

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

**A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA
313/761-4700 800/521-0600**



Université d'Ottawa • University of Ottawa

Nonlinear Analysis of Interspike Intervals

By
Daniel M. Racicot

A thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements
for the degree of Master of Science

Department of Physics
University of Ottawa

©Daniel M. Racicot, Ottawa, Canada, 1996.



**National Library
of Canada**

**Acquisitions and
Bibliographic Services**

**395 Wellington Street
Ottawa ON K1A 0N4
Canada**

**Bibliothèque nationale
du Canada**

**Acquisitions et
services bibliographiques**

**395, rue Wellington
Ottawa ON K1A 0N4
Canada**

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

The author retains ownership of the copyright in his/her thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced with the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

L'auteur conserve la propriété du droit d'auteur qui protège sa thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-20944-X

Abstract

Theories of neural coding and neural information processing rely on a knowledge of the correlations between firing events. We present a method of analyzing experimental interspike interval (ISI) sequences from neurons for the presence of nonlinear correlations. This is accomplished by comparing nonlinear predictions on experimental data sets and on their “surrogate” data sets. These surrogates have the same linear properties as the experimental data, but are otherwise stochastic. A difference in nonlinear predictability between both sets implies that nonlinear correlations are present in the experimental data. We also examine the relationship between, on one hand, neural firing times, and on the other, the dynamical variables of the neuron and its inputs. We focus on a recently proposed method to reconstruct the dynamics of the input to a neuron via a delay embedding of the interspike intervals. We show that this method works well for the simple integrate-and-fire type neuron models when the mean sampling rate is high. However, the method does not extend to more realistic models. Our analysis suggests that this method may be applicable to a small subset of real neurons and that, in general, we should not expect to easily extract information about input stimuli to neurons using ISI embeddings.

Acknowledgements

I would like to thank my research supervisor, André Longtin. It seems to me that he has put almost as many hours at the tortuous task of reading and improving this thesis as I have spent writing it. I would also like to thank him for the freedom he has given me as a research assistant, my time here has been an incredible learning experience due in large part to this freedom. I have benefited tremendously from working with him and I hope that our relationship has been beneficial to him in some way. Financial support from NSERC and the National Institute of Mental Health (USA) grant 1-R01-MH47184-01 is gratefully acknowledged.

I am grateful to Dr. Joseph Bastian, University of Oklahoma, Dr. Hans Braun, Philipps-University Marburg, Marburg, Germany and Dr. Leonard Maler from the University of Ottawa for providing the physiological data analyzed in this thesis.

I must thank two great teachers for caring about a lethargic student enough to get involved and encourage him, they are Profs. Peter Fraser and the late Roberto Mendel, I would probably not be here were it not for their involvement in my education.

Thanks must also go to the people whom I have befriended in the department, Madelaine, Lorraine and Mike for their friendship and their help in navigating the bureaucratic roadmap of the University. I must thank Dave Leblanc for putting up with my erratic (chaotic ??) dishwashing schedule for two years, and for the incredible amount interesting discussions we have had over that time. I must also thank the following friends in the department for making it enjoyable to come to work: Dave, Patrick, Grant, Lixin, Saeed, Dan, Bela, Dr. Song, Riccardo. Thanks also to my close friends not involved in physics, but rather, in golf: Luc, Paul, Sylvain and James. Perhaps the person who has helped me the most during my stay at the University is Bei Wah Chan. I thank him for the many hours he has spent helping.

I have had the good fortune of living close to members of my family here in Ottawa. I must thank Renée, Gordon, Geneviève and Renaud for their understanding, support, storage space, food, hours of fun, and their open door policy. I thank my mother for her support and never ending generosity, none of my education would have been possible without her.

Finally, I wish to thank my dearest friend Amanda Connell for her love, patience, and the joy she brings to my life.

**cette thèse est dédiée à ma mère, Germaine
qui m'a appris l'importance de la tolérance et de la générosité**

et

**à la mémoire de mon père, Eloi
tu seras toujours dans mes rêves**

Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	v
List of Figures	vii
1 Introduction	1
2 Neuron models	4
2.1 Introduction	4
2.2 Neuron models	6
2.3 Models of membrane excitation	6
2.3.1 The Hodgkin Huxley model	8
2.3.2 The FitzHugh-Nagumo model	9
2.3.3 Bursting models	10
2.3.4 Skipping	12
2.4 Integrate and fire (IF) models	13
2.4.1 A generic IF model	13
2.4.2 The Glass-Mackey IF model	14
2.4.3 Leaky IF model	14
3 Nonlinear dynamics and chaos	16
3.1 Introduction	16
3.2 Nonlinear dynamics	17

3.3	Dynamical systems	17
3.3.1	Attractors	21
3.4	Poincaré maps	22
3.5	Properties of dynamical systems	23
3.5.1	Dimensions	24
3.5.2	Lyapunov exponents	25
4	Time series analysis	27
4.1	Introduction	27
4.2	Reconstruction of attractors	28
4.3	Autocorrelation and dimension	29
4.4	Time series and point process prediction	31
4.5	Surrogate data	33
5	Interspike interval embeddings	39
5.1	Previous studies	39
5.2	ISI embeddings for chaotically forced neurons	42
5.2.1	Generic integrate and fire device	44
5.2.2	Glass-Mackey integrate and fire model	45
5.2.3	Leaky integrate and fire model	47
5.2.4	FitzHugh-Nagumo equations	49
5.3	Discussion	54
5.3.1	Integrate and fire models	54
5.3.2	FitzHugh-Nagumo model	57
5.4	Conclusion	57
6	Analysis of ISI sequences	59
6.1	Introduction	59
6.2	Surrogate data and point processes	60
6.2.1	Interpretation for point processes	60
6.2.2	Implementation of the surrogate data method	61
6.2.3	Analysis of ISI sequences from IF models	61
6.2.4	Predictability as a function of prediction time	63
6.3	Analysis of experimental ISI sequences	65

6.3.1	Data from a mammalian thermoreceptor	65
6.3.2	Data from weakly electric fish	67
6.4	Comparison of skipping models	70
6.5	Conclusion	71
7	Conclusion	73
A	Detailed description of the neuron models	77
A.1	Plant's model	77
A.2	Chay's model	79
B	Algorithms for generating surrogate data	81
B.1	Shuffled surrogates	81
B.2	Phase randomized surrogates	82
B.3	Amplitude adjusted surrogates	82

List of Figures

2.1	Common types of neurons	5
2.2	Schematic representation of an action potential	7
2.3	Sample dynamics from the FitzHugh-Nagumo equations	10
2.4	Bursting dynamics for Plant's model	11
2.5	Dynamics of the Glass-Mackey IF model	14
2.6	Dynamics of the leaky IF model	15
3.1	Sensitive dependence on initial conditions	18
3.2	Fixed points and limit cycles	20
3.3	The Hénon attractor	22
3.4	The Rössler attractor	23
3.5	The Poincaré section in three dimensions	24
4.1	Schematic representation of a 3-dimensional embedding	29
4.2	Rössler attractor and reconstruction by delay embeddings	30
4.3	ISI vs. ISI number for data and surrogates	36
4.4	ISI histograms for data and surrogates	37
4.5	Autocorrelation function for data and surrogates	38
5.1	Correlation integrals for ISI's generated from the Lorenz system	40
5.2	Rössler attractor of input system	44
5.3	ISI embeddings for a generic IF model	46
5.4	NPE for the two ISI sequences of the generic IF model	47
5.5	ISI embeddings for the Glass-Mackey IF model	48
5.6	NPE for the two ISI sequences of the Glass-Mackey IF model	49
5.7	ISI embedding for the leaky IF model: high sampling rate	50

5.8	ISI embedding for the leaky IF model: low sampling rate	50
5.9	NPE for leaky IF model	51
5.10	Dynamics of FHN equations with a chaotic input current	52
5.11	ISI return maps for FHN model with a chaotic input current	53
5.12	NPE for ISI's from the FHN model	53
5.13	Return map for large chaotic input current	54
5.14	NPE as a function of α for the Glass-Mackey IF model	56
5.15	Comparison of NPE for integrate and fire models	56
5.16	Comparison of correlation dimension for integrate and fire models	56
6.1	NPE for generic IF and surrogates	62
6.2	NPE for Glass-Mackey model and surrogates	62
6.3	NPE for the leaky IF model and surrogates	62
6.4	$\ln(1 - \tau(T_h))$ as a function of prediction horizon	64
6.5	Data from a cat thermoreceptor	66
6.6	NPE for thermoreceptor data and surrogates	67
6.7	Data from the primary afferent neuron of the weakly electric fish	68
6.8	NPE for data from a primary afferent neuron and surrogates	68
6.9	Data from the lateral segment of a weakly electric fish	69
6.10	NPE for data from the lateral segment and surrogates	70
6.11	NPE for ISI sequences from the Chay model (deterministic skipping) and the Plant model (noise induced skipping).	71

Chapter 1

Introduction

Many dynamical systems are characterized by the fact that they generate “events” in the course of their continuous time evolution. An event can be defined as a sudden change in one or more of the variables of the system. Epileptic seizures and cell division are examples of events in the time evolution of the corresponding biological system. In this thesis we are interested in the case of single neurons, where an event corresponds to a sudden increase in the membrane voltage of the neuron. These events are called “action potentials” or “voltage spikes”.

Since neurons are the basic building blocks of the nervous system, we must gain an understanding of their behaviour if we are ever to fully understand the functioning (and malfunctioning) of this system. One aspect of this understanding concerns the way in which neurons communicate with each other. It is known that this communication process is accomplished, in great part, through the propagation of voltage spikes. In these instances, the only information carried in “neural” messages is in the timing of the spikes, i.e. the interspike intervals (ISI’s). What is not well understood is the method by which the information is encoded into the ISI’s. In order to properly constrain theories of neural coding and information processing, it is important to characterize the correlations between ISI’s. In this thesis, we develop methods inspired from nonlinear dynamics to test for the possible presence of nonlinear correlations in sequences of ISI’s. These methods are applied successfully to ISI sequences from both neuron models and real neurons.

Neurons are excitable cells in which the transmembrane potential fluctuates. The preferred experimental method of studying the electrical activity of neurons is

to measure the voltage of the cell under controlled conditions. For example, we can control ionic concentrations and give a cell electrical stimulations. To monitor the voltage of the cell, we must put it in contact with a microelectrode. This electrode may damage the neuron or alter the behaviour of the neuron in an undesirable way. In many cases, the neurons are too small, or they are surrounded by specialized tissue, each of which prevents the neuron from being put into contact with the electrode. For these reasons we are often forced to take extracellular measurements with a nearby electrode which monitors only the times at which the neuron fires an action potential. We are thus left with only the ISI's as our experimental data set. Hence, there is great interest in relating the spiking times of the neurons to the dynamics of the underlying dynamical system which generated them (the neuron and its inputs).

Recently, a method has been proposed [45, 46] that allows the reconstruction of the dynamics of the input to a neuron from the ISI's. This method has been shown to work well for the simplest known neuron models¹. We extend it to more realistic models in order to determine the applicability of this method to real neurons. Our investigations suggest that it will not in general be possible to reconstruct dynamics of the underlying systems from interspike intervals. In other words, in general it is not possible to isolate the influence of the input stimulus from the ISI's of a neuron. Nevertheless, the method may be applicable to a limited class of real neurons.

A further motivation for studying ISI's using dynamical methods is for model validation. If we can determine that a certain experimental ISI sequence contains nonlinear correlations, we expect that a good model must be capable of producing ISI's with similar correlations. Finally, another reason for studying the dynamics of "inter-event" times in neurons is that the methods developed may be easily adaptable to the study of other systems that are characterized by the fact that they generate "events", such as the onset times of epileptic seizures and earthquakes.

This thesis is organized as follows. In Chap. 2 we give a general introduction to neurons and the neuron models used in this study. Chapters 3 and 4 present an overview of topics from nonlinear dynamics and nonlinear time series analysis which are pertinent to our study. In Chap. 5 we present our results on reconstructing

¹This is the simplest neuron model used in research on single neurons. There are simpler models used in artificial neural networks.

the attractor of the input to chaotically forced neuron models. The surrogate data method is applied to ISI sequences from neuron models and real neurons in Chap. 6. The thesis ends with a brief conclusion and outlook onto future research.

Publication notes

The results presented in this thesis have been published in:

D. M. Racicot and A. Longtin. Interspike interval attractors from chaotically driven neuron models. *Physica D*, 1996 (In Press).

A. Longtin and D. M. Racicot. Spike train Patterning and forecastability. *Biosystems* 40(1997)111–118.

A. Longtin and D. M. Racicot. Assessment of linear and nonlinear correlations between neural firing events. In: *Nonlinear Dynamics and Time Series: Building a Bridge Between the Natural and Statistical Sciences*. D. Kaplan and C. D. Cutler, editors. The Fields Institute Communications Vol. 11, 221–238, 1996. American Mathematical Society, Providence, RI.

D. M. Racicot and A. Longtin. Reconstructing dynamics from neural spike trains. In *Proceedings of the IEEE 17th Annual Conference on Engineering Methods in Biology and Medicine*, September 1995.

A. Longtin and D. M. Racicot. Distinguishing between chaotic and stochastic phase locking in mammalian thermoreceptors. To be submitted to *Biological Cybernetics*.

Chapter 2

Neuron models

Much of our analysis is performed on data generated by model neurons. There are many advantages to using this type of data: it is noise free and the length of the data set is limited only by computer resources. Data from real neurons is often of very short length and/or nonstationary. Clean data from reliable experiments can also be difficult to find. In the following chapter we give a brief introduction to the functioning of neurons. We follow with a description of neuron models, focusing on those we have used in this study.

2.1 Introduction

Neurons come in a variety of shapes and sizes, from the simple unipolar interneuron of invertebrates to the massively complex neurons in the hippocampal area of the brain. Fig. 2.1 shows some common types of neurons.

Like all living cells, neurons contain a cell body (or soma) which houses the nucleus and protein synthesis areas. Neurons also have dendrites branching out of them in various fashions, they act as surfaces for the reception of information from other cells. This information is in the form of chemical signals which cause an electrical signal to be propagated along the dendrites to the cell body. The cell body is what we are most interested in modeling since it is the part of the cell which can generate the electrical impulses which propagate to other neurons. In a typical neuron the information is input to the soma and processed, the result is a

series of *action potentials* (or spikes) which are sent down the axon to the synaptic region. At the synapse, the incoming action potential causes the release of chemicals known as neurotransmitters. These chemicals are the means by which the signal is transmitted to other cells. Some example of action potentials can be seen in Fig. 2.2.




Figure 2.1: Some common types of neurons, from [32] p. 18.

2.2 Neuron models

Our discussion of neuron models will be limited to the models that we have used in this study as well as the Hodgkin Huxley model. For a more detailed exposition see [20, 29, 59]. The models we cover can be grouped into two classes:

- (a) Models of the cell body as an excitable membrane with ionic channels immersed in an ionic solution.
- (b) Qualitative models which are meant to reproduce only a subset of the properties of the neuron. In this thesis, all models of this type are called *integrate and fire* (IF) models.

Type (a) models are generally expressed as a set of coupled ordinary differential equations for the membrane potential and other variables (e.g. ionic concentrations inside and outside the cell). Type (b) models are simple models with one or two variables. They are generally derived from type (a) models, for example by ignoring some of the less important ionic currents. Two very good examples of this type of reduction are presented in [2] and [11].

2.3 Models of membrane excitation

We can think of the cell body as an excitable membrane bathing in an ionic solution. The basic governing equation of such a membrane is the conservation of charge:

$$C_m \frac{dV}{dt} = -I_m + I_{ext}, \quad (2.1)$$

where V is the membrane voltage, C_m is the membrane capacitance, I_m is the sum of the membrane currents, and I_{ext} the collective external currents. The membrane currents arise primarily from the conduction of ions through the membrane by voltage dependent ion channels or by the leakage of certain ions through the membrane. The external currents are mainly those coming from the dendrites or, in an experimental setting, a stimulus delivered through a micro-electrode. We stimulate the models by applying various input currents. The most commonly used input currents

are constant bias currents and/or oscillating currents. We also force some of the models with a signal taken from a low dimensional chaotic system in Chap. 5.

In the absence of stimulation the typical cell maintains a constant negative resting potential $-V_0$. An input current can have two effects on the membrane. If the current is below a critical stimulus value, also called the firing threshold, the membrane responds passively by changing its potential and slowly recovering back to the resting potential. A current that is above the threshold of critical stimulus causes the membrane potential to depolarize (becoming more positive than $-V_0$) and then rapidly hyperpolarize (becoming more negative than $-V_0$). This is the action potential which is followed by a slow recovery back to the resting potential. This can be seen schematically in Fig. 2.2 A where the S denotes a depolarizing stimulus.




Figure 2.2: Typical action potentials from a frog sciatic nerve. A) The spike is caused by the stimulus given at S. B) and D) The effect of a second stimulus cannot cause another spike during the refractory period. C) This is an action potential where the nerve has been crushed between the two electrodes. Taken from [12], p. 443.

Consecutive stimuli can cause many responses in the membrane. If the previous stimulus caused an action potential there is a time when another stimulus, no matter how strong, cannot cause the cell to produce another action potential. This is the *absolute refractory period* and it is illustrated in Fig. 2.2D. After the absolute refractory period comes the *relative refractory period* when a stimulus can cause

another impulse but the threshold value for this excited state of the membrane is much higher than for the membrane at resting potential. The relative refractory period lasts until the cell has returned to its resting potential, and consequently the stimulus threshold has returned to its resting value. If the first stimulus was below threshold, there is a time where another weak current can cause the cell to spike. This enhanced phase is followed by a phase at which the threshold of the second stimulus is greater than the threshold at rest. Hence consecutive sub-threshold stimuli can combine to have the same effect of one above threshold stimulus but they must be close together in time [11].

