(0)$ set to the identity matrix. How this leads to a simplification can be seen by considering the first two steps:

$$D M(\vec{x}(1)) = Q(1) \cdot R(1), \quad (2.24)$$

$$D M(\vec{x}(2)) \cdot Q(1) = Q(2) \cdot R(2). \quad (2.25)$$

⁴ We let the trajectory start at $t = 1$ to have corresponding indices for $\vec{x}(j), Q(j)$ and $R(j)$.

The product of the first two Jacobians now simplifies to

$$\begin{aligned} \mathrm{D} M(\vec{x}(2)) \cdot \mathrm{D} M(\vec{x}(1)) &= \mathrm{D} M(\vec{x}(2)) \cdot Q(1) \cdot R(1) \\ &= Q(2) \cdot R(2) \cdot R(1). \end{aligned} \quad (2.26)$$

Carrying on in this manner, one can express $\mathrm{D} M^\tau(\vec{x}(1))$ as⁵

$$\mathrm{D} M^\tau(\vec{x}(1)) = \mathrm{D} M(\vec{x}(\tau)) \cdots \mathrm{D} M(\vec{x}(1)) = Q(\tau) \cdot \prod_{j=1}^{\tau} R(j), \quad (2.27)$$

In the same way one can treat the 2τ Jacobian matrices forming $O(\vec{x}, \tau)$ in Eq. 2.21, ending up with one orthogonal matrix $Q(2\tau)$ and 2τ upper right triangular matrices $R(j)$, none of which has elements much larger than e^{Λ_1} .

Ideally, if the matrix $Q(2\tau)$ were the identity matrix, the eigenvalues of $O(\vec{x}, \tau)$ could be calculated easily from the diagonal elements of the matrices $R(j)$. In general, this is not the case, but there is a way to ensure this by successive similarity transformations (which are known not to change the eigenvalues). Since this will be done recursively, we add a subscript k to the matrices Q and R . The values for the first iteration, $k = 1$, come from the QR -decomposition of the 2τ Jacobian matrices forming $O(\vec{x}, \tau)$. For clarity we repeat Eq. 2.23 with the new subscript:

$$\mathrm{D} M(\vec{x}(j)) \cdot Q_1(j-1) = Q_1(j) \cdot R_1(j), \quad j \in \{1, \dots, 2\tau\}, \quad (2.28)$$

with $Q_1(0)$ being the identity matrix. The matrix $O(\vec{x}, \tau)$ can be written analogously to Eq. 2.27 as

$$O(\vec{x}, \tau) = Q_1(2\tau) \cdot \prod_{j=1}^{2\tau} R_1(j). \quad (2.29)$$

To get to the second iteration, the 2τ matrices $R_1(j)$ take the role of the Jacobians $\mathrm{D} M(\vec{x}(j))$ in Eq. 2.28, and $Q_1(2\tau)$ serves as the initial orthogonal matrix $Q_2(0)$. Thus, Eqs. 2.24–2.25 describing the first two steps now become

$$R_1(1) \cdot Q_1(2\tau) = Q_2(1) \cdot R_2(1), \quad (2.30)$$

$$R_1(2) \cdot Q_2(1) = Q_2(2) \cdot R_2(2). \quad (2.31)$$

The general equations for the recursive QR -decomposition are given by

$$R_k(j+1) \cdot Q_{k+1}(j) = Q_{k+1}(j+1) \cdot R_{k+1}(j+1), \quad (2.32)$$

$$Q_{k+1}(1) = Q_k(2\tau), \quad j \in \{1, \dots, 2\tau\}. \quad (2.33)$$

⁵ Throughout this thesis products like that of the matrices $R(j)$ are understood to be read right-to-left as the product index goes from lower to upper bound: $\prod_{j=1}^{\tau} R(j) = R(\tau) \cdots R(1)$.

2.1. Exponents as a measure for divergence of trajectories

In order to use this recursion we do a similarity transformation of $O(\vec{x}, \tau)$ in the form of Eq. 2.29 with $Q_1(2\tau)$:

$$\begin{aligned} Q_1(2\tau)^\top \cdot O(\vec{x}, \tau) \cdot Q_1(2\tau) &= Q_1(2\tau)^\top \cdot Q_1(2\tau) \cdot \prod_{j=1}^{2\tau} R_1(j) \cdot Q_1(2\tau) \\ &= \prod_{j=1}^{2\tau} R_1(j) \cdot Q_1(2\tau), \end{aligned} \quad (2.34)$$

where we have used the fact that $Q_1(2\tau)^\top \cdot Q_1(2\tau) = 1$. We can use the matrices on the right hand side of the last equals sign as input for the next QR -decomposition as described above (Eqs. 2.32–2.33):

$$\begin{aligned} \prod_{j=1}^{2\tau} R_1(j) \cdot Q_1(2\tau) &= \prod_{j=2}^{2\tau} R_1(j) \cdot R_1(1) \cdot Q_1(2\tau) \\ &= \prod_{j=2}^{2\tau} R_1(j) \cdot Q_2(1) \cdot R_2(1) \\ &= \prod_{j=3}^{2\tau} R_1(j) \cdot R_1(2) \cdot Q_2(1) \cdot R_2(1) \\ &= \prod_{j=3}^{2\tau} R_1(j) \cdot Q_2(2) \cdot R_2(2) \cdot R_2(1) \\ &= \dots \\ &= R_1(2\tau) \cdot Q_2(2\tau - 1) \cdot R_2(2\tau - 1) \cdots R_2(1) \\ &= Q_2(2\tau) \cdot R_2(2\tau) \cdot R_2(2\tau - 1) \cdots R_2(1) \\ &= Q_2(2\tau) \cdot \prod_{j=1}^{2\tau} R_2(j). \end{aligned} \quad (2.35)$$

Combining the previous equation with Eq. 2.34 one can see that

$$Q_1(2\tau)^\top \cdot O(\vec{x}, \tau) \cdot Q_1(2\tau) = Q_2(2\tau) \cdot \prod_{j=1}^{2\tau} R_2(j), \quad (2.36)$$

the right hand side of which is almost the same as that of Eq. 2.29, except for the increase of the indices of $Q(2\tau)$ and $R(j)$. The next similarity transform with $Q_2(2\tau)$ gives

$$Q_2(2\tau)^\top \cdot Q_1(2\tau)^\top \cdot O(\vec{x}, \tau) \cdot Q_1(2\tau) \cdot Q_2(2\tau) = \prod_{j=1}^{2\tau} R_2(j) \cdot Q_2(2\tau). \quad (2.37)$$

Therefore, after k iterations the recursive QR -decomposition yields

$$\left(\prod_{j=k}^1 Q_j(2\tau) \right)^T \cdot O(\vec{x}, \tau) \cdot \prod_{j=k}^1 Q_j(2\tau) = \prod_{j=1}^{2\tau} R_k(j) \cdot Q_k(2\tau). \quad (2.38)$$

With increasing iteration number k , the matrix $Q_k(2\tau)$ converges quickly to the identity matrix (Stoer and Burlisch, 1980). Suppose that this has been achieved within a given tolerance after K iterations. The i -th maximum growth exponent $\rho_i^{(\tau)}$ can then be computed from the i -th diagonal elements $(R_K(j))_{ii}$ of the upper triangular matrices $R_K(j)$:

$$\rho_i^{(\tau)} = \frac{1}{2\tau} \sum_{j=1}^{2\tau} \ln (R_K(j))_{ii}. \quad (2.39)$$

In practice, no convergence problems with this algorithm were encountered. The orthogonal matrix Q was considered converged when the sum of the absolute values of the elements of $Q - 1$ was $< 10^{-6}$.

2.2 Model systems studied

The investigations conducted involve simple paradigmatic models exhibiting chaotic behaviour. One system is a conservative, Hamiltonian system, the others are dissipative systems. In all cases we use discrete maps instead of time-continuous flows. The latter are more common in practice, whereas the former are more easily manageable. The restriction to discretised time is not severe for the following reason: d -dimensional maps that are invertible can be understood as Poincaré sections of $d + 1$ -dimensional flows. Thus, general properties of chaotic systems can be studied by using the more easily manageable maps. It is emphasized, however, that calculations regarding local exponents or ensemble statistics can be done analogously for time-continuous systems.

The example systems used have long served as convenient models to investigate typical chaotic behaviour in dissipative and conservative systems. This is one reason for using them in our investigations. Another reason is of a more practical nature: The dynamical systems to be studied have to be simple enough to remain feasible even for long prediction times and possibly also for large ensemble sizes. Based on these premises, we chose the following two-dimensional chaotic maps.

2.2.1 Hénon map

One of the standard systems representative of dissipative chaotic motion is the well-known Hénon map (Hénon, 1976)

$$\begin{aligned}x_1(n+1) &= a - x_1(n)^2 + bx_2(n), \\x_2(n+1) &= x_1(n).\end{aligned}\tag{2.40}$$

It was introduced to simplify the numerical exploration of the now famous three-dimensional model of Lorenz (1963). One can think of this map as being composed of three simple maps, the first folding a given region of the plane, the second contracting the result of the first map, and the third swapping coordinates. The Hénon map reproduces characteristic features of the Lorenz system, like the constant negative divergence and the actions of stretching and folding. Also, the uniqueness of the three-dimensional trajectory is reflected in the fact that the map is invertible. Therefore, one can say that the Hénon map is “close” to the Poincaré map of the Lorenz system: It is not exact, but the features should be qualitatively similar.

The Jacobian of this map, $DM = -b$, is always constant, the parameter b can be considered as a dissipation. The usual values for b are $0 < b < 1$, we use $b = 0.3$. The parameter a is typically varied when studying bifurcation phenomena, we keep it fixed at $a = 1.4$. For the parameter values considered herein, the Hénon system exhibits the well-known chaotic attractor, shown in Fig. 2.3, with a global Lyapunov exponent of $\Lambda_1 \approx 0.42$.

2.2.2 Standard map

The class of conservative chaotic systems is represented by the standard map. It is well studied in the context of Hamiltonian chaos (Lichtenberg and Leiberman, 1992; Chirikov, 1979; Lichtenberg et al., 1980):

$$\begin{aligned}x_1(n+1) &= x_1(n) + x_2(n) \pmod{2\pi}, \\x_2(n+1) &= x_2(n) + K \sin(x_1(n) + x_2(n)) \pmod{2\pi}.\end{aligned}\tag{2.41}$$

Here the Jacobian is readily verified to be equal to 1, irrespective of the state of the system. This map derives from a model where a bar is fastened frictionlessly at one end and periodically kicked at the other; gravitational forces are neglected. Due to the resulting rotational motion around the pivot, this map is also called “kicked rotor map” (Zaslavsky, 1998). The variable x from Eq. 2.41 corresponds to the angle between the bar and some axis, while y denotes the angular momentum⁶. The kicks are exerted parallel to the axis determining the

⁶ The angular momentum is also taken modulo 2π for convenience as plots of system states are 2π -periodic in both x and y .

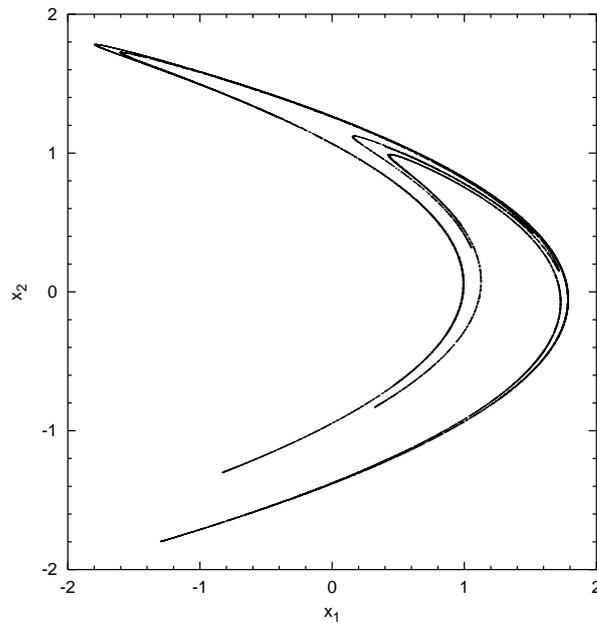


Figure 2.3: Attractor of the Hénon map for parameter values $a = 1.4$, $b = 0.3$.

value of x and have a strength of K . In order to arrive at a map, the original differential equations are integrated between two successive kicks. The values of the variables are examined immediately after each kick.

The dynamical behaviour of this system for one of the parameter values used ($K = 0.6$) is quite intricate (see Fig. 2.4) and consists of complexly intertwined regions of regular (periodic and quasiperiodic) and chaotic motion. If the initial condition is on an invariant KAM torus the dynamics traces out a closed curve corresponding to quasiperiodic motion. In Fig. 2.4, some KAM tori can be clearly seen running more or less horizontally from edge to edge. Apart from these, other tori are present, which are nested around periodic orbits like $(\pi, 0)$ and $(0, \pi) \leftrightarrow (\pi, \pi)$. These are the primary islands, around which there are secondary, tertiary islands and so on (see e. g. Lichtenberg and Lieberman, 1992). Such a hierarchical structure is seen in Fig. 2.4. Another possible dynamics is chaotic motion, wandering around an area and eventually densely filling that area. This behaviour is indicated by the regions near the corners in Fig. 2.4. These regions are called chaotic or stochastic layers (Zaslavsky et al., 1993). Overall, motion in this layer is weakly chaotic for $K = 0.6$, the global Lyapunov exponent is $\Lambda_1 \approx 0.07$. However, the global Lyapunov exponent calculated from one piece of trajectory, even a rather long one, need not result in this value due to stickiness near islands of regular motion. This is investigated in detail in the following

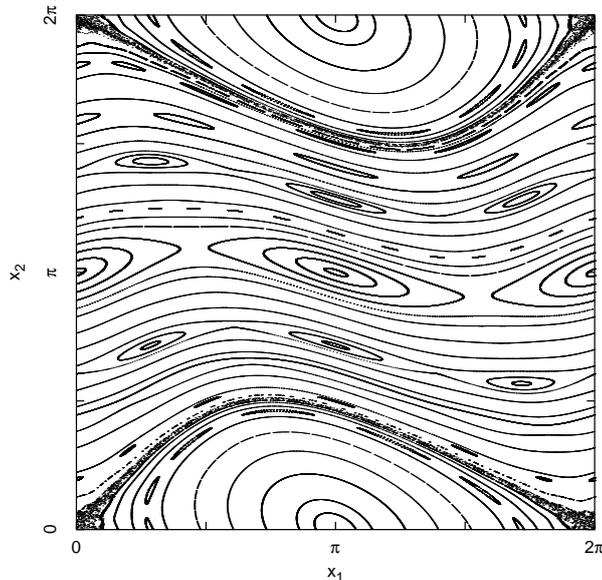


Figure 2.4: Dynamics of the standard map for $K = 0.6$ (47 initial positions).

chapter. It is important to note that not only the KAM curves themselves are invariant under the map, but also the area enclosed by any two of them.

For comparison, Fig 2.5 shows the dynamics at a relatively high value of $K = 4.2$. Now the chaotic layer occupies almost the whole state space, only a few islands are left. The motion is much more chaotic now with $\Lambda_1 \approx 0.85$.

For this system, we can restrict our study of local exponents to the larger one. The reason for this is that for Hamiltonian systems, $\det J(\vec{x}) = 1$, so that the sum over all exponents has to vanish at each point \vec{x} and the largest local exponent alone already contains all information.

2.2.3 Quasiperiodicity map

The quasiperiodicity map, defined in Grebogi et al. (1985), is more involved than the standard map and has the form

$$\begin{aligned} x_1(t_{n+1}) &= x_1(t_n) + c_1 + c_3 p_1(x(t_n)) \pmod{1}, \\ x_2(t_{n+1}) &= x_2(t_n) + c_2 + c_3 p_2(x(t_n)) \pmod{1}, \end{aligned} \tag{2.42}$$

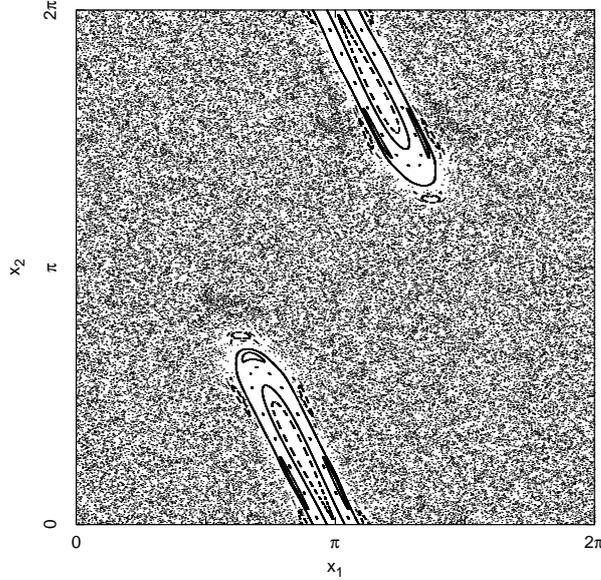


Figure 2.5: Dynamics of the standard map for $K = 4.2$ (19 initial positions).

where we used two sets of parameter values: $c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$ and $c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$. The nonlinearities p_1 and p_2 are

$$\begin{aligned}
 p_1(x(t_n)) &= \frac{1}{2\pi} \left(A_1 \sin(2\pi(x_1(t_n) + K_1)) + \right. \\
 &\quad A_2 \sin(2\pi(x_2(t_n) + K_2)) + \\
 &\quad A_3 \sin(2\pi(x_1(t_n) + x_2(t_n) + K_3)) + \\
 &\quad \left. A_4 \sin(2\pi(x_1(t_n) - x_2(t_n) + K_4)) \right), \\
 p_2(x(t_n)) &= \frac{1}{2\pi} \left(B_1 \sin(2\pi(x_1(t_n) + J_1)) + \right. \\
 &\quad B_2 \sin(2\pi(x_2(t_n) + J_2)) + \\
 &\quad B_3 \sin(2\pi(x_1(t_n) + x_2(t_n) + J_3)) + \\
 &\quad \left. B_4 \sin(2\pi(x_1(t_n) - x_2(t_n) + J_4)) \right),
 \end{aligned} \tag{2.43}$$

the values of the constants A_i , K_i , B_i , J_i are shown in Tab. 2.1. Since the parameter c_3 controls the strength of the nonlinear effects, the second set of parameter values is more chaotic ($\Lambda_1 \approx 0.377$) than the first ($\Lambda_1 \approx 0.0011$). At these parameter values, there exists a chaotic layer covering the whole state space of the system. An overview of the dynamics of the quasiperiodicity map at the

2.2. Model systems studied

i	A_i	K_i
1	-0.2681366365	0.9854608430
2	-0.9106755940	0.5044604561
3	0.3117202638	0.9470747252
4	-0.0400397884	0.2335010550

i	B_i	J_i
1	0.0881861167	0.9903072286
2	-0.5650288998	0.3363069701
3	0.1629954873	0.2980492123
4	-0.8039888198	0.1550646728

Table 2.1: Parameter values used for the quasiperiodicity map.

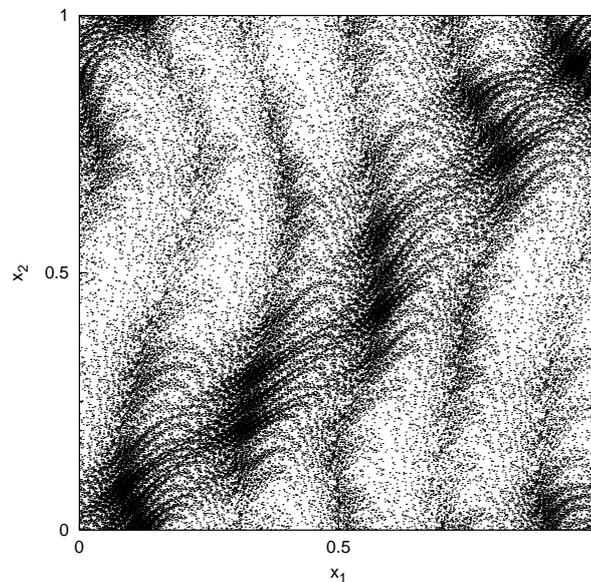


Figure 2.6: Dynamics of the quasiperiodicity map for $c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$.

parameter values used is given in Figs. 2.6 and 2.7. The fact that there is an extended chaotic layer, along with the structure of the quasiperiodicity map that is substantially more complex than the one of the standard map, made us choose this system for our investigation. Another reason for comparing these systems is the very different shape of their distributions of local exponents, as will be shown in Sec. 5.2. Comparing systems that differ so markedly in characteristic

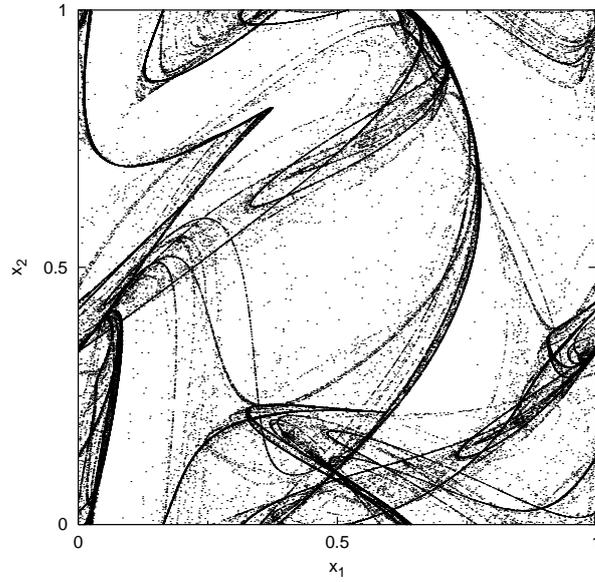


Figure 2.7: Dynamics of the quasiperiodicity map for $c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$.

properties allows one to draw more general conclusions than considering only one possibly special case.

2.2. *Model systems studied*

Chapter 3

Local exponents in chaotic layers of Hamiltonian systems

In this chapter, we investigate the properties of local exponents for trajectories inside the chaotic layer of a Hamiltonian system. Such motion is particularly interesting from the point of view of transport of passive tracers in a fluid. These are particles that do not interact with the fluid, so the flow remains unperturbed by adding passive tracer particles. One very important aspect is the mixing of these tracers. In nature, it can account for better nutrient supply for organisms floating on the surface of the sea (Abraham et al., 2000), for example, or for transitions in reaction-diffusion systems (Neufeld et al., 2002). Area-preserving two-dimensional maps are relevant for the motion of passive tracers in a two-dimensional time-periodic flow. The reason for this is related to the (also time-periodic) stream function $\psi(x_1, x_2, t)$ of the flow. The velocities of passive tracers in such a flow are given by

$$v_1 = \frac{\partial\psi}{\partial x_2}, \quad v_2 = -\frac{\partial\psi}{\partial x_1}. \quad (3.1)$$

This is structurally equivalent to a Hamiltonian system, with the stream function replacing the Hamilton function. If one now takes a Poincaré section by looking at the system in intervals of the period of the flow, one arrives at a two-dimensional Hamiltonian map.

Motion around islands in the chaotic layer is characterised by stickiness, i. e. once the trajectory comes close to an island of regular motion it will stay in its vicinity¹ for some time. This stickiness leads to particular statistical properties like anomalous diffusion and non-Gaussian transport in fluids (Zaslavsky et al.,

¹ More precisely, the trajectory will stay in the vicinity of islands of the same chain. The structure of islands and island chains will be discussed in the following section.

(1993); Leoncini and Zaslavsky, (2002)). Diffusion is called anomalous if the transport of particles is such that the second moment of the particle displacement R is

$$\langle R^2 \rangle \sim t^\gamma, t \rightarrow \infty \quad (3.2)$$

with an exponent that differs from the value of $\gamma = 1$ for the case of normal diffusion. This limits transport ($0 \leq \gamma < 1$, subdiffusion) or enhances it ($1 < \gamma \leq 2$, superdiffusion).

In the following sections, the dynamics of the system under consideration, the standard map, is described with the focus on island chains. A trajectory spends relatively more time near islands if the chaotic layer is rather small. Therefore we use a parameter value of $K = 0.6$, which gives us a better chance of finding stretches of a trajectory close to islands. One initial condition is iterated for a very long time in order to have acceptable statistics. Comparisons with shorter trajectories starting from different initial positions (in order to reach all parts of the chaotic layer) give no qualitatively different results.

As a simple way to get a first impression, the spectrum of local exponents is investigated in Sec. 3.2. Then, distributions of local exponents are presented in Sec. 3.3.

Throughout this chapter, we can restrict ourselves to the first of the two local exponents. The second one is always given by

$$\lambda_2(x_1, x_2) = -\lambda_1(x_1, x_2) \quad \text{and} \quad (3.3)$$

$$\rho_2(x_1, x_2) = -\rho_1(x_1, x_2) \quad (3.4)$$

because the standard map is area preserving. Thus, the distributions of λ_2 and ρ_2 are simply mirror images of those of the other exponents and contain no new information.

3.1 Island structure in chaotic layers

As has already been mentioned in Sec. 2.2.2, chaotic Hamiltonian systems can exhibit chaotic layers with intricate dynamics. If the nonlinearity is large enough the motion is no longer confined to regular dynamics. Starting from a suitable initial state, the trajectory can wander around in a part of the state space, but it is restricted to an area (in the case of two dimensions) between neighbouring KAM-curves. The reason for this is that the KAM-curves as well as the areas between any two such curves are invariant under the dynamics of the system. The area of this chaotic layer that is available for the trajectory increases with the size of the nonlinearity (s. Fig. 3.1).

