

Dynamical systems techniques in the analysis of neural systems

Kyle Wedgwood, MSc.

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

December 2013

Dedicated to my late grandparents, who have been a constant source of inspiration and support throughout my academic journey.

Abstract

As we strive to understand the mechanisms underlying neural computation, mathematical models are increasingly being used as a counterpart to biological experimentation. Alongside building such models, there is a need for mathematical techniques to be developed to examine the often complex behaviour that can arise from even the simplest models. There are now a plethora of mathematical models to describe activity at the single neuron level, ranging from one-dimensional, phenomenological ones, to complex biophysical models with large numbers of state variables. Network models present even more of a challenge, as rich patterns of behaviour can arise due to the coupling alone.

We first analyse a planar integrate-and-fire model in a piecewise-linear regime. We advocate using piecewise-linear models as caricatures of non-linear models, owing to the fact that explicit solutions can be found in the former. Through the use of explicit solutions that are available to us, we categorise the model in terms of its bifurcation structure, noting that the non-smooth dynamics involving the reset mechanism give rise to mathematically interesting behaviour. We highlight the pitfalls in using techniques for smooth dynamical systems in the study of non-smooth models, and show how these can be overcome using non-smooth analysis.

Following this, we shift our focus onto the use of phase reduction techniques in the analysis of neural oscillators. We begin by presenting concrete examples showcasing where these techniques fail to capture dynamics of the full system for both deterministic and stochastic forcing. To overcome these failures, we derive new coordinate systems which include some notion of distance from the underlying limit cycle. With these coordinates, we are able to capture the effect of phase space structures away from the limit cycle, and we go on to show how they can be used to explain complex behaviour in typical oscillatory neuron models.

Acknowledgements

First and foremost, I would like to thank my supervisors: Prof. Stephen Coombes and Dr. Rüdiger Thul. Throughout both my Masters and PhD studies, they have both been a constant source of inspiration, always willing to lend an ear and offer advice whenever called upon. As well as supervising my studies, they have both driven me to make important strides in establishing myself as a researcher, and have consistently provided opportunities for me to this end. Away from the scientific side of my PhD, they have also supported me emotionally during the difficult times during these past years. Truthfully, I can't imagine two supervisors who care more for their students, both as researchers, but also as people.

Secondly, I thank my parents and the rest of my family, who have stood behind me every step of the way down my journey, supporting and advising where necessary, and trusting me to find my feet in this world.

As well as my supervisors, I would like express my gratitude to Prof. Michael Tretyakov for discussions on numerical routines for SDE integration, Dr. Edward Hall for his input on finite element methods for solving the Fokker-Planck equation, Prof. Jonathan Rubin for encouraging my further study of the PWL-IF model, Dr. Kevin Lin for his help with the numerical simulations of the Morris-Lecar model and to all of the tutors at the Ottawa Summer School in Computational Neuroscience for giving me a grounded start in the field.

For their helpful input on how to improve this thesis and for an enjoyable viva, I would like to extend my gratitude to my examiners Prof. Peter Ashwin and Dr. Paul Matthews.

For their constant bubbly personalities, and for helping me out with a whole host of things, I must say a big thank you to the School Manager, Andrea Blackburn, and to all of the school secretaries, and in particular, to Hilary Lonsdale and Cathie Shipley, for putting up with my demands whenever I go abroad and to Helen Cunliffe for advice on preparing my thesis. For his assistance with all things related to my slowly dying Mac, I must also thank Dave Parkin.

For providing a outlet to complement my studies, I would like to thank all of the people I have worked with as part of the Students' Union and through dodgeball. I have had the pleasure of working with some incredible people throughout my time here, and it is incredibly refreshing to see the passion that students have and what they can achieve through it.

Last, but by no means least, I'd like to thank all of my friends here in Nottingham, for keeping my life fun and keeping me grounded for the past four years. I've had an incredible time here so far, and I hope that the next two will be just as good.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Thesis outline	4
2	Background	8
2.1	The Hodgkin–Huxley model	8
2.2	Dimension reduction of the Hodgkin–Huxley model	13
2.3	The integrate-and-fire neuron	21
2.4	Isochronal coordinates	27
2.5	Reduction of oscillatory models to a phase-only description	34
2.6	Noise in neural systems	36
2.7	Sources of noise	37
2.7.1	Electrical noise	37
2.7.2	Synaptic noise	38
2.8	Management of noise	38
2.9	Benefits of noise	39
2.10	Modelling noise	40
2.11	Random processes and sample paths	41
2.12	The Wiener process	41
2.13	Langevin equations	42
2.14	Convergence	44
2.15	Stochastic integrals	46
2.16	Comparison of the Itô and Stratonovich integrals	47
2.17	Itô’s formula	48
2.18	Coloured noise	49
2.19	Simulating stochastic processes	51
2.20	The Fokker–Planck equation	57
2.20.1	The forward Fokker–Planck equation	57
2.20.2	The backward Fokker–Planck equation	58

2.21	Transforming probability distributions	59
2.22	Phase reduction of stochastically forced oscillators	60
3	Non-smooth spiking models	63
3.1	A review of integrate-and-fire models	67
3.2	Nonlinear IF models	68
3.3	Planar IF models	70
3.4	The absolute IF model	72
3.5	A piecewise-linear IF model	73
3.6	Bifurcation structure	76
3.7	Solutions of the PWL-IF model	89
3.8	Numerical integration of planar IF models	91
3.9	Subthreshold orbits	93
3.10	Phase response curves	97
3.11	Weakly coupled networks	101
3.12	Firing map	108
3.13	Maximal Lyapunov Exponents	114
3.14	Discussion	123
4	Limitations of phase reductions	127
4.1	Stochastic forcing.	129
4.2	Breakdown of the phase reduction	130
4.3	Shear-induced chaos and related phenomena	133
4.3.1	Shear-induced chaos and isochrons	135
4.3.2	Summary and comparisons	138
4.4	Chaotic response to periodic kicking	140
4.4.1	Geometry of the ‘homoclinic regime’	140
4.4.2	Shear-induced chaos	144
4.5	Firing rates and interspike intervals	151
4.6	Comparison of firing rates.	153
4.7	Comparison of ISI distributions	154

4.8	Altered spiking patterns in bistable systems	156
4.9	A network of two cells	160
4.10	Discussion	165
5	A phase-amplitude description of neural oscillators	169
5.1	Phase-amplitude coordinates	171
5.1.1	Unforced systems	172
5.1.2	Computing A , f_1 and f_2 in terms of the original vector field .	175
5.1.3	Forced systems	181
5.2	Some neural examples	182
5.2.1	The ML model	182
5.2.2	The FitzHugh–Nagumo Model	186
5.2.3	A SNIC bifurcation	187
5.2.4	A 4D conductance-based model	191
5.3	Pulsatile forcing of phase-amplitude oscillators	194
5.4	Revisiting the two cell ML network	204
5.5	Discussion	204
6	Phase-amplitude descriptions of stochastically forced neural os-	
	cillators	208
6.1	Stochastically forced oscillators	210
6.2	Weakly forced systems in isochronal coordinates	210
6.2.1	Differences between the Itô and Stratonovich interpretations of the phase reduction	213
6.2.2	Phase distribution and frequency of the SL model	215
6.3	A new coordinate system for planar systems	217
6.3.1	Transformed dynamics	221
6.3.2	The SL oscillator	222
6.3.3	Steady state distribution for weak forcing	223
6.3.4	Steady state distributions for non-weak forcing	225
6.4	The subcritical Hopf	228

6.4.1	Solution for the 1D case	233
6.4.2	The two-dimensional case	242
6.5	Discussion	248
7	Discussion	254
8	Appendices	266
A	Numerical integration of the SDEs (4.17) and (4.18) for Chapt. 4	266
B	Numerical integration of the stochastic phase-amplitude equations as used in Chapt. 6	268
C	Steady state distributions for non-weak noise from Chapt. 6	272
C.1	Fokker–Planck equation for systems in the phase-amplitude coordinates from Chapt. 6	273
D	Solution of the Fokker–Planck equation by finite element methods	275
D.1	Small noise solutions using the finite element method	276
E	Solution of the Fokker–Planck equation using the finite difference method	278
E.1	Small noise problems using the finite difference method	280
F	FEniCS expressions for the Fokker-Planck equation for the phase-amplitude coordinates in Chapt. 6	282

1 Introduction

1.1 Motivation

The main focus of this thesis is the use of dynamical systems techniques in the analysis of neural behaviour. In particular, we consider the non-smooth dynamics that arise due to the reset mechanism in the widely used integrate-and-fire model, and the limitations and extensions of the phase reduction techniques used to study networks of neural oscillators. In the mathematical modelling of neural networks, there is a desire for models to be both biophysically realistic and mathematically tractable, so that they are amenable to analysis. This thesis addresses some of the issues regarding the use of simple phenomenological models, as well as the techniques used to simplify models by reducing the dimensionality of the underlying system.

Due to their low computational cost, integrate-and-fire type models have been widely used in computational studies of brain behaviour. With simple modifications, the simplest model can be shown to display the same characteristics as certain brain regions, for example in the response of ventral cochlear to periodic stimuli [1]. More recently, planar integrate-and-fire models have been fitted to data from real cortical neurons with great efficacy, highlighting their use as a computational tool [2]. Here, we develop a piecewise linear planar integrate-and-fire model. The piecewise linear nature of the model allows us to construct analytical solutions for the model and thus allows for a great deal of mathematical analysis. The non-smooth nature of the reset mechanism in integrate-and-fire models gives rise to a much richer bifurcation structure than for smooth models, by way of non-smooth bifurcations, the kind of which are reviewed in [3]. Using our simple model, we investigate these non-smooth bifurcations, the effect they have at the single cell and network level and their relevance to neuroscience.

Neurons in the cortex may receive synaptic input from up to tens of thousands of other neurons. The sheer number of connections and recurrence in such networks thus makes network analysis complicated. Biophysical models of neural

behaviour tend to be based on the Hodgkin–Huxley formalism, taking into account the dynamics of the ion channels in the neuron cell membrane and the effect that these have on the induced currents and potential difference across this membrane. Such models tend to have high dimension, particularly when they include cell specific ionic currents. Coupling between cells tends to be either through synapses or through gap junctions. Synaptic coupling is often modelled through further sets of equations. Constructing a large model with all of these factors quickly becomes prohibitively computationally expensive with the computing resources available to us at this time, though projects, such as the Blue Brain Project [4], are attempting to construct a virtual human brain using supercomputers. Owing to the vast number of model variables and parameters, such models are also difficult to analyse mathematically.

Phase reduction techniques reduce the complexity of these networks by assuming that the strength of coupling between cells is weak. The theory of weakly connected neural networks is now well established, and a nice review of techniques and results in this paradigm may be found in [5]. Under this assumption, the dynamics of a cell in the network may be reduced to dynamics on a circle, where the state variable is now the phase along an oscillation, and where the inputs from other cells manifest themselves through an object known as the phase response curve (PRC). Inputs from other cells thus affect the timing of the next oscillation through the PRC. Since all the information about the state of the cell is now captured through a single variable, the dimensionality of the system is thus reduced to the number of oscillators in the network. Under certain symmetry conditions, this can be reduced further, but this will be discussed elsewhere. Being an interdisciplinary subject, computational neuroscience requires a constant dialogue between experimentalists and theoreticians. Phase reductions are a nice example of how this can work in practice. The PRC can be constructed for a given cell experimentally by inducing a current in a voltage clamped cell. This allows a direct comparison between mathematical models and experimental results that can be exploited to bridge the gap between the disciplines. The connection between the

mathematics and experiments of PRCs is the subject of the new book [6].

The phase reduced systems are computationally cheap, since all of the complexity of the system has been reduced to a single object - the PRC, and are tractable as knowledge of the PRC and the coupling between cells will give conditions of the stability of network states. However, the notion of what exactly is ‘weak’ is a topic for debate. The weak assumption in phase reduction is a limiting case of low strength inputs, but it is one made typically for mathematical convenience rather than with any biological consideration. Even if individual inputs to a cell are weak, the summative effect across the network, and across time may result in non-weak effects. In addition to this, in terms of dynamical systems, nearby invariant structures can have a pronounced effect on the dynamics of a system. By considering a classical neural example, we show situations under which the phase reduced system is not a good approximation to the full system, and show how, by appending an extra variable to the system, we can circumvent these problems.

Noise is ubiquitous in biological systems. Exactly what we mean by biological noise is another subject for debate. Indeed, there are many potential sources of noise in neural systems, and it is possible that what appears as noise at a macroscopic level is simply a failure to capture all of the dynamics at the microscopic level. To confound matters, noise also factors into experimental observation, resulting in a fairly ambiguous sense of what exactly noise refers to in these systems. Using techniques from signal analysis, meaningful signals can be extracted from noisy ones. Understanding where this noise arises may will lead to the development of better techniques, both for collecting and analysing data.

Whilst noise has historically been seen as a problem for biological systems, there is a growing body of literature highlighting the potentially beneficial effects that noise can impart on a system. Thus, understanding its role is an area of growing interest in the field of mathematical neuroscience. In particular, in recent years there has been a spate of work on phase reductions for noisy oscillators [7, 8, 9]. Mathematically, noise can have a number of interesting effects on dynamical systems. Considering the noise as a perturbation to an otherwise deterministic

system, noise can shift bifurcations, change the stability of invariant structures, or induce switches between basins of attraction of different attracting sets, resulting in mixed-mode oscillations. Such phenomena are reviewed in [10] for fast-slow systems, which are common in neuroscience modelling, whilst specific neural examples are considered in [11]. We consider planar models of neural activity with noise and show under what conditions such models can be reduced to a phase only model. We go on to analyse the effects that noise can have on a system, and go on to derive a framework for when models cannot be reduced to a phase model.

1.2 Thesis outline

Here, we outline the content of the main chapters of this thesis.

Chapter 2

We begin with a general introduction on the development of mathematical models of neural activity and techniques used in their simplification. We review some of the main techniques from dynamical systems theory used in their analysis and review what kind of behaviours such models support. This chapter addresses some of the history behind the development of these models, and some of the criticisms of them.

We then go on to review some of the techniques from non-smooth analysis and show what pronounced effects non-smooth dynamics can have. Following on from this, we discuss phase reductions and their practical use in analysing network dynamics. Finally, we consider the sources and effects of noise on neural systems and the ongoing work in the field in this area.

Chapter 3

Considering a motivating example of bursting cells, this chapter focuses on the development and analysis of a piecewise linear integrate-and-fire model. We show how the simple, planar model can capture the same kinds of behaviour as more complex, high dimensional models. In particular, we look at the phenomenon of bursting, where a period of higher frequency oscillations is followed by a long period of quiescence. We go on to find the bifurcation structure, tracking both

smooth and non-smooth bifurcations. We derive the phase response curve for the model in a tonically spiking regime and in a bursting mode. Using techniques for non-smooth systems, we evaluate the maximal Lyapunov exponent, categorising where chaotic solutions exist and show that the model supports chaos between spike adding bifurcations and during the transition to oscillator death. Following in the same light, we analyse the behaviour of two cell network.

Chapter 4

Beginning with a review of some of the applications of phase models, including their use in experimental work, we consider the appropriateness of the phase reduction. We numerically examine three situations in which the phase reduced model is shown to miss important behaviour of the full model. In particular, we consider how nearby invariant structures can lead to shear induced chaos, how noisy forcing can result in phase models overestimating the neuronal firing rate and how switching between basins of attraction of attractors give rise to altered spiking patterns. These situations serve as a warning to computational neuroscience to take care when applying phase reduction techniques without first checking that the required assumptions are valid.

Chapter 5

To overcome the problems discussed in the previous chapter, we use a moving orthonormal coordinate system, building on ideas in [12]. This coordinate system appends the phase variable with amplitude variables, which capture the distance, in a Euclidean sense, from the limit cycle. We discuss how the coordinate system is established, and what considerations need to be made in order for it to be practical. We go through some typical neuron examples to show how to implement the system in practice. We proceed to compare the system to one studied in [13] to investigate how the different forms of phase, and now amplitude response functions, can give rise to chaotic dynamics. The orthonormal coordinate system retains information about other structures in the system, whilst preserving the nomenclature and ideas from phase reduction theory.

Chapter 6

Building on from the previous chapter, in which we discuss the effect that noise can have on neural oscillators, we proceed to study techniques for performing phase reductions in the stochastic setting. We review recent literature on the subject, comparing and contrasting results and discuss future direction. We highlight situations in which these techniques break down, in a similar vein to the previous chapter and demonstrate the differences in the alternate interpretations of the stochastic integral. We then present a new coordinate system, ideally suited to describing planar oscillators and show how it may be used to study oscillators away from the weak noise limit. Using techniques primarily used in engineering, we show how to efficiently solve the PDEs to find the steady state probability distribution around the limit cycle in the new coordinate system. Finally, we consider the case where the noise cannot be considered to be weak, and examine the effect of bistability between oscillatory and quiescent modes.

Chapter 7

The final chapter outlines the contribution of the work in this thesis and proposes future work, building upon the ideas contained within it.

Glossary

Below is a list of non-standard abbreviations used in this thesis:

- HH – Hodgkin–Huxley,
- ML – Morris–Lecar,
- FHN – FitzHugh–Nagumo,
- SL – Stuart–Landau,
- PRC – Phase response curve,
- iPRC – Infinitesimal phase response curve,
- IF – Integrate-and-fire,

- LIF – Leaky integrate-and-fire,
- QIF – Quadratic integrate-and-fire,
- LEIF – Linear-exponential integrate-and-fire,
- AIF – Absolute integrate-and-fire,
- PWL – Piecewise-linear,
- PWL-IF – Piecewise-linear integrate-and-fire,
- dHB – Discontinuous Hopf bifurcation,
- dSN – Discontinuous saddle node bifurcation,
- GB – Grazing bifurcation,
- SP – Spiking bifurcation,
- DB – Doublet bifurcation,
- OC – Orbit crisis bifurcation,
- OB – Onset of bistability bifurcation,
- HC – Homoclinic bifurcation,
- LE – Lyapunov exponent,
- MLE – Maximal Lyapunov exponent,
- SNIC – Saddle node on an invariant circle,
- CS – Connor–Stevens,
- SDE – Stochastic differential equation,
- SRK – Stochastic Runge–Kutta,
- PDF – Probability density function,
- CDF – Cumulative distribution function,
- OU – Ornstein-Uhlenbeck,

2 Background

Of the vast numbers of cells located in the brain, the neural cells, or neurons, are thought to be the most important in the transmission of information. Numbering in the hundreds of billions, neurons exist with a variety of morphologies, spatial location and electrical behaviour. Neurons communicate with one another electrically and chemically via gap junctions and synapses. Whilst the neural cell types themselves may be notably distinct, they share a common electrical behaviour: action potentials. Though they may differ from neuron to neuron, collectively, these action potentials, or spikes, are thought to be the key unit of neural information processing. Under the application of a direct current, neurons can be shown to exhibit this typified behaviour *in-vivo*. The action potential comprises two main features: a rapid depolarisation of the cell membrane followed by a slower repolarisation of the membrane potential back to its resting state, possibly following some overshoot. The duration of a single action potential is only a few milliseconds. [14, 15].

2.1 The Hodgkin–Huxley model

In describing the behaviour of neural dynamics at the single cell level, and in particular, the action potential, perhaps the most seminal work is due to Hodgkin and Huxley [16]. The model derived by the authors, for which they won a Nobel Prize in 1963, has become the base formalism for all conductance based neuron models, which are thought of as encompassing a primary class of biophysical models.

As the cell membrane of neurons is otherwise impermeable to ions (aside from through the ion channels we consider below), it acts as a barrier and thus separates charge between the cytosol and the extracellular space. This induces a difference in potential between the exterior and the interior of the cell. We refer to this potential difference as the membrane potential. The model essentially describes the dynamics of ions across the cell membrane and their effect on the membrane potential. Ions are transported through the cell membrane by ion channels, which

are pore forming proteins. Since the ions carry electric charge, this movement induces a current across the membrane, which affects the membrane potential. These ion channels are typically selective for different kinds of ion, and can exist in a variety of states, the most important being the closed and open states. In a closed state, no ions can pass through the channels, whereas ions can pass freely through them in the open state. Ion channels undergo conformational changes to switch between open and closed states. For general ion channels, such changes may be mediated by a number of factors, such as voltage, glutamate, cytosolic calcium concentration and others. The Hodgkin–Huxley (HH) model focuses on the voltage-gated ion channels for sodium Na^+ and potassium K^+ . For each ionic species, there exists a reversal difference between the intracellular and extracellular space, which the ion channels specific to that species will try to maintain. This is achieved by the generation of an electromotive force as the ion channels act to drive the system towards their respective reversal potential. This reversal potential ΔV , also known as the Nernst potential [17], is the potential at which there is no net flow of that ionic species across the membrane. It is given by

$$\Delta V = \frac{RT}{zF} \ln \left(\frac{[\text{ion}]_{\text{out}}}{[\text{ion}]_{\text{in}}} \right), \quad (2.1)$$

where R is the universal gas constant, T is the absolute temperature, z is the valence of the ionic species, F is the Faraday constant and the subscripts denote whether the concentration is inside or outside the cell. Due to active ion pumps, cell membranes are never truly in equilibrium, but the resting potential of a cell, that is, in the absence of any external current, can be calculated from the Goldman equation, written in the form below in [18], building upon earlier work in [19],

$$V_{\text{rest}} = \frac{RT}{F} \ln \left(\frac{\sum_i^N P_{\text{ion}_i^+} [\text{ion}_i^+]_{\text{out}} + \sum_i^M P_{\text{ion}_i^-} [\text{ion}_i^-]_{\text{in}}}{\sum_i^N P_{\text{ion}_i^+} [\text{ion}_i^+]_{\text{in}} + \sum_i^M P_{\text{ion}_i^-} [\text{ion}_i^-]_{\text{out}}} \right), \quad (2.2)$$

in which there are N monovalent positive ionic species and M negative species and P_{ion} denotes the permeability for that ionic species. For the HH model, the resting potential is thus given by

$$V_{\text{rest}} = \frac{RT}{F} \ln \left(\frac{P_{\text{K}}[\text{K}]_{\text{out}} + P_{\text{Na}}[\text{Na}]_{\text{out}}}{P_{\text{K}}[\text{K}]_{\text{in}} + P_{\text{Na}}[\text{Na}]_{\text{in}}} \right). \quad (2.3)$$

Drawing analogies from the world of electrical engineering, we may thus model the membrane as a capacitor. As the ion channels are selectively permeable to ions, we may model these as variable resistors. Finally, the propensity of the ion channels to drive the system to their respective reversal potentials generates an electromotive force, similar to a battery. Combining all of these analogies, we may derive an equivalent electrical circuit, as shown in Fig. 1, to model the dynamics at the membrane, after including an extra leak channel to describe the chloride and other ions not included in this model. We thus have

$$C \frac{dV}{dt} = -I + I_{\text{app}}, \quad (2.4)$$

where C is the capacitance of the membrane, V is the membrane potential, I is the current induced by the movement of the ions across the membrane and I_{app} is a current injected into the cell. As we here consider the effect of the Na^+ and K^+ channels, we apply Kirchoff's second law to the circuit to separate the current into the contribution from the sodium I_{Na} channels, potassium channels I_{K} and a leak current I_{L} , which accounts for all other currents not modelled. This gives us

$$I = I_{\text{Na}} + I_{\text{K}} + I_{\text{L}}. \quad (2.5)$$

It remains to describe the dynamics of these currents. Treating the membrane as being permeable to ions, we can model the currents using Ohm's law as

$$I_X = g_X(V - V_X), \quad X \in \{\text{Na}, \text{K}, \text{L}\}. \quad (2.6)$$

where g_X is the conductance, the inverse of resistance, for the ion X and V_X is the reversal potential of the ion X . Hodgkin and Huxley's great achievement was to model the dynamics of the conductance terms g_X . They postulated that these conductances were voltage sensitive. Using a solution of choline, rather than sodium chloride to perform their experiments, they effectively knocked out the Na^+ channel, and proceeded to measure the conductance of K^+ , independent of Na^+ , at different voltages of a voltage clamp experiment. In this way, they were able to make empirical fits of the data to an exponential. To model the Na^+ current,

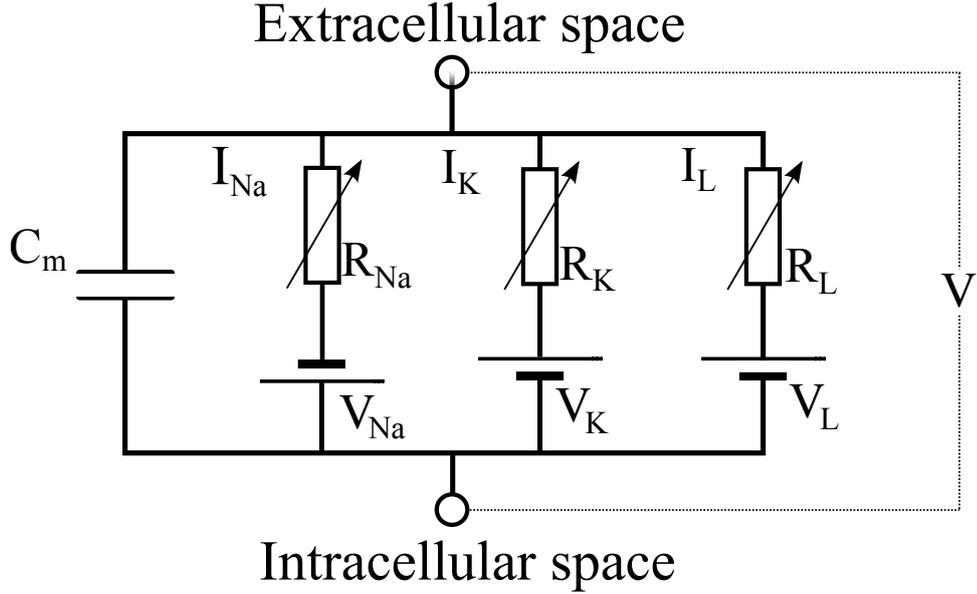


Figure 1: Equivalent circuit diagram of the HH model.

they compared data from the choline bath to that of a regular sea water bath to extract the change in Na^+ current, and again, fitted the data to an exponential. Hodgkin and Huxley found that the K^+ current could be modelled as

$$g_K = \bar{g}_K n^4, \quad (2.7)$$

where \bar{g}_K is the maximal conductance. In the above $n \in [0, 1]$ may be regarded as a *gating* variable. It was later shown that the nonlinear dependence of g_K on n was due to the gating charge [20], that is, the number of gating particles needed to be bound at the ion channel for the channel to be in an open state. In the case of the K^+ channel, the gating charge is four, hence the quartic relationship. The dynamics of n are given by

$$\frac{dn}{dt} = \alpha_n(V)(1 - n) - \beta_n(V)n, \quad (2.8)$$

where α_n is the rate of transfer of particles to the specified location to bind to the ion channel and β_n is the rate of transfer away from this location. Letting $t \rightarrow \infty$, we may define the steady state of (2.8) as

$$n_\infty(V) = \frac{\alpha_n(V)}{\alpha_n(V) + \beta_n(V)}. \quad (2.9)$$

By fitting to data, Hodgkin and Huxley found the following forms for α_n and β_n

$$\alpha_n = \frac{0.01(V + 10)}{\exp\left(\frac{V+10}{10}\right) - 1}, \quad (2.10)$$

$$\beta_n = 0.125 \exp\left(\frac{V}{80}\right). \quad (2.11)$$

Using these forms for α_n and β_n in (2.9), n_∞ can be seen to have a sigmoidal shape, suggesting that over a small range of V , the conductance of K^+ exhibits a large change. It is convenient to rewrite (2.8) as

$$\tau_n(V) \frac{dn}{dt} = n_\infty(V) - n, \quad (2.12)$$

where

$$\tau_n(V) = \frac{1}{\alpha_n(V) + \beta_n(V)}, \quad (2.13)$$

so that $\tau_n(V)$ acts as a voltage-dependent time constant for the dynamics of n . The equations governing the dynamics of Na^+ are written in a similar way. This time, we find the conductance of Na^+ rises as a cubic and that an extra, inactivating term is needed to describe the behaviour at large times. Thus, the conductance of Na^+ is written as

$$g_{Na} = \bar{g}_{Na} m^3 h, \quad (2.14)$$

where m may be thought of as an activating gating variable, whilst h is an inactivating gating variable and, as before, \bar{g}_{Na} is the maximal conductance of Na^+ . The dynamics of m and h are given in a similar form to those of n , though the direction of the sigmoid for h_∞ is reversed compared with those of n_∞ and m_∞ . These sigmoids are shown in Fig. 2. The details of the gating variables are completed as follows

$$\tau_X(V) \frac{dX(V)}{dt} = X_\infty(V) - X, \quad X \in \{m, h\} \quad (2.15)$$

where

$$X_\infty = \frac{\alpha_X(V)}{\alpha_X(V) + \beta_X(V)}, \quad (2.16)$$

with

$$\alpha_m = \frac{0.1(V + 25)}{\exp\left(\frac{V+25}{10}\right) - 1}, \quad \beta_m = 4 \exp\left(\frac{V}{18}\right), \quad (2.17)$$

$$\alpha_h = 0.07 \exp\left(\frac{V}{20}\right), \quad \beta_h = \frac{1}{\exp\left(\frac{V+30}{10}\right) + 1}. \quad (2.18)$$

Finally, the leak current, taking into account the total effect of all ion channels not considered, is simply modelled as

$$I_L = \bar{g}_L(V - V_L), \quad (2.19)$$

that is, we do not include any gating variables for this current. Overall, the current-balance equation for the membrane potential reads

$$C \frac{dV}{dt} = F(V, m, n, h) + I_{\text{app}}, \quad (2.20)$$

where,

$$F(V, m, n, h) = -\bar{g}_L(V - V_L) - \bar{g}_K n^4 (V - V_K) - \bar{g}_{\text{Na}} m^3 h (V - V_{\text{Na}}). \quad (2.21)$$

The HH model may then be shown to capture the behaviour of an isolated neuron under the application of a direct stimulus current, both in the form of the spike generated by the model, and the frequency of firing under direct applied current. We show some typical behaviour of the HH model in Fig. 3. The top panel in this figure shows the membrane potential over time when the neuron is firing. The rapid depolarisation and hyperpolarisation are easily identified, along with a refractory period in which no spiking occurs. The lower panel shows the evolution of the gating variables which, as for the membrane potential, evolve slowly between spikes and rapidly during them. It is clear that m is the fastest of the three variables, reaching a state in which all channels are open/closed very rapidly. It may also be seen that n and h follow similar time courses.

2.2 Dimension reduction of the Hodgkin–Huxley model

The HH model as defined above has four variables: a membrane potential, two activating gating variable and one inactivating gating variable. In general, due to a

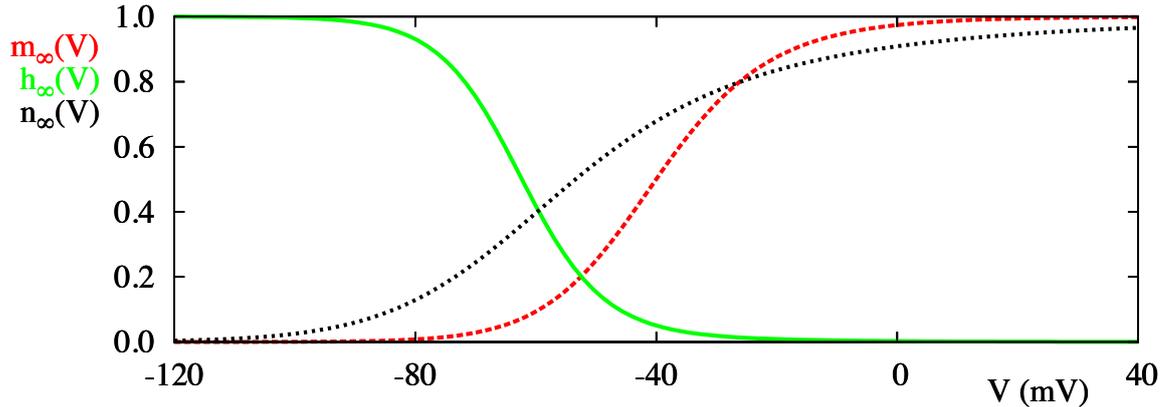


Figure 2: Sigmoidal shapes of $X_\infty(V)$, $X \in \{m, n, h\}$. The inverse direction of h_∞ compared with n_∞ and m_∞ belies its nature as an inactivating, rather than activating gate. Parameter values are $V_{\text{Na}} = 50$ mV, $V_{\text{K}} = -77$ mV, $V_{\text{L}} = -54.4$ mV, $\bar{g}_{\text{Na}} = 120$ mmho/cm², $\bar{g}_{\text{K}} = 36$ mmho/cm², $\bar{g}_{\text{L}} = 0.3$ mmho/cm², $C = 1$ μ F/cm².

high number of dimensions, phase space analysis can be cumbersome. Throughout this thesis, we shall take dimension to refer to the number of variables in a given model. The analysis of planar systems is much easier than that of systems with higher dimension. To gain a better understanding of the dynamics of the HH model, it is thus useful to reduce it to a planar system. The method of equivalent potentials, as described in [21], gives us the machinery to make this reduction for *fast-slow* systems, in which there exists a separation of time scales between subsets of variables.

We first note that $\tau_m(V)$ is small for all V , so that m approaches rapidly approaches its equilibrium value, compared with n and h . Secondly, we observe that n and h share similar time courses. In fact, the Na^+ channel closing occurs at the same rate, but in the opposite direction to K^+ . Hence, n and h may be slaved to another variable, U . By assuming that m is infinitely fast compared with the other model variables, we make a quasi-steady state approximation, $m(V) = m_\infty(V)$, and thus eliminate the dynamics of m . We mimic the slower approach of n and h

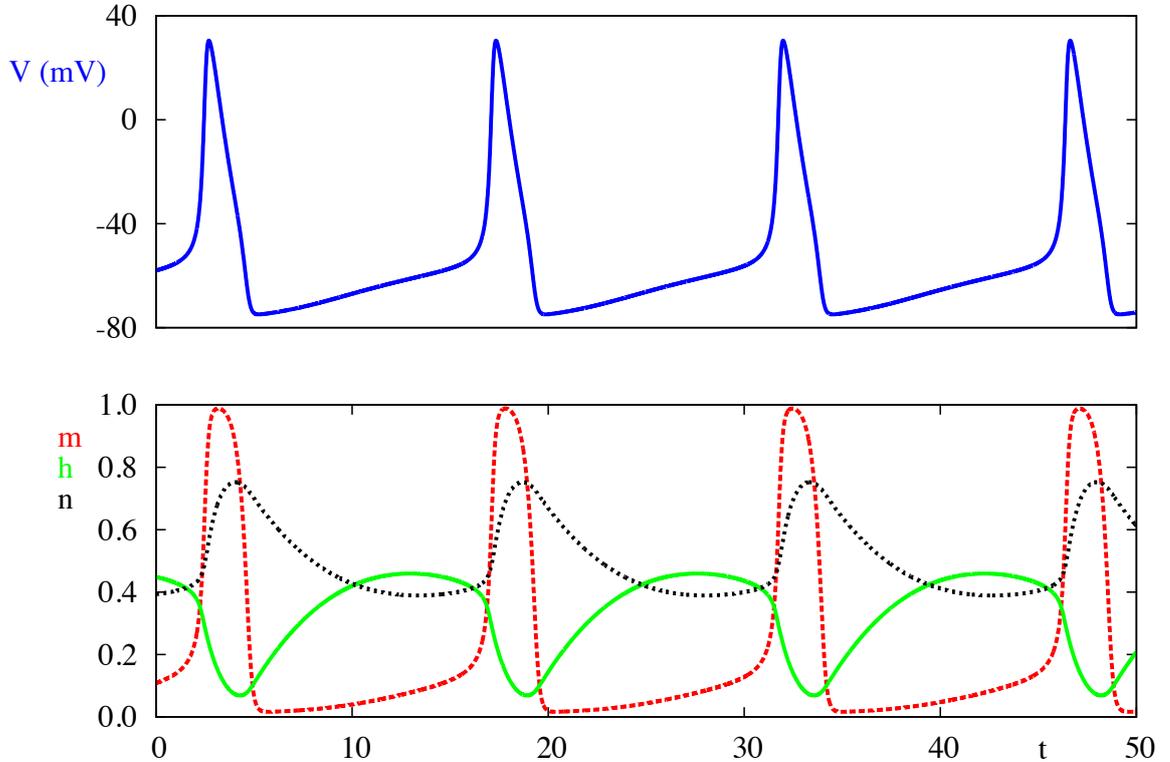


Figure 3: Spiking solutions of the full HH model. The top panel shows the membrane potential over time. For each action potential, we see a rapid depolarisation at the time of initiation, followed by a rapid hyperpolarisation and then a refractory period where no firing occurs. The lower panel shows the time evolution of the gating variables. As for the membrane potential, the gating variables evolve slowly between spikes, but rapidly during them. This is most pronounced for m , which quickly reaches a state in which all channels are open/closed. We also observe that n and h follow similar time courses. Parameters are as in Fig. 2, with $I_{\text{app}} = 10 \text{ mA/cm}^2$.

to their asymptotic values by the replacement

$$X = X_\infty(U), \quad X \in \{n, h\}. \quad (2.22)$$

This equation can always be solved for U , since the functions X_∞ are monotonic, and hence invertible. We can now rewrite the sodium current as a function of U , so that the entire system is defined by the variables V and U . The membrane current is now given by

$$f(V, U) = F(V, m_\infty(V), n_\infty(U), h_\infty(U)). \quad (2.23)$$

By demanding that the time dependence of the voltage dynamics of the original system, F , match those of the reduced system, f , for a given value of V , we have that

$$\frac{\partial F}{\partial h} \frac{dh(V)}{dt} + \frac{\partial F}{\partial n} \frac{dn(V)}{dt} = \left(\frac{\partial f}{\partial h_\infty} \frac{dh_\infty}{dU} + \frac{\partial f}{\partial n_\infty} \frac{dn_\infty}{dU} \right) \frac{dU}{dt}. \quad (2.24)$$

Under the approximation that $h \approx h_\infty(U)$ and $n \approx n_\infty(U)$, we may solve for the time evolution of U to obtain

$$C \frac{dV}{dt} = -f(V, U) + I_{\text{app}}, \quad (2.25)$$

$$\frac{dU}{dt} = g(V, U), \quad (2.26)$$

where

$$g(V, U) = \frac{\frac{dF}{dh} \left[\frac{h_\infty(V) - h_\infty(U)}{\tau_h(V)} \right] + \frac{dF}{dn} \left[\frac{n_\infty(V) - n_\infty(U)}{\tau_n(V)} \right]}{\frac{\partial f}{\partial h_\infty} \frac{dh_\infty(U)}{dU} + \frac{\partial f}{\partial n_\infty} \frac{dn_\infty(U)}{dU}} \quad (2.27)$$

and $\partial F/\partial h$ and $\partial F/\partial n$ are evaluated at $(h, n) = (h_\infty(U), n_\infty(U))$. The variable V describes the capacitive nature of the cell, whilst U describes the time-dependence of the membrane conductance. In fact, U may be regarded as the variable responsible for the refractory period of the cell. In many planar models of neurons, such as the Morris–Lecar [22] and FitzHugh–Nagumo [23, 24] models, the second variable is typically regarded as a recovery variable, which accounts for the dynamics of all of the gating variables, similar to U in the reduction by equivalent potentials.

Now that the system is planar, we may carry out phase-plane analysis in full. In Fig. 4, we plot the nullclines and direction field of the reduced model. The

V -nullcline is depicted in green and red, whilst the U -nullcline is shown as dashed black line. The arrows show the direction of the vector field at those points at which they are plotted. The V -nullcline has a roughly cubic shape. This nullcline may be thought of as being comprised of three branches, separated by extrema of the nullcline. The U -nullcline, meanwhile, is approximately linear. When $I_{\text{app}} < 0$, there exists a solitary stable fixed point, on the left hand branch of the V -nullcline. In fact, where they exist, any fixed point on the left or right hand branch of the V -nullcline is stable, whilst fixed points existing on the middle branch are unstable. The V -nullcline is coloured green where fixed points are stable and red where they are unstable. At $I_{\text{app}} = 0$, the system undergoes a Hopf bifurcation, signified by a zero crossing of the real part, with non-zero imaginary part, of the eigenvalues of the linearisation around the fixed point. This leads us to the situation depicted in Fig. 4, in which the fixed point is now on the middle, unstable branch. As there are no other fixed points in the system, an attracting limit cycle solution can easily be shown to exist by the Poincaré-Bendixson theorem. This limit cycle corresponds to the repetitive firing of the neuron under a constant injected current, of $I_{\text{app}} = 15 \text{ mA/cm}^2$ in this case. The limit cycle is shown in blue. At the bifurcation point, the frequency of these oscillations may simply be read off by examining the imaginary part of the eigenvalues of the linearisation around the fixed point. As I_{app} is increased further, these oscillations terminate in another Hopf bifurcation, so that there exists only a stable fixed point once again, now on the right hand branch of the V -nullcline. The HH model may thus be thought of as an excitable model, mimicking the electrical excitability of neurons. With no injected current, the neurons simply stay at a quiescent state. Within a certain range of injected current, we observe repetitive spiking. At the bifurcation point, the HH model supports oscillations with nonzero frequency. However, certain neurons, particularly those in the cortex, have been observed to spike with arbitrarily low frequency. We refer to this case as type I excitability, whilst the previous case is type II excitability.

The Morris–Lecar (ML) model [22] is a planar model, built on a HH formalism,

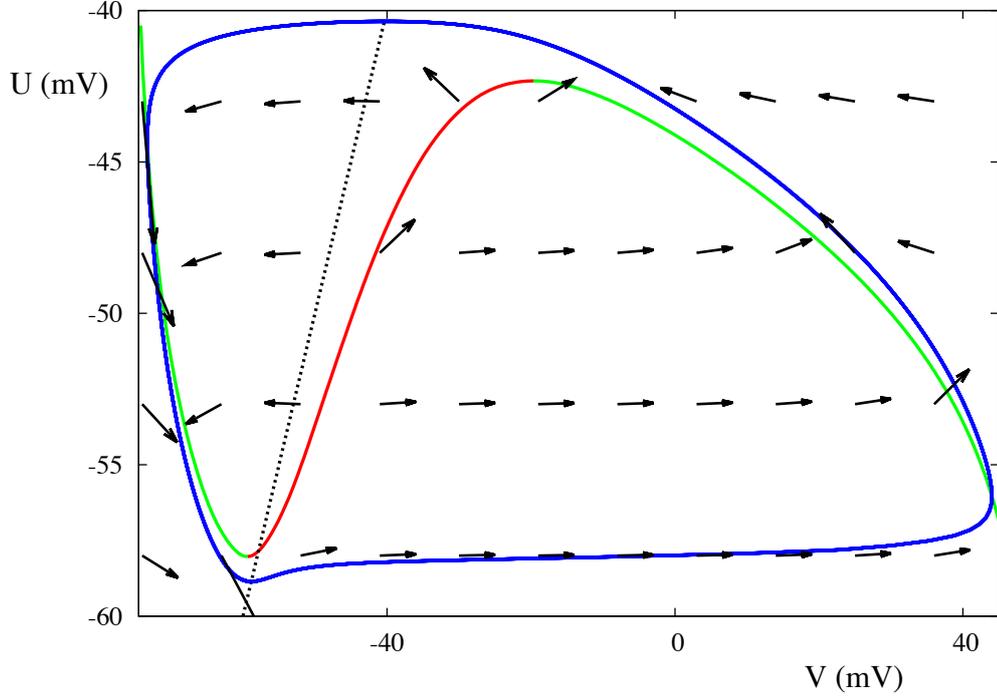


Figure 4: Phase plane of the reduced HH model (2.26). The red/green curve is the V -nullcline, along which $\dot{V} = 0$, whilst the black dashed line is the U -nullcline. The black arrows correspond to the direction of the vector field at the points at which they are plotted. These arrows point in the direction of the flow induced by the vector field. With an injected current of $I_{\text{app}} = 15 \text{ mA/cm}^2$, the only fixed point of the system (at the intersection of the two nullclines) is on the middle, unstable branch of the V -nullcline, and we observe limit cycle solutions, indicated by the blue curve. Note that the limit cycle tracks along the V -nullcline for a good portion of its orbit, and that the direction field primarily points to the left or right, apart from close to the V -nullcline. This is typical of excitable systems. Other parameters are as in Fig. 3.

including K^+ and calcium (Ca^{2+}) channels. The equations for this model are

$$C \frac{dV}{dt} = I_{app} - \bar{g}_L(V - V_L) - \bar{g}_K w(V - V_K) - \bar{g}_{Ca} m_\infty(V)(V - V_{Ca}), \quad (2.28)$$

$$\frac{dw}{dt} = \phi(w_\infty(V) - w)/\tau_w(V), \quad (2.29)$$

where

$$\begin{aligned} m_\infty(V) &= \frac{1}{2} [1 + \tanh((V - V_1)/V_2)], & \tau_w(V) &= 1/\cosh((V - V_3)/(2V_4)), \\ w_\infty(V) &= \frac{1}{2} [1 + \tanh((V - V_3)/V_4)]. \end{aligned} \quad (2.30)$$

As for the reduced HH model, the recovery variable w in (2.28)-(2.29) captures the dynamics of the gating variables for both K^+ and Ca^{2+} . The ML model was originally developed to describe the electrical activity in barnacle muscle fibres. However, due to its simplicity and ability to describe a range of electrical behaviours, it has seen widespread use in the modelling of cortical tissue. The ML model can display both type I and type II excitability, dependent on the parameter values as provided in [15]. In fact, the excitability of a given model is related to the bifurcation types leading to oscillatory behaviour. Hopf bifurcations give rise to type II excitability, whereas other bifurcations, such as the saddle node on an invariant circle (SNIC), as supported by the ML model will give rise to type I excitability. We will discuss the SNIC bifurcation in more detail later in this work. In Fig. 5, we show the dependence of frequency on I for the two different excitability types. Suppose that a neuron is in a quiescent state. Instead of forcing with a constant drive, we now inject a transient current into the neuron. This current mimics that of a synaptic current that may occur after the firing of a spike by a presynaptic cell. We keep the duration of the injected current fixed, whilst varying the magnitude of the input. If the magnitude of the current is small, we may see a small rise in the membrane potential before returning to its resting potential, but no spiking event, by which we mean a rapid, large amplitude excursion from the resting state. If the magnitude of the current is now increased, we may force the neuron to fire a spike before it returns to rest. Cells will typically only fire one spike, or one burst of spikes in response to a transient current of this

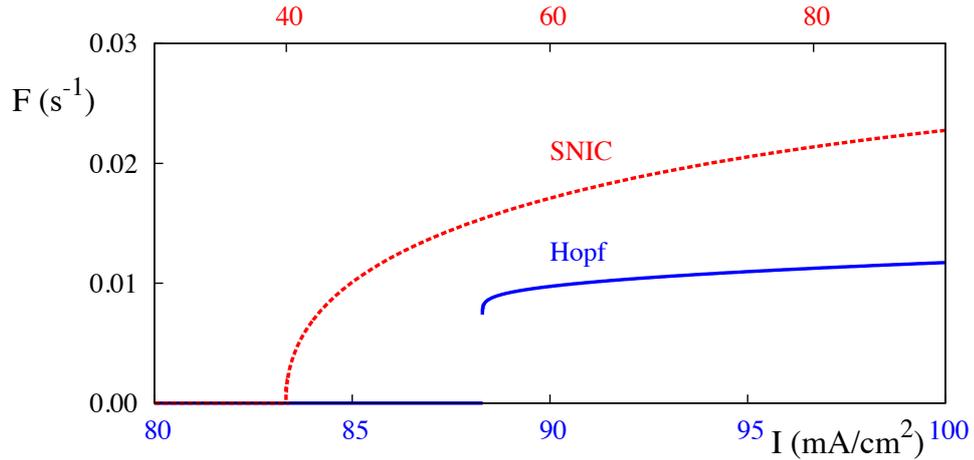


Figure 5: Comparison of $F - I$ curves for type I and II excitability. The SNIC bifurcation gives rise to type I excitability, marked by a zero firing frequency at the bifurcation point. The Hopf bifurcation gives rise to type II excitability, with a nonzero firing frequency at the bifurcation point. Parameters are $C = 20.0 \mu\text{F}/\text{cm}^2$, $\bar{g}_K = 8.0 \text{ mmho}/\text{cm}^2$, $\bar{g}_L = 2.0 \text{ mmho}/\text{cm}^2$, $V_K = -84.0 \text{ mV}$, $V_L = -60.0 \text{ mV}$, $V_{Ca} = 120.0 \text{ mV}$, $V_1 = -1.2 \text{ mV}$, $V_2 = 18.0 \text{ mV}$, with $\phi = 0.067$, $\bar{g}_{Ca} = 4.0 \text{ mmho}/\text{cm}^2$, $V_3 = 12.0 \text{ mV}$, $V_4 = 17.4 \text{ mV}$ for the SNIC case, and $\phi = 0.04$, $\bar{g}_{Ca} = 4.4 \text{ mmho}/\text{cm}^2$, $V_3 = 2.0 \text{ mV}$, $V_4 = 30.0 \text{ mV}$ for the Hopf case.

type, so that there are no self sustained oscillations. We denote this kind of behaviour *excitable*, as the neuron may be induced to fire under a transient current, even when the underlying dynamics are not oscillatory. There exists a minimum value for the magnitude of the current that will elicit a spike in response to the transient stimulus. This *threshold* value for the current may be seen more clearly through a phase-plane analysis of the reduced HH model, for example. We note that a cubic v -nullcline, as possessed by this reduced HH model is present in many mathematical descriptions of excitable systems.

In Fig. 6, we depict this situation. The V -nullcline is now shown only in green, in contrast to before, whilst the U -nullcline is still depicted by the dashed black curve. The base current in this situation is $I_{\text{app}} = -5.0 \text{ mA}/\text{cm}^2$, which is a

hyperpolarising current, keeping the neuron at a fixed point x^* on the stable, left hand branch of the V -nullcline. We choose initial conditions at this fixed point and approximate the transient stimulus by a Dirac delta given at time $t = 5$, scaled by a factor I_0 , so that the total stimulus current injected to the system is I_0 . For a sub-threshold current, indicated by the blue curve, trajectories simply decay back to x^* , but for a sufficiently large current, indicated by the red curve, the system will make a large amplitude excursion before returning to x^* . Examining the time course of V , shown in the lower panel of Fig. 6, during this excursion shows that it matches the time course of a spike. For planar models, with sufficiently large separation of timescales between variables, this threshold may be approximated by the middle branch of the V -nullcline. This notion of a threshold, leads us to consider a far simpler phenomenological model, which is more amenable to mathematical analysis.

2.3 The integrate-and-fire neuron

Originally considered as a model of frog nerve stimulation by Lapicque in 1907 [25], the integrate-and-fire (IF) model actually predates the HH model, although the model itself was not studied mathematically or even termed ‘integrate-and-fire’ until the 1960s in works by Knight and Stein [26, 27, 28]. The model takes a far simpler view to modelling nerve behaviour, omitting all of the details of the gating variables as considered in the HH model. One of the most basic model of this type is the leaky IF model (LIF), which is a simple application of Ohm’s law on the simple circuit shown in Fig. 7. Thus, the model is

$$\frac{dV}{dt} = -\frac{V}{\tau} + I(t), \quad (2.31)$$

where $\tau = RC$ is time constant of the membrane. The threshold, V_{th} , is introduced by demanding that when $V = V_{\text{th}}$, V is reset to a value V_r , with $V_r < V_{\text{th}}$. More formally:

$$V \mapsto V_r \quad \text{when} \quad V(t) = V_{\text{th}}. \quad (2.32)$$

Between threshold crossings, (2.31) simply acts as an integrator for the signal

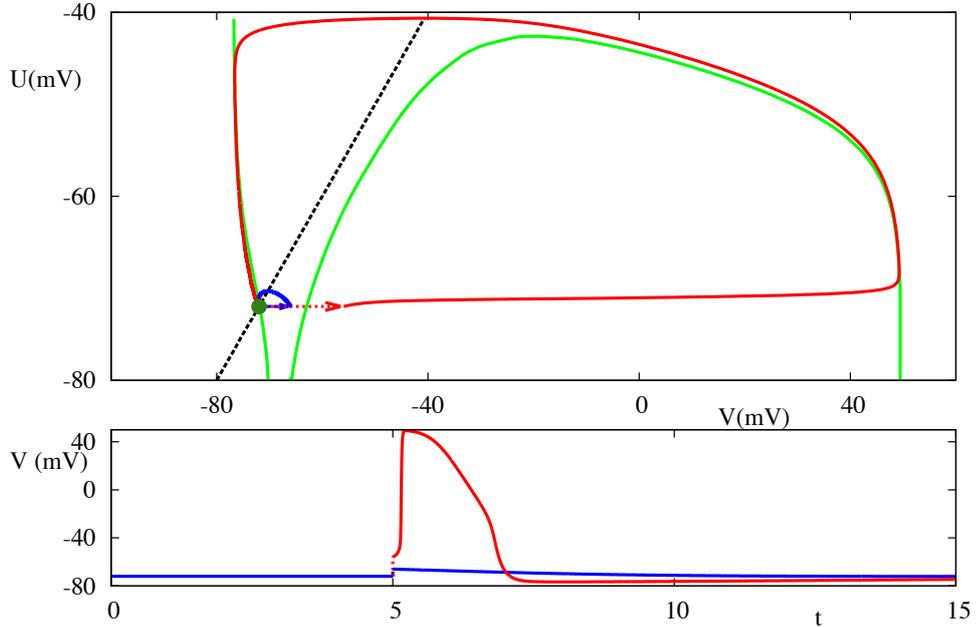


Figure 6: The top panel shows the phase plane of the reduced HH model in a quiescent mode, with the parameters matching those of Fig. 3, but with $I_{\text{app}} = -5.0 \text{ mA/cm}^2$. In the absence of any other inputs, the system has a stable fixed point, illustrated by a green circle on the left hand branch of the V -nullcline, depicted as a dashed orange curve, whilst the U -nullcline is shown by a dashed cyan line. Introducing an instantaneous, transient current may result in one of two scenarios. If the transient current is not strong enough, the resulting trajectory, shown in blue, will simply decay back to the fixed point. For a large enough injected current, shown by the red curve, the trajectory will make a large amplitude excursion before returning to the fixed point. An examination of the time course of the excursion, presented in the lower panel, shows that this excursion is indeed a spike. This behaviour is indicative of excitable behaviour that, although the neuron is at rest, it may be induced to fire by a transient stimulus. That the stimulus has to be strong enough, implies that there is some threshold for the stimulus to reach before any spikes will be fired.