2.3.1 The Hodgkin Huxley model

The general discussion above is one that is most often given when describing the model of Hodgkin and Huxley (HH) [19]. This model is based on experiments performed on the squid giant axon in the early 1950's. In this axon, as in many other cells there is a high concentration of sodium and a low concentration of potassium in the extracellular solution relative to the concentrations inside the cell. Hodgkin and Huxley proposed a conservation of charge equation as follows, with three different ionic currents;

$$C \frac{dV}{dt} = I_{ext} - g_{Na}(V - V_{Na}) - g_K(V - V_K) - g_l(V - V_L). \quad (2.2)$$

Here V_{Na} , V_K and V_L are the equilibrium voltages for sodium, potassium and leakage ions respectively, and g_{Na} , g_K and g_l are the conductances of the membrane to the different ions. We can write the conductances in terms of three time and voltage dependent variables

$$g_{Na} = \bar{g}_{Na} m^3 h, \quad g_K = \bar{g}_K n^4, \quad g_l = \text{constant}.$$

The \bar{g} terms are the experimentally determined maximum conductance values of the membrane to the respective ions, m and n are activation parameters which vary with voltage and time between 0 and 1, and h is an inactivation parameter which also varies in the unit interval. The time and voltage dependence of the conductance

variables are given by:

$$\begin{aligned}\dot{m} &= (1 - m)\alpha_m - m\beta_m, \\ \dot{n} &= (1 - n)\alpha_n - n\beta_n, \\ \dot{h} &= (1 - h)\alpha_h - h\beta_h.\end{aligned}$$

The α 's and β 's were determined by fitting to the experimental data:

$$\begin{aligned}\alpha_m &= \frac{0.1(25 - V)}{e^{0.1(25 - V)} - 1}, \quad \beta_m = 4e^{\frac{V}{18}}, \\ \alpha_n &= \frac{0.01(V + 10)}{e^{0.01(V + 10)} - 1}, \quad \beta_n = 0.125e^{\frac{V}{20}}, \\ \alpha_h &= 0.07e^{\frac{V}{20}}, \quad \beta_h = \frac{1}{1 + e^{\frac{V + 30}{10}}}.\end{aligned}$$

2.3.2 The FitzHugh-Nagumo model

One important property of the HH model is that there are very different time scales at work. The time scales for the processes associated with the action potential itself are very small compared to the time scales associated with the “refractoriness” and the relaxation. These in turn are small compared to the adaptation and recovery time scales. In the HH equations, the conductance variable m reaches its asymptotic value m_∞ much faster than the h and n conductance variables. If we are willing to forego accuracy in the model at small time scales we can simply make the m variable assume its asymptotic value for all time. This reduces the HH model from four to three dimensions. Another reduction can be accomplished by cleverly (and carefully) combining the h and n conductance variables into one. This type of reduction leads to the FitzHugh-Nagumo (FHN) equations [10, 31]:

$$\begin{aligned}\epsilon \frac{dv}{dt} &= v(v - a)(1 - v) - w + S(t), \\ \frac{dw}{dt} &= v - dw - b.\end{aligned}\tag{2.3}$$

This system while being only two dimensional still captures the basic mechanism for generating action potentials. The membrane potential is $v(t)$, and the variable $w(t)$ controls recovery by making the membrane potential temporarily more negative than

the resting potential after an action potential, thus producing a refractory period. We have added a term for the input currents $S(t)$. This model lies on the border of the type (a) and type (b) models discussed above. It would be the most realistic type (b) model or a simple type (a) models.




Figure 2.3: Sample dynamics for the FHN equations for $\epsilon = 0.005$, $a = 0.5$, $b = 0.175$, $d = 1$, and $S(t) = .01 \sin(0.75t)$.

2.3.3 Bursting models

Many nerve cells can display a behaviour called bursting. Instead of producing a single action potential these cells fire in bursts of activity which are separated by periods of inactivity. Typical bursting dynamics can be seen in Fig. 2.4. The most common models for bursting cells use an interaction of fast conductances (similar to the Hodgkin-Huxley type conductances) to produce action potentials, and much slower conductances to produce the alternation between firing and resting activity.

Plant's slow wave burster model

This is a model of the invertebrate *Aplysia* R 15 pacemaker cell [35]. It is called a slow wave burster because the bursts ride on the back of a slower oscillation as shown in Fig. 2.4. A slightly modified version of this model has been proposed to explain the mechanism of information encoding in mammalian cold receptors [25]. The model consists of spike generation mechanisms similar to the HH model. During a burst there is accumulation of calcium ions inside the membrane. When




Figure 2.4: Bursting dynamics for Plant's model.

this concentration reaches a certain level it activates an outward potassium current which ends the burst. During the interburst period the calcium is slowly removed and thus the potassium current diminishes until the spike generation currents again dominate and the bursting begins again [35]. The equations for Plant's model are as follows from [42]

$$C_m \frac{dV}{dt} = \bar{g}_{Na} m_{\infty}^3(V) h (V_{Na} - V) + \bar{g}_{Ca} x (V_{Ca} - V) + [\bar{g}_k n^4 + \frac{\bar{g}_{KCa} Ca}{Ca + 0.5}] (V_K - V) + \bar{g}_L (V_L - V), \quad (2.4)$$

$$\begin{aligned} \frac{dh}{dt} &= \lambda [h_{\infty}(V) - h] / \tau_h(V), \\ \frac{dn}{dt} &= \lambda [n_{\infty}(V) - n] / \tau_n(V), \\ \frac{dx}{dt} &= [x_{\infty}(V) - x] / \tau_x, \\ \frac{dCa}{dt} &= \rho [K_c x (V_{Ca} - V) - Ca]. \end{aligned}$$

Here V denotes the membrane potential, m and h are the HH like conductance variables for the sodium channels and n is the activation variable for the potassium channel. The slow outward current variable is x , and the calcium accumulation

is modelled as the voltage dependent Ca variable. Sect. A.1 gives details on the variables in the above equations as well as the parameter values.

Chay's model

The model proposed by Chay in [5] is similar to the above model. The difference is that there is no mechanism for calcium accumulation here. The model simply assumes that a combination of slow and fast currents are present. The conservation of charge equation is:

$$C_m \frac{dV}{dt} = -\bar{g}_{fast} m_\infty^3 h(V - V_{fast}) - \bar{g}_{slow} df(V - V_{slow}) - \bar{g}_K n(V - V_K) - \bar{g}_L(V - V_L), \quad (2.5)$$

where each of the variables h , d , f and n obey a differential equation of the form:

$$\frac{dx}{dt} = \frac{x_\infty - x}{\tau_x}.$$

The functions x_∞ and τ_x , as well as typical parameter values are given in section A.2.

2.3.4 Skipping

The models described so far are all capable of exhibiting a behaviour known as skipping. When the FHN equations are forced by an oscillating current (see Fig. 2.3) the spikes can become phase locked with the input. In other words the model will fire in phase with the oscillating input. By tuning the parameters [24], it is possible to get this phase locked behaviour but, rather than firing at every period, the neuron will fire at multiples of the period, skipping an apparently random number of cycles of the oscillation before firing again. The bursting models also can exhibit this type of behaviour by firing a single spike (or a burst) at multiples of the slow wave oscillation.

2.4 Integrate and fire (IF) models

Of all the models for biological rhythms the integrate and fire models are some of the simplest. These models assume that the activity of the system evolves following some preset rules¹ until it reaches a threshold at which point the activity is reset so that the cycle can begin again. When applying these models to neurons, we can think of the evolution as a “charge accumulation”. The charge accumulates (following the rules of the model) inside the cell until the membrane voltage reaches the firing threshold, when a spike is generated. This can be seen schematically in Figs. 2.5 and 2.6. Integrate and fire models are effective in duplicating the capacitive behaviour of a cell but they do not account for the refractory properties. These types of models have been used to study many different mechanisms such as circadian rhythms [14] and phase locking in the respiratory system [13]. The computational simplicity of these models makes it possible to use them in the study of networks of “biological” neurons.

2.4.1 A generic IF model

This is a simple IF model proposed in [45] to study the possibility of reconstructing the dynamics of the input to a neuron by embedding the interspike intervals. It is meant to be “a model which is as generic as possible” [45]. The membrane voltage $V(t)$ is the integral of an input signal. When the voltage reaches a threshold (Θ), a spike is generated, the voltage is reset to zero, and the integration continues. If $S(t)$ is the signal input to the IF device, then one can define the firing times T_1, T_2, \dots, T_N from the following:

$$V(t) = \int_{T_{n+1}}^{T_n} S(t) dt, \quad (2.6)$$

with the reset condition $V(t^*) = 0$ if $V(t) = \Theta$. The interspike intervals are defined as $I_n = T_n - T_{n-1}$. The threshold controls the firing rate of the model, as Θ increases the firing rate decreases.

¹For example linearly increasing, random walks or chaotic motion.

2.4.2 The Glass-Mackey IF model

This model is similar to the one used in [13] to study phase locking in the respiratory system. There, the threshold of the IF device was modulated with a sinusoidal signal. We modulate the threshold with a signal from a low dimensional chaotic system. This is equivalent to chaotic forcing of the system. The activity in this case is linearly increasing until it reaches the modulated threshold. A spike is generated and the activity is reset. Thus the activity of the system is:

$$\begin{aligned} V(t) &= \alpha t + \beta, \\ V(t^+) &= 0 \quad \text{if } V(t) = S(t). \end{aligned} \tag{2.7}$$

The integrate and fire dynamics are shown in Fig. 2.5. In this model the firing rate is controlled by the slope (α), as α increases so does the firing rate.

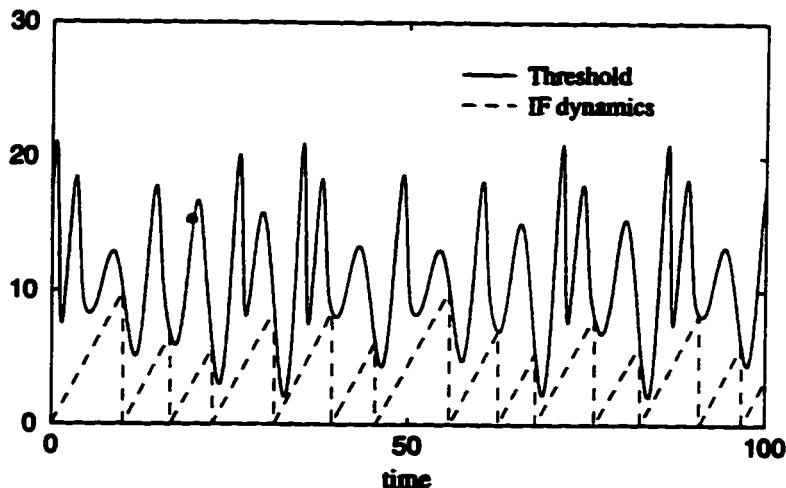


Figure 2.5: Dynamics of the Glass-Mackey IF model of Eq. 2.7 with $\alpha = 10, \beta = 0$.

2.4.3 Leaky IF model

This model is similar to the one used in [22]. The authors use it as a simplified version of HH type dynamics in order to analytically characterize its behaviour under periodic forcing. We will be studying this model under chaotic forcing. The model is also an extension of the generic IF model. The extension lies in the additional

dissipation term in the differential equation for the membrane potential. The model describes a leaky membrane in terms of its potential ($v(t)$), a dissipative constant σ and an applied stimulus $S_a(t)$. When $v(t)$ reaches a firing threshold a spike is generated and the voltage is reset to zero. This is described by:

$$\dot{v} = -\sigma v + S_0 + S(t), \quad (2.8)$$

with the reset condition $v(t^+) = 0$ if $v(t) = \Theta$. In this case the equations are normalized and $\Theta = 1$. The leakiness refers to the tendency of the membrane potential to decay back to the resting potential after a small perturbation. The roughly exponential decay is governed by the capacitative and resistive properties of the cell. This makes the leaky IF model the most realistic IF model used in this study. Dynamics of the leaky IF are shown in Fig. 2.6.




Figure 2.6: Dynamics of the leaky integrate and fire model with parameter values: $\sigma = 0.4$, $S_0 = 0.5$, and $\beta = 0.01$. $S(t)$ is the x component of the Rössler attractor.

Chapter 3

Nonlinear dynamics and chaos

“A very small cause which escapes our notice determines a considerable effect that we cannot fail to see, and then we say that the effect is due to chance. If we knew exactly the laws of nature and the situation of the universe at the initial moment, we could predict exactly the situation of that same universe at a succeeding moment. But even if it were the case that the natural laws had no longer any secret for us, we could still only know the initial situation approximately. If that enabled us to predict the succeeding situation with the same approximation, that is all we require, and we should say that the phenomenon had been predicted, that it is governed by laws. But it is not always so; it may happen that small differences in the initial condition produce very great ones in the final phenomena. A small error in the former will produce an enormous error in the latter. Prediction becomes impossible and we have the fortuitous phenomenon.” Henri Poincaré, 1903

3.1 Introduction

Traditionally, the tools for analysing temporal biological data have been based on linear and/or statistical analysis (e.g. Fourier analysis and linear prediction) [47]. During the past fifteen years the development of nonlinear time series analysis has generated several new methods of analyzing temporal data. Before we can move to an exposition of some of the methods of nonlinear time series analysis, we must define some key concepts from the theory of nonlinear dynamics and chaos.

For a good basic introduction to chaos at an easily accessible level, see the articles by Crutchfield *et al.* [6] and Jensen [21]. A more complete discussion is given in the review of Eckmann and Ruelle [7] and the collection of reprints by Bai-Lin Hao [17].

3.2 Nonlinear dynamics

Since the turn of the century it has been known that simple systems, like the three body problem analyzed by Poincaré, can exhibit complicated behaviour. For several years Poincaré ideas as stated in the quote above had little impact on the dynamics community. In the early sixties Lorenz [28] “re-discovered” *sensitive dependence on initial conditions* while studying a simple model of fluid convection in the atmosphere. Lorenz observed that simulations of his system with nearly identical starting points did not yield similar solutions. The Lorenz equations are a set of three coupled nonlinear ordinary differential equations:

$$\dot{x} = \sigma(y - z), \quad \dot{y} = x(\rho - z) - y, \quad \dot{z} = -\beta z + xy.$$

In Fig. 3.1, we compare two different solutions using the x component of the Lorenz equations. The initial conditions of these solutions are identical in two of the components (y and z) and differ by 10^{-4} in the x component. Notice that after 17 units of time the two solutions have diverged from each other and are distinguishable. This divergence is due to the sensitivity to initial conditions. The system is unpredictable in the long run, just like random systems, but predictable on short time scales. This raised the question of whether or not systems that were previously thought to be random might be deterministic. This discovery sparked a new research effort to understand the behaviour of nonlinear systems.

3.3 Dynamical systems

A *dynamical system* consists of two parts: a state, or set of variables which describes the system and a set of rules for advancing the state forward in time. The system’s evolution is given as a trajectory in *phase space*, a space spanned by the state variables of the system. Time parametrizes phase space trajectories. Nonlinear




Figure 3.1: Two x time series from the Lorenz systems. Their initial conditions differ by 10^{-4} in the x component.

dynamics studies the evolution of nonlinear dynamical systems in either discrete or continuous time. The former is referred to as a *flow* since the trajectories in phase space often resemble fluid flow. Discrete time dynamical systems are called *maps*; they can be useful to model systems that are measured in cycles (e.g. monthly or seasonally). For example, we could model the yearly population of a certain species as a function of the population in the previous year, $p_{i+1} = F(p_i)$, or the previous two years, $p_{i+1} = F(p_i, p_{i-1})$. When one has a higher order map (the two year population model would be a second order map) it is often advantageous to transform them into a set of coupled first order maps. This transformation often simplifies the analysis and allows to draw parallels with first order coupled differential equations. In the above model, if we took $x_i = p_{i+1}$ and $y_i = p_i$, the map would be transformed into two coupled first order maps:

$$\begin{aligned}x_{i+1} &= F(x_i, y_i), \\y_{i+1} &= x_i.\end{aligned}$$

Maps are also extremely important in the study of the geometrical properties of continuous time dynamical systems via Poincaré sections (see Sect. 3.4). The 2-dimensional map proposed by Hénon is closely related to the Poincaré section of the Lorenz equations [18]:

$$\begin{aligned}x_{n+1} &= a - x_n^2 + by_n, \\y_{n+1} &= x_n.\end{aligned}$$

Continuous time dynamical systems are most often expressed as a set of differential equations of the form $\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x})$ where the dot denotes differentiation with respect to time. For example, the damped forced harmonic oscillator obeys the second order differential equation: $m\ddot{x} + r\dot{x} + kx = A \cos(\omega t)$. We can transform this equation into the above form by taking $y = \dot{x}$, then:

$$\begin{aligned}\dot{x} &= y, \\ \dot{y} &= \frac{1}{m} [A \cos(\omega t) - ry - kx].\end{aligned}\tag{3.1}$$

Another example is the system of differential equations proposed by Rössler [44]:

$$\begin{aligned}\dot{x} &= -y - z, \\ \dot{y} &= x + ay, \\ \dot{z} &= b + (x - c)z.\end{aligned}\tag{3.2}$$

Rössler proposed his equations to study the geometric properties of a simple system that was capable of exhibiting chaotic behaviour. He based his equations on those studied by Lorenz. The Rössler system contains only one nonlinearity (the $x \cdot z$ term in the \dot{z} equation) whereas the Lorenz system contains two (one in each of the \dot{y} and \dot{z} equations). The Rössler model is the one we referred to in Chap. 2 as a low-dimensional chaotic system. It will be used to create chaotic input currents to neuron models in Chap. 5.




Figure 3.2: On the left the attractor is the circle, and is called a limit cycle. On the right the attractor is the point at the origin.

3.3.1 Attractors

Dynamical systems are classified into two categories: conservative and dissipative. Conservative systems preserve phase space volumes, in other words, if all points in an initial volume $V(0)$ of phase space are evolved forward in time according to the rules of the dynamical system, the resulting volume $V(t)$ is the same as $V(0)$ ¹. The equations of motion of a frictionless plane pendulum in a constant gravitational field is an example of a conservative system.

Dissipative systems shrink the phase space volume under time evolution over at least some parts of the phase space (the rate of volume contraction may not be constant throughout). Dissipative systems are commonly described in phase space in terms of their *attractors*. These are sets of points (or a single point) to which the trajectories evolve asymptotically with time. There are four common types of attractors: fixed points, limit cycles, tori, and strange. The first three are well behaved and correspond to predictable systems, while the fourth generally belongs to chaotic systems².

If we take the example of the damped forced oscillator (Eq. 3.1 above), it can exhibit both fixed point and limit cycle behaviour. If we set the constant $A = 0$ and give the system an initial push it will start oscillating, but lose energy to friction until it comes to rest. Thus the point $x = 0, y = 0$ is a stable *fixed point* of the system. If the constant A is nonzero, we can make the system oscillate. This would correspond to a *limit cycle* behaviour. It does not matter if we start the system with a little push or a large perturbation it will eventually come to oscillate on the limit cycle as in Fig. 3.2. A torus attractor (it has the shape of a donut) is slightly more complicated than the limit cycle but it is also associated with oscillating behaviour, in this case oscillations with two different frequencies. If a system is oscillating in more than one dimension such as a pendulum clock on a turntable, the attracting set is generally a torus, although it can also be a limit cycle or a fixed point depending on the parameters of the system.

We saw earlier that some equations can exhibit sensitive dependence on initial conditions. This implies that phase space trajectories that are close together in

¹The shape of the volume may change.

²Some strange attractors are associated with non chaotic systems [16].

space will diverge from each other with time, which in turn suggests that an initial volume of phase space eventually grows infinitely large. However, if the sensitive dependence on initial conditions is present within a dissipative system, we will often encounter strange attractors. Within these strange attractors nearby trajectories diverge away from each other on average. The phase space volume of the system is kept finite by the dissipative nature of the system. This behaviour is often referred to as stretching and folding since it is akin to making a pizza crust. The dough is kneaded and stretched to mix up the ingredients and folded over to keep it within a reasonable size. Strange attractors for the Hénon map and Rössler equations are shown in Figs. 3.3 and 3.4.




Figure 3.3: 10000 points of the Hénon attractor with parameters $a = 1.4$, $b = 0.3$.