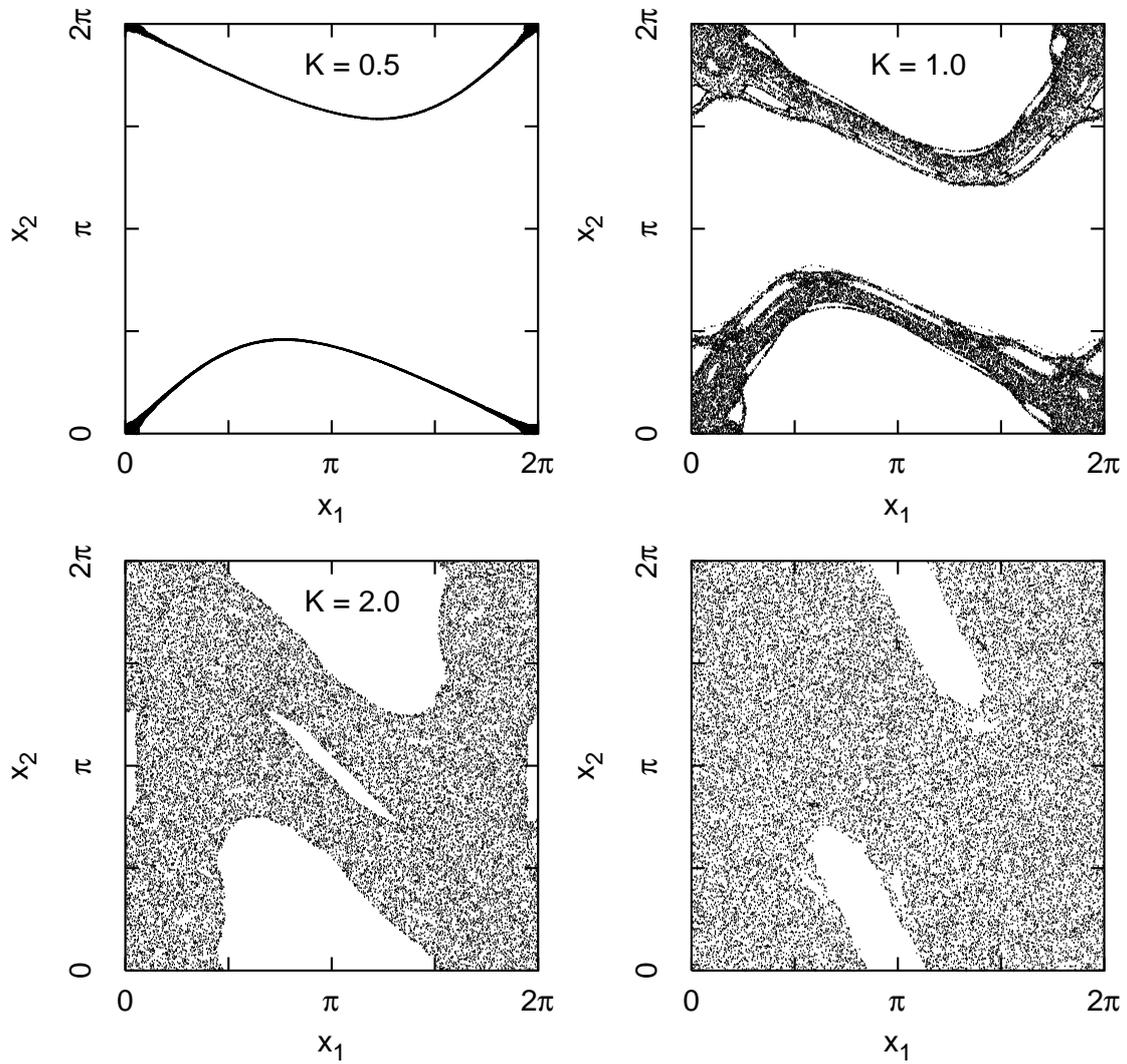


Figure 3.1: Dynamics of the standard map for various nonlinearity parameters: The area of the chaotic layer increases with K , the bottom right plot is for $K = 4.0$. Each plot shows one trajectory initialised at $(0.1, 0.1)$, inside the chaotic layer, iterated 25000 times.

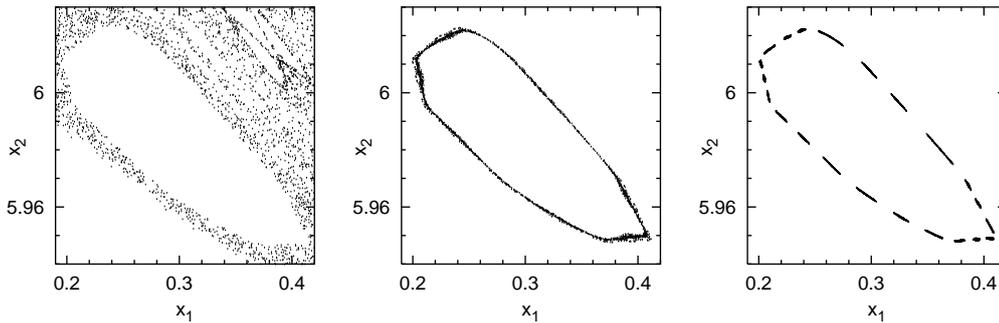


Figure 3.2: Dynamics around islands of different order. Left: order 1; middle: order 2; right: order 3.

It can be seen from the plots for higher values of K that islands exist within the chaotic layer. Such smaller islands can be found around larger islands of regular motion which are surrounded by the chaotic layer. Since they are located near the boundary of the large island, they are forming a so-called boundary islands chain (BIC) (Zaslavsky, 1998). Such a chain of islands can lie within the island of regular motion or outside of it, i. e. within the chaotic layer. We are interested here in the latter since the trajectory is confined to the chaotic layer. Each island in an island chain is surrounded by another island chain of higher order (see e. g. Lichtenberg and Lieberman, 1992).

We use a nonlinearity parameter that is rather small, $K = 0.6$, in order to have a chaotic layer that does not cover too much of state space. This helps to find the situation we are concerned with: a trajectory approaching an island of regular motion, passing through a series of cantori (Schellnhuber et al., 1986; Schellnhuber and Urbschat, 1988), the remains of destroyed KAM tori. At the beginning, the state of the system is not restricted within the chaotic layer, the trajectory moves around the primary, or order 1 island. If it passes through a cantorus, it will then stay in the vicinity of order 2 islands until it either goes back to move around the order 1 island again, or passes another cantorus to visit islands of higher order. An example of this behaviour can be seen in Fig. 3.2. Henceforth, we will call the surroundings of the order o islands the order o layer.

3.2 Spectrum of local exponents

An interesting feature distinguishing the different kinds of local exponents can already be seen by plotting the values of the local exponents against time, as in Fig. 3.3.

One can clearly distinguish different regimes in the case of the maximum

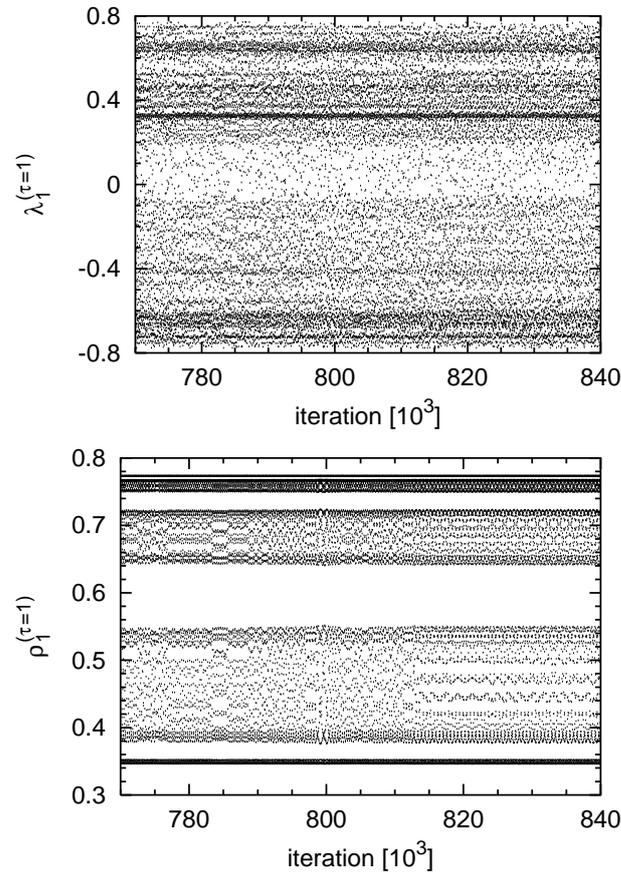


Figure 3.3: Values of the local exponents $\lambda_1^{(\tau=1)}$ (top) and $\rho_1^{(\tau=1)}$ (bottom) for a piece of trajectory. Only every third value is plotted.

growth exponent $\rho_1^{(\tau=1)}$, whereas the values of $\lambda_1^{(\tau=1)}$ take on the same range of values over the whole time span plotted. As a side note, the plot only shows the data for every third point along the trajectory. Such simplifications (for the sake of less crowded plots) have to be done carefully, though. If a piece of trajectory moves around an island chain of six islands, for example, then a plot of the trajectory in state space that takes into account only every third point will produce only two islands.

From Fig. 3.3 one can see that starting at around iteration 815000, for example, the spectrum of $\rho_1^{(\tau=1)}$ develops gaps that were previously not there. In order to understand this, one can look at lines $s(\rho)$ of constant $\rho_1^{(\tau=1)}$ in state space. In the present case of the standard map, the simple dependence of the Jacobian DM on x_1 and x_2 (through $x_1 + x_2$ alone) already hints at the simplicity of the

3.2. Spectrum of local exponents

lines $s(\rho)$:

$$DM(x_1, x_2) = \begin{pmatrix} 1 & 1 \\ K \cos(x_1 + x_2) & 1 + K \cos(x_1 + x_2) \end{pmatrix}. \quad (3.5)$$

The value of $\rho_1^{(\tau=1)}$ is the natural logarithm of the larger singular value σ_1 of the matrix $DM(x_1, x_2)$, i. e. the larger eigenvalue of the matrix $DM^T DM$. Using the abbreviation $z = K \cos(x_1 + x_2)$, one finds that

$$\sigma_1(z) = \sqrt{z^2 + z + \frac{3}{2}} + \sqrt{z^4 + 2z^3 + 4z^2 + 3z + \frac{5}{4}}. \quad (3.6)$$

Lines of constant $\rho_1^{(\tau=1)}$ can be found by first setting σ_1 to a constant value of $\tilde{\sigma}$ and then determining the relation between x_1 and x_2 . Since σ_1 depends on x_1 and x_2 only through z , it is easy to see that one possibility is $x_2 = -x_1$. To address this more rigorously, one can define an implicit function $S(x_1, x_2)$ that has to vanish:

$$S(x_1, x_2) = \sigma_1(K \cos(x_1 + x_2)) - \tilde{\sigma} \stackrel{!}{=} 0. \quad (3.7)$$

The slope of this implicit function is given by

$$\frac{dx_2}{dx_1} = -\frac{\frac{\partial S}{\partial x_1}}{\frac{\partial S}{\partial x_2}} = -\frac{\frac{\partial S}{\partial z} \frac{\partial z}{\partial x_1}}{\frac{\partial S}{\partial z} \frac{\partial z}{\partial x_2}} = -\frac{\frac{\partial z}{\partial x_1}}{\frac{\partial z}{\partial x_2}} = -\frac{-K \sin(x_1 + x_2)}{-K \sin(x_1 + x_2)} = -1. \quad (3.8)$$

The lines $s(\rho)$ of constant $\rho_1^{(\tau=1)}$ are therefore straight lines with a slope of -1 .

Knowing this, the lower part of Fig. 3.3 can be explained as follows. If a gap between $\rho_{1,\min}$ and $\rho_{1,\max}$ arises in the spectrum of ρ_1 , it is because the trajectory does no longer visit states that would contribute such values. This means that the trajectory is now restricted to a smaller region in state space, such that no point lies in between lines $s(\rho_{\min})$ and $s(\rho_{\max})$. The restriction to a smaller region of state space occurs when the trajectory passes through a cantorus and enters a layer of higher order, as in Fig. 3.2. In fact, pieces of the trajectory showing gaps in the values of $\rho_1^{(\tau=1)}$ do lie in the vicinity of islands of higher order than other pieces which have no gaps.

The occurrence of gaps in the spectrum of ρ does not alter the extreme values, so from this point of view it is not possible to speak of enhanced predictability. On the other hand, the trajectory is confined to a smaller region in state space, this can be argued for as better predictability. But this is more of an academical problem for the following reason: This prominent – and easily detectable – feature of gaps in the spectrum of ρ that are linked to movement in layers of different

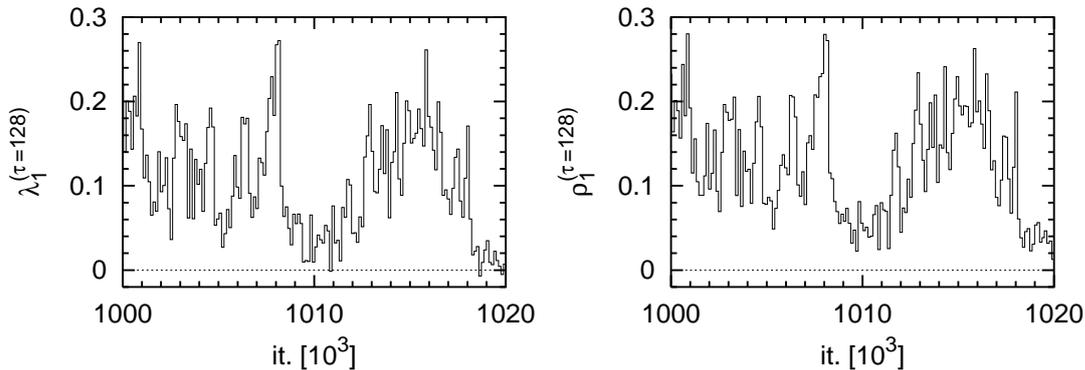


Figure 3.4: Local exponents for a long time horizon $\tau = 128$ along a segment of trajectory: finite time Lyapunov exponents $\lambda_1^{(\tau=128)}$ (left) and maximum growth exponents $\rho_1^{(\tau=128)}$ (right).

order can only be seen in simple systems. We do not expect to find anything similar in models with a medium to large number of variables, simply because the lines $s(\rho)$ will be much more complicated, and the positions of the islands cannot be expected to fall nicely into the “right” place.

In contrast to these findings for the maximum growth exponent $\rho_1^{(\tau=1)}$, the finite time Lyapunov exponent $\lambda_1^{(\tau=1)}$ does not show such a behaviour, since it depends not only on the Jacobians along the trajectory, but also on the past of the system.

To get a better feeling for the variability of local exponents over longer stretches of time, exemplary plots are shown in Figs. 3.4 and 3.5. As can be seen in Fig. 3.4, the larger finite time Lyapunov exponent can become negative even for large time horizons. This is only of limited value regarding predictability, however, since the other exponent, λ_2 , is then positive (s. p. 19 about the ordering of finite time Lyapunov exponents). So there does not necessarily occur a shrinking of a deviation from the true state. If the deviation has already developed earlier, then it will reach the initial state associated with a negative $\lambda_1^{(\tau)}$ pointing in the “right” direction, i. e. the one linked to $\lambda_1^{(\tau)}$, and thus shrink over τ time steps.

Judging by eye, one can see no significant differences between the plots for different local exponents, at least for very long time horizons (Fig. 3.5). This is actually to be expected if the time horizon τ reaches large values. Since we know that in the limit $\tau \rightarrow \infty$ both $\lambda_1^{(\tau)}$ and $\rho_1^{(\tau)}$ have to converge to the same value, it comes as no surprise that the distributions are almost the same for $\tau = 16384$. Therefore, we will only show results for both kinds of local exponents in this

3.2. Spectrum of local exponents

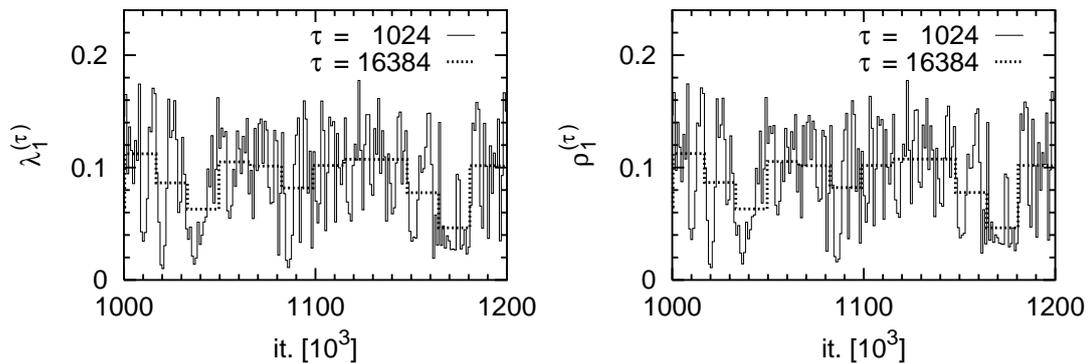


Figure 3.5: Local exponents for very long time horizons τ along a segment of trajectory: finite time Lyapunov exponents $\lambda_1^{(\tau)}$ (left) and maximum growth exponents $\rho_1^{(\tau)}$ (right).

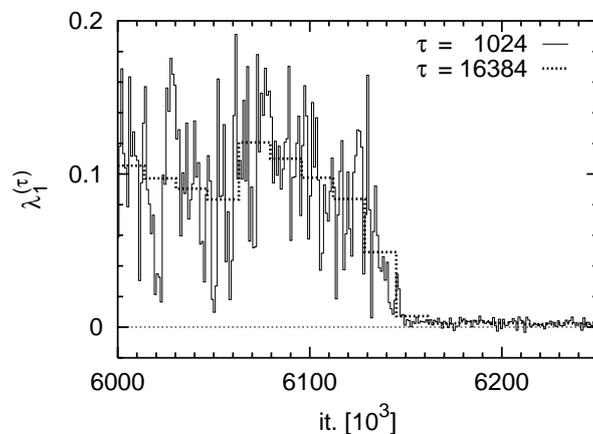


Figure 3.6: Finite time Lyapunov exponent $\lambda_1^{(\tau)}$ for very long time horizons along a segment of trajectory showing a change of the system dynamics.

chapter if there is a significant difference.

The plots in Figs. 3.4 and 3.5 give a good picture of what is going on most of the time. Sometimes, however, the trajectory passes through a cantorus and stays in the vicinity of islands of higher order. We expect this behaviour to be reflected in the values of the local exponents: The higher the order of the islands the trajectory stays close to, the closer to zero are the local exponents. Fig. 3.6 shows data for a stretch of time where the trajectory clearly changes its behaviour in some way. This dramatic change in the spectrum of $\lambda_1^{(\tau)}$ is indeed linked with a change of the regions in state space the trajectory visits.

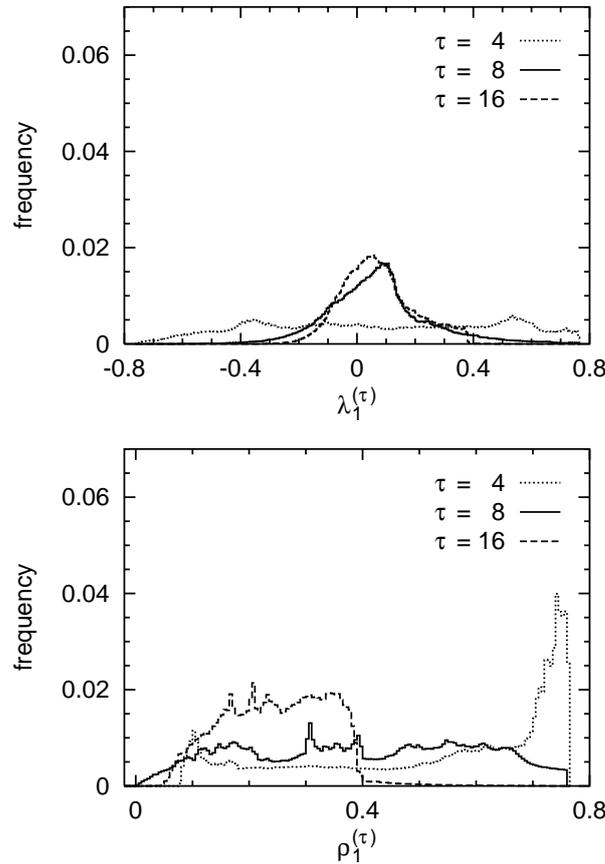


Figure 3.7: Distributions of finite time Lyapunov exponents $\lambda_1^{(\tau)}$ (top) and maximum growth exponents $\rho_1^{(\tau)}$ (bottom) for small to intermediate time horizons τ .

We will now move on to the distributions of local exponents for various values of the time horizon τ .

3.3 Distributions of local exponents

For small time horizons, Fig. 3.7 shows the histograms of the largest finite time Lyapunov exponent $\lambda_1^{(\tau)}$ and the largest maximum growth exponent $\rho_1^{(\tau)}$. It can already be seen from this that the curves become narrower and more peaked with increasing time horizon τ . This is expected, since both kinds of local exponents are known to converge to the values of the global exponents for $\tau \rightarrow \infty$.

Furthermore, the larger maximum growth exponent is restricted to positive

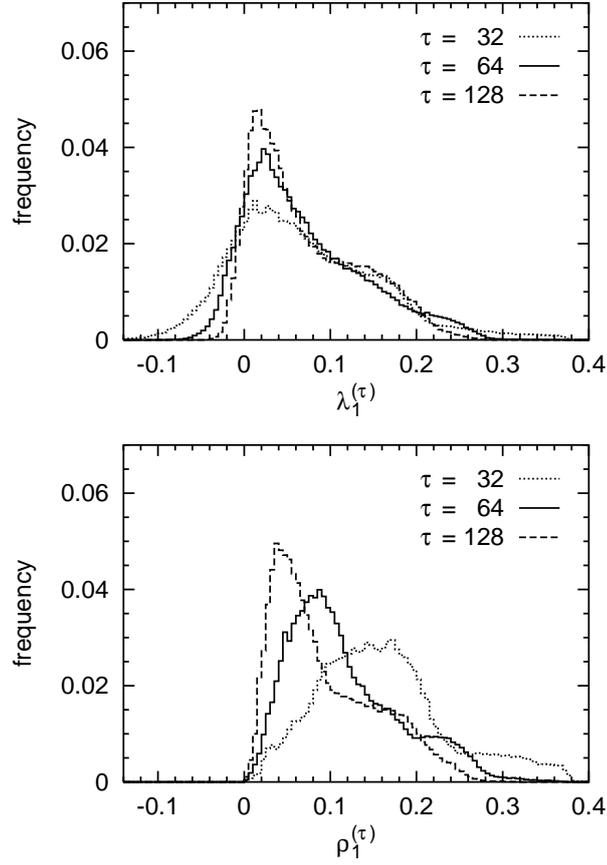


Figure 3.8: Distributions of finite time Lyapunov exponents $\lambda_1^{(\tau)}$ (top) and maximum growth exponents $\rho_1^{(\tau)}$ (bottom) for long time horizons τ .

values. This is a direct consequence of the definition via singular values (which are by convention ordered in a descending sequence $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d$) and the fact that $\rho_2(x_1, x_2) = -\rho_1(x_1, x_2)$. In contrast to this, the largest finite time Lyapunov exponent can attain negative values. The second finite time Lyapunov exponent is then positive, but $\lambda_1^{(\tau)}$ and $\lambda_2^{(\tau)}$ are not interchangeable, because they are linked to particular vectors $l_{1/2}$ of initial error orientation (s. Sec. 2.1.2). The fact that $\lambda_1^{(\tau)} < 0$ can be interpreted as enhanced predictability, since a typical deviation from the true system state will be stretched by $\lambda_1^{(\tau)}$ with probability 1. If one wants to be more demanding, one can also require both finite time Lyapunov exponents to be negative, meaning that any deviation will actually shrink within the time for which the exponents are calculated. This more demanding criterion will be used in the following chapter.

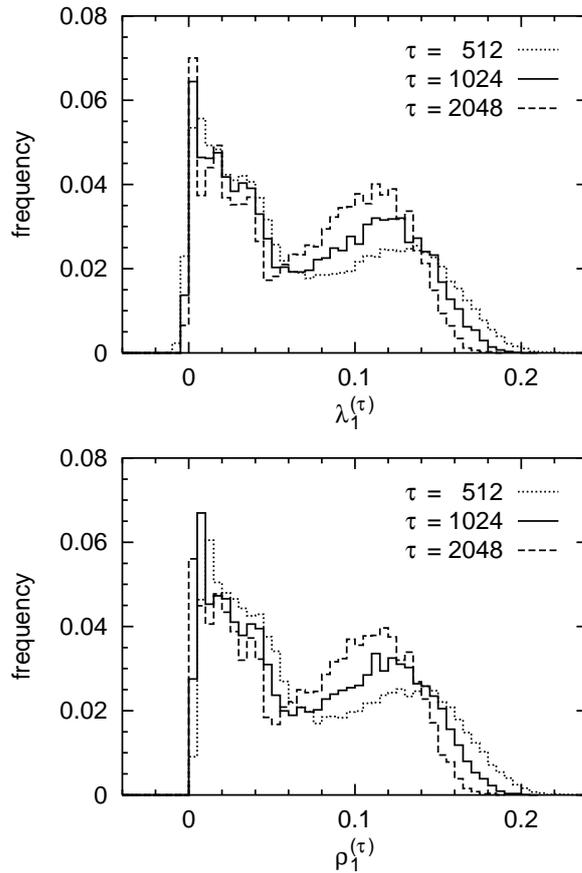


Figure 3.9: Distributions of finite time Lyapunov exponents $\lambda_1^{(\tau)}$ (top) and maximum growth exponents $\rho_1^{(\tau)}$ (bottom) for very long time horizons τ .

The distributions of local exponents within the chaotic layer change significantly upon further increasing τ (Fig. 3.8). For both kinds of local exponents, the distributions for $\tau = 128$ suggest the development of a shoulder towards higher values, while the mean of the ρ_1 -distribution shifts closer to zero. In order to test this hypothesis, we calculated the local exponents for even longer time horizons τ (Fig. 3.9). One can now clearly see the formation of a second peak at values greater than 0.1. The same qualitative behaviour is found for the distributions of maximum growth exponents $\rho_1^{(\tau)}$, starting at somewhat higher values of τ .

This behaviour stems from the fact that the trajectory visits different regions of the chaotic layer. Layers of higher order lead to smaller values of local exponents. Consequently, the rightmost peak in Fig. 3.9 is due to the dynamics around the order 1 island.