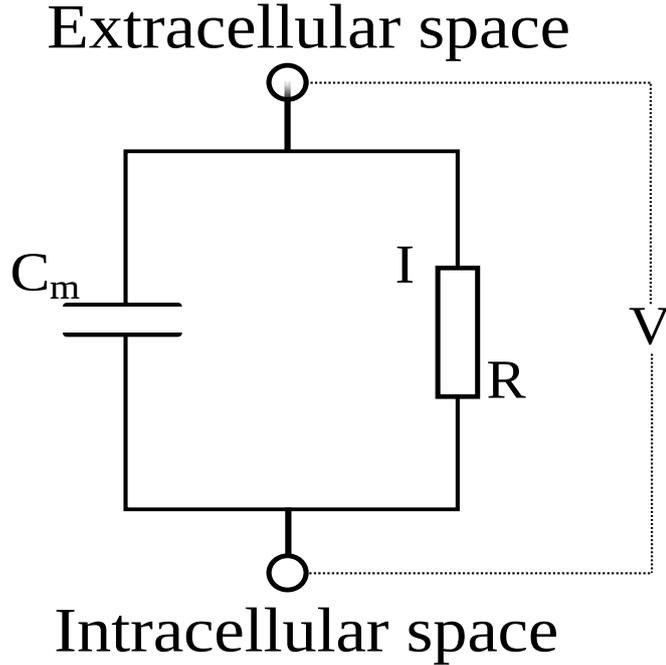


Figure 7: Circuit diagram of the IF model

$I(t)$. Using an integrating factor and with an initial condition of $V(t_0)$ at $t = t_0$, we may write down solutions to (2.31), as

$$V(t) = V(t_0)e^{-(t-t_0)/\tau} + \int_{t_0}^t e^{-(t-s)/\tau} I(s) ds. \quad (2.33)$$

For a constant drive $I(t) = I_0$, the system will simply tend towards $V = I_0\tau$. If $I_0\tau < V_{\text{th}}$, then the system will remain at this subthreshold membrane potential. If $I_0\tau > V_{\text{th}}$, then trajectories will reach $V = V_{\text{th}}$ and will successively be reset to $V = V_r$ at which point V will again increase towards $I_0\tau$, and the system will support non-smooth oscillations. The period of these oscillations may be calculated as

$$\Delta = \tau \ln \left(\frac{I_0\tau - V_r}{I_0\tau - V_{\text{th}}} \right) H(I_0\tau - V_{\text{th}}), \quad (2.34)$$

where H denotes the Heaviside step function. In such a regime, (2.31) may be thought of as representing a neuron undergoing repetitive firing under a constant drive, for example, system (2.26) in an oscillatory regime. Of course, the time course of the membrane potential, which we hereon refer to as the voltage trace, bears no resemblance to voltage traces of real neurons. In particular, there are no

spikes. In this regard, the LIF model may be thought as representing the timing of spikes, rather than their shape. Typically, a spike shape is appended to the voltage traces at the threshold crossing times to give something resembling a voltage trace from real neurons, for example as in [29].

The advantages of the LIF over biophysical models are two-fold. Firstly, given its simplicity, the LIF is mathematically tractable and amenable to deeper analysis. Secondly, the model is computationally cheap to solve. This allows for large network studies to be performed in reasonable CPU time, as considered in [30], for example. In spite of these advantages, one may ask in what way a model that does not innately describe the time course of the modelled variable, in this case voltage, can be a useful model. Whilst the spike shape of individuals may differ from one another, there is a remarkable similarity within the spike shapes from the same neuron. The factors contributing to the spike shape for individual neurons is considered in [31], wherein the authors find that it is the combination of the nonlinear ionic currents that is the key factor in determining the waveform of the spikes. This fact underpins spike sorting algorithms for separating time courses of membrane potentials that may arise from a number of different neurons being measured by the same electrode. The similarity of spikes from the same neuron leads to the hypothesis that spike shapes do not encode information. Instead, it is presumed that, for neurons, information is encoded either in the precise timing of spikes, or in the instantaneous firing rate of a neuron [32, 33, 34]. In this case, the LIF model is then a good model of a cell embedded within a network, in which we are interested in the way signals are communicated. The LIF model does, however, fail to capture physical processes that other models do, such as post-inhibitory rebound, in which a neuron fires a spike in response to a long, transient hyperpolarising current. Thus, the choice of which class of models to use to represent a network of cells is very much dependent on the phenomena we wish to study. An example of a typical trajectory for the LIF neuron (2.31) is shown in Fig. 8

Building upon the LIF, there is now a range of different models of the IF type.

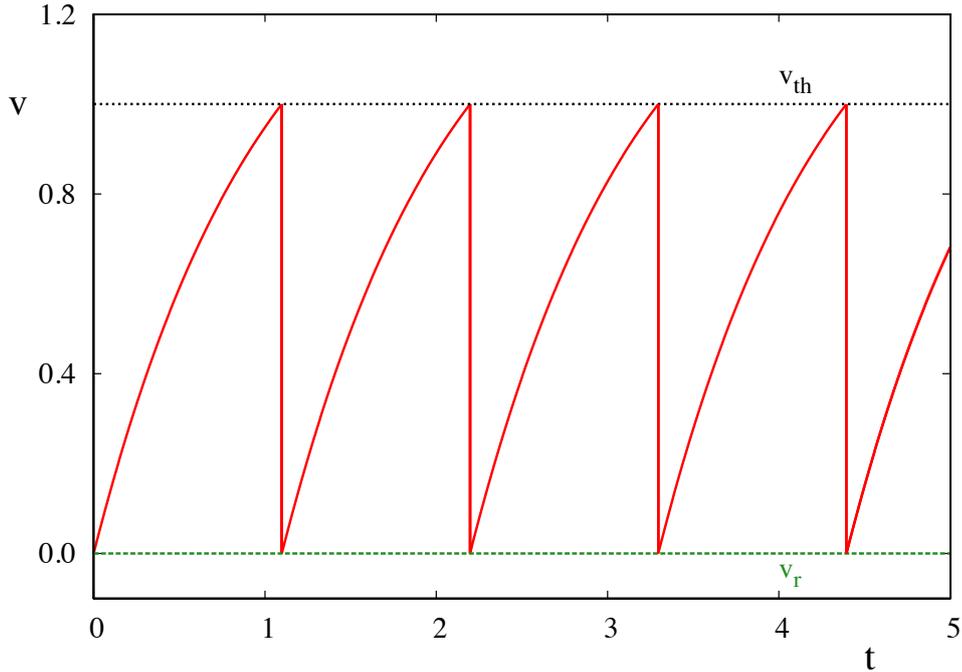


Figure 8: Trajectories of the LIF model (2.31) with $\tau = 1$, $I_0 = 1$, $V_{\text{th}} = 1$ and $V_r = 0$. Each time the membrane potential reaches V_{th} , a spike is fired and V is reset to V_r . The shape of the spike is typically appended to the voltage trace, as the LIF model does not produce the same kinds of spike shapes as produced in biophysical models, such as the HH model (2.21). The resulting trajectories are composed of a smooth orbit and a discontinuous ‘jump’.

Popular models include the Izhikevich model [35], a planar model that can support a variety of firing types, including bursting, which is an important firing pattern in the brain; the adaptive exponential IF model [36], which has been fitted to data with remarkable accuracy and the quartic IF model [37]. To reduce the number of dimensions in biophysical models, hard thresholds of the form (2.32) may be used to replace the recovery variables. For example, for the system (2.26), we may obtain a linear model by simply setting $V = U$ everywhere. This is quite a severe approximation, since V is not typically close to U except at the fixed point. The largest value of V for which $F(V) = f(V, V) = 0$ determines the onset of the oscillatory regime and can be used to determine the threshold for firing, since

$dV/dt > 0$ for large enough values of V .

As well as being used to model biological processes in the brain, IF models present novel mathematical problems. The non-smooth nature of the reset mechanism may give rise to interesting dynamics and bifurcations that are not observed in similar smooth systems. The classification of different kinds of nonsmooth behaviour is an active area of research outside of mathematical biology, as can be seen by the recent books on the topic: [3, 38, 39]. The effect of such behaviour on emergent dynamics at the network level is also an intriguing open topic along with efficient and accurate simulation of such networks [40].

We have already seen that, for a constant stimulus $I(t) = I_0\tau > V_{\text{th}}$, the LIF model (2.31) supports oscillations with period given by (2.34). Suppose we are interested in knowing how perturbations to the stimulus current affects the timing of the next spike. To this end, we split the injected signal into two components $I(t) = I_0 + I_s$, where I_s is the perturbation to the mean current I_0 . Since we are only interested in the times of spiking events, we can use the nonlinear transformation

$$\theta(t) \bmod 1 = \psi(V) = \frac{1}{\Delta} \int_0^{V(t)} \frac{dU}{I_0\tau - U} = \frac{1}{\Delta} \ln \left(\frac{I_0\tau}{I_0\tau - V} \right) H(I_0\tau - V_{\text{th}}), \quad (2.35)$$

to recast the dynamics in terms of a unitless *phase* variable $\theta \in [0, 1)$. Identifying $\theta = 0$, $V = V_r$, and $\theta = 1$, $V = V_{\text{th}}$ completes the specification of the new variable. Now, every time $\theta = 1$, the neuron fires a spike. Using (2.35), we may rewrite (2.31) as

$$\frac{d\theta}{dt} = \frac{1}{\Delta} + I_s(t)R(\theta), \quad (2.36)$$

where

$$R(\theta) = \frac{e^{\theta\Delta}}{\Delta I_0}. \quad (2.37)$$

The object R is the phase response curve (PRC). It quantifies the change in phase, and thus, the change in spike timing, in response to perturbations of the injected current. Thus, (2.31) has been transformed into dynamics around a circle with constant speed where we can clearly identify the change in spike timings in response to perturbations. If $I_s(t) = \delta(t - n\tau)$, then the new phase after the n 'th pulse is simply $R(\theta(n\tau))$. It is clear that we can uncover stability properties of

solutions for a forced IF neuron solutions using the PRC. The same idea carries over to networks of coupled IF cells by treating the signal received from each of the cells as a perturbation to its own natural periodic orbit. We will go on to use such techniques in the next chapter to study a network of two interacting IF type cells.

2.4 Isochronal coordinates

The phase variable considered in (2.35), is an exact, nonlinear transformation. Phase variables are not restricted to the analysis of IF models, however. In fact, they can be used for any system supporting limit cycles. Suppose we have a system of the form

$$\frac{dx}{dt} = f(x) + \varepsilon g(x), \quad x \in \mathbb{R}^n, \quad (2.38)$$

which possesses an asymptotically stable limit cycle, Γ , for $\varepsilon = 0$ having period Δ . The set Γ corresponds to the image of an interval $[0, \Delta)$ under x with $x(0) \in \Gamma$. This gives us a natural way to define a phase variable on the limit cycle itself. Setting $\theta = t/\Delta$ assigns a phase to each point, and thus Γ is now the image of $[0, 1)$ under x . The choice of the exact parametrisation of θ is a personal preference. Other common choices include $\theta = t$, and $\theta = 2\pi t/\Delta$, so that $\theta \in [0, 2\pi)$, as shown for a typical oscillatory solution of the HH model in Fig. 9. Each point on Γ is then parametrised by a unique phase. Of course, there are infinitely many different choices for this parametrisation, dependent on where we assign the zero-phase $\theta = 0$. As we are principally interested in the timing of spiking events, it is natural to assign $\theta = 0$ to the maximum value of the membrane potential. This now uniquely assigns a phase to each point in the limit cycle. Note that the points $\theta = 0$ and $\theta = 1$ are identified and that, on the limit cycle, the phase variable evolves with a constant velocity, $\dot{\theta} = 1/\Delta$, as before. On cycle, we have therefore reduced nonlinear dynamics in \mathbb{R}^n to constant velocity dynamics on \mathbb{S}^1 .

We now ask the question of how to consider the effect of perturbations, that is, of setting $\varepsilon \neq 0$. If we denote the flow of the unforced system by Φ_t , then the

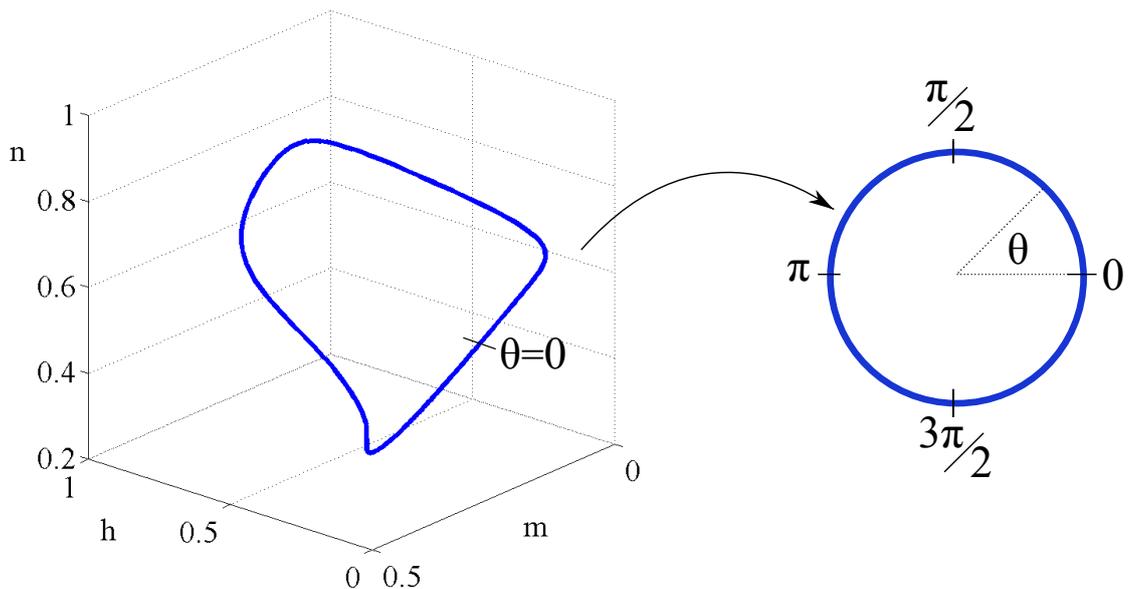


Figure 9: Parametrisation by θ of the orbit depicted in Fig. 3.

basin of Γ , denoted \mathcal{B}_Γ , is defined to be

$$\mathcal{B}_\Gamma = \{x \in \mathbb{R}^n | \Phi_t(x) \rightarrow \Gamma \text{ as } t \rightarrow \infty\}, \quad (2.39)$$

that is, the set of all points for which the flow under f tends to Γ in infinite time. We can assign to every point in \mathcal{B}_Γ an *asymptotic phase*. This asymptotic phase is the phase along the periodic orbit to which a trajectory evolved under the flow, $\Phi(t, x_0)$, of (2.38) with $\varepsilon = 0$, starting from a point $x_0 \in \mathcal{B}_\Gamma$, will tend to as $t \rightarrow \infty$. More formally, the asymptotic phase $\theta(x_0)$ of a point x_0 is given by the condition

$$\lim_{t \rightarrow \infty} \Phi(t, x_0) - \Phi(t + \theta(x_0)\Delta, \gamma_0) = 0, \quad (2.40)$$

with $\gamma_0 \in \Gamma$, being the point of zero phase on Γ . In essence, two points share the same asymptotic phase if the distance between trajectories with those points as initial conditions vanishes as $t \rightarrow \infty$. In this way, we can extend the notion of phase to points away from cycle in \mathcal{B}_Γ . *Isochrons*, as reviewed in [41], are a useful notion when studying perturbations to limit cycle oscillators; they are level sets of the asymptotic phase, i.e. they connect points of the same asymptotic phase across the whole basin of attraction of the limit cycle. Knowledge of what

isochron a point is on fully specifies the asymptotic behaviour of the flow from that point, in the absence of perturbations [42],[43]. In turn, this means that isochrons tell us whether perturbations advance or retard the phase of the oscillator, and in particular tell us whether the timing of the next spike is earlier or later. The isochron for a given phase is the level set of the asymptotic phase function over \mathcal{B}_Γ

$$\mathcal{I}(\gamma) = \{x \in \mathcal{B}_\Gamma \mid \theta(x) = \theta(\gamma), \gamma \in \Gamma\} \quad (2.41)$$

The set of isochrons $\forall \theta \in [0, 1)$ thus foliates the entire basin of attraction. Consider the simple planar example of the Stuart–Landau (SL) oscillator:

$$\begin{aligned} \dot{x} &= \frac{1}{2}\lambda x - \left(\frac{1}{2}\lambda c + \omega\right) y - \frac{1}{2}\lambda(x^2 + y^2)(x - cy), \\ \dot{y} &= \left(\frac{1}{2}\lambda c + \omega\right) x + \frac{1}{2}\lambda y - \frac{1}{2}(x^2 + y^2)(cx + y). \end{aligned} \quad (2.42)$$

The limit cycle for this system is the unit circle, and a coordinate transformation shows that the isochrons for this system can be given parametrically as $x = (1 + r) \cos(\theta + c \ln(1 + r))$, $y = (1 + r) \sin(\theta + c \ln(1 + r))$, for $r \in (-1, \infty)$ [7]. In a similar vein, the isochrons for the classical Hopf bifurcation normal form, which is essentially (2.42) with $c = 0$, can be shown to be radial lines emanating from the origin. In general, calculating isochrons is not analytically tractable and thus numerical methods are needed to find them. Typically, these methods start by picking points near to the limit cycle, and then by numerically integrating backwards for one period, taking advantage of the fact that isochrons are invariant under the flow generated by the system over one period. The accuracy of these methods can be improved by first finding linear approximations to the isochron near the limit cycle. This can be done, as described by Malkin in [44, 45], by solving the adjoint equation

$$\dot{Q} = -Df(x)^T Q, \quad (2.43)$$

where T denotes the transpose and Df is the Jacobian of f , such that Q is periodic (with period 1) and satisfies $\langle f, Q \rangle = 1/\Delta$ everywhere and $\langle \cdot, \cdot \rangle$ is the standard vector inner product. For a given phase θ_0 , the vector orthogonal to $Q(\theta_0)$ forms

a linear approximation to the isochron $\mathcal{I}(Q(\theta_0))$ at the limit cycle. Thus, points can be chosen near to the limit cycle along these approximation and then evolved backwards in time. More points can be added by noting that isochrons with an asymptotic phase difference of $\Delta\theta$ will map to one another under the flow over $\Delta\theta$ and $-\Delta\theta$ respectively [46]. By carefully inserting and removing points, smooth curves approximating isochrons may be found. Izhikevich has written a MATLAB code which performs such a routine [47].

The direct backward integration method for calculating isochrons suffers from two major drawbacks. Firstly, backward integration is numerically unstable, so that step size used during the integration must be very small. Secondly, and particularly for strongly attracting limit cycles, the trajectories generated by backward integration may quickly diverge to infinity, and thus extending the isochrons far from cycle becomes difficult. In 2010, Osinga and Moehlis [48] devised a continuation based algorithm to efficiently compute isochrons. This method overcomes the known problems with the instability associated with backwards integration. The basic idea is to first find the periodic orbit Γ and the local, linear isochrons segments by solving a periodic boundary value problem, which essentially amounts to solving (2.43) as before. For a given phase θ_0 , we then approximate the isochron by its linear segment, close to cycle. A new boundary value problem is then set up, having as the right end point, a point on this linear segment. The boundary value is then solved over for a time interval of Δ , where the left end point now gives a new point on the same isochron (now not a linear approximation). The solution to this boundary value problem may then be continued over a small length of the linear segment, using the distance from the limit cycle as a continuation parameter. This will thus give a curve in phase space that maps to the linear segment under the flow Φ_Δ , where we now express the time dependence of the flow operator as a subscript, and hence this curve forms part of the isochron. This process can be manipulated, taking the same segment, now solving over a window of $k\Delta$, $k = 2, 3, \dots$ to find sections of the isochrons further from the limit cycle. The concatenation of all of these continued solutions form an approxima-

tion to the isochron at θ_0 . Osinga and Moehlis' algorithm was coded and run in AUTO, a powerful numerical continuation package. The authors claim that, since AUTO measures the variation between orbit segments along the entire length of an orbit, the extreme sensitivities on initial conditions that are typical in multiple time scale problems, common in models of neurons, can be overcome. In addition, AUTO uses adaptive step sizes for its continuation, and this algorithm for finding isochrons is efficient.

Guillamon and Huguet gave, in [46], a new method to find high order approximations to isochrons for planar systems using geometric approaches. We will consider the general approach when Γ is hyperbolic, though similar results hold if this is not true. The problem essentially consists of finding a vector transverse to f , which we shall denote h , so that if the dynamics of f , expressed in isochronal coordinates (θ, r) , where r denotes the distance along the isochron, are

$$\dot{\theta} = \frac{1}{\Delta}, \quad \dot{r} = \frac{\lambda r}{\Delta}, \quad (2.44)$$

then we wish for the dynamics of h , written in the same coordinates, to be given by

$$\dot{\theta} = 0, \quad \dot{r} = 1. \quad (2.45)$$

Given these conditions, orbits of h crossing Γ will be the isochrons of f . Finding a vector field such that this is true is clearly a nontrivial problem. However, the authors show, that finding h is equivalent to solving for L in the PDE

$$\left(\frac{1}{\Delta} \partial_{\theta} + \frac{\lambda r}{\Delta} \partial_r \right) L(\theta, r) = f(L(\theta, r)), \quad (2.46)$$

where L now defines a map that will provide the isochrons of the system. As before, this method avoids any backward integration and allows one to find high order approximations to the isochrons, at least locally. However, the solutions only give the isochrons up to a finite value of r , after which the backward integration methods are needed to further extend the isochrons.

A recent paper by Mauroy and Mezić [49] presents a new method, using operator theoretic, rather than geometric approaches, to find isochrons using only

forward integration. Rather than focussing on the identification of invariant sets, operator theory is based on tracking observables, or measures on state space. For the calculation of isochrons, the Koopman operator, K , also known as the composition operator, defined by

$$K = z \circ \Phi_t(x), \quad (2.47)$$

where $z : \mathbb{R}^n \rightarrow \mathbb{R}$ is some observable of the state space and $\Phi_t(x)$ denotes the flow evolved for a time t of the vector field, starting at a point x . The Koopman operator is the left adjoint of the Frobenius-Perron operator, another important operator which encodes information about an iterated map. Mauroy and Mezić show that the level sets of the Fourier average, $z_{\omega_0}^*$, of an observable z

$$z_{\omega_0}^*(x) = \lim_{T \rightarrow \infty} \int_0^T (z \circ \Phi_t)(x) e^{-i\omega_0 t} dt, \quad (2.48)$$

where $\omega_0 = 2\pi/\Delta$ are, in fact, the isochrons of the system. Moreover, the Fourier averages are the eigenfunctions of the Koopman operator, so that

$$K z_{k\omega_0}^*(x) = e^{ik\omega t} z_{k\omega_0}^*(x), \quad k = 1, 2, \dots \quad (2.49)$$

Remarkably, this result holds for almost all observables within the set of possible functions. The only restriction is that the first Fourier coefficient of the Fourier observable evaluated along the limit cycle is nonzero over one period, i.e.

$$z^\Gamma = \frac{1}{\Delta} \int_0^\Delta (z \circ \Phi_t)(x) e^{-i\omega_0 t} dt \neq 0. \quad (2.50)$$

Usually, a linear combination of the state variables will suffice for the observable f . The limit in (2.48) may be replaced a large T , so that that Fourier averages may be computed in a finite time horizon. Then, an algorithm for calculating isochrons involves computing the Fourier averages for all points in \mathcal{B}_Γ and then extracting level sets of the resulting tensor. In fact, $z_{\omega_0}^*$ is precisely the phase function we wished to find earlier. Moreover, as this method requires only forward integration, it can be used to find isochrons in models with greater numbers of variables, which is difficult using the methods of Osinga and Guillamon. In Fig. 10, we plot the isochrons of (2.42) found using this method, superimposed

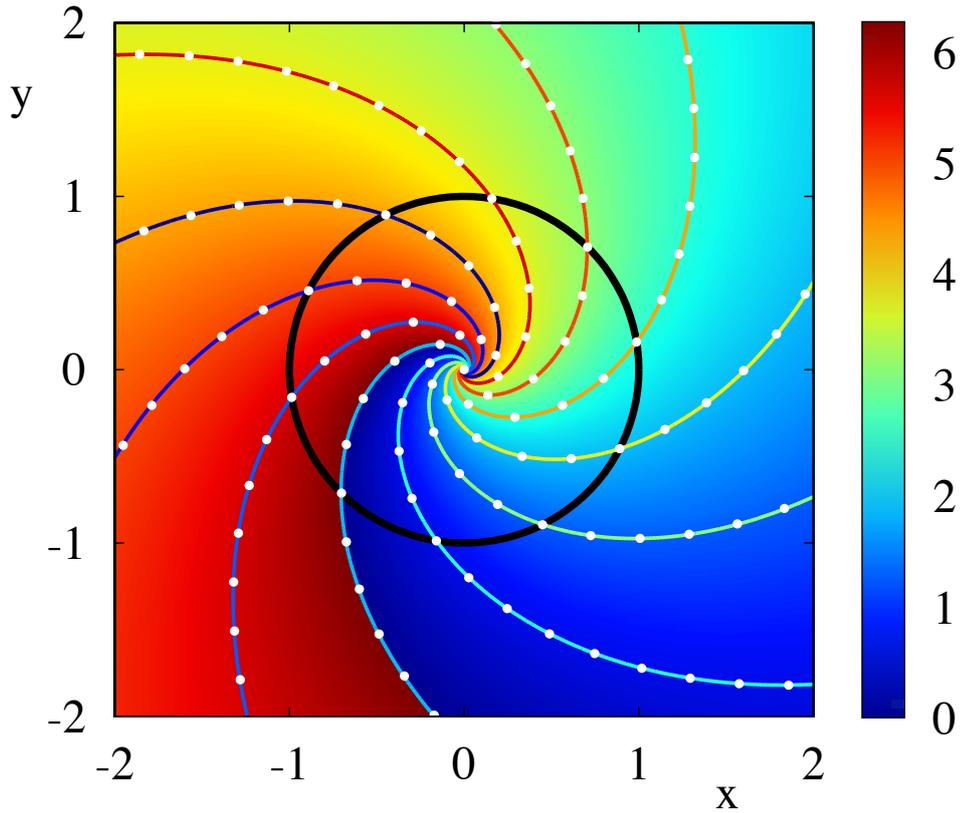


Figure 10: Isochrons founds using the method of Fourier averages (2.48) for the SL system (2.42), with $\lambda = 2.0, c = 1.0, \omega = 1.0$. The background colour represents the Fourier average, whilst the coloured lines are the isochrons, given as level sets of the Fourier average. The white dots are the actual isochrons, computed analytically.

on the computed Fourier averages. For comparison, the white dots represent the actual isochrons. Computing isochrons is time consuming, and is thus impractical for studies of large networks. However, using the linear approximations to the isochrons discussed earlier, and making certain assumptions about the system, we can use *phase reduction* techniques to reduce the number of dimensions in our model.

2.5 Reduction of oscillatory models to a phase-only description

Close to the limit cycle, we may consider the isochrons to be equal to the linear approximations we found earlier. If we assume that the perturbation to the system is weak, that is, $\varepsilon \ll 1$ in (2.38), and that perturbed trajectories relax instantly to the limit cycle, then we may rewrite (2.38) as

$$\dot{\theta} = \frac{1}{\Delta} + \varepsilon \langle Q(\theta), g(\Gamma(\theta)) \rangle, \quad (2.51)$$

using the parametrisation for Γ we considered earlier, and where Q is the solution to (2.43), using the same conditions as before. We call Q the *infinitesimal* PRC (iPRC), analogous to R in (2.37), except that it can only describe the change in phase for infinitesimal perturbations. We also have that $Q = \nabla_x \theta|_{x=\Gamma(\theta)}$, i.e. the phase response curve is equal to the change in θ given a change in the state variables, evaluated along the orbit. In fact, the phase response at a point x is given in general by $\nabla_x \theta|_x$. However, the phase function θ is, in general, not known. This is one of the reasons why there is such a desire to find efficient algorithms for computing isochrons, since these are level sets of the phase function. Referring back to the discussion on isochrons, it is clear that, at the orbit itself, the isochrons are orthogonal to the PRC. In neuroscience, we generally consider perturbations that induce additional currents through the membrane, either through injected current or signals from other cells. In this way, the perturbations affect only the membrane potential, V , which we generally take as the first component of a given neural model, so that g in (2.38) contains nontrivial entries only in the first component. Making this assumption, we need only consider the first component, $q = Q_1$ of Q , so that the system now reads

$$\dot{\theta} = \frac{1}{\Delta} + \varepsilon q(\theta) g_1(\Gamma(\theta)). \quad (2.52)$$

As for (2.37), we have reduced the dynamics of the full system (2.38) to movement with constant angular velocity on \mathbb{S}^1 modulated by perturbative effects, the strength of which is given by $\varepsilon q(\theta) g_1(\Gamma(\theta))$, so that a reduction in the number of

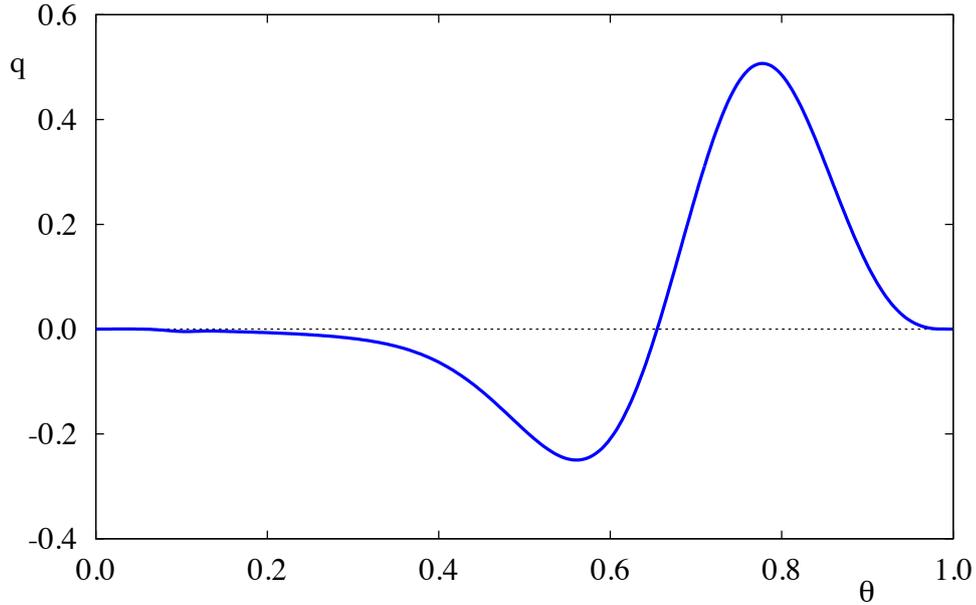


Figure 11: Infinitesimal PRC for the solution to the HH model as shown in Fig. 3. The PRC takes both positive and negative values, showing that this neuron is a type II neuron. This is to be expected as the bifurcation giving rise to oscillations in the HH model is a Hopf.

dimensions is achieved. Fig. 11 shows the first component of Q for the oscillatory solution to the HH model shown in Fig. 3. This solution was found by solving the adjoint equation, as described in (2.43). It is interesting to note that type I neurons support only positive values for q so that spikes may only be advanced in such models, whereas q can take both positive and negative values for type II neurons [50]. Fig. 11 clearly shows that this neuron is a type II neuron, as predicted from the fact that the underlying bifurcation is of Hopf type. PRCs are not only a mathematical construct. Through voltage clamp techniques, PRCs may be estimated from real neurons by artificially injecting a current through the membrane of a cell [51, 52]. This allows for a direct comparison of the PRCs of real neurons and their mathematical models. Using new technology, dynamic clamp techniques allow for the injection of arbitrary signals into a cell with real-time feedback. The advent of dynamic clamp techniques [53] is allowing for more advanced studies of

neural response to realistic stimuli. A comprehensive review of phase reduction techniques, both theoretical and experiment may be found in [6].

Using this phase reduction technique, we can analyse the effects of perturbations on the timing of spikes fired by a particular cell. The perturbations may be from some injected current, for example in experiments, or from other neurons in a network.

2.6 Noise in neural systems

Variability is ubiquitous at all levels of the central nervous system [54]. Even when factors are kept as constant as possible, trial to trial variability of neural response is observed when neurons are exposed to the same stimulus. There are a number of possible sources of this variability, and these variations have great importance when we ask questions about the reliability of a neural network to encode and transmit information reproducibly. The deterministic properties of the cells can give rise to variation in response. Prior to stimuli, neurons may be in a number of different states, in terms of the membrane potential and fraction of ion channels that are open. These differences in initial states can give rise to different emergent behaviour in response to exactly the same stimulus, particularly if the neuron is highly sensitive to its initial condition [55, 56].

Aside from these deterministic properties, stochastic effects can have a pronounced effect on neural dynamics [11]. The movement of ionic species is itself a thermodynamic process as molecules arrive at the receptor sites at different rates owing to the diffusive processes by which they travel. In addition, the receptor ligands are limited in their ability to count the number of signalling molecules arriving at the binding site [57, 58]. In fact, the equations governing the transition rate of the gating variables in the HH model may be thought of as a continuum limit of a discrete time Markovian process where the rates are the limit of the probabilities of transitioning between open and closed states [59]. At the individual neuron level, noise may result in variability of only a few milliseconds or smaller. However, since cortical neurons are able to distinguish firing events separated on

this order of magnitude, we cannot say for certain that these effects are negligible [60]. It is often difficult to tell what effect the noise has on information processing in larger networks as it is hard to distinguish a completely random process from one acting to transmit information optimally [61]. Although the magnitude of the individual effects is typically small, and averaging over the repertoire of stochastic processes tends to reduce variability as a whole, we must remember that neural networks perform highly nonlinear computations, and small effects may be amplified, for example, through signal transduction. In particular, these effects may become more pronounced when neurons are close to the firing threshold [62].

2.7 Sources of noise

As well as the previously discussed thermodynamic noise in ion transport, there exists a host of other biochemical and biophysical processes that contribute to the stochastic nature of the neuron.

2.7.1 Electrical noise

Variation in the membrane potential of a neuron can result in variability in the timing of both the initiation and propagation of spikes. By far, the dominant form of electrical noise is channel noise, that is, the aforementioned opening and closing of ion channels [63]. Stochastic studies show that the number of ion channels open at the spiking threshold is key in determining timing precision, so that, even though there exist a large number of ion channels at any given site, a much smaller number contribute to neural precision [62]. For weak signals and for smaller neurons, this can lead to increased variation in spike timing, as weaker signals are more susceptible to these effects, and smaller neurons have much higher resistance [64]. Spatial interactions along the axon down which the action potential propagates can cause memory effects lasting up to several hundred milliseconds and these effects may lead to the failure of the spike to reach the release sites [65].

Smaller sources of electrical noises include Johnson noise, generated by the thermal agitation of ions in the cytosol at equilibrium, regardless of any exter-

nal applied current and shot noise, where the finite number of ions arriving at a receptor is small enough to lead to statistical fluctuations. However, these noise sources are around three orders of magnitude smaller than channel noise in neurons [66, 67].

2.7.2 Synaptic noise

The sheer number of other cells with which each neuron communicates is enough to generate variability by itself. Neurons receive synaptic input from a large number of other cells, and this process is simply too complex to describe accurately, and so may only appear to be random [68]. Nevertheless, as with all biological systems, there are microscopic sources of true noise at each synapse that are likely to contribute to this background activity and, in turn, influence neural firing. Experiments in the absence of presynaptic currents show miniature postsynaptic currents, which provide strong evidence for synaptic noise [69].

The spontaneous opening of Ca^{2+} stores, Ca^{2+} channel noise, the spontaneous triggering of vesicle release or the spontaneous fusion of a vesicle with the cell membrane all contribute to synaptic noise [70]. Once the vesicle release pathway has been triggered, a postsynaptic current will be induced, which themselves display trial to trial variability in terms of their amplitude. The amplitude of this response is affected by such factors as the number of vesicles released from a presynaptic bouton, their size, diffusivity, the duration over which they are released and the spatial distribution of the receptors on the postsynaptic cell. In addition, synaptic receptor channel noise will also contribute to the synaptic noise [65].

2.8 Management of noise

Neural reliability refers to the ability of the neuron to reproduce a firing pattern in response to an identical stimulus [71]. In terms of information processing, it is clear that reliability is beneficial to neural systems for passing on information about stimuli. Noise acts to cause jitter in spike timing, which may propagate and amplify downstream. In extreme cases, the propagation of spike time jitter

can cause responses across trials to bear no resemblance to one another. Thus, the central nervous system must act to counter the harmful effects of noise. The management of noise, and the reduction of its negative effects imparts a high metabolic and structural cost. For example, a fly's photoreceptors account for around 10% of its resting metabolic consumption, whilst its eye's optics make over 20% of the fly's flight payload [72].

The principle of averaging can be applied whenever redundant information is present across sensory inputs. Averaging can act to peter out the effect of noise if multiple units carry the same information and are affected by independent noise sources [73]. Divergence, whereby one neuron synapses onto many others can support averaging; when signals are sent over long distances, it may be beneficial to send the same signal redundantly over many noisy axons and reconstruct the signal at the terminal axon rather than send the same signal along one large axon [74]. Prior knowledge of the structure signals or noise can be used to minimise the impact of noise as it can be used to distinguish signals and thus compensate for any noise [75]. Often, averaging is used in conjunction with prior knowledge about the statistical properties of signals to reduce variability of response and improve signal detection and transmission.

2.9 Benefits of noise

It has long been thought that noise is exclusively detrimental to brain function. This intuitive consideration is indeed backed up by experimental evidence. However, recent studies have identified situations in which noise may actually contribute to neural reliability and may act to reduced variation in neural response, referred to by some as stochastic facilitation [76]. Stochastic resonance, and the related phenomenon of coherence resonance are active interest areas exemplifying this. In stochastic resonance, a system which is quiescent but near threshold can gain an increased sensitivity to weak stimuli due to noise [77]. In a deterministic setting, the stimulus is not strong enough to induce neuronal firing by itself, but the combined effect of the signal and noisy fluctuations drive the neuron to fire

a spike. This mechanism allows for the detection and processing of weaker signals in the noisy environment. The related phenomenon of coherence resonance occurs when noise in a system makes the oscillatory response most coherent with the time dependent stimulus, as can be measured through correlations between the two or by examining the power spectrum of the neuronal response [78]. In this way, stochastic effects can, counter-intuitively, increase the reliability of spike timing, as well as aiding in the detection of weak signals.

As noise increases the likelihood of subthreshold to cause firing events, it may help to smooth out the nonlinearities associated with the threshold. As outputs are averaged over time, the result is an effectively smoothed nonlinearity. This nonlinearity helps facilitate spike initiation and can improve the behaviour of the network as a whole [54]. Finally, neural networks that have formed in noisy environments will be more robust, and will have explored more of the potential state space of the network, allowing for more dynamic behaviour, such as adaptation [79].

2.10 Modelling noise

We now consider the mathematical descriptions of stochastic effects. Beginning from first principles, we may model certain processes, such as the opening and closing of ion channels as Markov chains, and write down the state transition matrix between different states, and then go on to describe the effect on the neuron from these channel events [80]. From here, we may then write down the master equation associated with such a system and use, for example, the Gillespie algorithm [81] to find a typical solution to the problem, or Markov Chain Monte Carlo methods to sample the probability distribution of the system [82].

Considering the size and complexity of both single neuron models and networks, these descriptions tend to be too costly to solve in practice, even taking advantage of methods such as τ -leaping [83], which speeds up the Gillespie algorithm. Instead, we assume that the effects of stochasticity, rather than simply being averaged out completely, as in the deterministic models studied thus far,

give rise to additional terms in the model equations than describe the combined overall effect of the individual processes. This gives rise to a stochastic differential equation (SDE), which is computationally cheaper to solve. We now review some of the key components of modelling these processes, in a general setting, but introducing them into neuronal models where appropriate.

2.11 Random processes and sample paths

Since we typically consider random fluctuations to primarily affect the membrane potential, we may choose to reflect the effect of noise by considering this potential to be a random variable (RV), that is, one whose value is subject to chance. The repeated samplings of a random variable X over time gives rise to the notion of a random process $\{X_t\}$, where the subscript denotes the time dependence of the sampling. As an example of a continuous time random process, consider the movement of ionic species in a well mixed environment. The ions will move around by diffusive processes, but since the environment is well mixed, there is no directed movement down a concentration gradient. Instead, each ion will follow what manifests as an unbiased random walk. We refer to this as Brownian motion.

Each sampling of a random process, now across time, will give rise to a outcomes at each step that are likely to differ from other samplings. Thus, there is no true sampling of a random process across time. As such, we introduce the notion of a sample path, $X_\omega(t) : t \mapsto X(t)$, with some fixed $\omega \in \Omega$ which is one sampling of the random process at each time step.

2.12 The Wiener process

One of the key random processes in the construction of SDEs is that of the Wiener process. This might be more commonly thought of, outside the mathematical and physical sciences, as describing Brownian motion, considered above. The standard one dimensional Wiener process, $\{W_t\}$, where $t > 0$ are non-negative real numbers, has the following properties

1. $W_0 = 0$,
2. The function $t \mapsto W_t$ is almost surely continuous ,
3. The process $\{W_t\}$ has stationary, independent increments ,
4. The increment $W_{t-s} - W_s$ has the distribution $N(0,t)$,

where $N(0, t)$ is a normal distribution with zero mean and variance t . A standard d dimensional Wiener process can be obtained as a vector of one dimensional Wiener processes

$$W_t = (W_t^1, W_t^2, \dots, W_t^d). \quad (2.53)$$

The Wiener process has zero mean $\forall t$, however, it is clear to see the variance of the process scales linearly with time. As $t \rightarrow \infty$, the variance of the process also tends to infinity, leading to a high irregularity of different sample paths from the same Wiener process. In Fig. 12, we show five sample paths from the standard Wiener process in one dimension, which typify the behaviour of the Wiener process, with zero mean, but with variance increasing linearly with t .

2.13 Langevin equations

We now amend the deterministic system (2.38) with a random component to capture the stochastic effects in a noisy environment, and arrive at

$$dX = f(X)dt + \sum_{j=1}^d g_j(X)dW_j(t). \quad (2.54)$$

In the above, $dX = f(X)dt$ represents the underlying deterministic system, which may be the previously studied HH model, for example. This is referred to as the *drift* term. The remaining terms describe the role of the noise on the system. These terms are referred to as the *diffusion* terms. The dW_j represent increments of a d dimensional Wiener process, as described in Sect. 2.12, whilst the g_j describe the dependence of the noise on the state variables. In a neural environment, we might consider some of the background noise to be independent of the gating and

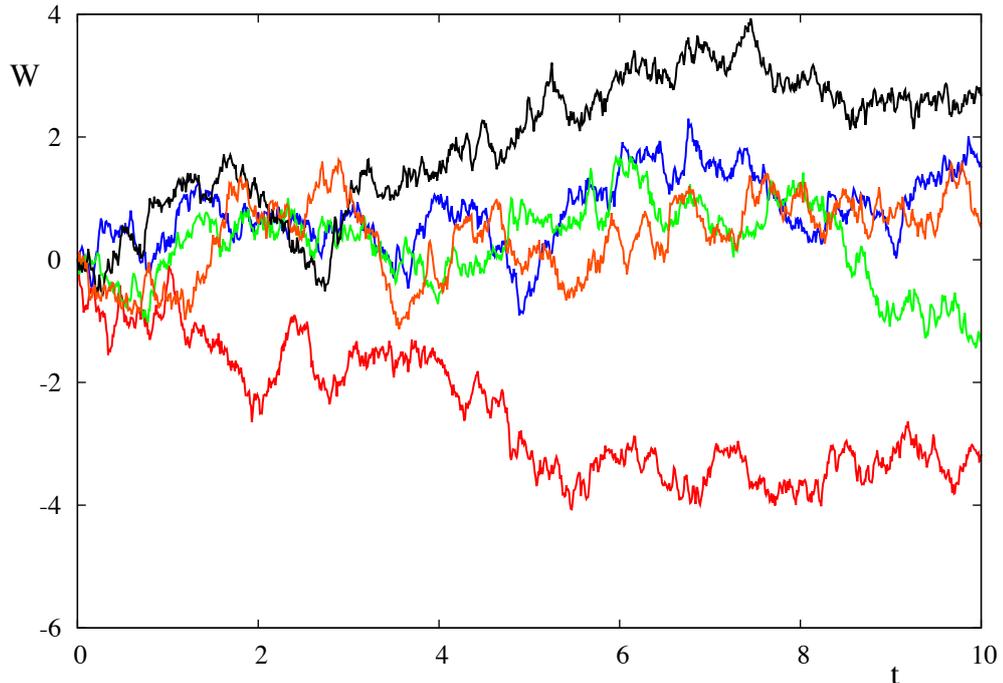


Figure 12: Five sample paths of the Wiener process. In each case, we observe different trajectories for $\{W_t\}$. It is clear to see that the trajectories are not converging to one another, and that the variance of $\{W_t\}$ is increasing with t .

membrane potential, in which case g_j would simply be constant vectors. In this case, the noise term would act solely on the membrane potential, and we would only have diffusion terms in the equation for the membrane potential. We might also consider channel noise, which might affect and be affected by the proportion of open channels, in which case the g_j would depend on the gating variables and would create diffusion terms in the equations for the gating variables. We refer to the first case as additive noise, since the noise simply acts to add on a noisy component to the deterministic process, whilst the second is referred to as multiplicative noise, since the noise effects are multiplied by some function of the state variables.

Equations of the type (2.54) are known as Langevin equations. Langevin equations can be formed from the master equation associated with a stochastic system, and have the advantage of being easier to work with, since the effect of the stochastic processes now simply appears as a product of a Wiener process and

some function of the state variables. In rare cases, an explicit solution to an SDE can be found, and here the advantage over the master equation is evident. In most cases, one must resort to numerical techniques, in which case, the computational cost is still smaller for integrating SDEs than for simulating the master equation. It is common to see (2.54) in the form

$$\frac{dX}{dt} = f(X) + \sum_j^d g_j(X)\xi_j(t), \quad (2.55)$$

where ξ_j are random processes such that $\langle \xi_j \rangle = 0$ and $\langle \xi(t)\xi(s) \rangle = \delta(t-s)$, where $\delta(t)$ is the Dirac δ function, $\langle \cdot \rangle$ denotes taking the ensemble average and $\langle xy \rangle$ is the correlation between x and y . We refer to ξ_j as a white noise process. To get from (2.54) to (2.55), one simply has to divide through by dt , so that the two are equivalent. However, dividing $dW_j(t)$ by dt involves taking the derivative of the Wiener process. Mathematicians dissent from using this equation since the derivative of the Wiener process is nowhere defined, and so this form of the Langevin equation is more commonly used by physicists.

If we rewrite (2.54) as

$$dX(t) = f(X(t))dt + g(X(t))dW(t), \quad (2.56)$$

where $d = 1$ and we drop the j subscript, we may solve (2.54) by integrating both sides with respect to t to give

$$X(t) - X(t_0) = \int_{t_0}^t f(X(s))ds + \int_{t_0}^t g(X(s))dW(s). \quad (2.57)$$

The first term on the right hand side, is easily recognised as the solution to the deterministic part of the system, whilst the second term is a stochastic integral with respect to a sample function $W(t)$.

2.14 Convergence

When considering solutions of SDEs, it is important to remark on some properties of convergence. There are a number of different notions of convergence that are relevant to the analysis of SDEs.

Mean-square convergence

This type of convergence deals with the mean-square deviation of a random process. We say that $\{X_n\}$ converges to X in the mean-square limit if

$$\lim_{n \rightarrow \infty} \langle (X_n - X)^2 \rangle = 0, \quad (2.58)$$

In such a case, we write $\text{ms-lim}_{n \rightarrow \infty} \{X_n\} = X$.

Limit in probability

We may also consider the possibility that $\{X_n\}$ approaches X because the probability of deviation from X approaches zero. This means that if, for any $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} P(|X_n - X| > \varepsilon) = 0, \quad (2.59)$$

then the limit in probability of X_n is X .

Limit in distribution

Finally, we may consider convergence in terms of probability distribution. This kind of convergence occurs if, for any continuous bounded function $a(x)$,

$$\lim_{n \rightarrow \infty} \langle a(X_n) \rangle = \langle a(X) \rangle, \quad (2.60)$$

so that the probability density of X_n approaches that of X .

Orders of convergence

Regarding the above differing notions of convergence of random processes, it may be shown that

$$\begin{aligned} \text{Convergence in mean-square} &\implies \text{convergence in probability.} \\ \text{Convergence in probability} &\implies \text{convergence in distribution.} \end{aligned} \quad (2.61)$$

Thus, we may think of convergence in mean-square as being *stronger* than convergence in probability and distribution, whilst convergence in distribution is *weaker* than convergence in mean-square and probability.

2.15 Stochastic integrals

In order to complete the details of the solution (2.57), we need to consider how to evaluate the stochastic integral. Suppose that $a(X, t)$ is an arbitrary function of time, $W(t)$ is the Wiener process and we wish to evaluate the integral $\int_{t_0}^t a(X, s)dW(s)$. In a similar way to numerical integration of ODEs, we divide the interval $[t_0, t]$ into n subintervals by a partition

$$t_0 \leq t_1 \leq t_2 \leq \cdots \leq t_{n-1} \leq t_n = t, \quad (2.62)$$

and define intermediate points such that

$$t_{i-1} \leq \tau_i \leq t_i, \quad (2.63)$$

The stochastic integral is then defined as the limit of partial sums as $n \rightarrow \infty$,

$$S_n = \sum_{i=1}^n a(X(\tau_i), \tau_i) [W(t_i) - W(t_{i-1})]. \quad (2.64)$$

It is plain to see that, in general, the integral written in this way depends on the particular choice of the intermediate points $\tau_i = \alpha t_i + (1 - \alpha)t_{i-1}$, where $\alpha \in [0, 1]$, and thus, there are an infinite number of possible interpretations of the stochastic integral, of which two have become the most popular.

Itô integral

The first interpretation sets $\alpha = 0$, so that $\tau_i = t_{i-1}$. This defines the *Itô stochastic integral* of $a(X, t)$ as

$$\int_{t_0}^t a(X(s), s)dW(s) = \text{ms-lim}_{n \rightarrow \infty} \left\{ \sum_{i=1}^n a(X(t_{i-1}), t_{i-1}) [W(t_i) - W(t_{i-1})] \right\}, \quad (2.65)$$

Two important consequences of the Itô definition of a stochastic integral are that $dW(t)^2 = dt$ and $dW(t)^{2+n} = 0$ for $n \in \mathbb{N}$.

Stratonovich integral

A different interpretation sets $\alpha = 1/2$, so that $\tau_i = 1/2(t_{i-1} + t_i)$. In this case, we define the *Stratonovich stochastic integral* to be

$$\int_{t_0}^t a(X(s), s) \circ dW(s) = \text{ms-lim}_{n \rightarrow \infty} \sum_{i=1}^n a\left(\frac{1}{2}(X(t_i) + X(t_{i-1})), t_{i-1}\right) [W(t_i) - W(t_{i-1})]. \quad (2.66)$$

It is worth noting, that although a choice of τ_i as indicated above leads to (2.66), this can also be reached by averaging X across the time points t_i and t_{i-1} . In fact, it is only the dependence of X on t that is averaged in this way, rather than the explicit dependence of a on t . Furthermore, if $a(X, t)$ is differentiable in t , then the integral may be shown to be independent of the particular choice of $t \in [t_{i-1}, t_i]$.

