Part I

Basic topics

Chapter 1

Introduction: why nonlinear methods?

You are probably reading this book because you have an interesting source of data and you suspect it is not a linear one. Either you positively know it is nonlinear because you have some idea of what is going on in the piece of world that you are observing or you are led to suspect that it is because you have tried linear data analysis and you are unsatisfied with its results.[1]

Linear methods interpret all regular structure in a data set, such as a dominant frequency, through linear correlations (to be defined in Chapter 2 below). This means, in brief, that the intrinsic dynamics of the system are governed by the linear paradigm that small causes lead to small effects. Since linear equations can only lead to exponentially decaying (or growing) or (damped) periodically oscillating solutions, all irregular behaviour of the system has to be attributed to some random external input to the system. Now, chaos theory has taught us that random input is not the only possible source of irregularity in a system's output: nonlinear, chaotic systems can produce very irregular data with purely deterministic equations of motion in an autonomous way, i.e., without time dependent inputs. Of course, a system which has both, nonlinearity *and* random input, will most likely produce irregular data as well.

Although we have not yet introduced the tools we need to make quantitative statements, let us look at a few examples of real data sets. They represent very different problems of data analysis where one could profit from reading this book since a treatment with linear methods alone would be inappropriate.

Example 1.1 (NMR laser data). In a laboratory experiment carried out in 1995 by Flepp, Simonet & Brun at the Physics Department of the University of Zürich, a **N**uclear **M**agnetic **R**esonance laser is operated under such conditions that the amplitude of the (radio frequency) laser output varies irregularly over time. From

[1] Of course you are also welcome to read this book if you are not working on a particular data set.

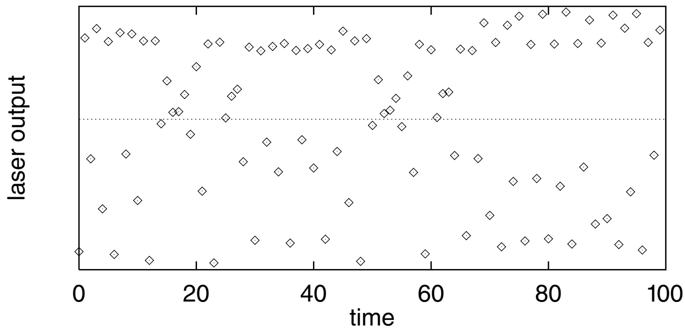

Figure 1.1 100 successive measurements of the laser intensity of an NMR laser. The time unit here is set to the measurement interval.

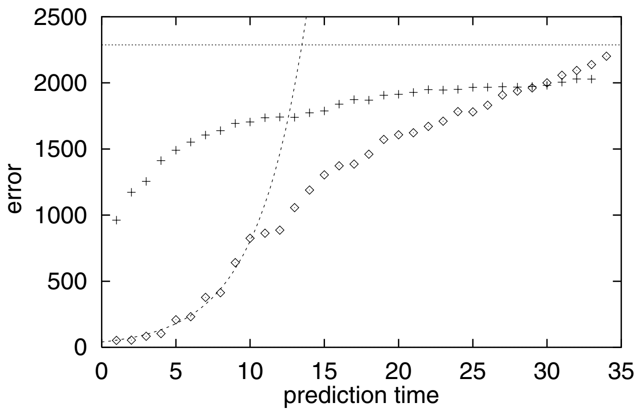

Figure 1.2 The average prediction error (in units of the data) for a longer sample of the NMR laser output as a function of the prediction time. For an explanation of the different symbols see the text of Example 1.1.

the set-up of the experiment it is clear that the system is highly nonlinear and random input noise is known to be of very small amplitude compared to the amplitude of the signal. Thus it is not assumed that the irregularity of the signal is just due to input noise. In fact, it has been possible to model the system by a set of differential equations which does not involve any random components at all; see Flepp *et al.* (1991). Appendix B.2 contains more details about this data set.