3.4 Poincaré maps

Since continuous time dynamical systems can be quite complicated, it is often desirable to simplify the dynamics in order to gain some insight into the behaviour of the system. One such method is called the *Poincaré section* method. A Poincaré section reduces the n -dimensional flow of a set of differential equations to an $(n-1)$ -dimensional map. This is accomplished by finding the intersection of the flow with

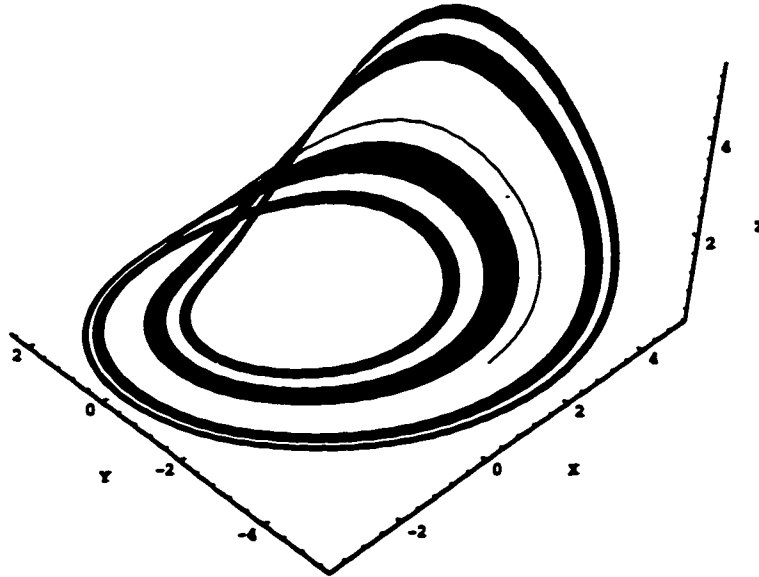


Figure 3.4: The Rössler attractor with $a = .386$, $b = 2$, $c = 4$, known as the four band chaotic attractor.

a suitably chosen hyperplane in phase space. For example, if we had a three dimensional trajectory $(x(t), y(t), z(t))$ and found the points of intersection of the flow with the plane $z = C$ (a constant) we would have a Poincaré section for the trajectory. The crossing points on the plane (x_i, y_i) constitute a two-dimensional map

$$\begin{aligned}x_{i+1} &= f(x_i, y_i) \\y_{i+1} &= g(x_i, y_i)\end{aligned}$$

which is a Poincaré map of the flow. This procedure is illustrated schematically in Fig. 3.5. It is important to note that f and g may be difficult or impossible to compute analytically; they sometimes can be approximated.

3.5 Properties of dynamical systems

We now turn our attention to quantifying some of the properties of dynamical systems and attractors. We first describe different definitions of dimensions and then we look at Lyapunov exponents, which quantify the sensitivity to initial conditions




Figure 3.5: The Poincaré section in three dimensions.

of a dynamical system.

3.5.1 Dimensions

When we see a fixed point or a limit cycle we can definitely say that they have, respectively, dimensions 0 (a point) and 1 (a line or curve). What about the Rössler or Hénon attractors? The Rössler is clearly not a two dimensional object but does not seem to be three dimensional either. The Hénon attractor is neither one nor two dimensional. These types of objects are called *fractals*, as they have non-integer dimensions. This property of strange attractors is the result of the stretching and folding process described earlier. The Hénon and Rössler attractors have dimensions of approximately 1.2 and 2.4 respectively.

A simple definition of dimension that allows for non-integer values is the *box counting dimension*, or capacity of a set. The box counting dimension is defined by covering the object with a grid of boxes, each with sides of length r . We then count the number of non-empty boxes $N(r)$. The box counting dimension is given by:

$$D_{\text{box}} = \lim_{r \rightarrow 0} \frac{\ln(N(r))}{\ln(1/r)}. \quad (3.3)$$

For a fixed point attractor, $N(r) = 1$ independently of r (i.e. it only takes one box to cover the point) and the limit gives $D_{\text{box}} = 0$. For a limit cycle attractor $N(r) = L/r$ where L is the total length of the curve, and taking the limit as $r \rightarrow 0$ gives $D_{\text{box}} = 1$. One of the disadvantages of the box counting dimension is that the function $N(r)$ assigns the same value to boxes with one point and boxes with many points.

Another type of dimension, which weighs boxes according to the probability that an orbit falls within them has been proposed. It is in fact a spectrum of dimensions $\{D_q \mid q = 0, 1, 2, \dots\}$:

$$D_q = \frac{1}{q-1} \lim_{r \rightarrow 0} \frac{\log(\sum_i P_i^q)}{\log(r)}, \quad (3.4)$$

where P_i is the probability that a trajectory falls within the i th box. D_0 , D_1 , and D_2 are called pointwise, information, and correlation dimensions respectively. The correlation dimension is the easiest to calculate from experimental data and is the most commonly used measure of dimension (see Sect. 4.3).

3.5.2 Lyapunov exponents

Lyapunov exponents are important in the study of dynamical systems. They provide a measure of the way trajectories move in time with respect to one another as opposed to the static measurements of dimension. For example, suppose we take a phase space sphere of radius $r(0)$ and evolve it in time according to the rules of a dynamical system. The sphere is transformed into an ellipsoid with axes $r_i(t)$, as some directions contract while others expand. We define the Lyapunov exponents as [1]:

$$\lambda_i \equiv \lim_{t \rightarrow \infty} \frac{1}{t} \ln \left(\frac{r_i(t)}{r(0)} \right). \quad (3.5)$$

For an N -dimensional dynamical system there are N Lyapunov exponents. The largest Lyapunov exponent (λ_1 by convention) measures the rate of separation in the direction that is expanding at the fastest rate. The smallest exponent measures the slowest rate of expansion³. We can also think of the largest Lyapunov exponent as measuring the way linear distances between neighbouring trajectories grow, averaged

³It is usually a contraction rate.

over the attractor. If two trajectories are initially separated by some distance δ their separation grows as $\delta e^{\lambda_1 t}$ on average. An element of 2-dimensional area a would grow as $a e^{(\lambda_1 + \lambda_2)t}$.

Lyapunov exponents are important in determining whether or not a specific system is chaotic. As we saw previously, for a dynamical system to have a strange attractor it must exhibit an exponential divergence of nearby trajectories in some regions of phase space. This statement is equivalent to saying that the system must have at least one positive Lyapunov exponent. The existence of a positive Lyapunov exponent is a necessary and sufficient condition for a system to be chaotic [1].

Chapter 4

Time series analysis

In this chapter we present methods from nonlinear time series analysis. We first look at a remarkable theorem which allows us to reconstruct the dynamics of a multi-dimensional system from the measurement of only one variable. A few commonly used methods of determining the properties of a time series are then discussed. We end by presenting *surrogate data analysis* which is an attempt to apply some of these methods in a more statistically rigorous manner. For more complete coverage, the reader is referred to reviews by Grassberger *et al.* [15] and Abarbanel *et al.* [1] and references therein.

4.1 Introduction

The emergence of deterministic chaos initiated a large research effort into the possibility of identifying chaos in experimental data. Much of the early work was done on model systems that were known to be chaotic (such as the Lorenz and Rössler systems). The focus quickly turned to analysing experimental signals for the possible presence of chaos. Many algorithms were developed to test time series, but these were often found to be unreliable [55]. Some of them require large amounts of data but were nonetheless applied by researchers to small data sets. Some other tests incorrectly identified noisy linear systems as chaotic. Experimental noise presented perhaps the largest problem for many of the tests, which worked well for noise-free data from model systems. Much was learned from resolving these difficulties and to-

day direct tests for chaos have given way to more conservative tests for determinism in time series as well as to surrogate data analysis [57, 56].

4.2 Reconstruction of attractors

When we are studying dynamical systems on a computer, we have access to all the variables of the system. In an experimental setting this is generally not the case as we often can only measure a single variable. How can we study the dynamics of a system when we only have access to time series in a single variable? In 1980, Packard *et al.* [33] proposed a remarkable idea which allows us to reconstruct the attractor of a system from the measurement of a single variable. It was proved by Takens [54] that the attractor reconstructed by the embedding method of Packard *et al.* is diffeomorphic to the attractor of the system, i.e. the two attractors have the same geometric and dynamical properties.

The idea behind the embedding theorem of Takens is to find natural coordinates which can be taken from the time series and plotted in phase space. The most common coordinates used today are delay coordinates, which were originally proposed by Ruelle [33]. Given a time series $\{x(t_0), x(t_0 + \delta t), x(t_0 + 2\delta t), \dots\}$, we reconstruct the attractor by creating delay coordinates and embedding them in a d_E -dimensional space (d_E is called the embedding dimension). The embedding is accomplished by creating d_E -dimensional vectors \vec{V}_i from delay values of the time series so that

$$\vec{V}_i = [x(t_i), x(t_i + \tau), x(t_i + 2\tau), \dots, x(t_i + (d_E - 1)\tau)],$$

where the $t_i = t_0 + i \cdot \delta t$ are the times at which the time series is sampled, δt is the sampling interval, and τ is the embedding delay. In three dimensions, we obtain vectors of the form:

$$\begin{aligned} \vec{V}_1 &= [x(t_0), x(t_0 + \tau), x(t_0 + 2\tau)], \\ \vec{V}_2 &= [x(t_0 + \delta t), x(t_0 + \delta t + \tau), x(t_0 + \delta t + 2\tau)], \\ &\vdots \\ \vec{V}_i &= [x(t_0 + i\delta t), x(t_0 + i\delta t + \tau), x(t_0 + i\delta t + 2\tau)]. \end{aligned}$$

These vectors are then plotted as points in space and joined together to reconstruct the attractor as shown schematically in Fig. 4.1. In Fig. 4.2 we show an example




Figure 4.1: Schematic representation of a 3-dimensional embedding.

of a reconstructed attractor for the Rössler system (Eqs. 3.3). The equations are integrated numerically, with a fixed step fourth order Runge-Kutta method [37]. The time step is $\delta t = 0.001$. An embedding delay of $\tau = 0.5$ is used to reconstruct the attractor. 2000 points from the reconstructed attractor are shown on the right and 5000 points from the actual attractor are shown on the left.

In chapter 5 we will investigate the usefulness of embedding interspike intervals to reconstruct the attractor of the input to a neuron.

4.3 Autocorrelation and dimension

The autocorrelation function of a time series is a measure of the linear correlations between the values at a time t and the values at some time lag (τ) in the past. The autocorrelation function is defined as:

$$A(\tau) = \langle (x(t) - \bar{x})(x(t - \tau) - \bar{x}) \rangle / \sigma^2,$$




Figure 4.2: Left panel: the Rössler attractor for the parameters $a = 0.386$, $b = 2.0$, $c = 4.0$. Right panel: three-dimensional embedding of the x variable with $\tau = 0.5$.

where the angle brackets denote averaging over the whole time series, \bar{x} is the mean and σ the standard deviation of $x(t)$. The autocorrelation function is normalized so that it lies in the unit interval $(-1, 1)$. If $A(\tau) = 1$ for some lag τ , then values of x separated in time by τ are perfectly correlated with each other, i.e. the value of x at time t is always the same as the value at time $t - \tau$. A periodic signal will have $A(\tau) = 1$ when τ is equal to the period of the signal. Positive (negative) values of $A(\tau)$ imply positive (negative) linear correlations between values at time t and $t + \tau$.

The easiest dimension to estimate from a time series is the correlation dimension (taking $q = 2$ in Eq. 3.4):

$$D_{cor} = \lim_{\tau \rightarrow 0} \frac{\log(\sum_i P_i^2)}{\log(\tau)}.$$

Consider the quantity:

$$C(r) = \lim_{N \rightarrow \infty} \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \Theta(r - \|\mathbf{x}_i - \mathbf{x}_j\|),$$

where Θ is the Heaviside step function. This quantity measures the pairs of points that have a distance between them of less than r . It can be shown (see [34] p. 173) that $\sum_i P_i^2 = C(r)$, which leads to a new definition of the correlation dimension:

$$D_{cor} = \lim_{r \rightarrow 0} \lim_{n \rightarrow \infty} \frac{\log(C(r))}{\log(r)}.$$

Thus the correlation dimension can be estimated from the slope of a log-log plot of $C(r)$ vs. r over a range of r where the slope is constant (if such a region exists). These regions are called "scaling" regions in the literature. To estimate the correlation dimension of a time series, one plots $\ln(C(r))$ against r for several values of the embedding dimension (d_E). For a chaotic time series the correlation dimension reaches a plateau as the embedding dimension is increased. This behaviour is not observed in stochastic data.

4.4 Time series and point process prediction

The idea of predicting time series is an old one. In 1927 Yule [60] attempted to predict the sunspot cycle by building a linear model from the data. The object of most statistical time series analysis is to eventually build a model of a system in order to make predictions of its evolution. We have seen that chaotic systems should be predictable in the short term, while their behaviour in the long term resembles that of a purely random process. If our data is from a chaotic system, then we should be able to make short term predictions. In 1987 Farmer and Sidorowich [9] introduced the idea of making predictions using the dynamics reconstructed from the time series by a delay embedding as in Sect. 4.2. Instead of the traditional global fitting of linear forecasting methods, which attempts to make predictions by fitting a model to the whole data set, the reconstructed state space allows us to make local fits to the data¹.

Linear and nonlinear forecasting techniques can be used to test for the presence of determinism in a time series. The ability to make accurate predictions indicates that there is some determinism in the data. Most of the work using nonlinear forecasting to detect the possible presence of chaos in data has concentrated on time series data. In this thesis, we are interested in making predictions on point process data, i.e. measurements of times at which events occur. We use a simple prediction algorithm similar to the one described in [45], which is a variant of the simplex method of Sugihara and May [52]. We make predictions on the interspike interval data by looking for patterns of m successive intervals. This is accomplished

¹Farmer and Sidorowich use a local linear fit; more complicated fits are discussed in [3].

by embedding the data, $\{I_1, I_2, \dots, I_N\}$ in an m -dimensional space. To make our prediction, we take a vector in the data $\vec{x}_n \equiv [I_{n-m+1}, I_{n-m+2}, \dots, I_n]$ and find βN ($\beta < 1$) neighbours $\vec{x}_n^j \equiv [I_{n-m+1}^j, I_{n-m+2}^j, \dots, I_n^j]$ in the data. For this study $\beta = 0.01$ unless otherwise noted, i.e. we are using 1% of the total number of points in the data set to make our predictions. We have found that, for the size of our data sets (a few thousand points), using more neighbours does not significantly increase the accuracy of our predictions.

The predicted interval I_p is the average of the last components of each neighbour iterated forward T_h steps, i.e.

$$I_p = \frac{1}{\beta N} \sum_{j=1}^{\beta N} I_{n+T_h}^j, \quad (4.1)$$

where T_h is defined as the prediction horizon. The predicted value is compared with the actual value of the interval I_{n+T_h} . We repeat the process for each vector in the embedding space. We then compute a normalized prediction error (NPE) by taking the average square error over the data set and dividing by the error we would make in simply predicting the mean of the data \bar{I} . Thus

$$NPE = \left[\frac{\langle (I_p - I_{n+T_h})^2 \rangle}{\langle (\bar{I} - I_{n+T_h})^2 \rangle} \right]^{\frac{1}{2}}. \quad (4.2)$$

If our NPE gives a result that is less than one, we are doing better than simply predicting the mean of the data set. If the NPE is close to zero then the system is likely deterministic.

Linear correlations in the data will obviously give rise to an NPE value that is less than one. To distinguish chaotic effects from linear ones it is useful to characterize the behaviour of the error for chaotic systems². Because chaotic systems exhibit exponential divergence of nearby trajectories the error we make in predicting the time series increases exponentially with prediction time (T_p), T_p is analogous to T_h but for time series. Farmer and Sidorowich have shown this to be true [9]. Linearly correlated stochastic data will not have the dependence of the error on the prediction time. Sugihara and May [52] have suggested a similar approach: instead of looking

²We also use surrogate data methods as described in Sect. 4.5.

at the dependence of the error on prediction time they analyzed the dependence of Pearson's correlation coefficient between the predicted values and the actual values, as a function of prediction time. If we define x as the predicted value and y as the actual value, then Pearson's correlation coefficient can be defined as:

$$r = \frac{\langle x \cdot y \rangle - \langle x \rangle \langle y \rangle}{\sigma(x)\sigma(y)}.$$

Where the angle brackets denote the average over all predicted and actual values, and σ is the standard deviation. The correlation coefficient varies from zero to one for uncorrelated and identical distributions respectively. Sugihara and May showed that for a chaotic system the correlation coefficient decreases with prediction time. They were able to distinguish chaotic systems from deterministic non chaotic systems with additive noise, for which the Pearson's correlation coefficient remained constant with prediction time.

It is possible to generate coloured or autocorrelated noise, which has the property that the correlations between predicted and actual values will decrease with time. Tsonis and Ellsner [58] have shown that it is possible to differentiate between coloured noise and chaotic systems by focussing on the dependence of the correlation coefficient r on prediction time. They have shown that for a chaotic time series:

$$r(T_p) \approx 1 - e^{KT_p}, \quad (4.3)$$

whereas for coloured noise $r(T_p) \approx T_p^\alpha$, where K and α are constants. A graph of $\ln(1 - r(T_p))$ against T_p then appears as a straight line for chaotic time series. In Chap. 6 we examine the possibility of extending this method to ISI sequences.

4.5 Surrogate data

We now want to address the problem of determining the reliability of the calculations used in nonlinear time series analysis. For example, if we find a Lyapunov exponent that is only slightly greater than zero, can we interpret this as a sign that our data is chaotic³? One way of determining the reliability of the computations of properties

³Recall that a positive Lyapunov exponent implies that the system is chaotic.

of a time series, also referred to as “statistics” of the time series in the context of surrogate data analysis, is to keep track of all possible systematic and statistical errors, in order to calculate error bars on the computed values. In this thesis the statistics we will use are the NPE and the correlation dimension.

Calculating error bars can be problematic for many reasons: [56]

- There is often only one short data set available. Hence we cannot get error estimates by repeating our calculations many times.
- Estimating the error requires some model of the underlying process. This is exactly what we do not know.
- The complexity of the algorithms makes finding an analytic expression for the error almost impossible.

Another method of determining the reliability of a statistic was proposed by Theiler *et al.* [57, 56]. They suggest a statistical approach similar to the bootstrapping methods described in [8]. Their method requires the generation of an ensemble of *surrogate data* sets which satisfy a specific *null hypothesis*. The chosen statistic is then computed for the original data and for each set in the ensemble. If the statistic differs significantly between the original and the surrogates then the null hypothesis corresponding to the surrogates can be rejected.

Theiler *et al.* describe several different types of surrogate data and accompanying null hypotheses. The simplest example is a *shuffled surrogate*. These surrogates are derived from the original data by shuffling it like a deck of cards in order to remove all temporal correlations. The surrogates have the same probability distribution as the original data but are otherwise random. The null hypothesis corresponding to this type of surrogate data is temporally uncorrelated noise, such as Gaussian white noise. It is often the case that this null hypothesis can be rejected easily even by linear methods.