2.16 Comparison of the Itô and Stratonovich integrals

To compare the effect of the difference in choice of interpretation of the stochastic integral, consider the case where $a(X, t) = W(t)$. In the Itô interpretation,

$$\int_{t_0}^t W(s) dW(s) = \frac{1}{2} [W(t)^2 - W(t_0)^2 - (t - t_0)], \quad (2.67)$$

whereas, in the Stratonovich interpretation

$$\int_{t_0}^t W(s) \circ dW(s) = \frac{1}{2} [W(t)^2 - W(t_0)^2]. \quad (2.68)$$

The difference between the two can be accounted for by the fact that, in the Itô sense, $dW(t)^2 = dt$ and so terms of second order in $dW(t)$ do not vanish on taking the limit. We note that the Stratonovich integral is precisely what we expect to obtain under the normal rules of calculus, ignoring the stochastic nature of the integral, and the extra term $(t - t_0)$ appears only in the Itô sense. In fact, using the Itô interpretation of integrals requires the using a different kind of calculus, surrounding Itô's formula. In the context of mathematical biology, the Itô interpretation may seem the most logical choice, since biological processes are non-anticipatory. That is, they depend only of values in the past or present, and not in the future. However, the Stratonovich interpretation may be more suited

if the assumption is that the noise is actually part of some other, non-observed biologically relevant process [84]. In practice, both interpretations are used, and there exist methods of moving between the two interpretations freely, as we shall see shortly. One must also take care, however, to set out clearly the assumptions made when choosing one or the other, and to remain consistent once this choice has been made.

2.17 Itô's formula

Itô's formula is essentially the chain rule for Itô calculus. For a function $h(x(t))$, where $x \in \mathbb{R}$, it tells us what stochastic differential equation it obeys. If

$$dx(t) = f(x(t))dt + g(x(t))dW(t), \quad (2.69)$$

then, upon expanding $h(x(t))$ to second order in $dW(t)$, we get

$$\begin{aligned} dh(x(t)) &= h(x(t) + dx(t)) - h(x(t)), \\ &= h'(x(t))dx(t) + \frac{1}{2}h''(x(t))dx(t)^2 + \dots, \\ &= h'(x(t))\{f(x(t))dt + g(x(t))dW(t)\} + \frac{1}{2}h''(x(t))g(x(t))^2dW(t)^2, \\ &= \left\{ f(x(t))h'(x(t)) + \frac{1}{2}g(x(t))^2h''(x(t)) \right\}dt + g(x(t))h'(x(t))dW(t), \end{aligned} \quad (2.70)$$

where in the last step we have used the fact that $dW(t)^2 = dt$ and in the penultimate step, we discard terms of higher order than $dW(t)^2$.

In the Stratonovich case, we observe the same chain rule as we would expect from ordinary calculus

$$dh(x(t)) = h'(x(t))\{f(x(t))dt + g(x(t))dW(t)\}. \quad (2.71)$$

In the case where $X \in \mathbb{R}^n$, we obtain the following multivariable version of Itô's formula

$$\begin{aligned} dh(X(t)) &= \left\{ \sum_i f_i(X(t))\partial_i h(X(t)) + \frac{1}{2} \sum_{i,j} [g(X(t))g^T(X(t))]_{ij} \partial_i \partial_j h(X(t)) \right\}dt \\ &\quad + \sum_{i,j} g_{ij}(X(t))\partial_i h(X(t))dW_j(t), \end{aligned} \quad (2.72)$$

For each of the respective cases, we may write an equivalent SDE in the other representation, taking advantage of the above formula. Writing the Itô SDE as

$$dX = f^I(X(t))dt + g^I(X(t))dW(t), \quad (2.73)$$

and the Stratonovich SDE as

$$dX = f^S(X(t))dt + g^S(X(t)) \circ dW(t), \quad (2.74)$$

we may derive the result connecting the solutions to the Itô and Stratonovich SDEs.

$$\int_{t_0}^t g^S(X(s)) \circ dW(s) = \int_{t_0}^t g^I(X(s))dW(s) + \frac{1}{2} \int_{t_0}^t g^I(X(s)) \partial_X g^I(X(s)) ds. \quad (2.75)$$

Using this integral, we may now write down the equivalent Stratonovich SDE corresponding to (2.73) by replacing

$$f_i^I \rightarrow f_i^S - \frac{1}{2} \sum_{jk} g_{jk}^S \partial_k g_{ij}^S \quad (2.76)$$

Similarly, the Itô SDE corresponding to (2.74) by replacing

$$f_i^S \rightarrow f_i^I + \frac{1}{2} \sum_{jk} g_{jk}^I \partial_k g_{ij}^I \quad (2.77)$$

This gives us a nice way of converting between equivalent representations of the same SDE. This can prove useful when analysing the probability distributions of the solution to SDEs, or when simulating SDEs numerically.

2.18 Coloured noise

Up to this point, we have only considered the case in which the noise sources are white. That is, where the sources of noise have no temporal correlations (other than at zero lag), so that the autocorrelation $R_x(\tau) = \langle x(t)x(t+\tau) \rangle$ of a white noise process is

$$R_x(\tau) = \delta(\tau), \quad (2.78)$$

where δ is the Dirac δ -function. White noise represents all temporal frequencies equally. This is represented by white noise processes having a flat power spectral

density (PSD), where the PSD describes how the power of a signal is distributed across different frequencies, and is given by the Fourier transform of $R_x(\tau)$. Noise processes in the real world are likely to have finite temporal correlations, and so a white noise description may not be the most appropriate for modelling biological processes and, as a neural example, it has been shown that input correlations in noise sources may impact interspike intervals [85]. We refer to any noise process that is not white as being a coloured noise process.

The Ornstein-Uhlenbeck (OU) process is a simple, mean-reverting noise process that has finite temporal correlations at non-zero lag [86]. The one dimensional OU process is written as

$$d\eta = -\gamma\eta dt + \sigma dW(t), \quad (2.79)$$

where $W(t)$ is the Wiener process as before. We refer to the OU process as being mean reverting, as sample paths will always tend to the mean $\langle\eta\rangle = 0$. The OU process can be solved analytically as

$$\eta(t) = \eta(t_0)e^{-\gamma t} + \sigma \int_{t_0}^t e^{-\eta(t-s)} dW(s). \quad (2.80)$$

If the initial condition is deterministic, or comes from a Gaussian distribution, then $\eta(t)$ is Gaussian. We may go on to calculate the temporal correlation function as

$$\langle\eta(t), \eta(s)\rangle = \left[\text{var}(\eta(t_0)) - \frac{\sigma^2}{2\gamma} \right] e^{-\gamma(t+s)} + \frac{\sigma^2}{2\gamma} e^{-\gamma|t-s|}, \quad (2.81)$$

so that it is clear that the correlation is nonzero at nonzero lag. We note that, if $\gamma > 0$, as $t, s \rightarrow \infty$, with finite $|t - s|$, the correlation function becomes stationary. In fact, if we set the initial time at $-\infty$, rather than 0, the solution (2.80) becomes

$$\eta(t) = \sigma \int_{-\infty}^t e^{-\gamma(t-s)} dW(s), \quad (2.82)$$

in which the correlation function and mean assume their stationary values. Since the process is Gaussian, it is stationary. At the expense of increasing the number of state variables in our system by one, we may go on to use the OU process as a noise source in our models to have noise terms with some temporal structure. We

amend our SDE (2.54) to include the OU process as

$$dX = [f(X(t)) + g(X(t))\eta] dt, \quad (2.83)$$

$$d\eta = -\gamma\eta dt + \sigma dW(t), \quad (2.84)$$

where we can tune $\gamma, \sigma > 0$ to best reflect the noise source we wish to model.

In Fig. 13, we plot the histogram of 10,000 sample paths of an OU process, with $\gamma = \sigma = 0.1$, as it evolves in time. The initial condition at $t = 0$ is the standard normal distribution $N(0,1)$, and so the resulting distribution of η will also be Gaussian $\forall t$, and is given by (2.80). The blue surface superimposed on the histogram is the solution to (2.80). As we can see, the histograms approximating the probability density match very well with the analytical solution.

2.19 Simulating stochastic processes

In almost all cases, we cannot find an analytical solution to SDEs and so numerical schemes to efficiently and accurately simulate sample paths of SDEs is an active area of research. In all numerical schemes, sample paths are generated by representing Wiener increments as samples from normal distributions. Clearly, each solution of the same SDE will only produce one sample path, and this may deviate from the ‘true’ solution. To circumvent this problem, many runs of the same SDE using different samples for the Wiener increments will produce different sample paths. Provided that the numerical scheme and SDE converge, in some notion described earlier, then ensemble averages of many sample paths should result in a path representing the mean path of the system. Monte Carlo methods may be used to approximate the probability distribution of the state variables of the system by generating histograms of many sample paths. As for numerical algorithms for ODEs, different routines have different rates of convergence of solutions to the exact solutions, and different regions of stability, that is where errors in the solution are attenuated as the numerical routine continues. Explicit or implicit schemes may be used, where explicit schemes are generally faster, since implicit schemes involving root finding at each step, but implicit schemes tend to have

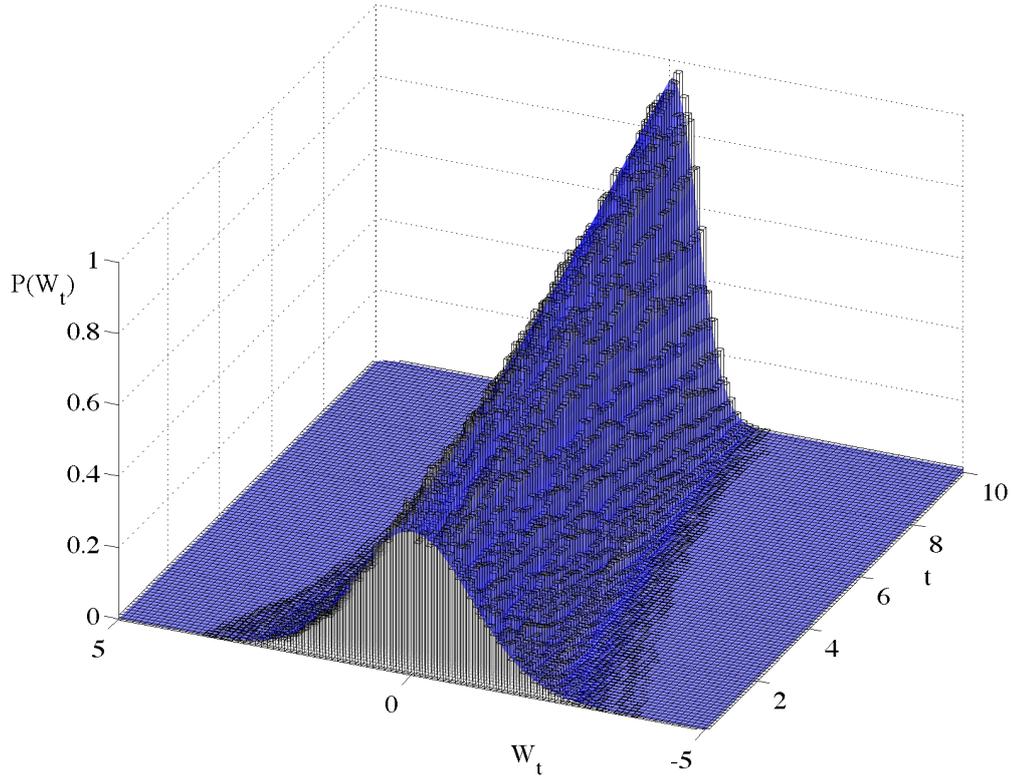


Figure 13: Approximate and exact solutions for the probability density function of (2.79) with $\gamma = \sigma = 0.1$. The blocks show the approximate density found by discretising time into $N = 1000$ bins, and plotting the histogram of η_n for each time bin for $n = 1, \dots, N$. The blue surface superimposed is the solution to (2.80). As we can see, the approximate density from the histograms matches very well with the analytical solution.

larger regions of stability. Finally, the numerical routines may differ dependent on whether the SDE to be simulated is of the Itô or Stratonovich type. For additive noise cases, there will be no difference between the two, but for multiplicative noise cases, they will differ. Of course, we also have the freedom to write an Itô equation in its equivalent Stratonovich form and then use a Stratonovich algorithm, or vice versa.

Convergence of numerical schemes

Numerical routines have different rates of convergence, corresponding to what notion of convergence is considered. When discussing numerical routines, the order of strong convergence corresponds to convergence in mean-square, whilst the order of weak convergence corresponds to convergence in distribution. Intuitively, strong convergence considers the deviation of solutions found using the numerical routine and the exact solution, whilst weak convergence considers the approach of the average of a smooth function of the variable to its exact value. Weak convergence is typically more rapid than strong convergence. In general, the weak and strong orders of convergence for a given scheme are not the same. Moreover, some schemes will give a better estimate of actual paths of a SDE, but may not necessarily improve estimates of averages. If we are interested only in averages and not sample paths, we may speed up the simulation of the SDE by approximating increments in the Wiener process by matching moments to a normal distribution. Thus, the choice of algorithm depends not only on the SDE in question, but also on what quantities we are interested in analysing.

Multiple integrals

When considering the case in which there are multiple noise sources in a model, we must take care to represent the combined effect of the noise terms correctly. Generally speaking, for an SDE, written in integral form as

$$X(t) = X(t_0) + \int_{t_0}^t f(X(s))dt + \int_{t_0}^t g(X(s))dW(s), \quad (2.85)$$

where $X \in \mathbb{R}^n$ and $W(t) \in \mathbb{R}^d$, for higher orders of expansion of functions, the solution to the SDE contains terms involving the general Itô stochastic multiple integral

$$I_{i_1, i_2, \dots, i_n}(t, t_0) = \int_{t_0}^t dW_{i_1}(s_1) \int_{t_0}^{s_1} dW_{i_2}(s_2) \dots \int_{t_0}^{s_{n-1}} dW_{i_n}(s_n). \quad (2.86)$$

There also exist similarly defined Stratonovich multiple integrals which are more convenient for the development of higher order algorithms. In general, these inte-

grals cannot be expressed in terms of Wiener increments. However, in the case that the noise is commutative, then this is not true.

Commutative noise

Noise is said to be commutative if the following condition is satisfied

$$\sum_{k=1}^d g_{k,j'} \frac{\partial g_{i,j}}{\partial X_k} = \sum_{k=1}^d g_{k,j} \frac{\partial g_{i,j'}}{\partial X_k}, \quad (2.87)$$

for $j \neq j'$. This condition is satisfied automatically in the following cases

1. When the noise sources are additive.
2. Where there is only one noise term, so that $j = j' = 1$.
3. When g can be written in the form

$$g_{i,j}(X) = X_i Z_{i,j}, \quad (2.88)$$

for some arbitrary constant matrix $Z_{i,j}$.

4. When $d = m$ and $g_{i,j} \propto \delta_{i,j}$, where $\delta_{i,j}$ is the Kronecker delta. This case is called *diagonal noise* by Kloeden and Platen [87].

In all of the cases we consider in this thesis, there will only be one noise source. We now go onto consider the most basic numerical algorithm for simulating SDEs, which is a simple extension of the one-step Euler algorithm for ODEs.

Euler-Maruyama scheme

Suppose we wish to simulate (2.54) over the interval $[0, T]$. We divide this interval into N subintervals of size $h = T/N$ at points $\tau_n = nh$ so that the function $X(t)$ is to be evaluated at the points

$$\tau_0, \tau_1, \dots, \tau_{N-1}, \tau_N, \quad (2.89)$$

and the corresponding Wiener increments are

$$\Delta W_n = W(\tau_{n+1}) - W(\tau_n) \sim \sqrt{h} N_d(0, 1), \quad (2.90)$$

where N_d is the d -dimensional normal distribution. Let us denote the solution of the algorithm by y_n . The Euler-Maruyama scheme gives the y_n as

$$y_{n+1} = y_n + f(y_n)h + \sum_{j=1}^d g_j(y_n)\Delta W_n, \quad (2.91)$$

It may be shown that this scheme has weak order of convergence of h , but a strong order of convergence of only $h^{1/2}$ due to the variance of the Wiener increments. A nice implementation of (2.91) is put forward in [88]. Although the scheme is efficient, the low order of convergence means that small h must be used to get good approximate solutions [87].

Other schemes

The simulation of numerical routines for approximating solutions to SDEs is an active area of research, and as such, there are now a whole suite of algorithms for this purpose. The Milstein algorithm uses one more term in the expansion in h of $x(t)$ and has both a weak and strong order of convergence of h . Higher order stochastic Runge–Kutta (SRK) methods have also been developed, but clearly have lower orders of convergence than their ODE counterparts. The general form for a s -stage SRK where $d = 1$ provides the y_n as

$$y_{n+1} = y_n + \sum_{i=1}^s a_i f(Y_i)h + \sum_{i=1}^s (b_i^1 \Delta W_n + b_i^2 J_{10}) g(Y_i), \quad (2.92)$$

$$Y_i = y_n + \sum_{j=1}^{i-1} \alpha_{ij} f(Y_j)h + \sum_{j=1}^{i-1} (\beta_{ij}^1 \Delta W_n + \beta_{ij}^2 J_{10}) g(Y_j), \quad i = 1, \dots, s, \quad (2.93)$$

where α and β are $s \times s$ constant matrices of coefficients and a and b are constant row vectors of coefficients $\in \mathbb{R}^s$. It had been thought that a strong order of convergence of $h^{3/2}$ was as good as could be obtained for such schemes. However, Burrage and Burrage [89] found a 4-stage SRK with a strong order of h^2 by using higher order multiple stochastic integrals. Other routines extend linear multistep methods, such as the Adams-Bashforth or Adams-Moulton methods, as discussed in [90, 91]. Two-step schemes following this philosophy give y_n as

$$\sum_{j=0}^2 \alpha_{n-j} y_{n-j} = h \sum_{j=0}^2 \beta_{2-j} f(y_{n-j}) + \sum_{j=1}^2 \sum_{r=1}^d g_r(y_{n-j}) \Delta W_{r,n-j}, \quad (2.94)$$

where α, β and γ are vectors with constant coefficients. Where $\beta_2=0$, these schemes are explicit, whereas where $\beta_2 \neq 0$, the schemes are implicit. Clearly, for these schemes, we need to start from $n = 2$, and more importantly, we need to specify y_1 as well as y_0 . In general, we can take y_1 from data, or we may simply use a one-step Euler-Maruyama scheme to generate y_1 from y_0 and use this as initial data for the routine.

Weak convergence

If we are only interested in averages (or some other functional) of the solution, rather than the sample paths themselves, we may speed up algorithm through the use of moment matching. Instead of using (2.90), we may simply use $\Delta W_{r,n} = \sqrt{h}\xi_{r,n}$, where the $\xi_{r,n}$ are i.i.d as follows

$$P(\xi_{r,n} = 0) = \frac{2}{3}, \quad P(\xi_{r,n} = \pm\sqrt{3}) = \frac{1}{6}. \quad (2.95)$$

Thus, instead of having to sample from a normal distribution, we are now only sampling from a uniform distribution, which provides a great speed increase to the algorithm.

Weak noise

Consider the SDE $dx = f(x(t))dt + \sigma \sum_{r=1}^d g_r(x)dW(t)$. In cases in which the contribution of the noise is small, that is, $\sigma \ll 1$, we can use other algorithms to provide efficient schemes with higher accuracy. For example, the standard fourth order Runge–Kutta scheme appended with the term

$$\sigma \sum_{r=1}^d g_r(y_n)dW_{r,n}, \quad (2.96)$$

has a strong order of convergence of $h^4 + \sigma h + \sigma^2 h^{1/2}$, which will allow for larger stepsizes to be taken in the routine to provide the same degree of accuracy, and thus will provide a speed up to the algorithm as a whole [92].

2.20 The Fokker–Planck equation

2.20.1 The forward Fokker–Planck equation

The Fokker–Planck equation, which can be constructed using terms from (2.54), can be used to derive many quantities of interest for a stochastic dynamical system [93, 80]. In particular, the forward Fokker–Planck equation tells us about the evolution of the probability density of the location of sample paths, given some past distribution. The large time solution for this problem will then give the steady state distribution of the system. From here, the ensemble average of quantities of interest can be computed by taking the integral of the product of the quantity with the steady state distribution over the entire domain. The forward Fokker–Planck equation (associated with the Itô interpretation) is written as

$$\frac{\partial P(X, t|Y, t')}{\partial t} = - \sum_i \frac{\partial}{\partial X_i} [f_i(X)P(X, t|Y, t')] + \frac{1}{2} \sum_{i,j} \frac{\partial^2}{\partial X_i \partial X_j} [G_{ij}(X)P(X, t|Y, t')], \quad (2.97)$$

where

$$G = gg^T, \quad (2.98)$$

is the outer product of g with itself. The Fokker–Planck equation may be seen to comprise of a drift term, given by the vector f , and a diffusion term, given by the matrix G . Since (2.97) is a PDE obeying the standard rules of calculus, we no longer need be concerned about the correct interpretation of integrals in its solution. However, except in simple cases, we do not expect analytical solutions to (2.97) to exist, and so we have exchanged the difficulties in solving SDEs for those involved in the solution of PDEs. In cases where X exists in a high dimensional space, it is often easier to use numerical techniques on sample paths of the Langevin equation to approximate the probability distributions. Notationally, we will drop the conditional dependence of P on (Y, t') in the remainder of this thesis.

In order for (2.97) to be well-posed, we need to define its associated initial and boundary conditions. Suppose that we have $X \in S \subset \mathbb{R}^n$, with boundary ∂S . In many cases, $S = \mathbb{R}^n$. However, we clearly cannot numerically solve on this domain so, in this case S merely approximates our domain. We require the normalisation

condition

$$\int_S P(X, t) dX = 1 \quad \forall t. \quad (2.99)$$

Thus, our initial condition must obviously be a probability distribution, and so we set

$$P(X, 0) = R(X), \quad \text{where} \quad \int_S R(X) dX = 1. \quad (2.100)$$

The probability current, $J \in \mathbb{R}^n$, associated with (2.97) is defined to be

$$J_i = -f_i P + \frac{1}{2} \sum_j G_{ij} \frac{\partial P}{\partial X_j}. \quad (2.101)$$

To ensure that our normalisation condition is always met, we impose a *reflecting* boundary condition on ∂S , so that

$$n \cdot J = 0, \quad X \in \partial S, \quad (2.102)$$

where n is the outward facing normal of ∂S at X . If we choose our domain S large enough, and we know that P and its spatial derivatives vanish at the boundary, it may be sufficient to impose the *absorbing* boundary condition, which is the Dirichlet condition

$$P(X) = 0, \quad X \in \partial S. \quad (2.103)$$

If we are only interested in the steady state distribution, we may set $\partial P / \partial t = 0$ in (2.97) and solve the resulting homogeneous PDE. There are difficulties associated with this, as this involves solving a homogeneous PDE with zero flux boundary conditions, for which solutions are constant over the whole domain. We will address the different ways to overcome this problem when we use the Fokker–Planck equation for this purpose later in this thesis.

2.20.2 The backward Fokker–Planck equation

Instead of propagating information forwards, we may instead work backwards and ask about the conditional dependence of probability distributions on later times. This gives rise to the backward Fokker–Planck equation (associated with the Itô

interpretation) [93, 80]

$$\frac{\partial P(X, t|Y, t')}{\partial t} = \sum_i f_i(Y) \frac{\partial}{\partial Y_i} [P(X, t|Y, t')] + \frac{1}{2} \sum_{i,j} G_{ij}(Y) \frac{\partial^2}{\partial Y_i \partial Y_j} [P(X, t|Y, t')]. \quad (2.104)$$

The absorbing boundary condition is the same as for the forward Fokker–Planck equation, and is given by (2.103). The reflecting boundary condition takes a slightly different form, and is given by the condition

$$\sum_{i,j} n_i G_{ij}(Y) \frac{\partial P}{\partial Y_j} = 0, \quad X \in \partial S. \quad (2.105)$$

By using variations of the backward Fokker–Planck equation and appropriate boundary conditions, we can find the spatial distribution of moments and distributions of first exit times from a specific subset of S . This can be useful to make observations about bistable systems. We will go through, in more detail, the steps involved in setting up and solving the appropriate Fokker–Planck equation when we address specific problems later in this thesis.

2.21 Transforming probability distributions

If we have a probability distribution P , say as a solution of a Fokker–Planck equation, given in terms of Cartesian variables x , we can easily find the probability distribution \tilde{P} in terms of some other variables y [80]. The general formula for transforming probability distributions is

$$\tilde{P} = \left| \det \frac{\partial x}{\partial y} \right| P. \quad (2.106)$$

As an example, consider the case in which $y = (\phi, r)$ are polar coordinates. We find \tilde{P} to be

$$\tilde{P} = \left| \det \begin{bmatrix} -r \sin \phi & \cos \phi \\ r \cos \phi & \sin \phi \end{bmatrix} \right| P = rP. \quad (2.107)$$

In some cases, it will be useful to find the probability distribution in one coordinate system and then transform it to another. These relations allow us to readily do this.

2.22 Phase reduction of stochastically forced oscillators

One of the recent interests in the community has been of performing the analogue of phase reduction for a stochastically forced oscillator, that is, one in which the deterministic drift part of the system supports a limit cycle solution which is strongly attracting and where the magnitude of the diffusion terms is small relative to the drift. In this way, the diffusion term may be thought of as a weak perturbation to the limit cycle and so we may use similar techniques as described in Section 2.5. We begin with system (2.55) forced by a single white noise process, which we model through increments of a Wiener process

$$dX = F(X)dt + \sigma G(X)dW_t, \quad X \in \mathbb{R}^n, \quad (2.108)$$

with $\sigma \ll 1$. We assume that this system has a periodic orbit solution, with period Δ , for $\sigma = 0$ and interpret the SDE in the Stratonovich sense. This system can be rewritten in terms of isochronal coordinates with asymptotic phase θ and $n - 1$ amplitude variables described by the vector $\mathbf{r} = (r_1, \dots, r_{n-1})$ as

$$d\theta = \Omega dt + \sigma h(\theta, \mathbf{r})dW_t, \quad (2.109)$$

$$dr_i = f_i(\theta, \mathbf{r})dt + \sigma g_i(\theta, \mathbf{r})dW_t, \quad (2.110)$$

where $\Omega = 1/\Delta$, $h(\theta, \mathbf{r}) = (\nabla_x \theta) \cdot G(x(\theta, \mathbf{r}))$, $f_i(\theta, \mathbf{r}) = (\nabla_x r_i) \cdot F(x(\theta, \mathbf{r}))$ and $g_i(\theta, \mathbf{r}) = (\nabla_x r_i) \cdot G(x(\theta, \mathbf{r}))$, for $i = 1, \dots, n$. A direct application of the standard phase reduction, as described in Sect. 2.5 results in the following phase-only system

$$d\theta = \Omega dt + \sigma Q(\theta)dW_t, \quad (2.111)$$

where $Q = h(\theta, 0)$, so that the iPRC for the deterministic case can be used to describe the phase response to stochastic perturbations. This notion underpins comparisons of mathematical iPRCs to those estimated in experiments, as in [52], and has been used in a number of studies of synchronisation in populations of neurons exposed to a common noise source [94, 95, 96].

If we instead consider (2.108) in the Itô sense, we must use Itô's formula when performing the change of variables. The phase reduction in this case is performed

in [97], whereupon the resulting phase equation is

$$d\theta = \Omega dt + \frac{\sigma^2}{2} Q(\theta)Q'(\theta)dt + \sigma Q(\theta)dW_t. \quad (2.112)$$

Using (2.112), the authors of [97] then derive analytical expressions for a number of quantities of interest. In particular, they give a perturbative formula (expanded in terms of the noise strength) for the firing rate of a stochastically forced oscillator. We use this phase equation, along with the formula for the perturbed firing rate in Sect. 4.5.

In [7], Yoshimura and Arai demonstrate examples in which (2.111) fails to capture the dynamics of the full system (for the Stratonovich system). The authors then go on to derive a different phase equation by considering the Fokker–Planck equation for the evolution of the probability density for system (2.109)–(2.110). Upon applying a Taylor expansion for solutions to the resulting PDE around the limit cycle and defining

$$Y(\theta) = \sum_{i=1}^n \frac{\partial h(\theta, 0)}{\partial r_i} g_i(\theta, 0), \quad (2.113)$$

they write down the reduced phase model for (2.108) as

$$d\theta = \Omega dt + \frac{\sigma^2}{2} Y(\theta)dt + \sigma Q(\theta)dW_t. \quad (2.114)$$

We note that this differs from (2.111) through the addition of the term $\sigma^2 Y/2$. Through a comparison with numerical results for a standard oscillator model, Yoshimura and Arai show that this modified phase equation better captures the dynamics of the full system, compared with (2.111), for the case that they consider.

Teramae *et al.* [8] respond to this result from Yoshimura and Arai by first assuming that the white noise process is the limit as $\tau_n \rightarrow 0$ of a noise process with time correlation τ_n . Suppose that we now consider the system

$$dX = F(X)dt + \sigma G(x)\eta(t)dt, \quad (2.115)$$

$$\tau_n d\eta = -\eta dt + dW_t, \quad (2.116)$$

which we rewrite as

$$d\theta = \Omega dt + \sigma h(\theta, \mathbf{r}) \frac{y}{\sqrt{\tau_n}} dt, \quad (2.117)$$

$$dr_i = -\frac{1}{\tau_r} f_i(\theta, \mathbf{r}) dt + \sigma g_i(\theta, \mathbf{r}) \frac{y}{\sqrt{\tau_n}} dt, \quad (2.118)$$

$$dy = -\frac{y}{\tau_n} dt + \frac{dW_t}{\sqrt{\tau_n}}, \quad (2.119)$$

where τ_r represents the timescale of the attraction back to the limit cycle. The phase reduction makes the assumption that $\tau_r \rightarrow 0$, that is, instantaneous attraction back to the limit cycle. Thus, to arrive at an equation including solely phase for a white noise process, we need to take the limit $\tau_n \rightarrow 0$ and $\tau_r \rightarrow 0$ in (2.117)-(2.119). Teramae *et al.* argue that the order in which these limits accounts for the difference between (2.114) and (2.111). If we let $k = \tau_n/\tau_r$ denote the ratio of the two model timescales, and take both limits, we may write down a general phase only model as

$$d\theta = \Omega dt + \frac{\sigma^2}{2} Q(\theta) Q'(\theta) dt + \frac{1}{1+k} \sigma Y(\theta) dt + \sigma Q(\theta) dW_t. \quad (2.120)$$

The case of $k = 0$ corresponds to (2.114) whereas $k = \infty$ corresponds to (2.111). In later work, Goldobin *et al.* [9] go on to derive a phase-only model for systems forced with coloured noise, but we do not consider that here. Armed with (2.120), we may derive the steady state phase distribution, diffusion coefficient for the stochastic process and demonstrate that there is a noise induced frequency shift in the firing of the neuron, even though the white noise process does not favour any particular frequency.

3 Non-smooth spiking models

The present chapter focuses on the mathematical analysis of a planar, piecewise-linear integrate-and-fire type model, building upon the leaky integrate-and-fire (LIF) model discussed in the introduction, and forms the bulk of the work in [98]. At present, the level of available technology prevents us from simulating full brains in any reasonable timeframe. Projects such as the Blue Brain Project [4] are attempts to construct a realistic model of neural tissue from first principles. At present, initial simulations of $\sim 10^4$ biophysically detailed neurons have been performed, setting the scale of the tissue at roughly one neocortical column. Of course, one does not have to resort to detailed biophysical models. Given that a whole human brain contains $\sim 10^{10}$ neurons, there has been a push in the computational neuroscience community to develop complimentary models that are reduced in their complexity, yet still able to generate the rich repertoire of behaviour seen in a real nervous system. The advantage of such models lies in their mathematical tractability, and their computational cheapness relative to their more detailed counterparts. Models such as the ML model, considered in the introduction, can capture many of the characteristics of higher dimension models. Furthermore, phenomenological models, such as the FitzHugh-Nagumo (FHN) model, are further simplifications, aiming only to capture the emergent behaviour of the system, without intrinsically describing any of the biophysical processes in a neuron. The FHN model has a phase plane mimicking that of the reduced HH model considered in Sect. 2.2, with a cubic v -nullcline and linear w -nullcline, where w represents a recovery variable, and reproduces the same kinds of behaviour, such as oscillatory spiking solutions borne out of a Hopf bifurcation, and excitability of a quiescent state.

Analytical progress has also been possible with one further step, namely, the introduction of piecewise-linear nullclines. The piecewise-linear paradigm takes caricatures of the underlying vector field by dividing up the state space into a finite number of regions, in which the vector fields governing the dynamics are

linear. For example, for models with a cubic nullcline, a natural way to divide the state space is to split it into three pieces, so that the extrema of the nullcline are preserved. The vector fields are then pieced together across the boundaries between them, which we call *switching manifolds*. For the FHN case, this gives rise to the so-called McKean model [99], for which a number of results about the existence and stability of periodic orbits are now known [100, 101, 102]. Across the switching manifold, there are a number of interesting behaviours that can occur, dependent on the relationship between the vector fields on either side of the switch. This can lead to discontinuities in the vector field, or in more extreme cases, sliding motion along the switching manifold, as considered in [103]. The study of such motions has been an area of much recent interest, particularly since it leads to a lack of uniqueness of trajectories in backwards time. However, we are typically interested in piecewise smooth systems in neural modelling, since these are generally the best caricatures of the underlying vector field. Indeed, there are now a number of planar PWL single neuron models for mimicking the behaviour of tonically firing neurons. A review of some of these models may be seen in [104]. The PWL nature of such models means that techniques from non-smooth dynamics are particularly relevant to their analysis and, as an example, recent progress on understanding canard explosions has been made by studying PWL models of FHN type [105]. However, the spiking patterns of such planar models are typically not as diverse as one needs to mimic realistic firing patterns, such as *bursting*. The main aim of this chapter is to extend the analysis of PWL neural models to a case with a threshold, as in the IF case, and to allow the study of bursting firing states.

We define bursting as a dynamic state in which a neuron fires two or more spikes rapidly, termed the *active* phase, followed by a period of quiescence, termed the *silent* phase [47]. Bursting patterns are observed in a number of brain regions, such as the neocortex [106] and pre-Bötzinger complex [107, 108] and may be thought of as another firing mode of a neuron. Regarding information transfer across the brain, bursting may be a useful firing modality, as bursts are more re-

liable than single spikes, are able to overcome synaptic transmission failure [109], are more robust to noise [110] and potentially encode more and different information than single spikes alone [111, 112]. Smooth dynamical systems descriptions of bursting must have dimension three or higher, and typically involve a separation of timescales between a fast subsystem and a slow subsystem. The fast subsystem exists either in an oscillatory regime or a quiescent regime, and may take describe the membrane potential and the gating variables as described in the introduction. The slow subsystem drives the fast subsystem through bifurcations initiating and terminating the oscillatory mode. The slower intracellular calcium processes in cells have been put forward as candidates for the slow subsystem of such models [47]. Systems with such a separation of timescales are amenable to *fast-slow* decomposition [113]. This decomposition assumes an infinite separation of timescales to study the two subsystems separately, with the variables of the slow subsystem considered as parameters of the fast subsystem. The dynamics of the slow subsystem are then obtained by describing motion on a critical manifold of the fast subsystem in the slow subsystem. Mathematically, if we write the system as

$$\begin{aligned}\frac{dx}{dt} &= f(x, w), & \text{fast subsystem,} \\ \frac{dw}{dt} &= \varepsilon g(x, w), & \text{slow subsystem,}\end{aligned}\tag{3.1}$$

where $\varepsilon \ll 1$, $x \in \mathbb{R}^n$ are the fast variables and $w \in \mathbb{R}^m$ are the slow variables. The fast subsystem may be studied by setting $\varepsilon = 0$, to give

$$\begin{aligned}\frac{dx}{dt} &= f(x, w), \\ \frac{dw}{dt} &= 0\end{aligned}\tag{3.2}$$

treating w as parameters. In order to study dynamics in the slow subsystem, we rescale time as $\tau = \varepsilon t$. After some rearranging, and again setting $\varepsilon = 0$, we now have

$$\begin{aligned}0 &= f(x, w), \\ \frac{dw}{d\tau} &= g(x, w),\end{aligned}\tag{3.3}$$

where $f(x, w) = 0$ is the critical manifold in the above differential algebraic equation. Studying bursting models in this way allows us to study two simpler models and piece together solutions at the ‘jump’ points between subsystems, using matched asymptotics. Finally, the relaxation of an infinite separation of time scales may be achieved using geometric singular perturbation theory, as reviewed in [114]. This is performed by treating ε as an expansion parameter to piece together slow manifolds of the system, adhering to results from Fenichel [115]. Since the active phase of the burst is initiated and terminated by bifurcations of the fast subsystem, this has led to a classification of bursting types by these bifurcations [47, 116], where the different bursting types possess different properties.

The currently most successful class of minimal models that satisfy the criterion of being able to generate realistic firing patterns, such as bursting, are those of IF type as considered in Sect. 2.3, where a simple threshold unit is used to caricature the excitable aspect of real cells that gives rise to an action potential spike. In these models, the spike shape is discontinuous. The LIF model has only one dimension, and does not support bursting. However, appending a recovery variable akin to that for the ML or FHN model to the LIF model will enable it to support bursting [117, 118]. In this case, the discontinuous reset mechanism replaces the need for a third variable required in the smooth case. One key aspect of any IF model is the discontinuous reset of a state variable upon reaching some threshold for spiking. It is this particular harsh nonlinearity in the dynamics that endows these models with interesting dynamics and precludes their description using the machinery of smooth dynamical systems. Indeed, they have much in common with models of impacting systems that have been developed for the study of mechanical structures such as rocking blocks [119], rattling gear boxes [120] and print hammers [121]. Many results for impacting systems may be found in the recent book by di Bernardo *et al.* [122].

This chapter combines both the ideas of using piecewise-linear systems, along with IF type dynamics in the form of a new IF model, namely the PWL-IF model. This model is able to reproduce a range of spiking patterns, from tonic to burst

firing, yet is analytically tractable. In essence, the model below the threshold for firing evolves according to a planar PWL dynamical system. We present an original bifurcation analysis of this model in response to constant current injection, focusing on local discontinuity-induced bifurcations. Next, in Sect. 3.7, we show how to construct periodic orbits and determine their stability as well as calculate the phase response curve (by adapting techniques originally developed for the analysis of limit cycles in smooth dynamical systems). Following this, spike-adding bifurcations (for bursting orbits) are described in terms of bifurcations of an associated one-dimensional return map. Finally, the notion of Lyapunov exponents for this IF system is developed, using techniques originally developed for the analysis of impact oscillators.

3.1 A review of integrate-and-fire models

We begin by briefly reviewing the 1D integrate-and-fire models discussed in Sect. 2.3. The general one-dimensional nonlinear IF model takes the form

$$\frac{dv}{dt} = f(v) + I(t), \quad (3.4)$$

such that v is reset to v_R just after reaching the threshold value $v_{th} > v_R$. Here, v is interpreted as a voltage variable and $I(t)$ is an external drive (that might be under the control of an experimentalist or arise from the activity of other neurons to which a cell is coupled). Firing times are defined iteratively according to

$$T_n = \inf\{t \mid v(t) \geq v_{th} ; t \geq T_{n-1}\}. \quad (3.5)$$

One-dimensional IF models with a fixed voltage threshold are caricatures of excitable neural systems and as such it is well to mention that they cannot adequately capture the refractory properties of real neurons. This is often achieved with the introduction of an absolute time during which they cannot fire after reaching threshold or by the introduction of a time dependent threshold that increases after a firing event and makes it harder for the neuron to subsequently fire (mimicking a relative refractory period), as reviewed in [32]. Moreover, real

neurons (and HH style models) do not possess a fixed voltage threshold, and firing ultimately depends on the state of receptors within a membrane. Although differential equations for the threshold in IF models can be found that mimic more closely the properties of real neurons [123], we only consider a fixed threshold in this thesis.

Electrophysiologists often classify neuron response in terms of the so-called $F - I$ curve, which shows the frequency of oscillation as a function of the time independent drive I . For the LIF model this is easily constructed from (2.34) using $F = \Delta^{-1}$, showing a sharp rise in F (from zero) as I increase through the critical value v_{th}/τ . The tractability of this single neuron model, which has linear dynamics between firing events, means that it is particularly suited to analysis at the network level with event based models of chemical synapses. Indeed, a theory of phase-locked behaviour for strong coupling has been developed for just this scenario [124]. However, since the model does not intrinsically produce anything mimicking the shape of an action potential, it is not well suited for all cases. In particular, to effectively model gap junctions, we need to have a model in which the v variable mimics the voltage of a real cell, particularly in the spiking regime.

3.2 Nonlinear IF models

The quadratic IF (QIF) neuron is the simplest generalisation of the LIF model that captures qualitatively the behaviour of the $F - I$ curve of a large family of more realistic models [125]. Interestingly, this model was apparently already known to Alan Hodgkin, and used to fit some of his data (and also subsequently analysed by Bruce Knight). Up to shifts and constant factors it is defined by

$$f(v) = v^2. \tag{3.6}$$

Unlike the LIF model, the QIF does allow a representation of an action potential shape (for $I > 0$ the voltage rises sharply to threshold), as shown in Fig. 14. For $I < 0$ there are two equilibria (one stable and the other unstable) and for $I > 0$ these disappear via a saddle-node bifurcation at $I = 0$. In the oscillatory regime

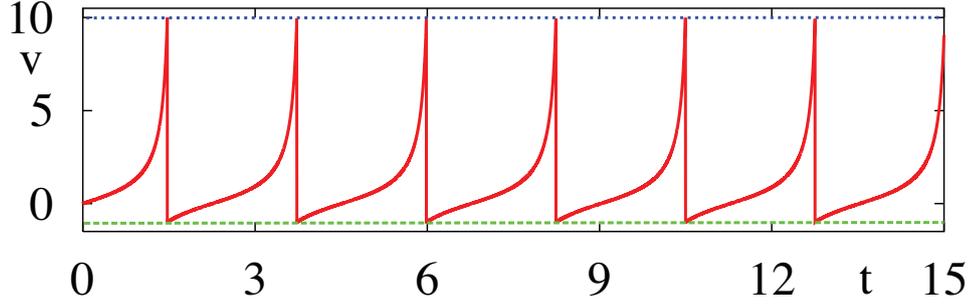


Figure 14: Voltage trace for the QIF oscillator, (3.4) with f given by (3.6), whose solution is given by (3.7), with constant drive $I = 1$ with $v_{\text{th}} = 10$ and $v_{\text{R}} = -1$.

($I > 0$), the trajectory (for constant drive) can be integrated for $T_n < t < T_{n+1}$ to give

$$v(t) = \sqrt{I} \tan \left(\tan^{-1} \left(\frac{v_{\text{R}}}{\sqrt{I}} \right) + \sqrt{I}(t - T_n) \right). \quad (3.7)$$

The period of oscillation is calculated by setting $v(T_{n+1}) = v_{\text{th}}$ with $v(T_n) = v_{\text{R}}$ giving

$$\Delta = \frac{1}{\sqrt{I}} \left(\tan^{-1} \left(\frac{v_{\text{th}}}{\sqrt{I}} \right) - \tan^{-1} \left(\frac{v_{\text{R}}}{\sqrt{I}} \right) \right) H(I).$$

In the limit $v_{\text{th}} \rightarrow \infty$ and $v_{\text{R}} \rightarrow -\infty$ we see that $\Delta = \pi/\sqrt{I}$ (and we have blowup of the voltage trajectory in finite time), and the $F - I$ curve shows a \sqrt{I} dependence, which matches many cortical neurons much better than the LIF $F - I$ curve. For further discussion of this model, we refer the reader to the book by Izhikevich [47].

With the improvement in neuronal modelling by simply changing the shape of the nonlinearity from (2.31) to (3.6), this raises the question as to whether more judicious choices can improve things further still. Interestingly, Fourcaud-Trocme *et al.* [126] have shown that choosing $f(v) = \exp(v)$ (up to shifts and scaling) can act as an approximation of a more detailed conductance-based spiking model. In fact, it has now been shown that real cortical data (from layer-5 pyramidal cells) can be very accurately fit with the following choice [2]:

$$f(v) = -\frac{1}{\tau}(v - v_{\text{L}}) + \frac{\kappa}{\tau} e^{(v - v_{\kappa})/\kappa}, \quad (3.8)$$

with $v_{\text{th}} = 30.0$ mV, $v_{\text{R}} = -71.2$ mV, $v_{\text{L}} = -68.5$ mV, $\tau = 3.3$ ms, $v_{\kappa} = -61.5$ mV

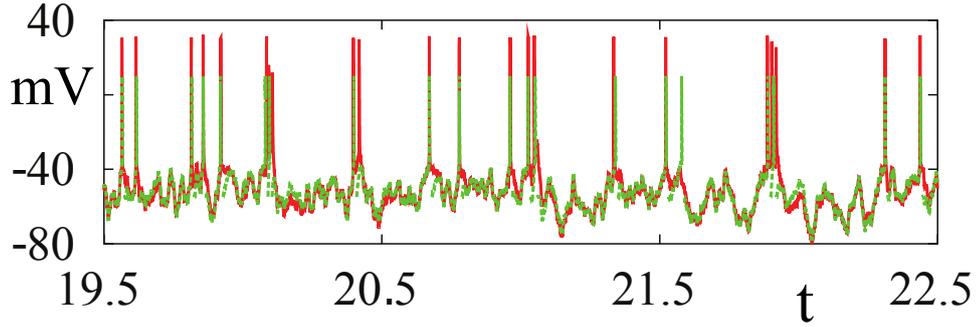


Figure 15: Sample voltage traces (mV) as a function of time (seconds) from the linear-exponential IF model (green dashed line), (3.4) with f given by (3.8), and data (red solid line) from a layer-5 pyramidal cell in response to a noisy current injection (see [2] for further details).

and $\kappa = 4\text{ mV}$. A fit to sample data using this model may be seen in Fig. 15. In a similar fashion to the QIF model, the linear-exponential IF (LEIF) model obtained using (3.8) has two equilibria (defined by $f(v) + I = 0$), which disappear in a saddle-node bifurcation when $I = -f(v^*)$, where v^* is defined by $f'(v^*) = 0$. In common with the QIF model, it is able to support oscillations with arbitrarily low frequency just beyond the bifurcation point. Both the QIF and LEIF models have only a weak dependence on the choice of threshold value since they both blow up in finite time (in the absence of a threshold).

3.3 Planar IF models

Unfortunately, one-dimensional nonlinear IF models, as they stand, are unable to reproduce bursting patterns of activity. One way to incorporate a slow process that would give rise to bursting is by coupling the voltage dynamics to a recovery process:

$$\begin{aligned} \frac{dv}{dt} &= f(v) - a + I, \\ \frac{1}{\omega} \frac{da}{dt} &= \beta v - a. \end{aligned} \tag{3.9}$$

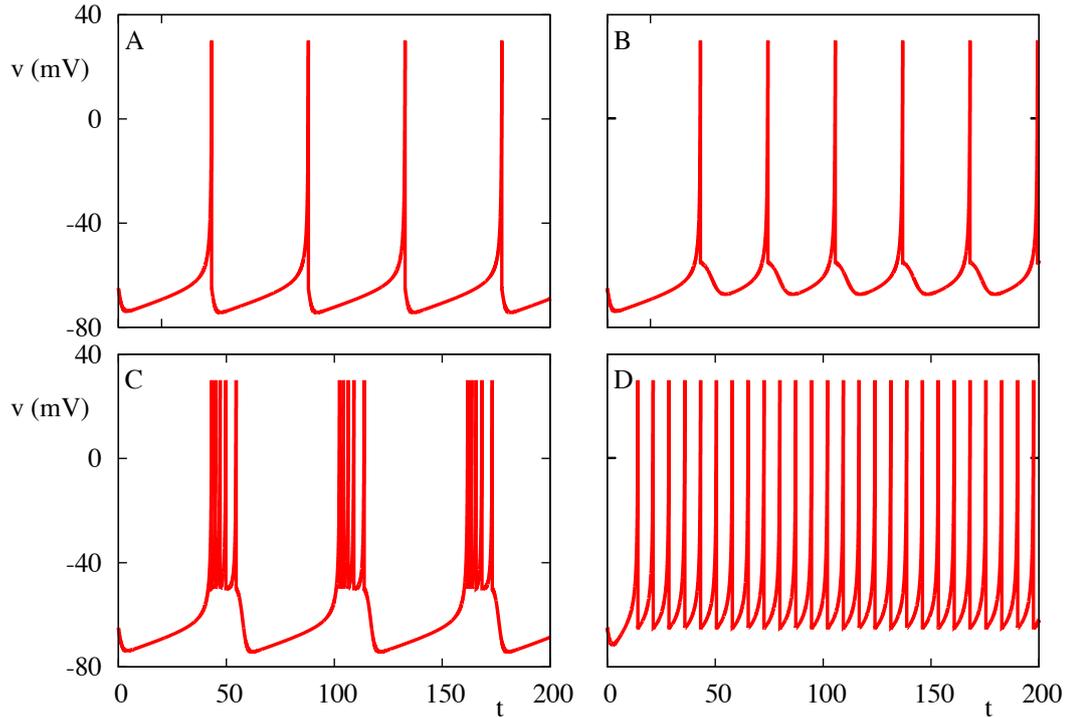


Figure 16: Firing patterns in the Izhikevich model (3.9) with f given in the text below, with $I = 10$ and $v_{\text{th}} = 30$. Voltage traces as a function of time for A: tonic spiking, with $\omega = 0.02$, $\beta = 0.2$, $v_{\text{R}} = -65$, $k = 8$, B: tonic spiking, with $\omega = 0.02$, $\beta = 0.2$, $v_{\text{R}} = -55$, $k = 4$, C: bursting, with $\omega = 0.02$, $\beta = 0.2$, $v_{\text{R}} = -50$, $k = 2$, and D: fast spiking, with $\omega = 0.1$, $\beta = 0.2$, $v_{\text{R}} = -65$, $k = 2$.

Here, the parameters β and ω respectively describe the sensitivity and decay rate of the adaptive process. Upon reaching threshold, the voltage is reset ($v \rightarrow v_{\text{R}}$) and a is adjusted according to $a \rightarrow a + k$. The Izhikevich model [127, 128] is one such model with $f(v) = 0.04v^2 + 5v + 140$. Interestingly, this model can capture a number of neuronal firing patterns including tonic (repetitive) spiking, bursting and fast spiking as illustrated in Fig. 16, despite its sensitivity to the choice of threshold value [129]. It is worth noting that a similar model to that of Izhikevich was independently introduced by Gröbner *et al.* [130] as a model of a pyramidal cell in hippocampus CA3. The adaptive exponential integrate-and-fire model is obtained using a linear-exponential term for $f(v)$ (as in equation (3.8))

[131, 132], whilst the quartic model is obtained by choosing $f(v) = v^4 + 2\omega v$ [133]. Both are able to produce a wide variety of firing patterns, and the quartic model in particular has a very nice repertoire of responses ranging from tonic spiking to bursting as well supporting phasic responses, rebound, spike frequency adaptation, sub-threshold oscillations and much more, all of which are discussed in detail in [133]. Since all of these models are planar, we may study them using phase-plane analysis, without any requirement for a separation of timescales.

3.4 The absolute IF model

Apart from the LIF model, none of the models described above admits closed form solutions for arbitrary (non-constant) drive. A somewhat overlooked, tractable (one-dimensional) nonlinear IF model is that of Karbowski and Kopell [134], with a nonlinearity given by $f(v) = |v|$, which we shall call the absolute IF model (AIF). Because of the choice of a PWL form of the nonlinearity, the AIF model can be explicitly analysed. Moreover, it is also capable of generating behaviour consistent with that of a fast-spiking interneuron [135]. The generalisation of the model to allow for bursting behaviour is easily achieved by extending it to the form of (3.9). A minimal AIF model with adaptation is obtained for $f(v) = |v|$ and $\beta = 0$. For sufficiently small k , the model fires tonically and for larger values of k the model can also fire in a burst mode. The mechanism for this behavior in the AIF model (and indeed all planar models discussed here) is most easily understood in reference to the geometry of the phase-plane. We illustrate, in Fig. 17, the phase plane for the AIF model. The analysis of how parameter space partitions into tonic, 1-spike per burst, 2-spike per burst, etc. firing patterns is an open mathematical (classification) challenge. It is worth noting that all the planar models considered here have much in common and can generate a very similar repertoire of firing behaviours, though the AIF model does not have trajectories that blow up in finite time (in the absence of a threshold).

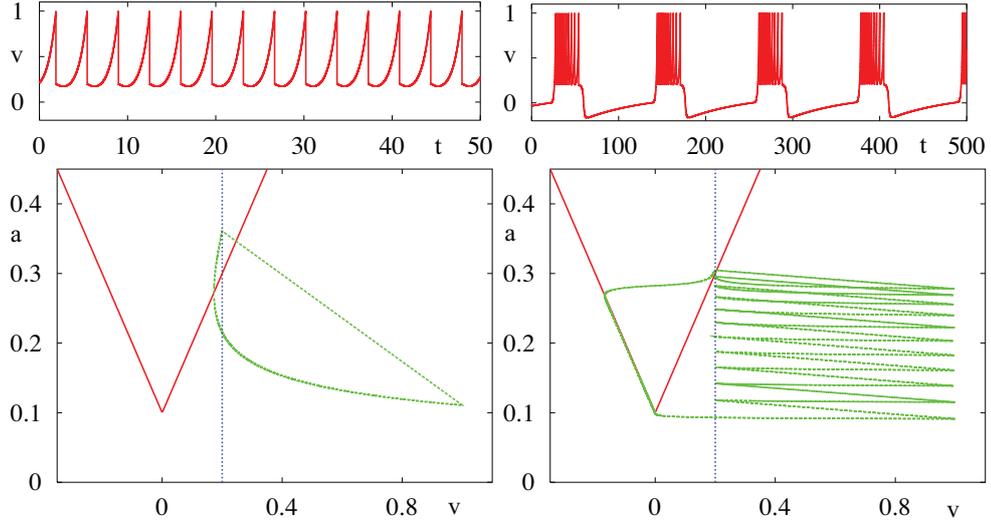


Figure 17: Top left: Tonic firing in the AIF model with spike adaptation. Here, $\omega = 1/3$ and $k = 0.75\omega$. Top right: Burst firing in the AIF model with spike adaptation. Here, $\omega = 1/75$ and $k = 2\omega$. Bottom left: A periodic orbit in the (v, a) plane corresponding to the tonic spiking trajectory shown above (green curve). Also shown is the voltage nullcline (red lines) as well as the value of the reset (black dashed line). Bottom right: Burst firing in the AIF model with spike adaptation. Other parameters are $\beta = 0$, $v_R = 0.2$, $v_{th} = 1$ and $I = 0.1$.