Successive values of the signal appear to be very erratic, as can be seen in Fig. 1.1. Nevertheless, as we shall see later, it is possible to make accurate forecasts of future values of the signal using a nonlinear prediction algorithm. Figure 1.2 shows the mean prediction error depending on how far into the future the forecasts are made. Quite intuitively, the further into the future the forecasts are made, the larger

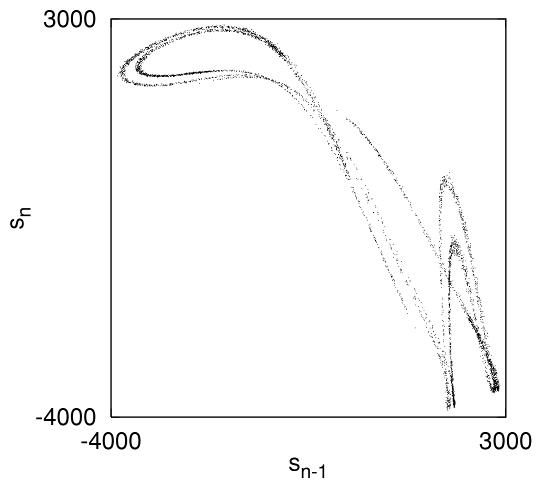

Figure 1.3 Phase portrait of the NMR laser data in a stroboscopic view. The data are the same as in Fig. 1.1 but all 38 000 points available are shown.

will the uncertainty be. After about 35 time steps the prediction becomes worse than when just using the mean value as a prediction (horizontal line). On short prediction horizons, the growth of the prediction error can be well approximated by an exponential with an exponent of 0.3, which is indicated as a dashed line. We used the simple prediction method which will be described in Section 4.2. For comparison we also show the result for the best *linear* predictor that we could fit to the data (crosses). We observe that the predictability due to the linear correlations in the data is much weaker than the one due to the deterministic structure, in particular for short prediction times. The predictability of the signal can be taken as a signature of the deterministic nature of the system. See Section 2.5 for details on the linear prediction method used. The nonlinear structure which leads to the short-term predictability in the data set is not apparent in a representation such as Fig. 1.1. We can, however, make it visible by plotting each data point versus its predecessor, as has been done in Fig. 1.3. Such a plot is called a *phase portrait*. This representation is a particularly simple application of the *time delay embedding*, which is a basic tool which will often be used in nonlinear time series analysis. This concept will be formally introduced in Section 3.2. In the present case we just need a data representation which is printable in two dimensions. □

Example 1.2 (Human breath rate). One of the data sets used for the Santa Fe Institute time series competition in 1991–92 [Weigend & Gershenfeld (1993)] was provided by A. Goldberger from Beth Israel Hospital in Boston [Rigney *et al.* (1993); see also Appendix B.5]. Out of several channels we selected a 16 min

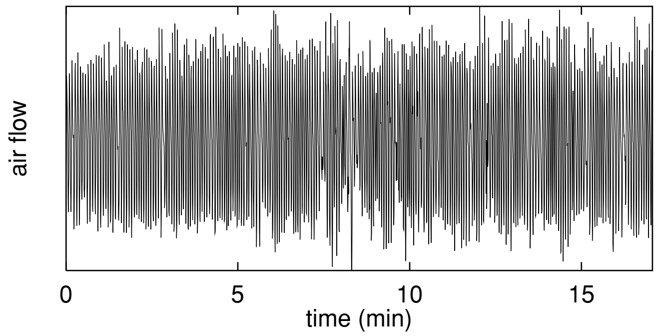

Figure 1.4 A time series of about 16 min duration of the air flow through the nose of a human, measured every 0.5 s.

record of the air flow through the nose of a human subject. A plot of the data segment we used is shown in Fig. 1.4.

In this case only very little is known about the origin of the fluctuations of the breath rate. The only hint that nonlinear behaviour plays a role comes from the data itself: the signal is not compatible with the assumption that it is created by a Gaussian random process with only linear correlations (possibly distorted by a nonlinear measurement function). This we show by creating an artificial data set which has exactly the same *linear* properties but has no further determinism built in. This data set consists of random numbers which have been rescaled to the distribution of the values of the original (thus also mean and variance are identical) and filtered so that the power spectrum is the same. (How this is done, and further aspects, are discussed in Chapter 4 and Section 7.1.) If the measured data are properly described by a linear process we should not find any significant differences from the artificial ones.