Another null hypothesis we would like to reject is linearly correlated noise. An example of linearly correlated noise would be the process:

$$y_i = a_0 + \sum_{j=1}^N (a_j y_{i-j}) + \sigma \epsilon_i,$$

where ϵ_i is an uncorrelated noise process. In order to generate surrogate data corresponding to this null hypothesis, we could fit the parameters a_i and σ to our data and iterate the difference equation. Theiler *et al.* suggest a simple algorithm that does not require us to fit the data. Their algorithm consists of taking the Fourier transform of the original data, then shuffling the phases and taking the inverse transform as the surrogate data. These surrogate data have a Gaussian distribution and the same power spectrum as the original data. By the Wiener-Khinchin theorem, they have the same autocorrelation function. This type of surrogate data is called a *phase randomized surrogate*.

The last null hypothesis we look at is described in [57] as a static nonlinear transform of linearly correlated noise. Under this null hypothesis the data is linearly correlated noise, but the nonlinearity comes from the observation. These surrogate data preserve the linear correlations in the data *and* the probability distribution; thus they are called *amplitude adjusted surrogates*⁴. To make this type of surrogate data, we generate Gaussian random numbers that have the same amplitude ordering as the data, then generate phase randomized surrogates of these Gaussian numbers. The surrogates are simply the data shuffled to have the same amplitude ordering as these phase randomized surrogates. Algorithms for generating surrogate data are given in Appendix B.

These surrogate data methods may be applied to point process data. However, we have found that a modification of the interpretation of the results of surrogate data analysis is required due to the nature of the point process data [26] (see Sect. 6.2). We have chosen to illustrate the difference between raw and surrogate data sets using ISI data from a primary afferent sensory neuron of an electric fish. We are grateful to Dr. Joseph Bastian, University of Oklahoma, for making this data available to us. We will look more closely at this data in Sect. 6.3.2; it is also analyzed in [26, 27].

Fig. 4.3 shows a plot of ISI against ISI number. Notice that the shuffled and amplitude adjusted surrogates are nearly indistinguishable from the original data by eye. In Fig. 4.4 we present the interspike interval histograms (ISIH). The shuffled and amplitude adjusted surrogates have exactly the same ISIH since they are both

⁴These are often called amplitude adjusted phase randomized surrogates as well.

shuffled versions of the original data. The phase randomized surrogates have a Gaussian distribution. Fig. 4.5 is a plot of the autocorrelation function for each set. The amplitude adjusted surrogates have almost exactly the same autocorrelation function as the original data set whereas the phase randomized data has exactly the same AF. This similarity is the reason why these three versions of the data have the same linear correlations. We can see that the shuffled surrogates clearly no longer have any of the correlations that were present in the data.

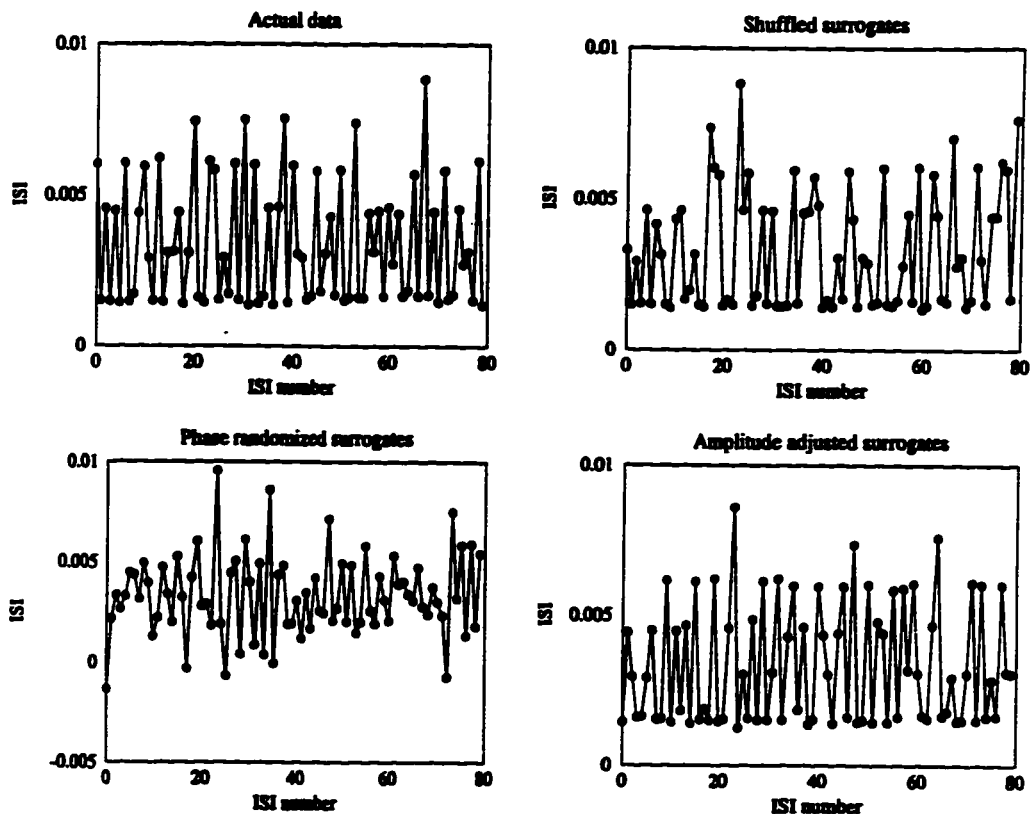


Figure 4.3: ISI vs. ISI number for the data and three surrogates.

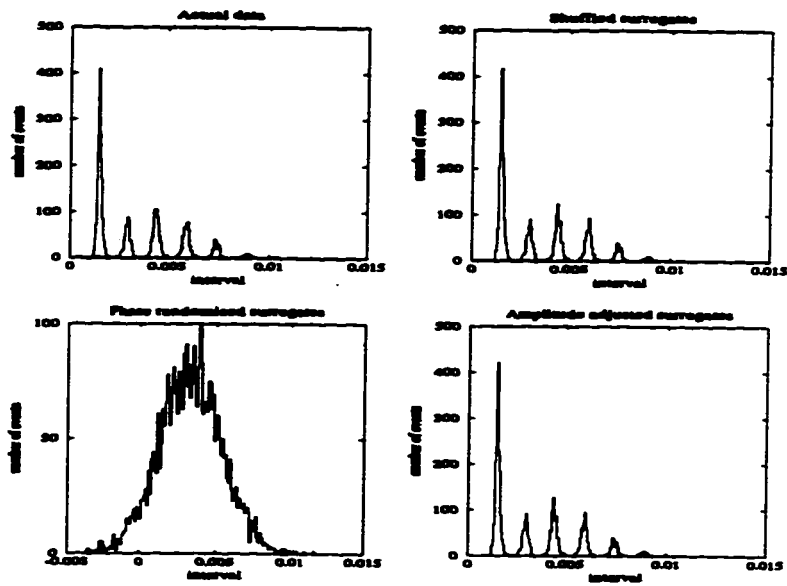


Figure 4.4: ISI histograms for the data and three surrogates.

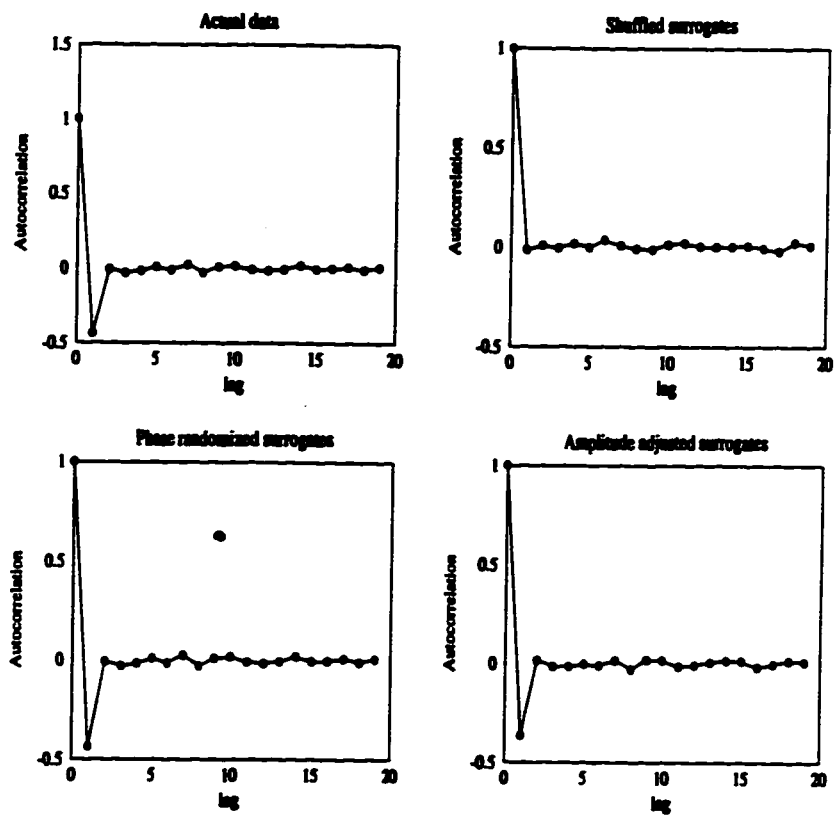


Figure 4.5: The autocorrelation functions of the data and three surrogates.

Chapter 5

Interspike interval embeddings

5.1 Previous studies

Within the past ten years there have been several exploratory studies of the applicability of methods from nonlinear time series analysis to point process data [50, 40, 53, 36, 61, 23, 49, 48, 45, 46]. In particular, there has been much work on the analysis of interspike interval data from neurons and other excitable cells as well as on point process data taken from various well known dynamical systems, such as the dripping faucet [50]. In that study, return maps and various dynamical information measures were applied successfully to the study of inter-drop intervals. The common goal of these studies was to determine whether nonlinear methods are useful tools for characterizing point process data which is suspected to be generated by low-dimensional dynamical system. In this section we present in more detail four of the above studies. We wish to convey some of the reasons for our interest in this field, and give a “historical” perspective to the results presented in this and the following chapters.

In 1985 Rapp *et al.* [40] calculated the correlation dimension of interspike interval data taken from ten spontaneously firing neurons in the brain of a squirrel monkey. They found a small correlation dimension (< 5) in three of their data sets. The other seven data sets had correlation dimensions which could not be determined or which were very high (> 20). The authors concluded that their results were encouraging because they showed that neurons might exhibit low-dimensional behaviour.

This data later failed to pass the surrogate data tests [41] and thus evidence of nonlinearity, let alone chaos, was found to be insufficient.

In 1990 Preißl *et al.* [36] compared the correlation dimensions of the Lorenz attractor to that of ISI embeddings of intervals generated from the Lorenz equations. They generated ISI's by setting a threshold on the z component of the Lorenz equations. Even though the Lorenz model is not a neuron model, they considered that a "spike" was generated when z crosses the threshold going upwards. Downward crossings of the threshold were not counted. In Fig. 5.1 we show the correlation dimension results for the "interspike" intervals. This figure is taken from [36].




Figure 5.1: Correlation integrals for different spike trains generated from the z component of the Lorenz system. The thresholds are as follows: (a) $Z = 35$ (b) $Z = 25$ (c) $Z = 15$ (d) $Z = 12$.

For each value of the threshold, the correlation dimension was calculated for embedding dimension (d_E) from one to six. Larger values of d_E were not taken since the dimension of the attractor is known to be small (it is ~ 2.06 for the chosen parameters). In panel (a) and (b) the correlation dimension was calculated to be 1.92 and 1.79, respectively. In panel (c) and (d), there are two scaling regions, denoted

by $sr1$ and $sr2$. In panel (c) the dimension in $sr1$ is 1.52, whereas in $sr2$ the correlation dimension increases with the embedding dimension. For panel (d) they were unable to determine the correlation dimension in either of the scaling regions. The correlation dimension is still increasing at $d_E = 6$. The authors conclude that the correlation dimensions calculated from the ISI data are not equal to the correlation dimension of the Lorenz attractor, and argued that the discontinuity of the point process does not allow for the calculation of dynamical invariants of a system from embeddings of ISI data. This implies that the attractor reconstructed by embedding the ISI's is not diffeomorphic to the attractor of the underlying dynamical system. This result is not surprising because the attractor reconstructed from these "ISI embeddings" bears very little resemblance to the Lorenz attractor (not shown). It is still an important step in discovering the relationship between point processes and the underlying dynamical system from which the point process is generated.

Zimmerman *et al.* [61] studied neurons from the brain of a rat in which behaviour similar to a focal epileptic seizure was induced by applying a highly concentrated drop of penicillin to the surface of the brain. By standard statistical measures they were unable to differentiate the mean of the pre-penicillin and post-penicillin data sets. When the ISI's are embedded, the pre-penicillin and post-penicillin data sets are clearly differentiated by their geometrical structure. While the results of this study are related to those of this chapter, they are also important to us for another reason: the study illustrates simply that the geometrical nature of the data is sensitive to the presence of penicillin without claiming fractal and/or chaotic structures, a popular trend at the time. The authors were careful in drawing their conclusions, and they tried to understand the limitations of their analyses. This has been a guiding principle in our study.

Longtin [23] has studied the predictability of interspike intervals from periodically forced neurons. Two data sets were studied. The first is taken from a neuron in the somato-sensory cortex of the cat. The neuron is being forced by stimulating the mechanoreceptor in the paw with a periodic signal. The second data set is from a cat auditory nerve fiber in the presence of a pure tone stimulus of fixed frequency and amplitude. Both of these data sets are examples of skipping data discussed in Sect. 2.3.4. Based on surrogate data methods the author finds no evidence for nonlinear determinism in the data sets. The first data set is then transformed into

a series of phases. If the stimulus has a period of T_0 , then the neuron fires near integer multiples of T_0 . The phase is defined as the time between the firing and the closest multiple of T_0 . For these phases the shuffled surrogates are rejected for all embedding dimensions, and the amplitude adjusted surrogates are rejected only for high embedding dimensions. The author concludes that the behaviour of the phase variable is more characteristic of linear stochastic dynamics for two reasons:

- (i) The null hypothesis of linearly correlated noise cannot be rejected for all embedding dimensions.
- (ii) The error decreases as the number of neighbours used to make the forecast is increased. This is characteristic of linear stochastic dynamics [4].

The results of the present chapter deal with the geometrical and dynamical properties of ISI embeddings for chaotically forced neurons. Chapter 6 is an extension of the work of [23].

5.2 ISI embeddings for chaotically forced neurons

The studies quoted in the previous section make extensive use of interspike interval embeddings. ISI embeddings are a straightforward analogy to the delay embeddings discussed in Sect. 4.2 for time series. We begin with a series of N interspike intervals $\{I_1, I_2, \dots, I_N\}$. The ISI's are then embedded into a d_E -dimensional space by creating d_E -dimensional vectors using the delay coordinate method. For $d_E = 3$ we would make vectors:

$$\begin{aligned}\vec{V}_1 &= [I_1, I_2, I_3]^T, \\ \vec{V}_2 &= [I_2, I_3, I_4]^T, \\ &\vdots \\ \vec{V}_{N-2} &= [I_{N-2}, I_{N-1}, I_N]^T,\end{aligned}$$

using an embedding delay of one. These vectors are plotted in d_E -space and their end points are joined by straight lines as in Fig. 4.1.

In Sect. 4.2 we saw that Takens' theorem allows us to reconstruct the attractor of a dynamical system from measurements of only one variable. In 1994, Sauer [45]

demonstrated that it is possible to reconstruct the attractor of a chaotic input to the simple integrate and fire (IF) neuron (described in Sect. 2.4.1) by performing a delay embedding of the ISI's. This result caused some confusion among scientists, for example the following quote from [30]:

“Sauer has recently shown how to analyze such intervals, generated by the threshold crossings of an underlying chaotic process.”

The intervals being referred to in the previous quote are ISI intervals measured between spikes by a neuron in a network of rat hippocampal cells. There is a significant difference between a hippocampal cell of a rat and the generic IF model used by Sauer in [45]. In this chapter we wish to determine if the method of Sauer can be extended to more realistic IF models and the FitzHugh-Nagumo (FHN) model in order to determine its possible applicability to data from real neurons.

The chaotic input to the various neuron models is derived from the x variable of the Rössler equations (Eq. 3.3) with parameters $a = 0.36$, $b = 0.4$ and $c = 4.5$. The input is of the form:

$$S(t) = \delta \cdot (x(t) + C).$$

This type of scaling and shifting of the input is necessary in order to match the Rössler signal to the neuron model. The reconstructed attractors should be compared to the Rössler attractor of Fig. 5.2 below. The main characteristics to note are the funnel shape and the banded structure of the attractor. In this chapter we use the terms “reconstructed attractor”, “reconstruction” and “ISI embedding” interchangeably.

To allow comparison of the reconstructed attractors from different models, we define a sampling rate which tells us how often the neuron model is sampling its input. We use the simplest possible definition of sampling rate which is the mean number of spikes per unit of time. This rate is determined by dividing the sum of the interspike intervals (i.e. the total time of the data set) by the number of spikes. We express this sampling rate in Hz^1 throughout the chapter.

In the following sections we present the results of ISI embeddings for three IF

¹Note that real neuron firing time scales can be as short as msec, which would make the sampling rate in kHz.




Figure 5.2: The Rössler attractor from which the input to the neurons is obtained. Every fiftieth point of the numerical integration (with a timestep of 0.001) is shown.

models and the FHN model. We then compare the quality of the reconstruction and discuss some of the implications of our results.

5.2.1 Generic integrate and fire device

This is the model used in [45] and further studied in [46]. We introduced it as a generic IF device in Sect. 2.4.1. The neuron activity $V(t)$ is a time integration of the input signal $S(t)$:

$$V(t) = \int_{T_{n-1}}^t S(t') dt'.$$

This equation is valid for $t > T_{n-1}$, the last firing time. A spike is generated when $V(t)$ reaches a threshold Θ : $V(T_n) = \Theta$; at this point the voltage is reset to zero and the integration continues. The interspike intervals are defined as $I_n = T_n - T_{n-1}$. We use the same input signal as in [45]: $S(t) = x(t) + 40$. The constant is added to the Rössler x signal to make $S(t)$ positive everywhere, hence the integral is always increasing towards the threshold. In order to determine the spiketimes we integrate the Rössler equations using a fourth order Runge-Kutta method with time step 0.001. The contribution to the integral is computed every time step. When $V(t)$

crosses the threshold the time step is reduced (multiplied by a random number in $(0, 1)$) and the step is repeated until $|V(t) - \Theta| < \Delta$ (in general $\Delta = 10^{-6}$). The threshold is fixed for each simulation, but it can be varied to change the sampling rate: as Θ increases the sampling rate decreases.

Fig. 5.3 shows the results of the ISI embeddings for $\Theta = 35$ and $\Theta = 90$. In the left panel we have a small portion of $x(t)$, the black circles represent the times at which a spike was generated. The attractor reconstructed by embedding the ISI's is shown in the right panels. The reconstructed attractor for $\Theta = 35$ is a good reconstruction of the Rössler attractor, it has the banded structure and the funnel shape of the original attractor. The sampling rate for this reconstruction is 1.15 Hz. The reconstructed attractor for $\Theta = 90$ bears some geometrical resemblance to the Rössler attractor, however it is not a good reconstruction. The sampling rate is 0.45 Hz.