3.5 A piecewise-linear IF model

The aspect of the LIF model that allows one to perform an analysis such as the one above is obviously its linearity (below threshold). A similar analysis for say the QIF, LEIF or Izhikevich model would be much harder, owing to the inherent nonlinear nature of these models. However, the AIF model described in Sect. 3.1 is a natural starting point for the development of a more general PWL spiking neuron model that can be explicitly analysed. We begin by advocating a new type of PWL IF model, that we shall call the PWL-IF model. It is a generalisation of the AIF model with adaptation that we write in the form of (3.9) with

$$f(v) = \begin{cases} v & v \geq 0 \\ -sv & v < 0 \end{cases}, \quad (3.10)$$

with $s > 0$, where we now treat all variables and parameters as being dimensionless. Throughout this chapter, we set $v_R = 20, v_{th} = 60$, and note that a realistic range for the voltage component of this model may be obtained as $V = v - 40$. For a constant drive I , the model may exhibit a number of different periodic attractors, and in particular we distinguish between those that remain sub-threshold, and those that cross threshold, which we shall term *spiking solutions*. We make further distinctions between solutions as follows.

- **Fast spiking orbits:** Attracting limit cycles which have $v > 0$ along the entire orbit and which have $v(t^*) = v_{th}$ at precisely one value of $t^* \in [t, t + \Delta] \quad \forall t$, where Δ is the period of the limit cycle.
- **Regular (or tonic) spiking orbits:** Attracting limit cycles which have $v < 0$ for some segment of the orbit and which have $v(t^*) = v_{th}$ at precisely one value of $t^* \in [t, t + \Delta] \quad \forall t$, where Δ is the period of the limit cycle.
- **n -Spike bursting orbits:** Attracting limit cycles which have $v < 0$ for some segment of the orbit and which have $v(t^*) = v_{th}$ at precisely n values of $t^* \in [t, t + \Delta] \quad \forall t$, where Δ is the period of the limit cycle.
- **Sub-threshold oscillations:** Attracting limit cycles which have $v < v_{th}$ along the entire orbit.

The fast spiking orbits are so called as they may have arbitrarily fast frequency, whereas the frequency of regular spiking orbits must be finite. With increasing I , the model can make a transition from regular to fast spiking. Contrary to the case for smooth systems, periodic orbits in discontinuous systems need not enclose a fixed point. In fact, the reset mechanism of the PWL-IF model allows for periodic orbits of (3.9) in the absence of any fixed points. For $\beta < 1$, the $F - I$ curve (regular spiking) reaches a maximum value before a bifurcation to fast spiking occurs. The switch between the two modes for $\beta > 1$ may have a further signature of doublet (2-spike burst) firing (which we shall consider in more detail below), and leads to a discontinuous $F - I$ curve. We depict, in Fig. 18,

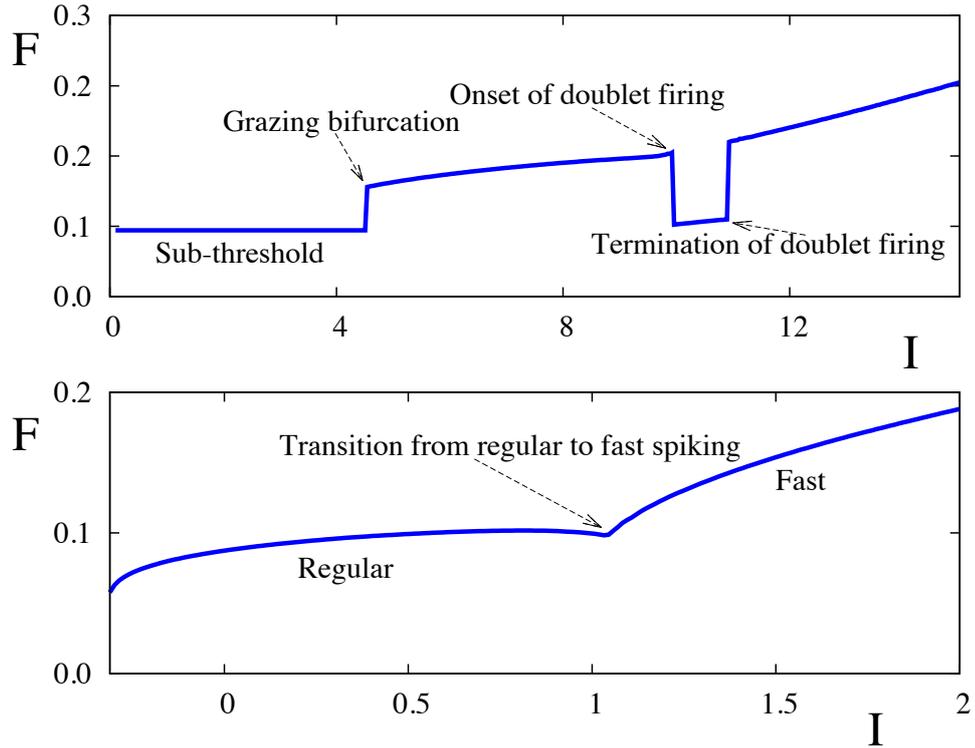


Figure 18: Firing frequency under variation of the drive I in the PWL-IF model, (3.9) with f given by (3.10), for: Top: $\beta = 1.2$, Bottom: $\beta = 0.9$. We can clearly see how the firing rate changes as we move between solution types, and that the firing rate of the model during fast spiking is much more sensitive to changes in I than in the regular spiking mode. Other parameters are $s = 0.35$ and $\omega = 0.9$. The top curve may be thought of as moving through Fig. 22 along a horizontal section at $\beta = 1.2$, though the blue (GB), black (SP) and pink (DB) bifurcation curves, whereas the bottom curve shows the transition through the SP bifurcation, this time along a horizontal section at $\beta = 0.9$.

the $F - I$ curve for differing values of β under variation of I . We can clearly see the transitions between the different oscillatory regimes, particularly for $\beta = 1.2$, where we observe discontinuities in the frequency response at a grazing bifurcation, in which a subthreshold orbit tangentially touches threshold, and at the onset and termination of doublet firing. We will review grazing bifurcations in more detail in the next section.

3.6 Bifurcation structure

In order to characterise where in parameter space different types of solution exist, it is useful to consider the different types of bifurcation that can occur. The v -nullcline has a characteristic ‘V’ shape, whilst the a -nullcline is a straight line with slope β . By inspection, we see that there may exist one, two or no fixed points of (3.9) with f defined as in (3.10). There is a slight subtlety, in that the nullclines may intercept where $v > v_{\text{th}}$, generating a *virtual* fixed point. From here on, we refer to the branch of the v -nullcline with $v < 0$ ($v > 0$) as the left (right) v -branch. Since the system is PWL we may easily construct the eigenvalues of fixed points, where they exist, as

$$2\lambda_{\pm} = \begin{cases} 1 - \omega \pm \sqrt{(1 - \omega)^2 - 4\omega(\beta - 1)}, & v > 0, \\ -s - \omega \pm \sqrt{(s + \omega)^2 - 4\omega(\beta + s)}, & v < 0. \end{cases} \quad (3.11)$$

Thus, fixed points on the left v -branch are always stable, and the stability of fixed points on the right v -branch depends on the sign of $1 - \omega$. The exact nature of the fixed points is determined by the sign of the expression under the square root. Since β must be less than 1 to have two fixed points, the fixed point on the right v -branch is a saddle whenever the system supports two fixed points.

The sub-threshold dynamics are described by a continuous but non-differentiable system, so that the Jacobian matrix (around a fixed point) is not defined at the *switching manifold*, that is, the border separating linear subsystems at $v = 0$. Non-smooth bifurcations can occur as fixed points or limit cycles touch the switching manifold under parameter variation. Importantly, the presence of a firing threshold in IF systems means that other non-smooth bifurcations, can arise. We illustrate, in Fig. 25, the stable solutions in the various regions of parameter space indicated in Fig. 22. The curves in Fig. 22 are generated by numerical continuation of solutions obtained from the firing map discussed later in Sect. 3.12.

The PWL-IF can generate periodic behaviour via a Hopf bifurcation (HB) of a fixed point on the right v -branch when $\omega = 1$ (with $\beta > 1$) or through a

discontinuous Hopf-like (dHB, black line in Fig. 22) bifurcation at $I = 0$ (with $\omega < 1$). The conditions for a Hopf bifurcation in the smooth case are also met by the dHB, namely that the eigenvalues of the Jacobian around the fixed point have nontrivial real part, with real part crossing through zero. For planar systems, this is equivalent to the condition that the trace of the Jacobian (around the fixed point) passes through a zero, whilst its determinant is non-zero. We describe the dHB as being discontinuous since the Jacobian around the fixed point changes discontinuously. Borne out of the dHB is a sub-threshold limit cycle, which crosses through the switching manifold $v = 0$. We plot in Fig. 19, a plot of the phase plane, and a few sample trajectories just before ($I = -0.5$) and just after ($I = 0.5$) the dHB. Additionally, the inset in this figure shows the trace of the Jacobian at the fixed point, evaluated under smooth variation of I , so that we see that it clearly changes discontinuously at the bifurcation at $I = 0$. The limit cycle produced by the dHB bifurcation is similar to those in smooth systems, with a few noticeable differences. In fact, these sub-threshold limit cycles may be thought of as being composed of two distinct spiral sections pieced together across the switching manifold [136]. This may be seen more clearly in Fig. 20. In this figure, we plot the extension of the spiral sections which form the limit cycle in each PWL region using the corresponding vector field. That is, we extend the spiral section for $v > 0$, into $v < 0$, propagating with the vector field for $v > 0$, instead of using the vector field for $v < 0$ (and vice versa). The red curve shows the extension of the spiral for $v < 0$ into $v > 0$. This trajectory tends towards the fixed point, shown by the intersection of the dashed orange (representing the extension of the left v -branch into $v > 0$) with the cyan line, representing the a -nullcline. Of course, this fixed point does not exist in the full PWL system. The blue curve is the extension of the spiral for $v > 0$ into $v < 0$. Now, the only fixed point in this linear system is unstable (it is an unstable focus), and so trajectories spiral outwards. We can clearly see that, the limit cycle solution is the union of these spiral in each respective PWL region, with one section spiralling outwards, the other complementing this by spiralling inwards (towards a virtual fixed point of

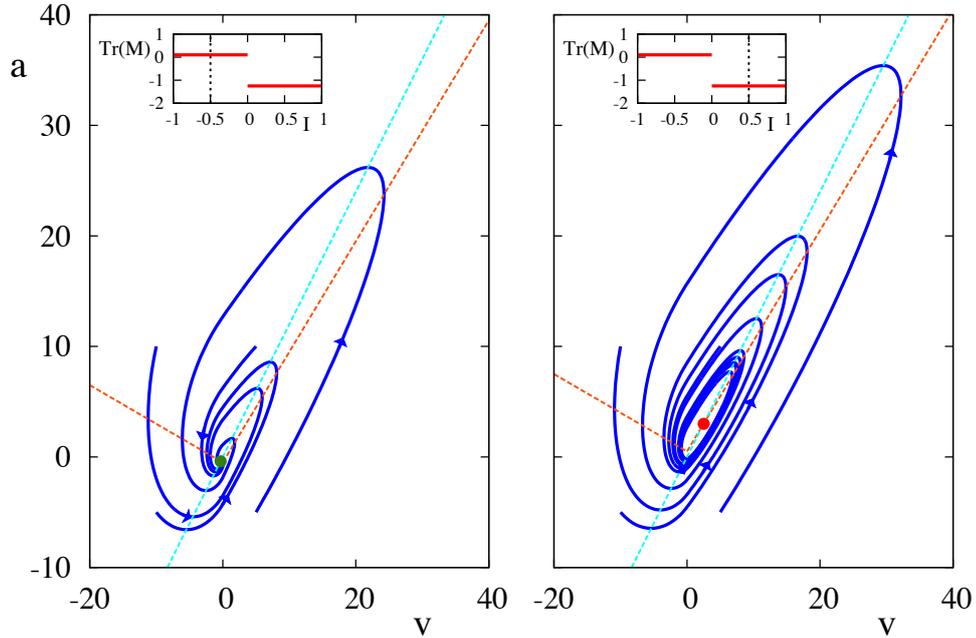


Figure 19: Phase plane and sample solutions just before (left) and just after (right) the dHB bifurcation. In the left panel, we set $I = -0.5$. We see that the fixed point, depicted by a green circle, exists on the left v -branch and so is stable. All of the sample trajectories spiral in towards the fixed point. For $I = 0.5$, in the right panel, the fixed point, now shown as a red circle, is now on the right v -branch and is unstable. Trajectories now tend towards a limit cycle solution, which may be thought of as a union of spiral sections in $v < 0$ and $v > 0$. This may be seen more clearly in Fig. 20. The inset shows the trace of the Jacobian evaluated around the fixed point under smooth variation of I , with the vertical black lines indicating the value of I in each of the respective panels. We observe a discontinuity in the trace as it passes through zero, at the bifurcation at $I = 0$. This is why we refer to this bifurcation as the dHB, in that it meets the conditions for a Hopf bifurcation, but the trace of the Jacobian (and hence its eigenvalues) is discontinuous at the bifurcation point. Also shown are the v - and a -nullclines in orange and cyan respectively. We set $\beta = 1.2$, with other parameters as in Fig. 22, so that this figure shows the transition across the dHB (black) curve at $\beta = 1.2$.

the $v < 0$ system). Limit cycle solutions of this type differ from the smooth case in that, firstly, the amplitude of the oscillation grows linearly with I , rather than growing proportional to \sqrt{I} and secondly, the criticality of the Hopf bifurcation is determined by linear, rather than cubic terms. The previous two statements are true for general piecewise smooth systems with discontinuous Hopf bifurcations. For PWL systems, we also have that the period of the oscillations is invariant under changes in I (whilst the oscillation remains subthreshold). We will go on to prove this result in Sect. 3.9.

As the limit cycle grows, it can tangentially touch the firing threshold. We refer to this scenario as a grazing bifurcation, at $v = v_{\text{th}}$, whereupon sub-threshold oscillations are replaced by regular spiking solutions. We later illustrate, in the phase plane, this transition from sub-threshold to spiking orbits in Fig. 30. In Fig. 22, we may observe both the dHB (black) and the grazing bifurcation (blue) in (I, β) parameter space. Grazing bifurcations may also occur at $v = 0$, where a graze results in the transition from fast to regular spiking, which may occur after a window of doublet firing. The black curve (SP), in Fig. 22 corresponds to the transition to regular spiking, either from fast spiking, or from doublet firing, whereas the pink curve (DB) marks the onset/termination of doublet firing. We note that in order to have a graze at v_{th} , we require that $\beta > \beta_c$, since we need the v -nullcline to be below the a -nullcline for $\dot{v} = 0$ in this part of the phase-plane. The preceding discussion considers transitions to and from non-bursting solutions. We shall consider the transitions between bursting states in due course.

Bistability can arise between a stable fixed point on the left v -branch and a fast spiking orbit when $\beta < 1$ and $I < 0$. In this parameter regime, there exists a saddle node on the right v -branch, which is key in delineating the basins of attraction of the two attractors, as shown in Fig. 24. The basin of attraction of the stable fixed point is the union of the set of initial data such that trajectories reach threshold and are subsequently reset to the right of the separatrix of the saddle on the right v -branch. A homoclinic bifurcation (HC), indicated by the red curve in Fig. 22, will occur when the spiking limit cycle collides with the

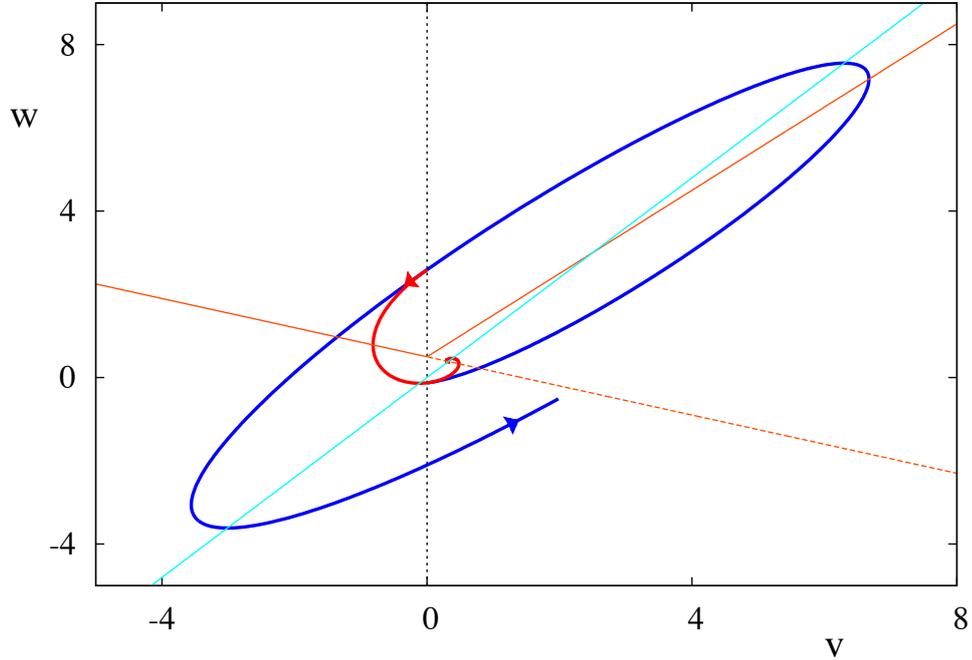


Figure 20: The spiral sections which form the limit cycle borne out of the dHB. In this figure, we extend the spiral section for $v < 0$, into $v > 0$, but under the evolution of the vector field in $v < 0$, rather than for $v > 0$ (and vice versa). As such, we extend each the spirals by making the system linear everywhere. The red curve shows the extension of the spiral for $v < 0$ into $v > 0$. This trajectory tends towards the fixed point, shown by the intersection of the dashed orange (representing the extension of the left v -branch into $v > 0$) with the cyan line, representing the a -nullcline. The blue curve is the extension of the spiral for $v > 0$ into $v < 0$. The only fixed point in this case is an unstable focus and so trajectories spiral outwards. The limit cycle is seen to be the union of these two spiral sections in the respective PWL regions. Parameters values are as in the right panel of Fig. 19.

saddle, resulting in a homoclinic orbit from the saddle at the bifurcation point. The homoclinic orbit at this point may be thought of as being comprised of two straight lines (in addition to the reset condition), which are sections of the stable and unstable manifolds of the saddle. The homoclinic orbit at its corresponding bifurcation is shown in Fig. 21. Another form of bistability is also possible in this parameter regime, namely when a regular spiking limit cycle encloses the stable fixed point. The basin of attraction of this limit cycle is the set of points such that trajectories reach threshold and are reset to the right of the separatrix of the saddle (which is also enclosed by the stable spiking orbit). The basin of attraction is thus the union of disconnected sets. Numerical studies suggest that the regular spiking orbit is lost as the basin of attraction of the stable fixed point grows and touches the orbit, and as such we shall call this an orbit crisis. As with the HC bifurcation, after this point all trajectories will tend towards the stable fixed point. A plot of the basins of attraction of the two attractors is shown in Fig. 24, whilst a plot of parameter values for which we have an orbit crisis is depicted by the green (OC) curve in Fig. 22.

For $\beta < 1$, we have a discontinuous saddle node bifurcation (dSN, orange line in Fig. 22) at $I = 0$, where the saddle and stable fixed point come together and annihilate one another. We refer to this as a discontinuous bifurcation owing to the fact that the Jacobian of the system is undefined at the bifurcation point, $I = 0$. For $I > 0$, there are no fixed points, and the only attractor is either the regular spiking or fast spiking orbit, dependent on the value of β . We show, in Fig. 23, the phase plane and sample trajectories just before and just after the dSN bifurcation, using the parameter values as in Fig. 19, with $\beta = 0.7$. In the left panel, we observe two fixed points, one on each of the v -branches. The fixed point on the left v -branch is stable, whilst the one on the right v -branch is a saddle. As we have bistability between between the stable fixed point and a fast spiking solution, trajectories will head towards one of these two attractors. After the bifurcation, at $I = 0.5$, we no longer have any fixed points, and so all trajectories will tend towards the fast spiking solution. For planar systems, the condition that

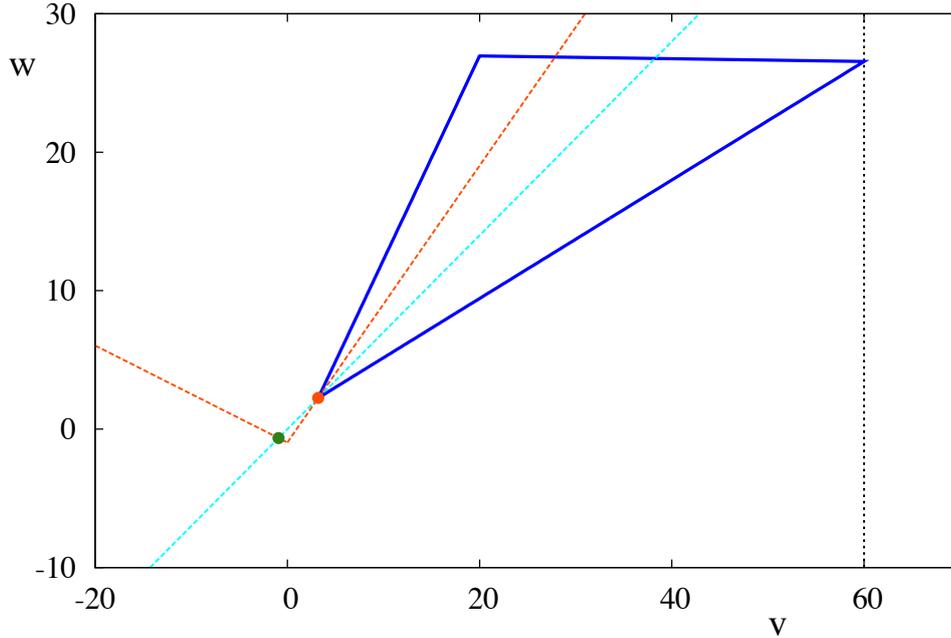


Figure 21: Homoclinic orbit at the homoclinic bifurcation point, with $\beta = 0.7, I \approx -0.9568$. Other parameters are as in Fig. 22. We observe that the homoclinic connection comprises two straight lines, which are sections of the stable and unstable manifolds of the saddle (depicted as an orange circle), in addition to the reset condition. Also shown are the v - and w -nullclines, in orange and cyan respectively, and the stable fixed point (green circle). This bifurcation corresponds to crossing the red (HC) curve in Fig. 22.

a saddle node bifurcation occurs is a passing through zero of the determinant of the Jacobian around the fixed point. The inset of Fig. 23 shows the determinant of this Jacobian, M . We see that, as in the case for the dHB, the determinant passes through zero discontinuously, which is why we refer to the saddle node bifurcation as the dSN. If $\beta > 1$, then the system only possesses one fixed point, which may be on the left or right v -branch dependent on the sign of I . As I crosses 0 from below, there are three scenarios: either $\omega > 1$, in which case no change of stability occurs and trajectories tend to the fixed point, else $\omega < 1$ and the fixed point becomes unstable. We either may observe sub-threshold oscillations or spiking oscillations (either bursting or tonic) depending on the other parameter values.

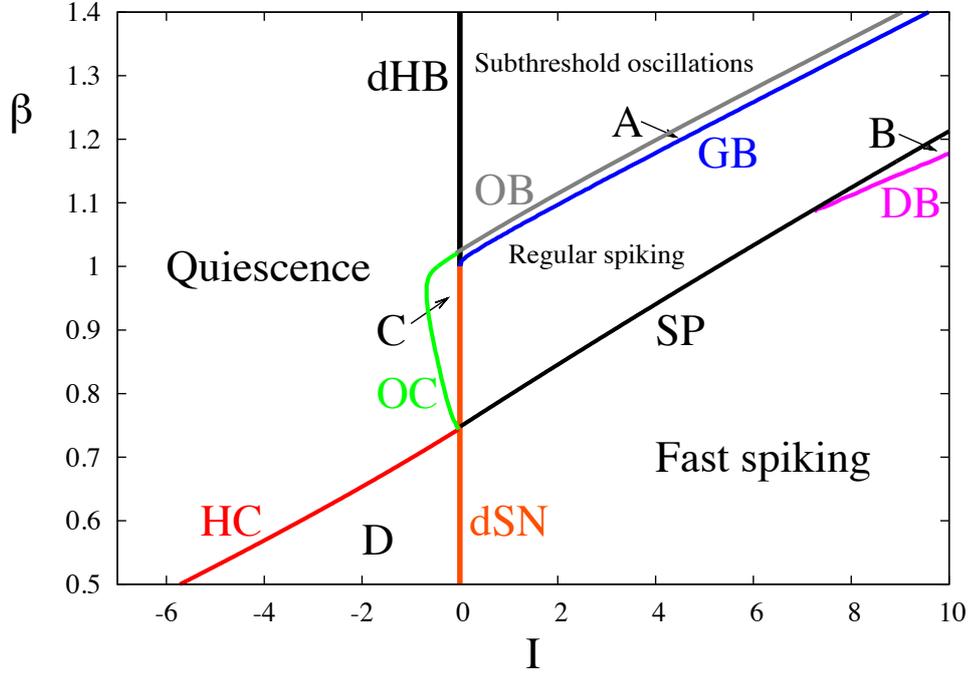


Figure 22: Bifurcation curves showing where solution types exchange stability in the (I, β) parameter plane. Other parameters are $\omega = 0.9$, $s = 0.35$, and $k = 0.4$. The dHB refers to the discontinuous Hopf bifurcation, dSN refers to the discontinuous saddle node bifurcation, GB is the grazing bifurcation between sub-threshold oscillations and regular spiking ones, SP is the bifurcation between the regular and fast spiking solutions, HC is the homoclinic bifurcation occurring when the fast spiking orbit collides with the saddle, OC is the orbit crisis, marking the loss of the regular spiking solution, OB is the bifurcation marking the onset/termination of bistability between sub-threshold oscillations and spiking ones, DB is the bifurcation marking the end of doublet firing, the onset of which occurs along the SP curve. Regions A,B,C,D correspond to bistable parameter regimes, the solutions of which are depicted in Fig. 25. Solution types in the other regimes are marked.

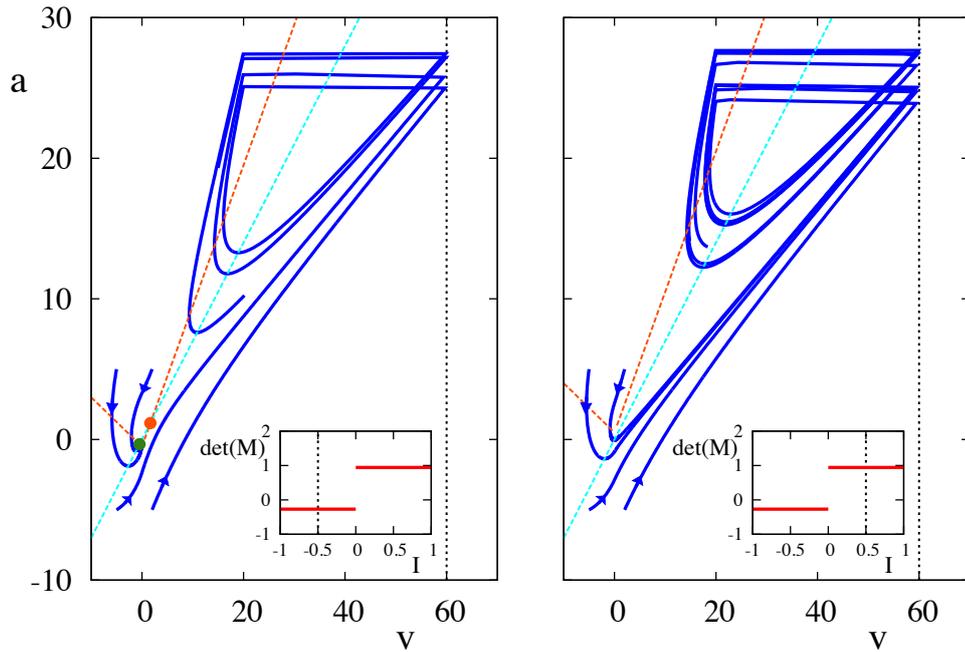


Figure 23: Phase plane and sample solutions just before (left) and just after (right) the dSN bifurcation. In the left panel, we set $I = -0.5$. We see that two fixed points exist, one on each of the v -branches, with the left one, depicted by a green circle, being stable, and the right one, shown by the orange circle, being a saddle. Since the system exhibits bistability between the fixed point and a fast spiking solution, trajectories will tend towards either of these attractors. For $I = 0.5$, in the right panel, there are no fixed points. Trajectories now tend towards the fast spiking solution, which is now the only attractor in the system. The inset shows the determinant of the Jacobian evaluated around the fixed point under smooth variation of I , with the vertical black lines indicating the value of I in each of the respective panels. We observe a discontinuity in the determinant as it passes through zero, at the bifurcation at $I = 0$. This is why we refer to this bifurcation as the dSN, in that it meets the conditions for a saddle-node bifurcation, but the determinant of the Jacobian is discontinuous at the bifurcation point. Also shown are the v - and a -nullclines in orange and cyan respectively. We set $\beta = 0.7$, with other parameters as in Fig. 22, so that this figure shows the transition across the dSN (orange) curve at $\beta = 0.7$.

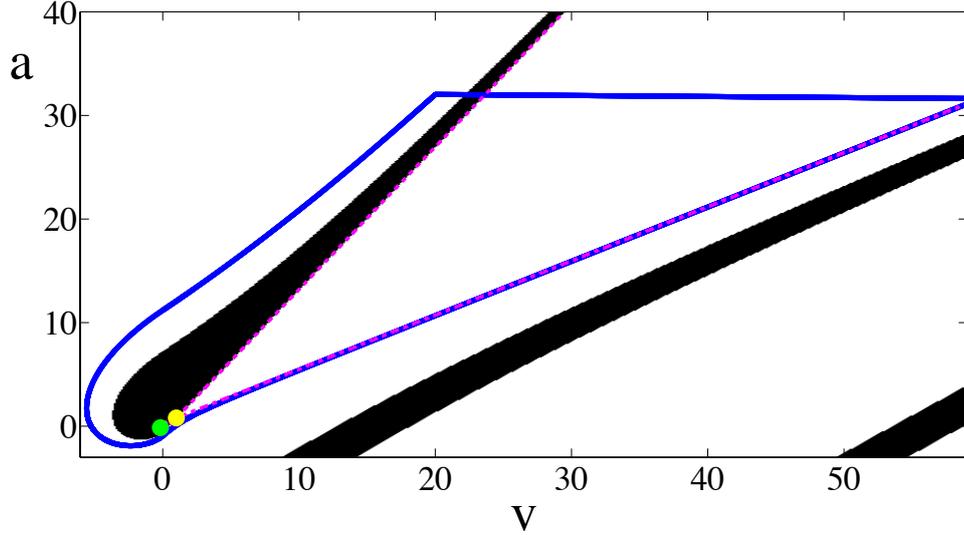


Figure 24: Basins of attraction for the stable fixed point and limit cycle for $\beta = 0.8$, $I = -0.2$. Other parameters are as in Fig. 22. We note that this solutions is equivalent to the one in panel C of Fig. 25. Black denotes the basin of attraction of the stable fixed point whereas white denotes the basin of attraction of the limit cycle. We see that both basins are the union of disconnected sets. The green and yellow circles depict, respectively, the stable fixed point and saddle whilst the purple dashed lines are the separatrices of the saddle node, given by the eigenvectors of the Jacobian there. The large amplitude limit cycle is lost at the point where it touches the basin of attraction of the stable fixed point.

As β decreases through $\beta_c = (v_{\text{th}} - I)/v_{\text{th}}$, the fixed point no longer exists and we see spiking solutions only, regardless of the values of the other parameters.

In parameter regimes where bursting orbits are stable, spikes are added when the a value after reset of the last spike of a bursting orbit crosses some value a_c , resulting in a grazing bifurcation. The graze either occurs at $v = 0$, when the fixed point of (3.9), with f as in (3.10) is to the right of v_R , or at $v = v_{\text{th}}$ if the fixed point is to the left of v_R . After this point, trajectories will be forced up to threshold, so that the orbit gains an additional spike. For the case where the graze occurs at $v = 0$, the value of a_c may be found by integrating backwards from $(v, a) = (0, I)$, the point at which $(v, \dot{v}) = (0, 0)$, a time T , such that $v(-T) = v_R$.

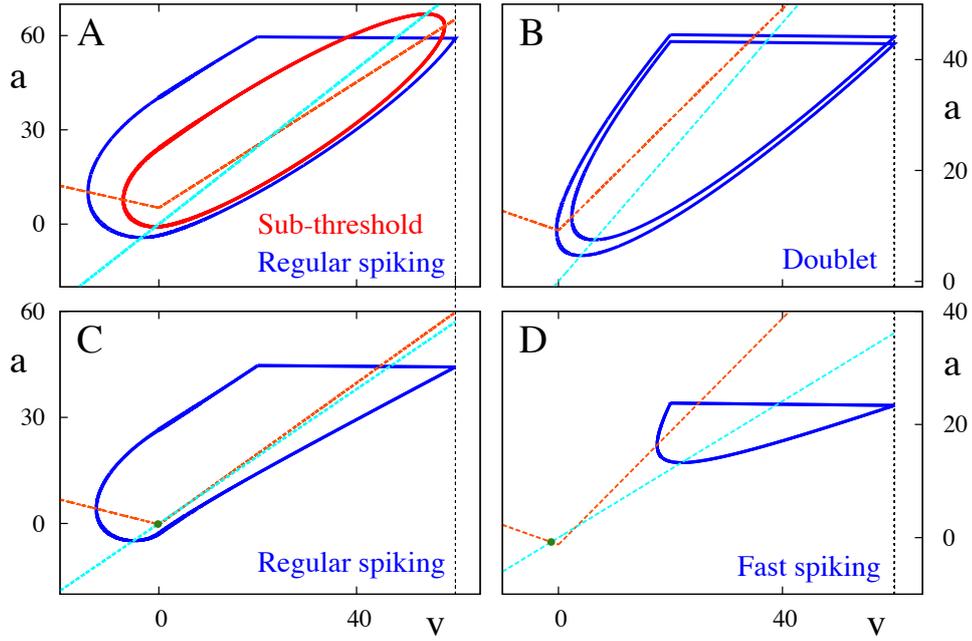


Figure 25: Solution types in the regions indicated in Fig. 22. The blue and red solid curves indicate the periodic solutions; all solutions are stable. The orange dashed lines are the branches of the v -nullcline, whilst the cyan dashed depicts the a -nullcline. The green circles in the lower two figures are stable fixed points.

The value of a_c is then equal to $a(-T)$. T is the flight time (in backwards time) from $v = 0$ to $v = v_R$ and may be found numerically. For the case where the graze occurs at $v = v_{th}$, the same method can be used, this time integrating from $(v, a) = (v_{th}, v_{th} + I)$. Interestingly, for bursting orbits, the value of a_c may also be found by finding the curves of inflection of the vector field. These curves separate trajectories that ‘bend’ rightwards, towards threshold, and those which ‘bend’ leftwards, towards the switching manifold, and are given by the solution to the equation $d^2a/dv^2 = 0$. Substituting $v = v_R$ in the resulting equation will give a_c . Following [105], we may calculate the curves of inflection for $v > 0$ as

$$a_{\pm}(v) = \frac{v(\beta\omega\sigma_{\pm} - 1) - I}{\omega\sigma_{\pm} - 1}, \quad (3.12)$$

with

$$\sigma_{\pm} = \frac{1 + \omega \pm \sqrt{(1 - \omega)^2 - 4\omega(\beta - 1)}}{2\beta\omega}. \quad (3.13)$$

Note that the inflection curves only exist if solutions (3.12) are real, which is the case if the eigenvalues (3.11) of the Jacobian of the system in the region $v > 0$ are real. We comment that, since $v_R > 0$, we need not calculate the inflection curves for $v < 0$. In fact, it turns out that the eigenvalues for $v < 0$ are complex, and so the inflection curves do not exist here. In the singular-limit, as $\omega \rightarrow 0$, the inflection curve for $v > 0$ is precisely the right branch of the v -nullcline. Upon substituting $v = v_R$ into (3.12), and by simulating trajectories, we find that the solution using σ_- is the relevant one, and hence

$$a_c = \frac{v_R (\beta \omega \sigma_- - 1) - I}{\omega \sigma_- - 1}. \quad (3.14)$$

To demonstrate the importance of this line for bursting trajectories, we plot the behaviour of two solutions with initial conditions near (v_R, a_c) in Fig. 26. The first of these solutions, in blue, has $a(t_0) < a_c$ and the resulting trajectory bends rightwards, toward threshold and another spike is fired. The second trajectory, in red, has $a(t_0) > a_c$, and the change in curvature causes the resulting trajectory to bend leftwards, away from threshold, forcing the neuron to enter the quiescent phase of the bursting. Spike adding bifurcations can occur as the value of a following the reset of the last spike in a burst falls below a_c , and the trajectory reaches threshold once more before falling into the quiescent phase.

The number of spikes in a burst is controlled by varying either ω , I , v_R or v_{th} . Decreasing any of these parameters will result in bifurcations that decrease the number of spikes in a burst. Where $v_R < 0$, the system is unable to burst as trajectories are always reset to the left of the right v -branch and are attracted to the left v -branch. We also note that we observe bursts for larger values of ω in the case where $\beta > \beta_c$ than where $\beta < \beta_c$, and that large values of I may prohibit bursting, so that we observe only fast spiking, so that I and β may be used as control parameters to switch between fast spiking and burst modes.

Owing to the discontinuous nature of the flow at reset, we may observe spiking orbits that enclose all other stable attractors, be they fixed points or sub-threshold oscillations. The emergence of such orbits is controlled by the parameter k . Where k is too small, trajectories will simply tend towards the attractors whose basin of

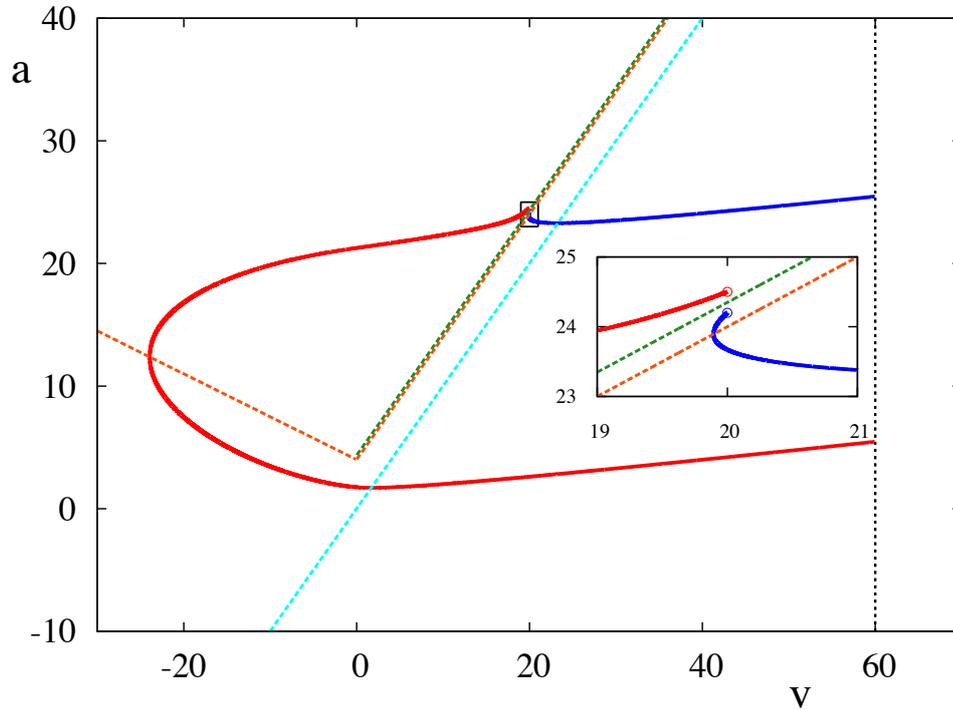


Figure 26: The green dashed line shows the inflection curve for $v > 0$, which is the curve along which $d^2a/dv^2 = 0$. This curve separates trajectories that ‘bend’ to the right, and those that ‘bend’ to the left. The blue trajectory has an initial condition with $a < a_c$, where a_c is the given by (3.14). This trajectory hits threshold and so a spike is fired before the neuron enters its quiescent phase. In contrast, the red trajectory has an initial condition with $a > a_c$ and does not reach threshold, entering the silent phase, where $v < 0$. The inset shows a close up of the region of interest, enclosed by the rectangle in the main figure, around the initial conditions, which are depicted by the open circles. For reference, the nullclines of the system are shown. Note that the inflection curve is near parallel to the v -nullcline, shown as a dashed orange line. In the singular-limit, as $\omega \rightarrow 0$, the inflection curve is the v -nullcline. Parameters are $\beta = 1, \omega = 0.08, s = 0.35, I = 4.0$.

attraction they are in. However, when k is large enough, we see the emergence of large amplitude limit cycles. These occur as the flows get ‘interrupted’ as they head towards an attractor in the sub-threshold system. All trajectories starting outside these limit cycles are in the basin of attraction of such orbits.

3.7 Solutions of the PWL-IF model

To solve the PWL-IF model, that is (3.9) with f given by (3.10), it is useful to recast the dynamics in matrix form so that:

$$\dot{X} = \begin{cases} A_1 X + \mu, & X_1 \geq 0, \\ A_2 X + \mu, & X_1 < 0, \end{cases} \quad (3.15)$$

where

$$A_1 = \begin{bmatrix} 1 & -1 \\ \omega\beta & -\omega \end{bmatrix}, \quad A_2 = \begin{bmatrix} -s & -1 \\ \omega\beta & -\omega \end{bmatrix}, \quad \mu = \begin{bmatrix} I \\ 0 \end{bmatrix}, \quad (3.16)$$

with X_i referring to the i 'th component of X (i.e. $X_1 = v$ and $X_2 = a$). The solution to the equation $\dot{X} = MX + \mu$ can be written using matrix exponentials in the form

$$X(t) = G(t)X(0) + K(t)\mu, \quad (3.17)$$

where

$$G(t) = e^{Mt}, \quad K(t) = \int_0^t G(s)ds. \quad (3.18)$$

Following [137], we may easily construct explicit solutions for G and K . For a 2×2 matrix M with real eigenvalues $\lambda_{+,-}$ the entries of the matrix G are given by

$$\begin{aligned} G_{11}(t) &= \frac{1}{\lambda_+ - \lambda_-} \{ \lambda_+ e^{\lambda_+ t} - \lambda_- e^{\lambda_- t} - m_{22} [e^{\lambda_+ t} - e^{\lambda_- t}] \}, \\ G_{12}(t) &= -\frac{\lambda_+ - m_{22} \lambda_- - m_{22}}{\lambda_+ - \lambda_-} \frac{m_{21}}{m_{21}} [e^{\lambda_+ t} - e^{\lambda_- t}], \\ G_{21}(t) &= \frac{m_{21}}{\lambda_+ - \lambda_-} [e^{\lambda_+ t} - e^{\lambda_- t}], \\ G_{22}(t) &= \frac{1}{\lambda_+ - \lambda_-} \{ \lambda_+ e^{\lambda_+ t} - \lambda_- e^{\lambda_- t} + m_{22} [e^{\lambda_+ t} - e^{\lambda_- t}] \}, \end{aligned} \quad (3.19)$$

whilst, for K

$$\begin{aligned}
K_{11}(t) &= \frac{1}{\lambda_+ - \lambda_-} \left\{ e^{\lambda_+ t} - e^{\lambda_- t} - m_{22} \left[\frac{e^{\lambda_+ t} - 1}{\lambda_+} - \frac{e^{\lambda_- t} - 1}{\lambda_-} \right] \right\}, \\
K_{12}(t) &= -\frac{\lambda_+ - m_{22}}{\lambda_+ - \lambda_-} \frac{\lambda_- - m_{22}}{m_{21}} \left[\frac{e^{\lambda_+ t} - 1}{\lambda_+} - \frac{e^{\lambda_- t} - 1}{\lambda_-} \right], \\
K_{21}(t) &= \frac{m_{21}}{\lambda_+ - \lambda_-} \left[\frac{e^{\lambda_+ t} - 1}{\lambda_+} - \frac{e^{\lambda_- t} - 1}{\lambda_-} \right], \\
K_{22}(t) &= \frac{1}{\lambda_+ - \lambda_-} \left\{ \frac{\lambda_+}{\lambda_-} [e^{\lambda_- t} - 1] - \frac{\lambda_-}{\lambda_+} [e^{\lambda_+ t} - 1] + m_{22} \left[\frac{e^{\lambda_+ t} - 1}{\lambda_+} - \frac{e^{\lambda_- t} - 1}{\lambda_-} \right] \right\}.
\end{aligned} \tag{3.20}$$

If M has complex, eigenvalues $\lambda_{\pm} = \nu \pm i\omega$, these matrices are instead given by

$$G(t) = \frac{e^{\nu t}}{\hat{\omega}} \begin{bmatrix} \hat{\omega} \cos \omega t - \hat{\nu} \sin \omega t & \sin \omega t \\ -(\hat{\nu}^2 + \hat{\omega}^2) \sin \omega t & \hat{\omega} \cos \omega t + \hat{\nu} \sin \omega t \end{bmatrix} \tag{3.21}$$

and

$$K(t) = \frac{1}{\hat{\omega}} \begin{bmatrix} \hat{\omega} K_R(t) - \hat{\nu} K_I(t) & K_I(t) \\ -(\hat{\nu}^2 + \hat{\omega}^2) K_I(t) & \hat{\omega} K_R(t) + \hat{\nu} K_I(t) \end{bmatrix} \tag{3.22}$$

where $\hat{\omega} = \omega/m_{12}$, $\hat{\nu} = (\rho - m_{11})/m_{12}$ and

$$K_R(t) = \frac{1}{\nu^2 + \omega^2} \left\{ \nu [e^{\nu t} \cos(\omega t) - 1] + \omega e^{\nu t} \sin(\omega t) \right\}, \tag{3.23}$$

$$K_I(t) = \frac{1}{\nu^2 + \omega^2} \left\{ \omega [1 - e^{\nu t} \cos(\omega t)] + \nu e^{\nu t} \sin(\omega t) \right\}. \tag{3.24}$$

Hereafter, we refer to G^i, K^i as the above expressions with the respective matrix $M = A_i$. To find a fast spiking orbit of period Δ (in response to constant forcing), we need only solve $(X_1(\Delta), X_2(\Delta)) = (v_{\text{th}}, a_0 - k)$ subject to $(X_1(0), X_2(0)) = (v_{\text{R}}, a_0)$, which gives a pair of simultaneous equations for (Δ, a_0) as:

$$v_{\text{th}} = G_{11}^1(\Delta) v_{\text{R}} + G_{12}^1(\Delta) a_0 + K_{11}^1(\Delta) I, \tag{3.25}$$

$$a_0 = \frac{G_{21}^1(\Delta) v_{\text{R}} + K_{21}^1(\Delta) I + k}{1 - G_{22}^1(\Delta)}, \tag{3.26}$$

which may be solved numerically. Bursting orbits may be constructed using similar ideas, though with more book-keeping to keep track of the sub-trajectories (each determined by a linear system) that build the full periodic orbit. For example, for an orbit with ‘times-of-flight’ T_i , $i = 1, \dots, N$, (defined by the time spent in

a region of phase space before meeting $v = 0$ or $v = v_{\text{th}}$) describing a bursting orbit with $N - 2$ spikes, we have to solve for the unknowns (T_1, \dots, T_N, a_0) using a system of equations of the form

$$\begin{aligned}
\begin{bmatrix} 0 \\ a_1 \end{bmatrix} &= G^1(T_1) \begin{bmatrix} v_{\text{R}} \\ a_0 \end{bmatrix} + K^1(T_1)\mu, \\
\begin{bmatrix} 0 \\ a_2 \end{bmatrix} &= G^2(T_2) \begin{bmatrix} 0 \\ a_1 \end{bmatrix} + K^2(T_2)\mu, \\
\begin{bmatrix} v_{\text{th}} \\ a_3 \end{bmatrix} &= G^1(T_3) \begin{bmatrix} 0 \\ a_2 \end{bmatrix} + K^1(T_3)\mu, \\
&\vdots \\
\begin{bmatrix} v_{\text{th}} \\ a_n \end{bmatrix} &= G^1(T_n) \begin{bmatrix} v_{\text{R}} \\ a_{n-1} + k \end{bmatrix} + K^1(T_n)\mu,
\end{aligned} \tag{3.27}$$

for $n = 4, \dots, N$ subject to $a_0 = a_N + k$. The period of the orbit is simply $\Delta = \sum_{i=1}^N T_i$. From here onwards, we shall denote periodic orbits by $z(t)$ where $z(0) = [v_{\text{R}}, a_0]$.

3.8 Numerical integration of planar IF models

Amongst the standard techniques to numerically integrate systems of the form $\dot{x} = f(x)$, one of the most common is the Runge-Kutta scheme; a multistage predictor-corrector scheme that involves subdividing the time step and at each stage refining the estimated solution. Such algorithms are included, for example, in the standard suite of ODE solvers provided by MATLAB. Of these, the fourth order scheme is popular amongst modellers for its speed and accuracy, for smooth systems, of $O(h^4)$, where h is the stepsize used by the numerical routine. For IF systems, Hansel showed in [138] that, owing to the error accrued at reset, the fourth order scheme is no more accurate than its equivalent second order scheme, and thus for numerical integration, a second order scheme should be used given its increased speed compared to the fourth order scheme. To detect threshold crossing, we integrate until $v_n > v_{\text{th}}$ for some n in the numerical solution. At this

point, we use linear interpolation to approximate both the time and the value of a at the threshold crossing. Supposing that we have $v_n > v_{\text{th}}$ at t_n , but $v_{n-1} < v_{\text{th}}$ at t_{n-1} . We can then estimate the time, t^* of the threshold crossing by interpolating between these two values to give

$$t^* = t_{n-1} + h \frac{v_{\text{th}} - v_{n-1}}{v_n - v_{n-1}}. \quad (3.28)$$

We may use the values (a_{n-1}, a_n) to estimate the value of a^* at the threshold crossing as

$$a^* = a_{n-1} + \frac{v_{\text{th}} - v_{n-1}}{v_n - v_{n-1}}(a_n - a_{n-1}). \quad (3.29)$$

Since (3.15) is linear, its solutions will cross threshold exponentially, and so we will make errors by using linear interpolation to estimate the threshold crossing time. However, we may use the transformation $u = \log(v + 1)$ and solve the new system

$$\frac{du}{dt} = 1 + (-1 + I - a)e^{-u}, \quad (3.30)$$

$$\frac{1}{\omega} \frac{da}{dt} = \beta(e^u - 1) - a, \quad (3.31)$$

whenever $v > 0$, leaving the system for $v < 0$ unchanged. The v component of these solutions now crosses threshold linearly, so that linear interpolation can be applied with more accuracy. We note that we could also transform a in this way, but since the v dynamics are faster than the a dynamics, we did not feel that there is much of an error accrued by leaving a untransformed.

Of course, the advantage of the piecewise-linear regime is that we have solutions in closed form. Thus, instead of resorting to any form of numerical scheme, we may simply use symbolic dynamics to denote the sign of v and simply solve for the times-of-flight, to either the switching manifold, if the start point has $v < 0$, or threshold, if the start point has $v > 0$, ensuring that the flight times are never zero. Since this step only involves an implementation of a root finding algorithm, we can find solutions rapidly and with accuracy up to that of the machine.

3.9 Subthreshold orbits

We may use the eigenvalues computed earlier as (3.11) to find the criticality of the discontinuous Hopf. For this kind of Hopf, the criticality is given by linear terms rather than the cubic ones in the smooth case. In fact, from [136], the criticality is given by

$$\Lambda = \frac{(1 - \omega)}{\sqrt{4\omega(\beta - 1) - (1 - \omega)^2}} - \frac{s + \omega}{\sqrt{4\omega(\beta + s) - (s + \omega)^2}}. \quad (3.32)$$

When $\Lambda < 0$, the bifurcation is supercritical, and we can expect stable subthreshold solutions for $I > 0$, in the absence of a threshold. Conversely, if $\Lambda > 0$, the bifurcation is subcritical and we no longer have periodic solutions for $I > 0$, but repelling periodic orbits exist for $I < 0$. At $\Lambda = 0$, there is a codimension two bifurcation, analogous to the Bautin bifurcation. We may solve (3.32) for β to find the equation of a surface separating the sub- from the supercritical case, as

$$\beta = \frac{\omega^2 + s}{s + 2\omega - 1}. \quad (3.33)$$

Of course, if I is too large, we may not observe subthreshold oscillations, since the orbits may cross threshold and we will instead observe spiking solutions. Similarly, whilst we may not observe subthreshold orbits in the case where $\Lambda > 0$, we will see spiking orbits, since the orbits will tend to infinity and will thus cross threshold. We show in Fig. 27 the criticality surface for the discontinuous Hopf bifurcation. The value of β is plotted on the z -axis. Below this criticality surface, we have $\Lambda > 0$, and so the system does not support subthreshold oscillations, whereas we do expect them for parameter values above the surface. It is clear that it is harder for the system to support subthreshold orbits as ω and s decrease, meaning that the neuron is more likely to fire spikes in regions with low ω and s .