Let us again use a *time delay embedding* to view and compare the original and the artificial time series. We simply plot each time series value against the value taken a *delay time* τ earlier. We find that the resulting two *phase portraits* look qualitatively different. Part of the structure present in the original data (Fig. 1.5, left hand panel) is not reproduced by the artificial series (Fig. 1.5, right hand panel). Since both series have the same linear properties, the difference is most likely due to nonlinearity in the system. Most significantly, the original data set is statistically asymmetric under time reversal, which is reflected in the fact that Fig. 1.5, left hand panel, is non-symmetric under reflection with respect to the diagonal. This observation makes it very unlikely that the data represent a noisy *harmonic oscillation*. □

Example 1.3 (Vibrating string data). Nick Tufillaro and Timothy Molteno at the Physics Department, University of Otago, Dunedin, New Zealand, provided a couple of time series (see Appendix B.3) from a vibrating string in a magnetic field.

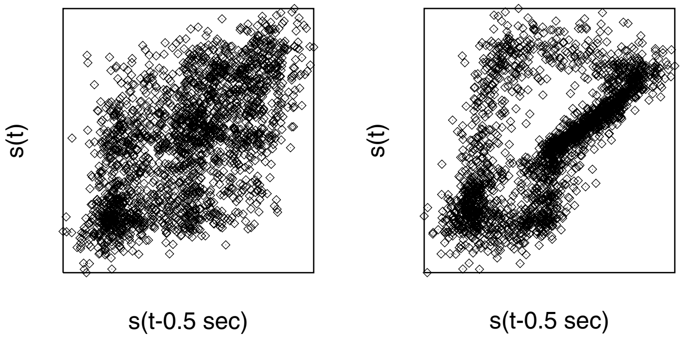

Figure 1.5 Phase portraits of the data shown in Fig. 1.4 (left) and of an artificial data set consisting of random numbers with the same linear statistical properties (right).

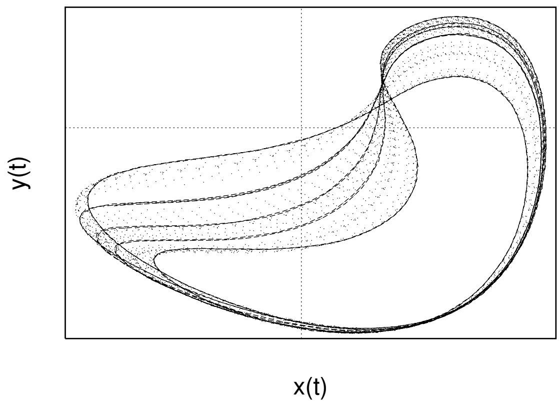

Figure 1.6 Envelope of the elongation of a vibrating string. Line: the first 1000 measurements. Dots: the following 4000 measurements.

The envelope of the elongation of the wire undergoes oscillations which may be chaotic, depending on the parameters chosen. One of these data sets is dominated by a period-five cycle, which means that after every fifth oscillation the recording (approximately) retraces itself. In Fig. 1.6 we plot the simultaneously recorded x- and y-elongations. We draw a solid line through the first 1000 measurements and place a dot at each of the next 4000. We see that the period-five cycle does not remain perfect during the measurement interval. This becomes more evident by plotting one point every cycle versus the cycle count. (These points are obtained in a systematic way called the *Poincaré section*; see Section 3.5.) We see in Fig. 1.7 that the period-five cycle is interrupted occasionally. A natural explanation could be that there are perhaps external influences on the equipment which cause the system to leave the periodic state. However, a much more appropriate explanation

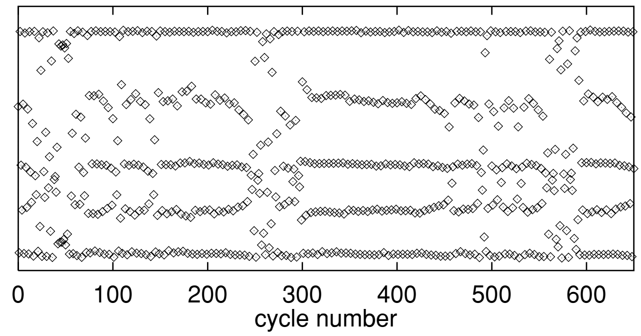

Figure 1.7 Data of Fig. 1.6 represented by one point per cycle. The period-five cycle is interrupted irregularly.