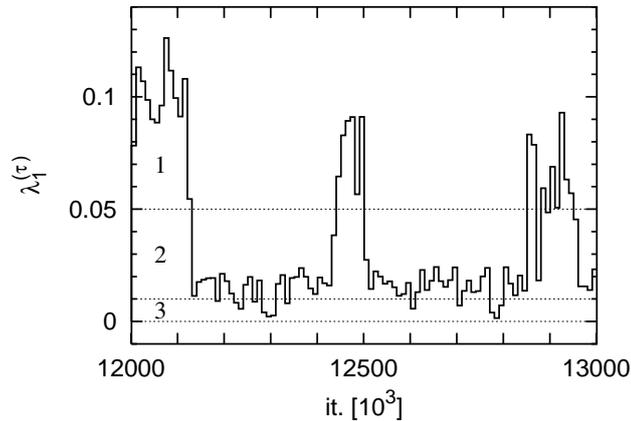


Figure 3.10: Values of finite time Lyapunov exponents $\lambda_1^{(\tau)}$, averaged over $\tau = 10000$ iterations, for a piece of trajectory. The thin dotted lines separate stretches associated with motion in layers of different order, the orders are given by the numbers (see text for details).

In order to handle the large amount of data efficiently, we did not use the appearance of gaps in the spectrum of $\rho_1^{(\tau)}$ to identify the order of the current layer, or looking at the actual distribution of points in state space. Instead, we use finite time Lyapunov exponents averaged over a very long time, like $\tau = 10000$. This allows one to distinguish times with rather high values of $\lambda_1^{(\tau=10000)}$ and times with low values, as can be seen exemplarily in Fig. 3.10. More precisely, time slices with $\lambda_1^{(\tau=10000)} > 0.05$ are taken to be associated with motion in layer 1, the lower bound for layer 2 is 0.01, and even lower values of $\lambda_1^{(\tau=10000)}$ are labeled layer 3 (see Fig. 3.10). This is, of course, not strictly true, since all higher order layers contribute to this part of the $\lambda_1^{(\tau=10000)}$ spectrum. However, trying to determine a further bound, separating layer 3 from higher order layers, is difficult. Furthermore, reasonable choices do not alter the result for layer 3 and do not give sufficient statistics for higher order layers either. Even if the distinction of different layers by a simple boundary value of $\lambda_1^{(\tau)}$ is correct, the use of long-term average values will lead to errors. Parts of the trajectory will be categorised as belonging to the wrong layer. On the whole, however, they will not change the statistics too much.

Taking all time slices belonging to a particular layer, distributions of finite time Lyapunov exponents for the different orders of layers can be calculated. The result is shown in Figs. 3.11 and 3.12. As can be seen, the distributions are shifted towards smaller values for increasing order. What is also indicated is a

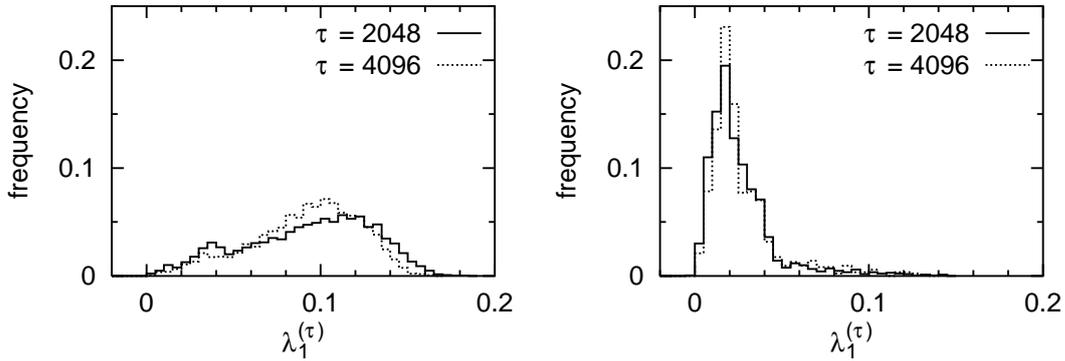


Figure 3.11: Distributions of finite time Lyapunov exponents $\lambda_1^{(\tau)}$ for orders 1 (left) and 2 (right).

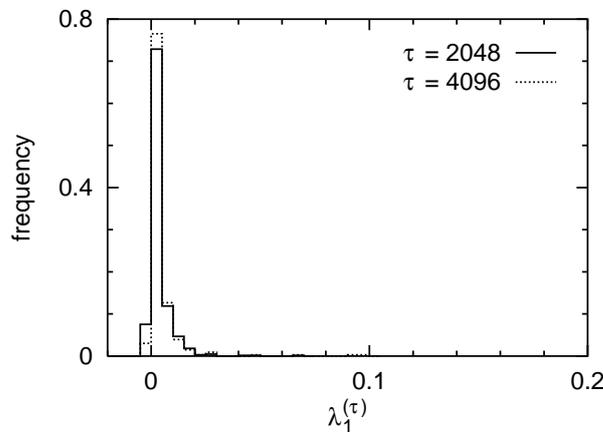


Figure 3.12: Distributions of finite time Lyapunov exponents $\lambda_1^{(\tau)}$ for order 3.

minute second peak at smaller values for order 1. This is very probably because of pieces of trajectory that are actually of higher order, but are erroneously classified as order 1. Similarly, the distributions for orders 2 and 3 extend farther to the right.

To compare the results for different orders with the distribution for the whole trajectory, the distribution for each order is scaled by the number of data points and plotted in Fig. 3.13. The distributions of the different layers, taken together, explain the shape of the global distribution very well.

Zaslavsky et al. (1993) found anomalous (i. e. non-Gaussian) transport in chaotic layers, for which the scaling exponent of the mean squared particle displacement is $\gamma \neq 1$. They showed that this exponent can be expressed through

3.3. Distributions of local exponents

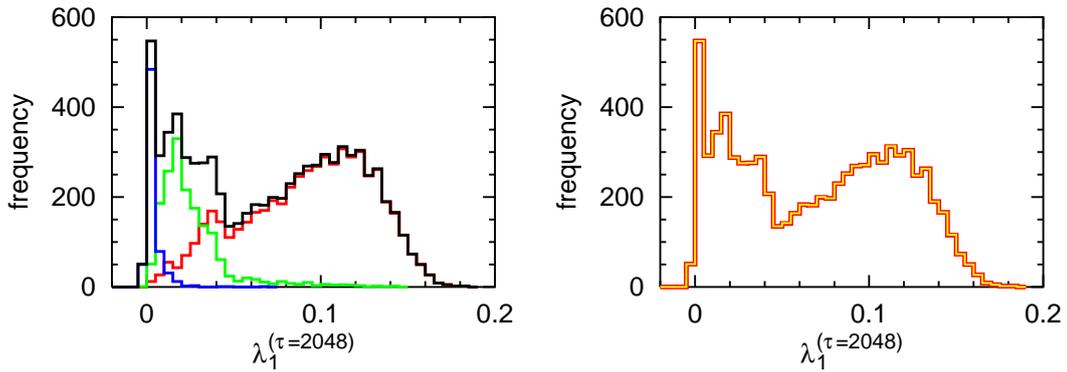


Figure 3.13: Distributions of finite time Lyapunov exponents $\lambda_1^{(\tau=2048)}$. Left: orders 1 (red), 2 (green) and 3 (blue), weighted sum (black). Right: weighted sum over orders (thick red line) and global (yellow).

local constants of the dynamics. One of these is a scaling coefficient describing the dependence of the total size of the islands belonging to one order on their order. The other is the characteristic time scale for the island of a particular order, which is related to the Lyapunov exponent for motion in the corresponding chaotic layer. Such studies involve investigations of the topology of state space, which of course plays an important role, but was beyond the scope of this thesis.

Chapter 4

Local exponents and homoclinic tangencies

In this chapter we take a look at the position of regions of extreme predictability in state space, the most predictable and the least predictable regions. This is done using one conservative as well as one dissipative system, namely the standard map and the Hénon map.

It would be advantageous if one were able to relate regions of particularly good (or bad) predictability to other properties of the dynamical system. Then, one could hope to gain further insights by studying these related properties. Ideally, observing an easily calculated variable gives us all the information about the predictability of the system we are interested in. This ideal cannot be attained in practice, of course. Nevertheless, one can hope to identify at least some regions in state space that stand out by providing starting positions for forecasts of either very good or very bad predictability. This information can then be used by checking whether the current system state lies within such an identified region.

As a candidate for this kind of property of dynamical systems we investigated homoclinic tangencies, which will be introduced in the following section.

4.1 Definition of homoclinic tangencies

In a point of homoclinic tangency, stable and unstable manifolds become tangent to each other. Formulated more mathematically, the definition of homoclinic tangencies incorporates the stable manifolds W^s and the unstable manifolds W^u of a periodic orbit. In the case of a two-dimensional map M , the manifolds of a

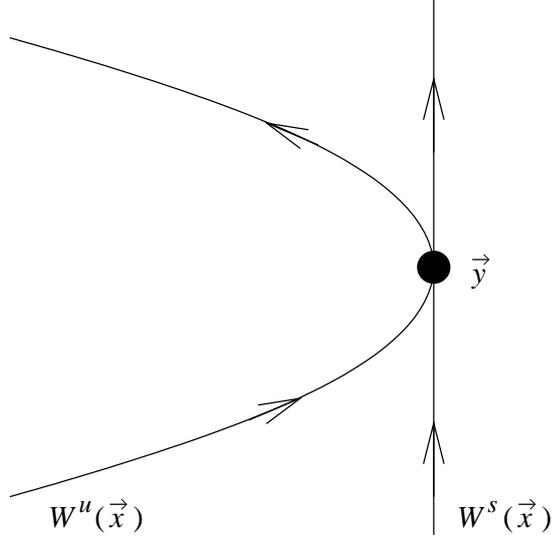


Figure 4.1: Schematic situation near a homoclinic tangency: Lines show the stable and unstable manifolds, the arrows give the direction in which a point is mapped (i. e. where $M(\vec{z})$ will be in relation to \vec{z}). Note that the point of homoclinic tangency, \vec{y} , corresponds to some point \vec{y} in Eqs. 4.1–4.2, not to \vec{x} .

point \vec{x} are:

$$W^s(\vec{x}) = \{ \vec{y} \in \mathbb{R}^2 \mid \lim_{n \rightarrow \infty} \|M^n(\vec{x}) - M^n(\vec{y})\| = 0 \}, \quad (4.1)$$

$$W^u(\vec{x}) = \{ \vec{y} \in \mathbb{R}^2 \mid \lim_{n \rightarrow -\infty} \|M^n(\vec{x}) - M^n(\vec{y})\| = 0 \}. \quad (4.2)$$

This means that the stable manifold of \vec{x} consists of those points \vec{y} whose future images $M^n(\vec{y})$ under the map M will come arbitrarily close to the corresponding image $M^n(\vec{x})$ of the reference point \vec{x} . For the unstable manifold, the same is true for reversed time, i. e. the preimages of \vec{x} and \vec{y} approach each other. A point where $W^s(\vec{x})$ and $W^u(\vec{x})$ touch but do not intersect transversally is called a homoclinic tangency. A sketch of this situation is shown in Fig. 4.1. The fact that the stable manifold is shown as a straight line and the unstable manifold as a parabola is not arbitrary, but can be shown to be the general structure of these manifolds in the vicinity of a homoclinic tangency (Palis and Takens, 1993).

In general, infinitely many unstable periodic orbits are embedded in a chaotic attractor. Each of these orbits possesses stable and unstable manifolds and, thus, allows for homoclinic tangencies to occur.

Due to ergodicity, there has to be a dense set of homoclinic tangencies on the attractor if one point of homoclinic tangency exists on the attractor (Grassberger et al., 1988). Some of these points form sequences, they are related in the

following sense: If \vec{x}_{HT} is a point of homoclinic tangency, so are its predecessor $M^{-1}(\vec{x}_{\text{HT}})$ and its successor $M(\vec{x}_{\text{HT}})$, since the stable and unstable manifolds are invariant sets. Among each such sequence of homoclinic tangencies, one of them stands out by having a minimal (along this sequence) sum of the curvatures of stable and unstable manifolds. These are called “primary homoclinic tangencies” and are the easiest to find numerically (see Jaeger and Kantz, 1997, and references therein): With increasing n , one has to come ever closer to the points $M^n(\vec{x}_{\text{HT}})$ and $M^{-n}(\vec{x}_{\text{HT}})$ in order to numerically identify a homoclinic tangency.

Homoclinic tangencies play an important role in bifurcations, in particular in crises of chaotic attractors (Grebogi et al., 1982, 1983) and so-called metamorphoses, sudden changes in basin boundaries (Grebogi et al., 1986, 1987). Crises occur when varying a system parameter leads to a collision of the attractor with an unstable periodic orbit. One can differentiate between multiple types of crises that affect the attractor in different ways. Boundary crises, for example, lead to the destruction of the attractor. They happen when the attractor collides with an unstable periodic orbit on the attractor’s basin of attraction boundary. Interior crises, on the other hand, cause sudden changes of the attractor size and occur when the attractor collides with an unstable periodic orbit within the basin of attraction. These phenomena are also studied experimentally, for example with dripping faucets (Pinto and Sartorelli, 2000). More examples can be found in Robert et al. (2000).

It is important to note that there is a connection between homoclinic tangencies and the concept of hyperbolicity. Exact results and theorems are mainly known for Axiom-*A* systems as they are mathematically more easily manageable. Hyperbolicity is the primary feature of Axiom-*A* systems (Abraham and Smale, 1968). It requires that stable and unstable manifolds intersect transversally in any point. Thus the concept of hyperbolicity and the existence of homoclinic tangencies are mutually exclusive. Realistic models of the physical world are almost always nonhyperbolic, implying the presence of homoclinic tangencies.

Furthermore, homoclinic tangencies have the advantageous property of being invariant under coordinate transformations. Since this is not the case for the typical local measures of predictability like local exponents, homoclinic tangencies can help to identify regions of enhanced predictability regardless of the particular measure of local predictability.

One can now ask for the effect of the approach of the unstable manifold towards the stable one on the predictability of the dynamics. The following plausibility argument is the reason for choosing to study homoclinic tangencies.

When a trajectory comes close to a point of homoclinic tangency \vec{x}_{HT} , the

orientation of a perturbation can no longer point away from the stable direction at a large angle, since the unstable manifold has to approach the stable one by definition. Instead, the perturbation has to become nearly aligned with the direction of the stable manifold, if the trajectory is close enough to \vec{x}_{HT} .

This becomes more important when the unstable manifold stays close to the stable manifold in a larger surrounding of \vec{x}_{HT} , which can be expressed by the curvatures of the manifolds in the point \vec{x}_{HT} . The smaller the sum of the curvatures, the larger should the effect of the homoclinic tangency be on perturbation growth. Actually one defines those points of homoclinic tangency \vec{x}_{HT} that have a minimal sum of these curvatures as *primary* homoclinic tangencies (Grassberger and Kantz, 1985).

Fortunately, the homoclinic tangencies with a small sum of curvatures, which are the ones we are most interested in, are the easiest to find numerically. This is even more fortunate considering the fact that points of homoclinic tangency are dense on the attractor. This could lead one to assume that there is no gain in linking certain properties to homoclinic tangencies. However, there *is* something to be gained if we do not consider all homoclinic tangencies, but only those with a rather small sum of manifold curvatures. These are the homoclinic tangencies we mean when we talk about connections between homoclinic tangencies and regions of enhanced predictability in Sec. 4.3.

4.2 Calculation of homoclinic tangencies

The method used to calculate homoclinic tangencies derives directly from the definition. Along with the system state \vec{x} , a vector (randomly initialised) is iterated with the Jacobians $DM(\vec{x})$ at the point of the trajectory. This vector will point in the direction of the unstable manifold after a transient time. From then on, vector and point of the trajectory are stored. By inverting time (i. e., by using the stored trajectory in reverse order), adopting the same procedure yields the direction of the stable manifold at each point of the trajectory. Whenever the angle between the two directions is less than some predetermined small value, a homoclinic tangency is identified. The application of this method of finding homoclinic tangencies is restricted to two-dimensional maps, whereas the calculation of the local exponents described in section 2.1 has no such limitation in principle.

An alternative method of computing homoclinic tangencies is presented by Jaeger and Kantz (1997). It is based on the fact that the most expanding direction in tangent space over n iterations is perpendicular to the tangent of the attractor, provided that n is large enough.

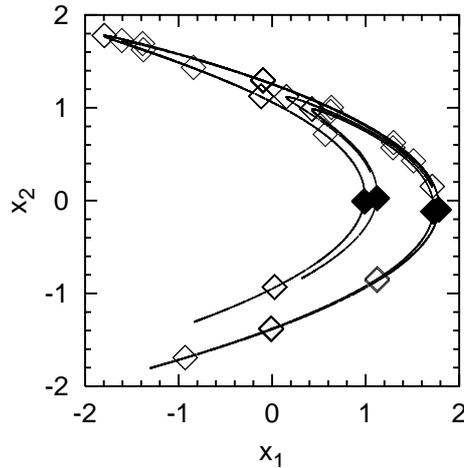


Figure 4.2: Points of homoclinic tangency (diamonds) and attractor for the Hénon map ($a = 1.4$, $b = 0.3$). Filled symbols denote primary homoclinic tangencies.

Some of the homoclinic tangencies of the Hénon system are plotted together with the attractor in Fig. 4.2. The primary homoclinic tangencies are represented by the four filled diamonds with $x_2 \approx 0$ or $x_2 \approx -0.1$. In the case of the standard map, the homoclinic tangencies for one initial position are shown in Fig. 4.3. The corresponding trajectory eventually fills the chaotic layer. Near the origin, i. e., in the corners of the left part of Fig. 4.3, not so many homoclinic tangencies were found, whereas the connecting arcs are quite densely covered. Furthermore, one can see some structure in the region around the origin (s. right part of Fig. 4.3). This feature will be discussed in greater detail in the following section.

4.3 Connection between homoclinic tangencies and regions of enhanced predictability

Based on the argument presented in Sec. 4.1, the predictability can be expected to be better when the system is in a state near a homoclinic tangency. Conversely, the most deteriorated predictability should occur when the angle between stable and unstable manifold is largest, i. e., close to a right angle. In the following, we will compare these hypotheses with the results obtained for the Hénon map and the standard map.

Overview images of local exponents (and in the following chapter also of error growth factors, in the case of finite initial errors) are all created by the same

4.3. Connection between homoclinic tangencies and regions of enhanced predictability

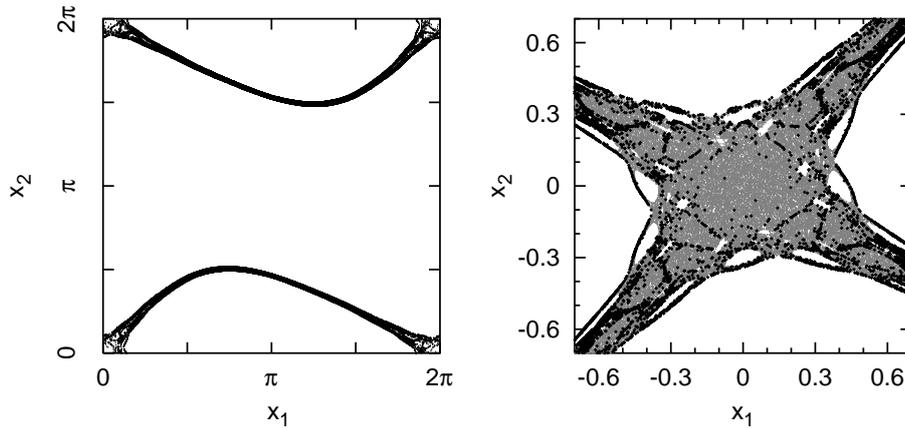


Figure 4.3: Points of homoclinic tangency for the standard map ($K = 0.6$). The left part shows only points of homoclinic tangency inside the chaotic layer. The right part shows a magnification of the region around the origin, with ordinary points inside the chaotic layer shown in grey and homoclinic tangencies in black.

procedure. This is explained here lest we interrupt the discussion of results. For good examples of what such an image looks like the reader is referred to Figs. 4.10 (p. 60) and 4.12 (p. 61).

These images are made up of many small cells, formed by a rectangular grid laid over the region of interest. For each cell, the number of points of a given trajectory that fall into this cell is recorded, along with the corresponding values of $\lambda_1^{(\tau=1)}$ or $\rho_1^{(\tau=1)}$. In the end each cell is coloured according to the average value recorded for this cell. If a cell is never visited there are no values associated with it, so it remains white.

Ideally, each point along the trajectory would give one point in the image. A straightforward plotting of the values, however, would lead to one of the following situations: Either many points will overlap other points, and only the colour of the topmost point, i. e., the last one drawn and thus the last one along the trajectory, will be visible. Or the points are small enough not to overlap, but in this case they are typically too small to give a good impression of the structure and the distribution of values in state space. This is why we have chosen the visualisation using averages over a fine rectangular grid.

Not surprisingly, the method we use has its drawbacks, too. Apart from the finite graining, another possible source for misleading images lies in extreme values that can change the average. To prevent this, the cells should be made small enough. In the figures shown herein, the average number of input points

per cell is in the low hundreds for overview pictures like Figs. 4.10 and 4.11. For more detailed figures showing magnifications, there are typically a few dozen points in each cell. One can make the cells even smaller, of course, but this leads to more cells being empty (white), since no point of the trajectory happens to lie inside. The final image will then have “holes” where in fact the trajectory would eventually pass by — it would not only be less pleasant to the eye, but also more difficult to interpret correctly. Such “holes” in the image can also arise when comparing different prediction times τ . The number of input points can be calculated as the number of points along the trajectory divided by τ . Consequently, taking a given trajectory and region of state space, increasing the prediction time τ will typically produce white holes inside of the coloured parts. This has to be kept in mind when comparing images for different τ .

From what has just been said the obvious idea to remove white cells is to start trajectories in them and then calculate the desired value. More precisely, one should start the trajectories in a preimage of a point in the cell for the finite time Lyapunov exponents, since the algorithm needs some iterations to find the right directions of errors associated with the Lyapunov exponents (s. Sec. 2.1.3). Alternatively, one can start with a given point and do some n iterations backwards as well as forwards in time. Then, the information gained on the directions of error growth can be used to approximate the Lyapunov vectors \vec{l}_i . This has not been done in this work, but is possible in principle. The removal of white gaps is even easier for the maximum growth exponents, since they do not depend on the past of the trajectory in any way. The same holds for the error growth factors to be discussed in the next chapter.

4.3.1 Hénon map

The distribution of local exponent values in state space for the Hénon map with parameter values of $a = 1.4$ and $b = 0.3$ can be seen in Fig. 4.4. However, since the attractor does not cover a large portion of state space, this kind of picture is only useful as a general overview.

In this section we understand enhanced predictability to mean that both values of the finite time Lyapunov exponents are negative, $\lambda_{1,2}^{(\tau)} < 0$, ensuring that any infinitesimal error will shrink over the time span τ . Regions with such enhanced predictability can be seen to exist in Figs. 4.5 and 4.6; indeed, they are related to points of homoclinic tangency. By comparing different values of the prediction time horizon τ , one can see that around some homoclinic tangencies the regions of good predictability do not show up for all τ , e. g., at $x_1 \approx 0$ for $\tau = 1$.

Some of the homoclinic tangencies cannot be linked with a region of good predictability for any values of τ studied. This might be due to the statistics

4.3. Connection between homoclinic tangencies and regions of enhanced predictability

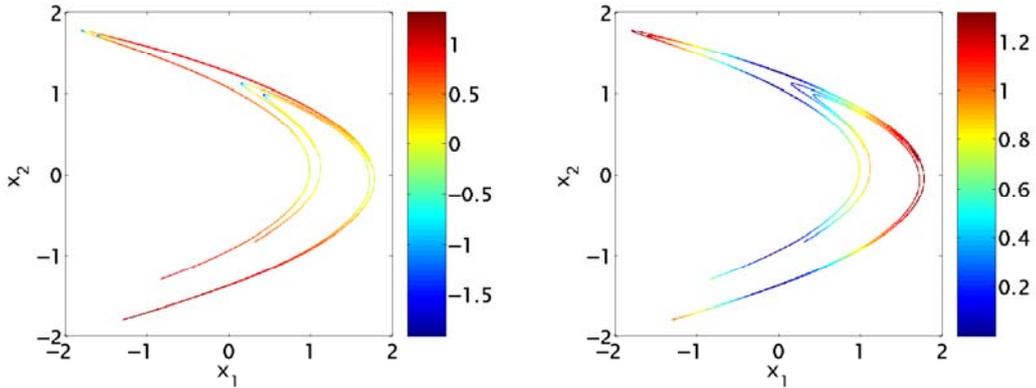


Figure 4.4: Hénon map ($a = 1.4$, $b = 0.3$): Local averages of $\lambda_1^{(\tau=1)}$ (left) and $\rho_1^{(\tau=1)}$ (right). See p. 53 for details about how this and similar plots are produced.

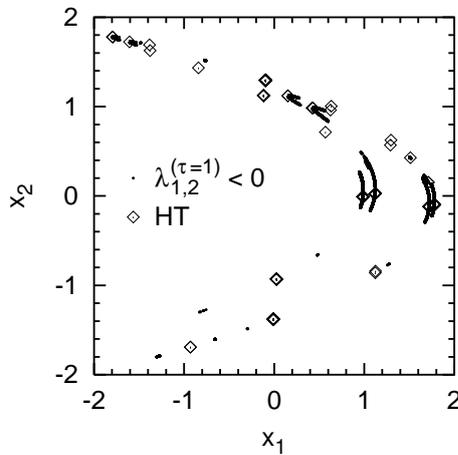


Figure 4.5: Hénon map ($a = 1.4$, $b = 0.3$): Comparison of regions of enhanced predictability ($\lambda_{1,2}^{(\tau=1)} < 0$, filled circles) with points of homoclinic tangency (diamonds) for $\tau = 1$ time step.

being not good enough, since the regions of good predictability get smaller upon increasing τ . This reduces the probability of hitting such a region with a finite piece of trajectory.