Finding subthreshold orbits for the PWL-IF system follows the same approach as finding spiking orbits, except this time using $v = 0$ rather than $v = v_{\text{th}}$ to piece together solutions. The subthreshold orbits are fully specified by the identification of the times-of-flight: $T_{1,2}$ in $v > 0$ and $v < 0$ respectively, and a start point $X(t_0) = [0, a_1]$ on the switching manifold. It may be shown that $T_{1,2}$ may be

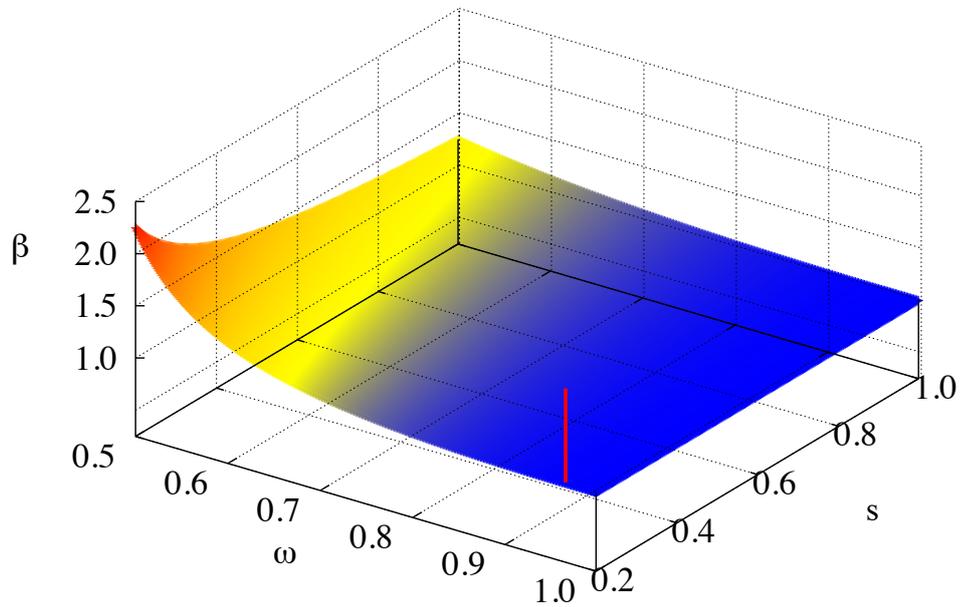


Figure 27: Criticality surface for the discontinuous Hopf bifurcation. This surface separates the criticality of the discontinuous Hopf bifurcation at $I = 0$, with β plotted on the z -axis against ω and s . Below the surface, subthreshold orbits are not supported, whilst we may observe them for parameter values above the surface. Whether or not subthreshold orbits are actually supported will depend on the threshold value, since it is not explicitly included here. We see that, as s and ω decrease, the β values required for subthreshold oscillations increases, indicating that the neuron is more likely to fire spikes in these regions. The red line indicates where in (s, ω) space we evaluate the bifurcation curves in Fig. 22. In particular, the red line is the line $I = 0$ (with $\beta > 1$) in said figure.

found as the solution to the system of simultaneous equations

$$0 = -G_{12}^2(T_2)G_{22}^1(T_1)K_{11}^1(T_1) + G_{12}^1(T_1)G_{12}^2(T_2)K_{21}^1(T_1) + G_{12}^1(T_1)K_{11}^2(T_2), \quad (3.34)$$

$$\frac{I(G_{22}^2(T_2)K_{21}^1(T_1) + K_{21}^2(T_2))}{1 - G_{22}^2(T_2)G_{22}^1(T_1)} = a_1 = -\frac{K_{21}^1(T_1)I}{G_{12}^1(T_1)}. \quad (3.35)$$

To check the stability of the subthreshold oscillations born out of the Hopf bifurcations, we may construct the Floquet multipliers $\gamma_{1,2}$ of the orbit. Following the work in [137], we use the fact that $\gamma_1\gamma_2 = \exp\left(\int_0^t \text{Tr}(Df(s))ds\right)$ (where DF denotes the Jacobian of f) and that one of the multipliers is equal to one. For the PWL-IF model, $DF \in \{A_1, A_2\}$. We thus have that the Floquet exponents, $\sigma_{1,2} = \ln(\gamma_{1,2})$, are given by $(\sigma_1, \sigma_2) = (0, \sigma)$ where σ is given by:

$$\begin{aligned} \sigma &= \frac{1}{T} \left(\int_0^{T_1} \text{Tr} A_1 ds + \int_0^{T_2} \text{Tr} A_2 ds \right), \\ &= \frac{1}{T} (T_1 \text{Tr} A_1 + T_2 \text{Tr} A_2), \\ &= \frac{T_1(1 - \omega) + T_2(-s - \omega)}{\Delta}, \\ &= \frac{T_1 - sT_2}{\Delta} - \omega, \\ &= 1 - \frac{(1 + s)T_2}{\Delta} - \omega, \end{aligned} \quad (3.36)$$

remembering that $T_1 + T_2 = \Delta$. Requiring that $\sigma < 0$ results in the following stability condition on T_2 :

$$T_2 > \frac{1 - \omega}{1 + s} \Delta, \quad (3.37)$$

which can be checked easily after solving (3.34)-(3.35). We note that the dependence of the Floquet multiplier on parameters other than s, ω is wrapped up in the definition of T_2 . Now that we have found explicit subthreshold solutions of (3.15), we can show that the period of subthreshold orbits is independent of the drive I . To start, let us define a_1 and a_2 to be the distinct values of a when the orbit crosses $v = 0$, with $a_1 < a_2 = G_{22}^1(T_1)a_1 + IK_{21}^1(T_1)$. We have that

$$[0, a_2]^T = G^1(T_1)[0, a_1]^T + IK^1(T_1)[1, 0]^T, \quad (3.38)$$

$$[0, a_1]^T = G^2(T_2)[0, a_2]^T + IK^2(T_2)[1, 0]^T. \quad (3.39)$$

Dividing these equations by I gives:

$$\begin{aligned}\frac{a_1}{I} &= F_1^1(T_1), & \frac{a_2}{I} &= F_2^1(T_1), \\ \frac{a_2}{I} &= F_1^2(T_2), & \frac{a_1}{I} &= F_2^2(T_2).\end{aligned}$$

where

$$F_1^i(T_i) = -\frac{K_{11}^i(T_i)}{G_{21}^i(T_i)}, \quad F_2^i(T_i) = \frac{G_{22}^{j \neq i}(T_{j \neq i})}{I} + K_{21}^{j \neq i}(T_{j \neq i}), \quad i, j \in \{1, 2\}. \quad (3.40)$$

Equating terms gives:

$$F_1^1(T_1) = F_1^2(T_2), \quad (3.41)$$

$$F_2^1(T_1) = F_2^2(T_2). \quad (3.42)$$

Since both these equations are independent of I , both T_1 and T_2 are independent of I and so is their sum. In Fig. 28, we plot the orbits of the system borne out of the discontinuous Hopf bifurcation, with I as the control parameter. The plot clearly shows that the amplitude, as expected, grows linearly with I whilst the frequency of the oscillations is constant. This is in contrast to the situation shown in Fig. 29, in which ω is the control parameter, and the oscillations emanate from a smooth Hopf bifurcation from a focus on the right-branch of the v -nullcline. In this case, there exists a value of ω for which the frequency is highest, whilst the amplitude of oscillations continues to grow in an exponential fashion with increasing ω .

As these subthreshold orbits continue to grow, they will eventually graze against the threshold, and will thus cease to exist. It is interesting to note that, prior to this point, the system supports bistability between a regular spiking solution and the subthreshold oscillation, so that perturbations may switch trajectories between the basins of attractions of these solutions. Fig. 30 depicts this scenario, with the left panel showing the oscillation just before and the right panel showing the spiking solution just after the grazing bifurcation, with solutions of the bistable window represented in the middle panel.

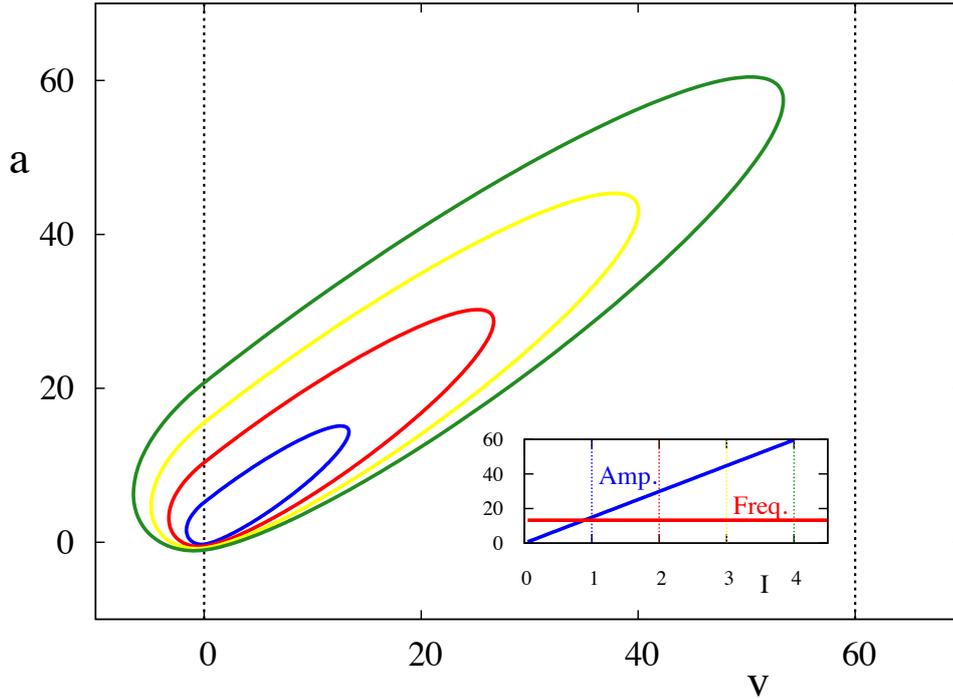


Figure 28: Subthreshold orbits for the PWLIF system. The amplitude (blue) and frequency (red) are plotted against the control parameter I in the inset. We clearly observe that the amplitude grows linearly with I , whilst the frequency does not vary. We may think of the subthreshold orbits as being two spiral sections, joined at the switching manifold $v = 0$. Parameter values are $\omega = 0.9, \beta = 1.2, s = 0.35, I = 1, 2, 3, 4$. The vertical lines in the inset graph correspond to the location (in parameter space) of the depicted orbits. For reference, the threshold and switching manifolds are also shown in the main figure as a vertical dashed line.

3.10 Phase response curves

Unlike the LIF model, which can be fully characterised by a phase variable, the PWL-IF is planar and so a phase only description will not be sufficient to describe the behaviour of the full system. However, we may still wish to study the effect of weak forcing, and perform a phase reduction of the type considered in Sect. 2.5. Here, we will solve the adjoint equation for the PWL-IF model, which gives the iPRC, after a trivial rescaling. Recapping from Sect. 2.5, the adjoint equation is

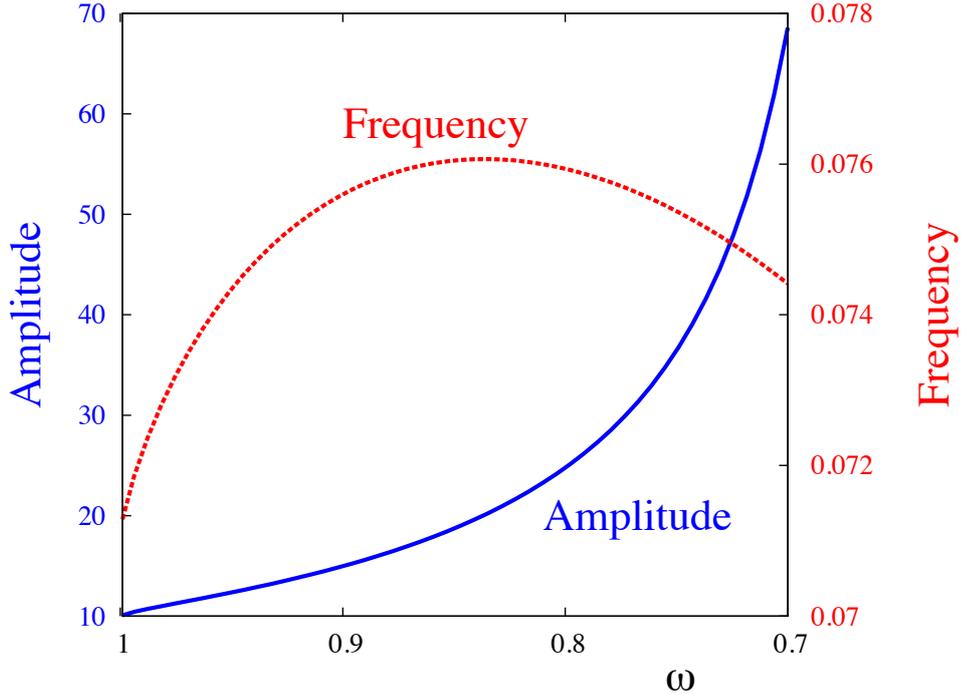


Figure 29: A plot of the amplitude (blue solid) and frequency (red dashed) of a sub-threshold limit cycle born from a Hopf bifurcation. The bifurcation, at $\omega = 1$, in this case is not a discontinuous one as the Jacobian is continuous at the bifurcation point. We clearly see that the frequency of the oscillations varies with the control parameter, ω and that the amplitude of the oscillations does not grow linearly with ω , in contrast to the discontinuous Hopf using I as the control parameter. Other parameter values are $\beta = 1.2$, $I = 4$ and $s = 0.35$.

given by

$$\frac{dQ}{dt} = -DF^T(Z(t))Q, \quad (3.43)$$

subject to the conditions $Q^T(0)F(z(0)) = 1$ and $Q(t) = Q(t + \Delta)$. The first condition simply guarantees that $\dot{\theta} = 1$ (at any point on the periodic orbit), and the second enforces continuity (and periodicity). The (vector) iPRC, R , is related to Q according to the scaling $R = Q\Delta$. In general, (3.43) must be solved numerically to obtain the iPRC, say, using the *adjoint* routine in XPP [139]. However, for PWL models $DF(Z)$ is piecewise constant, and we can obtain a solution in closed form [104]. Moreover, it is also possible to extend the Malkin method to treat an IF

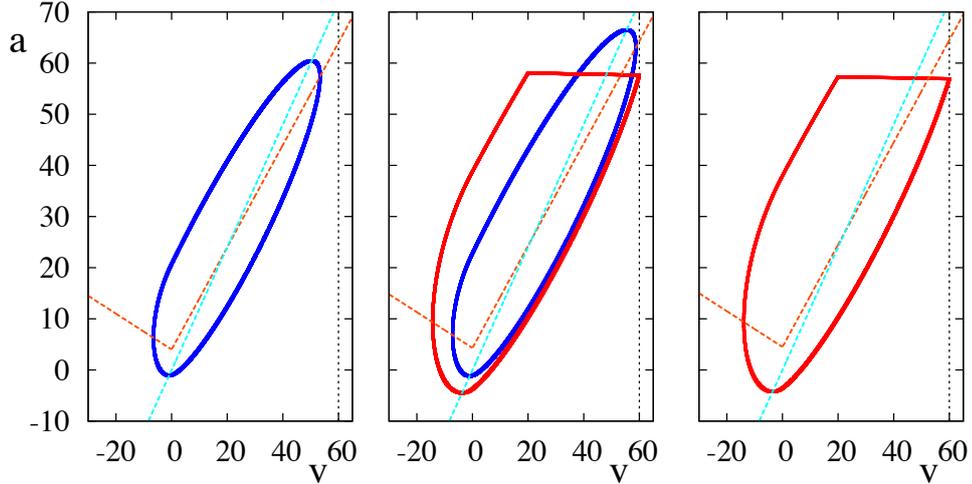


Figure 30: The transition from a subthreshold to a spiking orbit. In the left panel, we observe only a subthreshold orbit, shown in blue, for $I = 4.0$. Moving to the middle panel, we now observe bistability between a subthreshold orbit and a spiking orbit, shown in red, for $I = 4.5$. Finally, the subthreshold orbit is lost via a grazing bifurcation as it tangentially touches threshold, and so we are left with only a spiking orbit for $I = 5.0$. Other parameter values are as in Fig. 28. Also shown are the nullclines of the system. These transitions can be seen by moving through Fig. 22 horizontally, at $\beta = 1.2$ through the grey (OB) and blue (GB) bifurcation curves.

process [135], which would give rise to a discontinuous iPRC (at the spike time). In this latter case, the continuity condition is swapped in favour of enforcing the normalisation condition $Q^T(t)F(z(t)) = 1$ for all t , so that

$$Q^T(T_n)F(z(T_n)) = 1, \quad \forall n. \quad (3.44)$$

For the PWL-IF model, we may construct Q in given regions of phase space according to the prescription $Q(t) = J(T_i - t)Q(T_i)$, where $J = G^T$ [104]. Enforcing the normalisation condition at the times T_i is enough to define a periodic (yet discontinuous) form for Q . For example, for a simple tonic spiking orbit we see that solving (3.43) and imposing the normalisation condition at $t = 0$ and $t = \Delta$ gives a system of two linear equations in (q_1, q_2) , where q_i are the components of

Q as

$$\begin{aligned} q_1(\Delta)(v_{\text{th}} + I - a_0 + k) + q_2(\Delta)(\omega(\beta v_{\text{th}} - a_0 + k)) &= 1, \\ q_1(0)(v_{\text{r}} + I - a_0) + q_2(0)(\omega(\beta v_{\text{r}} - a_0)) &= 1. \end{aligned} \quad (3.45)$$

Using the further result that $Q(0) = \Gamma Q(T)$ where $\Gamma = J^1(\Delta)$ for fast spiking orbits or $\Gamma = J^1(T_3)J^2(T_2)J^1(T_1)$ for regular spiking orbits, gives

$$\begin{aligned} q_2(\Delta) &= \frac{r_1 - r_2\Gamma_{11} - r_4\Gamma_{21}}{r_1(\Gamma_{12}r_2 + r_4\Gamma_{22}) - (r_3r_2\Gamma_{11} + r_3r_4\Gamma_{21})}, \\ q_1(\Delta) &= \frac{1}{r_1}(1 - r_3q_2(\Delta)), \end{aligned} \quad (3.46)$$

where

$$\begin{aligned} r_1 &= v_{\text{th}} + I - a_0 + k, & r_2 &= v_{\text{R}} + I - a_0, \\ r_3 &= \omega(v_{\text{th}} - a_0 + k), & r_4 &= \omega(v_{\text{R}} - a_0). \end{aligned} \quad (3.47)$$

Hence, for a fast spiking orbit, the adjoint is given by $Q(t) = J(\Delta - t)Q(\Delta)$ and for a regular spiking orbit the corresponding Q is

$$Q(t) = \begin{cases} J^1(T_1 - t)J^2(T_2)J^1(T_3)Q(\Delta), & 0 \leq t \leq t_1 \\ J^2(T_2 - (t - t_1))J^1(T_3)Q(\Delta), & t_1 \leq t \leq t_2, \\ J^1(T_3 - (t - t_2))Q(\Delta), & t_2 \leq t \leq \Delta \end{cases} \quad (3.48)$$

where $t_j = \sum_{i=1}^j T_i$. In both cases, the form of $Q(\Delta)$ is given by (3.46). iPRCs for bursting solutions may be constructed in the same way, except that discontinuities are now not isolated to the ends of the periodic orbit, and so we must enforce both the normalisation condition just before and just after each threshold crossing. For subthreshold orbits, with T_1, T_2 being the times of flight in $v > 0$ and $v < 0$ respectively, Q takes the form

$$Q(t) = \begin{cases} J^1(T_1 - t)J^2(T_2)Q(\Delta), & 0 \leq t \leq t_1 \\ J^2(T_2 - (t - t_1))Q(\Delta), & t_1 \leq t \leq \Delta \end{cases}, \quad (3.49)$$

To find $Q(\Delta)$ for these orbits, we now need to enforce the periodicity condition, as the orbits themselves do not possess a discontinuity. Setting $\Gamma = J^2(T_2)J^1(T_1)$, the periodicity condition takes the form

$$(\Gamma_{11} - 1)q_1(\Delta) + \Gamma_{12}q_2(\Delta) = 0, \quad (3.50)$$

and we now only need to enforce the first normalisation condition in (3.46). This normalisation condition and (3.50) define a pair of linear equations:

$$\Psi \begin{bmatrix} q_1(\Delta) \\ q_2(\Delta) \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \Psi = \begin{bmatrix} I - a_1 & -\omega a_1 \\ \Gamma_{11} - 1 & \Gamma_{12} \end{bmatrix}. \quad (3.51)$$

This system can be solved using Cramer's rule, giving $q_i(\Delta) = \det \Psi_i / \Psi$ where

$$\Psi_1 = \begin{bmatrix} 1 & -\omega a_1 \\ 0 & \Gamma_{12} \end{bmatrix}, \quad \Psi_2 = \begin{bmatrix} I - a_1 & 1 \\ \Gamma_{11} - 1 & 0 \end{bmatrix}. \quad (3.52)$$

Typically, when studying neural oscillators, we are primarily concerned with the first (voltage) component of Q , since perturbations to the system are usually given by changes in the external current, which acts only on the voltage variable. As an example, we plot in Fig. 31 the voltage component of Q for a regular spiking orbit and a burst containing three spikes, whilst in Fig. 32, we show this for a subthreshold orbit. Knowledge of the iPRC is fundamental in building network descriptions of weakly coupled oscillators [140].

3.11 Weakly coupled networks

As an example, let us consider a network of two identical, interacting cells. We shall consider gap junction coupling between the two cells. We remark that, although the LIF model is not usually well suited to studying this form of coupling, since it does not intrinsically include a spike shape, the PWL-IF model does include a representation of a spike and so we may consider gap junction coupling. Through this channel, currents are induced from one cell to another as ions move down electrochemical gradients. We will consider the case where the channel is always open, and will assume, in this case that the conductance of the channel $\varepsilon \ll 1$, so that we may apply weakly coupled oscillator theory. For simplicity, we will only consider the fast spiking case, although it is a straightforward extension to consider regular spiking or bursting orbits. For ease of notation, in this section

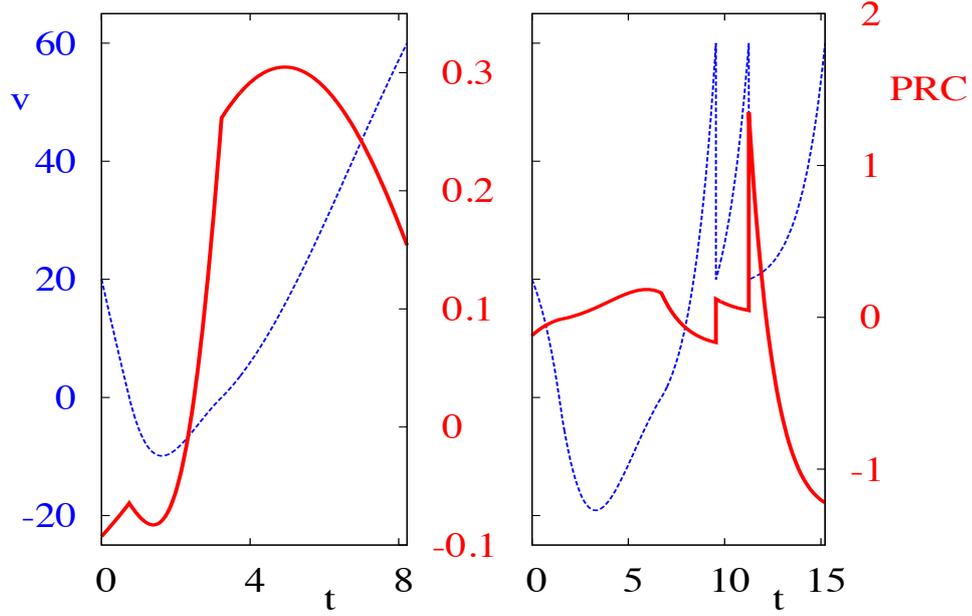


Figure 31: Left: Voltage component of the phase response curve for a regular spiking orbit (red, solid). Right: Voltage component of the phase response curve for a 3-spike bursting orbits (red, solid). Parameter values are $\beta = 1.1$, $s = 0.35$, $k = 0.4$, $I = 4.0$ and $\omega = 1$ for the regular spiking orbit and $\omega = 0.25$ for the bursting orbit. Corresponding orbits are shown in dashed blue.

we take $G = G^1$ and $J = J^1$ wherever they are used. In this case, the network dynamics may be written as

$$\frac{dv_i}{dt} = f(v_i) + I - a_i + \varepsilon(v_j - v_i), \quad (3.53)$$

$$\frac{1}{\omega} \frac{da_i}{dt} = \beta v_i - a_i, \quad j \neq i, \quad i, j \in \{1, 2\}. \quad (3.54)$$

Using the iPRC as defined in the previous section, we may rewrite the dynamics of the network as dynamics on a torus $\mathbb{S}^1 \times \mathbb{S}^1$, where $\theta_i = t \bmod \Delta$, $i = 1, 2$ as [141]

$$\frac{d\theta_i}{dt} = 1 + \varepsilon R(\theta_i) C(\theta_j, \theta_i), \quad (3.55)$$

where C describes the coupling between the cells. For gap junction coupling, it takes the form $C(\theta_j, \theta_i) = v_0(\theta_j) - v_0(\theta_i)$, where v_0 is the voltage component of the unperturbed oscillation. We may move to a rotating frame by introducing the

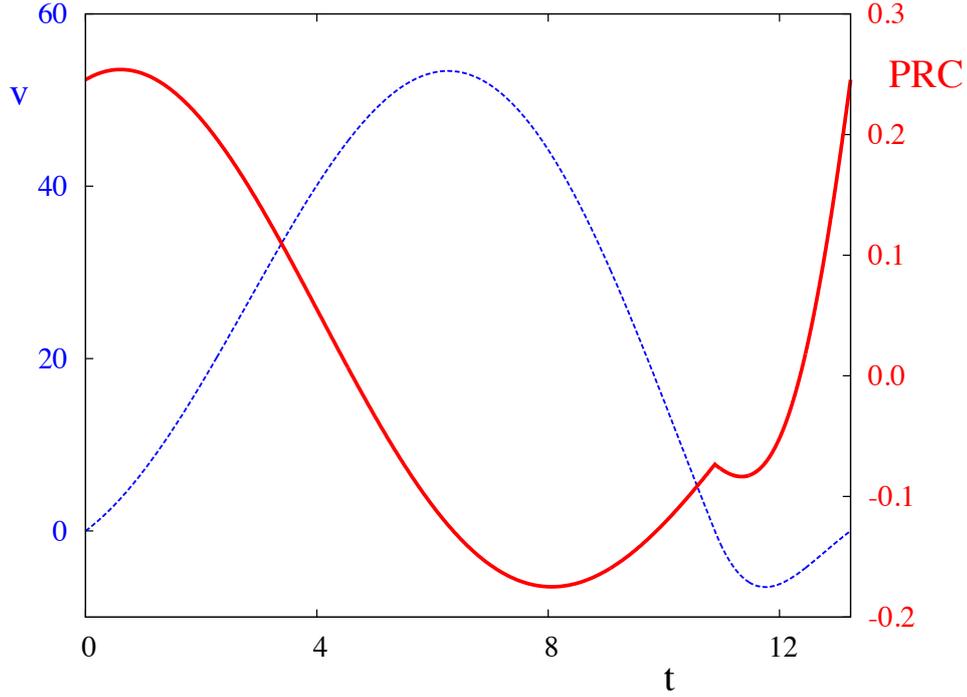


Figure 32: Voltage component of the phase response curve for a subthreshold orbit. In particular, the subthreshold orbit is the green one depicted in Fig. 28. The corresponding orbit is shown in dashed blue.

coordinate $\theta_i = t + \phi_i(t)$, $\phi_i \in [0, \Delta)$, resulting in

$$\frac{d\phi_i}{dt} = \varepsilon R(\phi_i + t)C(\phi_j + t, \phi_i + t). \quad (3.56)$$

For weak coupling, each neuron still fires approximately at the unperturbed rate, but ϕ_i slowly drifts. Since $\varepsilon \ll 1$, we can take advantage of averaging theory, taking the average of this drift over one period to eliminate the explicit time dependence, and arrive at an equation for the phase alone as

$$\frac{d\phi_i}{dt} = \frac{\varepsilon}{\Delta} \int_0^\Delta R(t)C(t + (\phi_j - \phi_i), t)dt = \varepsilon H(\phi_j - \phi_i). \quad (3.57)$$

We refer to H as the phase interaction function [5], and note that it depends only on the phase difference $\psi = \phi_j - \phi_i$ between the two cells. This means we can exploit the form of H to tell us about the existence and stability of solutions for the network. Specifically, the solutions of $H_{\text{odd}} = 1/2[H(\psi) - H(-\psi)] = 0$

give the fixed points for the evolution of ψ . For symmetrically coupled cells, the existence of both a synchronous solution, with $\psi = 0$, and the anti-phase solution, with $\psi = 1/2$ is guaranteed. For any solution for ψ , the sign of the derivative of $H_{\text{odd}}(\psi)$ at these solutions is the sole piece of information we need to classify its stability: if the sign of H'_{odd} is positive, the solution is unstable, whilst we have stable solutions if the sign of H'_{odd} is negative. Upon making the substitution $R = Q\Delta$, the phase interaction function, $H(\psi)$ may be written as

$$H(\psi) = \int_0^\Delta Q(t)C(t + \psi, t)dt. \quad (3.58)$$

We must take care, when computing the phase interaction function to handle the discontinuity at $v = v_{\text{th}}$, or equivalently at $t = \Delta$. As such, we decompose H as

$$H(\psi) = \int_0^{\Delta-\psi} Q(t)C(t + \psi, t)dt + \int_0^\psi Q(t + \Delta - \psi)C(t, t + \Delta - \psi)dt. \quad (3.59)$$

It is convenient to first suppose that $C(t + \psi) = X(t + \psi) - X(t)$ and simply take the first component of this function. For a fast spiking solution, we may write $Q(t)C(t + \psi, t)$, as

$$Q(t)C(t + \psi, t) = J(\Delta - t)Q(\Delta) \left(G^1(t + \psi)X_0 + K(t + \psi)\mu - G^1(t)X_0 - K(t)\mu \right), \quad (3.60)$$

where $X_0 = (v_{\text{R}}, a_0)$. Since multiples of the matrix A commute with one another, we note that

$$G(t + \psi) - G(t) = G(t)G(\psi) - G(t) = G(t)[G(\psi) - \mathbb{I}_2], \quad (3.61)$$

where \mathbb{I}_2 is the identity matrix in \mathbb{R}^2 . We also note that

$$\begin{aligned} K(t + \psi) &= \int_0^{t+\psi} G(s)ds \\ &= \int_0^t G(s)ds + \int_t^{t+\psi} G(s)ds \\ &= K(t) + \int_0^\psi G(s + t)ds. \end{aligned} \quad (3.62)$$

We let

$$\tilde{K}(t, \psi) = \int_0^\psi G(s + t)ds. \quad (3.63)$$

Using these relations, and remembering that $J = G^T$, we may rewrite (3.60)

$$Q(t)P(t + \psi, t) = G^T(\Delta - t)Q(\Delta) \left(G(t)[G(\psi) - \mathbb{I}_2]X_0 + \tilde{K}(t, \psi)\mu \right). \quad (3.64)$$

Since the forcing is in the voltage component only, we are only interested in the first component of H . Thus, we require the first component of

$$\begin{aligned} H(\psi) &= \int_0^{\Delta - \psi} \left(G^T(\Delta)G^T(-t)Q(\Delta) \right) \bullet \left(G(t)[G(\psi) - \mathbb{I}_2]X_0 + \tilde{K}(t, \psi)\mu \right) dt \\ &\quad + \int_0^\psi \left(G^T(\psi)G^T(-t)Q(\Delta) \right) \bullet \left(G(t)[\mathbb{I}_2 - G(\Delta)G(-\psi)]X_0 - \tilde{K}(t, \Delta - \psi)\mu \right) dt, \end{aligned} \quad (3.65)$$

where \bullet denotes elementwise multiplication. For regular spiking solutions, a similar result holds, taking care to piece together solutions across the switching manifold. For bursting solutions, we must take into account the multiple threshold crossings, and split the integral for H into n distinct pieces, where n is the number of spikes in the burst.

To demonstrate the use of the phase interaction function, we now show the results for a network of two identical neurons. In Fig. 33, we consider the case where the coupling is through gap junctions and the neurons, in the absence of inputs, are in a fast spiking regime. We observe that the odd part of H , shown in red, has fixed points at $\psi = 0$ and $\psi = \Delta/2$ as expected. The fixed point at 0 is unstable, whilst the fixed point at $\Delta/2$ is stable, having $H'_{\text{odd}} < 0$. As such, we expect to see the anti-phase solution persist after transients have decayed. This is precisely what is observed in the computed trajectories, shown in the bottom two panels of this figure. In another case with gap junction coupling, this time where the cells are in a bursting regime (with 3 spikes), we see that the phase interaction predicts bistability of the synchronous solution and the anti-phase solution, separated by an unstable fixed point with $\psi \approx 1$. For a network solution with an initial phase difference of $\psi(0) < 1$, we thus expect to see synchrony between the two cells as $t \rightarrow \infty$, as is observed in the trajectories in Fig. 34.

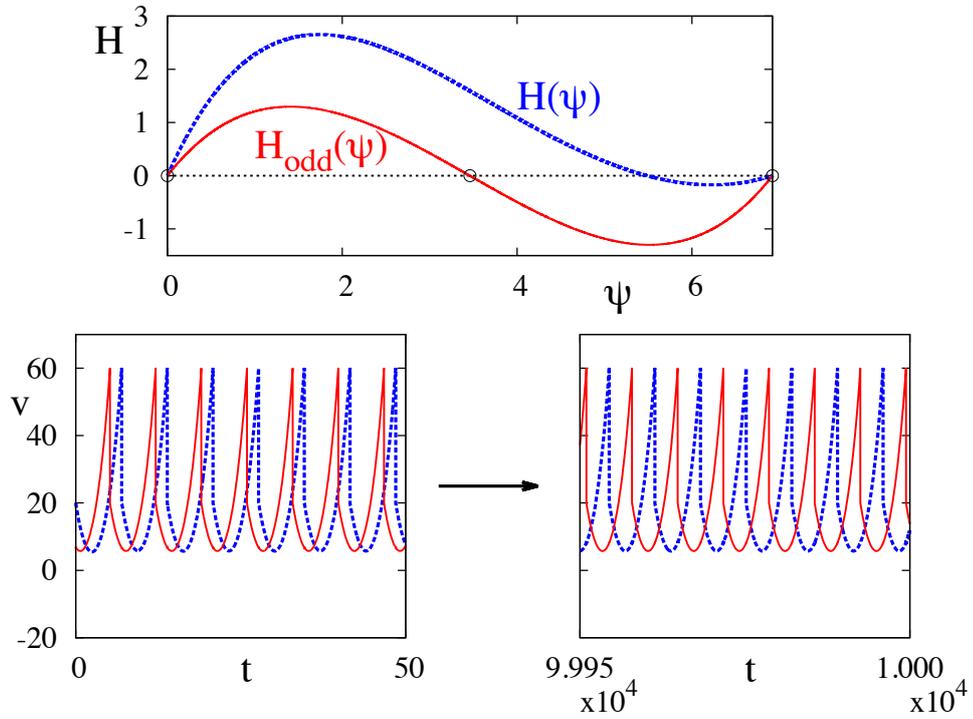


Figure 33: Top: Phase interaction function (blue) along with its odd component (red) for a network of two gap junction-coupled neurons in an intrinsically fast spiking regime. The open circles show the solutions of $H_{\text{odd}} = 0$. The phase interaction function predicts stability of the anti-phase solution, with $\psi = \Delta/2$, which is observed by trajectories of the solution (bottom) after transients have decayed. Parameters are for the regular spiking orbit in Fig. 31, except $\beta = 0.9$. The strength of the coupling between cells is set to $g = 0.001$.

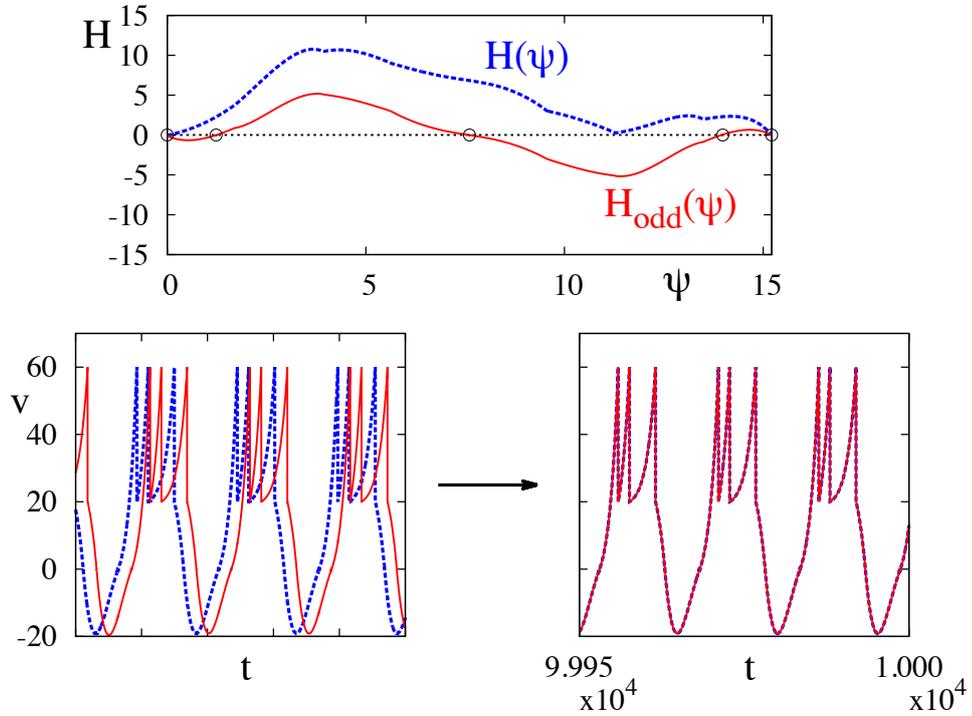


Figure 34: Top: Phase interaction function (blue) along with its odd component (red) for a network of two gap junction-coupled neurons in an intrinsically 3 spike bursting regime. The open circles show the solutions of $H_{\text{odd}} = 0$. In contrast to the situation showing in Fig.33, the phase interaction function predicts bistability of the synchronous solution with $\psi = 0$, and the anti-phase solution. For initial conditions with $\psi < 1$, we expect to see synchrony between the cells after transients have decayed, which is observed by trajectories of the solution (bottom) after transients have decayed. Parameters are are for the bursting spiking orbit in Fig. 31, with the strength of the coupling between cells set to $g = 0.001$

3.12 Firing map

We now consider an approach to study the existence and stability of behaviour of a single neuron subject to constant forcing. Due to the nature of the non-smooth dynamics of the system at reset, it is useful to consider a map of the adaptation variable at successive firing times. This will collapse the dynamics of the full system to a one-dimensional return map. This has previously been considered by Touboul and Brette [37] for a broad class of planar nonlinear IF models. Here, we focus on the construction of such a map for the PWL-IF model. We consider a set, called the Poincaré section, $\Sigma = \{(v, a) | v = \bar{v} \in \mathbb{R}\}$ which is transverse to the flow for all $(\bar{v}, a) \in \Sigma$. The value of \bar{v} above is arbitrary, so that our section may be placed anywhere in the phase plane. The first return map is function which gives, for each value $a_0 \in \mathbb{R}$, the value of a at the next intersection with Σ , of a trajectory starting from $(\bar{v}, a_0) \in \Sigma$. The second return map gives the second intersection of such a trajectory with Σ and so forth. We refer to the firing map as the first return map with $\bar{v} = v_R$. We note that trajectories will not intersect v_R upon reaching threshold, but are reset discontinuously to it, and thus we may consider intersection of the trajectory with $\Sigma_1 = \{(v, a) | v = v_{th}\}$ and apply the reset conditions to give the value of a we seek. Suppose that the trajectory starting from (v_R, a_n) reaches threshold at time Δ_n . Defining the firing map as the unique function $P : \mathbb{R} \rightarrow \mathbb{R}$ such that $P(a_n) = a(\Delta_n) + k$, we have

$$P(a) = \begin{cases} G_{21}^1(\Delta_n)v_R + G_{22}^1(\Delta_n)a + K_{21}^1(\Delta_n)I + k, & a < a_c, \\ G_{22}^1(T_3)a^* + K_{21}^1(T_3)I + k, & a > a_c, \end{cases} \quad (3.66)$$

where T_3 is the flight time from $v = 0$ to $v = v_{th}$, $a < a_c$ is the same as the one described in section 3.5, which separates trajectories, starting from (v_R, a) , which cross the switching manifold from those which do not, and

$$\begin{aligned} a^* &= G_{22}^2(T_2)(G_{21}^1(T_1)v_R + G_{22}^2(T_1)a \\ &\quad + K_{21}^1(T_1)I) + K_{21}^2(T_2)I, \end{aligned} \quad (3.67)$$

where T_1 is the flight time from $v = v_R$ to $v = 0$ and T_2 is the total flight time in the region $v < 0$. The flight times are given by the solutions to transcendental

equations and hence are not available in an explicit form, and so we find the values of T_1, T_2 and T_3 numerically.

At $a = a_c$, the map may have a discontinuity, dependent on the pair (β, ω) . For the map to be discontinuous, we require that β to larger than the value in (3.33), so that the dHB at $I = 0$ is subcritical. An example of such a map is depicted in Fig. 35.

Fixed points of the map are found by solving $\bar{a} = P(\bar{a})$, and the points are stable if $|J(\bar{a})| < 1$ where $J(\bar{a}) = P'(\bar{a})$. Fixed points of the map may lose stability via a tangent bifurcation where $J(\bar{a}) = 1$ or a period-doubling bifurcation where $J(\bar{a}) = -1$. They can also be lost as they pass through the discontinuity at $a = a_c$.

In order to characterise the stability of the fixed points, we first need to find an expression for $J(a)$. We have, upon setting $a_n = a, \Delta_n = \Delta$, for $a < a_c$ that

$$\begin{bmatrix} v_R \\ P(a) \end{bmatrix} = G^1(\Delta) \begin{bmatrix} v_R \\ a \end{bmatrix} + K^1(\Delta)\mu + \begin{bmatrix} v_R - v_{th} \\ k \end{bmatrix}. \quad (3.68)$$

where G_1, K_1 and μ are as in (3.15). Differentiating this equation with respect to a yields:

$$\begin{bmatrix} 0 \\ J(a) \end{bmatrix} = G^1(\Delta) \begin{bmatrix} 0 \\ 1 \end{bmatrix} + \frac{d\Delta}{da} \left\{ A_1 G^1(\Delta) \begin{bmatrix} v_R \\ a \end{bmatrix} + G^1(\Delta)\mu \right\}. \quad (3.69)$$

We may solve the first of the above equations to find an expression for $d\Delta/da$ as

$$\frac{d\Delta}{da} = \frac{G_{12}^1(\Delta)}{(G_{21}^1(\Delta) - G_{11}^1(\Delta))v_R + (G_{22}^1(\Delta) - G_{12}^1(\Delta))a - G_{11}^1(\Delta)I}, \quad (3.70)$$

after which we may then define $J(a)$ in terms of $d\Delta/da$ as

$$\begin{aligned} J(a) = G_{22}^1(\Delta) + \frac{d\Delta}{da} \left(\omega \left\{ (\beta G_{11}^1(\Delta) - G_{21}^1(\Delta)) v_R \right. \right. \\ \left. \left. + (\beta G_{12}^1(\Delta) - G_{22}^1(\Delta)) a \right\} + G_{21}^1(\Delta)I \right). \end{aligned} \quad (3.71)$$

A similar process determines $J(a)$ for $a > a_c$ (taking care to piece together solutions and derivatives across $v = 0$). We observe a qualitatively similar form of the firing

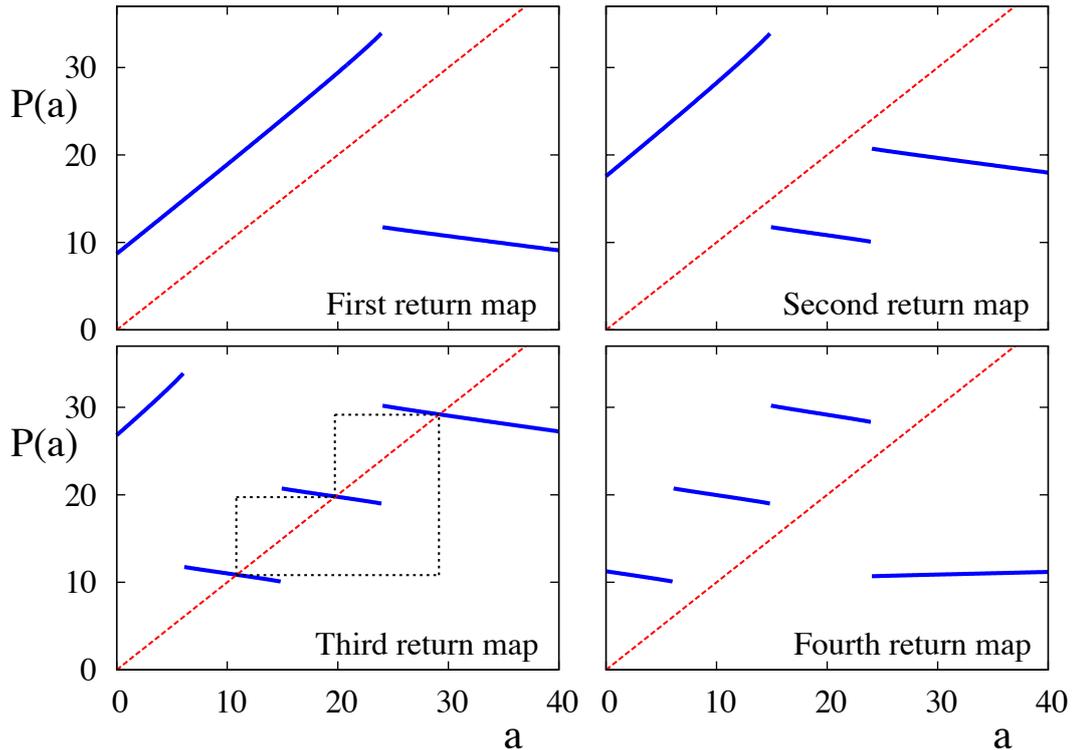


Figure 35: First, second, third and fourth return maps at $\omega = 0.19$, $\beta = 1.2$, $I = 4.0$, $s = 0.35$ and $k = 0.4$. We see three stable fixed points on the third return map, (dashed cobwebs) corresponding to a 3-spike burst.

map to that found by Medvedev [142] for the Chay-Keizer model [143, 144] (for bursting in a pancreatic β -cell). The map may be thought of as divided into two portions at a_c , with the left hand portion, with $a < a_c$, attaining some global maximum value and the right hand portion having a small and negative slope. Geometrically, fixed points of the firing map correspond to the intersection of the map P with a line of unit slope and zero origin. Fixed points may exist in either portion, and it is easy to construct scenarios in which fixed points ‘disappear’ across the discontinuity. As an example, we plot in Fig. 35, the first, second, third and fourth return maps in a parameter regime that supports a stable burst with 3 spikes. We see three stable fixed points on the third return map, corresponding to the a -values at the spike times.

Stable fixed points on the left hand portion of the map correspond to fast

spiking solutions, whilst those on the right correspond to regular spiking solutions. We note that, under parameter variation, it is possible to generate unstable fixed points in the right hand portion of the map. In this parameter regime, we observe doublet firing, the onset of which is marked by a period-doubling bifurcation. Shown in Fig. 36, is the representation of doublets in the return maps. We plot both the first and second return maps, along with their respective first derivatives. There exists an unstable fixed point in the first return map, and two stable fixed points in the second, corresponding to the doublet. As I is increased (decreased), the fixed point in the first return map will move leftward (rightward) and stabilise so that the fast (regular) spiking solution becomes stable and we lose the doublet.

The bifurcation to doublet firing occurs as fast spiking orbits approach the switching manifold. We may track the onset and termination of doublet firing in (I, β) parameter space by continuing the steady states for which $J(\bar{a}) = -1$. We find that for a given value of β , there are necessarily two bifurcation branches; one to mark the onset and one to mark the termination of doublet firing. We also observe that below some value of β , the model no longer fires in doublets, and the transition from regular to fast spiking occurs exactly as the fast spiking orbit grazes the switching manifold. Here, all of the steady states have $J(\bar{a}) > -1$.

As well as doublet firing, we often have bistability of periodic attractors near bifurcations, as can be seen in the top panel of Fig. 25, in which the sub-threshold oscillation and spiking orbit are both stable. Since we cannot always ‘see’ sub-threshold attractors with the firing map, we may repeat the same methodology, setting $\bar{v} = 0$, thus moving the Poincaré section to $\Sigma_2 = \{(v, a) | v = 0\}$. The emergence of the spiking orbit is marked, as for doublet firing, by the passing of a steady state through $J(\bar{a}) = -1$ so we may track this point in parameter space to give us the boundary on which this occurs. The grazing bifurcation, resulting in the loss of the stable sub-threshold oscillation occurs as the fixed point corresponding to the sub-threshold oscillation crosses the discontinuity in the return map. Thus, unlike the spiking solution, the sub-threshold oscillation is always stable where it exists. This does not, however, preclude the existence of unstable sub-threshold

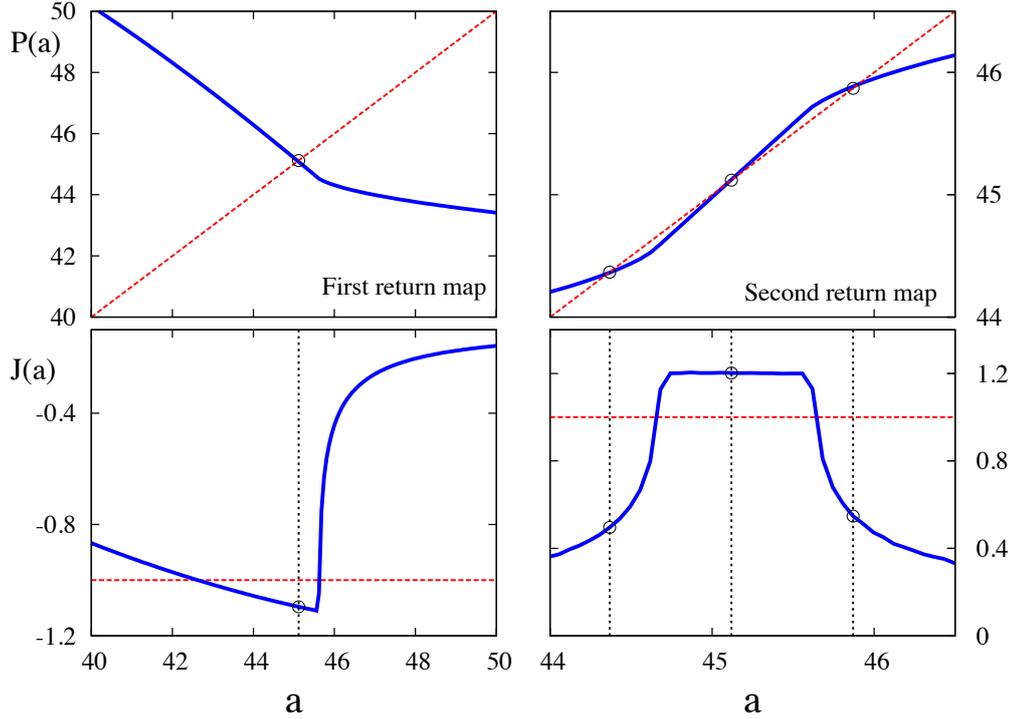


Figure 36: First and second return maps (top), together with their first derivatives (bottom), for the doublet firing parameter regime with $\omega = 0.9, \beta = 1.2, I = 10.0, k = 0.04$ and $s = 0.35$. The vertical dashed lines in the lower figures indicate where the fixed points of the maps are. In all cases, the fixed point locations are shown by open circles. We can see the fixed point in the first return map is unstable. Of the three fixed points in the second return map, we observe one unstable fixed point corresponding to the unstable fixed point in the first return map, along with a pair of stable fixed points corresponding to a doublet.

limit cycles, which we may expect when the dHB is subcritical. We may then observe where the fixed point ‘disappears’ to track where in parameter space the grazing bifurcation occurs. Interestingly, we note that the system may already be in a bistable regime as the dHB occurs. Depicted in Fig. 37 are the first return maps just before and after the bistability between subthreshold orbits and spiking ones, as considered in Fig. 30. In both panels, we see that there are two fixed points. Of these fixed points, the leftmost one corresponds to the spiking orbit,

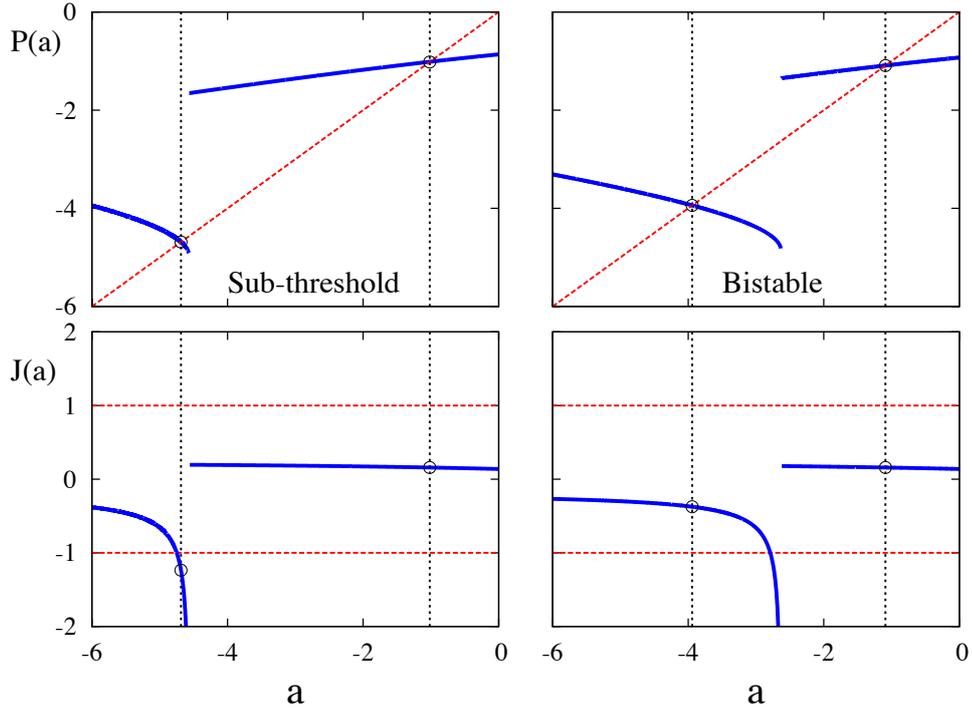


Figure 37: First return maps (top), together with their first derivatives (bottom), for the parameter sets in the left two panels of Fig. 30. The vertical dashed lines in the lower figures indicate where the fixed points of the maps are, which are also shown with open circles. We can see the leftmost fixed point in the left return map, corresponding to a spiking solution is unstable, whilst the rightmost fixed point on the same map is stable. As we increase I we arrive at the right pair of figures, where both fixed points are now stable. These two firing maps reflect the situations in the left and middle panels of Fig. 30.

whilst the rightmost corresponds to the subthreshold orbit. In the left panel we see that the fixed point corresponding to subthreshold orbits is stable, but the one corresponding to spiking orbits is not, having $J(\bar{a}) < -1$. Moving to the right panel, we now observe that both fixed points are stable, and so we have bistability between these solutions.