We have calculated the nonlinear predictability of the intervals as a way to assess the quality of the reconstruction. The predictability results are shown in Fig. 5.4. We find that the better reconstruction has a lower normalized prediction error (NPE) than the poor one, as expected. Nevertheless, the poor reconstruction still has a low NPE value which implies that there is still much information about the deterministic system in these ISI's.

5.2.2 Glass-Mackey integrate and fire model

In this section we present ISI embedding results for the Glass-Mackey IF model described in Sect. 2.4.2. A preliminary account of the results in this section can be found in [38]. For this model the neural activity increases in a linear fashion towards a chaotic threshold $S(t)$:

$$\begin{aligned} V(t) &= \alpha t + \beta, \\ V(t^+) &= 0 \text{ if } V(t) = S(t). \end{aligned}$$

We use $S(t) = x(t) + 10$ as our chaotic threshold. The constant is added to make $S(t)$ positive everywhere since the threshold must always be above the resetting point of the voltage. The spiketimes (t^*) are determined by integrating the Rössler equations with a time step of 0.001. When $V(t)$ crosses the threshold, a linear interpolation

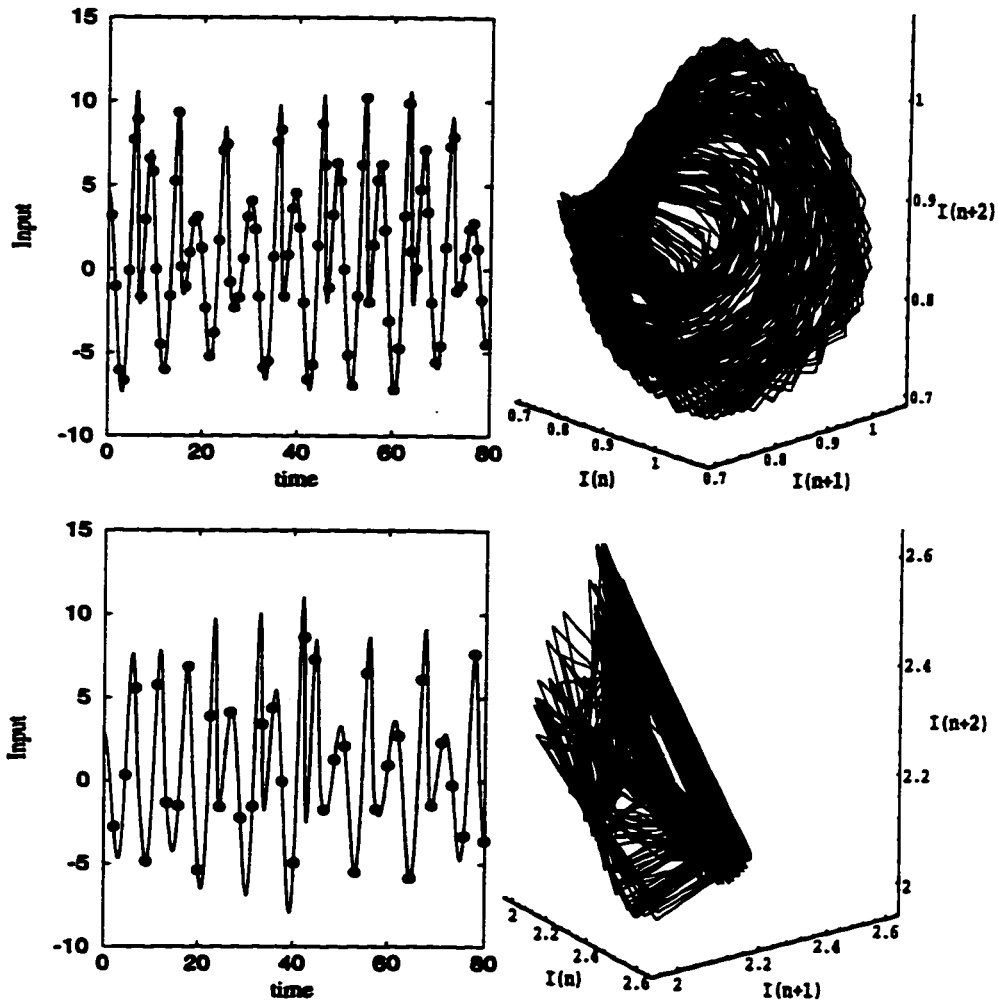


Figure 5.3: ISI embeddings for the generic IF model. Top: $\Theta = 35$. Bottom: $\Theta = 90$. Left panels: the x time series (line) with the spiking times superimposed as black circles. Right panels: 500 points from the reconstructed attractors joined by straight lines.




Figure 5.4: The NPE for the two ISI sequences of the generic IF model.

of $S(t)$ is performed to find the intersection t^* , such that $S(t^*) = V(t^*)$.

Reconstructions are shown for $\alpha = 10$ and $\alpha = 1$ ($\beta = 0$ in both cases) in Fig. 5.5. Both left panels show the dynamics $V(t)$ and the threshold. The sampling rate for these two reconstructions are 1.15 Hz and 0.15 Hz respectively. The reconstruction for the high sampling rate still looks like the Rössler attractor, it has the funnel shape and a somewhat banded structure. The reconstruction for $\alpha = 1$ does not look like the Rössler attractor, a consequence of the small sampling rate. In Fig. 5.9 we show the NPE results for both reconstructions. Once again the better reconstruction has a lower NPE, however the poor reconstruction still has $\text{NPE} < 1$.

5.2.3 Leaky integrate and fire model

In this model, described in Sect. 2.4.3, the activity of the cell is governed by:

$$\dot{v} = -\sigma v + S_0 + S(t),$$

with the chaotic input $S(t) = \gamma(x(t) + 10)$, $\sigma = 0.5$ and $S_0 = 0.4$. A spike is generated whenever $v(t) = 1$ (the threshold is set to 1). Recall that the dissipative term is to account for the tendency of the membrane potential to always decay to the resting potential when the membrane is excited (as long as the threshold is not crossed).

To determine the spiketimes, we numerically integrate the coupled differential equations for $S(t)$ and $v(t)$ with a fixed time step of 0.001. When the threshold is

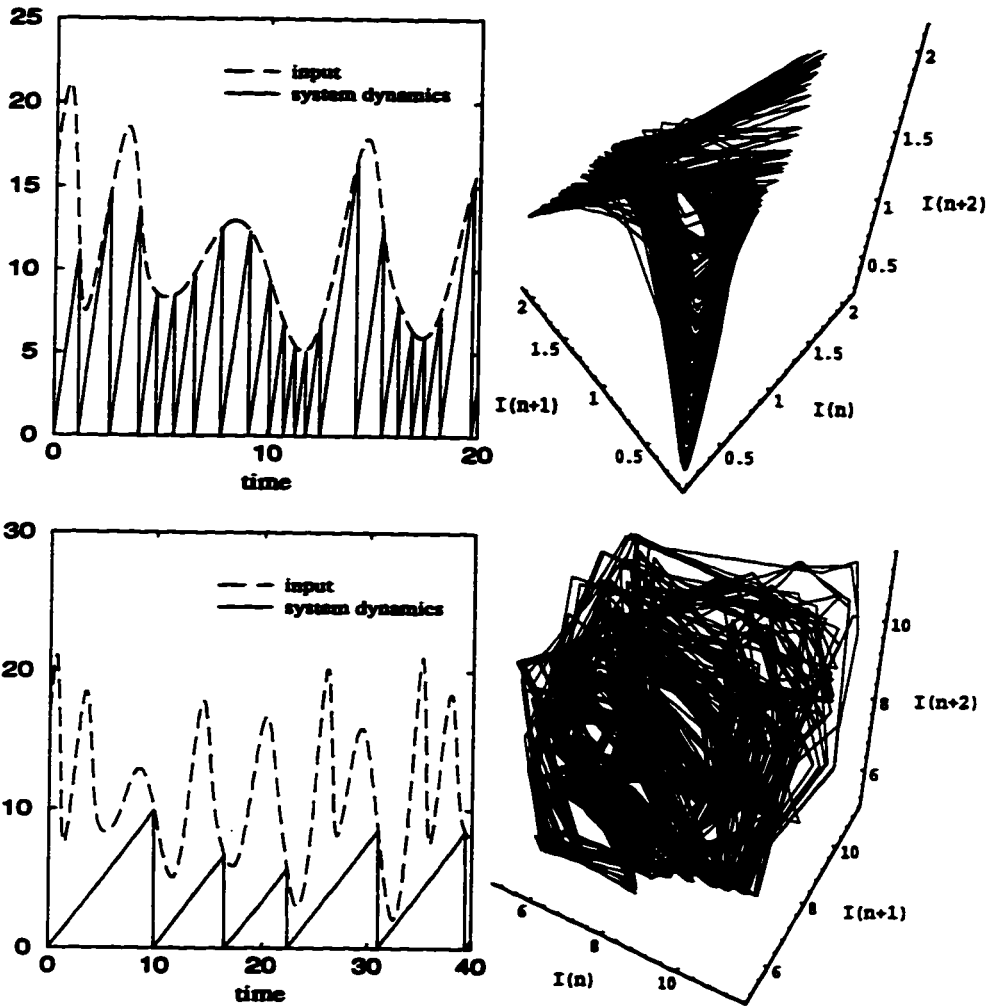


Figure 5.5: ISI embeddings for the Glass-Mackey IF model. Top: $\alpha = 10$. Bottom: $\alpha = 1$. Left panels: dynamics of the model $V(t)$ and threshold $S(t)$. Right panel: 500 point attractor reconstructions.




Figure 5.6: The NPE for the two ISI sequences of the Glass-Mackey IF model.

crossed we perform a linear interpolation of the two points straddling the threshold to determine t^* such that $v(t^*) = 1$. The forcing parameter γ is varied to increase the amount of influence the Rössler system has on the interspike intervals. As γ increases, so does the sampling rate since the enhanced signal brings the voltage to threshold more often.

Figures 5.7 and 5.8 show the ISI embeddings for $\gamma = 0.154$ (a sampling rate of 1.15 Hz) and $\gamma = 0.01$ (sampling rate of 0.30 Hz). For $\gamma = 0.154$ we see that the reconstructed attractor has a certain resemblance to the Rössler attractor, but this is not as good a reconstruction as the previous examples. The attractor seems to be evolving almost entirely in a planar region and there is no banded structure. In Fig. 5.8 the results for $\gamma = 0.01$ are presented. The left panel shows 2000 points whereas the right panel shows 500 points with lines joining them. The object has some interesting geometrical structure that may be related to the Rössler attractor but it does not reconstruct dynamics of the input attractor.

In Fig. 5.9 we present the results for forecasting the ISI sequences above. Once again the better reconstruction has a lower NPE but the NPE for the sequence with the lower sampling rate is still less than one.

5.2.4 FitzHugh-Nagumo equations

We now shift our focus to a more complicated model, the FitzHugh-Nagumo model of membrane excitation, described in Sect. 2.3.2. This model is a set of two coupled differential Eqs. 2.3. We modify the equations by adding a chaotic input current

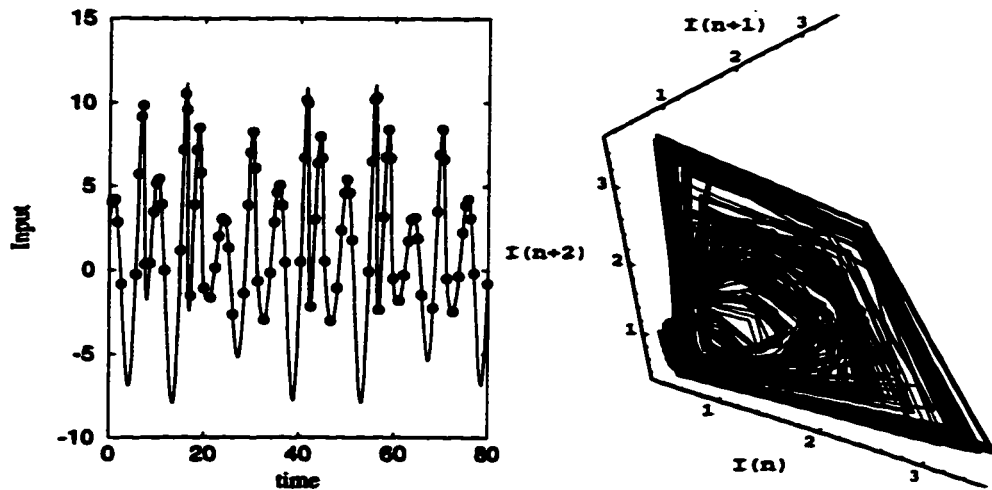


Figure 5.7: ISI embeddings for the leaky IF model with $\gamma = 0.154$. Left panel: Rössler x time series with spike times superimposed. Right panel: 500 point attractor reconstructions.




Figure 5.8: ISI embeddings for the leaky IF model with $\gamma = 0.01$. Left panels: 2000 point reconstruction. Right panel: 500 points of the same attractor joined by straight lines.




Figure 5.9: NPE for the two ISI embeddings of the leaky IF model.

$S(t)$:

$$\begin{aligned} \epsilon \frac{dv}{dt} &= v(v-a)(1-v) - w + S(t), \\ \frac{dw}{dt} &= v - dw - b, \end{aligned}$$

where $S(t) = \mu x(t)$ ($x(t)$ is the x-component of the Rössler equation). The modified model is a set of five coupled ordinary differential equations. To determine spiking times we numerically integrate the equations and find the times at which the v variable crosses a threshold in the upwards direction via a linear interpolation as previously discussed. In this case we set the threshold to 0.7. The scaling factor μ is varied to increase the effect of the chaotic input on the neuron model. The other variables are fixed at $a = 0.5$, $d = 1.0$, $b = 0$, $\epsilon = 0.005$.

For these choices of parameters the FHN equations cannot generate spikes without the input current. We have found a minimum value for spike generation is $\mu \approx 0.007$. An upper limit on μ is imposed in order to keep the magnitude and range of $S(t)$ from being much larger than those of v . The range of $v(t)$ when spikes are being generated is approximately $(-0.4, 1.1)$. A value of $\mu = 0.075$ gives $S(t) \in (-0.6, 0.9)$. At this value of μ some of the spikes (the first one in Fig. 5.10(b)) are already distorted. These spikes are not action potentials (see Figs. 2.2 and 2.3) because the voltage remains at a high value for a long time compared to a normal action potential for this model. Since we want to investigate the effect of the chaotic input current on the FHN model we take the value of $\mu = 0.075$ as our largest.




Figure 5.10: Voltage dynamics of the FHN model forced by a chaotic input current; (a) $\mu = 0.04$, (b) $\mu = 0.075$.

In Fig. 5.11 we show the ISI return maps (a 2-dimensional embedding) for $\mu = 0.04$ and $\mu = 0.075$. Neither of these bear any resemblance to the Rössler attractor. The result of NPE calculations are shown in Fig. 5.12. Notice that both still have NPE less than one which indicates determinism in the ISI sequence. It is also important to note that the NPE is smaller for the ISI sequence with a lower forcing amplitude. This is an indication that the neuron in this case is playing a more important role than simply passing along information to the next neuron: its dynamics are affecting the output (as is generally expected from a more realistic model). If we forget about model neurons for a moment and let μ increase, there is a value ($\mu \approx 0.3$) beyond which the return map does not change. The return map for $\mu = 0.4$ is shown in Fig. 5.13. We see that this map also does not look like the Rössler attractor; rather it looks like the attractor of a two dimensional map (recall the Hénon attractor of Fig. 3.3) or the projection of an attractor of higher dimension. The attractor shown in Fig. 5.13 is similar to those that are generated from intervals between threshold crossings of the x variable of the Rössler equations. This is not surprising since the dynamics of the FHN equations are dominated by the forcing current from the Rössler system. In this case it seems advantageous to think of the ISI embedding as some kind of map akin to the Poincaré map of Sect. 3.4.




Figure 5.11: ISI return maps for FHN model with a chaotic input current. Top: $\mu = 0.04$. Bottom: $\mu = 0.075$. The right panels show the same points as the left without lines joining them. Note the scale difference from (c) to (d).




Figure 5.12: NPE results for the FHN model with low (realistic) forcing.




Figure 5.13: ISI return map for the FHN equations with large magnitude chaotic forcing, $\mu = 0.4$.

5.3 Discussion

5.3.1 Integrate and fire models

We have studied the possibility of reconstructing the dynamics of the input to four neuron models by embedding their interspike intervals. In each case the input is a chaotic signal generated from the x variable of the Rössler equations 3.3. It was shown in [45] that this procedure worked well for the simple integrate and fire device of Sect. 2.4.1. We have applied the method to two other integrate and fire models as well as the FHN model. For the integrate and fire models we showed that the attractor of the input can be reconstructed if the sampling rate is high (in our case it is 1.15 Hz). As the sampling rate is reduced the ISI's become uncorrelated and the attractor cannot be reconstructed. Nevertheless the ISI's show some level of predictability even for the smallest sampling rate (0.15 Hz for the Glass-Mackey IF model of Sect. 5.2.2). Surrogate data analysis to assess whether the predictability has linear or nonlinear origin is performed in Chap. 6.

In order to compare the reconstruction across the different integrate and fire models, we have adjusted the parameters of each model to give one example of a “good” reconstruction with the same sampling rate (1.15 Hz). The first test one should always apply to a reconstruction is the ocular inspection test (i.e. look at the data). To judge the quality of the reconstruction we are looking for the banded structure and the funnel shape of the Rössler attractor as well as for a

smooth distribution of points on the reconstructed attractor. Sharp edges are more prominent in a poor reconstruction. If we apply this method to the ISI embeddings in Figs. 5.3, 5.5, and 5.7 we can see that the best reconstruction of the Rössler attractor is achieved for the generic IF device proposed in [45]. The quality of the reconstruction diminishes for the Glass-Mackey IF model and for the leaky IF model.

Quantitatively, we examine two properties of the attractor in order to determine the quality of the reconstruction: the normalized prediction error (NPE) and the correlation dimension. Recall that for a chaotic system the NPE should be small at short prediction time. Hence the better reconstructions should have the lowest prediction error. In Fig. 5.14 we show the NPE as a function of α for the Glass-Mackey IF model. As α increases, the quality of the reconstruction increases (according to the ocular inspection method) and the NPE decreases. A similar behaviour is obtained with the leaky IF model by varying the variable σ . However when σ is varied past a certain level the NPE does not drop monotonically. There are certain values of σ for which the NPE drops significantly, apparently due to some kind of quasi-phase locked behaviour. In fact, the ISIH's are particularly simple (they have one or two well defined peaks) in these cases. This is not surprising since the leaky IF model is more complicated than the other two models. Results for the three IF models with a sampling rate of 1.15 Hz are shown in Fig. 5.15. The generic IF device has significantly lower NPE than the other two models, hence it is the better reconstruction. By this method it is difficult to distinguish the Glass-Mackey model from the leaky model.