To get more insight into the dependence of the size of the region of enhanced predictability on τ , we study the region around an arbitrarily chosen homoclinic

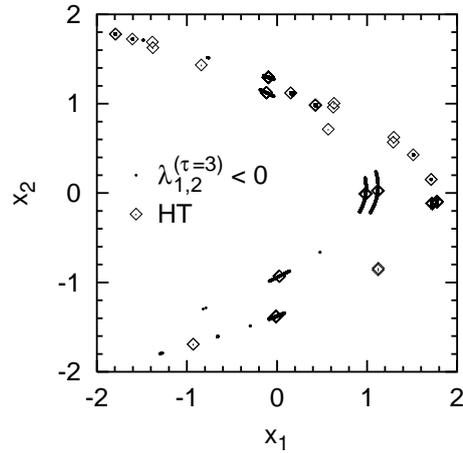


Figure 4.6: Hénon map ($a = 1.4$, $b = 0.3$): Comparison of regions of enhanced predictability ($\lambda_{1,2}^{(\tau=3)} < 0$, filled circles) with points of homoclinic tangency (diamonds) for $\tau = 3$ time steps.

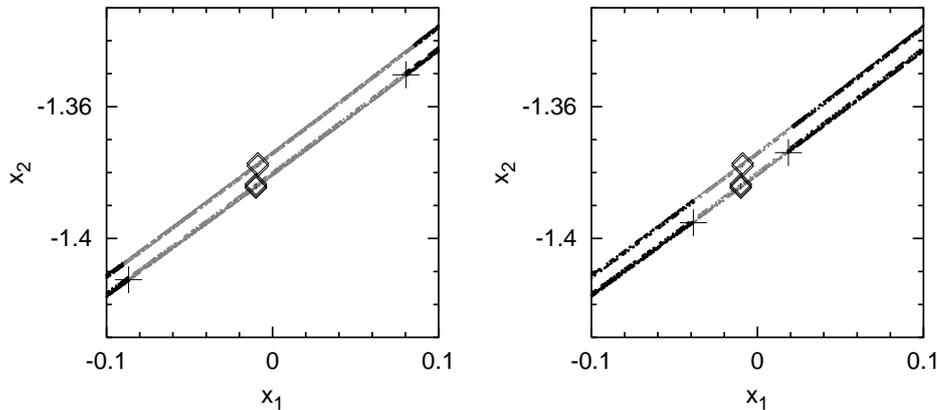


Figure 4.7: Enlargement of a region of enhanced predictability for the Hénon attractor ($a = 1.4$, $b = 0.3$) for $\tau = 3$ (left) and $\tau = 5$ (right): Points with $\lambda_1^{(\tau)} < 0$ and $\lambda_2^{(\tau)} < 0$ are shown in grey, other points in black. Diamonds show homoclinic tangencies, the crosses mark the points used to calculate the length of the region of enhanced predictability.

tangency. To this end, we examine the region around $x_1 \approx 0$, $x_2 \approx -1.4$ with respect to the length of the segment of the attractor showing enhanced predictability. Figure 4.7 shows a magnification of this region, with points \vec{x} for which $\lambda_1^{(\tau)}(\vec{x}) < 0$ and $\lambda_2^{(\tau)}(\vec{x}) < 0$ shown in grey.

4.3. Connection between homoclinic tangencies and regions of enhanced predictability

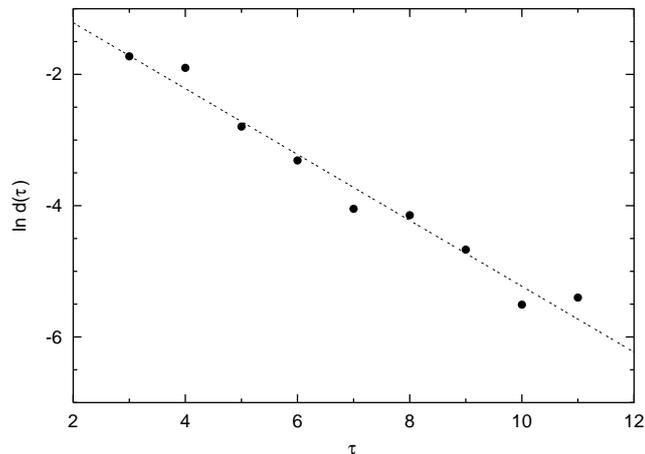


Figure 4.8: Scaling of the size of a region of enhanced predictability with τ for the Hénon map ($a = 1.4$, $b = 0.3$). The ordinate shows the logarithm of the length d of the region (see text for a description of how d is calculated). The dashed line represents a linear least squares fit: $\ln d(\tau) = -0.50\tau - 0.21$.

The size of the region of enhanced predictability is measured by considering the “border” values to the left and right of this region (marked with crosses in Fig. 4.7), i. e. the first points for which at least one finite time Lyapunov exponent is positive. The Euclidean distance d of these border points is what we call the size of the region of enhanced predictability. Even though it seems like there are only two “lines” of points in the region considered, one has to keep in mind that there are actually infinitely many¹. For practical reasons we restrict the border points to the lower “line” in Fig. 4.7.

The logarithm of the distance of border points is plotted versus τ in Fig. 4.8 and shows that the size of the regions of enhanced predictability shrinks exponentially with increasing prediction time τ .

A possible connection between regions of extraordinarily small values of the maximum growth exponents $\rho^{(\tau)}$ and homoclinic tangencies was investigated, but could not be established as for the finite time Lyapunov exponents.

The opposite extremes – nearly orthogonal intersections of stable and unstable manifolds and deteriorated predictability – are found to be not related, regardless of the kind of local exponents used.

¹The transversal structure of the Hénon attractor is a Cantor set (Hénon, 1976).

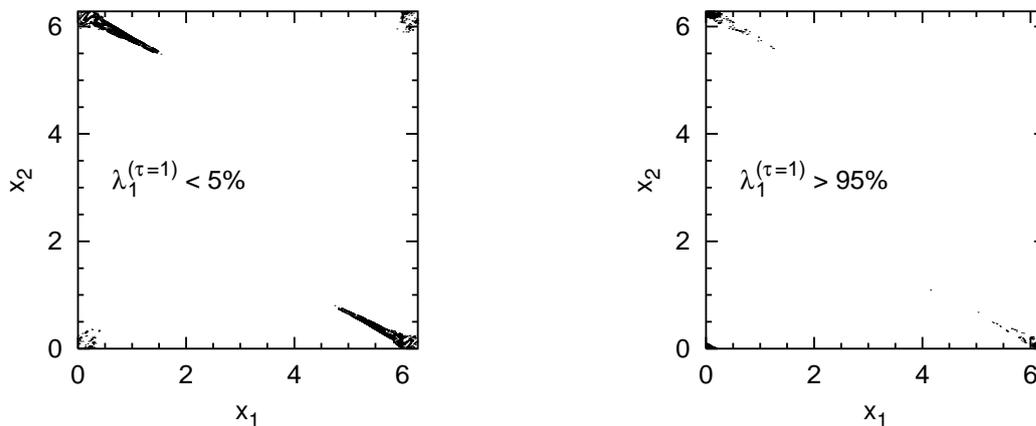


Figure 4.9: Regions with enhanced and poor predictability for the standard map ($K = 0.6$), according to finite time exponents $\lambda_1^{(\tau=1)}$ for a time horizon of one step. Plotted points have a largest exponent of less than the 0.05-percentile (left panel) or more than the 0.95-percentile (right panel).

4.3.2 Standard map

Let us now discuss the results in conservative systems. The regions with enhanced and poor predictability are shown in Fig. 4.9 for the standard map at $K = 0.6$. An important fact to note is that the regions for good and bad predictability occur very close to each other. In fact they are complexly intertwined, which makes it rather difficult to develop a measure for the prediction of predictability. Slight changes in the considered region in state space may lead to a dramatic change in predictability.

To get an idea of the distribution in state space of the local exponent values, Figs. 4.10 and 4.11 show an overview of the whole $2\pi \times 2\pi$ unit cell that determines all values due to periodicity.

The relation between homoclinic tangencies and points of enhanced predictability is exemplarily illustrated for the standard map in Fig. 4.12, which shows a magnification of the region around the origin. The homoclinic tangencies are marked by black dots. The colour shows the local average of the largest finite time exponent $\lambda_1^{(\tau=1)}$ using grid cells of side length 0.005. In the upper left and lower right branches, the homoclinic tangencies correspond well with the meandering pattern exhibited by the dark blue parts, which represent states with good predictability for one time step ($\tau = 1$). Similar curves of black points in

4.3. Connection between homoclinic tangencies and regions of enhanced predictability

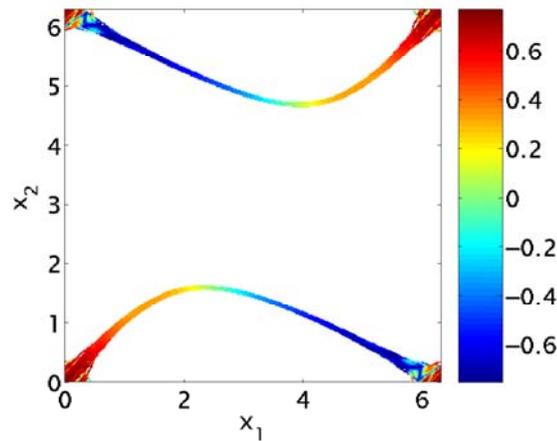


Figure 4.10: Standard map ($K = 0.6$): local average of largest finite time Lyapunov exponent $\lambda_1^{(\tau=1)}$. See p. 53 for details about how this and similar plots are produced.

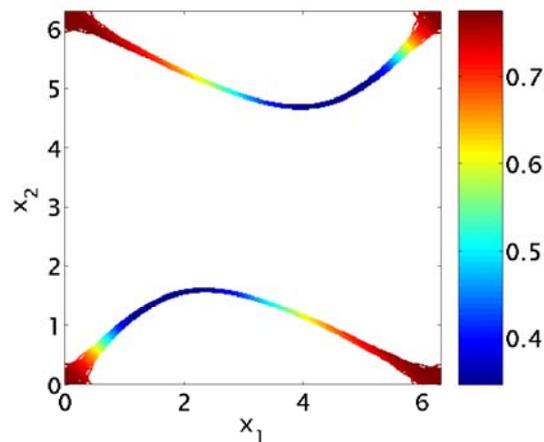


Figure 4.11: Standard map ($K = 0.6$): local average of largest maximum growth exponent $\rho_1^{(\tau=1)}$.

the lower left and upper right parts of Fig. 4.12 correspond to a pattern visible for $\tau = 8$ (Fig. 4.13). Note that the colour scale is the same as in Fig. 4.12, the absence of intense colours is due to the smaller width of the distribution of $\lambda_1^{(\tau=8)}$ compared to that of $\lambda_1^{(\tau=1)}$.

The structure of these homoclinic tangencies clearly resembles the behaviour

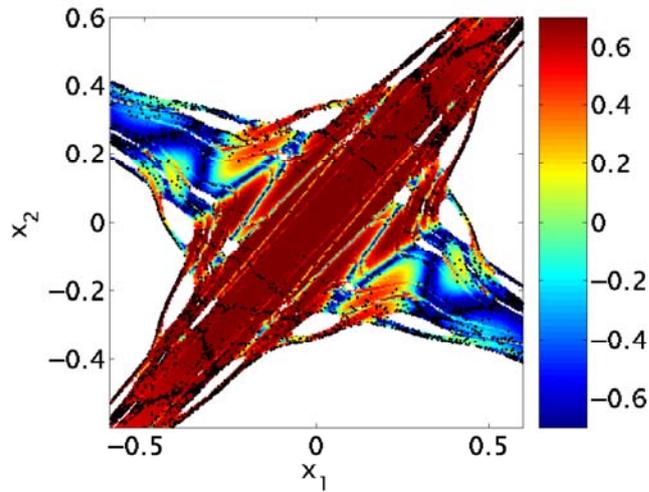


Figure 4.12: Standard map ($K = 0.6$): points of homoclinic tangency (black) and local average of $\lambda_1^{(\tau=1)}$ (colour) around the origin.

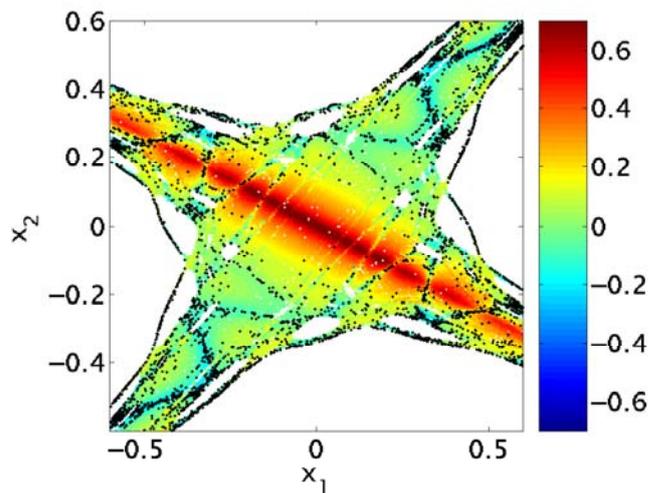


Figure 4.13: Standard map ($K = 0.6$): points of homoclinic tangency (black) and local average of $\lambda_1^{(\tau=8)}$ (colour). The colour code is the same as in Fig. 4.12, the smaller width of the distribution of $\lambda_1^{(\tau=8)}$ leads to less intense colours.

of the stable and unstable manifolds of a fixed point. In our case, the point $(0, 0)$ is the fixed point and the local direction of its stable manifold leaves the picture

4.3. Connection between homoclinic tangencies and regions of enhanced predictability

in the lower right and upper left, while the unstable manifold near the origin lies almost on the diagonal $x_2 = x_1$. After going along the arcs visible in Fig. 4.3, the stable manifold approaches the fixed point from the directions in which the unstable manifold leaves the fixed point, oscillating more and more wildly. The pattern of the black dots is reminiscent of these oscillations.

As in the case of the Hénon map, nearly orthogonal intersections of stable and unstable manifolds and deteriorated predictability are found to be unrelated to each other either, independent of the kind of local exponents used.

Chapter 5

Predictability studies using ensembles

In the two previous chapters we investigated error growth using local exponents, i. e. we implicitly assumed the errors to remain small enough for the linearisation to remain valid. In this chapter we drop this assumption and follow the evolution of finite errors as governed by the full nonlinear system.

Such studies are necessary whenever the initial error size is too large or the prediction time horizon is too long for the local exponents to yield an accurate description of error growth. This is the common situation in practice, where even in the case of a perfect model the initial state cannot be determined to infinite accuracy by measurements. This is the reason why operational weather forecasting models of today, for example, use methods involving the full nonlinear system (e. g. Buizza, 1996; Toth and Kalnay, 1993).

One such method is the study of an ensemble of trajectories. In a nutshell, this method can be described as follows: initialise a bunch of trajectories $\vec{y}_j(t_0)$ around the system state $\vec{x}(t_0)$, let all these states evolve for an amount of time τ and compare these states $\vec{y}_j(t_0 + \tau)$ with the “true” state $\vec{x}(t_0 + \tau)$. A detailed description of which quantities are used to describe predictability will be given in the following section.

Using the nonlinear system for an ensemble of trajectories in order to determine error growth increases the computational cost. This is the price one has to pay for the increase in reliability, or for the possibility of estimating the “predictability” of the forecast.

The larger the ensemble one uses, the more costly the calculations will be. From this it is obvious that choosing a “good” size of the ensemble is vital for finding an acceptable trade-off between reliability of results and computational cost. Therefore, it would be very helpful to have a tool at hand which provides

an estimate of the minimal ensemble size needed for a certain desired reliability.

This chapter is organised as follows: In Sec. 5.1 we introduce the concept of ensembles. We also define the quantities that are now used in place of local exponents to quantify the predictability of the initial state, the error growth factors. After this technical part we come to the results. In Sec. 5.2 histograms of error growth factors are presented for the example systems studied, the standard map and the quasiperiodicity map. A comparison of results for local exponents and error growth factors is given in Sec. 5.3. Then we turn to the systematics found in the results for the nonlinear error growth factors in Secs. 5.4 and 5.5. In particular, we show that the error growth in the worst case depends on prediction time in a systematic way (Sec. 5.4). Furthermore, the distributions of error growth factors converge with a characteristic dependence on ensemble size, a fact that can help to determine the minimal ensemble size needed for a given error tolerance (Sec. 5.5). In this last section results for the Hénon map are also presented, in the rest of this chapter we focus on the standard map and the quasiperiodicity map.

5.1 Error growth factors as a measure of divergence of trajectories

In ensemble studies, not only a single trajectory $\vec{x}(t)$ of the dynamical system is followed over a certain period of time, but a whole bunch, or ensemble, of trajectories. These are typically initialised at some time t_0 around the current state $\vec{x}(t_0)$ of the reference trajectory (also called the “true” trajectory):

$$\vec{y}_j(t_0) = \vec{x}(t_0) + \vec{\varepsilon}_j, \quad j \in \{1, \dots, N_{\text{ens}}\}, \quad (5.1)$$

where N_{ens} is the size of the ensemble, \vec{y}_j denote the states of the ensemble members and $\vec{\varepsilon}_j$ are the initial separations. All states, the reference as well as all ensemble members, evolve according to the full nonlinear system dynamics. After some chosen prediction time τ , for each ensemble member the error growth factor $m_j^{(\tau)}$ of the distance from the true state with respect to the initial distance is calculated:

$$m_j^{(\tau)}(x(t_0)) = \frac{\|\vec{y}_j(t_0 + \tau) - \vec{x}(t_0 + \tau)\|}{\|\vec{y}_j(t_0) - \vec{x}(t_0)\|}. \quad (5.2)$$

These factors are then used to measure the predictability of the original state $x(t_0)$. In particular, the minimum $m_{\text{min}}^{(\tau)}(x(t_0))$, maximum $m_{\text{max}}^{(\tau)}(x(t_0))$ and average $m_{\text{avg}}^{(\tau)}(x(t_0))$ of the ensemble are considered. To simplify notation, we will

from now on suppress the explicit dependence of the magnification factors $m_j^{(\tau)}$ on $x(t_0)$.

The maximum error growth factor $m_{\max}^{(\tau)}$ obviously describes the worst case regarding predictability over the time span τ for the given starting position. Correspondingly, $m_{\min}^{(\tau)}$ stands for the best case one can encounter, and $m_{\text{avg}}^{(\tau)}$ can be interpreted as a kind of “average” error growth that is to be expected when nothing about the initial direction of the error is known in advance. When comparing such computations to real world situations, the results for “worst-case error growth”, for example, can only be reliable if the underlying ensemble reflects the frequency of various initial errors in the real world. Since we are not concerned with such comparisons in this work, we can choose our ensembles in the simplest possible way. Due to this no results stemming from a preferred orientation of initial errors will occur, for example. In general, the actual increase of the error for a given ensemble member \vec{y}_j depends, of course, not only on the size, but also on the orientation $\alpha(t_0)$ of the initial error $\vec{\varepsilon}_j$.

It should be noted that the modulo operations contained in the equations of motion of the dynamical systems studied within this chapter are only applied after the distances of ensemble states from the true state have been calculated, i. e. after τ iterations. This ensures that no artefacts are produced by shifting some ensemble members that happened to cross the boundary of the unit cell.

There are a number of free parameters to be set, some of which will be varied in order to study the dependence of predictability on their values. All results presented in the following sections were obtained with a particular way of initialising the distribution of ensemble states. Namely, all initial separations are of the same size, d_0 , the direction is uniformly random. Other initial distributions were compared against this choice: Firstly, a uniformly random size $\leq d_0$ with uniformly random direction. Secondly, Gaussian distributions with standard deviation d_0 in all coordinates. These alternative ways of initialising the ensemble lead to the same distributions of the magnification factors m . Examples of histograms are shown in Figs. 5.1 and 5.2. For $\tau = 1$ one can see small differences, but the distributions are almost identical. For the higher prediction time of $\tau = 3$ the bin width was increased, because otherwise the distributions turn out to be too ragged to be easily comparable. In this case, too, the distributions are virtually identical. This result justifies the procedure of using only one kind of initialisation for the random initial perturbations.

Especially in the case of high-dimensional models of real-world processes, certain directions of initial error can be more probable than others and can have a larger magnitude. For example, if wind velocities are more difficult to determine than temperatures, then the initial distribution of the ensemble members

5.1. Error growth factors as a measure of divergence of trajectories

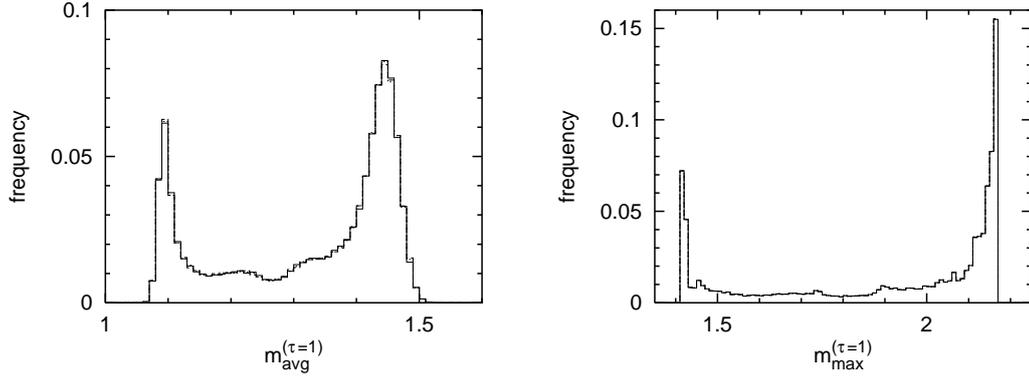


Figure 5.1: Standard map ($K = 0.6$, $N_{\text{ens}} = 1000$, $d_0 = 0.001$): Distributions of one-step error growth factors $m_{\text{avg}}^{(\tau=1)}$ (left) and $m_{\text{max}}^{(\tau=1)}$ (right) for different initial error distributions: fixed size d_0 (solid line), Gaussian (dashed) and uniformly random size (dotted).

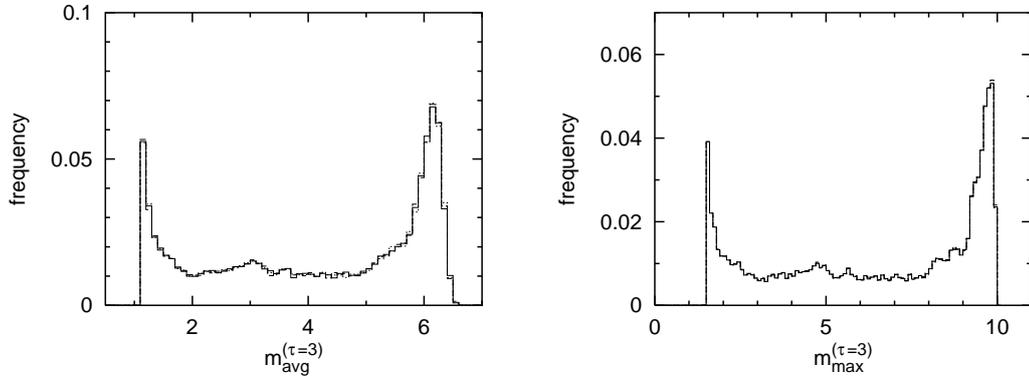


Figure 5.2: Same as Fig. 5.1, but for $\tau = 3$.

should be changed accordingly. In the case of Gaussian distributions of errors in each direction of state space, this can be accomplished by simply increasing the standard deviation of the appropriate component.

The value of d_0 , together with the time horizon τ , are parameters that are varied to study predictability. It turned out that both of them have a strong influence on the predictability. In practice, the size N_{ens} of the ensemble is very important since it is limited by the available computing power. Therefore, results for different values of N_{ens} are compared with each other to determine a minimal value of N_{ens} needed to ensure a reasonable estimate of predictability.

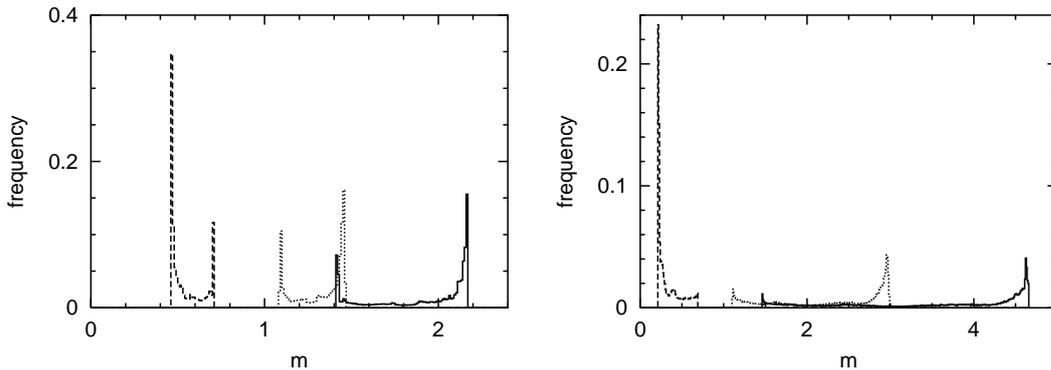


Figure 5.3: Standard map ($K = 0.6$, $N_{\text{ens}} = 20000$, $d_0 = 0.001$): Distributions of error growth factors $m_{\text{max}}^{(\tau)}$ (solid line), $m_{\text{min}}^{(\tau)}$ (dashed) and $m_{\text{avg}}^{(\tau)}$ (dotted) for $\tau = 1$ (left) and $\tau = 2$ (right).