The firing map is of the type that allows for a snap-back repeller, and as such will support chaotic solutions [145]. Typically, systems have to possess at least three dimensions to be chaotic, whilst the PWL-IF model has only two. However,

the same is not true for non-smooth systems. In the case of the PWL-IF mode, the hard reset condition can act to make the system chaotic. To define such a snap-back repeller suppose \bar{a} is a fixed point of P with $|P'(\bar{a})| > 1$, and suppose there exists a point $a \neq \bar{a}$ in a repelling neighbourhood of \bar{a} , such that $a_M = \bar{a}$ and $P'(a_k) \neq 0$ for $1 \leq k \leq M$, where $a_k = P^k(a_0)$. Then, \bar{a} is called a snap-back repeller of P . Essentially, this definition means that the fixed point itself is unstable, but the map acts in such a way such trajectories starting from a neighbourhood of \bar{a} are brought back to \bar{a} . Zheng and Tonnelier [146], have shown the presence of snap-back repellers in the QIF model with adaptation. Given the similarities between the PWL-IF model and the QIF with adaptation, we may expect similar properties. We show an example of such a point in the PWL-IF model in Fig. 38, along with an associated chaotic orbit in Fig. 39. We shall now use the notion of Lyapunov exponents in non-smooth systems to demonstrate where chaotic solutions exist in the PWL-IF model.

3.13 Maximal Lyapunov Exponents

Lyapunov exponents (LEs) measure the exponential rates of divergence of nearby orbits of an attractor in state space and serve as a useful measure to identify regions of parameter space with differing emergent behaviour. Stable equilibria have only negative LEs, periodic attractors have one zero exponent, whilst the rest are negative. Chaotic attractors, however, have at least one positive LE. Conversely, where the attractor has at least one positive LE, we expect chaotic behaviour and thus, the presence of chaos in a dynamical system may be characterised in terms of LEs.

The spectrum of LEs, λ_i , for a general system $\dot{X} = f(X)$, $X \in \mathbb{R}^n$ is defined for an initial condition $X(t_0)$ as:

$$\lambda_i = \lim_{t \rightarrow \infty} \frac{1}{t - t_0} \ln \frac{\|\delta X(t)\|}{\|\delta X_i(t_0)\|}, \quad (3.72)$$

where $\delta X(t)$ is separation vector giving the distance $X(t) - \tilde{X}(t)$ between $X(t)$ and the trajectory $\tilde{X}(t)$ with a perturbed initial condition $X(t_0) + \delta X(t_0)$, and the

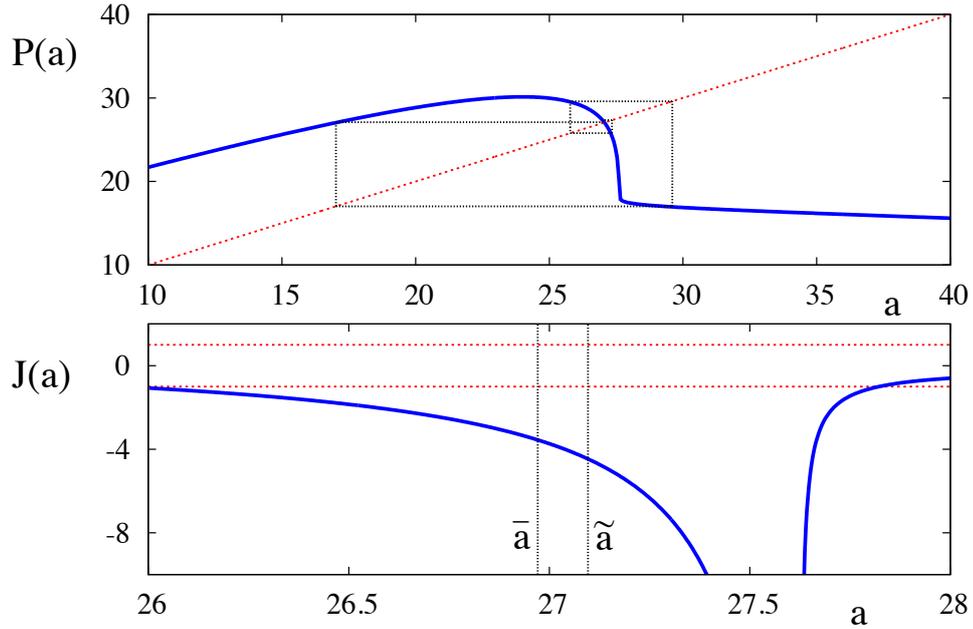


Figure 38: A snap-back repeller. Top: Firing map, Bottom: First derivative of the firing map. We see the presence of an unstable fixed point \bar{a} , in conjunction with a point \tilde{a} in its repelling neighbourhood such that $P^4(\tilde{a}) = \bar{a}$. It may be shown that the first derivative of the evolution of \tilde{a} under P is nowhere equal to zero. Parameter values here are $I = 4.0$, $\beta = 0.9$, $\omega = 0.4$, $s = 0.35$ and $k = 0.4$. The vertical dashed lines on the bottom plot indicate the location of \bar{a} and \tilde{a} . The chaotic orbit with these parameters is shown in Fig. 39.

subscript denotes the respective components of X .

LEs for continuously differentiable dynamical systems are generally calculated using the Jacobian of the system along the orbit that the flow produces, by solving a variational equation. This gives rise to a state transition matrix $\Phi(t, t_0)$ as the solution to the equation

$$\dot{\Phi}(t, t_0) = DF(X(t))\Phi(t, t_0), \quad X \in \mathbb{R}^n, \quad (3.73)$$

$$\Phi(t_0, t_0) = \mathbb{I}_n, \quad (3.74)$$

where \mathbb{I}_n is the identity matrix in \mathbb{R}^n . By evolving trajectories for a given system for large times, and using Gram-Schmidt orthonormalisation [147, 148], the full

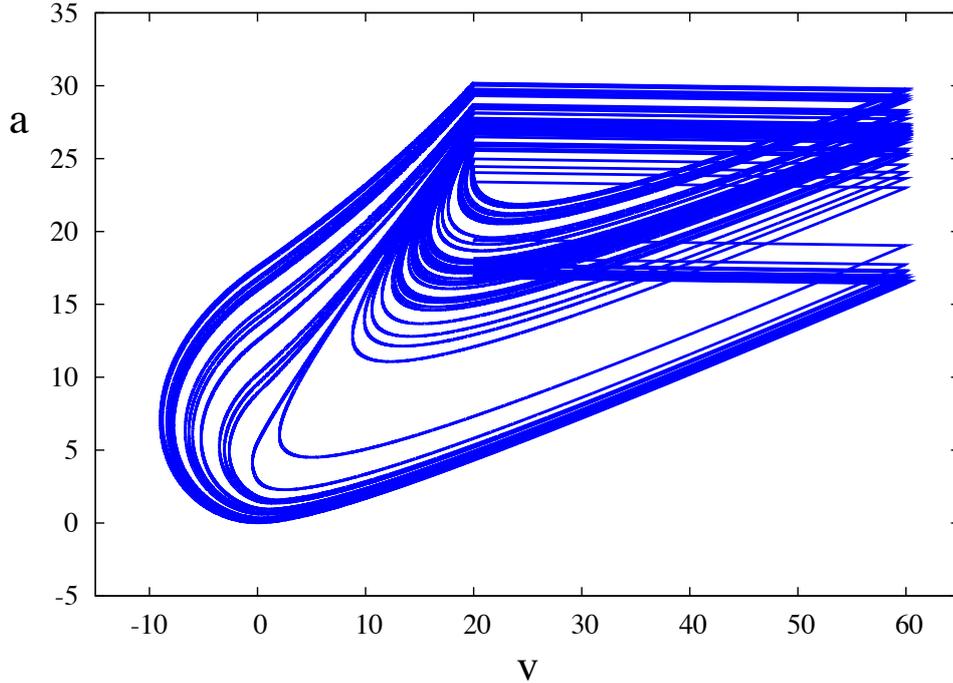


Figure 39: The chaotic orbit associated with the snap-back repeller in Fig. 38.

spectrum of LEs can be evaluated, by extracting eigenvalues of the state transition matrix solution, or by considering the log increments of the L_2 norm of the n separation vectors forming the columns of Φ . If we are only interested in calculating the maximal LE (MLE), the reorthonormalisation step is not required. This is because the LEs essentially measure the deformation of a sphere with radius $\varepsilon \ll 1$, centred on $X(t_0)$, along the n principle axes of the ellipsoid it becomes. These deformations tend to be dominated along the axis with the MLE, so that as time goes by, evaluating the rates of contraction/expansion for the other axes becomes numerically inaccurate. For this reason, if we are only interested in the MLE, we only need to consider the evolution of one typical separation vector, δX . In all cases, whenever LEs are evaluated, a number of initial conditions should be chosen to ensure that a range of phase space is sampled, and that the emergent dynamics are global, and not just some local basin of attraction for a certain attractor.

The PWL-IF system is everywhere linear, except at $v = 0$, so that the Jacobian is piecewise constant. Owing to the discontinuity in the PWL-IF model, we must

be careful when considering what happens to δX at reset, recalling that δX is a small perturbation to some orbit, X . In [149], a framework for studying the evolution of δX in impacting systems, in a model for which $\dot{X} = F(X)$ between impacts, was developed. This approach was applied to one-dimensional IF models in [150]. We now use this framework to develop the notion of LEs for the PWL-IF model.

In a sub-threshold regime, the linearised equations of motion around a trajectory $X(t)$ satisfy

$$\frac{d\delta X}{dt} = DF(X(t))\delta X. \quad (3.75)$$

Since our system is piecewise-linear, $DF(X(t)) = M$ so that

$$DF(X(t)) = \begin{cases} A_1, & v \geq 0 \\ A_2, & v < 0 \end{cases}. \quad (3.76)$$

In order to compute the change in δX across reset, we first introduce $X^-(T) = \lim_{\varepsilon \rightarrow 0^+} X(T - \varepsilon)$, so that the superscript indicates that we evaluate X at the firing event *before* $t = T$. Similarly, we introduce $X^+(T) = \lim_{\varepsilon \rightarrow 0^+} X(T + \varepsilon)$ where the superscript now means that we evaluate X just *after* $t = T$. We define an indicator function h as

$$h(X) = X_1 - v_{\text{th}} \quad (3.77)$$

so that the discontinuity in the system occurs at time T where $h(X(T)) = 0$. We also define a vector function

$$g(X) = \begin{bmatrix} v_{\text{R}} \\ X_2 + k \end{bmatrix}, \quad (3.78)$$

which governs the transition condition across the discontinuity so that $X^+(T) = g(X^-(T))$. Suppose that we have two trajectories: an unperturbed trajectory $X(t)$ and a perturbed trajectory $\tilde{X}(t)$ such that $\delta X(t) = \tilde{X}(t) - X(t)$, and that the unperturbed trajectory crosses threshold at time T , and the perturbed trajectory crosses at $\tilde{T} = T + \delta t$. Writing $\delta X^- = \delta X^-(t)$ and $X^- = X^-(t)$, we have, from [149], that

$$\mathcal{H}(X^-)[\delta X^- + (A_1 X^- + \mu)\delta t] = 0, \quad (3.79)$$

where

$$\mathcal{H}(X^-) = \left. \frac{\partial h(X)}{\partial X^T} \right|_{X=X^-(T)}, \quad (3.80)$$

is the Jacobian of the indicator function. For our choice of h , this is simply the row vector $\mathcal{H}(X^-) = [1, 0]$. We then solve (3.81) to give:

$$\delta t = -\frac{\mathcal{H}(X^-)\delta X^-}{\mathcal{H}(X^-)(A_1 X^- + \mu)} = -\frac{\delta v^-}{v_{\text{th}} + I - a^-(T)}, \quad (3.81)$$

where $\delta X^- = (\delta v^-, \delta a^-)$ and $X^- = (v^-, a^-)$. Defining the Jacobian of the transition condition as

$$\mathcal{G}(X^-) = \left. \frac{\partial g(X)}{\partial X^T} \right|_{X=X^-(T)}, \quad (3.82)$$

we arrive at

$$\delta X^+ = \mathcal{G}(X^-)\delta X^- + [\mathcal{G}(X^-)(A_1 X^- + \mu) - (A_1 X^+ + \mu)]\delta t, \quad (3.83)$$

where $\delta X^+ = \delta X^+(T + \delta t)$. For $v_{\text{R}} < 0$, we would replace (3.83) by

$$\delta X^+ = \mathcal{G}(X^-)\delta X^- + [\mathcal{G}(X^-)(A_1 X^- + \mu) - (A_2 X^+ + \mu)]\delta t. \quad (3.84)$$

For the PWL-IF model the matrix \mathcal{G} is simply

$$\mathcal{G}(X^-) = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad (3.85)$$

so that (3.83), using (3.81) becomes

$$\delta X^+ = \begin{bmatrix} 0 \\ \delta a^- \end{bmatrix} + \frac{\delta v^-}{v_{\text{th}} + I - a^-} \begin{bmatrix} v_{\text{R}} + I - a^- - k \\ \omega(\beta(v_{\text{R}} - v_{\text{th}}) - k) \end{bmatrix}. \quad (3.86)$$

This is linear in δv^- and δa^- , so we may write this in matrix form as

$$\delta X^+ = S(a^-(T))\delta X^-, \quad (3.87)$$

where

$$S(a) = \begin{bmatrix} k_1(a) & 0 \\ k_2(a) & 1 \end{bmatrix}, \quad (3.88)$$

with

$$k_1(a) = \frac{v_R + I - a - k}{v_{\text{th}} + I - a}, \quad (3.89)$$

$$k_2(a) = \frac{\omega(\beta(v_R - v_{\text{th}}) - k)}{v_{\text{th}} + I - a}. \quad (3.90)$$

Thus, overall the separation vector δX evolves as

$$G(t - T_k)S(a^-(T_k)) \dots S(a^-(T_1))G(T_1)\delta X(0), \quad (3.91)$$

for $k = 1, 2, 3, \dots$, with $G(t) = \exp(Mt)$. We remark that S is what Filippov referred to as the *saltation matrix* in [103]. The MLE is then given by the natural logarithm of the modulus of the largest eigenvalue of the matrix

$$L = \lim_{k \rightarrow \infty} \frac{1}{\Delta_k} S(a^-(T_k))G(T_k) \dots S(a^-(T_1))G(T_1), \quad (3.92)$$

where $\Delta_k = \sum_{i=1}^k T_i$. A plot of the MLE in the (I, ω) plane is shown in Fig. 40. In this region of parameter space, we see bursting orbits for smaller values of I whereas larger values of I prohibit bursting, so that we have ‘burst death’ under variation of I . We observe the presence of chaotic solutions both at this boundary where burst solutions cease to exist, marked by the large sweeping vertical arc and at the boundaries of transitions between solutions with differing numbers of spikes, marked by the thin horizontal arcs. In these regions of parameter space, we see that the firing map possesses snap-back repellers, so that the map also predicts chaos.

For a one-dimensional nonlinear IF model, the above analysis becomes simpler. Assume that, below threshold, a perturbed and unperturbed trajectory, \tilde{v} and v respectively, are related by the equation

$$\delta v(t) = \Phi(t, T_m)\delta v(T_m), \quad (3.93)$$

where $\delta v = v - \tilde{v}$. Following the propagation of a perturbation through threshold gives (cf equation (3.87))

$$\delta v^+ = \frac{\dot{v}(T_m^+)}{\dot{v}(T_m^-)} \delta v^-. \quad (3.94)$$

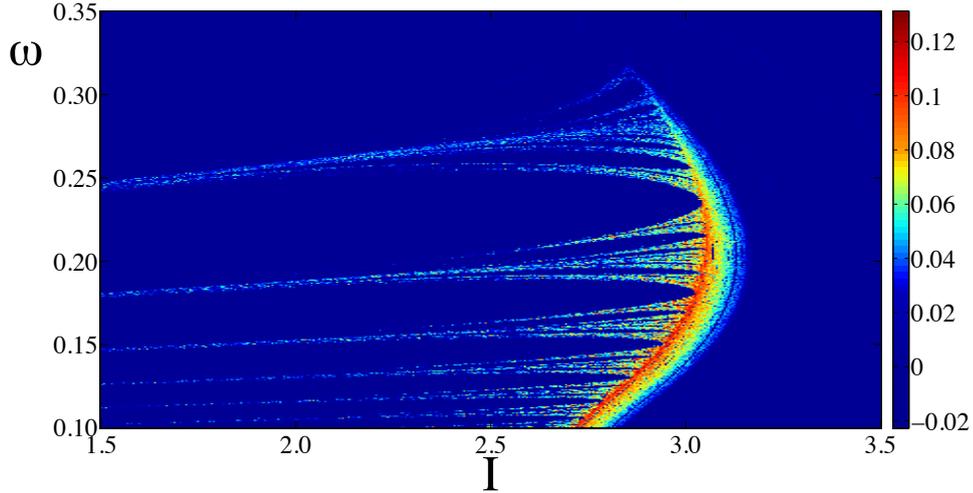


Figure 40: MLE for the PWL-IF system, (3.9) with f given by (3.10), evaluated at $\beta = 0.8$, $k = 0.4$, $v_{\text{th}} = 60$, $v_{\text{R}} = 20$ and $s = 0.35$. Warmer colours indicate positive values, whereas cooler colours correspond to zero or negative values. We see a marked boundary of chaotic solutions. This boundary marks the transition between burst firing and fast spiking as we increase I to its critical value. We also observe chaos in transitions between different burst states.

Hence the LE is

$$\begin{aligned} \Lambda &= \lim_{t \rightarrow \infty} \frac{1}{t} \ln \left| \frac{\delta v(t)}{\delta v(0)} \right| \\ &= \lim_{k \rightarrow \infty} \frac{1}{T_k - T_1} \sum_{m=1}^k \ln \left| \Phi(T_{m+1}, T_m) \frac{\dot{v}(T_m^+)}{\dot{v}(T_m^-)} \right|. \end{aligned} \quad (3.95)$$

It is informative to calculate the LE for the example of a Δ -periodic orbit in a nonlinear IF model under constant input where $\dot{v} = f(v) + I$. In this case, below threshold, a perturbed and unperturbed trajectory are related by the equation

$$\int_{v(0)}^{v(t)} \frac{dv}{f(v) + I} = \int_{\tilde{v}(0)}^{\tilde{v}(t)} \frac{dv}{f(v) + I}. \quad (3.96)$$

For small deviations between the two trajectories, we may expand (3.96) to obtain the result $\Phi(t, s) = \Phi(t - s)$, where

$$\Phi(t) = \frac{f(v(t)) + I}{f(v(0)) + I}. \quad (3.97)$$

Hence, the LE is

$$\Lambda = \frac{1}{\Delta} \ln \left[\frac{f(v(\Delta)) + I \dot{v}(\Delta^+)}{f(v(0)) + I \dot{v}(\Delta^-)} \right] = 0, \quad (3.98)$$

as expected for a periodic orbit. This result nicely illustrates that although a non-linear IF model may have a positive exponent in the absence of a firing threshold (as would be the case for the QIF model), the firing and reset mechanism can inhibit the exponential divergence of nearby trajectories.

We now proceed to study LEs of networks of linearly coupled IF systems. Suppose we have a network of N cells coupled globally through gap junctions. Writing the state vector of the matrix as $Y = (v_1, a_1, \dots, v_n, a_n)$, the subthreshold dynamics of the network may be written as

$$\frac{dY}{dt} = \mathcal{M}Y + \mathcal{F}, \quad (3.99)$$

where

$$\mathcal{M} = \mathbb{I}_N \otimes M + g\mathcal{C} \otimes B, \quad (3.100)$$

$$\mathcal{F} = \mathbf{1}_N \otimes \mu, \quad (3.101)$$

where \mathbb{I}_N is the identity matrix in \mathbb{R}^N and $\mathbf{1}_N$ is the N dimensional vector whose entries are all ones. $\mathcal{C} \in \mathbb{R}^{N \times N}$ encodes the topology of the network, while $B \in \mathbb{N}^{2 \times 2}$ determines the variables through which the coupling, of strength g , occurs. For example, in the case of a globally connected network all entries of \mathcal{C} are equal to one except those on the diagonal, which are given by $(1 - N)$. Since, in the present model only the voltage equations are coupled and the voltage variable is the first in each state vector of a single neuron, we have

$$B = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}. \quad (3.102)$$

The Jacobian for the variational equation is simply $DF = \mathcal{M}$. We now consider the evolution of the network separation vector δY across the discontinuity associated with one of the neurons firing. The firing and subsequent reset of a cell i will have

an effect on all of the elements of δY . It is convenient to reformulate the reset conditions as N mappings $\mathcal{R}_i : \mathbb{R}^{2N} \rightarrow \mathbb{R}^{2N}$, $i = 1, \dots, N$

$$\begin{aligned}\mathcal{R}_i(Y_i) &= v_R \\ \mathcal{R}_i(Y_{i+1}) &= Y_{i+1} + k, \\ \mathcal{R}_i(Y_{j \notin \{i, i+1\}}) &= Y_j\end{aligned}\tag{3.103}$$

We now have N distinct indicator functions, one for each neuron in the network

$$h_i(Y) = Y_{2i-1} - v_{\text{th}}.\tag{3.104}$$

Each of the indicator functions admits a normal vector η_i , which has a one in the $(2i - 1)$ 'th position and zeros everywhere else. Following [103], we may write the matrix \mathcal{S}_i to describe the evolution of δY across the discontinuity associated with the firing of the i 'th cell in the network as

$$\mathcal{S}_i(Y) = D\mathcal{R}_i(Y) + \frac{[\mathcal{M}\mathcal{R}_i(Y) - D\mathcal{R}_i(Y)\mathcal{M}Y]\eta_i^T}{\eta_i^T \mathcal{M}(Y)},\tag{3.105}$$

so that

$$\delta Y^+(T) = \mathcal{S}_i(Y^-(T))\delta Y^-(T)\tag{3.106}$$

upon the firing of the i 'th cell, where $D\mathcal{R}_i$ may be evaluated as the identity matrix in $\mathbb{R}^{2N \times 2N}$ with the $(2i - 1)$ 'th entry along the leading diagonal set to zero. For the gap junction system, we may evaluate \mathcal{S}_i as essentially being the identity matrix with the $(2i - 1)$ 'th column replaced by a vector whose $(2i - 1)$ 'th element is

$$\frac{(1 - Ng)v_R + I - a_i - k + g \sum_j v_j}{(1 - Ng)v_{\text{th}} + I - a_i + g \sum_j v_j^-},\tag{3.107}$$

whose $2i$ 'th element is

$$\frac{\omega(\beta(v_R - v_{\text{th}}) - k)}{(1 - Ng)v_{\text{th}} + I - a_i + g \sum_j v_j^-},\tag{3.108}$$

whose $j_{\text{odd}} \neq (2i - 1)$ entries are given by

$$\frac{g(v_R - v_{\text{th}})}{(1 - Ng)v_{\text{th}} + I - a_i + g \sum_j v_j^-},\tag{3.109}$$

with all other entries being zero.

To illustrate the use of these matrices, we perform calculations of a network of two cells coupled via gap junctions with $g = 0.1$. We note that these scenarios differ from those covered in Sect. 3.10, since the coupling between the cells is not weak, so that we cannot apply averaging techniques. In both cases, we solve the variational equation associated with (3.101) and compute the MLE by summing increments of the natural logarithm of a separation vector, rescaling as we go. We compare the MLE obtained rescaling δY using the matrices \mathcal{S}_i with those obtained without rescaling. In all cases, we observe differences between the two, and in particular, note the erroneous nature of the latter, which we shall refer to as the naïve case.

In Fig. 41, we see the time evolution of the MLE of a system of two gap junction-coupled cells, where the cells intrinsically burst. The large time behaviour of the network is shown below the MLE. The MLE using the above matrices converges towards zero, which we expect for system with periodic orbits, as we have here. Conversely, the MLE calculated in the naïve case, without rescaling, converges to a positive value, suggesting the presence of chaos in the system, which is not true. A non-periodic network orbit is shown in Fig. 42. In this case, both the blue and red curves converge to a positive value, but a larger MLE is predicted in the naïve case.

3.14 Discussion

In this chapter, we have provided an overview of nonlinear IF models that are currently of interest to the computational neuroscience community. Their obvious discontinuous nature means that much can be gained from mathematical studies that make use of the growing variety of tools and techniques being developed for the study of non-smooth and PWL systems. Through the introduction of a particular form of PWL IF model, we have demonstrated how tools for studying smooth neuron models can be carried over into the non-smooth models. Single neuron studies in response to natural stimuli are highly relevant for understanding sensory processing and unravelling the neural code, whilst theoretical analysis of

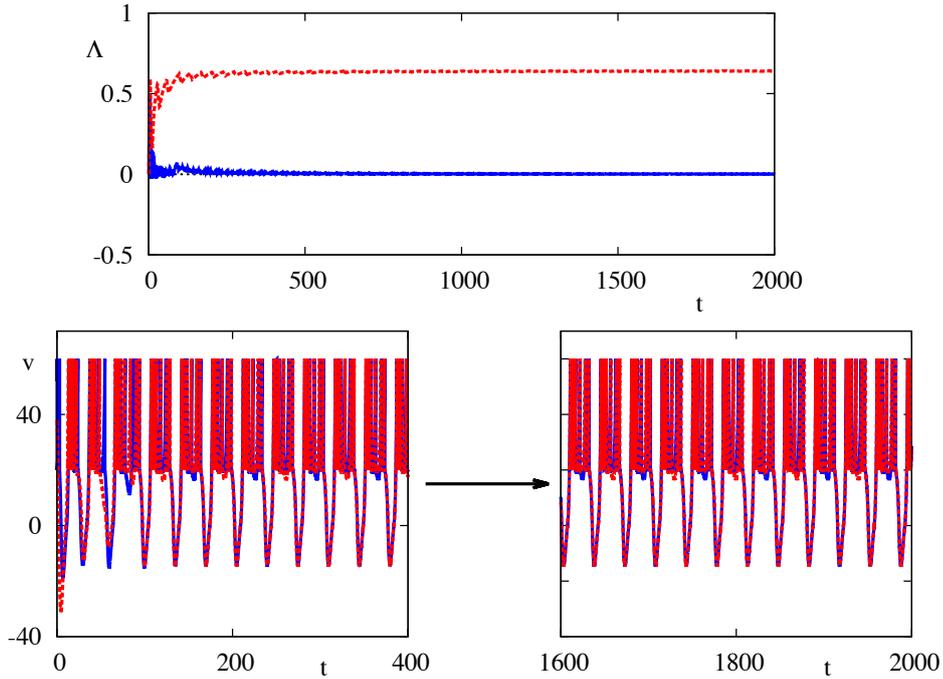


Figure 41: Top: Time evolution of the MLE computed using the appropriate rescaling across reset (blue), along with the exponent (red) calculated without rescaling across reset, for a network of two gap junction-coupled neurons in an intrinsically 3 spike bursting regime. Bottom: The large time solution shows that the network supports a periodic solution, for which we expect a MLE of 0, which the blue curve is converging towards, whilst the red curve is not. Parameters are for the bursting spiking orbit in Fig. 31, with the strength of the coupling between cells set to $g = 0.1$

spiking networks is relevant to brain studies at the highest level, and in particular, for neural computation and cognition. It is clear that, although simple, planar IF models support a rich repertoire of behaviour, as well as interesting bifurcation structures. Building up large networks of such models can be used to describe whole brain regions, and may elucidate how information is processed and transmitted in these regions. In [98], we continue building on ideas in this chapter, to study the existence and stability of solutions in large scale networks of PWL-IF neurons.

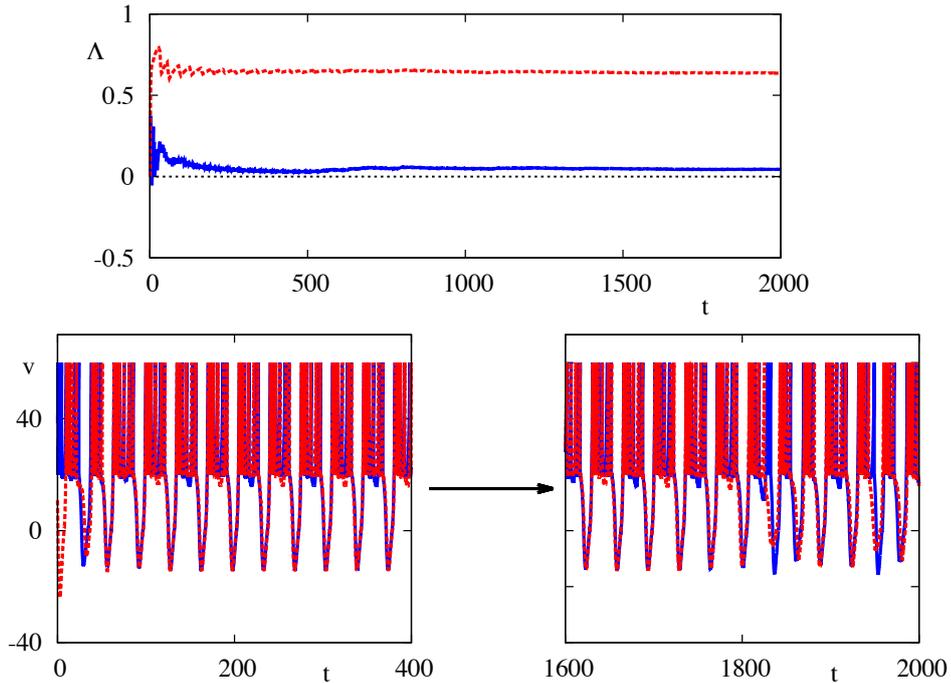


Figure 42: Top: Time evolution of the MLE computed using the appropriate rescaling across reset (blue) along with the exponent (red) calculated without rescaling across reset for a network of two gap junction coupled-neurons in an intrinsically 3 spike bursting regime. Bottom: The large time solution shows that the network is chaotic, for which we expect a positive MLE. Even though both curves are converging towards positive values, we note that there is a significant difference between these values. Parameters as in Fig. 41, except with $\omega = 0.2, I = 3.0$

There remain a number of open questions regarding these models. The construction of Arnol'd tongues for more general nonlinear IF models and in particular the PWL-IF model is one obvious next step, generalising the approach used for the LIF model, in which the non-smooth grazing bifurcations play a large role. This question has been addressed recently, for the PWL adaptive exponential IF model in [151], for example. Indeed, developing a way to describe the response of such models to more complicated signals than just periodic ones is also vital for understanding how neurons process natural stimuli. At the network level it is quite

common to first consider the behaviour of a set of weakly interacting oscillators, where knowledge of the iPRC is key to making progress. We have shown how to do this here for the PWL-IF model (and indeed, it is straightforward to do this for any single variable nonlinear IF model). However, we did not discuss the notion of isochronal coordinates that underpins the usefulness of a coupled oscillator theory. We shall go on to consider such coordinates later in this thesis.

In this chapter, we have focused on gap junction coupling between cells. Even though this may be an important form of coupling between cells, chemical synapses are by far the most dominant form of communication between neurons. In this case it is natural to consider event based coupling, as in [124], and it then remains a challenge to develop a theory of strong interactions. Once again, focusing on models with sub-threshold PWL dynamics may allow for progress. It has been shown recently that, by making the threshold fluctuate as a noisy process, that the LIF neuron can be made to fit real data from stellate neurons in the ventral cochlear nucleus [152]. Given the increased dynamic range of the PWL-IF model compared with the LIF model, this leads us to believe that the PWL-IF can be made to fit a wide range of data from experiment. The contribution of dynamics of threshold noise on such systems is an interesting mathematical problem in its own right [85]. The PWL nature of these models may additionally make them amenable to finding analytical solutions for networks in a noisy environment, for example, considering the effect of noise on reliability [153], for mixed-mode oscillations, as considered for the noisy PWL FHN model in [154], or where the coupling between units is noisy. Finally, building upon work in [155], we hope to analyse networks of heterogeneous bursting neurons to analyse the existence and stability of fully and partially synchronised states.

4 Limitations of phase reductions

The remaining chapters of this thesis are dedicated to the phase reduction techniques outlined in Sect. 2.5. Such methods are useful to reduce the number of dimensions of a system, and can be applied to systems which support oscillations in the absence of any external forcing. We refer to the oscillations in such a system as the unperturbed oscillations. As we have seen in the previous chapter, the iPRC gives information about the modulation of spike timing by external stimuli, which can subsequently be used to make predictions for weakly coupled networks, as is the focus of [5] and [156]. In the present chapter, we investigate the discrepancies between the reduced system and the full system caused by structures in phase space not considered by the iPRC. This chapter forms the majority of the work in [157]. We recall first that the assumptions needed for the appropriate application of a phase reduction are that the inputs to a given cell, which may be considered to be perturbations to the intrinsic dynamics of that cell, are weak, and that the relaxation back to the unperturbed oscillation is fast, at least with respect to the timescale of the input. In terms of a dynamical system description, we treat the inputs as perturbations to the limit cycle of the unperturbed system. If these assumptions are valid, the phase variable captures all of the information we need to describe the state of the network. In particular, the set of differences between phases of oscillators in the network will illuminate the dynamics of the network.

It is pertinent to ask about the validity of such assumptions in a neural context, taking into consideration the way in which cells interact. By far the most common form of communication between cells is through synaptic coupling. Neurons typically possess two kinds of far reaching processes: axons and dendrites. Dendrites typically receive inputs from other cells, whilst axons typically transmit action potentials. Synapses form at the junction of the axon of one cell, termed the pre-synaptic cell, and the dendrite of another, termed the post-synaptic cell.

Neurotransmitters are stored in vesicles close to the cell membrane of the axon near the synapse. When an action potential reaches a synapse along an axon, these

neurotransmitters are exocytosed into the synaptic cleft, which is simply a small gap between the axon of the presynaptic cell and the dendrite of the postsynaptic cell. The neurotransmitter diffuses across the cleft and binds to ligands, which are receptors on the dendrite of the postsynaptic cell. This induces a conformational change in the ligand, which then allows ions to flow into the dendrite from the extracellular space, inducing a current and thus a change in membrane potential as it does so. In this way, electrical signals are passed from the pre- to the post-synaptic cell.

Whilst the typical post-synaptic potential induced by a single spike may be small, the dendritic tree, made of all the dendrites of the post-synaptic cell may synapse on to large numbers of other cells, sometimes up to $O(10^4)$, so that the post-synaptic potential received by the soma of the cell is the summation of these many inputs. In addition, some synapses are slower than others, meaning that signals can persist for longer durations in the post-synaptic cell. Thus, although individual inputs are weak, the summative effect of inputs across a large number of cells, and across time may not be.

Perturbation methods, including the phase reduction technique, do not always correctly reflect the true dynamical picture, as they systematically overlook aspects of the dynamics, as we will outline in this chapter. Here, we will identify a number of scenarios in which the predictions are significantly different from those using the full model. One situation is when the perturbation causes the trajectory to leave the basin of attraction of the limit cycle, which can occur even with moderately weak forcing. Even without leaving the basin, subtler effects can result in incorrect predictions. These situations include the presence of ‘sticky’ invariant phase-space structures near a limit cycle, by which we mean structures along which trajectories spend a long time, which can cause perturbation theory to overestimate the regularity and frequency of a stimulated oscillator. We will also show that dynamical shear in a neighbourhood of the oscillator can cause it to behave chaotically when forced. These scenarios cannot be captured by infinitesimal phase reductions.

As our model of choice, we use the ML model, introduced in Sect. 2.2. We find

that, for this model, the phenomena of interest are robust to parameter variation, that is, they do not require fine tuning of parameters. We point out also that while we focus here primarily on single neuron models, our findings remain relevant for oscillators operating within networks, and we later illustrate some examples of a two cell network.

For systems that are near a bifurcation, the phase reduction technique procedure can be used to derive analytical approximations of the iPRC via normal forms [158]. In more general situations, such approximations are not available analytically. However, one can obtain iPRCs through numerical computation, for example, by solving the adjoint equation, as done in Sect. 3.10, or even directly from experimental measurements (see, e.g., [159, 160, 41] and references therein).

We recall that we are interested in studying systems of the form

$$\dot{x} = f(x) + \varepsilon g(x, t), \quad (4.1)$$

where $\varepsilon \in \mathbb{R}$, and the system with $\varepsilon = 0$ has an attracting hyperbolic limit cycle Γ , with period Δ , that we may parametrise through the phase variable $\theta \in [0, \Delta)$. The function g describes the forcing to the system, which may depend on the state of the system. For neural systems, the forcing is generally introduced through additional currents, and so this acts on the voltage variable, typically considered as the first component of x , so that $g(x, t)$ has nontrivial entries in the first component only. Since the forcing typically represents currents, we set $g(x, t) = (I(t), 0, \dots, 0)^T$. We shall denote the iPRC of Γ , found by solving the adjoint equation, by $R(\theta)$, so that the dynamics of θ are given by [161]

$$\dot{\theta} = 1 + \varepsilon R(\theta)I(t), \quad \theta \in \mathbb{S}^1. \quad (4.2)$$

4.1 Stochastic forcing.

Given the noisy environment in which neurons operate, we may be interested in studying systems in which the inputs are random in nature. In this case, the basic methodology of iPRCs can be extended to systems driven by stochastic forcing. We may think of this, for example, of electrical background noise affecting the

membrane potential. That is, suppose that in addition to a deterministic forcing, we add a white-noise term, $\sigma\xi(t)$

$$\dot{x} = f(x) + (\varepsilon I(t) + \sigma\xi(t)) \hat{k}, \quad \sigma \in \mathbb{R}, \quad (4.3)$$

where \hat{k} has a one in the first component and zeros otherwise, $\langle \xi(t) \rangle = 0$, $\langle \xi(t)\xi(s) \rangle = \delta(t-s)$, with $\langle \cdot \rangle$ denoting the ensemble average. In [97], Ly and Ermentrout show that with the above forcing, the phase variable θ satisfies the equation

$$\dot{\theta} = 1 + \varepsilon R(\theta)I(t) + \frac{\sigma^2}{2}R(\theta)R'(\theta) + \sigma R(\theta)\xi(t) + O(\varepsilon^3), \quad (4.4)$$

assuming that ε and σ are small.

Eq. (4.4) can be used to derive a number of quantities of interest. For example, if we view (4.4) as describing a neuron that spikes whenever $\theta = 0$, then a result of Ly and Ermentrout states that the firing rate, r , resulting from a constant forcing $I(t) \equiv \varepsilon$ is

$$r(\varepsilon) = 1 + \varepsilon \bar{R} + \varepsilon^2 \int_0^\Delta (\bar{R}^2 - R^2(\theta)) d\theta + \frac{\sigma^4}{4} \int_0^\Delta R^2(\theta)(R'(\theta))^2 d\theta + O(\varepsilon^3), \quad (4.5)$$

where $\bar{R} = \int_0^\Delta R(\theta) d\theta$. We will go on to consider a number of results regarding stochastically forced oscillators in Chapt. 5.

4.2 Breakdown of the phase reduction

The iPRC is defined in terms of f in (4.1) and its derivative along the limit cycle Γ . As such, it can only contain information about the flow in an infinitesimal neighbourhood of Γ . In particular, since the iPRC only uses linear information around the limit cycle, it is only accurate up to $O(\varepsilon)$. We discuss, in this subsection, a few scenarios in which the behaviour of the flow a finite distance away from Γ can have a dramatic effect on the oscillator's response.

We will go on to provide concrete examples of the effects of these scenarios in a neural model later in the chapter.

(1) *Leaving the basin of attraction of Γ .* The simplest possible way for the phase reduction to break down is when the forcing causes trajectories to leave the basin

of attraction of Γ , which we shall denote \mathcal{B}_Γ . If a perturbation causes a trajectory to leave \mathcal{B}_Γ , the behaviour of the trajectory may depend on dynamical structures far away from Γ_T . For example, if the system is bistable, or multi-stable, then a trajectory can end up in the basin of attraction of some other attractor in the system, and in time, it may or may not get forced back into \mathcal{B}_Γ . Needless to say, the behaviour of such a trajectory bears little resemblance to that predicted by the iPRC. In fact, even if the trajectory re-enters \mathcal{B}_Γ , the iPRC gives no information about the phase at which this would occur. This is because isochrons of Γ are only defined in \mathcal{B}_Γ , and not outside it. A consequence of this is that, even if we had perfect information about the isochrons, we would still not be able to capture the dynamics of this trajectory. Of course, if the other attracting sets in the phase space are also limit cycles, with their own basins of attraction, then they will also have isochronal coordinates defined in these basins. However, in general, these isochronal systems will not follow the same dynamics. This scenario must be taken into consideration when the forcing is strong relative to the distance of $\partial\mathcal{B}_\Gamma$ to Γ , where $\partial\mathcal{B}_\Gamma$ is the boundary of \mathcal{B}_Γ .

(2) *Invariant structures and ‘trapping’*. Even without leaving \mathcal{B}_Γ , perturbed trajectories passing near to $\partial\mathcal{B}_\Gamma$ can be affected by certain dynamical structures within \mathcal{B}_Γ . These structures are non-attracting, and so, at first, may seem innocuous. However, in the perturbed case, they can have a pronounced impact on the dynamics. As an example, consider a saddle of the unperturbed flow. An orbit passing near to this fixed point will remain near it for some time, since $f \approx 0$ in the neighbourhood of the fixed point. The exact time spent near the fixed point will depend of the ratio of the eigenvalues of the linearised flow at the point. Of course, the fixed point is not truly invariant for the perturbed flow, and so trajectories will ‘escape’ the fixed point, and will return to Γ after some time. However, the time taken can be long, and this effect will not be captured by phase reductions. This effect can be magnified if the forcing acts to move trajectories to regions of phase space near the saddle. In fact, the forcing may act to create trapping regions by itself, within \mathcal{B}_Γ , if the forcing is strong in comparison to the attraction to Γ .

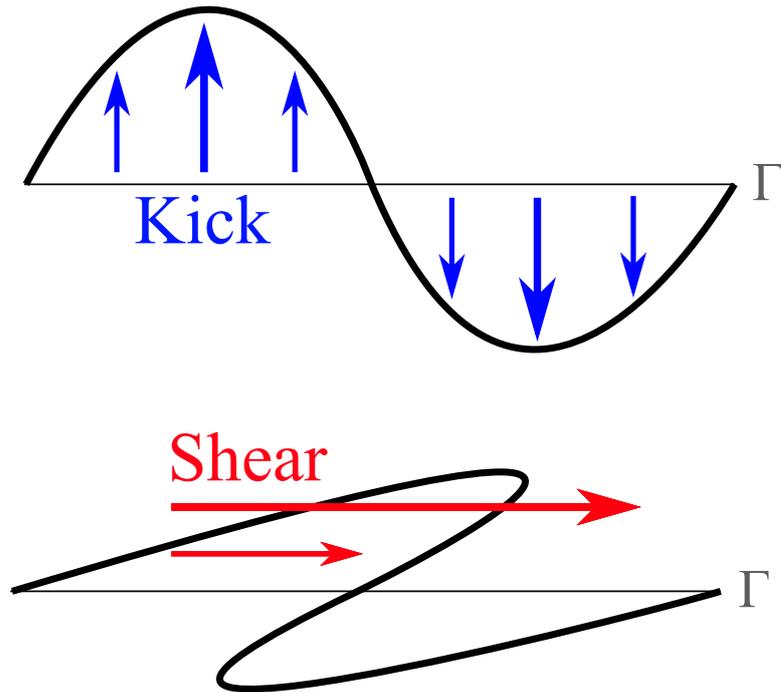


Figure 43: The stretch-and-fold action of a kick followed by relaxation in the presence of shear. Top: A non-constant kick, shown in blue, moves Γ to the black curve. Bottom: As the image of Γ under the kick relaxes back to Γ , under the unperturbed flow, the action of shear, shown in red, causes folds to appear. The size of the arrows corresponds to the magnitude of the effect. In both, the horizontal grey line depicts Γ .

(3) *Shear-induced chaos*. This is yet another dynamical phenomenon that cannot be captured by iPRCs, considered in some depth in [71]. In this phenomenon, shear forces act to speed up (or slow down) trajectories further away from Γ compared to those closer to it. This phenomenon is illustrated in Fig. 43. In each of the two pictures, Γ is represented by the horizontal grey line. Shear refers to the differential in the horizontal component of the velocity as one moves vertically up in the phase space. Here, points above Γ move around Γ_T faster than those below. If we apply an impulsive perturbation, or *kick*, shown in blue, we can produce a ‘bump’ in Γ . As the flow relaxes, this ‘bump’ is attracted back to Γ . As it evolves, it is folded and stretched by the flow if sufficient shear, indicated by the red lines, is

present. This stretching and folding of phase space is associated with complex dynamical behaviour such as horseshoes and ‘strange attractors’. The term ‘strange attractor’ essentially refers to an attractor which has a positive LE, starting from almost every initial condition in the basin of attraction of the attractor. This is precisely the kind of chaotic solution we observed in the previous chapter, for the PWL-IF model, in Sect. 3.12. This picture suggests that perturbing a limit cycle with strong shear can lead to chaotic dynamics as these folds accumulate over successive kicks. It has been demonstrated, in [71], that this can indeed happen.

We mention that invariant structures within \mathcal{B}_Γ can be a contributing factor to shear, but shear can also arise for many other reasons. We will now provide a brief overview of shear in the context of dynamical systems, reviewing some of the geometric ideas put forward in [162, 163] (and also [164, 165] and [166]). We focus here on periodically-kicked oscillators, since in this setting the various dynamical mechanisms are most transparent.

4.3 Shear-induced chaos and related phenomena

Suppose we have a system of the form

$$\dot{x} = f(x) + A \cdot H(x) \sum_{k=-\infty}^{\infty} \delta(t - kT), \quad k = 1, 2, \dots \quad (4.6)$$

where $(A, T) \in \mathbb{R} \times \mathbb{R}^+$ are parameters and $H : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a smooth function. We assume, as before, that $\dot{x} = f(x)$ has an attracting hyperbolic limit cycle Γ . This system thus models an oscillator that is given a sharp ‘kick’ every T units of time. At the times $t = nT$, we thus apply a mapping \mathcal{R} , constructed in terms of H , to the system. The dynamics of (4.6) can thus be obtained by iterating the time- T map $F_T = \Phi_T \circ \mathcal{R}$. If there is a neighbourhood \mathcal{U} of Γ , such that $\mathcal{R}(\mathcal{U}) \subset \mathcal{B}_\Gamma$, and T is long enough that points in $\mathcal{R}(\mathcal{U})$ return to \mathcal{U} , *i.e.*, $F_T(\mathcal{U}) \subset \mathcal{U}$, then $\Gamma_T = \bigcap_{n \geq 0} F_T^n(\mathcal{U})$ is an attractor for the periodically kicked system F_T . We may consider $\Gamma_T = \Gamma_T(\mathcal{R}, T)$ to be the attractor corresponding to Γ_T when the oscillator is periodically kicked.

The structure of Γ_T and the associated dynamics depend strongly on the kick parameters A and T , as well as on the relation between the kick map \mathcal{R} and the flow near Γ . When A is small, we generally expect Γ_T to be a slightly perturbed version of Γ , since hyperbolic limit cycles are robust under small perturbations [167]. In this case, Γ_T is known as an *invariant circle*, and the restriction of F_T to Γ_T is equivalent to a diffeomorphism on S^1 . Circle diffeomorphisms are well known to exhibit essentially two distinct types of behaviour: *quasiperiodic motion*, in which the mapping is equivalent to rotation by an irrational angle, that is, where the rotation number is irrational, and *gradient-like* behaviour characterised by fixed points on the invariant circle. In terms of the kicked oscillator dynamics, the former corresponds to the forced oscillator drifting in and out of phase relative to the kicks, while the latter corresponds to stable phase-locking, as considered previously.

To summarise, the above discussion suggests that, when kicks are weak, we should expect fairly regular behaviour. To obtain more complicated behaviour, it is necessary to ‘break’ the invariant circle. The main idea is best illustrated in the following *linear shear model*, a version of which was first studied in [168]:

$$\dot{\theta} = 1 + \beta\rho, \tag{4.7}$$

$$\dot{\rho} = -\lambda\rho + A \cdot H(\theta) \sum_{n=0}^{\infty} \delta(t - nT) \tag{4.8}$$

where $(\theta, \rho) \in S^1 \times \mathbb{R}$ are coordinates in the phase space, λ , β , and A are constants, and $H : S^1 \rightarrow \mathbb{R}$ is a non-constant smooth function. When $A = 0$, the unforced system has a limit cycle $\Gamma_T = S^1 \times \{0\}$. If the quantity

$$\frac{\beta}{\lambda} \cdot A \equiv \frac{\text{shear}}{\text{contraction rate}} \cdot (\text{kick ‘amplitude’}) \tag{4.9}$$

is sufficiently large, then there is a positive measure set $\mathcal{T} \subset \mathbb{R}^+$ such that for all $T \in \mathcal{T}$, Γ_T is a ‘strange attractor’ of F_T . We remark that exactly how large this parameter grouping must be depends on the function H . We also note that it is important that H be non-constant, since $H(\theta)$ is what creates the bumps in Fig. 43. The geometric meaning of the term involving β , the *shear*, is as depicted

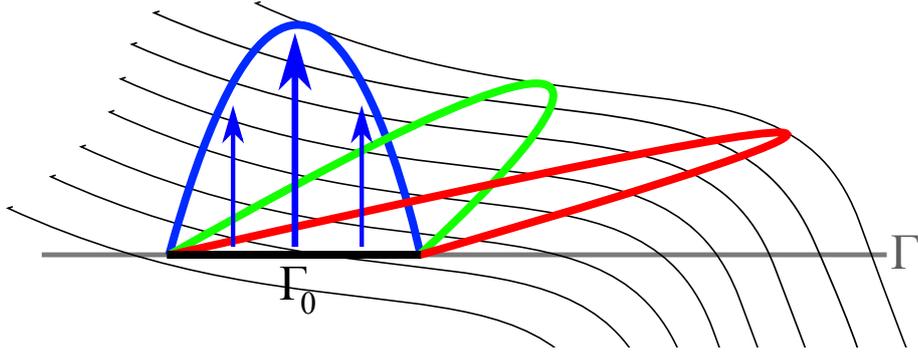


Figure 44: Geometry of folding in relation to the isochron-foliation. As in Fig. 43, the horizontal grey line represents Γ , whilst the curved black lines are the isochrons of the system. An initial segment, Γ_0 , is kicked to the blue curve, and is then allowed to relax back to Γ , passing through the green and red curves, which are images of the blue curve under the unperturbed flow. The red curve represents the image of $F_T = \Phi_T \circ \mathcal{R}$ of Γ . We observe that the points of the blue curve move down towards the isochrons as they relax back towards Γ .

in Fig. 43. It is easy to see why $\beta A/\lambda$ is key to production of chaos by fixing two of these quantities and varying the third; the larger β or A are, the larger the fold. Note also that weaker limit cycles are more prone to shear-induced chaos: the closer λ is to 0, the slower $R(\Gamma_T)$ returns to Γ_T , and the longer the shear acts on it, assuming T is large enough.