The correlation dimension of the Rössler attractor shown in Fig. 5.2 is approximately 2.2. If we look at the correlation dimension of the reconstructed attractors, the better reconstruction will be the one with the dimension closest to the true value². In Fig. 5.16 we show the correlation dimension obtained by averaging the values calculated by two different algorithms [43, 51]. One of the algorithms gives lower dimension estimates than the other. We can see that the reconstruction for the generic IF has a higher dimension than the other two models.

²Recall that a true reconstruction should have the same correlation dimension as the original attractor.




Figure 5.14: NPE as a function of α for the Glass-Mackey IF model. Calculations were made with ISI sequences containing 2000 intervals.

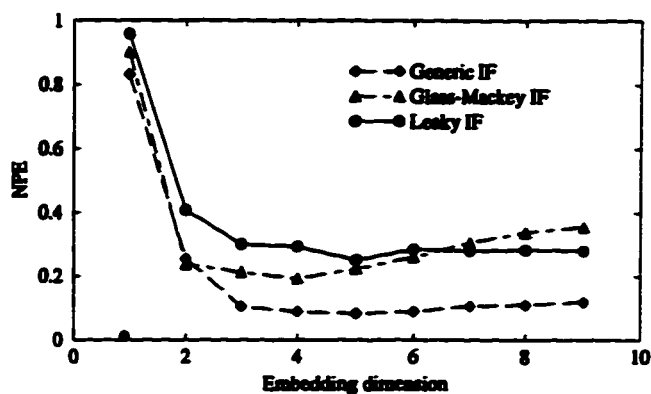


Figure 5.15: A comparison of the NPE for the reconstructed attractors of the three integrate and fire models.




Figure 5.16: A comparison of the correlation dimension for the reconstructed attractors of the three integrate and fire models.

5.3.2 FitzHugh-Nagumo model

Our results for the FHN model are very different. For realistic values of the forcing parameters (i.e. such that the model is still considered a neuron model) we are unable to reconstruct the Rössler attractor. When we allow the forcing term to become large we are still unable to reconstruct the Rössler attractor, but we get ISI return maps that are similar to those one would get by simply setting a threshold on the x variable of the Rössler equations.

5.4 Conclusion

Our main goal in this chapter was to examine the possibility of extending the results presented in [45] to more realistic IF models and to the excitable FitzHugh-Nagumo model. We have shown that the input attractor to an IF neuron can be reconstructed if the sampling rate is large enough. The quality of the reconstruction deteriorates as the IF models become more realistic. For the FitzHugh-Nagumo model we were unable to reconstruct the dynamics of the chaotic input to the neuron. The combination of results above severely restricts the hope of reconstructing input dynamics using ISI embeddings with experimental ISI sequences, since real neurons generally exhibit behaviour that is more complex than that of the FHN model neuron. Reconstructing the attractor of the input may be possible for a small subset of real neurons (known as rate coders) which fire regularly at a frequency that depends on the magnitude of the input. The leaky IF shown in Fig. 5.7 behaves roughly like a rate coder, all the spikes are generated when the input signal is “high”. These neurons are very simple and tend to be “information transformers”, i. e. they take a continuous time signal and encode it into a series of ISI’s which are passed to other neurons. These ISI’s follow the input signal amplitude, i.e. when the input signal is strong the ISI’s are small (the neuron is firing fast), and vice versa. Unfortunately, such neurons also tend to be noisy, their spontaneous firing activity (i.e. with no input) is high and apparently random. Even for constant input there would be a distortion of the ISI’s. The possibility that noisy rate coders produce ISI’s from which inputs can be reconstructed will be further explored in the near future.

The ISI embeddings for the IF models at small sampling rate and for the FHN

model all point in a similar direction, that is, to think of these ISI embeddings as a map of the dynamics similar to a Poincaré map. A Poincaré map is known to reduce the dimension of the system by one. It is yet unclear how the map produced by extracting the ISI's is related to the dynamics of the system. At high sampling rates the results in Fig. 5.16 show that the combination of the simple IF dynamics and thresholding process are not equivalent to a Poincaré map since the dimension values differ from $2.2 - 1 = 1.2$. Also, the attractor from ISI's are very similar to those of the input. Analytical work on simple systems, analogous to that in [22], may shed some light on this question. We may also be able to approximate the dimension of the map as the embedding dimension which yields the lowest NPE. If we look at the NPE results for the low sampling rate IF data sets and for the FHN equations, we see that all give estimates for the dimension from 1 to 4, but there is no clear rule as to the dimension reduction.

The predictability results we have shown are more encouraging. We can see that even for embeddings that do not reconstruct the attractor, the NPE is still low (for example see Fig. 5.6). This means that there is still information about the deterministic process formed by the input and the neuron dynamics in the interspike intervals. In the next chapter we use the NPE and surrogate data methods to characterize the correlations in the data from the IF models and to test for the presence of determinism and nonlinearity in experimental ISI sequences.

Chapter 6

Analysis of ISI sequences

6.1 Introduction

In the previous chapter we saw that all the data sets had normalized prediction error (NPE) less than one. This indicates that the ISI sequences contain information about the deterministic process which generated them. In this chapter we will first investigate the nature of the correlations in the ISI sequences using the surrogate data method described in Sect. 4.5, with the NPE as our discriminating statistic. We previously defined three kinds of surrogates: shuffled, phase randomized, and amplitude adjusted. Recall that the phase randomized surrogates preserve the linear correlations in the data and have a Gaussian distribution. The amplitude adjusted surrogates preserve the linear properties of the data as well as the distribution of the original data set. All these surrogates were conceived with a specific null hypothesis in mind. In Sect. 6.2 we will discuss the meaning of these null hypotheses for point process data such as ISI's.

We will also investigate the possible extension of the method of Tsonis and Ellsner [58], for detecting chaos in time series, to ISI sequences. They argue that for a chaotic time series the Pearson's correlation coefficient (r) varies with respect to prediction time according to Eq. 4.3. We try to draw a parallel from prediction time (T_p) to prediction horizon (T_h) and accordingly verify that for chaotic ISI sequences:

$$r(T_h) \approx 1 - e^{-KT_h}.$$

Finally we will apply the NPE and surrogate data methods to real ISI data sets from neurophysiological experiments.

6.2 Surrogate data and point processes

6.2.1 Interpretation for point processes

The surrogate data method was originally proposed for traditional time series. Because time series are measurements of a physical system (or of one component of a physical system) evolving in continuous time, the surrogates were designed having in mind a specific null hypothesis for the system's dynamics. Thus, if we can distinguish the surrogates from the actual time series, we can reject the null hypothesis (e.g. linearly correlated noise in the case of phase randomized surrogates) as a process for the generation of the time series. In the case of neural point processes, we are not measuring the variables of the system directly, we only have measurements of the times at which one of the variables reaches a certain threshold. In general we do not expect to be able to reject specific processes for the generation of spikes with these types of surrogates, especially when the ISIH for the raw data is significantly different from a Gaussian. Distinguishing these surrogates from the raw ISI sequence allows us to only say something about the ISI's themselves rather than about the neuron dynamics. This may seem like a trivial point, but it is extremely important to realize the limitations of this method for ISI sequences [26]. On the other hand one can generate more biophysically motivated surrogates by e.g. simulating ISI sequences from specific neuron models [27]. One can then compare the behaviour of statistics such as the NPE on the raw and model ISI sequences, to see if the model is compatible with the experimentally observed point process. An example of this type of analysis is given in Sect. 6.4.

The modified null hypotheses for point processes are quite simple: for the shuffled surrogates the null is that the ISI's are uncorrelated. We will use the phase randomized and amplitude adjusted surrogates to test the effect of the linear correlations in the data on the NPE. The mechanics of generating surrogate data are still the same, e. g. to generate a shuffled surrogate we still shuffle the data set. These methods are described in detail in Sect. 4.5 and in Appendix B.

6.2.2 Implementation of the surrogate data method

From the ISI sequences we create 10 sets of each surrogate. We calculate the NPE for each surrogate set. The NPE is the average for the 10 sets with an error bar equal to $\pm 2\sigma$, where σ , is the standard deviation estimated from the 10 surrogate NPE values. This procedure is repeated for several “embedding dimensions” (m). Embedding the ISI’s in an m -dimensional space is a convenient method for finding similar patterns of m successive intervals. If we think of the ISI embeddings as a map of the dynamics similar to the Poincaré map, then m does take on the meaning of a dimension. If there is a minimal value of m which yields a lower NPE value than the others, then we may take that m as a first approximation to the dimension of the map governing the generation of successive ISI’s. For example a low NPE for $m = 2$ suggests that I_{n+1} is a function of I_n and I_{n-1} .

There are a few important notes on the graphs of this chapter:

- The legend of Fig. 6.1 applies to all graphics in the section. Solid squares denote real data, filled diamonds, triangles and circles represent shuffled, amplitude adjusted and phase randomized surrogates respectively.
- There is an error bar on all NPE points for surrogates. If it cannot be seen, the size of the symbol should be taken as the error.
- The scales are not all the same. It is important to note the scale when comparing NPE for different data sets.

6.2.3 Analysis of ISI sequences from IF models

The first data sets we analyze are taken from the same IF models as in Chap. 5. Each of the six sequences is composed of 3000 ISI’s. Recall that for each model, there were two data sets, one of which represented a “good” reconstruction of the attractor of the input to the neuron, and one which represented a “poor” reconstruction. In the following four figures, we show the NPE results for all the data sets with the three types of surrogates. In every case we are able to distinguish the data from the shuffled surrogates. It has been our experience that for any shuffled surrogates the NPE is always slightly above one (it would converge to one as the data set is




Figure 6.1: NPE for generic IF and surrogates. Left: $\Theta = 35$. Right: $\Theta = 90$. NB: The legends also apply to all other graphs in this chapter.




Figure 6.2: The NPE for the Glass-Mackey model and surrogates. Left: $\alpha = 10$. Right: $\alpha = 1$.




Figure 6.3: The NPE for the leaky IF model and surrogates. Left: $\gamma = 0.154$. Right: $\gamma = 0.01$.

made larger). For this reason we will no longer include the shuffled surrogates in our analysis.

The amplitude adjusted and phase randomized surrogates are also distinguishable from the data sets in all cases. For some of the data sets we are not able to distinguish the surrogates and the data at $m = 1$. This is because the $m = 1$ case is not a good predictor, it is simply the average of the 30 (1% of the total) closest intervals iterated forward one step.

Some of the power of the surrogate data method can be seen in Fig. 6.2 for $\alpha = 1$. The data set is shown in Fig. 5.5 on the bottom right panel. By simply inspecting the embedding there seems to be some determinism in the data. The NPE calculations (Fig. 5.6) reveal that there is some predictability in the ISI sequence. Now, after the surrogate data test, we can see that the structure in the attractor and the predictability are not caused simply by linear correlations in the data. We would not be able to tell that the data set was generated by a nonlinear system without this test.

6.2.4 Predictability as a function of prediction time

The results for Pearson's correlation coefficient (r) are shown in Fig. 6.4. Recall that for a chaotic time series a plot of $\ln(1 - r)$ as a function of the prediction time will give a straight line. In this case we have plotted $\ln(1 - r)$ as a function of prediction horizon for three known chaotic ISI sequences.

Predictions were made on each set using the embedding dimension which yielded the lowest NPE in the calculations above. We see that for all the poor reconstructions ($\Theta = 90$, $\alpha = 1$, $\gamma = 0.01$) there is a tendency for $\ln(1 - r)$ to climb towards zero and level off. The sampling rate for these data sets is much smaller than for each of their counterparts, thus making successive ISI's less correlated than those for the higher sampling rate. Our ability to predict the sequence falls off rapidly with prediction horizon (from an already low level). The set of predicted and actual values is less correlated which makes the correlation coefficient and $\ln(1 - r)$ tend to zero. For the ISI sequences obtained with a high sampling rate, the application of this method to point processes shows some promise. The leaky and generic IF models have a linearly increasing tendency, whereas the first four points for the

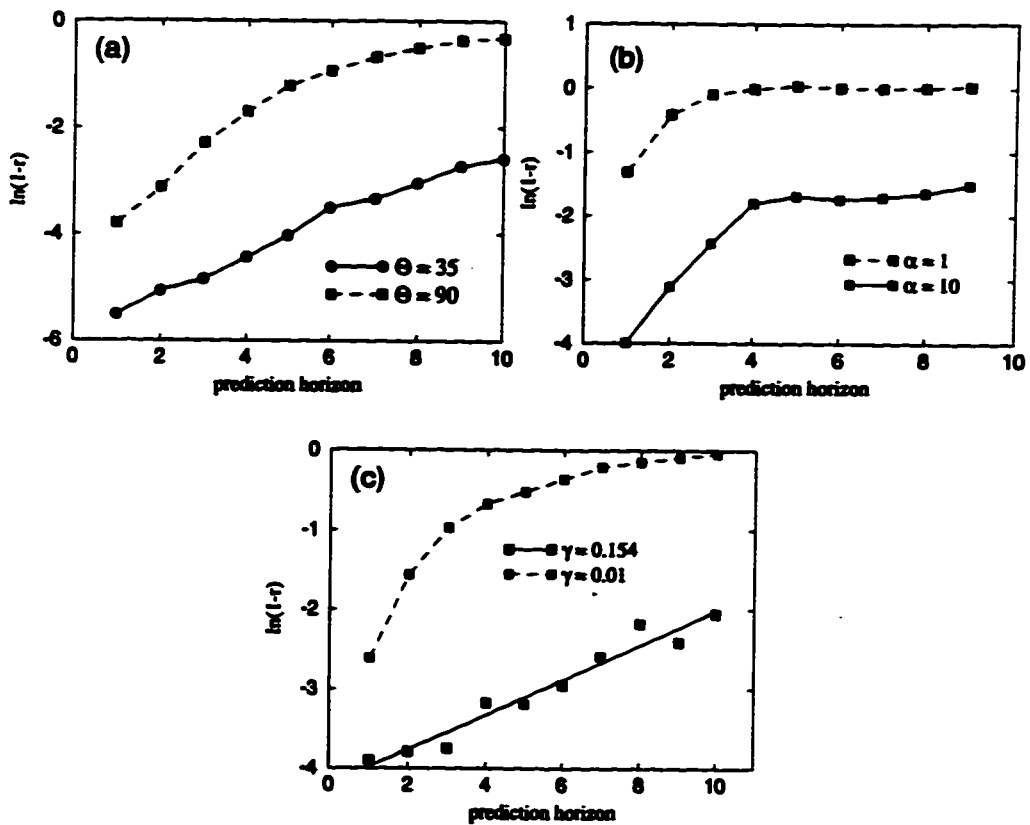


Figure 6.4: $\ln(1 - r(T_h))$ as a function of prediction horizon for ISI's from the IF models. (a) generic, (b) Glass-Mackey, (c) leaky.

Glass-Mackey model are also on a straight line.

This shows that there is promise for extending this method directly to ISI sequences by replacing the prediction with the prediction horizon. We must again caution that a straight line in the graph is only an indication that the data set *may* be chaotic and does not provide a proof of chaos.

6.3 Analysis of experimental ISI sequences

6.3.1 Data from a mammalian thermoreceptor

We now turn our attention to real ISI sequences from neurophysiological experiments. The first data sets we analyze are taken from thermoreceptors (the lingual nerves in the cat). These nerves encode a steady state temperature into a characteristic pattern of action potentials. We analyze two data sets from this experiment. Data set A is shown on the top two panels of Fig. 6.5 and contains 2679 ISI's. Data set B occupies the bottom panels of Fig. 6.5 with 3807 data points. The graphs that we show to describe the data are plots of ISI vs. ISI number and a return map (a 2-dimensional embedding without lines joining the points) of the data. These are also known as "raster plots" or "joint interval histograms" in the physiology literature. Both data sets have a bimodal distribution of long and short intervals as can be seen from the two distinct "stripes" in the raster plot. The return map for data set A shows a clear preference for long intervals to be followed by short intervals and for shorts to be followed by longs. The cluster of points in the center of the map shows that there are a number of long-long events. There are also a small number of short-short events. Data set B is different because, as the return map shows, it contains no short-short events; all short intervals are followed by a long one. There are three times as many long intervals as short ones. These linear correlations will obviously make the data sets predictable; we would like to know if there is any predictability that may be due to other types of correlations among the intervals.

The predictability results for both data sets are shown in Fig. 6.6. We can reject the phase randomized surrogates in both cases because the error bars do not overlap with the data set. Because the values for the phase randomized surrogates and

the real data are so close, we know that the linear correlations accounts for most of the predictability in these data sets. Nevertheless there is some evidence that there are nonlinear correlations in these data sets. Analysis on longer ISI sequences (with more short-short and long-long events) will clarify this. The predictability for the amplitude adjusted surrogates is different from that of the phase randomized surrogates. Since the latter have the same linear correlations with a different distribution, this discrepancy may indicate that the distribution of the ISI's can influence the forecastability. This effect may also be caused by the slight nonstationarity of the data sets which may affect the algorithms for generating the surrogates.

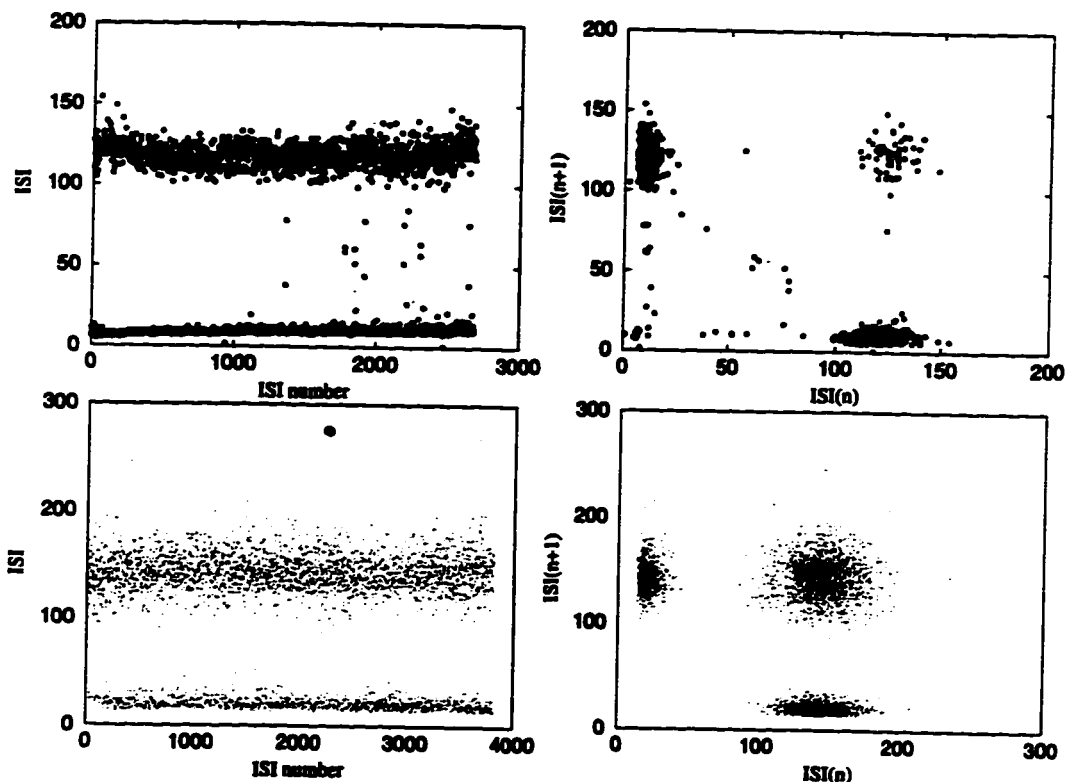


Figure 6.5: Raster plot (left) and return map (right) for the data sets from a cat thermoreceptor. Top: data set A. Bottom: data set B. Data provided by Dr. Hans Braun, Philipps-University Marburg, Marburg, Germany (ISI units are *ms*).