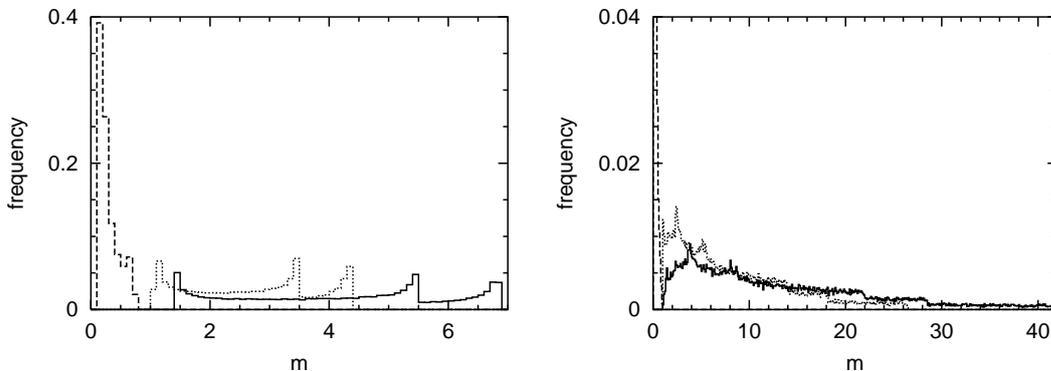


Figure 5.4: Same as Fig. 5.3, but for the standard map at $K = 4.2$.

5.2 Distributions of error growth factors

We calculate the error growth factors for many initial states $x(t_0)$, taken from a long piece of trajectory. The number of initial states used was typically 10^5 . One can get a first impression of the occurring values by looking at the distribution of $m_{\text{min}}^{(\tau)}$, $m_{\text{avg}}^{(\tau)}$ and $m_{\text{max}}^{(\tau)}$. The corresponding histograms are shown in Figs. 5.3 and 5.4 for the standard map and in Figs. 5.5 and 5.6 for the quasiperiodicity map.

In the case of the standard map the distribution of $m_{\text{min}}^{(\tau)}$ does not change its general structure when going from $\tau = 1$ to $\tau = 2$, but the small values become more frequent. The distributions of $m_{\text{avg}}^{(\tau)}$ and $m_{\text{max}}^{(\tau)}$ generally become broader and flatter and are shifted towards higher values. For the quasiperiodicity map the broadening and shifting is much more pronounced for $m_{\text{avg}}^{(\tau)}$ and $m_{\text{max}}^{(\tau)}$. For

5.2. Distributions of error growth factors

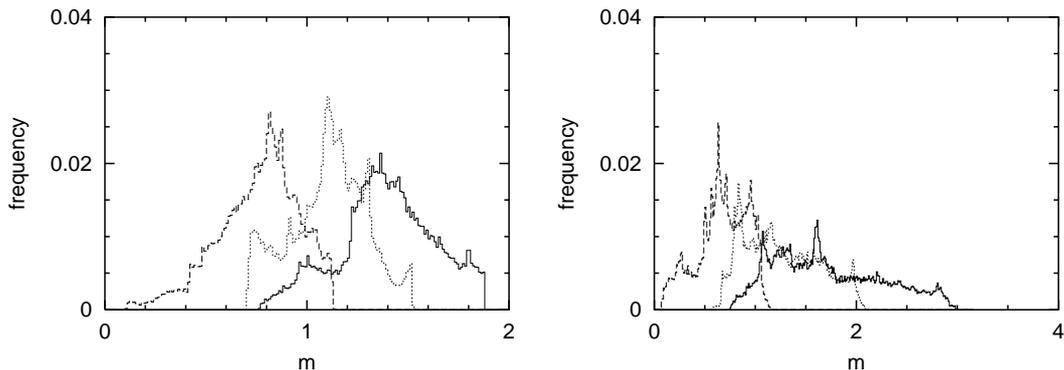


Figure 5.5: Quasiperiodicity map ($c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$; $N_{\text{ens}} = 20000$, $d_0 = 0.001$): Distributions of error growth factors for $m_{\text{max}}^{(\tau)}$ (solid line), $m_{\text{min}}^{(\tau)}$ (dashed) and $m_{\text{avg}}^{(\tau)}$ (dotted), $\tau = 1$ (left) and $\tau = 2$ (right).

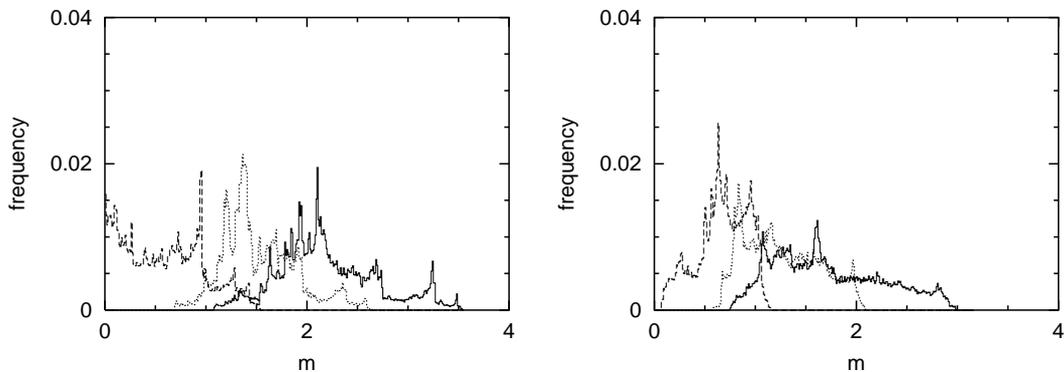


Figure 5.6: Same as Fig. 5.5, but for the quasiperiodicity map with $c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$.

$m_{\text{min}}^{(\tau)}$, the same behaviour can be seen as in the standard map, small values are found more often.

These figures clearly show an important difference between the two systems: The distributions for the standard map have very sharp peaks near the ends of the distributions. On the other hand, the quasiperiodicity map exhibits more Gaussian-like distributions, cut off at $m = 0$ in the case of $m_{\text{min}}^{(\tau)}$. This difference expresses a fundamental difference of the underlying dynamical systems, leading us to believe that our results are general, valid for many chaotic dynamical systems.

One should keep in mind that the system sizes ℓ_{std} for the standard map

and ℓ_{qp} for the quasiperiodicity map are different. For example, the relative size of an initial error of $d_0 = 0.001$ is $\approx 0.0159\%$ of ℓ_{std} and 0.1% of ℓ_{qp} . These are still quite small errors, but considerably higher values of d_0 were studied as well. The different typical sizes of the systems do not facilitate the comparison of results, but this is not a grave disadvantage since we are interested in the general behaviour, not so much in the precise numbers for this or that system.

5.3 Comparison of error growth factors and local exponents

In this section we compare the error growth factors with the corresponding values of the local exponents. We do this by looking at the maximum error growth factor $m_{\text{max}}^{(\tau)}$ and the largest local exponents $\lambda_1^{(\tau)}$ and $\rho_1^{(\tau)}$. Actually these values cannot be compared as such, because the latter are defined as exponents (cf. definitions (2.10), (2.11) and (5.2)). The same values should show up in these quantities (we omit the analogous expressions for $\lambda_1^{(\tau)}$):

$$\begin{aligned} m_{\text{max}}^{(\tau)} &\longleftrightarrow \exp(\tau \rho_1^{(\tau)}) && \text{or} \\ \frac{1}{\tau} \ln m_{\text{max}}^{(\tau)} &\longleftrightarrow \rho_1^{(\tau)}. \end{aligned} \tag{5.3}$$

If they differ significantly, we can conclude that the approximation inherent in the use of local exponents has broken down.

Additionally, this provides a check to whether the error growth factors give sensible values. In the case of small initial deviation d_0 and short prediction time τ , they should describe the same amount of error growth as the local exponents.

In order to compare the local properties in state space, we need to go beyond the distributions discussed in the previous section. Instead, we use two-dimensional plots of state space to visualise the results, coloured according to the local value of error growth factor or local exponent.

How these plots are obtained has already been explained in detail at the beginning of chapter 4.3. To recapitulate: For each point along a trajectory we determine the grid cell to which it belongs and add this point's value (error growth factor or local exponent) to this grid cell's memory. After all points have been dealt with we calculate the average value for each grid cell and colour the cell according to this average value. Some grid cells may never be visited, which results in these grid cells remaining white.

5.3. Comparison of error growth factors and local exponents

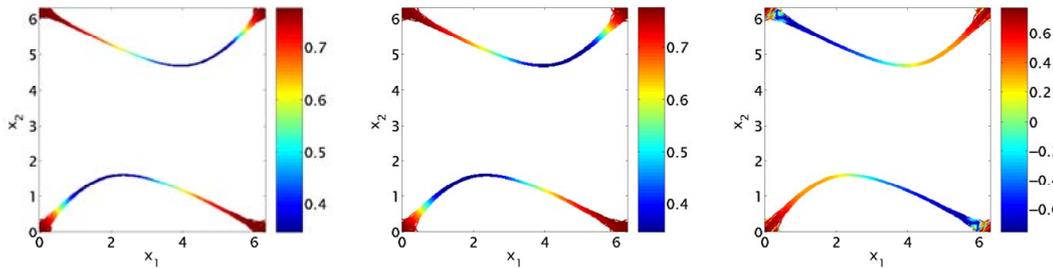


Figure 5.7: Standard map ($K = 0.6$), $\tau = 1$: Comparison of transformed maximum error growth factor $m_{\max}^{(\tau)}$ (left; $N_{\text{ens}} = 10000$, $d_0 = 0.001$) and largest local exponents $\rho_1^{(\tau)}$ (middle) and $\lambda_1^{(\tau)}$ (right).

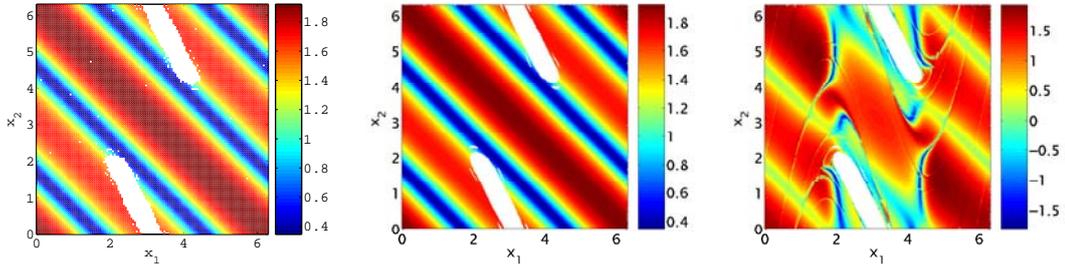


Figure 5.8: Same as Fig. 5.7, but for $K = 4.2$.

Let us first take a look at the standard map for short prediction times τ and small initial errors d_0 . We get the results shown in Fig. 5.7 for $K = 0.6$ and Fig. 5.8 for $K = 4.2$. Already here it is obvious (at least for $K = 4.2$) that one should only compare the transformed values of the error growth factors with the maximum growth exponents, since the finite time Lyapunov exponents show many features that are not present in the $m_{\max}^{(\tau)}$ data. These features are related to the past of the system via its dynamics, but the past states do not influence the maximum error growth. Furthermore, it is only reasonable to compare the error growth factors with the maximum growth exponents, because of their very definition. They are defined to capture the maximum growth of an infinitesimal initial error possible over a given period of time.

For the study of the effects of longer prediction times τ and larger initial errors d_0 we focus on the higher parameter value of $K = 4.2$. The reason for this is simply that a much larger part of state space is visited and the figures are therefore much easier to compare.

From Fig. 5.9 it is clear that for small initial errors like $d_0 = 0.001$, the

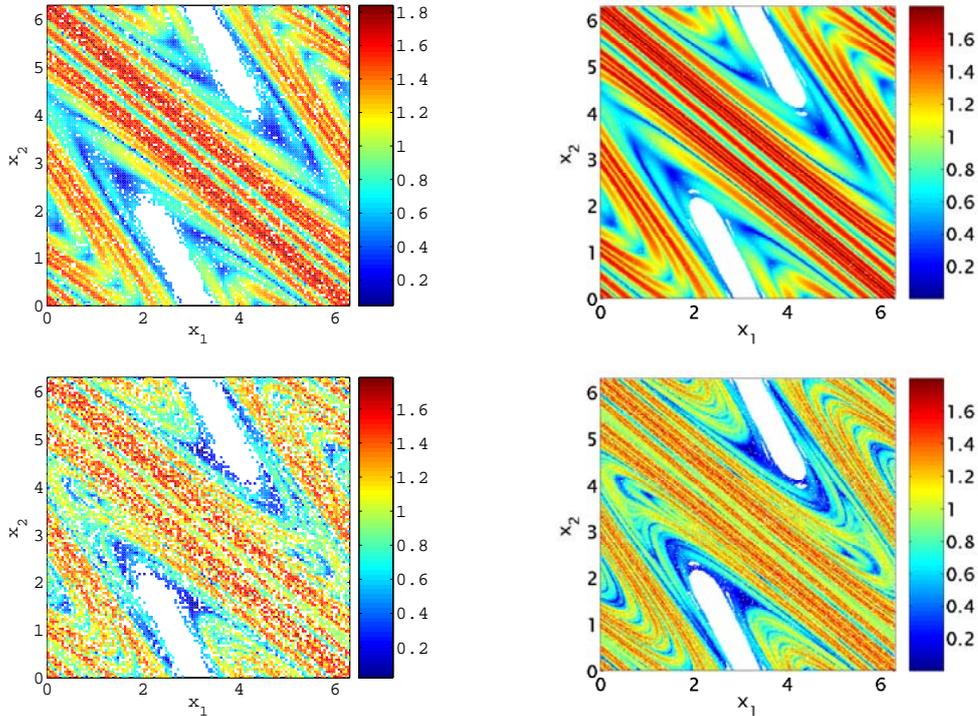


Figure 5.9: Standard map ($K = 4.2$): Comparison of transformed maximum error growth factor $m_{\max}^{(\tau)}$ (left; $N_{\text{ens}} = 10000$, $d_0 = 0.001$) and largest maximum growth exponent $\rho_1^{(\tau)}$ (right). Top row is for $\tau = 3$, bottom row for $\tau = 5$.

maximum growth exponents produce good results up to moderate values of τ . The approximately constant range of the $\rho_1^{(\tau)}$ values when varying τ stems from the fact that the local exponents are defined as an “average” over the number of time steps considered.

When increasing the initial error size d_0 the linearisation becomes less and less accurate. This results in the plots showing more and more discrepancies. For $\tau = 1$ this is difficult to see because of the simple spatial structure. Therefore, Fig. 5.10 shows data for $\tau = 2$. The main difference between these plots is in the neighbourhood of the diagonal $x_2 = 2\pi - x_1$. This can be understood qualitatively by considering the mapping involved. A point $(x_1, 2\pi - x_1)$ is mapped to $(0, 2\pi - x_1)$, since the term $K \sin(x_1 + x_2)$ does not contribute. If, however, a small error is present then the effect of the additional sine term will be maximal when the derivative is maximal, i. e. when the argument of the sine is zero. In other words, initial errors are most harmful when the original point is of the form $(x_1, 2\pi - x_1)$.

5.3. Comparison of error growth factors and local exponents

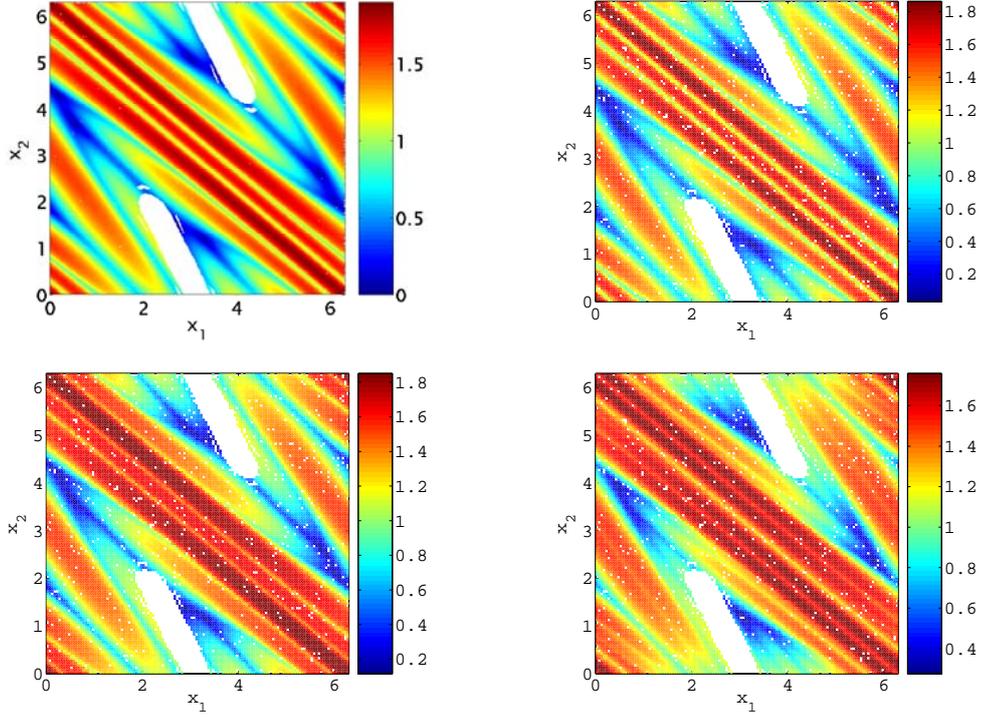


Figure 5.10: Standard map ($K = 4.2$), $\tau = 2$: Comparison of the largest maximum growth exponent $\rho_1^{(\tau=2)}$ (top left) and transformed maximum error growth factors $m_{\max}^{(\tau)}$ ($N_{\text{ens}} = 10000$) for $d_0 = 0.001$ (top right), $d_0 = 0.1$ (bottom left), and $d_0 = 0.3$ (bottom right).

In the case of the quasiperiodicity map we start with the parameter values of $c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$, for which the dynamics is less chaotic and fills more of state space.

The comparison of the transformed maximum error growth factors $m_{\max}^{(\tau)}$ and the largest local exponents for $\tau = 1$ (Fig. 5.11) is even clearer than the results for the standard map. There is a very precise correspondence between $\log m_{\max}^{(\tau)}$ and $\rho_1^{(\tau)}$, whereas $\lambda_1^{(\tau)}$ seems to be completely unrelated to the former two.

Increasing the prediction time τ leads to the results shown in Fig. 5.12. If the initial error size is small (here $d_0 = 0.001$) there is no significant difference even at $\tau = 8$. The reason for this is that the nonlinearity is relatively weak ($c_3 = 0.6$); the other, more chaotic parameter set draws a different picture (s. below).

The effect of increasing the initial error size d_0 is shown in Fig. 5.13. Up to a value of $d_0 = 0.01$ the increase of initial error size does not show at all. Drastic changes only emerge at higher values of d_0 that already constitute a

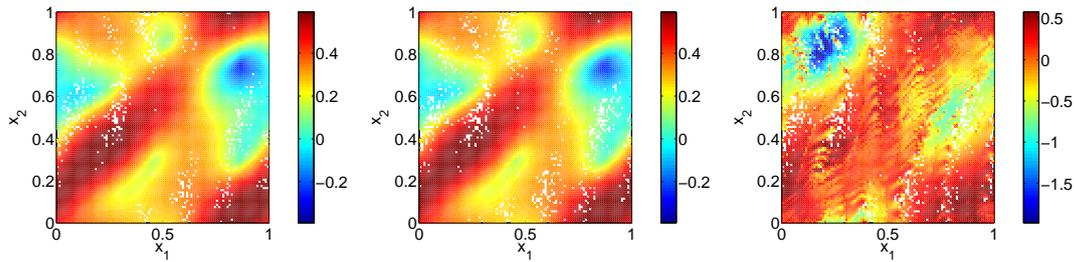


Figure 5.11: Quasiperiodicity map ($c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$), $\tau = 1$: Comparison of transformed maximum error growth factor $m_{\max}^{(\tau)}$ (left; $N_{\text{ens}} = 10000$, $d_0 = 0.001$) and largest local exponents $\rho_1^{(\tau)}$ (middle) and $\lambda_1^{(\tau)}$ (right).

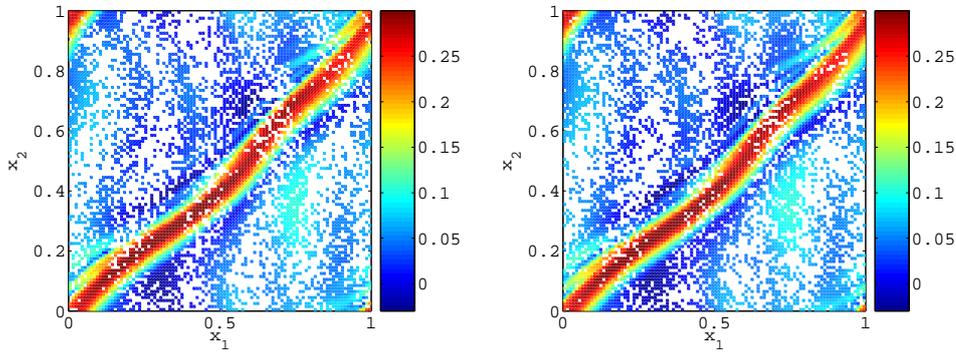


Figure 5.12: Quasiperiodicity map ($c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$), $\tau = 8$: Comparison of transformed maximum error growth factor $m_{\max}^{(\tau)}$ (left; $N_{\text{ens}} = 10000$, $d_0 = 0.001$) and largest maximum growth exponent $\rho_1^{(\tau)}$ (right).

considerable fraction of the system size. The absolute value of error growth per iteration decreases for increasing d_0 as can be seen from the colours. This due to the folding inherent in the mapping (not the modulo operation): Points starting close together diverge at first because of the stretching, but later they are folded together again since the attractor is finite.

The more chaotic parameter set for the quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$) yields the following results. Comparing the transformed maximum error growth factors $m_{\max}^{(\tau)}$ and the largest local exponents for $\tau = 1$ (Fig. 5.14), one finds that the finite time Lyapunov exponents again show many features that are not present in the other plots, whereas the maximum growth exponents and

5.3. Comparison of error growth factors and local exponents

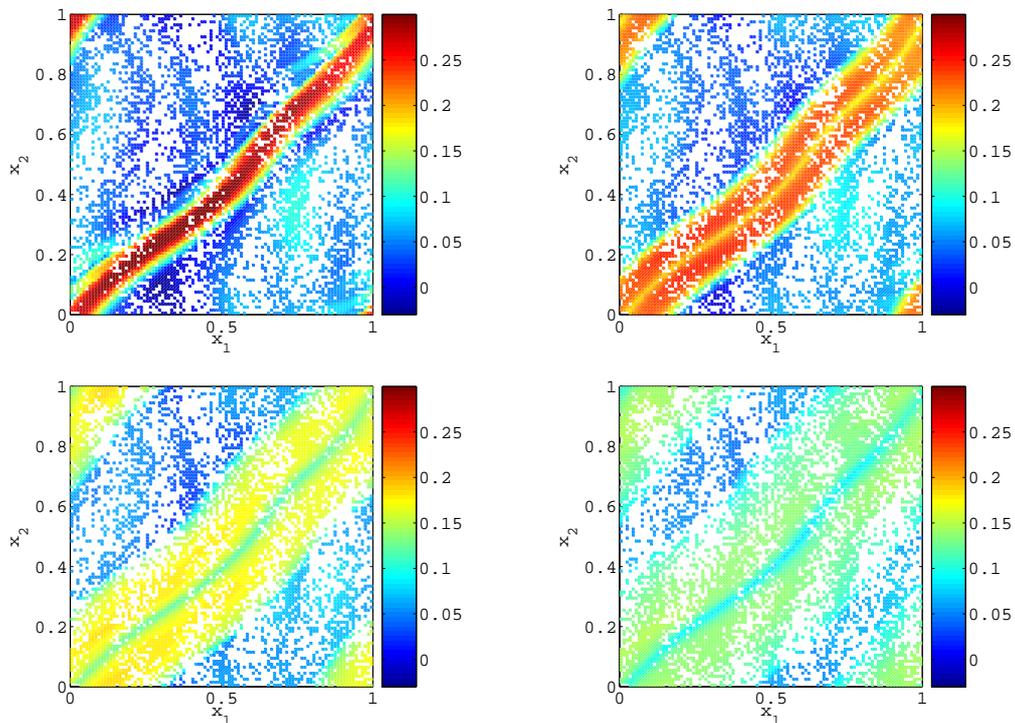


Figure 5.13: Quasiperiodicity map ($c_1 = 0.2$, $c_2 = 0.3$, $c_3 = 0.6$), $\tau = 8$: Comparison of transformed maximum error growth factors $m_{\max}^{(\tau)}$ ($N_{\text{ens}} = 10000$) for $d_0 = 0.01$ (top left), $d_0 = 0.1$ (top right), $d_0 = 0.2$ (bottom left) and $d_0 = 0.3$ (bottom right). The corresponding plots for $d_0 = 0.001$ and $\rho_1^{(\tau)}$ look exactly like the one for $d_0 = 0.01$.

the transformed error growth factors give almost identical results.

Increasing the prediction time τ leads to the results shown in Fig. 5.15. For a small initial error size (here $d_0 = 0.001$) there is no big difference even at $\tau = 5$, but the resemblance is worse than for the less chaotic parameter set at $\tau = 8$.

Figure 5.16 shows the effect of increasing the initial error size d_0 for $\tau = 2$ (as in Fig. 5.10 for the standard map). Deviations from the maximum growth exponent values become obvious at $d_0 = 0.01$ in this case.

We have shown that the maximum error growth factors $m_{\max}^{(\tau)}$ compare well with the largest maximum growth exponent $\rho_1^{(\tau)}$, but not at all with the largest finite time Lyapunov exponent $\lambda_1^{(\tau)}$. In both systems, the correspondence between $m_{\max}^{(\tau)}$ and $\rho_1^{(\tau)}$ remains close up to moderate values of τ even for the more chaotic parameter set. Furthermore, an increase in initial error size d_0 invalidates the

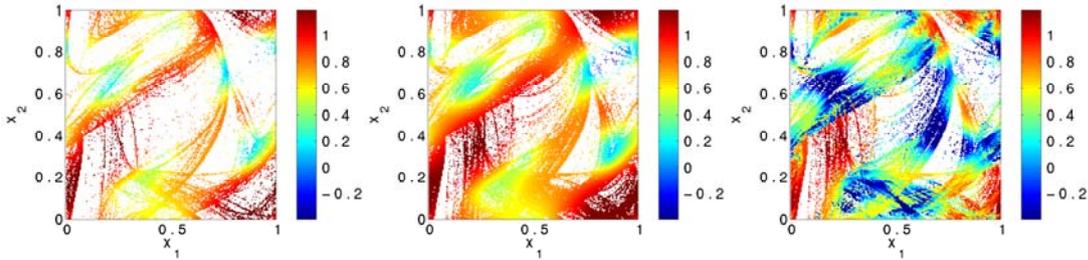


Figure 5.14: Quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$), $\tau = 1$: Comparison of transformed maximum error growth factor $m_{\max}^{(\tau)}$ (left; $N_{\text{ens}} = 10000$, $d_0 = 0.001$) and largest local exponents $\rho_1^{(\tau)}$ (middle) and $\lambda_1^{(\tau)}$ (right). The actual values of $\lambda_1^{(\tau)}$ go as low as -3 , but the same colour mapping is used for ease of comparison.

linear approximation sooner for the more chaotic parameter set.