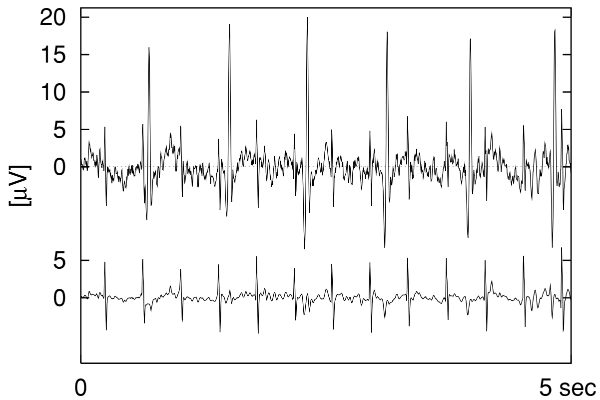

Figure 1.8 Upper trace: electrocardiographic recording of a pregnant woman. The foetal heart causes the small spikes. Lower trace: the foetal electrocardiogram has been extracted by a nonlinear projection technique.

of the data set is given in Example 8.3 in Section 8.4. The irregular episodes are in fact due to *intermittency*, a typical phenomenon found in nonlinear dynamical systems. □

Example 1.4 (Foetal electrocardiogram). Let us consider the signal processing problem of extracting the tiny foetal electrocardiogram component from the (small) electrical potentials on the abdomen of a pregnant woman (upper trace in Fig. 1.8). The faster and smaller foetal signal cannot be distinguished by classical linear techniques because both the maternal and foetal electrocardiograms have broad band power spectra. However, using a nonlinear phase space projection technique that was originally developed for noise reduction in chaotic data (Section 10.3.2), it is possible to perform the separation in an automated way. The lower trace in Fig. 1.8 shows the extracted foetal component. More explanations are given in

Section 10.4. Note that we do not need to assume (nor do we actually expect) that the heart is a deterministic chaotic system. □

There is a broad range of questions that we have to address when talking about nonlinear time series analysis. How can we get the most precise and meaningful results for a clean and clearly deterministic data set like the one in the first example? What modifications are necessary if the case is less clear? What can still be done if, as in the breath rate example, all we have are some hints that the data are not properly described by a linear model with Gaussian inputs? Depending on the data sets and the analysis task we have in mind, we will choose different approaches. However, there are a number of general ideas that one should be acquainted with no matter what one's data look like. The most important concepts analysing complex data will be presented in Part One of the book. Issues which are either theoretically more advanced, require higher-quality data, or are of less general interest will be found in Part Two.

Obviously, in such a diverse field as nonlinear time series analysis, any selection of topics for a single volume must be incomplete. It would be naive to claim that all the choices we have made are exclusively based on objective criteria. There is a strong bias towards methods that we have found either conceptually interesting or useful in practical work, or both. Which methods are most useful depends on the type of data to be analysed and thus part of our bias has been determined by our contacts with experimental groups. Finally, we are no exception to the rule that people prefer the methods that they have been involved in developing.

While we believe that some approaches presented in the literature are indeed useless for scientific work (such as determining Lyapunov exponents from as few as 16 data points), we want to stress that other methods we mention only briefly or not at all, may very well be useful for a given time series problem. Below we list some major omissions (apart from those neglected as a result of our ignorance). Nonlinear generalisations of ARMA models and other methods popular among statisticians are presented in Tong (1990). Generalising the usual two-point autocorrelation function leads to the *bispectrum*. This and related time series models are discussed in Subba Rao & Gabr (1984). Within the theory of dynamical systems, the analysis of unstable periodic orbits plays a prominent role. On the one hand, periodic orbits are important for the study of the topology of an attractor. Very interesting results on the template structure, winding numbers, etc., have been obtained, but the approach is limited to those attractors of dimension two or more which can be embedded in three dimensional phase space. (One dimensional manifolds, e.g. trajectories of dynamical systems, do not form knots in more than three dimensional space.) We refer the interested reader to Tufillaro *et al*. (1992). On the other hand, periodic orbit expansions constitute a powerful way of computing characteristic quantities which are defined as averages over the natural measure, such as dimensions, entropies, and Lyapunov exponents. In the case where a system is *hyperbolic*, i.e., expanding and contracting directions are nowhere tangent to each other, exponentially converging expressions for such quantities can be derived. However, when hyperbolicity is lacking (the generic case), very large numbers of periodic orbits are necessary for the cycle expansions. So far it has not been demonstrated that this kind of analysis is feasible based on experimental time series data. The theory is explained in Artuso *et al*. (1990), and in "The Webbook" [Cvitanović *et al.* (2001)].