4.3.1 Shear-induced chaos and isochrons

In the general setting of (4.6), the shear does not appear explicitly as a parameter of the system. To demonstrate the effect of shear in this setting, we can appeal to the isochrons of the system to elucidate what plays the role of the shear. Let Γ and Φ_t be as at the beginning of Sect. 4.3. We may think of the isochrons, \mathcal{I}_γ , as the *strong stable manifold* of Γ , as defined in Sect. 2.4.

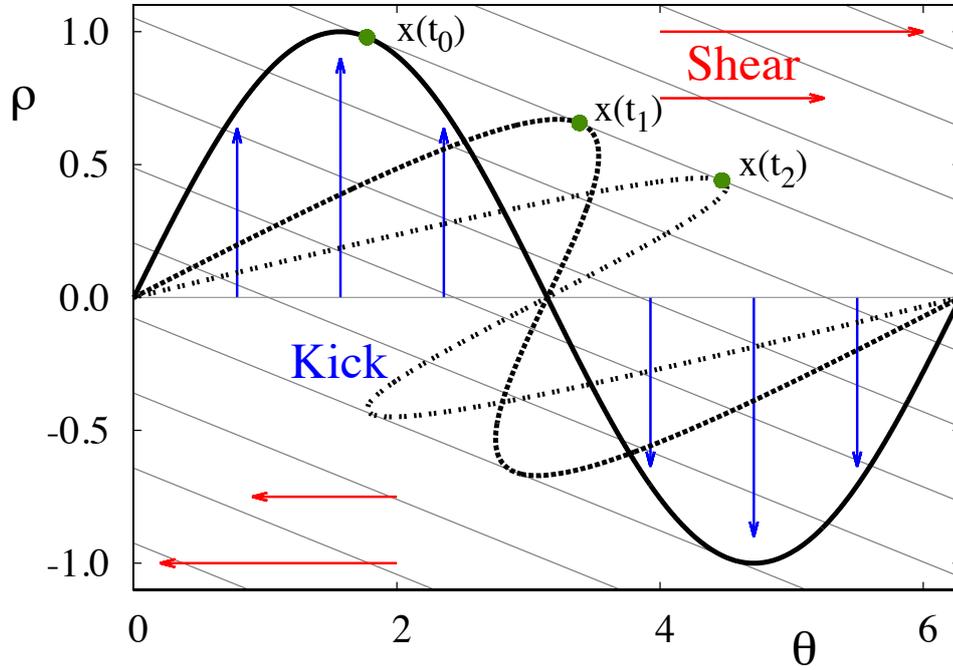


Figure 45: Stretch-and-fold action of a kick in the linear shear model (4.8). The thin black lines are the isochrons of the system, which in the case of the linear model (4.8), are simply straight lines with slope $-\lambda/\beta$. The thin grey line at $\rho = 0$ represents the limit cycle, which is kicked, at $t = t_0$, by $H(\theta) = \sin(\theta)$, with strength $A = 1$ to the solid curve. After this, the orbits are allowed to evolve under the flow generated by the continuous part of the system. The dashed and dotted curves represent the image of the kicked solid curve under this flow, at times t_1 and t_2 respectively. The green marker shows how one point, $x(t_0)$ evolves under the flow, first to $x(t_1)$ and then to $x(t_2)$, following an isochron as it relaxes back to the limit cycle. The effect of the shear forces and the subsequent folding, caricatured by the red arrows, can clearly be seen.

We now examine the action of the kick map \mathcal{R} in relation to the isochronal foliation. Fig. 44 is analogous to Fig. 43; it shows the image of a segment Γ_0 of Γ under $F_T = \Phi_T \circ \mathcal{R}$. The red curve shows one application of F_T , whilst the green curve is $F_{T/2}$. For illustration purposes, we assume Γ_0 is kicked upward with its end points held fixed, and assume $T = np$ for some $n, p \in \mathbb{Z}^+$ (otherwise the picture is shifted to another part of Γ but is qualitatively similar). Since Φ_{np} leaves each isochron invariant, we may imagine that, during relaxation, the flow moves each point of the kicked curve $\mathcal{R}(\Gamma_0)$ back toward Γ along the isochrons. In the situation depicted, the effect of the shear is evident: as the blue curve moves to the red curve, through the green curve, we clearly observe the fold forming.

This figure sheds some light into what types of kicks are conducive to the formation of strange attractors. Kicks along isochrons, or in directions roughly parallel to the isochrons, will not produce strange attractors, nor will kicks that essentially carry one isochron's manifold to another, since no folding will occur in this case. In order to have stretching and folding, there must be movement of points $\gamma \in \Gamma$ in a direction transverse to the isochrons, so that the ordering of points, in terms of asymptotic phase, is altered by the action of the kick. In the linear shear model (4.8), because of the linearity of the unforced equation, the isochrons are straight lines with slope $-\lambda/\beta$ in (θ, ρ) -coordinates. Variation in the sense above is created by any non-constant H ; the larger the ratio $\beta A/\lambda$, the greater this variation. This situation is depicted in Fig. 45. As our kick, we choose $H(\theta) = \sin(\theta)$, denoted by the blue arrows, mapping Γ to the leftmost curve. As these points relax back towards Γ , they are clearly seen to move down their respective isochrons: the point marked in green moves from $x(t_0)$, through $x(t_1)$ to $x(t_2)$ along its isochron, for example.

The effects of shear forces were discovered independently, in the rigorous work of Wang and Young, who proved, under suitable geometric conditions on phase variations, the existence of strange attractors having many of the properties commonly associated with chaos [164, 165]. These ideas have since been applied to various situations; rigorous results include [162, 163, 169, 166, 170] and [171],

and numerical results indicate the occurrence of shear-induced chaos in broader dynamical settings [13, 71, 172].

4.3.2 Summary and comparisons

The key differences between results surrounding the standard phase reduction technique and Sect. 4.3 are due to the assumptions we make about the forcing. The phase reduction method amounts to viewing the perturbation as a sequence of infinitesimal kicks, which project instantaneously the kicked trajectory towards Γ along linearised approximations to the isochrons, following each kick. In Sect. 4.3, a more global picture is painted, without any requirement of weakness of forcing. This approach allows perturbed orbits to move farther from Γ , such that they have the potential to interact, or at least pass close to, other structures. One must take care to consider the geometry, and in particular, the curvature of the isochrons in relation to the kick in assessing its impact.

We remark here that the phase reductions rules out *a priori* any possibility of chaotic behaviour, since flows in one dimension cannot be chaotic. However, the analysis considered in Sect. 4.3 suggests that kicks applied in a particular way, in the full model, can give rise to positive LEs for large sets of initial conditions. As has been stressed previously, phase reductions are very useful computational tools, allowing for analysis and simulation of a dimension-reduced model. These approaches have even further benefit when assessing network dynamics. Clearly, simulation and analysis of the full model is more costly than using phase alone. However, we feel that the messages we present in this chapter are important to emphasise that phase-only models miss notable behaviour of full systems. Calculating the isochrons of the system and then comparing them to the form of the kick being supplied will give a good indication of the presence of shear in a full. Whilst there are a number of techniques to perform such analysis, we remark that, this can more directly be achieved by simply observing successive F_T images of Γ to see whether folds develop. We will consider this for a neural model in the next section.

In order to illustrate these ideas in a neural setting, we will use the ML model, as considered in Sect. 2.2, which we recall is given by

$$C \frac{dV}{dt} = I_{\text{app}} - \bar{g}_L(V - V_L) - \bar{g}_K w(V - V_K) - \bar{g}_{Ca} m_\infty(V)(V - V_{Ca}), \quad (4.10)$$

$$\frac{dw}{dt} = \phi(w_\infty(V) - w)/\tau_w(V), \quad (4.11)$$

where

$$\begin{aligned} m_\infty(V) &= \frac{1}{2} [1 + \tanh((V - V_1)/V_2)], & \tau_w(V) &= 1/\cosh((V - V_3)/(2V_4)), \\ w_\infty(V) &= \frac{1}{2} [1 + \tanh((V - V_3)/V_4)]. \end{aligned} \quad (4.12)$$

From a quiescent state, by increasing $I(t) \equiv I_0$, from 0 to a large value, we may drive the ML model, through a bifurcation, into an oscillatory regime, in which the model supports limit cycles, meaning that we can then apply the phase reduction technique. More generally, if a neuron is operating in a ‘mean-driven’ regime, in which the stimulus it receives consists of a large DC component plus a weaker, fluctuating AC component, one can view the AC component of the stimulus as a perturbation of the oscillator [15]. Mathematically speaking, we treat I_0 as the DC component, whilst the other perturbations form the AC component.

Properties of interest

In order to compare the dynamics of the ML model to its phase-only analogue, we first need to describe some features of interest. One of the most basic statistics associated with a spiking neuron is that of its firing rate, corresponding to the frequency of the limit cycle. Indeed, many models postulate that information is encoded in the firing rates of neurons, rather than the specific spike times [14]. A related feature is that of the inter-spike interval (ISI), which is simply the difference in time between one spike and the next. The distribution of ISIs will thus contain information about the mean firing rate, as well as the variability of regular spiking activity.

We may also be interested in more specific properties of spike response. In particular, we may be interested in the ability of a neuron to respond reliably

in terms of its spike train, to repeated applications of the same stimulus. We say that a neuron, or network of neurons, is reliable if its response does not vary significantly over trials under such forcing. Reliability is of interest because it constrains a neuron's (or network's) ability to encode information via temporal patterns of spikes, spike coding being the other main theory in the debate on neural information processing, as reviewed in [173]. Mathematically, a stimulus-driven system can be viewed as a non-autonomous dynamical system of the form $\dot{x} = f(x, I(t))$, where $I(t)$ represents the stimulus. We then reformulate our discussion of reliability by asking whether the system's response to a signal $I(t)$ depends in an essential way on the initial condition $x(0)$. If not, then we can regard the system as being reliable, otherwise it is unreliable. In order to address this question, we may compute the MLE of the system. If the exponent is positive, the system is unreliable, whilst if it is negative, the system is reliable. This is due to the fact that a negative LE corresponds to phase space contraction, in which information regarding initial conditions is lost quickly, whilst positive LEs correspond to the dynamics amplifying small differences in initial conditions.

4.4 Chaotic response to periodic kicking

We will now show how shear can lead to chaos in the periodically forced ML model near a homoclinic bifurcation, resulting in a lack of reliability. As mentioned in Sect. 4.3.2, this possibility is ruled out in the case of the phase only system. To put the ML model in a homoclinic regime, we use parameter values, taken from [15] as shown in Table 1.

4.4.1 Geometry of the 'homoclinic regime'

We begin by examining the phase plane of the ML model in this regime, considering (2.28)-(2.29) with $I(t) \equiv I_0$ for some fixed I_0 as the unperturbed system, and apply to it, periodic kicks in the v -variable.

At $I_0 = I_{\text{crit}} \approx 35$, the system possesses a homoclinic loop: that is, there is a connection of one of the unstable and one of the stable branches of the saddle,

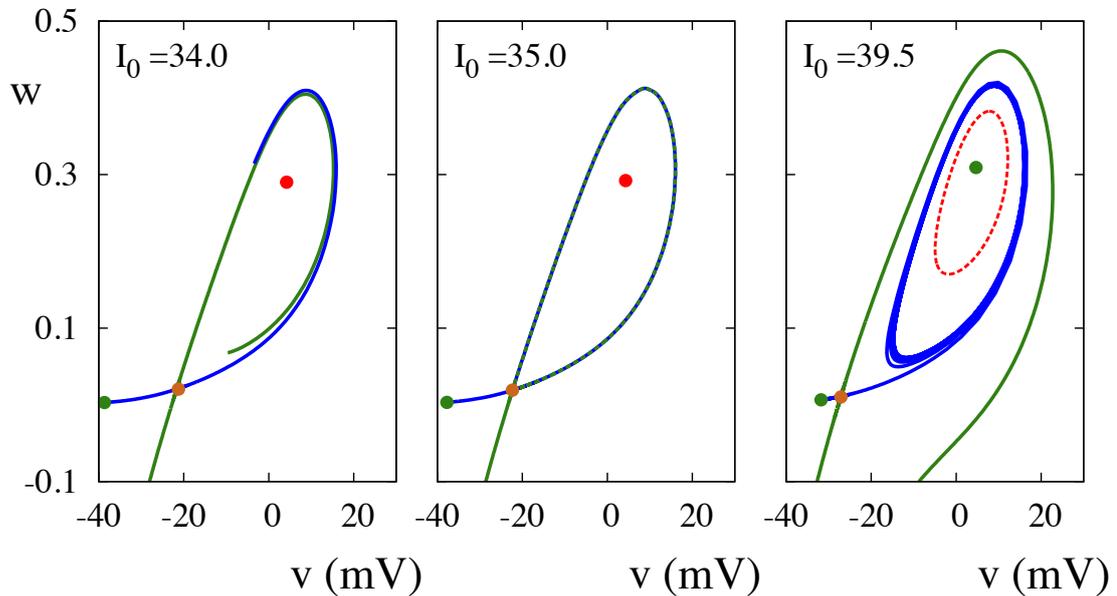


Figure 46: A homoclinic bifurcation and the ‘homoclinic regime’. In this parameter regime, the system has 3 fixed points. Stable fixed points are shown in green, saddles are shown in orange, whilst unstable focuses are red. Left: The situation prior to the homoclinic bifurcation, where $I_0 < I_{\text{crit}}$. We see that the unstable manifold of the saddle (blue curve) passes the stable manifold (green curve) of the same saddle on the outside. In this scenario, the stable fixed point is the only attractor. Middle: A homoclinic loop anchored at the saddle, with an unstable fixed point on the right and a stable fixed point on the left. As I_0 increases, the homoclinic loop breaks apart and a limit cycle emerges as the unstable manifold of the saddle now passes the stable manifold on the inside. Right: Phase portrait at the parameter regime we consider, which is well past the homoclinic bifurcation. In addition to the homoclinic bifurcation, a Hopf bifurcation has occurred, leading the fixed point on the right to become stable surrounded by an unstable periodic orbit (dashed red curve).

Parameter	Value	Parameter	Value
I_0	39.5 mA/cm ²	v_L	-60.0 mV
C	20.0 μ F/cm ²	v_{Ca}	120.0 mV
\bar{g}_{Ca}	4.0 mmho/cm ²	v_1	-1.2 mV
\bar{g}_K	8.0 mmho/cm ²	v_2	18.0 mV
\bar{g}_L	2.0 mmho/cm ²	v_3	12.0 mV
v_K	-84.0 mV	v_4	17.4 mV
ϕ	0.23		

Table 1: Parameters of ML model near a homoclinic bifurcation.

shown in orange in Fig. 46. The homoclinic loop is the blue/green curve. To the left of the saddle lies a stable fixed point, shown in green, which attracts the left branch of the unstable manifold, shown as a green curve. Inside the homoclinic loop, there is an unstable focus, shown in red.

For $I_0 < I_{\text{crit}}$, the unstable manifold (blue curve) of the saddle passes along the ‘outside’ of the stable manifold (green curve), and there is no connection, as shown in the left panel of Fig. 46. In this case, the fixed points are preserved, but there are no other invariant structures of the system, and so all trajectories will tend towards the stable fixed point as $t \rightarrow \infty$.

For $I_0 > I_{\text{crit}}$, the homoclinic loop is broken, with the unstable manifold now ‘inside’ the stable manifold, as seen in the right panel of Fig. 46. The unstable manifold wraps around a newly emerged limit cycle, the presence of which can be confirmed through an application of Melnikov’s method, as reviewed in [167]. We will take this limit cycle to be our Γ . Both the saddle and the stable fixed points to the left remain unchanged in this scenario, but we note both the change of stability of the unstable fixed point to the right to a stable one. This occurs via a Hopf bifurcation, which gives rise to an unstable limit cycle, depicted by the dashed red curve, which delimits \mathcal{B}_Γ and the basin of attraction of the fixed point it encloses. For the remainder of this section, we set $I_0 = 39.5$ mA/cm², so that the unforced flow will be similar to that shown in the right panel of Fig. 46. The stable

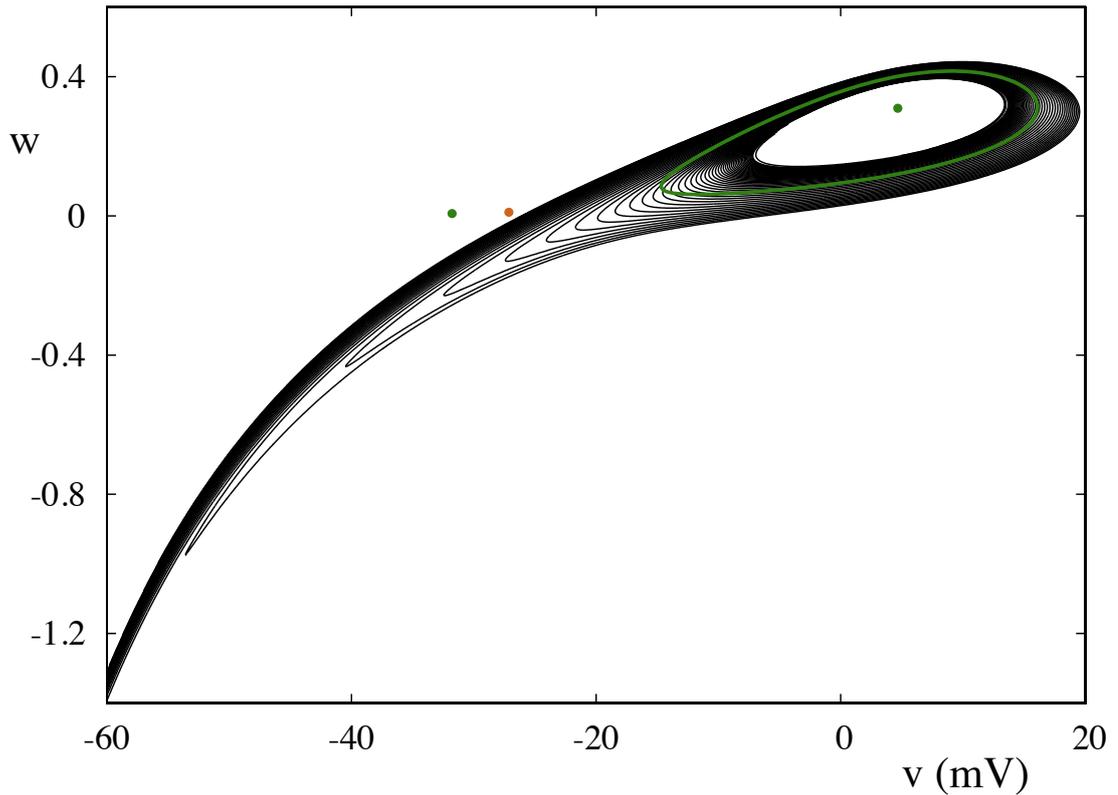


Figure 47: Isochrons of the ML model near a homoclinic bifurcation. The green curve depicts Γ , the attracting limit cycle, whilst the black curves are the isochrons of the system, with 20 distinct ones plotted. We observe that the isochrons wrap tightly around the limit cycle, with high curvature near to the saddle. This suggests that even small perturbations in this vicinity can have a pronounced affect on the asymptotic phase of an orbit. Also plotted are the two stable fixed points in green, and the saddle, in orange.

limit cycle can be found by simply solving the periodic boundary value problem associated with (2.28)-(2.29), with a suitable phase condition, whilst the unstable limit cycle is found in a similar way, but by reversing time, $t \mapsto -t$. An inspection of the isochrons of the system, as depicted in Fig. 47, in this regime shows them to be wrapped tightly around the limit cycle, with high curvature near to the saddle. This suggests that even small perturbations in this vicinity can have a pronounced affect on the asymptotic phase of an orbit. To apply the periodic kicks, we consider

input of the form $I(t) = I_0 + A \cdot \sum_k \delta(t - kT)$, in (2.28)-(2.29) where $|A|$ is the kick amplitude and T the kick period. Geometrically, this kick corresponds to shifting simultaneously all phase points by A in the horizontal direction.

If $|A|$ is sufficiently large, a number of situations can arise. Large negative kicks can drive points on Γ beyond the stable manifold of the saddle. These points will now be outside \mathcal{B}_Γ and will be attracted towards the stable fixed point on the left. Depending on the subsequent kicks, they may then never return to \mathcal{B}_Γ . In a similar fashion, kicks with an appropriate size and timing can drive points into the basin of attraction of the fixed point inside the unstable limit cycle. We will show in the next subsection that something more subtle can happen, even with small to medium kicks which do not drive points outside of \mathcal{B}_Γ .

4.4.2 Shear-induced chaos

We turn to the question of what mechanisms may give rise to shear in such a setting. Let us suppose that $A < 0$, so that the kick \mathcal{R} moves phase points to the left. We assume here that $|A|$ is not large enough to force points out of \mathcal{B}_Γ . Applying the kick uniformly to Γ , will cause about half of Γ to move inside itself, with the remainder moving outside. If we evolve this perturbed orbit under the flow generated by (2.28)-(2.29) for a period, with no kicks, we find that points close to the stable manifold of the saddle are likely to follow it for a nontrivial amount of time, taking them closer to the saddle, lengthening the time spent away from Γ . It is clear that the closer orbits come to the saddle, the longer they remain in its vicinity. The differences in time in the vicinity of the saddle between different points along Γ may lead to a fold in the $\Phi_t \circ \mathcal{R}$ -image of Γ . Although this reasoning tells us whether or not shear is present, it does not tell us whether it is sufficient to cause chaos.

As mentioned in Sect. 4.3.2, the most direct way to detect the presence of shear-induced chaos is to plot the images of $\mathcal{R}(\Gamma)$ under the unperturbed flow, and to see if folds develop in time. Fig. 48 shows the evolution of such an image. By $t = 10$, a ‘tail’ has developed: some points in this tail are evidently stuck near

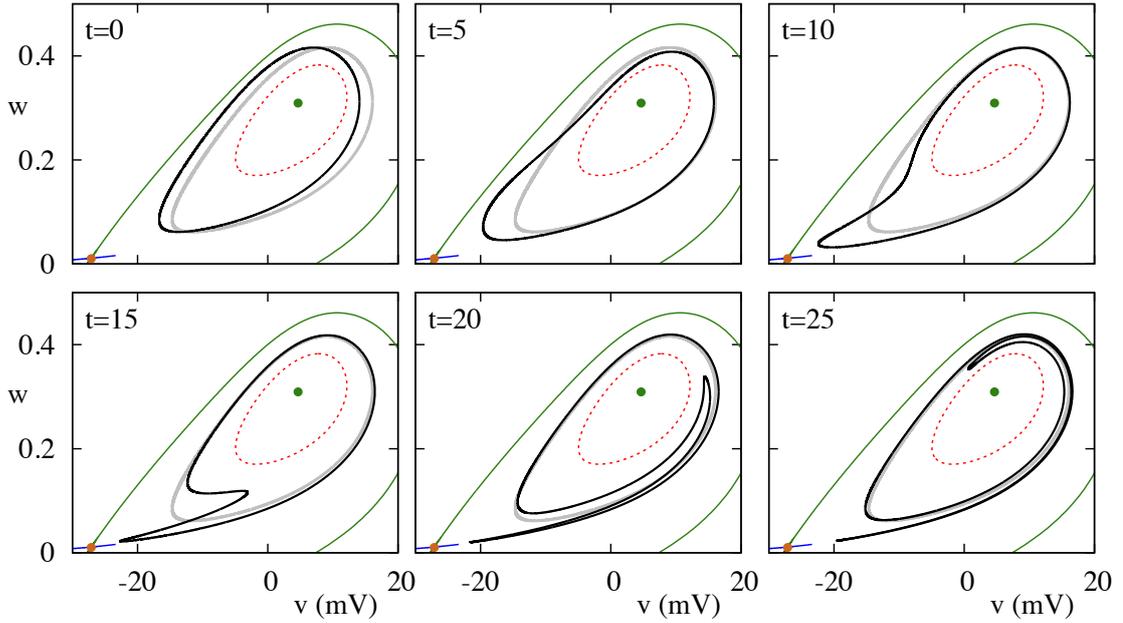


Figure 48: Shear-induced folding in the ML model. Here, the stable cycle, Γ , shown in grey is given one kick with $A = -2.0$, moving Γ to the black curve on the left, and then allowed to relax back to the cycle under the unperturbed flow. We can clearly see a fold forming as the image of Γ under the kick relaxes back to the limit cycle. Also shown are the unstable limit cycle as a dashed red curve, the stable manifold of the saddle, in green, and a small portion of the corresponding unstable manifold, in blue.

the saddle, while some other points, evidently Φ_t -images of points kicked inside the cycle, have remained inside, and at $t = 10$ they are beginning to gain on points in the tail in terms of their angular position (or phase) around Γ . At $t = 15$, these points have overtaken those in the tail, and the difference is further exaggerated in the last two frames. One would conclude that, for the parameters shown, the system very likely supports shear-induced chaos.

We may continue this process further to examine the action of iterations of the $\mathcal{R} \circ \Phi_T$ map. In Fig. 49, we show the image of the kick trajectories after they are evolved using the unperturbed flow for $T = 27$ units of time. After the second kick, we can clearly see the tail formed during the relaxation back to cycle. Following

the third iteration, there are now two visible folds, and we may observe three folds about the fourth iteration. The accumulation of these folds under repeat iterations of the map gives rise to chaotic trajectories around the strange attractor, shown in Fig. 50.

Since this step can be done quickly for the ML model, we may use it to locate parameters with the desired properties. The kick size used in Fig. 48 is $A = -2.0$, which is physiologically reasonable: it takes at least 10-15 kicks of this size to push a neuron over threshold. Furthermore, Fig. 48 also tells us that it takes on the order of 15 units of time for the fold to begin to form, so that for a periodically-kicked system to produce chaos, the kick period should probably be upwards of 20 units of time. We note that kicks with $A < 0$ that are delivered too frequently may drive points to the basin of attraction of the stable fixed point to the left of the saddle.

Now that we have a feel for the parameters for which we expect chaos, we can use LEs to explore this region in a more systematic fashion. We evaluate MLEs by solving the variational equation

$$\dot{\delta x} = Df(x) \delta x, \quad (4.13)$$

associated with (2.28)-(2.29), where, as before $Df(x)$ is the Jacobian of (2.28)-(2.29) evaluated around the orbit x and δx is a tangent, or separation vector. We form an estimate for the MLE of the system, by summing increments of the natural logarithm of the l_2 norm of $\delta x(nT)$, $n \in \mathbb{T}$:

$$\Lambda \approx \frac{1}{NT} \sum_{n=1}^N \log |\delta x(nT)|, \quad (4.14)$$

for large N , taking care to normalise δx after finding the log-increment (so that the norm of δx is reset to one just after a kick). In many, but not all cases, Λ is independent of the choice of x_0 . To overcome any potential problem of dependence of Λ on initial conditions, for each choice of parameters, we take six random initial conditions for x_0 . We also choose a random initial condition for the tangent vector δx_0 , by drawing its two components independently from a standard normal

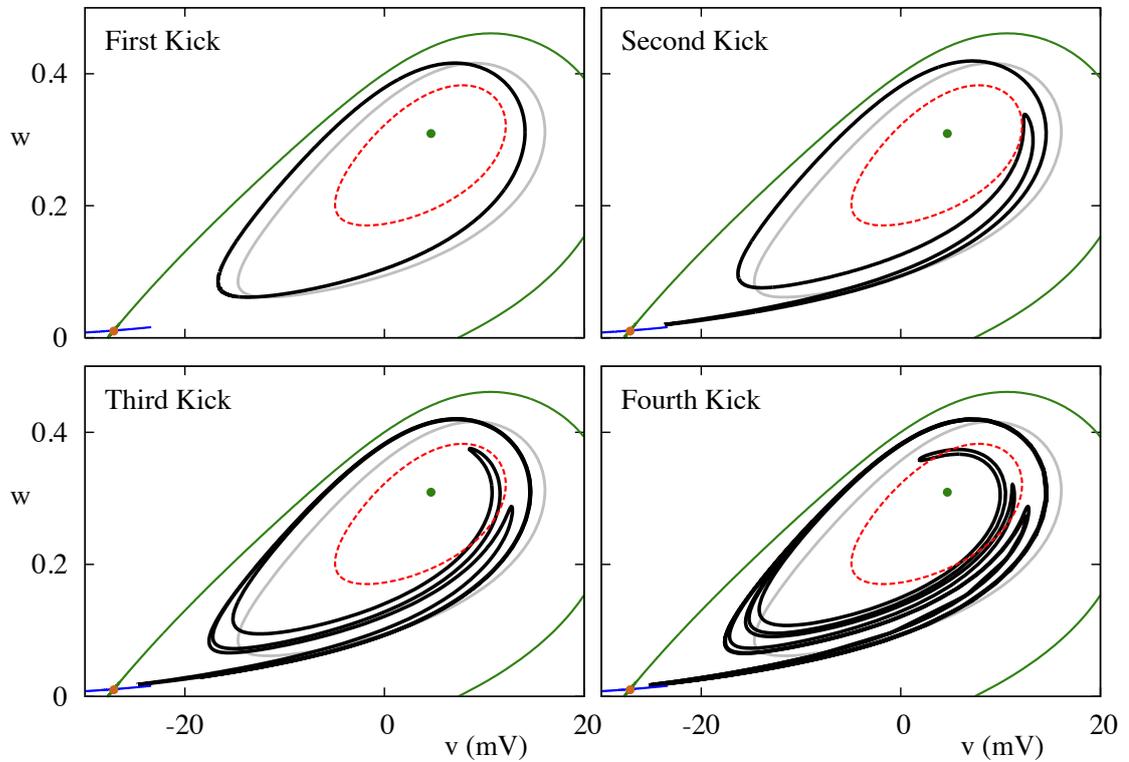


Figure 49: Shear induced folding in the ML model under repeated iterates of the kick and flow map. The grey curve in all panels represents the limit cycle of the unperturbed system, whilst the green curve represents the stable manifold of the saddle, indicated by the orange marker. The top left panel shows the image of Γ under the first kick, which essentially moves all phase points to the left, to the black curve. The successive panels show the image of this set of points after letting points evolve freely under the system defined by the ML equations, and then apply the kick again. The curves shown are the images of the initial phase points just after each kick, as indicated in the figure. We can clearly observe the shear induced folding, with one fold evident just after the second kick, two after the third kick and three just after the fourth kick.

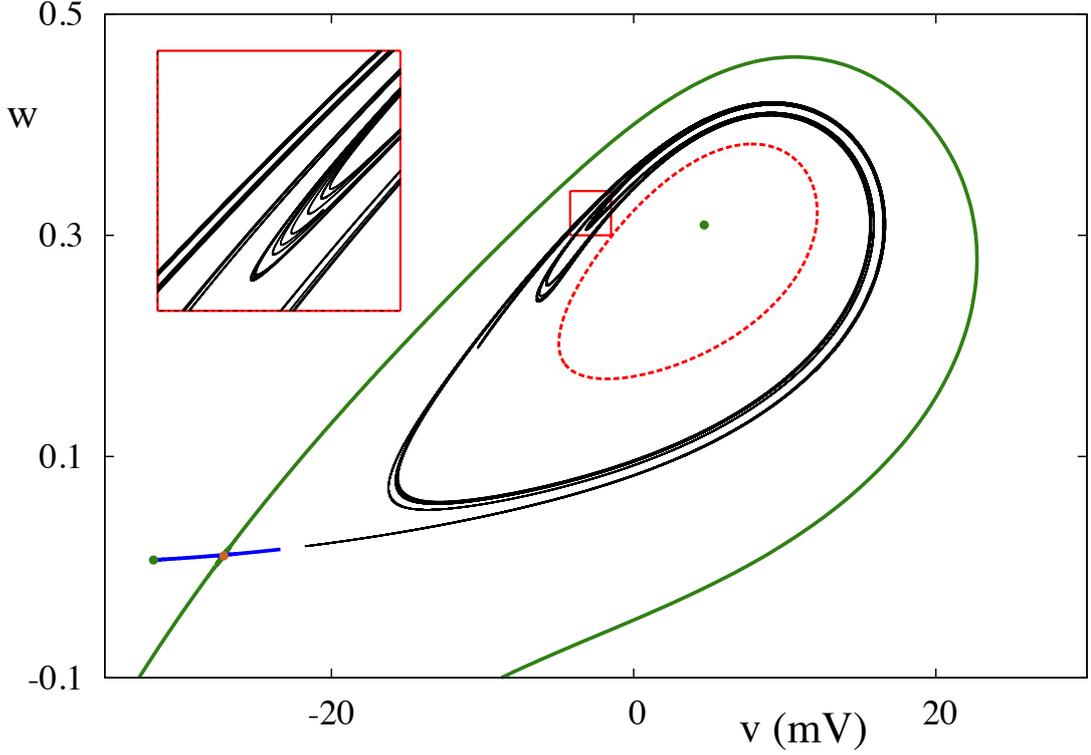


Figure 50: The strange attractor created by shear-induced chaos. The red dashed curve shows the basin boundary of the limit cycle Γ , which is the unstable limit cycle borne out of the Hopf bifurcation. Here, $A = -2.0, T = 27$. The inset shows a close up of the region enclosed by the red rectangle.

distribution and then normalising, so that $|\delta x_0| = 1$. For each set of initial conditions, (4.13) is then solved in conjunction with (2.28)-(2.29) numerically, using a fourth-order Runge–Kutta scheme for T units of time, after which the kick map \mathcal{R} is applied to the solution at the end point and the process is repeated. From the six resulting approximations for Λ , the largest and smallest are removed, and the median of the remaining ones is taken as our approximation to the MLE of the system.

In Fig. 51, we show Λ as a function of T for a variety of kick amplitudes. For small A , the MLEs are predominantly negative, as expected. As we increase A , the tendency to form positive exponents becomes greater, so that for $A = -2.0$ and sufficiently large T , most exponents sampled are positive, confirming the presence

of sustained and robust chaos discussed in Sect. 4.3. We note, however, that this is only a tendency as the fluctuations seen in the MLEs as we vary T are likely to be real, and reflect the competition between the strange attractors in the system and fixed points of the kicked map.

Although we know that the iPRC model cannot support chaos, and thus we know that the LEs are non-positive, we may still wish to evaluate them. We find the iPRC by solving the adjoint equation (2.43), with the appropriate periodicity and normalisation conditions. The iPRC for the ML model near the homoclinic bifurcation is shown in the top panel of Fig. 52, whilst, for later reference, we also depict the iPRC for the system near a Hopf bifurcation in the bottom part of this figure. Since the phase model is

$$\dot{\theta} = 1 + A \cdot R(\theta) \sum_k \delta(t - kT), \quad (4.15)$$

where R is the iPRC, θ simply increases with unit velocity between kicks. At the point of the kick, we may integrate over the Dirac δ -function and use asymptotic methods, matching solutions just prior to and just after the kick to evaluate θ after the kick. Since the derivative of a constant is zero, there will be no increase in separation, $\delta\theta$ between kicks. Across the kick, it may be shown that

$$\delta\theta^+ = \frac{R(\theta^+)}{R(\theta^-)} \delta\theta^-, \quad (4.16)$$

where the superscript $+$ denotes that we are evaluating just after the kick, and the superscript $-$ denotes that we are evaluating just before the kick. We will go through the derivation of (4.16) and show how to evaluate θ^+ in the next chapter. The results of computing the LEs are shown in Fig. 53, where we see, in all cases, that the MLE is never positive, and that some structure of the dependence of Λ on T is preserved across the values of A .

Whilst we have focused on the effect of periodic kicks, it has been shown, for the linear shear model, (4.8), that the same observations can be made for a system with similarly sized (but different) kicks, arriving at random times that are sufficiently far away, as shown in [13, 71]. The reasons for this are essentially the

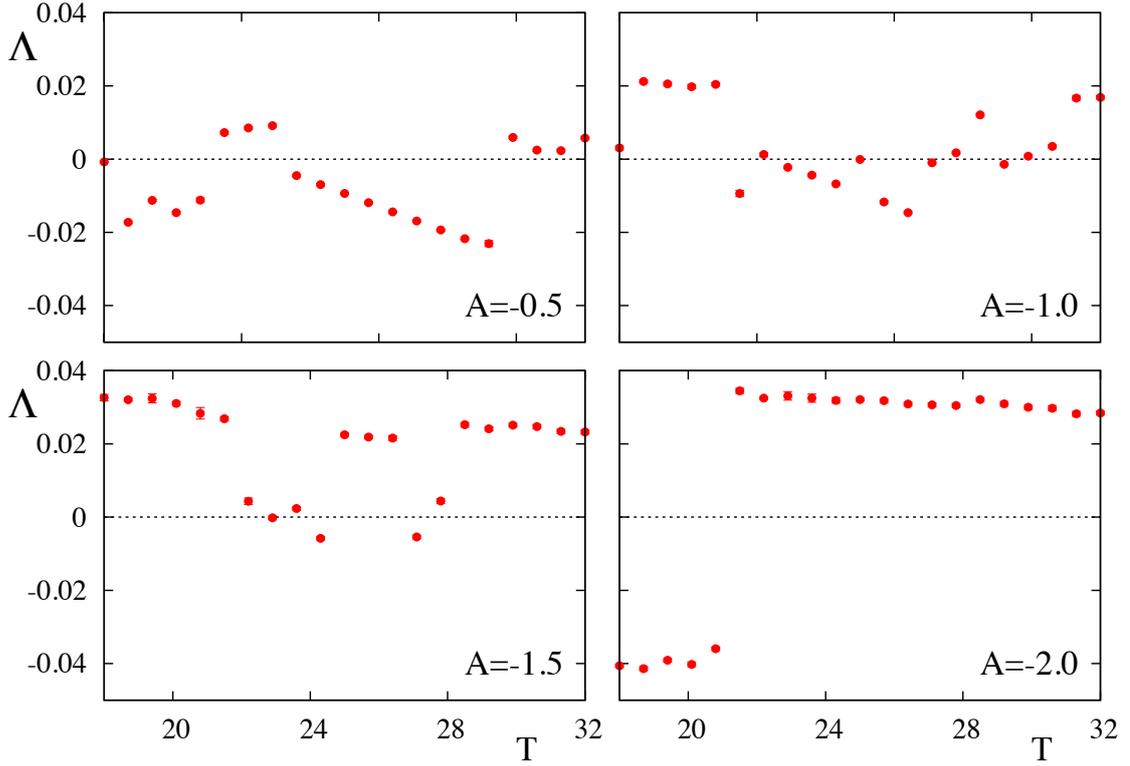


Figure 51: MLEs for the ML system, found by solving the variational equation associated with the forced system. Exponents are computed as follows: For each choice of kick amplitude and kick period, 6 random initial conditions are used to estimate exponents via long-time simulations. The maximum and minimum values are treated as outliers and discarded; whilst the dots mark the medians of the remaining exponents. We see that as T increases, the tendency for the system to support positive MLEs increases, and furthermore, this tendency increases with increasing $|A|$. We remark that the fluctuations in Λ as we vary A are most likely due to competition between the strange attractor and fixed points of the forced system.

same as the reasons discussed in this chapter. To summarise, we remark that it is the saddle (and its manifolds) that are the cause of the folding in this system. It causes a reordering of points of $\mathcal{R}(\Gamma)$ in terms of their phase. In the parameter regime that we have chosen, the saddle is not close to the limit cycle. This means that these effects would not be captured by the iPRC. Furthermore, since the

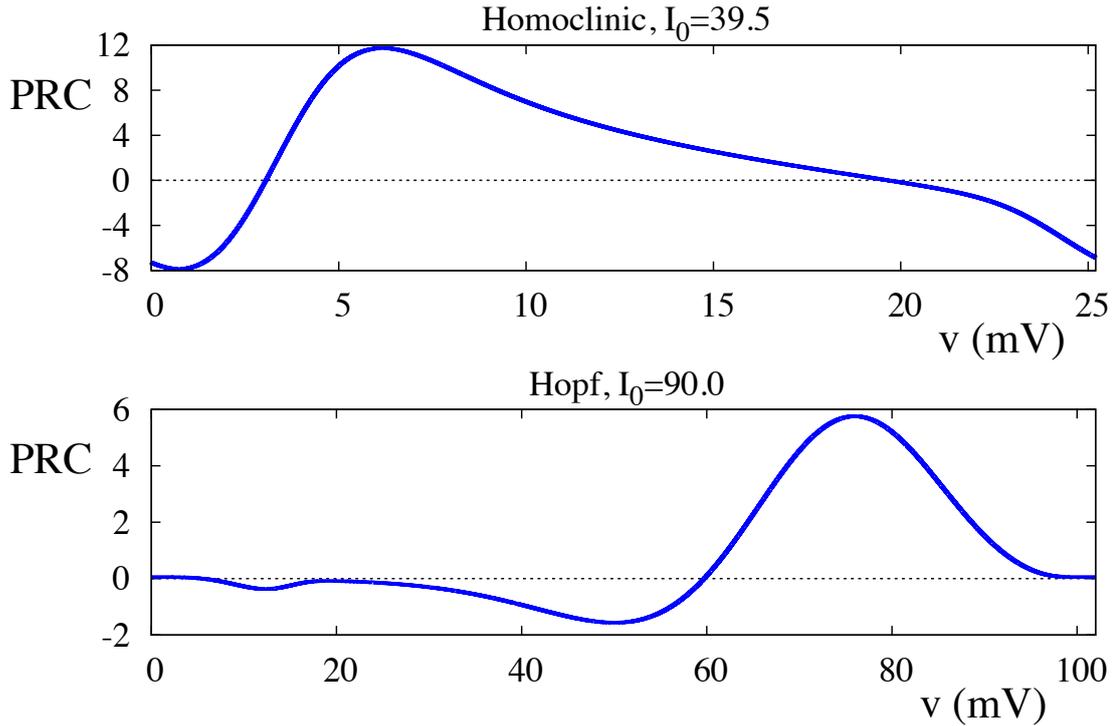


Figure 52: iPRC of the ML model, found by solving the adjoint equation (2.43) with the appropriate normalisation and periodicity conditions. Top: iPRC near a homoclinic bifurcation. Bottom: iPRC near a Hopf bifurcation.

iPRC uses information locally around Γ , it would not provide that any indication that a breakdown had occurred.

4.5 Firing rates and interspike intervals

We now move to consider the effect of ‘sticky’ sets in the interior of $\partial\mathcal{B}_\Gamma$ and probe the differences in results between the full model and those found using the iPRC. In this section, we will be primarily concerned with the firing rates and ISIs, rather than the reliability, to test overall differences between the two. We remain, for the time being, focused on the homoclinic parameter regime, where the parameter values are as in Table 1. The boundary of \mathcal{B}_Γ is given by the stable manifold of the saddle and the unstable limit cycle borne out of the Hopf enclosed by Γ . We have already seen the effect that the ‘stickiness’ of the saddle can have on the dynamics

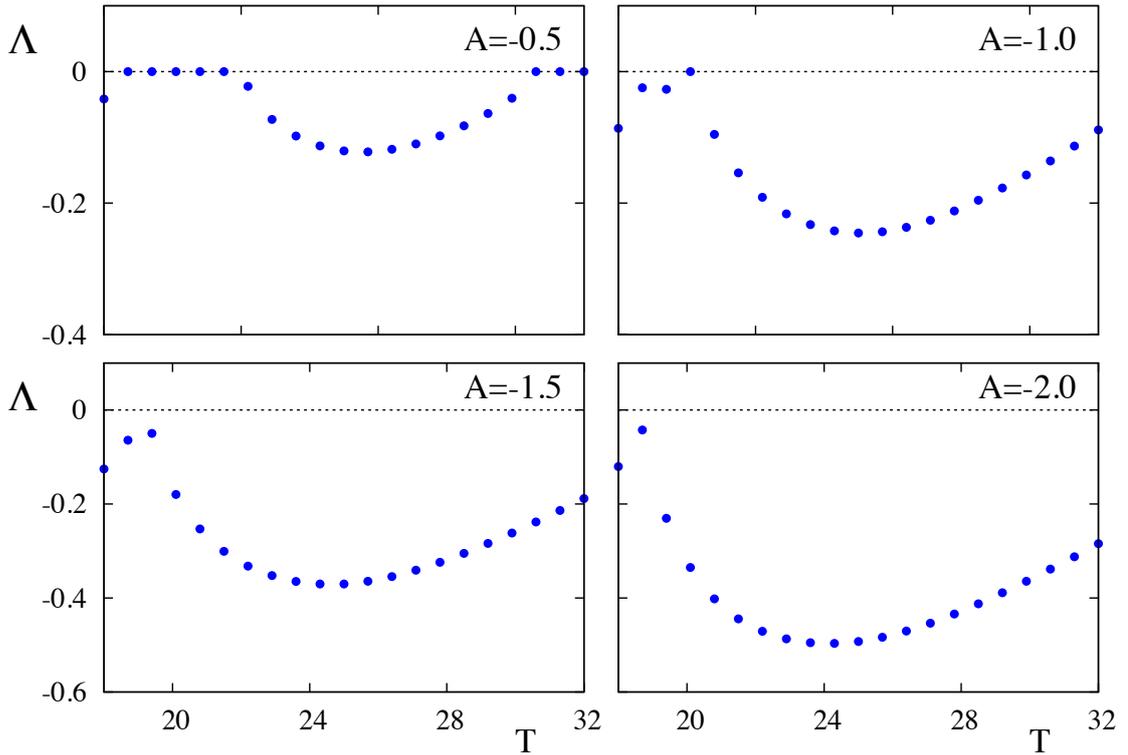


Figure 53: MLEs for the phase-only system system, found by solving the variational equation associated with (4.15). The exponents are computed in the same fashion as in Fig. 51. We see that exponents are never positive, and so there is no chaos in the system.

of the full system. In a similar fashion, we can consider the unstable limit cycle as a sticky set, since expansion away from it will be weak, as it has just emerged from a Hopf bifurcation. Additionally, orbits coming close to it may easily be forced to cross the unstable limit cycle, into the basin of attraction of the stable node.

We now move to consider stochastic forcing, rather than periodic kicks. This gives rise to a system of SDEs, in which the forcing is now continuous in time. In addition to the steady current I_0 , we add a white noise component ξ , so that $I(t) = I_0 + \sigma\xi(t)$, using the notation from physics, where the units of σ are mV/s, and study the resulting random dynamical system resulting from analysing one realisation of the white noise process at a time.

To perform a comparison between the full and phase only models, we thus need to study the system of SDEs corresponding to (2.28)-(2.29), given by

$$\begin{aligned} C dv &= \left(I_0 - \bar{g}_L(v - v_L) - \bar{g}_K w(v - v_K) - \bar{g}_{Ca} m_\infty(v)(v - v_{Ca}) \right) dt + \sigma dW_t, \\ dw &= \phi(w_\infty(v) - w) / \tau_w dt, \end{aligned} \tag{4.17}$$

and that corresponding to (4.2) given by

$$d\theta = \left(1 + \frac{\sigma^2}{2} R(\theta) R'(\theta) \right) dt + \sigma R(\theta) dW_t, \tag{4.18}$$

where we have written both systems using standard notation from mathematics (rather than from physics) to facilitate the use of the numerical integration scheme detailed in Sect. 2.19. The precise details of the numerical scheme that we use to integrate (4.17) and (4.18) can be found in Appx. A. We remark that we interpret both (4.17) and (4.18) in the Itô sense.

In order to analyse firing rates and ISIs, we first need to define precisely what we mean by a spike in each of the two settings. For a phase only model, we may define a spiking event as one where the phase θ goes through some point $\in [0, \Delta)$, since it would pass through this point once per period in the unperturbed case. For the full model, we use an artificial definition of spiking events, whereupon we say that the neuron resets if v falls below -10 mV and we define a spiking event to occur at the first crossing of v through 12.5 mV following reset. The ISIs are then given as the differences between successive times that v passes through 12.5 mV following reset. It may be seen in Fig. 50 how these points correspond to the location of the limit cycle.

4.6 Comparison of firing rates.

We begin with a comparison of the firing rates of the full and reduced model, as shown in the left panel of Fig. 54. Plotted are firing rates as a function of the drive amplitude σ for three different systems: (i) the full ML system; (ii) the phase reduction, computed empirically; and (iii) the phase reduction as computed

by the perturbative formula (4.5) of Ly and Ermentrout. We plot both (ii) and (iii) because (4.5) is itself an approximate result based on the phase reduction (4.4), valid only for small σ . To obtain the firing rates for the phase model empirically, we solve (4.18) numerically, then find the firing times based on the passing of θ through some point $\in [0, \Delta)$ and use these to compute the mean firing rate. As one would expect, for smaller forcing amplitudes ($\sigma < 0.1$), all three agree. As we increase σ , the perturbative formula tracks the empirical firing rate of the phase reduction fairly well, but neither of these iPRC-based predictions capture the dramatic drop in firing rate of the full system occurring around $\sigma = 0.2$. Shown in the right panel of Fig. 54, is a sample path of (4.17), in the case where $\sigma = 0.2$, where we observe that the drop in firing rate is due to trajectories spending more time near the saddle.

4.7 Comparison of ISI distributions

We now move to consider the distribution of ISIs in the full and reduced models. These distributions are obtained by taking the differences in spike times along sample trajectories and then by binning these data into 100 bins to form histograms, which may be normalised to give approximations for the probability density of the ISIs. We show, in Fig. 55, the numerically computed ISI distributions for the full ML model and for the phase reduction (4.4). As before, for small σ , we see the full system and phase reduction agree fairly well: the ISI distribution is concentrated around the period of the cycle (about 25.2), and that the shape is roughly Gaussian. As σ increases, we see in the full ML system both a broadening of the ISI distribution, and an overall shift toward larger ISIs. Neither of these effects are captured by the iPRC, as they involve structures a finite distance away from Γ . Finally, to compare the tails of the ISI distributions, we show in Table. 2, the fraction of ISIs greater than twice the cycle period. We observe that the iPRC systematically under predicts the probability of a long inter-spike interval, consistent with the ‘trapping’ effect of nearby invariant structures.

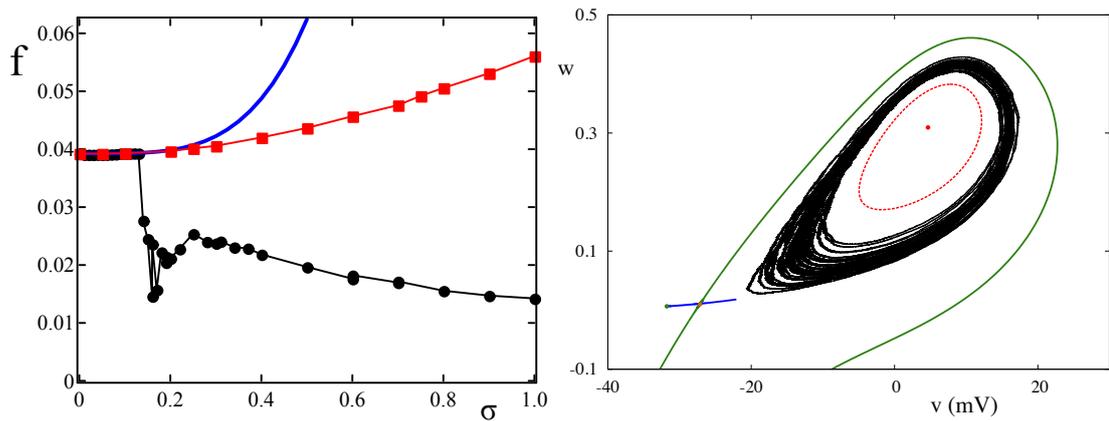


Figure 54: Left: Firing rates of the stochastically-forced ML model under variation of the noise amplitude. We compare of firing rates for the full ML system and its phase reduction driven by white noise, where the black dots are the rates for the full model, the red squares are from the empirical phase reduction and the blue curve is the approximation given by the perturbative formula (4.5). We see that all the curves are close for small σ , but that there is a significant drop in the firing rate of the full model at $\sigma \approx 0.2$. Right: Sample path of the full model for $\sigma = 0.2$. We observe that the sample path drifts near to the saddle often. In the vicinity of the saddle, the evolution of the path will slow down, explaining the drop in firing rate observed in the left panel.