5.4 Dependence of the worst-case error growth on prediction time

We now go back to the histograms of $m_{\max}^{(\tau)}$ as seen in Figs. 5.3–5.6 and focus on their right ends, i. e. the largest error growth factor found when comparing all initial positions. We will denote this by $M^{(\tau)}$:

$$M^{(\tau)} = \max_{x_0} m_{\max}^{(\tau)}(x_0). \quad (5.4)$$

This quantity represents the worst case for a prediction of time τ when the initial state is not known. The dependence of $M^{(\tau)}$ on τ is shown in Fig. 5.17 for the standard map and in Fig. 5.18 for the quasiperiodicity map, both with the more chaotic parameter values (the number of initial states used is $10^5/\tau$). The semi-logarithmic plots clearly show an exponential increase for small values of τ , in particular for the standard map. For comparison, the dotted line shows $\exp(\rho\tau)$, where $\exp(\rho)$ is the largest singular value for one time step. This means that for $\tau = 1$, the error grows at most by a factor of $\exp(\rho)$, provided that the implicit linearisation is valid. If, after one iteration, the error lies in the direction of maximum error growth of the new state $x(t_0 + 1)$, the growth over two iterations will be $\exp(2\rho)$. Thus, the fact that the solid lines are below the dotted lines in Figs. 5.17 and 5.18 is reasonable, because errors will very rarely happen to point in the “most unpredictable” directions over several iterations.

5.4. Dependence of the worst-case error growth on prediction time

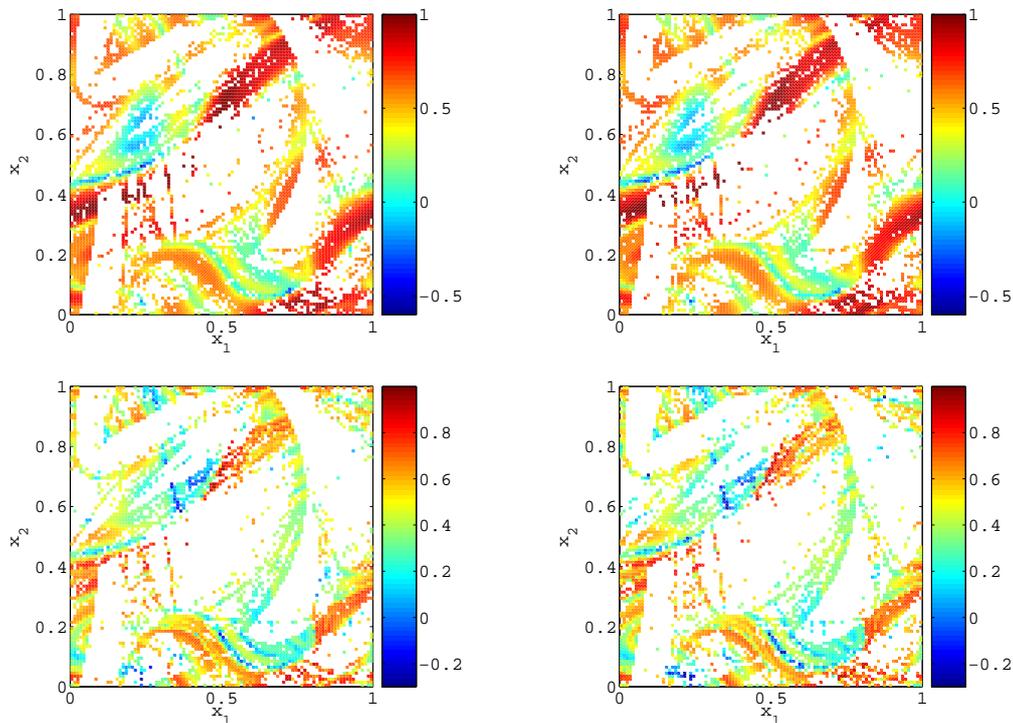


Figure 5.15: Quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$): Comparison of transformed maximum error growth factor $m_{\max}^{(\tau)}$ (left; $N_{\text{ens}} = 10000$, $d_0 = 0.001$) and largest maximum growth exponent $\rho_1^{(\tau)}$ (right). Top row is for $\tau = 3$, bottom row for $\tau = 5$.

Actually it is more appropriate to compare $M^{(\tau)}$ with $\exp(\rho^{(\tau)}\tau)$, the corresponding maximum error growth of the linearised system over time τ . This gives an indication where the linearisation breaks down. For example, in the case of the quasiperiodicity map this does not happen yet at $\tau = 5$ for $d_0 = 10^{-3}$ but already at $\tau = 2$ for $d_0 = 10^{-1}$ (s. Fig. 5.19).

These findings compare well with Figs. 5.9–5.10 for the standard map and Figs. 5.15–5.16 for the quasiperiodicity map: The highest values over the whole unit cell in these figures are the values shown in Fig. 5.19. As we have noted before, for $d_0 = 0.001$ the maximum growth exponents produce very good results up to $\tau = 5$ for both systems (s. Figs. 5.9, 5.15). Comparing different values of d_0 for $\tau = 2$ (s. Figs. 5.10, 5.16), we find that for the standard map a significant difference of the maximum error growth factors from $\rho_1^{(\tau)}$ is first seen for $d_0 = 0.3$. For the quasiperiodicity map, the differing values show up already for $d_0 = 0.01$.

Interestingly, the influence of the nonlinearity does not lead to an increase of

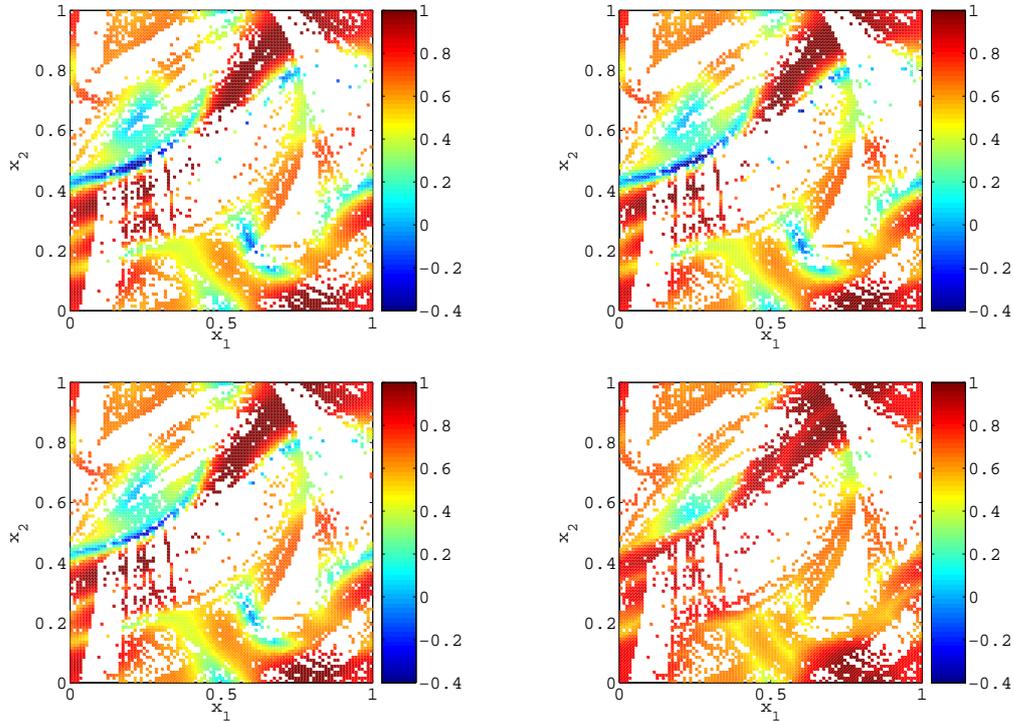


Figure 5.16: Quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$), $\tau = 2$: Comparison of the largest maximum growth exponent $\rho_1^{(\tau)}$ (top left) and transformed maximum error growth factors $m_{\max}^{(\tau)}$ ($N_{\text{ens}} = 10000$) for $d_0 = 0.001$ (top right), $d_0 = 0.01$ (bottom left), and $d_0 = 0.1$ (bottom right).

error growth. In fact the errors do not grow as fast as in the linear system, which can already be seen by looking at Figs. 5.17 and 5.18. This is due to the folding in state space that occurs because of the nonlinearity.

The behaviour of $M^{(\tau)}$ for large times is qualitatively different. The right, double-logarithmic plots of Figs. 5.17 and 5.18 suggest a power-law

$$M^{(\tau)} \sim \tau^\alpha. \quad (5.5)$$

Values of the exponent α are shown in Fig. 5.20. They were obtained from linear least squares fits of the interval $\tau \in [10; 20]$, except for the smallest d_0 in the case of the quasiperiodicity map (cf. Fig. 5.18). A power-law dependence of α on d_0 is found to agree well with the data.

5.4. Dependence of the worst-case error growth on prediction time

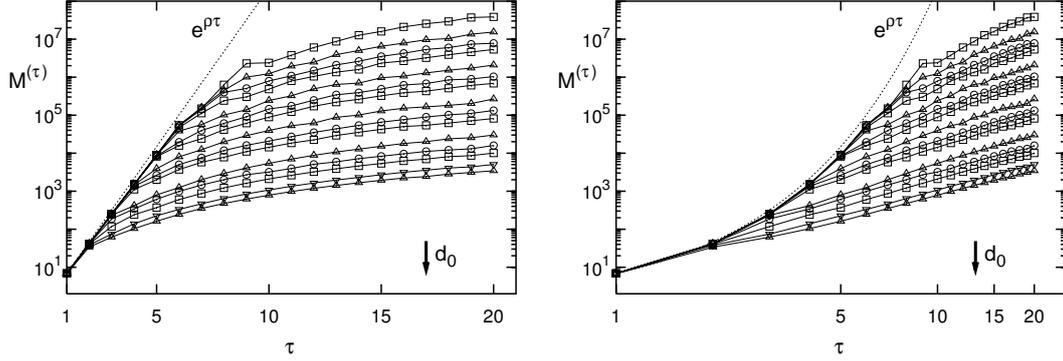


Figure 5.17: Standard map ($K = 4.2$, $N_{\text{ens}} = 10000$): Dependence of $M^{(\tau)}$ on τ . The curves represent different values of d_0 , from 10^{-4} (top) to 3×10^{-1} (bottom). The dotted line shows an exponential increase, see text for details. τ is shown on a linear scale (left) and on a logarithmic scale (right).

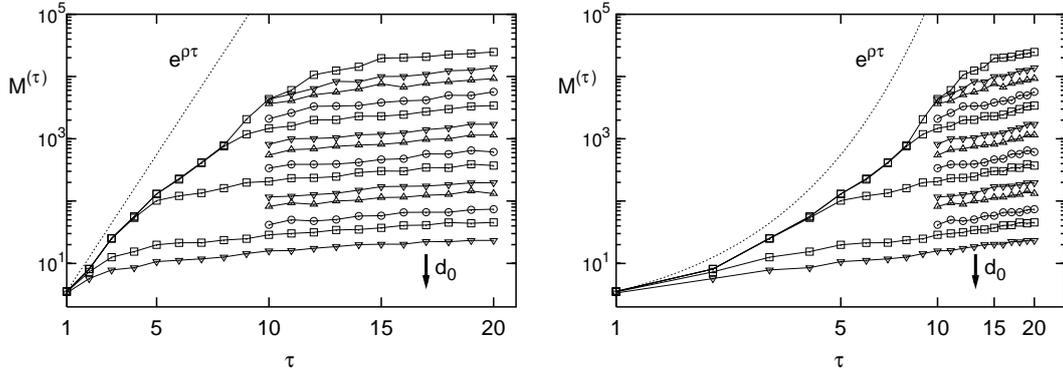


Figure 5.18: Same as Fig. 5.17 for the quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$; $N_{\text{ens}} = 10000$). The curves represent different values of d_0 , from 10^{-4} (top) to 2×10^{-1} (bottom).

We have shown a systematic dependence of $M^{(\tau)}$, representing the worst-case error growth for any initial state, on prediction time horizon τ . For short times we find an exponential increase of $M^{(\tau)}$ consistent with the maximum growth exponents. For long prediction times $M^{(\tau)}$ grows according to a power law, $M^{(\tau)} \sim \tau^\alpha$. The scaling exponent α depends systematically on the initial error size d_0 , again according to a power law.

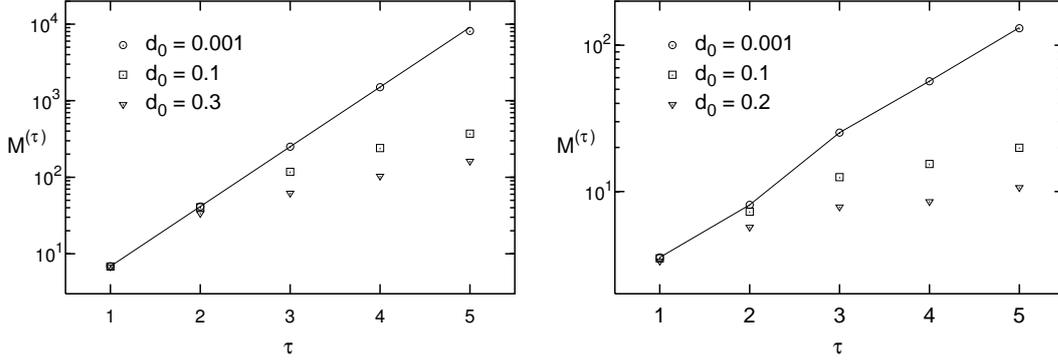


Figure 5.19: Comparison of $M(\tau)$ (symbols) and $\exp(\rho(\tau)\tau)$ (solid line) for the standard map ($K = 4.2$; left) and the quasiperiodicity map ($c_1 = 0.42, c_2 = 0.3, c_3 = 1.6$; right).

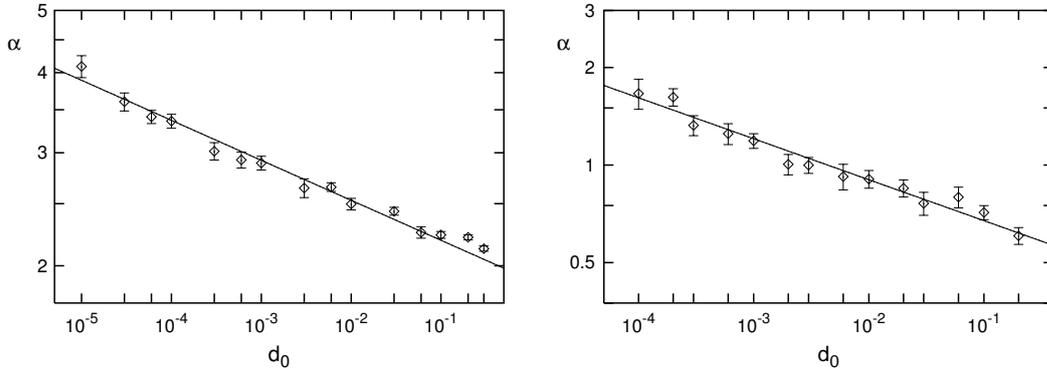


Figure 5.20: Double-logarithmic plot of scaling exponent α of $M(\tau)$ for large τ . Results for the standard map ($K = 4.2$; left) and the quasiperiodicity map ($c_1 = 0.42, c_2 = 0.3, c_3 = 1.6$; right), based on the right parts of Figs. 5.17 and 5.18. Linear least-squares fits were made for $\tau \geq 10$, except for the quasiperiodicity map at $d_0 = 10^{-4}$, where only $\tau \geq 12$ were used; error bars represent one standard deviation.

5.5 Minimal ensemble size needed

In this section we present results that are useful for finding a minimal ensemble size. Here, we focus on the more chaotic parameter sets of the standard map ($K = 4.2$) and the quasiperiodicity map ($c_1 = 0.42, c_2 = 0.3, c_3 = 1.6$). Additionally, we also present results for the dissipative Hénon map.

In order to define a minimal size $N_{\text{ens}}^{\text{min}}$ of an ensemble, one needs a quanti-

5.5. Minimal ensemble size needed

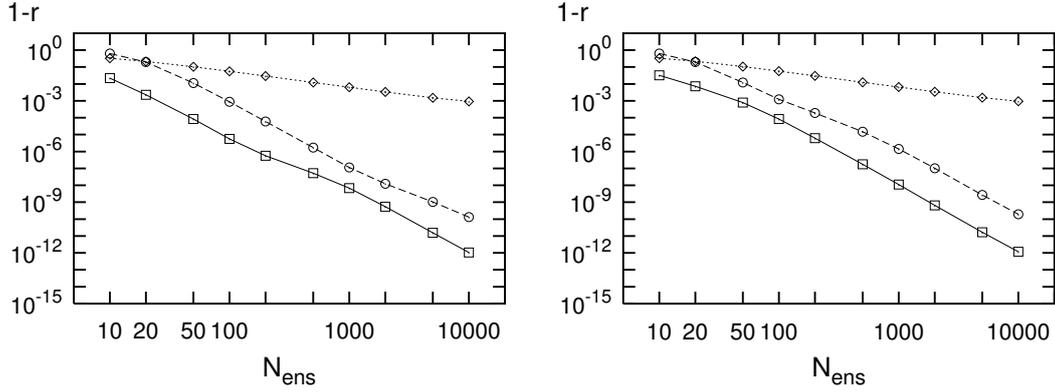


Figure 5.21: $1 - r$ for the standard map ($K = 0.6$); $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right): Convergence of data with ensemble size as calculated by linear correlation r . Points are for $m_{\max}^{(\tau)}$ (squares), $m_{\text{avg}}^{(\tau)}$ (diamonds) and $m_{\min}^{(\tau)}$ (circles). Lines connecting the points are only guides to the eye.

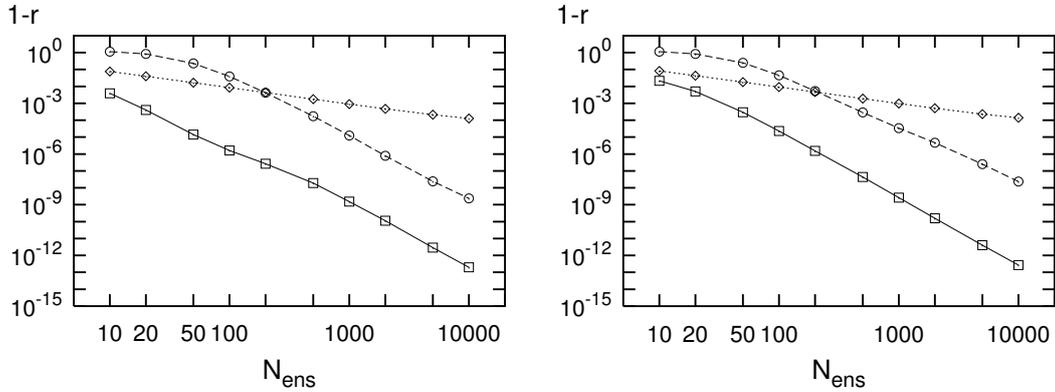


Figure 5.22: $1 - r$ for the standard map at $K = 4.2$; $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right).

tative measure of how close the results for a given ensemble size N_{ens} are to the hypothetical results for an infinite ensemble. This number of ensemble members $N_{\text{ens}}^{\text{min}}$ will of course depend on the prediction time τ and the size of the initial error d_0 . In the following, we will use the largest ensemble ($N_{\text{ens}} = 20000$) as a reference, because it is our best approximation to an infinite ensemble. Results for smaller ensembles are judged by how close they are to the reference. Measurement of this closeness or “convergence” of ensembles is done in two ways, one using the full spatially-resolved data, and one using histograms.

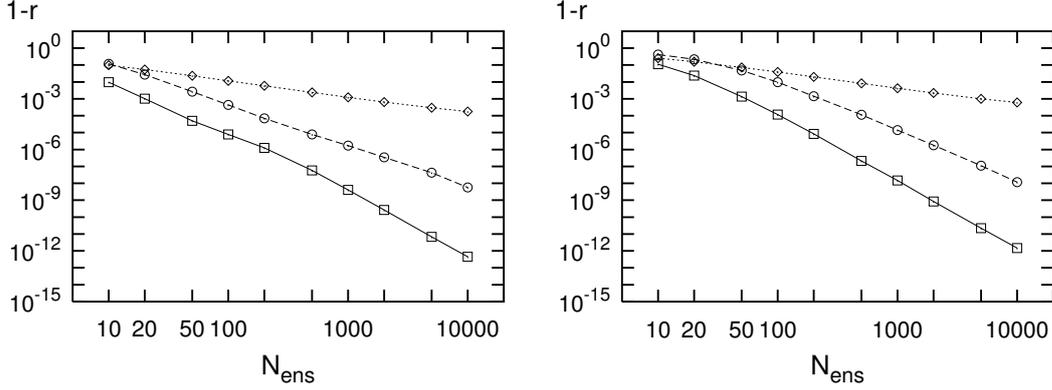


Figure 5.23: $1 - r$ for the quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$); $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right).

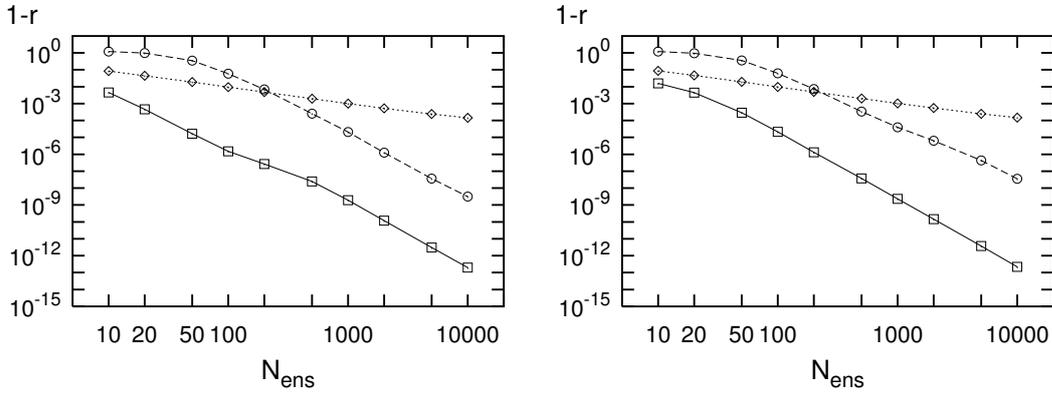


Figure 5.24: $1 - r$ for the Hénon map ($a = 1.4$, $b = 0.3$); $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right).

The first approach looks at each initial state x_0 and compares results for, say, $m_{\max}^{(\tau), N_{\text{ens}}}(x_0)$ obtained by using different ensemble sizes with the corresponding value for the reference ensemble, $\hat{m}_{\max}^{(\tau)}(x_0) = m_{\max}^{(\tau), N_{\text{ens}}=20000}(x_0)$.

Measurement of convergence of the ensemble is done by calculating the linear correlation r between $m_{\max}^{(\tau), N_{\text{ens}}}(x_0)$ and $\hat{m}_{\max}^{(\tau)}(x_0)$ for all x_0 , giving an ideal value of $r = 1$ for identical values and $r = 0$ for uncorrelated data. Thus, $1 - r$ is the error introduced by using a finite ensemble size. Its dependence on N_{ens} is shown exemplarily in Figs. 5.21 and 5.22 for the standard map at $K = 0.6$ and $K = 4.2$, respectively, in Fig. 5.23 for the quasiperiodicity map and in Fig. 5.24 for the Hénon map. For ensembles that are large enough, a power-law decrease of $1 - r$ can be seen in all cases. For $m_{\text{avg}}^{(\tau)}$, this is valid over the entire range

5.5. Minimal ensemble size needed

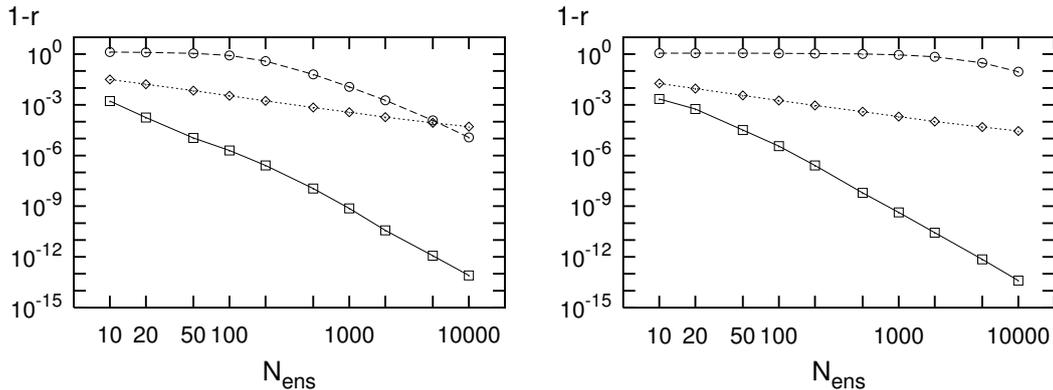


Figure 5.25: $1 - r$ for the standard map ($K = 4.2$); $d_0 = 0.001$, $\tau = 2$ (left) and $\tau = 4$ (right).

of N_{ens} , but the error is typically orders of magnitude larger than the errors for $m_{\text{min}}^{(\tau)}$ and $m_{\text{max}}^{(\tau)}$. The range over which $m_{\text{min}}^{(\tau)}$ and $m_{\text{max}}^{(\tau)}$ obey the power-law seems to increase with d_0 . Whether this is a general behaviour and how it could be interpreted is not clear to us.