Figure 6.6: NPE for thermoreceptor data and surrogates. Left: data set A. Right: data set B.

6.3.2 Data from weakly electric fish

In this section we examine two data sets from different neurons of the weakly electric fish *Apteronotus leptorhynchus*. The fish produces an AC electric field, known as the electric organ discharge (EOD). In order to electrolocate and electrocommunicate, the fish monitors modulations of this electric field caused by surrounding objects or by EODs of other fish. The detection and monitoring is accomplished by a network of neurons. The first layer of the network consists of electroreceptors which transduce modulations and distortions of the electric field into temporal patterns of action potentials. The EOD is driven by one of the most accurate biological pacemakers known. Once the signal is processed by the electroreceptors, action potentials are sent along their axons to the next layer (known as segments).

Data from the primary afferent neuron

In this section we will be looking at data from the primary afferent neuron of the fish which is known as an electroreceptor. The measurement is taken from an axon which is driven by an electroreceptor; under controlled conditions where the only input to the electroreceptor is the EOD stimulus. Under these conditions the receptor is phase locked to the EOD but exhibits the skipping behaviour described in Sect. 2.3.4. Plots of the data as well as its surrogates are shown in Figs. 4.3, 4.4, and 4.5 and 6.7. The data set contains 9165 intervals.

When predictions are made on the data we see (Fig. 6.8) that there is some

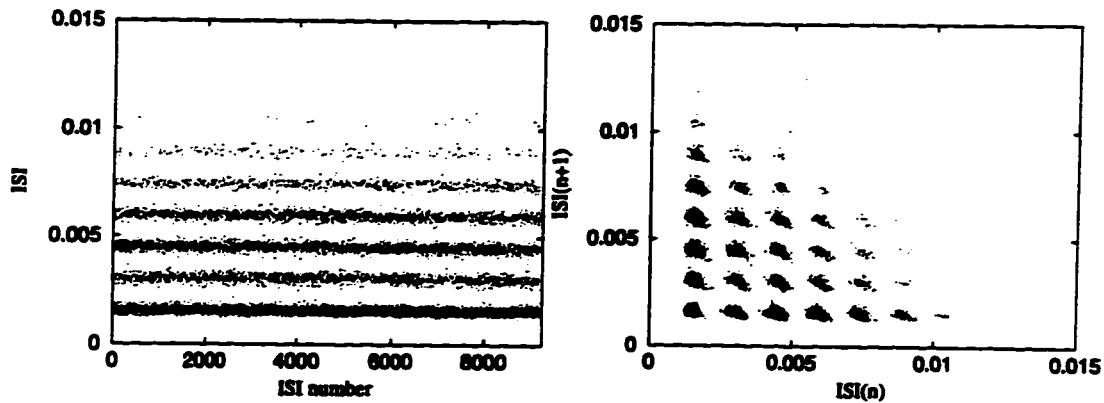


Figure 6.7: Raster plot and return map for data from the primary afferent neuron of the weakly electric fish. Provided by Dr. Joseph Bastian, University of Oklahoma.

predictability. We also know from Fig. 4.5 that there are linear correlations in the data. Fig. 6.8 shows that we can reject both the amplitude adjusted and phase randomized surrogates. We can say that there is more to the correlations in the data than linear ones. Investigation with the method of Tsonis and Ellsner was inconclusive.




Figure 6.8: NPE for data from a primary afferent neuron and surrogates. Data provided by Joseph Bastian, U. of Oklahoma (ISI units are *ms*).

Data from the lateral segment

This is the data taken from the target neurons of the electroreceptors, in the electrosensory lateral line lobe of the same species of weakly electric fish. In other words

these cells are driven by the “skipping” firing patterns of the primary afferents. The data appears in Fig. 6.9, which shows 5000 ISI’s from a stationary firing pattern. The mean value of the intervals is 0.034, hence the majority of intervals are very short: 75% are less than 0.004. In the bottom panels of Fig. 6.9 we show an expansion of the region near the origin. The return map shows some structure among the short intervals. The predictability results are shown in Fig. 6.10. We are able to distinguish both types of surrogates from the actual data set. This predictability and the structure observed in the return map is an indication that the data set has some nonlinear correlations. An analysis using the method of Tsonis and Ellsner to reveal evidence of chaos was inconclusive.

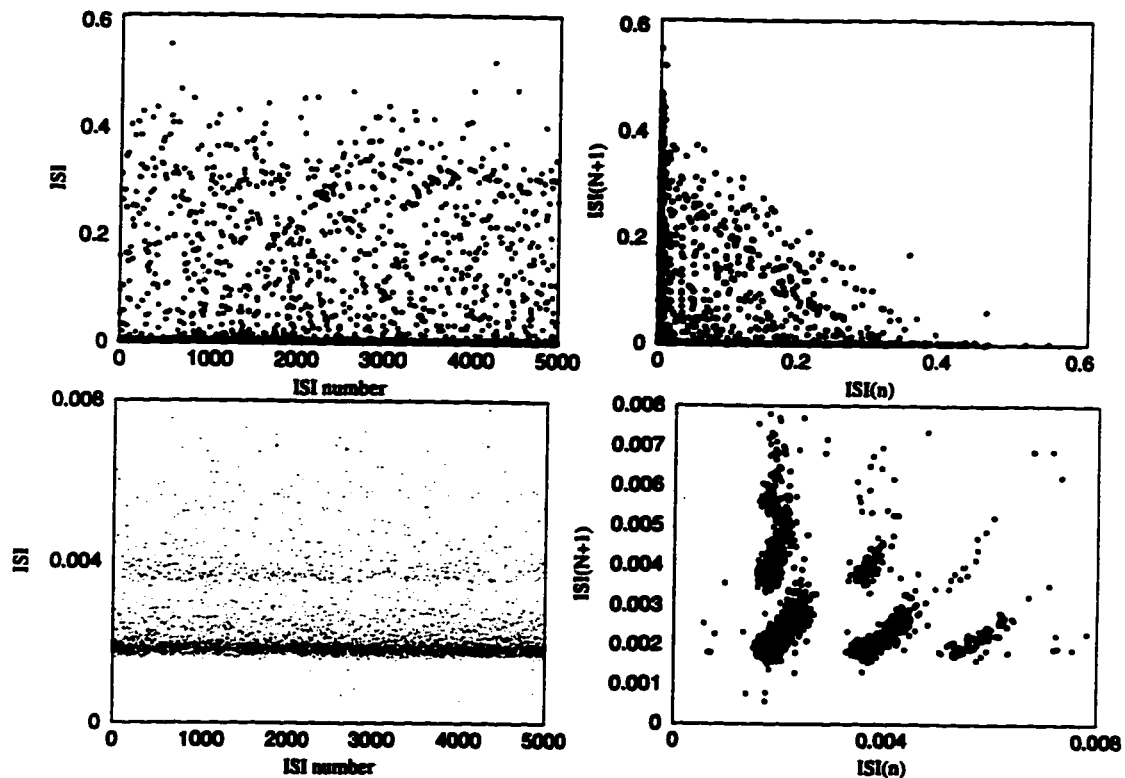


Figure 6.9: Raster plot (left) and return map (right) for data from the lateral segment of the weakly electric fish. The bottom panel is an expansion of a region of the top panel. Data provided by Leonard Maler, U. of Ottawa.




Figure 6.10: NPE for data from the lateral segment and surrogates.

6.4 Comparison of skipping models

Thermoreceptor activity has been modeled by both the Chay and Plant models of Sect. 2.3.3 [25]. It is possible to “tune” these models to reproduce the firing patterns of the thermoreceptors. Here, we investigate the predictability of ISI sequences from both models in order to help constrain the proper model for bursting and skipping in thermoreceptors. The ISI sequence for the Chay model is known to originate from deterministic skipping, whereas for Plant’s model the skipping is noise induced (the model does not skip without the added noise). In Fig. 6.11, we compare the predictability of both models. The left panel shows the NPE for all embedding dimensions; the right panel zooms in on the higher embedding dimensions. The Chay model generates ISI’s which are predictable at low embedding dimension and apparently random at high embedding dimension ($NPE \approx 1$). ISI’s from the Plant model generate ISI’s which give an NPE less than, but close to one for all embedding dimensions, but which increases slightly with embedding dimension. When these results are compared to those of Fig. 6.6 we see that the Chay ISI’s do not behave in a manner consistent with the ISI’s from thermoreceptors. This is a strong indication that the skipping mechanism of the Chay model is not the one which produces the skipping in thermoreceptors. Our preliminary analysis (not shown) indicates however that the combination of Chay’s deterministic skipping with noise may be a suitable candidate model. We should note that this analysis does not indicate that the Plant model involves the correct mechanism, it simply is a more appropriate model from the point of view of the predictability of the ISI’s.




Figure 6.11: NPE for ISI sequences from the Chay model (deterministic skipping) and the Plant model (noise induced skipping).

6.5 Conclusion

Our purpose in this chapter was to analyze ISI sequences with the method of surrogate data, to examine the possibility of extending the method of Tsonis and Ellsner to such sequences, and to apply these methods to data sets from neurophysiological experiments. We examined the ISI sequences from the three IF models of Chap. 5 and showed that the surrogate data method with the NPE as a statistic correctly identified the nonlinear correlation of these chaotic data sets.

The method of Tsonis and Ellsner was applied to the known chaotic ISI sequences from Chap. 5. By replacing the prediction time with the prediction horizon the method identified the chaotic nature of the sequences. Thus it appears that this method may be extendable from time series to ISI sequences.

We analyzed data from three experiments. The first data sets were from cat thermoreceptors. We were able to distinguish the surrogates from the data sets within the error bars we had set on the experiment which provides an indication that there are nonlinear correlations in the data set. The data sets from two neurons in the weakly electric fish Sect. 6.3.2 were distinguishable from the surrogate data sets. This implies that these ISI sequences contain nonlinear correlations. Analysis of the experimental data sets with the method of Tsonis and Ellsner was inconclusive. This suggests that the system which generated the ISI's are not chaotic or it may be that the nonlinear correlations in the data are not strong enough to be picked out by the method or that the noise level is too high for this method to work. Further

investigation must be carried out on known chaotic ISI sequences before we can determine the usefulness of this method.

Our analysis with surrogate data made it possible to identify what portion of the NPE was attributable to linear correlations in the data. However it did not allow us to reject physically meaningful mechanisms associated with null hypotheses. This is an indication of the need for biophysically motivated surrogates as advocated in [23, 26, 27]. These biophysical surrogates are ISI sequences which have similar properties as the data sets of interest and are generated by models with well known dynamics. For example in [27] a biophysical surrogate for the data of Sect. 6.3.2 is generated which has a similar ISI histogram as the data. This biophysical surrogate is taken from the FitzHugh-Nagumo model and is a manifestation of noise induced skipping. This surrogate was found to have many of the properties of the raw ISI data, which suggests that the associated biophysical picture ("noise induced skipping from the FHN system") is to a large extent correct. In Sect. 6.4 we have shown another example of the biophysical surrogate method whereby we generated surrogates for the thermoreceptor data of Sect. 6.3.1 with the Plant and Chay model of Sect. 2.3.4. These surrogates corresponded respectively: noise induced skipping and deterministic skipping. We were able to reject the deterministic skipping by comparing the NPE of both ISI sequences with the NPE of the thermoreceptor data. This type of analysis is still in the early stages of development but shows promise in the area of rejecting physically meaningful null hypotheses.

Chapter 7

Conclusion

The work presented in this thesis concerns the analysis of dynamical systems which generate events in the course of their continuous time evolution. In particular, we have concentrated on the case of neurons and their corresponding events: voltage spikes. Our study consists of two parts. The first addresses the possibility of relating interspike interval (ISI) sequences to the dynamical variables of the neural system (the neuron and its input). The second part concerns the assessment of linear and nonlinear correlations in ISI sequences.

Due to experimental limitations, it is often not possible to make transmembrane voltage measurements on neurons. This means that we are limited to making extracellular measurements, which only yield the neuron's spiking times. We have concentrated on the relationship between these spiking times and the dynamics of the input to a neuron. In [45] it was shown that, for the simplest neuron model, it is possible to reconstruct the dynamics of the input to the neuron by embedding the ISI's. We have extended this method to more realistic integrate and fire (IF) models and to the FitzHugh-Nagumo (FHN) model of membrane excitation. We have shown that the method works well for the IF models when the firing rate is high, i.e. when many samplings occur during a characteristic cycle of the input. We have defined the sampling rate as the mean number of spikes per unit time and have compared the quality of the reconstructions keeping this sampling rate constant. We have found that, as the models become more complex, the quality of the reconstruction decreases. For the FHN model, the ISI embeddings do not reconstruct the dynamics of the input to the system. This appears to be a consequence of the fact

that firing intervals no longer reflect the value of the input, and also, that firings are not seen over significant lower portions of the input.

Our results suggest that we should not, in general, expect ISI embeddings from real neurons to reconstruct the attractor of the input to the neuron. Nevertheless, attractor reconstruction may be possible for a small subset of real neurons known as rate coders. Since rate coders are known to be noisy neurons, future work on this problem must clearly investigate the influence of noise on the reconstructions. This understanding will be crucial in order to make precise statements about real neurons. It should be noted that, due to our results, experimental studies of reconstruction using known rate coding neurons have been undertaken at Boston University.

In [39] we have shown that, for a high sampling rate, the ISI's can be written as a map, or nonlinear transformation, of the input values at which firings occur. Since the input is a linear function of the x variable of the Rössler system, the ISI embeddings are in fact a transformation of an embedding of the x variable. This explains why the ISI embeddings reconstruct the Rössler attractor, in the high sampling rate limit. It would be interesting to extend this result to arbitrary firing rates and also to more complicated neuron models such as FHN. Once again, the influence of noise on such mappings will need to be assessed. Further, the FHN model studied here does not have a real threshold, but rather a pseudo-threshold [11]. This means that the heights of neural firing responses do not change abruptly, but rather are graded over a narrow range of input amplitudes. Other more realistic neuron models, such as the Hodgkin-Huxley model of Sect. 2.3.1, can have real or pseudo-thresholds. It would be interesting to study how our conclusions depend on the type of threshold.

In order to properly constrain theories of neural coding and neural information processing, it is necessary to characterize the correlations in ISI sequences. A knowledge of the correlations in the ISI's can also be used as a tool for the validation of neuron models, as was shown in Sect. 6.4. There, we compared two models for the thermoreceptors of Sect. 6.3.1. These models both reproduced the ISI histogram of the data and the skipping behaviour. One of the models (Chay) produced deterministic skipping, whereas the other (Plant) was a manifestation of noise induced skipping. These models served as biophysical surrogates and we showed that it was possible to eliminate the Chay model by comparing the predictability of both models

with the predictability of the thermoreceptor data.

In Chap. 6 we have presented a method, based on surrogate data methods from nonlinear time series analysis, which allows us to detect nonlinear correlations in ISI sequences. This is accomplished by investigating the predictability of the ISI sequences and comparing it to the predictability of “surrogate” data sets. These surrogates have the same linear properties as the original sets, but are otherwise stochastic. We analyzed ISI sequences from the integrate and fire models and showed that the method of surrogate data was able to detect the nonlinear correlations in these data sets. We then analyzed several experimental ISI sequences. Our analysis allowed us to determine what portion of the predictability of these data sets was attributable to the linear correlations in them. We found significant nonlinear correlations in the data from the weakly electric fish.

The surrogate data method was originally developed for time series. In that case, when a surrogate is rejected, we are able to reject a physically meaningful null hypothesis for the generation of the time series (for example Gaussian white noise or coloured noise). In the case of point process data, such as ISI sequences, we can not associate a physically meaningful null hypothesis to the surrogates. This immediately points us in the direction of future research. It would be interesting to devise new kinds of surrogates for point process data which would allow us to reject a null hypothesis along with a physically meaningful mechanism associated with this null. Also, we should continue the work on biophysical surrogates as described in Sect. 6.2 and 6.4.

We have also extended the method of Tsonis and Ellsner [58] to point process data by drawing a parallel between prediction time, in the time series case, to prediction horizon, in the point process case. The method worked well for the ISI sequences from the integrate and fire models. The method was also applied to ISI sequences from real neurons and proved to be inconclusive. This suggests that either the underlying continuous dynamics are not chaotic, or that if they are, the noise destroys the nonlinear forecastability. It may also be the case that firings are too far apart in comparison with the prediction horizon (related to the Lyapunov exponent) of the underlying system. We feel that the positive results with the IF models at high sampling rates are a consequence of the fact that only in this case do the ISI sequences contain sufficient information about the chaotic dynamics. Applying

this method to known chaotic point process data will allow us to further assess the applicability of this method to ISI sequences.

Throughout this study we have used a simple nonlinear prediction algorithm to calculate the predictability of the ISI sequences. The prediction algorithm currently considers all neighbours equally when making a prediction. An exponential weighting of the neighbours according to their distance from the point which is being predicted would improve the prediction. Kernel estimation, where all points in the space are considered rather than just the nearest neighbours may also yield a more powerful predictor. It would also be interesting to adapt other “statistics” inspired from nonlinear dynamics and information theory, such as dynamical entropies, to point process data.

Appendix A

Detailed description of the neuron models

We give a more detailed version of the differential equations for some of our model neurons. For the sake of clarity some of the information from Chap. 2 is repeated.