After a brief review of the basic concepts of linear time series analysis in Chapter 2, the most fundamental ideas of the nonlinear dynamics approach will be introduced in the chapters on phase space (Chapter 3) and on predictability (Chapter 4). These two chapters are essential for the understanding of the remaining text; in fact, the concept of a phase space representation rather than a time or frequency domain approach is the hallmark of nonlinear dynamical time series analysis. Another fundamental concept of nonlinear dynamics is the sensitivity of chaotic systems to changes in the initial conditions, which is discussed in the chapter about dynamical instability and the Lyapunov exponent (Chapter 5). In order to be bounded and unstable at the same time, a trajectory of a dissipative dynamical system has to live on a set with unusual geometric properties. How these are studied from a time series is discussed in the chapter on attractor geometry and fractal dimensions (Chapter 6). Each of the latter two chapters contains the basics about their topic, while additional material will be provided later in Chapter 11. We will relax the requirement of determinism in Chapter 7 (and later in Chapter 12). We propose rather general methods to inspect and study complex data, including visual and symbolic approaches. Furthermore, we will establish statistical methods to characterise data that lack strong evidence of determinism such as scaling or self-similarity. We will put our considerations into the broader context of the theory of nonlinear dynamical systems in Chapter 8.

The second part of the book will contain advanced material which may be worth studying when one of the more basic algorithms has been successfully applied; obviously there is no point in estimating the whole spectrum of Lyapunov exponents when the largest one has not been determined reliably. As a general rule, refer to Part One to get first results. Then consult Part Two for *optimal* results. This rule applies to embeddings, Chapter 9, noise treatment, Chapter 10, as well as modelling and prediction, Chapter 12. Here, methods to deal with nonlinear processes with stochastic driving have been included. An advanced mathematical level is necessary for the study of invariant quantities such as Lyapunov spectra and generalised dimensions, Chapter 11. Recent advancements in the treatment of non-stationary signals are presented in Chapter 13. There will also be a brief account of chaotic synchronisation, Chapter 14.

Chaos control is a huge field in itself, but since it is closely related to time series analysis it will be discussed to conclude the text in Chapter 15.

Throughout the text we will make reference to the specific implementations of the methods discussed which are included in the TISEAN software package.[2] Background information about the implementation and suggestions for the use of the programs will be given in Appendix A. The TISEAN package contains a large number of routines, some of which are implementations of standard techniques or included for convenience (spectral analysis, random numbers, etc.) and will not be further discussed here. Rather, we will focus on those routines which implement essential ideas of nonlinear time series analysis. For further information, please refer to the documentation distributed with the software.

Nonlinear time series analysis is not as well established and is far less well understood than its linear counterpart. Although we will make every effort to explain the perspectives and limitations of the methods we will introduce, it will be necessary for you to familiarise yourself with the algorithms with the help of artificial data where the correct results are known. While we have almost exclusively used experimental data in the examples to illustrate the concepts discussed in the text, we will introduce a number of numerical models in the exercises. We urge the reader to solve some of the problems and to use the artificial data before actually analysing measured time series. It is a bad idea to apply unfamiliar algorithms to data with unknown properties; better to practise on some data with appropriate known properties, such as the number of data points and sampling rate, number of degrees of freedom, amount and nature of the noise, etc. To give a choice of popular models we introduce the logistic map (Exercise 2.2), the Hénon map (Exercise 3.1), the Ikeda map (Exercise 6.3), a three dimensional map by Baier and Klein (Exercise 5.1), the Lorenz equations (Exercise 3.2) and the Mackey–Glass delay differential equation (Exercise 7.2). Various noise models can be considered for comparison, including *moving average* (Exercise 2.1) and *Brownian motion* models (Exercise 5.3). These are only the references to the exercises where the models are introduced. Use the index to find further material.

Further reading

Obviously, someone who wants to profit from time series methods based on chaos theory will improve his/her understanding of the results by learning more about chaos theory. There is a nice little book about chance and chaos by Ruelle (1991). A readable introduction without much mathematical burden is the book by Kaplan & Glass (1995). Other textbooks on the topic include Bergé *et al*. (1986), Schuster

[2] The software and documentation is available at `www.mpipks-dresden.mpg.de/~tisean`.