If the ensemble is very small, the error of $m_{\text{min}}^{(\tau)}$ can remain almost constant at a high value, especially in the case of the standard map at $K = 4.2$ and the Hénon map. The range of ensemble sizes for which this behaviour occurs extends towards increasing N_{ens} for larger times τ (s. Figs. 5.25–5.27), whereas it depends only very weakly on d_0 . In the quasiperiodicity map, this almost constant error can only be found for higher values of $\tau \geq 4$ and for rather large $d_0 \geq 0.1$. In the Hénon map this behaviour is most pronounced.

The error of $m_{\text{max}}^{(\tau)}$ shows a kink at intermediate values of N_{ens} for small values of d_0 and τ in both systems. This kink separates two domains with an approximately linear dependence of $1 - r$ on N_{ens} in the double logarithmic plot. We do not know where this kink stems from or how it should be interpreted.

The effect of longer prediction times τ can be seen exemplarily by comparing Figs. 5.22 and 5.25 for the standard map, Figs. 5.23 and 5.26 for the quasiperiodicity map and Figs. 5.24 and 5.27 for the Hénon map. When the initial errors are small (Figs. 5.25 and 5.27), only $m_{\text{min}}^{(\tau)}$ depends considerably on τ . This is most likely due to the dependence of the error growth factors m on the initial direction α_0 . If the smallest value $m_{\text{min}}^{(\tau)}$ is attained only for a narrow interval of α_0 , small ensembles can be expected to produce a bad estimate of $m_{\text{min}}^{(\tau)}$. When initial errors are not small (Fig. 5.26), $m_{\text{max}}^{(\tau)}$ also shows a strong dependence on τ . This comes from the fact that longer prediction times together with larger initial

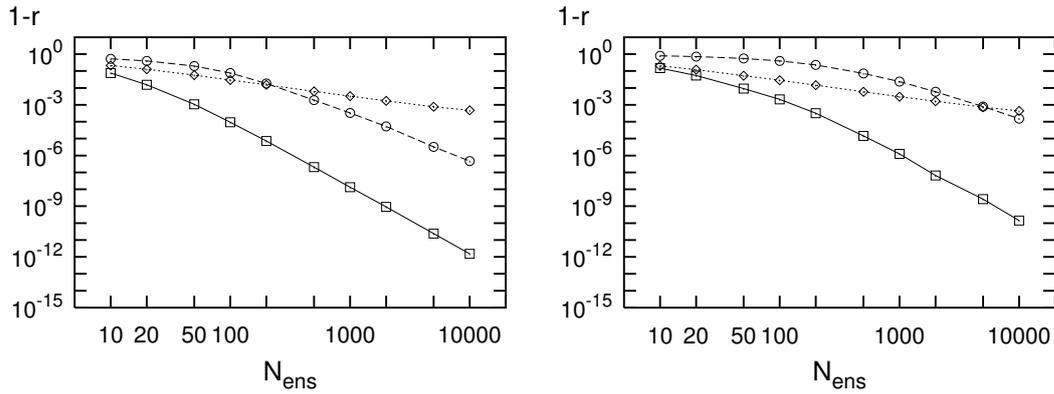


Figure 5.26: $1 - r$ for the quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$); $d_0 = 0.2$, $\tau = 2$ (left) and $\tau = 4$ (right).

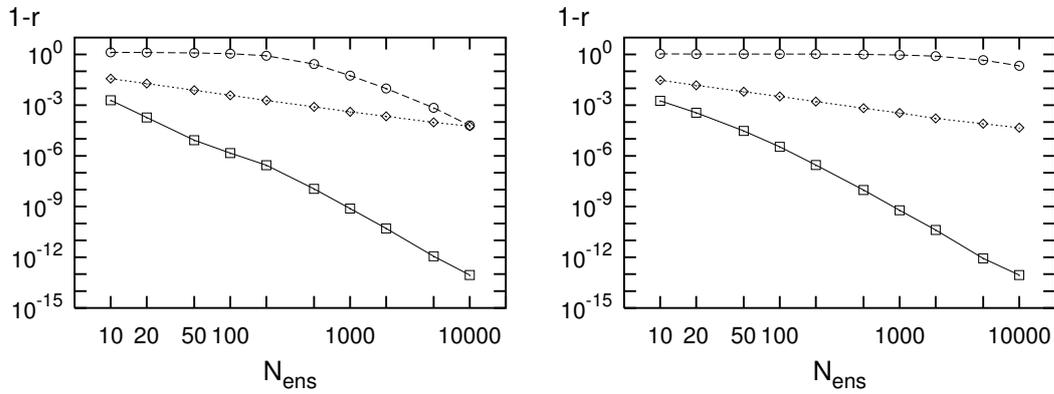


Figure 5.27: $1 - r$ for the Hénon map ($a = 1.4$, $b = 0.3$); $d_0 = 0.001$, $\tau = 2$ (left) and $\tau = 4$ (right).

errors create many more peaks in the distribution of error growth factors m over direction α_0 (not shown). Therefore it becomes more difficult to find the global maximum, i. e. $m_{\max}^{(\tau)}$. Preliminary investigations have been done regarding the distributions of error growth factors over initial error direction α_0 . These could not be pursued further, however, for lack of time.

The power-law behaviour mentioned above, obeyed for large N_{ens} , can be expressed as a scaling exponent γ :

$$1 - r \sim N_{\text{ens}}^{\gamma}. \quad (5.6)$$

This exponent is extracted from the data shown in Figs. 5.22–5.27, for example, by applying a linear least squares fit over a suitable interval of N_{ens} . The values

5.5. Minimal ensemble size needed

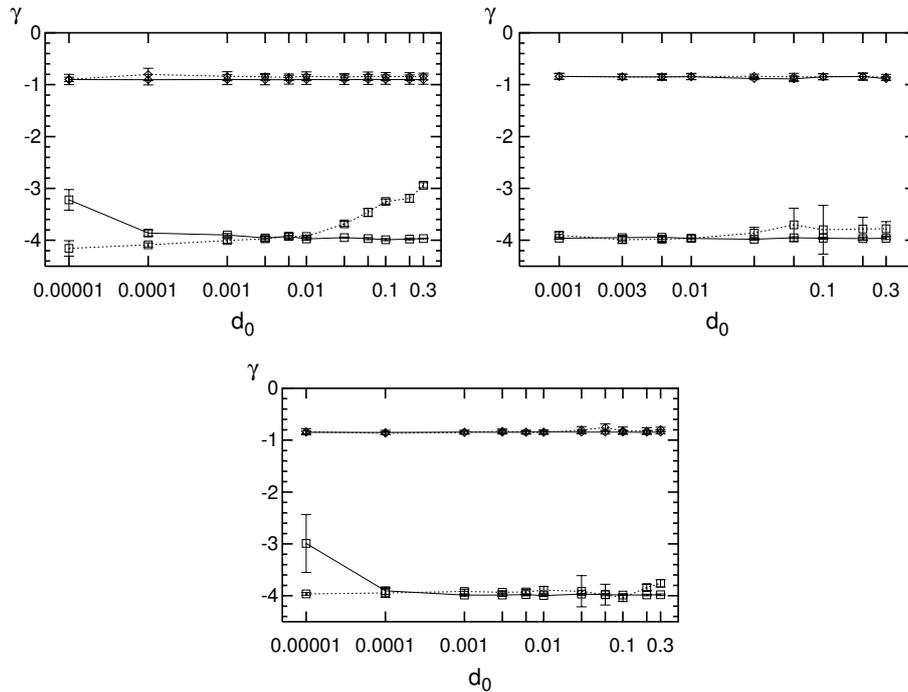


Figure 5.28: Scaling exponent γ of $1 - r$ for large N_{ens} . Top left: standard map ($K = 4.2$); top right: quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$); bottom: Hénon map ($a = 1.4$, $b = 0.3$). The points show $m_{\text{max}}^{(\tau)}$ (squares) and $m_{\text{avg}}^{(\tau)}$ (diamonds), solid lines connect points for $\tau = 1$, dotted lines stand for $\tau = 5$. Error bars give an indication of the validity of the linear fit.

of γ for $m_{\text{avg}}^{(\tau)}$ and $m_{\text{max}}^{(\tau)}$ are plotted in Fig. 5.28 for $\tau = 1$ and $\tau = 5$. We find that for $m_{\text{avg}}^{(\tau)}$, the scaling exponent is constant and close to $\gamma = -1$ for a wide range of initial errors d_0 and times τ . For $m_{\text{max}}^{(\tau)}$, however, the behaviour for larger times is different, showing a possible increase for quite large errors d_0 . This can be clearly seen for the standard map, but only guessed for the Hénon map, whereas the quasiperiodicity map does not show such a clear increase. For small prediction times, the scaling exponent of $m_{\text{max}}^{(\tau)}$ remains close to $\gamma = -4$ for almost the whole range of d_0 . One has to keep in mind that fitting a straight line to the data is not always straightforward if d_0 and τ are too large. The error bars in Fig. 5.28 reflect this uncertainty.

The second approach adopted to measure “convergence” of ensembles does not use the spatial information contained in the data, i. e. values of $m_{\text{max}}^{(\tau)}$, for example, are not associated with the state to which they originally belong. Instead,

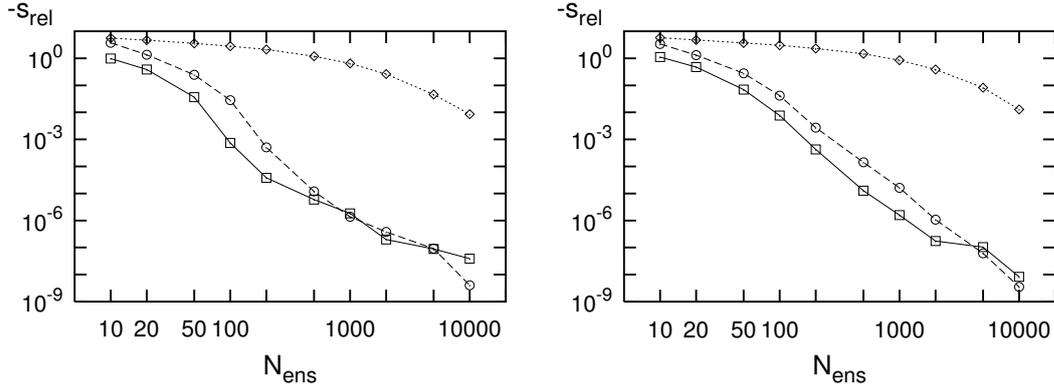


Figure 5.29: Relative entropy s_{rel} for the standard map ($K = 0.6$); $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right). Points are for $m_{\text{max}}^{(\tau)}$ (squares), $m_{\text{avg}}^{(\tau)}$ (diamonds) and $m_{\text{min}}^{(\tau)}$ (circles).

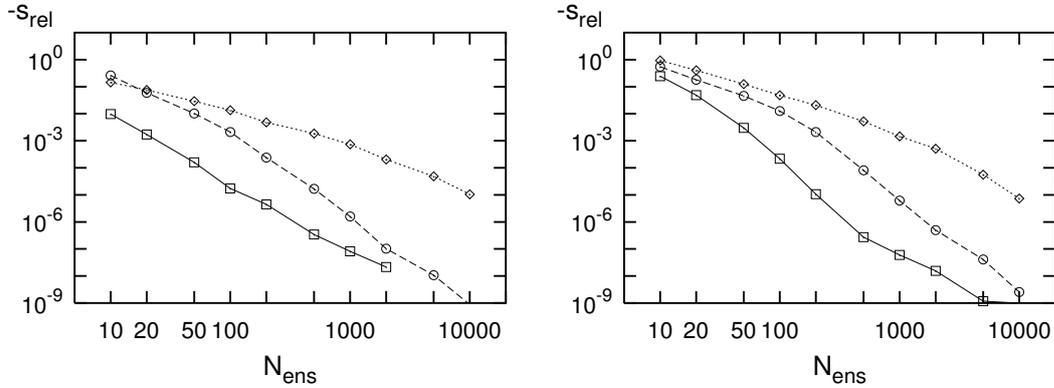


Figure 5.30: Relative entropy s_{rel} for the quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$); $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right).

the data is stored in histograms. These can then be compared for different ensemble sizes N_{ens} , where the largest ensemble is again used as a reference. In order to determine the convergence of the ensemble, we calculate the relative entropy s_{rel} of the histograms. This measure can be interpreted as the amount of information lost due to the finiteness of the ensemble. It is widely used and has a sound theoretical basis (Honerkamp, 1994). For identical distributions a value of $s_{\text{rel}} = 0$ is obtained, for differing distributions $s_{\text{rel}} < 0$.

The results for the standard map are shown in Fig. 5.29 (cf. Fig. 5.21 for the corresponding linear correlation data). The analogous graphs for the quasiperiodicity map are plotted in Fig. 5.30 (cf. Fig. 5.23) and for the Hénon map in

5.5. Minimal ensemble size needed

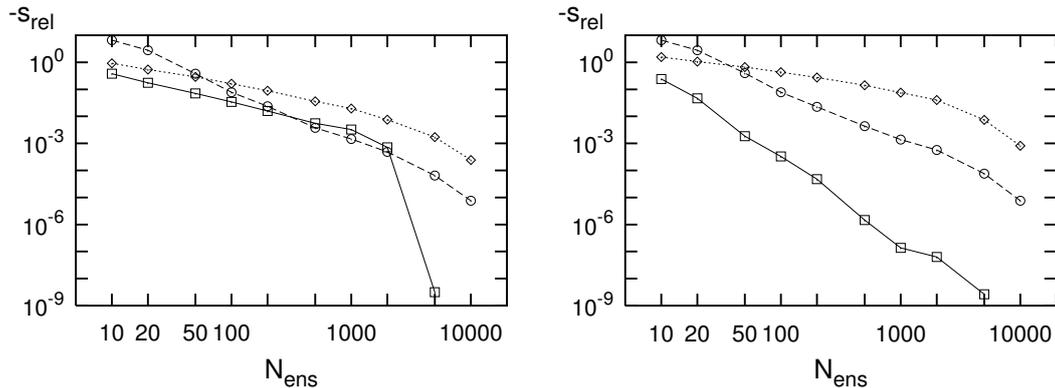


Figure 5.31: Relative entropy s_{rel} for the Hénon map ($a = 1.4$, $b = 0.3$); $\tau = 1$, $d_0 = 0.001$ (left) and $d_0 = 0.2$ (right). The sudden drop for $m_{\text{max}}^{(\tau)}$ in the left plot is explained by the distributions for $N_{\text{ens}} = 5000$ and $N_{\text{ens}} = 20000$ (the reference) being almost equal already.

Fig. 5.31 (cf. Fig. 5.24). Overall, the behaviour of s_{rel} is less systematic than that of the linear correlation coefficient. The power law-behaviour found for the latter can be seen here, too, but not as clearly (especially for small errors in the standard map, Fig. 5.29). This fact can be attributed to the loss of information that occurs when one puts all data in one histogram and forgets about the spatial information. For the relative entropy, the convergence of $m_{\text{min}}^{(\tau)}$ is almost as good as the one of $m_{\text{max}}^{(\tau)}$, but this holds only for sufficiently small d_0 and τ .

In addition, χ^2 -statistics were computed from the histograms. These turned out to behave in the same qualitative way as the relative entropy. Therefore, they are considered no more useful for our purpose than the measures already introduced.

For practical purposes, the asymptotic behaviour of the error growth factors with respect to ensemble size may not be very helpful. What is helpful, however, is the actual ensemble size needed for the error $1 - r$ to drop below a predefined tolerance, say, $\epsilon = 10^{-3}$. Defining this error again as the deviation from the case $N_{\text{ens}} = 20000$, we arrive at the values shown in Figs. 5.32 and 5.33 for $m_{\text{max}}^{(\tau)}$. For the standard map, one finds indication of yet another power law,

$$N_{\text{ens}}^{\text{min}} \sim d_0^\beta, \quad (5.7)$$

for larger initial errors d_0 and small to intermediate prediction times τ . The result for the quasiperiodicity map also shows such an increase, whereas the Hénon map shows almost no change for higher values of d_0 and no clear dependence on τ .

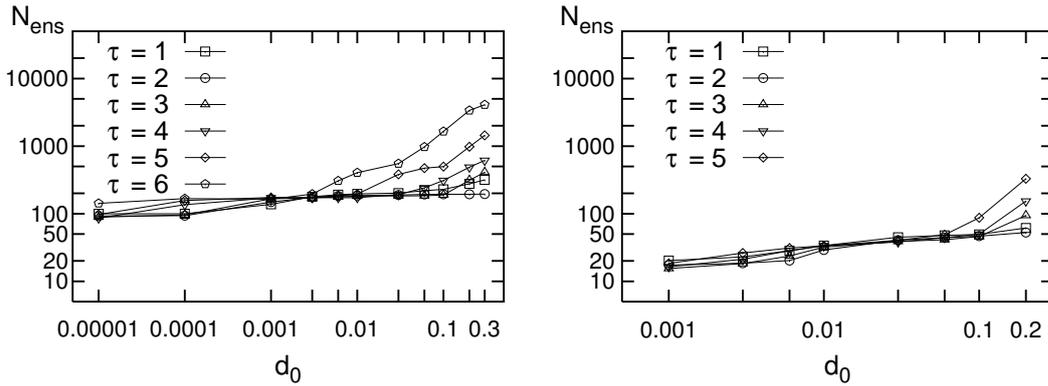


Figure 5.32: Minimal ensemble size N_{ens}^{\min} needed for the error in $m_{\text{max}}^{(\tau)}$ to drop below a given tolerance ϵ , calculated using linear interpolation in the double-logarithmic plots like Figs. 5.21–5.26. Left: standard map at $K = 4.2$ with $\epsilon = 10^{-6}$; right: quasiperiodicity map ($c_1 = 0.42$, $c_2 = 0.3$, $c_3 = 1.6$) with $\epsilon = 10^{-3}$.

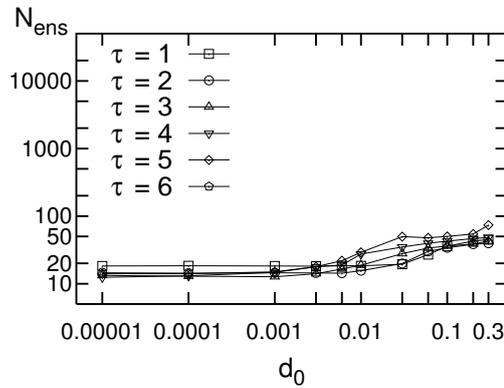


Figure 5.33: As Fig. 5.32, but for the Hénon map ($a = 1.4$, $b = 0.3$) with $\epsilon = 10^{-3}$.

Even smaller error tolerances, like $\epsilon = 10^{-9}$, were also considered, but these eventually lead to a levelling off at the largest ensemble size used. In all cases considered we found a region of d_0 -values indicative of the power law in Eq. 5.7. From the investigations done, a simple dependence of β on τ or ϵ could not be established.

We have used various statistics to compare error growth factors for different ensemble sizes. The linear correlation coefficient r produced the best results, enabling us to obtain a power law behaviour for the “convergence” of $m_{\text{max}}^{(\tau)}$ and

5.5. Minimal ensemble size needed

$m_{\text{avg}}^{(\tau)}$ for larger ensemble sizes. The scaling exponent γ remains almost constant at $\gamma \approx -1$ for $m_{\text{avg}}^{(\tau)}$, $\gamma \approx -4$ for $m_{\text{max}}^{(\tau)}$, as long as the prediction time τ and the initial error size d_0 are not too large. Results have also been shown for a direct calculation of the minimal ensemble size needed for convergence up to a given tolerance ϵ , with a significant dependence on prediction time only present for large initial errors.

Chapter 6

Summary

We would like to end this treatise with a summary of the main results, together with suggestions for possible future extensions of this work.

We have introduced local measures to characterise heterogeneity in state space of the dynamics of conservative and dissipative chaotic systems. Our aim was twofold: We want to quantify the heterogeneity in state space and we want to assess the predictability of the system in a given state.

In chapter 3 we have investigated the distribution of finite time Lyapunov exponents of a trajectory in a chaotic layer of a conservative map. The distributions of finite time Lyapunov exponents exhibit multiple maxima when the finite piece of trajectory is long enough. We have shown that this multimodality is a result of the particular structure of the dynamics in state space. There exists a hierarchy of islands embedded in the chaotic sea, or, more specifically, in the chaotic layer under consideration. This hierarchy of islands is reflected in the distribution of finite time Lyapunov exponents. While the motion of the trajectories around the primary islands is highly chaotic, it becomes less chaotic as the order of the island in the hierarchy increases. The higher the order of the island, the closer is the largest Lyapunov exponent to zero. Using a simple but easily implementable method, we have separated pieces of the trajectory moving in layers of different orders. The individual distributions of finite time Lyapunov exponents for layers of different orders show clearly separated peaks. The overall motion is characterised by jumps between the layers of different orders, which leads to an overlap of all individual distributions of finite time Lyapunov exponents. As a result we find multimodality where each extremum corresponds to the motion in a particular chaotic layer.

Recently, Szezech et al. (2005) have reported their finding of a bimodal distribution of “finite time Lyapunov exponents” (FTLEs). It has to be stressed that their FTLEs are what we call maximum growth exponents. It can be argued

that one of their figures indicates that there are not only two peaks in the FTLE distribution. One of the peaks, the one closer to zero, could in fact consist of multiple peaks as a result of motion in layers of higher order. We think that our study and the work of Szezech et al. (2005) complement each other in the following sense: Their focus was on distinguishing regions with and without stickiness by means of the distribution of (their) FTLE. We, on the other hand, set out to explain the in fact multimodal distribution as stemming from islands of different orders, using finite time Lyapunov exponents for a high number of steps.

Future research on this topic could investigate how the distributions of finite time Lyapunov exponents in layers of different orders are affected by parameters like: strength of the nonlinearity, length of the finite pieces of trajectory, and possibly also the actual island structure (number of islands at the various orders). The statistics will of course be of particular importance for such studies, especially in order to establish quantitative results.

The results of chapter 4 relate the predictability (quantified via local exponents) of dissipative and conservative dynamical systems, to the orientation of its stable and unstable manifolds. An interrelation of points of homoclinic tangency and regions of enhanced predictability corresponding to low values of the finite time exponents $\lambda^{(\tau)}$ was found for the dissipative Hénon map and for the conservative standard map, even though it is more difficult to identify in the latter case.

When the trajectory comes close to a point of homoclinic tangency \vec{x}_{HT} , the orientation of a perturbation can no longer point away from the stable direction at a large angle. The reason for this is that the unstable manifold has to approach the stable one due to the definition of a homoclinic tangency. Instead of pointing away from it, the perturbation has to become nearly aligned with the direction of the stable manifold, if the trajectory is close enough to \vec{x}_{HT} . This argument in favour of a connection between homoclinic tangencies and regions of enhanced predictability applies, in principle, to heteroclinic tangencies as well as to homoclinic ones. Since the calculation of the former is much more demanding, we restricted our investigation to the latter. The size of a region of enhanced predictability around a homoclinic tangency was found to shrink exponentially with increasing time horizon of prediction.

In both systems under investigation, there was no connection with extraordinarily small values of the maximum growth exponents $\rho^{(\tau)}$. This finding suggests that maximum growth exponents are not an ideal measure to quantify predictability, at least when comparing with other properties of the dynamical system that depend on past states. Still, they are widely used to quantify predictability in atmospheric science and to determine low-dimensional approximations (Farrell

and Ioannou, 2001). This practice is justified by the fact that observational data is repeatedly assimilated into the state of the system, i. e., the state is changed in order to better agree with the new measurements of the actual state of the physical system. Therefore, there is not enough time for perturbations to point along the globally most expanding direction $\vec{l}_1(\vec{x})$, because data assimilation can make the state of the system jump from one trajectory to another. The finite time Lyapunov exponents are not an adequate measure of predictability in this case, but the maximum growth exponents are. The latter are completely defined by a finite piece of trajectory that one can choose to be as short as is desirable. Furthermore, the largest maximum growth exponents $\rho_1^{(\tau)}$ compare very well with the maximum growth factors $m_{\max}^{(\tau)}$ of finite errors for small initial errors and short prediction times.

As an important point we note again that the characterisation of regions of enhanced predictability by homoclinic tangencies is independent of the particular coordinate system used. While homoclinic tangencies are invariant under coordinate transformations, finite time Lyapunov exponents and maximum growth exponents are not. The relation between homoclinic tangencies and regions of enhanced predictability can be used to identify those regions as invariants of the corresponding dynamical system. However, it has to be kept in mind that homoclinic tangencies are not the reason for the appearance of regions of enhanced predictability. The latter can also appear in hyperbolic systems, where no homoclinic tangencies are present.

A connection between regions of enhanced predictability and homoclinic tangencies was examined, but could not be established for the maximum growth exponents ρ . The opposite extremes – nearly orthogonal intersections of stable and unstable manifolds and deteriorated predictability – are found to be not related to each other either, regardless of the system under consideration and the kind of local exponents used.

In the case of the conservative system, there is no clear large-scale distinction between regions of extremely high and extremely low values of the finite time exponents. This means that good or bad predictability can occur for points in state space that are very close to each other. Therefore, a simple quantification of predictability relying on local exponents in state space cannot be expected to be feasible. However, there seem to be regions with better predictability than their surroundings. Homoclinic tangencies were found in such regions, following the same pattern.

Another point to consider is the following: The attractor of the system can be considerably altered by perturbations like numerical errors occurring near homoclinic tangencies, as the trajectory can be driven away from the attractor (Jaeger and Kantz, 1997). This raises the question why enhanced predictability

is nevertheless found near homoclinic tangencies. It might be due to the fact that the effect of noise was studied using the full nonlinear system, while the local exponents are defined using the linear approximation. Further studies like an examination of the behaviour of perturbations governed by the nonlinear system could be conducted in the future to better understand the relation between perturbations, homoclinic tangencies and predictability.

The linear approximation inherent in using local exponents to measure trajectory divergence must break down in practice sooner or later, i. e. once the error reaches a certain size. Then one has to calculate the error growth using the nonlinear system.