A.1 Plant's model

Taken from [42]

$$C_m \frac{dV}{dt} = \bar{g}_{Na} m_\infty^3(V) h(V_{Na} - V) + \bar{g}_{Ca} x(V_{Ca} - V) \\ + [\bar{g}_k n^4 + \frac{\bar{g}_{KCa} Ca}{Ca + 0.5}] (V_K - V) + \bar{g}_L (V_L - V),$$

$$\dot{h} = \lambda [h_\infty(V) - h] / \tau_h(V),$$

$$\dot{n} = \lambda [n_\infty(V) - n] / \tau_n(V),$$

$$\dot{x} = [x_\infty(V) - x] / \tau_x,$$

$$\dot{Ca} = \rho [K_c x (V_{Ca} - V) - Ca].$$

The parameter values and voltage dependencies for \dot{V} are given by:

$$C_m = 1 \mu F/cm^2, \quad V_{Na} = 30 mV, \quad V_{Ca} = 140 mV,$$

$$\begin{aligned}
V_K &= -75 \text{ mV}, & V_L &= -40 \text{ mV}, & \bar{g}_{Na} &= 7.84 \text{ mmho/cm}^2, \\
\bar{g}_{Ca} &= 0.01 \text{ mmho/cm}^2, & \bar{g}_K &= 0.363 \text{ mmho/cm}^2, & \bar{g}_{KCa} &= 0.03 \text{ mmho/cm}^2, \\
\bar{g}_L &= 0.003 \text{ mmho/cm}^2, & \lambda &= \frac{1}{12.5},
\end{aligned}$$

$$m_\infty(V) = \frac{\alpha_m(\tilde{V})}{\alpha_m(\tilde{V}) + \beta_m(\tilde{V})},$$

$$\tilde{V} = C_1 V + C_2, \quad C_1 = 127/105, \quad C_2 = 8265/105,$$

$$\alpha_m(V) = \frac{50 - V}{10} \left[\exp\left(\frac{50 - V}{10}\right) - 1 \right]^{-1},$$

$$\beta_m = 4 \exp\left(\frac{25 - V}{18}\right).$$

For \dot{h} :

$$h_\infty(V) = \frac{\alpha_h(\tilde{V})}{\alpha_h(\tilde{V}) + \beta_h(\tilde{V})},$$

$$\tau_h(V) = [\alpha_h(\tilde{V}) + \beta_h(\tilde{V})]^{-1},$$

with

$$\alpha_h(V) = 0.07 \exp\left(\frac{25 - v}{20}\right),$$

$$\beta_h(V) = \left[1 + \exp\left(\frac{55 - V}{10}\right) \right]^{-1}.$$

For \dot{n} :

$$n_\infty(V) = \frac{\alpha_n(\tilde{V})}{\alpha_n(\tilde{V}) + \beta_n(\tilde{V})},$$

$$\tau_n(V) = [\alpha_n(\tilde{V}) + \beta_n(\tilde{V})]^{-1},$$

where

$$\alpha_n(V) = \frac{50 - V}{100} \left[\exp\left(\frac{55 - V}{10}\right) - 1 \right]^{-1},$$

$$\beta_n = 0.125 \exp\left(\frac{45 - V}{80}\right).$$

For \dot{x} :

$$x_\infty(V) = [\exp(A(B - V)) + 1]^{-1},$$

$$\tau_x = 235 \text{ msec},$$

where $A = 0.15$ and $B = -50$.

For $\dot{C}a$:

$$\rho = 0.0003 \text{ msec}^{-1} \quad \text{and} \quad K_c = 0.0085 \text{ mV}^{-1}.$$

A.2 Chay's model

The basic differential equation for the membrane potential is:

$$C_m \frac{dV}{dt} = -\bar{g}_{fast} m_\infty^3 h (V - V_{fast}) - \bar{g}_{slow} d f (V - V_{slow}) \\ - \bar{g}_K n (V - V_K) - \bar{g}_L (V - V_L).$$

Each of the variables h , d , f and n obeys a differential equation of the form:

$$\frac{dx}{dt} = \frac{x_\infty - x}{\tau_x},$$

with

$$m_\infty = \left[1 + \exp \left(\frac{V_m - V}{S_m} \right) \right]^{-1},$$

$$h_\infty = \left[1 + \exp \left(\frac{V_h - V}{S_h} \right) \right]^{-1}, \quad \tau_h = \bar{\tau}_h \left[\exp \left(\frac{V_h - V}{2S_h} \right) + \exp \left(\frac{V - V_h}{2S_h} \right) \right]^{-1},$$

$$d_\infty = \left[1 + \exp \left(\frac{V_d - V}{S_d} \right) \right]^{-1}, \quad \tau_d = \bar{\tau}_d \left[\exp \left(\frac{V_d - V}{2S_d} \right) + \exp \left(\frac{V - V_d}{2S_d} \right) \right]^{-1},$$

$$f_\infty = \left[1 + \exp \left(\frac{V_f - V}{S_f} \right) \right]^{-1}, \quad \tau_f = \bar{\tau}_f \left[\exp \left(\frac{V_f - V}{2S_f} \right) + \exp \left(\frac{V - V_f}{2S_f} \right) \right]^{-1},$$

$$n_\infty = \left[1 + \exp \left(\frac{V_n - V}{S_n} \right) \right]^{-1}, \quad \tau_n = \bar{\tau}_n \left[1 + \exp \left(\frac{V - V_n}{2S_n} \right) \right]^{-1}.$$

Values for the constants used to generate the ISI's of Sect. 6.4 are given below.

For \dot{V} :

$$C_m = 1 \mu\text{F}/\text{cm}^2, \quad V_{fast} = 80 \text{ mV}, \quad V_{slow} = 140 \text{ mV}, \\ V_K = 75 \text{ mV}, \quad V_L = 70 \text{ mV}, \quad \bar{g}_{fast} = 800 \text{ mmho}/\text{cm}^2,$$

$$\bar{g}_{slow} = 15.5 \text{ mmho/cm}^2, \quad \bar{g}_K = 500 \text{ mmho/cm}^2, \quad \bar{g}_L = 20 \text{ mmho/cm}^2,$$
$$V_m = 35.0 \text{ mV}, \quad S_m = 5 \text{ mV}.$$

For \dot{h} , \dot{d} , \dot{f} , and \dot{n} :

$V_h = 50 \text{ mV},$	$S_h = 9 \text{ mV},$	$\bar{\tau}_h = 0.2 \text{ sec},$
$V_d = 50 \text{ mV},$	$S_d = 7 \text{ mV},$	$\bar{\tau}_d = 1.0 \text{ sec},$
$V_f = 55 \text{ mV},$	$S_f = 10 \text{ mV},$	$\bar{\tau}_f = 40.0 \text{ sec},$
$V_n = 15 \text{ mV},$	$S_n = 15 \text{ mV},$	$\bar{\tau}_n = 0.3 \text{ sec}.$

Appendix B

Algorithms for generating surrogate data

This appendix addresses the issue of transforming a time series $x[t]$, $t = 1, \dots, N$ into various surrogates $y[t]$. A good discussion on sorting, indexing and ranking (Chap. 8), and random number generation (Chap. 7) can be found in [37].

B.1 Shuffled surrogates

This algorithm consists of simply shuffling the data like a deck of cards. The surrogates have the same mean and standard deviation as the original but they are temporally uncorrelated. We generate some random numbers, sort them, and assign the order of the sorted random numbers to the data.

1. Generate an array of N uniformly distributed random numbers $r[t]$.
2. Create an index table for the sorted array $index[t]$. This creates an array of integers such that if $r[5]$ and $r[33]$ are the two smallest integers then $index[1] = 5$, and $index[2] = 33$.
3. The surrogates are then generated by $y[t] = x[index[t]]$.

B.2 Phase randomized surrogates

These surrogate data are constructed to have the same power spectrum as the original data.

1. Compute the discrete Fourier transform of the original data: $Z[t] = DFT(x[t])$ ¹. Note that $Z[t]$ is a complex array.
2. Randomize the phases $Z'[t] = Z[t]e^{i\phi(t)}$ where $\phi(t)$ are random numbers uniformly distributed on $[0, 2\pi)$.
3. In order for the inverse Fourier transform to be real we must symmetrize the phases of Z' . This is accomplished by creating another array Z'' such that:

$$\begin{aligned} \text{Re}(Z''[t]) &= \frac{1}{2}\text{Re}(Z'[t] + Z'[N+1-t]) \\ \text{Im}(Z''[t]) &= \frac{1}{2}\text{Im}(Z'[t] - Z'[N+1-t]), \end{aligned}$$

where N is the number of points in the time series.

4. The surrogate data are given by computing the inverse discrete Fourier transform of this new series, $y[t] = DFT^{-1}(Z''[t])$.

B.3 Amplitude adjusted surrogates

The idea here is to first rescale the data so that it has a Gaussian distribution. We then take phase randomized surrogates of this Gaussian data. The new surrogates are then generated by rescaling these so that they have the same distribution as the original data.

1. Generate an array of N random numbers from a Gaussian distribution $g[t]$.
2. Rearrange the $g[t]$ into an array $g'[t]$ so that the ranks of $g'[t]$ and $x[t]$ agree. This is accomplished by the following steps.

¹If N is a power of 2 we can use the Fast Fourier Transform.

- (i) Sort the Gaussian array in ascending order $g_s[t] = Sg[t]$.
 - (ii) Create a rank table for the original data array. This creates an array of integers such that if $x[1]$ is the third largest element of the x array and $x[2]$ is the twentieth then $rank[1] = 3$ and $rank[2] = 20$ and so on.
 - (iii) $g'[t] = g_s[rank[t]]$.
3. Note that if $x[n]$ is the n th smallest of all the x 's then $g'[n]$ is the n th smallest of all the g 's. Hence we say that the new series $g'[t]$ "follows" the original $x[t]$.
 4. Generate phase randomized surrogates of $g'[t]$, call it $g''[t]$.
 5. Rank the original time series $x[t]$ so that it "follows" $g''[t]$, using steps (i),(ii), and (iii) above, this produces an amplitude adjusted surrogate data set $y[t]$.

Bibliography

- [1] H. D. I. Abarbanel, R. Brown, J. L. Sidorowich, and L. S. Tsimring. The analysis of observed chaotic data in physical systems. *Rev. Mod. Phys.*, 65:1331–1392, 1993.
- [2] L. F. Abbott and T. B. Kepler. Model neurons: from Hodgkin Huxley to Hopfield. Proceedings of the eleventh Sitges Conference on Neural Networks, June 1990.
- [3] M. Casdagli. Nonlinear prediction of chaotic time series. *Physica D*, 35:335–356, 1989.
- [4] M. Casdagli. Chaos and deterministic versus stochastic nonlinear modeling. *J. Roy. Stat. Soc. B*, 54:303–328, 1992.
- [5] T. R. Chay. Bursting excitable cell models by a slow Ca^{2+} current. *J. Theor. Biol.*, 142:305–315, 1990.
- [6] J.P. Crutchfield, J.D. Farmer, N.H. Packard, and R.S. Shaw. Chaos. *Scientific American*, 254:46–57, 1986.
- [7] J.P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57:617–656, 1985.
- [8] B. Efron. Computers and the theory of statistics: thinking the unthinkable. *SIAM Review*, 21:460–480, 1979.
- [9] J. D. Farmer and J. J. Sidorowich. Predicting chaotic time series. *Phys. Rev. Lett.*, 59:845–848, 1987.

- [10] R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophys. J.*, 1:445–457, 1961.
- [11] R. FitzHugh. Mathematical models of excitation and propagation in nerve. In H. P. Schwann, editor, *Biological Engineering*, chapter 1, pages 1–85. McGraw-Hill, New York, 1962.
- [12] A. C. Giese. *Cell Physiology*. W. B. Saunders Company, Philadelphia, PA, fifth edition, 1979.
- [13] L. Glass and M. C. Mackey. A simple model for phase locking of biological oscillators. *J. Math. Biology*, 7:339–352, 1979.
- [14] L. Glass and M. C. Mackey. *From Clocks to Chaos*. Univ. Press, Princeton NJ, 1988.
- [15] P. Grassberger, T. Schreiber, and C. Schaffrath. Nonlinear time sequence analysis. *Int. J. of Bifurcation and Chaos*, 1(3):521–547, 1991.
- [16] C. Grebogi, E. Ott, S. Pelikan, and J. A. Yorke. Strange attractors that are not chaotic. *Physica D*, 13:261–268, 1984.
- [17] Bai-Lin Hao. *Chaos II*. World Scientific, Singapore, 1990.
- [18] M. Hénon. A two-dimensional mapping with a strange attractor. *Commun. Math. Phys.*, 50:69–77, 1976.
- [19] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.*, 116:500–544, 1952.
- [20] F. C. Hoppensteadt. *An Introduction to the Mathematics of Neurons*, volume 6 of *Cambridge Studies in Mathematical Biology*. Cambridge University Press, Cambridge, 1986.
- [21] R.V. Jensen. Classical chaos. *American Scientist*, 75:168–181, 1987.

- [22] J. P. Keener, F. C. Hoppensteadt, and J. Rinzel. Integrate and fire models of nerve membrane response to oscillatory input. *SIAM J. Appl. Math.*, 41:503–517, 1981.
- [23] A. Longtin. Nonlinear forecasting of spike trains from sensory neurons. *Int. J. of Bifurcation and Chaos*, 3(3):651–660, 1993.
- [24] A. Longtin. Stochastic resonance in neuron models. *Journal of Statistical Physics*, 70(1/2):309–327, 1993.
- [25] A. Longtin and K. Hinzer. Encoding with bursting, subthreshold oscillations, and noise in mammalian cold receptors. *Neural Computation*, 8:215–255, 1996.
- [26] A. Longtin and D. M. Racicot. Assessment of linear and nonlinear correlations between neural firing events. In D. Kaplan and C. D. Cutler, editors, *Nonlinear Dynamics and Time Series: Building a Bridge Between the Natural and Statistical Sciences.*, pages 221–238. Fields Institute For Research in the Mathematical Sciences Communications Series, Vol. 11, American Mathematical Society, Providence, RI, 1996.
- [27] A. Longtin and D. M. Racicot. Spike train patterning and forecastability. *Biosystems*, 40:111–118, 1997.
- [28] E. N. Lorenz. Deterministic nonperiodic flow. *J. Atmos. Sci.*, 20:130, 1963.
- [29] R. J. MacGregor and E. R. Lewis. *Neural Modeling, Electrical Signal Processing in the Nervous System*. Plenum Press, New York, 1977.
- [30] F. Moss. Chaos under control. *Nature (News and Views: Neurobiology)*, 370:596–597, 1994.
- [31] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. In *Proc. IRE*, volume 50, pages 2061–2070, 1962.
- [32] J.G. Nicholis, A.R. Martin, and B.G. Wallace. *From Neuron to Brain: a cellular and molecular approach to the function of the nervous system*. Sinauer Associates Inc., Sunderland, MA, third edition, 1992.

- [33] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Phys. Rev. Lett.*, 45:712–716, 1980.
- [34] T. S. Parker and L. O. Chua. *Practical numerical algorithms for chaotic systems*. Springer, Berlin, 1989.
- [35] R. E. Plant. Bifurcation and resonance in a model for bursting nerve cells. *J. Math. Biology*, 11:15–32, 1981.
- [36] H. Preißl, A. Aertsen, and G. Palm. Are fractal dimensions a good measure for neuronal activity. In R. Eckmiller, R. Hartmann, and G. Hausske, editors, *Parallel Processing in Neural Systems and Computers*, pages 83–86. Elsevier Science Publisher, Amsterdam, 1990.
- [37] W. H. Press, S. T. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: the art of scientific computing*. Cambridge University Press, Cambridge, (England), second edition, 1992.
- [38] D. M. Racicot and A. Longtin. Reconstructing dynamics from neural spike trains. In *Proc. IEEE 17th Ann. Conf. Eng. Med. Biol. Soc.*, 1995.
- [39] D. M. Racicot and A. Longtin. Interspike interval attractors from chaotically driven neuron models. *Physica D*, 1996. In Press.
- [40] P. E. Rapp, I. D. Zimmerman, A. M. Albano, G. C. deGuzman, and N. N. Greenbaun. Dynamics of spontaneous neural activity in the simian motor cortex: The dimension of chaotic neurons. *Phys. Lett. A*, 110:335–338, 1985.
- [41] P.E. Rapp, A.M. Albano, I.D. Zimmerman, and M.A. Jiménez-Montano. Phase-randomized surrogates can produce spurious identifications of non-random structure. *Phys. Lett. A*, 192:27–33, 1994.
- [42] J. Rinzel and Y.S. Lee. Dissection of a model for neuronal parabolic bursting. *J. Math. Biology*, 25:653–675, 1987.
- [43] M. T. Rosenstein. *MTR Chaos, Time Series Analysis Tools*. Dynamical Research, 15 Pecunit Street, Canton MA, USA, 02021-1219, 1993.

- [44] O. E. Rössler. An equation for continuous chaos. *Phys. Lett. A*, 57:397–398, 1976.
- [45] T. Sauer. Reconstruction of dynamical systems from interspike intervals. *Phys. Rev. Lett.*, 72:3811–3814, 1994.
- [46] T. Sauer. Interspike interval embeddings of chaotic signals. *Chaos*, 5(1):127–132, 1995.
- [47] B. McA. Sayers. Inferring significance from biological signals. In M. Clynes and J. Milsum, editors, *Biomedical Engineering Systems*, chapter 4, pages 84–164. McGraw-Hill, 1970.
- [48] S.J. Schiff, K. Jerger, T. Chang, T. Sauer, and P.G. Aitken. Stochastic versus deterministic variability in simple neuronal circuits II: Hippocampal slice. *Biophys. J.*, 67:684–691, 1994.
- [49] S.J. Schiff, T. Sauer, and T. Chang. Discriminating deterministic versus stochastic dynamics in neuronal activity. *Integrative Physiology and Behavioural Science*, 29:246–261, 1994.
- [50] R. Shaw. *The dripping faucet as a model chaotic system*. Aerial Press, Santa Cruz, CA, 1984.
- [51] J. C. Sprott and G. Rowlands. *Chaos Data Analyzer*. Physics Academic Software, American Institute of Physics, New York, 1992.
- [52] G. Sugihara and R. May. Nonlinear forecasting as a way of distinguishing chaos from measurement errors in time series. *Nature*, 344:734–741, 1990.
- [53] N. Takahashi, Y. Hanyu, T. Musha, R. Kubo, and G. Matsumoto. Global bifurcation structure in periodically stimulated giant axons of squid. *Physica D*, 43:318–334, 1990.
- [54] F. Takens. Detecting strange attractors in turbulence. In Rand D. A. and Young L.-S., editors, *Dynamical Systems and Turbulence (Warwick 1980) (Lecture Notes in Mathematics)*, volume 898, pages 366–381. Springer, Berlin, 1980.

- [55] J. Theiler. Spurious dimensions from correlation algorithms applied to limited time series data. *Phys. Rev. A*, 34:2427–2432, 1986.
- [56] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer. Testing for nonlinearity in time series: the method of surrogate data. *Physica D*, 58:77–94, 1992.
- [57] J. Theiler, B. Galdrikian, A. Longtin, S. Eubank, and J. D. Farmer. Using surrogate data to detect nonlinearity in time series. In M. Casdagli and S. Eubank, editors, *Nonlinear modeling and forecasting. Proceedings of the workshop held in Santa Fe, New Mexico, September 1990*, pages 163–188, Reading Mass., 1992. Addison Wesley.
- [58] A. A. Tsonis and J. B. Elsner. Nonlinear prediction as a way of distinguishing chaos from random fractal sequences. *Nature*, 359:217–220, 1992.
- [59] H. Tuckwell. *Introduction to Theoretical Neurobiology*. Cambridge University Press, Cambridge, 1988.
- [60] G. U. Yule. On a method of investigating periodicities in disturbed series with special reference to Wolfer’s sunspot numbers. *Philos. Trans. Roy. Soc. London A*, 226:267–298, 1927.
- [61] I.D. Zimmerman, P.E. Rapp, and A.I. Mees. The geometrical characterization of neural activity displays a sensitivity to convulsants. *Int. J. of Bifurcation and Chaos*, 1(1):253–259, 1991.