In chapter 5 we studied the growth of non-infinitesimal errors in three exemplary chaotic systems, the standard map, the quasiperiodicity map and the Hénon map. Non-infinitesimal, finite errors cannot be treated with local exponents. Therefore we followed ensembles of size N_{ens} for a time τ , with all ensemble members initialised with a fixed error magnitude of d_0 and uniformly random direction. Each ensemble member evolves according to the full nonlinear system. The final error sizes divided by d_0 yield the error growth factors $m^{(\tau)}$ for each ensemble member, from which one can compute quantities like the worst-case error growth $m_{\text{max}}^{(\tau)}$, for example.

Comparing the maximum error growth factor $m_{\text{max}}^{(\tau)}$ with local exponents, we have found that the data of the largest maximum growth exponents $\rho_1^{(\tau)}$ compares very well with $m_{\text{max}}^{(\tau)}$ for small initial errors d_0 and short prediction times τ . The reason for this is that the maximum growth exponents are a purely local quantity, in the sense that they depend only on the system states visited during the time span τ . Since they are defined through the singular values, there is a clear correspondence between $m_{\text{max}}^{(\tau)}$ and the largest maximum growth exponent ρ_1 for the same initial state. The finite time Lyapunov exponents, however, show many features not present in the $m_{\text{max}}^{(\tau)}$ data, since they indirectly rely on the past of the system via the Lyapunov vectors \vec{l}_i . This striking difference makes it clear that the choice of local exponents is very important, especially considering the results of chapter 4. Depending on one's interest, one or the other type of local exponents may be the better choice.

For the error growth of the worst case, i. e. the largest error growth factor $M^{(\tau)}$ for all initial conditions, we have demonstrated a systematic dependence on the prediction time τ . For small times, there is an exponential increase according to the largest maximum growth exponent $\rho_{1,\text{max}}^{(\tau)}$. This does no longer hold if the prediction time is too long. Then one observes a crossover to a regime of power-law growth, $M^{(\tau)} \sim \tau^\alpha$. The scaling exponent α of this power law has been found to depend on the size d_0 of initial errors, presumably according to a

power law as well.

In practice, it is desirable to keep the computational cost as low as possible. In the context of ensemble studies, doing so requires one to use a minimal ensemble, one that extracts the required information with the fewest members. To find such a minimal ensemble size $N_{\text{ens}}^{\text{min}}$ needed for reliable statements about predictability, we have studied various measures, namely the linear correlation coefficient r , the relative entropy s_{rel} and the χ^2 -statistics. Here, r was computed using spatially resolved data, comparing results for different ensemble sizes at each initial point separately. The other two measures, however, were using histograms for each ensemble size. This means that only the overall frequency of a certain range of error growth factor values is used; to which initial point these values are connected is not used. With these measures, the difference between a given ensemble size and the largest ensemble (used as a reference) is quantified to understand how good the approximation of the given, possibly small ensemble is.

For the linear correlation coefficient, a power-law has been shown to exist for large ensemble sizes, $1 - r \sim N_{\text{ens}}^\gamma$. The scaling exponent γ is constant over a wide range of initial error sizes d_0 and prediction times τ , at values of $\gamma \approx -1$ for the mean error growth $m_{\text{avg}}^{(\tau)}$ and $\gamma \approx -4$ for the maximum error growth $m_{\text{max}}^{(\tau)}$. In the latter case a dependence of γ on d_0 only arises for τ and d_0 large enough.

These results for γ deal with the asymptotic behaviour of error growth factors for increasing ensemble size. However, the absolute value of the error, the error made by using a rather small ensemble, need not directly relate to the value of γ . For the cases studied herein, the maximum error growth $m_{\text{max}}^{(\tau)}$ is the one converging most quickly. We have shown that it can reach very small values even for intermediate ensemble sizes. Or, conversely, for a given error tolerance ϵ , the minimal ensemble size $N_{\text{ens}}^{\text{min}}$ is not very large. This minimal ensemble size was calculated from the double-logarithmic plots of $1 - r$ vs. N_{ens} as the value of N_{ens} where $1 - r$ drops below ϵ . When trying to use this procedure in practice, a difficulty will probably be the determination of the reference data, which we produced using a very large ensemble. It may be possible, though, to “guess” the general form of the reference histogram used with the relative entropy. To validate or invalidate the guess one can use a few runs with large ensembles that are only followed for a short period of time. In any case, the fact that the maximum error growth $m_{\text{max}}^{(\tau)}$ converges most quickly remains encouraging, since one can expect the worst case to be handled better than the average or minimal error growth.

It is quite obvious that a method using the full spatially resolved data is superior to others based on histograms. However, while using binned data in a histogram involves a loss of information, it has its advantage as well: One does not need to store the same (possibly huge) amount of data as in the first case. So, from a practical point of view, it may be preferable – or even unavoidable – to use

binned data. Even if one is forced to do an analysis without spatially resolved data, one can find the same general variation of error growth with ensemble size as in the theoretically preferable case, as we have shown using the relative entropy. Conclusions regarding a systematic behaviour are much more difficult to draw on the basis of histograms alone, however.

We stress that the determination of a minimal ensemble size for a given purpose and, consequently, for a given acceptable error tolerance is always system-dependent. In practice, the choice of ensemble size is often based on feasibility and experience. There does not seem to be any hard and fast rule about this. Still, the behaviour of the scaling exponent γ is a general one that can be expected to be found in other systems as well. Furthermore, minimal ensemble sizes needed for a given error tolerance can be estimated using a few ensembles at small prediction times τ and typical error values d_0 . These findings can be of help in deciding on the usefulness of increasing the ensemble size.

The dependence of error growth on the initial direction α_0 of the error was studied briefly in a qualitative way. The main conclusion that can be drawn is that things get very involved as d_0 and τ increase. Many maxima of $m^{(\tau)}$ (as a function of α_0) can appear for each initial point x_0 , making it difficult to find the global maximum, i. e. $m_{\max}^{(\tau)}(x_0)$. A possible direction of future work could be to study the dependence of $\alpha_{\max}^{(\tau)}(x_0)$, the direction leading to $m_{\max}^{(\tau)}(x_0)$, on N_{ens} . Another possibility is to quantify the “ruggedness” of the curve $m_{\max}^{(\tau)}(\alpha_0)$ in a suitable way, like simply calculating what proportion of initial angles α_0 lead to a value of $m^{(\tau)} \geq 0.9m_{\max}^{(\tau)}$. Of course such studies will become infeasible for medium dimensional systems since huge ensembles would be needed to more or less cover the hypersphere that was just a circle in our simple, two-dimensional case.

In conclusion, one has to face the fact that it is not possible to characterise the predictability of a complex, chaotic system by a simple number, or even a few numbers. It is possible, however, to approach this goal in a certain sense.

Firstly, one has to be aware of the huge difference that the choice of the quantity or method used to quantify predictability makes. Finite time Lyapunov exponents can be the method of choice when trying to link predictability with dynamical properties of the system, since the past states of the system play a role. Maximum growth exponents are better when one is concerned with the growth of small errors that are unrelated to the past. What makes this kind of exponents particularly appealing is that they approximate the growth of (small) finite size errors very well. Ensembles of finite size errors should be used when the initial errors are too large to be adequately handled by local exponents.

Secondly, one has to consider the scale (in state space) on which changes

in predictability (as quantified by the method chosen) take place. As one expects from the continuous stretching and folding, the structures get smaller and smaller as the prediction time τ increases. Therefore it may be possible to detect some regions in state space that exhibit enhanced predictability for short prediction times and accordingly divide the state space into regions with a certain predictability. We think that even this will be difficult, however, in the case of high-dimensional systems.

Thirdly, if one chooses to use ensembles one should consider spending some time and effort on studying the dependence of the results on ensemble size. One can then set the number of ensemble members to the smallest possible value that one can expect to yield reliable results. In this context, an encouraging result of this thesis is the fact that for the worst-case error growth a smaller ensemble is needed than for the average error growth, for example.

We hope that our work is contributing to highlighting the importance of the choice of the adopted method, helping to determine a reasonable ensemble size and motivating future research in this area.

Bibliography

- Abarbanel, H. D. I., Brown, R., and Kennel, M. B.: Lyapunov exponents in chaotic systems: their importance and their evaluation using observed data, *Int. J. Mod. Phys. B*, 5, 1347–1375, 1991.
- Abraham, E. R., Law, C. S., Boyd, P. W., Lavender, S. J., Maldonado, M. T., and Bowie, A. R.: Importance of stirring in the development of an iron-fertilized phytoplankton bloom, *Nature*, 407, 727–730, 2000.
- Abraham, R. and Smale, S.: Nongenericity of Ω -stability, in *Global Analysis I*, *Proc. Symp. Pure Math. A. M. S.*, 14, pp. 5–8, 1968.
- Adler, R. L., Konheim, A. G., and McAndrew, M. H.: Topological entropy, *Trans. Amer. Math. Soc.*, 114, 309–319, 1965.
- Akin, E. and Kolyada, S.: Li-Yorke sensitivity, *Nonlinearity*, 16, 1421–1433, 2003.
- Anosov, D. V.: Geodesic flows on closed Riemannian manifolds with negative curvature, *Proc. Steklov Inst. Math.*, 90, 1–235, 1967.
- Assaf, D. and Gadbois, S.: Definition of Chaos, letter in *Amer. Math. Mon.*, 99, 865, 1992.
- Aurell, E., Bofetta, G., Crisanti, A., Paladin, G., and Vulpiani, A.: Predictability in the large: an extension of the concept of Lyapunov exponent, *J. Phys. A: Math. Gen.*, 30, 1–26, 1997.
- Auslander, J. and Yorke, J. A.: Interval Maps, Factors of Maps, and Chaos, *Tohoku Math. J.*, 32, 177–188, 1980.
- Banks, J., Brooks, J., Cairns, G., Davis, G., and Stacey, P.: On Devaney’s Definition of Chaos, *Amer. Math. Mon.*, 99, 332–334, 1992.
- Barkmeijer, J.: Constructing fast-growing perturbations for the nonlinear regime, *Q. Atmos. Sci.*, 53, 2838–2851, 1997.

- Benettin, G., Galgani, L., Giorgilli, A., and Strelcyn, J.-M.: Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; A method for computing all of them. Part I: Theory., *Meccanica*, 15, 9–20, 1980a.
- Benettin, G., Galgani, L., Giorgilli, A., and Strelcyn, J.-M.: Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; A method for computing all of them. Part II: Numerical application., *Meccanica*, 15, 21–30, 1980b.
- Bowen, R.: ω -limit sets for axiom A diffeomorphisms, *J. Differential Equations*, 18, 333–339, 1975.
- Buizza, R.: Optimal perturbation time evolution and sensitivity of ensemble prediction to perturbation amplitude, *Q. J. R. Meteorol. Soc.*, 121, 1705–1738, 1996.
- Buizza, R., Tribbia, J., Molteni, F., and Palmer, T.: Computation of optimal unstable structures for numerical weather prediction models, *Tellus*, 45A, 388–407, 1993.
- Chirikov, B. V.: A universal instability of many-dimensional oscillator systems, *Phys. Rep.*, 52, 263–379, 1979.
- Dawson, S., Grebogi, C., Sauer, T., and Yorke, J. A.: Obstructions to Shadowing When a Lyapunov Exponent Fluctuates about Zero, *Phys. Rev. Lett.*, 73, 1927–1930, 1994.
- Devaney, R.: A first course in chaotic dynamical systems, Perseus, New York, 1992.
- d’Ovidio, F., López, C., Hernández-García, E., and Fernández, V.: Mixing structures in the Mediterranean Sea from Finite-Size Lyapunov Exponents, *Geophys. Res. Lett.*, 31, L17 203, 2004.
- Eckhardt, B. and Yao, D.: Local Lyapunov exponents in chaotic systems, *Physica D*, 65, 100–108, 1993.
- Eckmann, J.-P. and Ruelle, D.: Ergodic theory of chaos and strange attractors, *Rev. Mod. Phys.*, 57, 617–656, 1985.
- Eckmann, J.-P., Kamphorst, S. O., Ruelle, D., and Ciliberto, S.: Liapunov exponents from time series, *Phys. Rev. A*, 34, 4971–4979, 1986.

- Farmer, J. D., Ott, E., and Yorke, J. A.: The dimension of chaotic attractors, *Physica D*, 7, 153–180, 1983.
- Farrell, B. F. and Ioannou, P. J.: Accurate Low-Dimensional Approximation of the Linear Dynamics of Fluid Flow, *J. Atmos. Sci.*, 58, 2771–2789, 2001.
- Grassberger, P. and Kantz, H.: Generating partitions for the dissipative Hénon map, *Phys. Lett. A*, 113, 235–238, 1985.
- Grassberger, P., Badii, R., and Politi, A.: Scaling laws for invariant measures on hyperbolic and non-hyperbolic attractors, *J. Stat. Phys.*, 51, 135–178, 1988.
- Grebogi, C., Ott, E., and Yorke, J. A.: Chaotic Attractors in Crisis, *Phys. Rev. Lett.*, 48, 1507–1510, 1982.
- Grebogi, C., Ott, E., and Yorke, J. A.: Crises, Sudden Changes in Chaotic Attractors and Chaotic Transients, *Physica D*, 7, 181–200, 1983.
- Grebogi, C., Ott, E., and Yorke, J. A.: Attractors on an N -torus: Quasiperiodicity versus Chaos, *Physica D*, 15, 354–373, 1985.
- Grebogi, C., Ott, E., and Yorke, J. A.: Metamorphoses of Basin Boundaries in Nonlinear Dynamical Systems, *Phys. Rev. Lett.*, 56, 1011–1014, 1986.
- Grebogi, C., Ott, E., and Yorke, J. A.: Basin boundary metamorphoses: Changes in accessible boundary orbits, *Physica D*, 24, 243–262, 1987.
- Greene, J. M. and Kim, J.-S.: The calculation of Lyapunov spectra, *Physica D*, 24, 213–225, 1987.
- Gringorten, I. I.: On the comparison of one or more sets of probability forecasts, *J. Meteorol.*, 15, 283–287, 1958.
- Harle, M. and Feudel, U.: On the relation between predictability and homoclinic tangencies, *Int. J. Bif. Chaos*, 15, 2523–2534, 2005.
- Harle, M. and Feudel, U.: Hierarchy of islands in conservative systems yields multimodal distributions of FTLEs, *Chaos, Solitons & Fractals*, 31, 130–137, 2007.
- Harle, M., Kwasniok, F., and Feudel, U.: Growth of finite errors in ensemble prediction, *Nonlin. Processes Geophys.*, 13, 167–176, 2006.
- Hénon, M.: A two-dimensional mapping with a strange attractor, *Commun. Math. Phys.*, 50, 69–77, 1976.

- Honerkamp, J.: Stochastic Dynamical Systems, Wiley-VCH, New York, 1994.
- Jaeger, L. and Kantz, H.: Homoclinic tangencies and non-normal Jacobians – Effects of noise in nonhyperbolic chaotic systems, *Physica D*, 105, 79–96, 1997.
- Kalnay, E., Corazza, M., and Cai, M.: Are bred vectors the same as Lyapunov vectors?, in AMS Symposium on observations, data assimilation and probabilistic prediction, pp. 173–177, Amer. Meteor. Soc., 2002.
- Knopf, B., Held, H., and Schellnhuber, H. J.: Forced versus coupled dynamics in Earth system modelling and prediction, *Nonlin. Processes Geophys.*, 21, 311–320, 2005.
- Kolyada, S. F.: Li-Yorke sensitivity and other concepts of chaos, *Ukr. Math. J.*, 56, 1242–1257, 2004.
- Koussis, A. D., Lagouvardos, K., Mazi, K., Kotroni, V., Sitzmann, D., Lang, J., Zaiss, H., Buzzi, A., and Malguzzi, P.: Flood Forecasts for Urban Basin with Integrated Hydro-Meteorological Model, *J. Hydrologic Engrg.*, 8, 1–11, 2003.
- Leoncini, X. and Zaslavsky, G. M.: Jets, stickiness, and anomalous transport, *Phys. Rev. E*, 65, 046 216–1–046 216–16, 2002.
- Li, T. Y. and Yorke, J. A.: Period 3 implies chaos, *Am. Math. Mon.*, 82, 985–992, 1975.
- Lichtenberg, A. J. and Leiberman, M. A.: *Regular and Chaotic Dynamics*, Springer, Berlin, second edn., 1992.
- Lichtenberg, A. J., Lieberman, M. A., and Cohen, R. H.: Fermi acceleration revisited, *Physica D*, 1, 291–305, 1980.
- Lorenz, E. N.: Deterministic nonperiodic flow, *J. Atmos. Sci.*, 20, 130–141, 1963.
- Lorenz, E. N.: A study of the predictability of a 28-variable atmospheric model, *Tellus*, 17, 321–333, 1965.
- McConchie, J. A., Toleman, I. E. J., and Hawke, R. M.: Effect of flow regulation on near-bank velocities and sediment transport potential: A case study from the Waikato River, New Zealand, *J. Hydrol. (New Zealand)*, 44, 45–72, 2005.
- Milnor, J.: On the concept of attractor, *Commun. Math. Phys.*, 99, 177–195, 1985.

- Moron, V., Philippon, N., and Fontaine, B.: Skill of Sahel rainfall variability in four atmospheric GCMs forced by prescribed SST, *Geophys. Res. Lett.*, 30, 2221, 2003.
- Nese, J. M.: Quantifying local predictability in phase space, *Physica D*, 35, 237–250, 1989.
- Neufeld, Z., Haynes, H., and Tél, T.: Chaotic mixing induced transitions in reaction-diffusion systems, *Chaos*, 12, 426–438, 2002.
- Nicolis, J. S., Meyer-Kress, G., and Haubs, G.: Non-Uniform Chaotic Dynamics with Implications to Information Processing, *Z. Naturforsch.*, 38a, 1157–1169, 1983.
- Oprocha, P.: Relations between distributional and Devaney chaos, *Chaos*, 16, 033 112–1–033 112–5, 2006.
- Oseledec, V. I.: A multiplicative ergodic theorem: Lyapunov characteristic numbers for dynamical systems, *Trans. Moscow Math. Soc.*, 19, 197–231, 1968.
- Ott, E.: *Chaos in dynamical systems*, Cambridge University Press, Cambridge, 1993.
- Palis, J. and Takens, F.: *Hyperbolicity and Sensitive-Chaotic Dynamics at Homoclinic Bifurcations, Fractal Dimensions and Infinitely Many Attractors*, Cambridge University Press, Cambridge, 1993.
- Palmer, T. N., Buizza, R., Molteni, F., Chen, Y.-Q., and Corti, S.: Singular vectors and the predictability of weather and climate, *Phil. Trans. R. Soc. Lond. A*, 348, 459–475, 1994.
- Patil, D. J., Hunt, B. R., Kalnay, E., Yorke, J. A., and Ott, E.: Local Low Dimensionality of Atmospheric Dynamics, *Phys. Rev. Lett.*, 86, 5878–5881, 2001.
- Pinto, R. and Sartorelli, J.: Homoclinic tangency and chaotic attractor disappearance in a dripping faucet experiment, *Phys. Rev. E*, 61, 342–347, 2000.
- Robert, C., Alligood, K. T., Ott, E., and Yorke, J. A.: Explosions of chaotic sets, *Physica D*, 144, 44–61, 2000.
- Sanders, F.: On Subjective Probability Forecasting, *J. Appl. Meteorol.*, 2, 191–201, 1953.

- Sauer, T., Grebogi, C., and Yorke, J. A.: How Long Do Numerical Chaotic Solutions Remain Valid?, *Phys. Rev. Lett.*, 79, 59–62, 1997.
- Schellnhuber, H. J. and Urbschat, H.: Analytical study of Cantori: Gap structure, initial conditions, and dimensions, *Phys. Rev. A*, 38, 5888–5901, 1988.
- Schellnhuber, H. J., Urbschat, H., and Block, A.: Calculation of “Cantori”, *Phys. Rev. A*, 33, 2856–2858, 1986.
- Schweizer, B. and Smítal, J.: Measure of chaos and a spectral decomposition of dynamical systems of interval, *Trans. Amer. Math. Soc.*, 344, 737–754, 1994.
- Shimada, I. and Nagashima, T.: A numerical approach to ergodic problem of dissipative dynamical systems, *Progr. Theor. Phys.*, 61, 1605–1616, 1979.
- Smith, L. A.: Disentangling Uncertainty and Error: On the Predictability of Nonlinear Systems, in *Nonlinear Dynamics and Statistics*, edited by A. Mees, chap. 2, pp. 31–64, Birkhäuser, Basel, 2000.
- Smith, L. A., Ziehmann, C., and Fraedrich, K.: Uncertainty dynamics and predictability in chaotic systems, *Q. J. R. Meteorol. Soc.*, 125, 2855–2886, 1999.
- Sornette, D. and Zhou, W.-X.: Predictability of large future changes in major financial indices, *Int. J. Forec.*, 22, 153–168, 2006.
- Stoer, J. and Burlisch, R.: *Introduction to numerical analysis*, Springer, New York, 1980.
- Szezech, J. D., Lopes, S. R., and Viana, R. L.: Finite-time Lyapunov spectrum for chaotic orbits of non-integrable Hamiltonian systems, *Phys. Lett. A*, 335, 394–401, 2005.
- Tang, X. Z. and Boozer, A. H.: Finite time Lyapunov exponent and advection-diffusion equation, *Physica D*, 95, 283–305, 1996.
- Thomson, M. C., Doblas-Reyes, F. J., Mason, S. J., Hagedorn, R., Connor, S. J., Phindela, T., Morse, A. P., and Palmer, T. N.: Malaria early warnings based on seasonal climate forecasts from multi-model ensembles, *Nature*, 439, 576–579, 2006.
- Toth, Z. and Kalnay, E.: Ensemble forecasting at NMC: the generation of perturbations, *Bull. Amer. Meteorol. Soc.*, 74, 2317–2330, 1993.
- Toth, Z. and Kalnay, E.: Ensemble Forecasting at NCEP and the Breeding Method, *Mon. Wea. Rev.*, 125, 3297–3319, 1997.

Bibliography

Zaslavsky, G. M.: *Physics of chaos in Hamiltonian systems*, Imperial College Press, London, 1998.

Zaslavsky, G. M., Stevens, D., and Weitzner, H.: Self-similar transport in incomplete chaos, *Phys. Rev. E*, 48, 1683–1694, 1993.

Ziehmann, C.: Skill prediction of local weather forecasts based on the ECMWF ensemble, *Nonlin. Processes Geophys.*, 8, 419–428, 2001.

Ziehmann, C., Smith, L. A., and Kurths, J.: The bootstrap and Lyapunov exponents in deterministic chaos, *Physica D*, 126, 49–59, 1999.

Ziehmann, C., Smith, L. A., and Kurths, J.: Localized Lyapunov exponents and the prediction of predictability, *Phys. Lett. A*, 271, 237–251, 2000.

Bibliography

Acknowledgements

My first thanks go to my supervisor Ulrike Feudel for welcoming me to her (then still small) research group and of course for guiding and supporting my work. Not only her knowledge and insights did she share, but also her great enthusiasm. Looking back today, it feels like her encouragement may have been as necessary as her help in scientific matters.

I also sincerely thank GKSS for providing the opportunity to work on such a fascinating topic. In particular, I would like to express my gratitude to Hans von Storch and Ralf Weisse, who have helped me, encouraged me and allowed me great freedom regarding the direction of my research.

This work has also profited a lot from discussions with other scientists, especially Celso Grebogi, Jürgen Kurths, Antonio Politi, Leonard A. Smith and Christine Ziehmann. I am very thankful for their time, opinion and ideas.

The members of Ulrike's group have also helped me a lot. Be it problems with computers, printers, operating systems, programs, programmings, typesettings, whatever – there was always someone ready to help.

Quite many people have helped me in one way or another. There are a few people I would like to thank in particular, starting with Thilo Groß. To name but one reason to thank him: he played a vital role in making me rediscover the joys of pen-and-paper role-playing games.

Martin Baurmann and Dirk Stiefs helped in preparing the manuscript of what you have been reading. Martin has also been an excellent source of inspiration, especially during the coffee breaks. Speaking of coffee breaks: A big thanks goes to the all the people who have helped to make the coffee break what it was. I hope the special tradition of the Friday's coffee break will remain alive and well long after I have left the group.

Acknowledgements

Other people have also helped me a lot to enjoy my time in Oldenburg, which in turn helped my scientific work. In particular the lads (and lasses) of the role-playing groups and those of the *Flying Burex* have to be mentioned. During my visits to my “homeland,” it was always a pleasure to meet my Schafkopf friends and my chess mates again.

My most heartfelt thanks go to my family, my parents, sister and brother-in-law, who have strongly supported and encouraged me. Without them all this would not have been possible.

Selain itu, saya mau berterima kasih kepada pacar saya. Dia bantuan besar untuk saya sampai sekarang. Makasih banyak sayangku cintaku!

Lebenslauf

Name: Markus Harle
Geburtsdatum: 15.02.1976
Geburtsort: Wertingen
Staatsangehörigkeit: deutsch

Wissenschaftlicher Werdegang:

Juni 1995: Abitur am Gymnasium Wertingen
Oktober 1995 – September 2001: Studium der Physik an der Universität Augsburg, Abschluß: Diplom-Physiker
Thema der Diplomarbeit: Untersuchungen des Mehrband-Hubbard-Modells mit Spin-Bahn-Wechselwirkung
Oktober 2004 – Juli 2007: Promotionsstudium an der Carl von Ossietzky Universität Oldenburg

Veröffentlichungen:

M. Harle und U. Feudel: On the relation between predictability and homoclinic tangencies, *Int. J. Bif. Chaos*, **15**, 2523–2534, 2005.

M. Harle, F. Kwasniok, und U. Feudel: Growth of finite errors in ensemble prediction, *Nonlin. Processes Geophys.*, **13**, 167–176, 2006.

M. Harle und U. Feudel: Hierarchy of islands in conservative systems yields multimodal distributions of FTLEs, *Chaos, Solitons & Fractals*, **31**, 130–137, 2007.