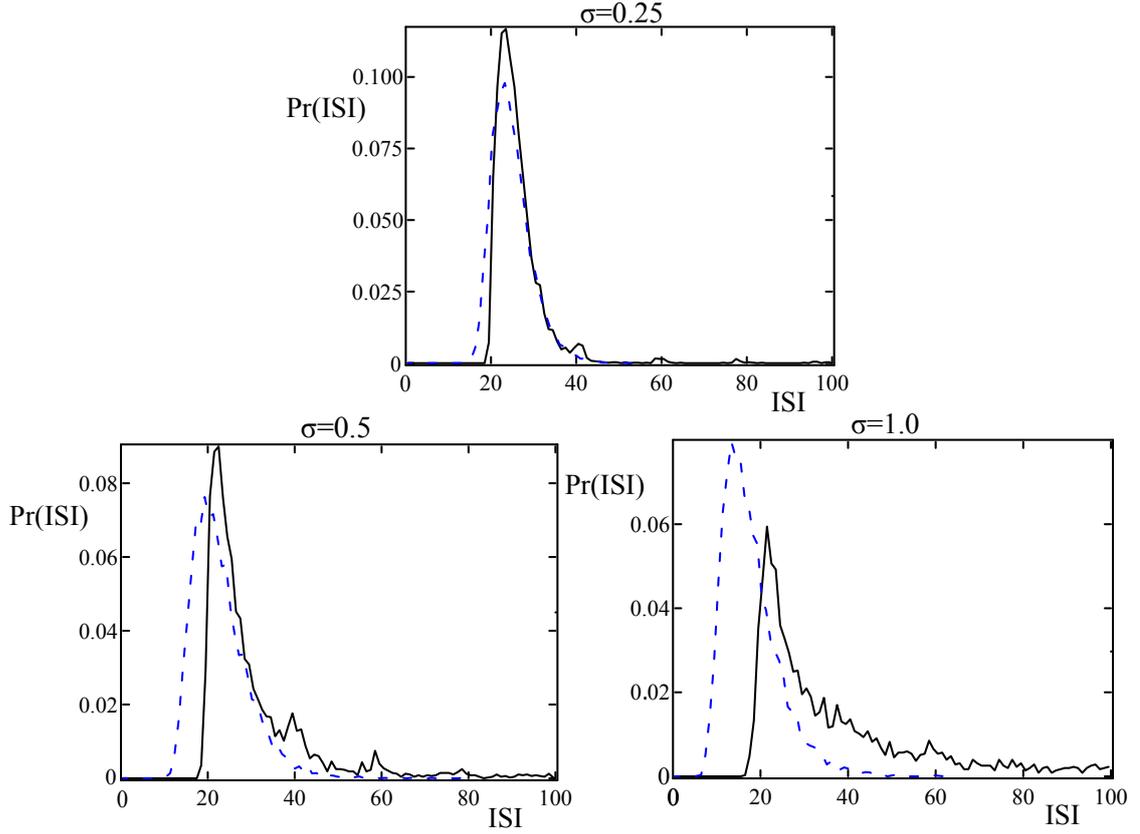


Figure 55: ISI distributions for various σ . The solid black lines are histograms for the full ML system, whilst the dashed blue lines are those for the phase reduction, with the noise amplitude shown above the respective graphs. For small σ , the histograms are very similar. However, as the noise amplitude is increased, the ISI distributions for the full model are broader, and have a mean shifted to higher ISIs compared to those predicted by the iPRC.

4.8 Altered spiking patterns in bistable systems

In order to consider an example in which the forcing causes trajectories to cross $\partial\mathcal{B}_\Gamma$, we now consider the ML system near a Hopf bifurcation, in which the iPRC is given by the bottom panel of Fig. 52. Of course, there is no reason that we would expect predictions based on the iPRC to provide any useful information in this situation, particularly since the asymptotic phase, which is essential to reducing models to a phase only description is not defined outside \mathcal{B}_Γ . It is however, notable

σ	Full ML	iPRC
0.25	0.0225	0.0003
0.5	0.127	0.0032
1.0	0.322	0.0013

Table 2: Fraction of ISIs greater than twice the period of the limit cycle. We observe that the iPRC underestimates the tails of the ISI distributions, and thus underestimates the effect of ‘trapping’ of nearby invariant structures.

Parameter	Value	Parameter	Value
ϕ	0.04	v_3	2.0 mV
\bar{g}_{Ca}	4.4 mmho/cm ²	v_4	30.0 mV
I_0	90.0 mA/cm ²		

Table 3: Changes to Table. 1 to put the ML model near a Hopf bifurcation.

that even a weak drive can lead to spike patterns that are nontrivially different. In order to put the system near the Hopf bifurcation, we use the parameters as in Table. 1, exchanging values with those indicated in Table. 3 where necessary.

We show, in Fig. 56, the relevant dynamical structures in this scenario. In this regime, the system has a stable limit cycle, shown in blue, which we take as Γ . Each time an orbit passes near the right-most point of Γ , we think of it as producing a ‘spike’. The system is bistable; the fixed point, depicted by the green point, from which the limit cycle is borne out of, is also stable in this regime, and is surrounded by an unstable limit cycle, shown in red, which separates the respective basins of attractions. The repelling orbit is close to Γ , but since it is a positive distance away, it will not show up in iPRC considerations. In a similar vein to Sect. 4.5, we use a white noise perturbation of the form $I(t) = I_0 + \sigma\xi(t)$.

To more clearly observe the difference in spike patterns, we plot the voltage traces, rather than the trajectories in phase space, as this allows us to see the time dependence of v . In Fig. 57, we plot the voltage traces of the full model and iPRC

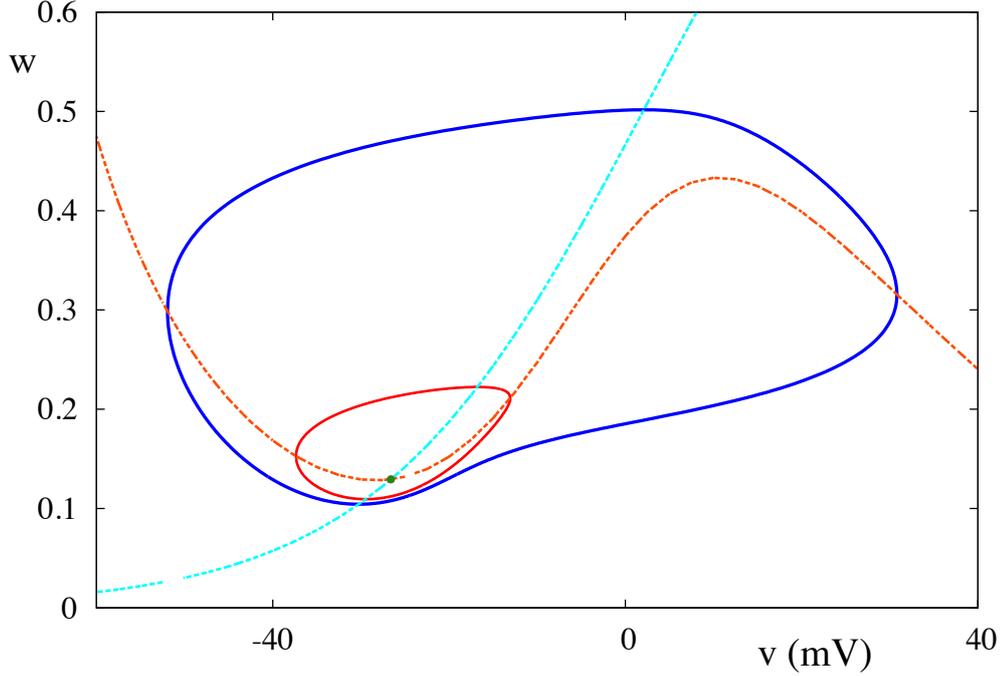


Figure 56: Phase portrait for the ML model in the Hopf regime. The system possesses a stable limit cycle (blue) surrounding an unstable cycle (red) which in turn surrounds a stable fixed point (green circle). Also shown are the nullclines of the system with the dashed orange curve representing the v -nullcline and the dashed cyan curve representing the w -nullclines. Here, $I_0 = 90 \text{ mA/cm}^2$, corresponding to an unforced period of about 102 ms.

predictions, computed at $I_0 = 90 \text{ mA/cm}^2$, under this forcing paradigm. The top panel shows the full ML simulation with a relatively weak forcing with $\sigma = 0.2$, with initial conditions starting on Γ . Since the basin of the fixed point is so close to Γ , a trajectory following Γ can easily get pushed into the basin, and then be attracted towards the fixed point. This can happen even with very small σ . The fixed point itself is further away from the unstable limit cycle than Γ , meaning that it is more difficult for a trajectory near the sink to escape, into \mathcal{B}_Γ , under weak σ . This is why, for weak noise such as $\sigma = 0.2$, it is easy to observe a transition from dynamics around Γ to those around the fixed point, but not easy to see transition in the reverse direction.

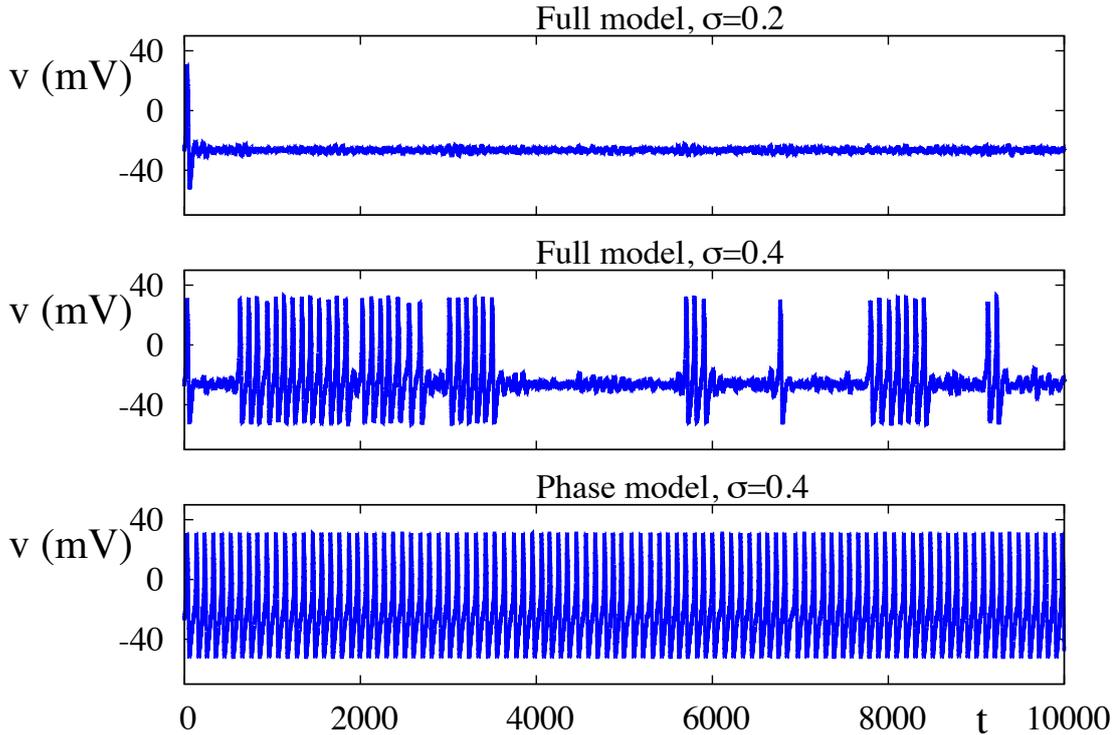


Figure 57: Voltage traces for the ML system in Hopf regime. The top two panels show simulation using the full ML model with the indicated level of stochastic forcing. The bottom panel shows the iPRC using $\sigma = 0.4$. We observe that, for the full model, the neuron falls almost immediately quiescent for $\sigma = 0.2$ and alternates between quiescence and spiking for $\sigma = 0.4$, neither of which captured by the phase-only model. Results for the phase-only model with $\sigma = 0.2$ are not shown as they are qualitatively similar to those for $\sigma = 0.4$. Here, $I_0 = 90$ mA/cm².

If we now increase the strength of the noise to be a little larger, letting $\sigma = 0.4$, the trajectory jumps back and forth more readily, as shown in the middle panel of Fig. 57. It roughly tracks Γ for some length of time, before crossing the unstable limit cycle, and staying near the fixed point. In terms of the voltage traces, this manifests itself as moving alternately between a spiking and quiescent mode, albeit at random times. Unsurprisingly, these perturbations do not have a pronounced effect on the voltage traces from the phase-only model, as can be seen in the

bottom panel of Fig. 57, with the spike patterns alternating only between regular and slightly irregular patterns. Finally, we note that in this bistable situation, a neuron can exhibit substantial sub-threshold activity. The iPRC underestimates the extent of such activity.

4.9 A network of two cells

Since the arguments presented here focus on a single forced oscillator, we may ask the question whether our results carry over to the network level where, in the weakly coupled regime, identical cells tend to either synchronise behaviour, or be in anti-phase. Considering the full dynamics of the system leads to the possibility of other patterns of activity emerging, due to effect of the other invariant structures, or shear effects in phase space. In this setting, our results preempt the failure of the iPRC to capture these effects, and we now conduct simulations of a network of two identical cells to illustrate these effects. We note that the product of the limit cycles for the unperturbed system may not actually be an invariant set for the coupled system, just as Γ_T is the attractor, rather than Γ for the forced ML oscillator considered earlier. However, since the limit cycle is hyperbolic, we expect it to be robust to weak coupling between units as so we may use it as a reference here.

First, we consider pulsatile coupling between cells in the homoclinic regime, mimicking that of the Sect. 4.4.2. Using the same spiking and reset conditions as in Sect. 4.5, neuron i sends a pulse of strength $|A|$ to neuron $j \neq i$ as v_i passes through $v_{\text{th}}=12.5$ mV, so that the dynamics of the two cell network are given by

$$C \frac{dv_i}{dt} = I_0 - \bar{g}_L(v_i - v_L) - \bar{g}_K w(v_i - v_K) - \bar{g}_{Ca} m_\infty(v_i)(v_i - v_{Ca}) + A \cdot \sum_{j \neq i} \delta(v_j - v_{\text{th}}), \quad (4.19)$$

$$\frac{dw_i}{dt} = \phi(w_\infty(v_i) - w_i) / \tau_w(v_i), \quad i = 1, 2. \quad (4.20)$$

For small $|A|$, the network tends to synchronise behaviour, as is predicted by the phase analogue. However, if we set $A = -2.0$, we enter the chaotic regime

as shown in Fig. 51, remembering that the natural period of the cells is around 25 ms. In this scenario, we observe a chaotic response, with the oscillators not phase locking and find a positive MLE of $\Lambda = 0.3937$. The trajectories in this situation are depicted in Fig. 58, as both voltage traces and trajectories in the phase plane. If we consider gap junction coupling, rather than pulsatile coupling between the cells, we can observe a more subtle effect of the saddle. Making this change simply replaces the last term in (4.19) with $A \cdot (v_j - v_i)$, where $A > 0$. Predictions from averaging theory suggest the anti-phase solution is stable for weak forcing. This can be seen by examining the odd part of the phase interaction function, H , as considered in Sect. 3.11. As before, the derivative of H_{odd} at the fixed points for the phase difference, ψ , will give the stability of these fixed points. Positive values for this derivative indicate unstable solutions, whereas negative values imply stable solutions. As shown in Fig. 59, the anti-phase solution, with $\psi = \Delta/2$ is stable. Indeed, this is what we observe for small values of A (say with $A = 0.001$). We show, in Fig. 60, the phase plane in this situation. Even though trajectories do not come close to the saddle of the unperturbed system, we can still see the effect it has on the synchronisation properties of the network. Examining closer the portion of the trajectories closest to the saddle, we see that the trajectory of one neuron, shown in orange, favours the ‘outside’ of the orbit, whilst the other neuron, shown in black, favours the ‘inside’ of the orbit (by orbit, we mean the orbit of the unperturbed system). We note that the tendency to phase-lock in an anti-phase mode increases the closer we are, in parameter space, to the homoclinic bifurcation, [174], further highlighting the role of the saddle in this case. As the strength of the coupling is increased, it is not unusual to observe an exchange of stability between the synchronous and anti-phase state in gap junction-coupled systems near a homoclinic bifurcation [175].

In our system, we have an even more drastic effect. Around $A \sim 0.04$, the coupling between the cells results in oscillator death for the network. Past this point, neither cell exhibits oscillatory activity, and trajectories will tend towards a fixed point of the network equations. In Fig. 61, we show the voltage traces for

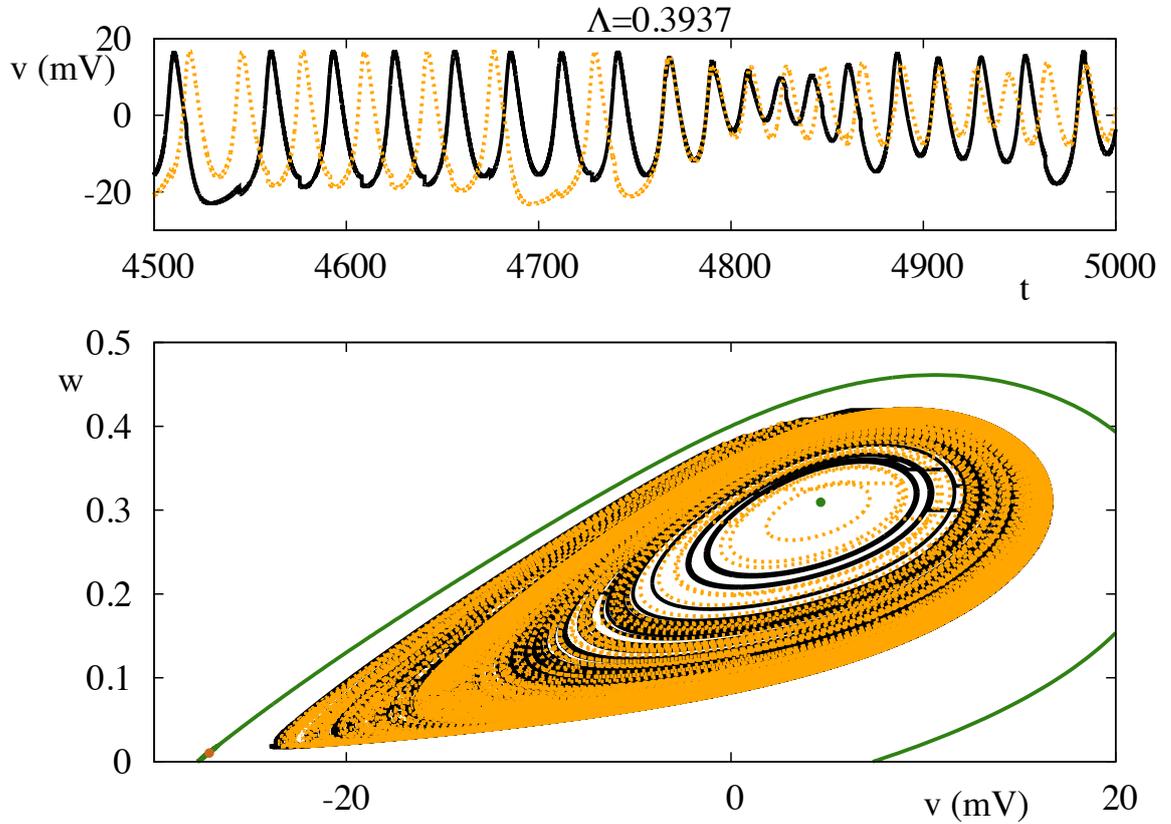


Figure 58: Top: Trajectories of the network of pulse-coupled ML cells. Bottom: Corresponding trajectories in the phase plane. Whilst the network is synchronous, as predicted by the phase model, for small $|A|$, we see that for larger values, here $A = -2.0$, the system is chaotic, with a positive MLE. We note that, since the natural period of the unperturbed oscillator is around 25 ms, we are in the parameter window for which we would expect chaos in the case of a single cell, forced periodically, as in Fig. 51.

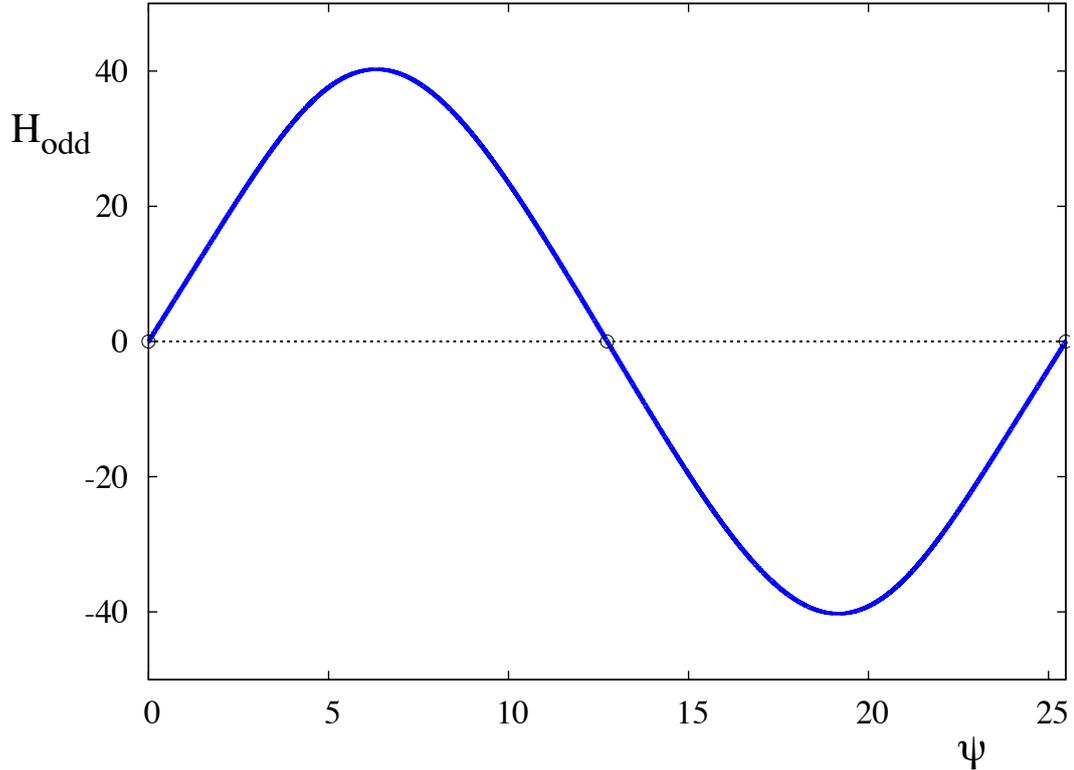


Figure 59: Odd part of the phase interaction function, H , as used in Sect. 3.11, for the case of two gap junction-coupled ML neurons in the homoclinic regime. The weakly coupled oscillatory theory predicts a stable anti-phase solution for this network, as can be seen by the negative derivative at the fixed point at $\psi = \Delta/2$.

$A = 0.04$ for both the full network model and the phase-only network. As predicted by averaging theory, the phase-only model exhibits a stable anti-phase solution, whilst the full model, initially oscillatory, becomes quiescent around $t = 3000$. Clearly, iPRCs will fail to capture both the change of stability between network states as well the death of oscillatory activity in the network.

Finally, if we make the gap junctions stochastic, exchanging the previous term for $A \cdot \xi \cdot (v_j - v_i)$, and use the parameters for the ML model in the Hopf regime, we can see the effect of leaving the basin of attraction of the full system. In this situation, phase theory predicts synchrony between the two cells, and this is precisely what we observe for small $A \sim 0.001$. As we increase A to 0.01, we still

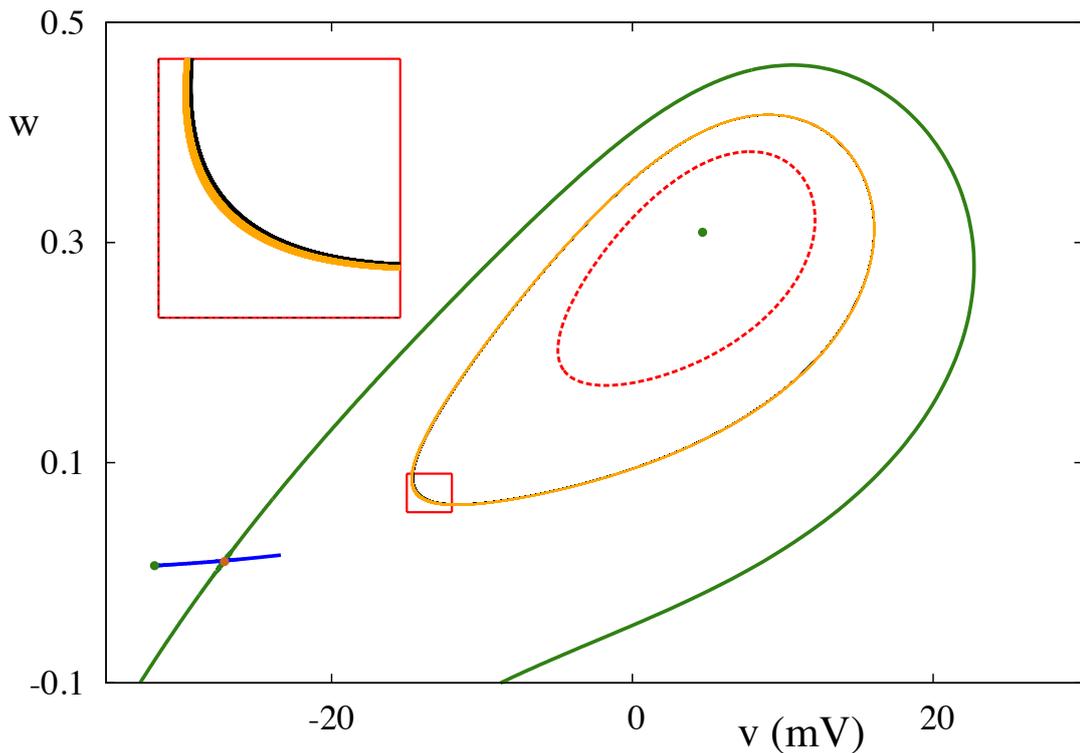


Figure 60: Trajectories of the network of ML cells with gap junction coupling near a homoclinic bifurcation. Although the orbits of neither cell come close the saddle, we can still see the effect it has on the dynamics of the network. The trajectory of one cell are consistently ‘outside’ the other as the network settles down to its phase locked state. The difference in times taken to complete one revolution promote asynchronous behaviour.

observe the tendency of the network to synchronise, but this is interrupted as one or both of the cells enters a quiescent phase, as shown in Fig. 62. Provided that the other cell remains in the spiking regime, the quiescent cell is likely to reenter the spiking regime, but the synchrony between the cells will have been lost. If both cells simultaneously enter the quiescent phase, the oscillations may be lost, since the forcing between the two cells is dependent on the voltage difference between them, which will approach zero if both cells are quiescent. Increasing A further makes it more likely that cells will reenter the spiking phase, but at a further loss

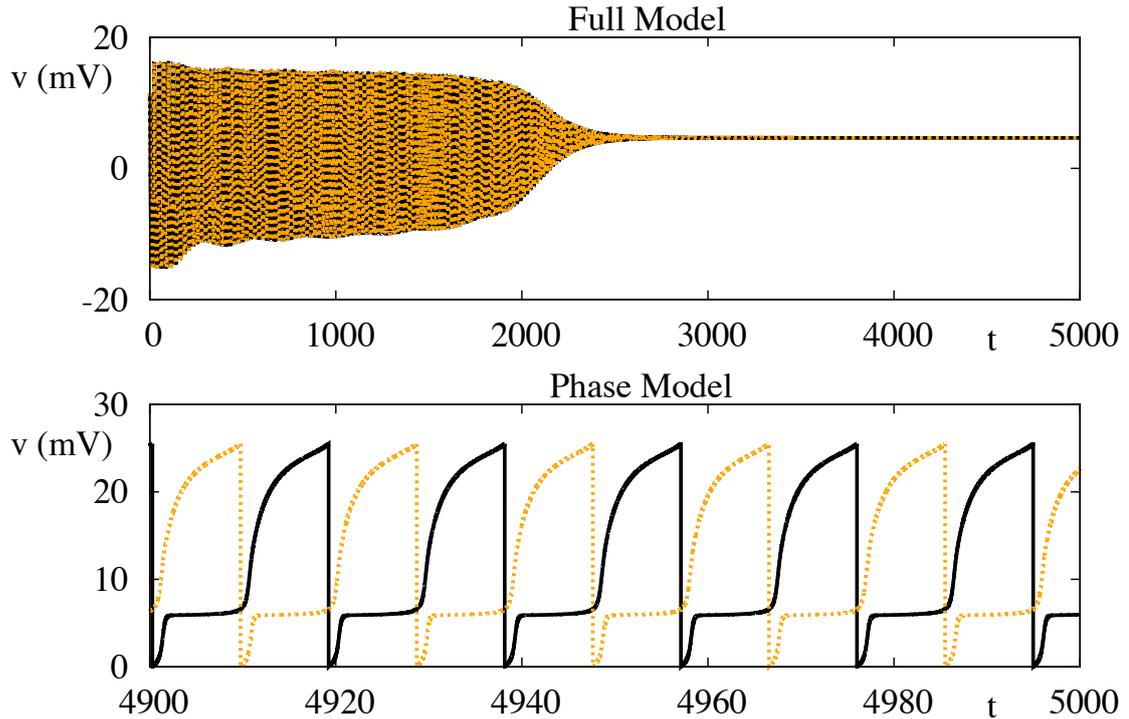


Figure 61: Top: Voltage traces for a gap junction-coupled ML network with $A = 0.04$. Bottom: Voltage traces for the phase-only model of the same system. We see that the phase model predicts a stable anti-phase solution between the two oscillators after transients have decayed, as shown in Fig. 59. However, in the full model, we see that a quiescent state, in which neither cell oscillates, is asymptotically stable.

of regularity of the network.

4.10 Discussion

This chapter compares the phase reduced models of an oscillating neuron's response to perturbation to the response of the full model. We have shown that, whilst the iPRC has the virtue of being analytically explicit, straightforward to compute and results in a large reduction in terms of dimension of network models, it may fail to capture more exotic behaviour of the full model, even under weak to moderate forcing. In these cases, the iPRC predicts regular oscillatory behaviour

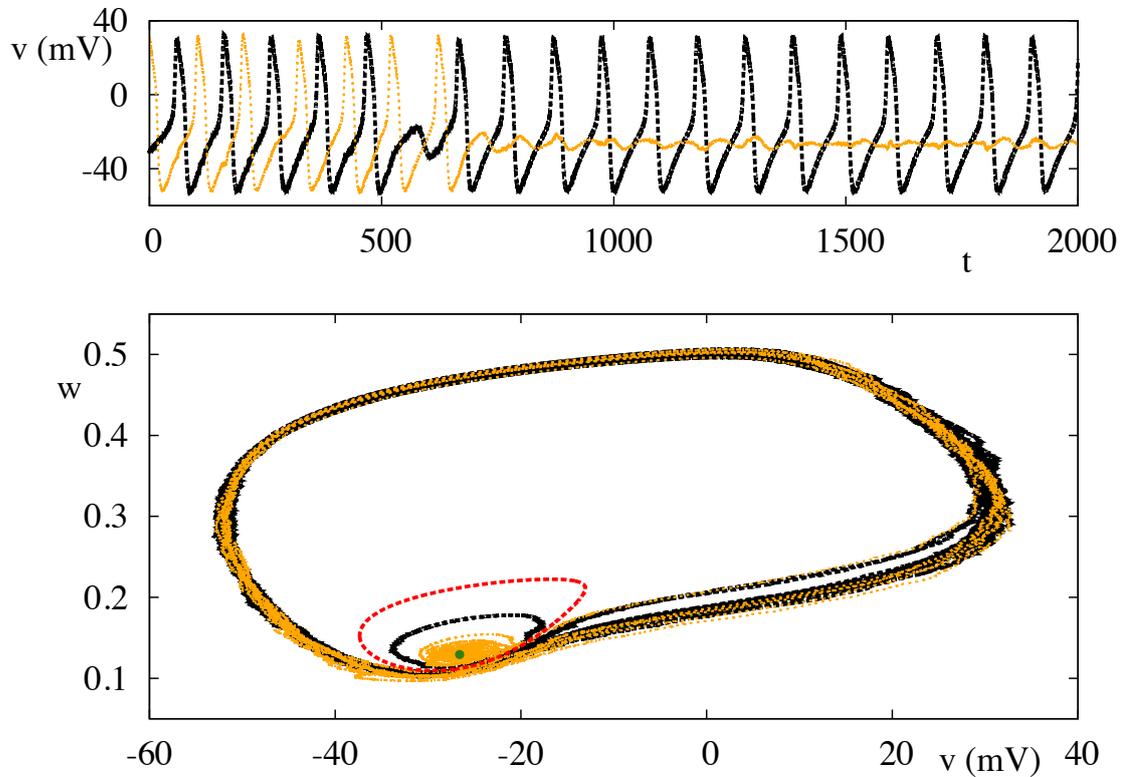


Figure 62: Top: Voltage traces of a network of ML cells with stochastic gap-junction coupling, now near a Hopf bifurcation. Bottom: Corresponding trajectories in the phase plane. In this regime, the cells tend to synchronise. However, given the proximity of \mathcal{B}_Γ to Γ in this case, trajectories may get forced into the basin of attraction of the stable node, as has occurred in this case. The quiescent cell now performs subthreshold oscillations, as the spiking cell still provides a perturbative current through the gap junction coupling, however, in this case, it is not strong enough to return to the cell to spiking. Since the forcing between cells is moderated by a stochastic process, the quiescent cell may reenter the spiking mode, at which point the cells will try to synchronise again.

when the full system is anything but. In moderate cases, this may simply mean that the phase model misrepresents the firing pattern and ISI distribution, and in more extreme cases, may predict regular behaviour when the full system is quiescent or chaotic. This is of particular importance when considering the neuroscience applications of these oscillatory models. In biological systems, the summative effect of inputs may not be weak and the attraction to the underlying oscillation may not be strong. In models in which we need either precise spike timing or a good estimate of firing rates, these perturbative effects may thus lead to predictions to be inaccurate, depending on the underlying phase space of the full model.

Specifically, we presented examples from the ML neuron model to show that

1. Periodic kicking of the ML system can lead to unreliable response in the full model via the mechanism of shear-induced chaos, contrary to iPRC predictions.
2. When stochastically driven, stickiness of nearby invariant structures can lead to lower firing rates and longer ISIs compared to iPRC predictions.
3. The forcing need not be strong to bring about serious discrepancies in firing patterns between full and phase-reduced models.

Moreover, in all the situations examined, the phase reduction itself offers no hint that any break down has occurred.

In terms of neural response, this means that iPRCs may overestimate spike-time reliability and firing rates; they may also underestimate the mean and variance in ISIs, and have a tendency to downplay sub-threshold activity. Thus, caution needs to be exercised when interpreting results that come from phase reduction arguments, especially for systems near bifurcation points. We have also shown simple numerical techniques, involving the examination of the evolution of a kick image under the unperturbed flow and straightforward numerical integration of stochastic systems, that can flag up the presence of shear and stickiness in a system geometrically, so that these can be examined before relying on predictions from iPRCs. While this chapter focuses on the ML model, the geometric ideas

are quite general. In particular, homoclinic and Hopf bifurcations are present in a range of neural and other biological models, and we predict similar phenomena to occur in such models. Simple numerical experiments on a network of two cells demonstrate that these phenomena carry over to the network level, and highlight the importance of prior considerations before phase reductions are applied to the network.

We will go on, in the remaining chapters, to try to introduce other methods, taking inspiration from iPRCs, to define coordinate systems, set up around the limit cycle, that explicitly include variables capturing the distance of points away from the limit cycle. In this case, we can capture the effect of invariant structures near to, but at a positive distance away from Γ , and analyse their effect on trajectories of the system.

5 A phase-amplitude description of neural oscillators

In the previous chapter, we highlighted scenarios in which the phase-only description of a neural oscillator was not sufficient to describe the dynamics of the full system. Of course, one could argue that the phase reduction is only appropriate when the strength of external forcing, or of coupling to other cells is weak, and the cells quickly relax back into their unperturbed rhythm. However, in the situations we considered, the forcing was not physiologically unreasonable. It was shown that the presence of other invariant structures in phase space can have a pronounced effect on the emergent behaviour of the system, suggesting that one must consider the geometry of phase space before performing the reduction. These effects are likely to impact the dynamics at the network level too, so that predictions on network behaviour using iPRCs quickly become erroneous. The failure of a phase description is in itself no surprise and underlies why the community emphasises the use of the word *weakly* in the phrase ‘weakly connected neural networks’. Indeed, the assumption that phase alone is enough to capture the essentials of neural response is one made more for mathematical convenience rather than being physiologically motivated. In spite of this, the computational advantage of phase-based models means that the study of many synchronisation phenomena are carried out under this assumption, leading to possibly erroneous predictions as we move away from the assumption of weak coupling. In the worst cases, using the phase-only model may miss oscillator death [176], after which the network ceases to oscillate. In other cases, the phase-only model may make errors in the prediction of stability of phase locked states, as discussed in Sect. 4.9. In this chapter, we generalise the phase description that allows one to track the evolution of distance from the cycle as well as phase on cycle. We use a classical technique from the theory of ordinary differential equations that makes use of a moving coordinate system to analyse periodic orbits. The subsequent *phase-amplitude* description is shown to be very well suited to understanding the response of the oscillator to external

stimuli (which are not necessarily weak). We consider a number of examples of neural oscillator models, ranging from planar through to high dimensional models, to illustrate the effectiveness of this approach in providing an improvement over the standard phase-reduction technique. As an explicit application of this phase-amplitude framework, we consider in some detail the response of a generic planar model where the strong-attraction assumption does not hold, and examine the response of the system to periodic pulsatile forcing.

Progress has previously been made to alleviate the effect of using only a phase description, by taking higher order approximations to the isochrons themselves to facilitate the construction of higher order iPRCs [177]. Even using perfect information about isochrons, the phase reduction still assumes persistence of the limit-cycle and instantaneous relaxation back to cycle. However, as we saw in the preceding chapter, the presence of nearby invariant phase-space structures such as (unstable) fixed points and invariant manifolds may result in trajectories spending long periods of time away from the limit cycle. Moreover, for gap junction coupling between ML neurons in the homoclinic regime, as considered in Sect. 4.9, increasing the coupling strength too much results in the destruction of any oscillatory behaviour of the network, a phenomenon known as oscillator death [176].

Strong forcing will necessarily take one away from the neighbourhood of a cycle where a phase description is expected to hold. Thus, developing a reduced description which captures some notion of *distance from cycle* is a key component of any theory of forced limit cycle oscillators. The development of phase-amplitude models that better characterise the response of popular high dimensional single neuron models is precisely the topic of this chapter. Given that it is a major challenge to construct an isochronal foliation, we use non-isochronal phase-amplitude coordinates as a practical method for obtaining a more accurate description of neural systems. Recently, Medvedev has used this approach to understand in more detail the synchronisation of linearly coupled stochastic limit cycle oscillators [178].

First, we consider a general coordinate transformation which recasts the dynamics of a system in terms of phase-amplitude coordinates. This approach is

directly taken from the classical theory for analysing periodic orbits of ODEs, originally considered for planar systems in [179], and for general systems in [12]. We advocate it here as one way to move beyond a purely phase-centric perspective. We illustrate the transformation by applying it to a range of popular neuron models. Following this, we consider how inputs to the neuron are transformed under these coordinate transformations and derive the evolution equations for the forced phase-amplitude system. Importantly, we show that the behaviour of the phase-amplitude system is much more able to capture that of the original single neuron model from which it is derived. Finally, we discuss the relevance of this work to developing a theory of network dynamics that can improve upon the standard weak coupling approach. The bulk of this chapter forms the recently published [180].

5.1 Phase-amplitude coordinates

Throughout this chapter, we remain focused on the study of the dynamics prescribed by the system $\dot{x} = f(x) + \varepsilon g(x, t)$, $x \in \mathbb{R}^n$, with solutions $x = x(t)$ that satisfy $x(0) = x_0 \in \mathbb{R}^n$. As before, we will assume that the system with $\varepsilon = 0$ admits an attracting hyperbolic periodic orbit, Γ , (namely, a periodic orbit having one zero Floquet exponent whilst its other Floquet exponents have negative real part), with period Δ , such that $\Gamma(t) = \Gamma(t + \Delta)$. We parametrise Γ by a phase variable $\theta \in [0, \Delta)$, so that $\dot{\theta} = 1$ for $x \in \Gamma$. It has long been known in the dynamical systems community how to construct a coordinate system based on this notion of phase as well as a *distance* from cycle, see [12] for a discussion. In fact, Ermentrout and Kopell [181] made good use of this approach to derive the phase-interaction function for networks of weakly connected limit-cycle oscillators in the limit of infinitely fast attraction to cycle. However, this assumption is particularly extreme and unlikely to hold for a broad class of single neuron models. Thus, it is interesting to return to the full phase-amplitude description. In essence, the transformation to these coordinates involves setting up a *moving orthonormal system* around the limit cycle. One axis of this system is chosen to be in the direction of

the tangent vector along the orbit, and the remaining are chosen to be orthogonal. We introduce the normalised tangent vector ξ as

$$\xi(\theta) = \frac{d\Gamma}{d\theta} / \left| \frac{d\Gamma}{d\theta} \right|. \quad (5.1)$$

The remaining coordinate axes are conveniently grouped together as the columns of an $n \times (n - 1)$ matrix ζ . In this case we can write an arbitrary point x as

$$x(\theta, \rho) = \Gamma(\theta) + \zeta(\theta)\rho, \quad (5.2)$$

where $\theta \in [0, \Delta)$ and $\rho \in \mathbb{R}^{n-1}$. In the above, $|\rho|$ represents the Euclidean distance from the limit cycle. A caricature, in \mathbb{R}^2 , of the coordinate system along an orbit segment is shown in Fig. 63. Through the use of the variable ρ , we can consider points away from the periodic orbit. Rather than being isochronal, lines of constant θ are simply straight lines that emanate from a point on the orbit in the direction of the normal. We remark that, for planar systems, we choose to have ζ pointing outside of Γ so that points outside Γ have positive ρ . Using (5.2), we may then go on to derive the equations of motion for the variables (θ, ρ) .

5.1.1 Unforced systems

We first consider the case of an unforced oscillator. That is, we set $\varepsilon = 0$, so that we start from

$$\dot{x} = f(x), \quad x \in \mathbb{R}^n. \quad (5.3)$$

We make the transformation $x(t) = \Gamma(\theta(t)) + \zeta(\theta(t))\rho(t)$, giving

$$\left[\frac{d\Gamma(\theta)}{d\theta} + \frac{d\zeta(\theta)}{d\theta}\rho \right] \dot{\theta} + \zeta(\theta)\dot{\rho} = f(\Gamma(\theta) + \zeta(\theta)\rho). \quad (5.4)$$

We proceed by projecting (5.4) onto $\xi(\theta)$, using (5.1). The left hand side of (5.4) now reads:

$$\left[\left| \frac{d\Gamma}{d\theta} \right| + \xi^T \frac{d\zeta}{d\theta} \rho \right] \frac{d\theta}{dt}, \quad (5.5)$$

where ξ^T denotes the transpose of ξ and the right hand side of (5.4) becomes

$$\xi^T f(\Gamma + \zeta\rho) = \left[\left| \frac{d\Gamma}{d\theta} \right| + \xi^T \frac{d\zeta}{d\theta} \rho \right] + \xi^T f(\Gamma + \zeta\rho) - \xi^T f(\Gamma) - \xi^T \frac{d\zeta}{d\theta} \rho. \quad (5.6)$$

Thus,

$$\dot{\theta} = 1 + f_1(\theta, \rho), \quad (5.7)$$

where

$$f_1(\theta, \rho) = -h^T(\theta, \rho) \frac{d\zeta}{d\theta} \rho + h^T(\theta, \rho) [f(\Gamma + \zeta\rho) - f(\Gamma)], \quad (5.8)$$

with

$$h(\theta, \rho) = \left[\left| \frac{d\Gamma}{d\theta} \right| + \xi^T \frac{d\zeta}{d\theta} \rho \right]^{-1} \xi(\theta). \quad (5.9)$$

Upon projecting both sides of (5.4) onto ζ , the left hand side reads

$$\zeta^T \left[\frac{d\Gamma}{d\theta} + \frac{d\zeta}{d\theta} \rho \right] \frac{d\theta}{dt} + \frac{d\rho}{dt} = \zeta^T \frac{d\zeta}{d\theta} \rho \frac{d\theta}{dt} + \frac{d\rho}{dt} = \zeta^T \frac{d\zeta}{d\theta} \rho [1 + f_1(\theta, \rho)] + \frac{d\rho}{dt}, \quad (5.10)$$

whilst the right hand side becomes

$$\zeta^T f(\Gamma + \zeta\rho) = -\zeta^T f(\Gamma) + \zeta^T Df \zeta \rho - \zeta^T Df \zeta \rho + \zeta^T f(\Gamma + \zeta\rho), \quad (5.11)$$

since $\zeta^T f(\Gamma) = \zeta^T d\Gamma/d\theta = 0$ and where Df denotes the Jacobian of f . Putting together (5.10) and (5.11) yields

$$\dot{\rho} = A(\theta)\rho + f_2(\theta, \rho), \quad (5.12)$$

where

$$A(\theta) = \zeta^T \left[-\frac{d\zeta}{d\theta} + Df \zeta \right], \quad (5.13)$$

and

$$f_2(\theta, \rho) = -\zeta^T \frac{d\zeta}{d\theta} \rho f_1 + \zeta^T [f(\Gamma + \zeta\rho) - f(\Gamma) - Df \zeta \rho]. \quad (5.14)$$

It may be easily seen that $f_1(\theta, \rho) = O(\rho)$ as $\rho \rightarrow 0$, that $f_2(\theta, 0) = 0$ and $\partial f_2(\theta, 0)/\partial \rho = 0$. Overall, combining (5.7) and (5.12) we arrive at the transformed equations of the unforced system:

$$\begin{aligned} \dot{\theta} &= 1 + f_1(\theta, \rho), \\ \dot{\rho} &= A(\theta)\rho + f_2(\theta, \rho). \end{aligned} \quad (5.15)$$

In the above, f_1 captures the *shear* present in the system, that is, whether the speed of θ increases or decreases dependent on the distance from cycle. Additionally,

$A(\theta)$ describes the θ -dependent rate of attraction or repulsion from cycle, whilst f_2 represents a correction term to the otherwise linear ρ dynamics captured by A . We note that even though we refer to f_2 as a correction term, (5.15) is an exact transformation of (5.3), written in (θ, ρ) coordinates in a ‘tube’ around Γ . The width of the tube is set by a θ -dependent value of $|\rho|$, at which point the transformation ceases to be invertible, and so breaks down. We go on to discuss the break down of the transformation in Sect. 5.1.2.

Although, the (θ, ρ) coordinates are not isochronal, they offer an easy, practical way to study the phase and amplitude of a trajectory relative to the underlying periodic orbit. Furthermore, this coordinate system is still defined outside the basin of attraction of Γ . In fact, we may represent other stable attractors in (θ, ρ) coordinates. This may be useful when considering the role of, say, the saddle and stable nodes of the ML model in the homoclinic regime, or another stable limit cycle of the ML model in the Hopf regime.

Since ρ captures the Euclidean distance from cycle along the various coordinates which comprise ζ , it is sensible to rescale variables so that the original variables are of the same order of magnitude. In many neural models, some of the variables represent gating variables, indicating the fraction of ion channels that are open. As such, these variables can only take values on $[0, 1]$, whereas variations in the variable for membrane potential along an orbit are typically $O(100)$. To deal with this discrepancy in magnitude amongst the variables, we rescale the components using the affine transformation $X_i = \alpha_i x_i$, where the index indicates the i 'th component of the state variable and

$$\alpha_i = \frac{\max\{\Gamma_1(\theta)\} - \min\{\Gamma_1(\theta)\}}{\max\{\Gamma_i(\theta)\} - \min\{\Gamma_i(\theta)\}}, \quad (5.16)$$

where Γ_i refers to the i 'th component of the orbit Γ . Typically, Γ_1 corresponds to the variable representing membrane potential. The units of the α_i are chosen to be consistent with the original system (5.3).

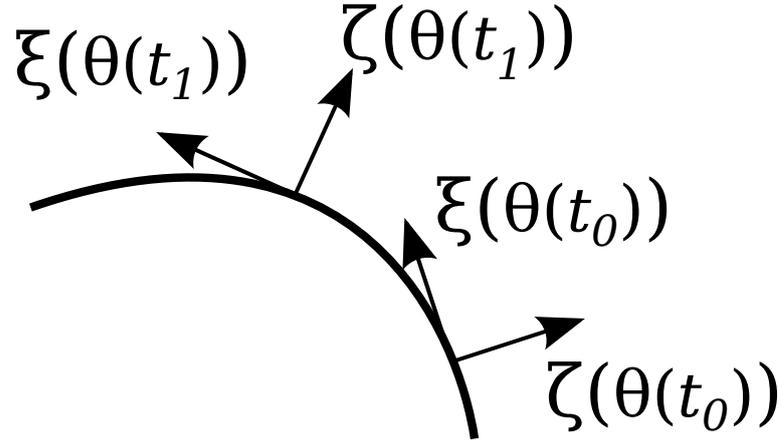


Figure 63: Demonstration of the moving orthonormal coordinate system along an orbit segment. As t evolves from t_0 to t_1 , the coordinates vary smoothly. In this planar example, ζ always points to the outside of the orbit.

5.1.2 Computing A , f_1 and f_2 in terms of the original vector field

For computational ease when simulating the system, we may wish to write A , f_1 and f_2 in terms of the original vector field f (and its derivatives). We do this for planar models, though the extension to higher dimensional systems is straightforward. Denoting the components of the original state variable as $x = (x_1, x_2)$, we first define an automorphism $\mathcal{F} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ as

$$\mathcal{F}(x) = (x_2, -x_1), \quad (5.17)$$

noting that the derivative $\partial\mathcal{F}/\partial x$ is given by

$$D\mathcal{F} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} = \mathcal{K}, \quad (5.18)$$

and then defining a scalar function $\mathcal{H} : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$

$$\mathcal{H}(x) = |x|, \quad (5.19)$$

whose derivative is given by

$$D\mathcal{H} = \frac{x^T}{\mathcal{H}(x)}. \quad (5.20)$$

We may now write ζ as

$$\zeta = \mathcal{F}(\xi(\theta)) \quad (5.21)$$

so that

$$\frac{d\zeta}{d\theta} = D\mathcal{F}(\xi(\theta)) \frac{d\xi}{d\theta}, \quad (5.22)$$

by the chain rule. We have that $D\mathcal{F} = \mathcal{K}$, so we simply need to evaluate $d\xi/d\theta$.

We proceed by writing ξ in the form:

$$\xi = \left(\frac{f^1(x_1, x_2)}{\mathcal{H}(f^1(x_1, x_2), f^2(x_1, x_2))}, \frac{f^2(x_1, x_2)}{\mathcal{H}(f^1(x_1, x_2), f^2(x_1, x_2))} \right)^T, \quad (5.23)$$

where f^i is the i 'th component of the original vector field f . Consider the first component ξ_1 of ξ . We have that

$$\frac{d\xi_1(x)}{d\theta} = D\xi_1(x) \frac{dx}{d\theta}. \quad (5.24)$$

By the quotient rule for vectors,

$$D\xi_1 = \frac{\mathcal{H}Df^1 - f^1D\mathcal{H}}{\mathcal{H}^2}. \quad (5.25)$$

Define $J(x) = Df(x)$ as the Jacobian of the original vector field. Then Df^1 is simply given by the first row of J , which we denote by $J_{1,\bullet}$. Similarly, Df^2 would be given by $J_{2,\bullet}$ and so on. Finally, we need to evaluate $D\mathcal{H}$. By the chain rule, the components of $D\mathcal{H}$ are given by

$$\begin{aligned} D\mathcal{H} &= \left[\frac{\partial \mathcal{H}}{\partial f^1} \frac{\partial f^1}{\partial x_1} + \frac{\partial \mathcal{H}}{\partial f^2} \frac{\partial f^2}{\partial x_1}, \frac{\partial \mathcal{H}}{\partial f^1} \frac{\partial f^1}{\partial x_2} + \frac{\partial \mathcal{H}}{\partial f^2} \frac{\partial f^2}{\partial x_2} \right] \\ &= \left[\frac{1}{\mathcal{H}} \frac{d\Gamma}{d\theta} \cdot J_{\bullet,1}, \frac{1}{\mathcal{H}} \frac{d\Gamma}{d\theta} \cdot J_{\bullet,2} \right]. \end{aligned} \quad (5.26)$$

Writing $d\zeta/d\theta$ as above, we may then write A as

$$\begin{aligned} A(\theta) &= \frac{1}{\mathcal{H}(\Gamma(\theta))} \left\{ f^2(\Gamma(\theta)) \frac{d\zeta_1}{d\theta} - f^1(\Gamma(\theta)) \frac{d\zeta_2}{d\theta} \right\} \\ &\quad + \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \left\{ f^2(\Gamma(\theta)) \left(J_{11}(\Gamma(\theta)) f^2(\Gamma(\theta)) - J_{12}(\Gamma(\theta)) f^1(\Gamma(\theta)) \right) \right. \\ &\quad \left. - f^1(\Gamma(\theta)) \left(J_{21}(\Gamma(\theta)) f^2(\Gamma(\theta)) + J_{11}(\Gamma(\theta)) f^1(\Gamma(\theta)) \right) \right\}, \end{aligned} \quad (5.27)$$

so that A is expressed only in terms of the original vector field f and the periodic orbit Γ . We may proceed, in a similar fashion, to find expressions for f_1 and f_2 . We first note that h may be written as

$$h(\theta, \rho) = \frac{\left(f^1(\Gamma(\theta)), f^2(\Gamma(\theta))\right)}{\mathcal{H}(\Gamma(\theta))^2 + \rho \left(f^1(\Gamma(\theta)) \frac{d\zeta_1}{d\theta} + f^2(\Gamma(\theta)) \frac{d\zeta_2}{d\theta}\right)}, \quad (5.28)$$

and we shall denote the denominator of this expression by $D(\theta, \rho)$. We may then compute f_1 as

$$f_1(\theta, \rho) = \frac{1}{D(\theta, \rho)} \left\{ f^1(\Gamma(\theta)) \left(-\rho \frac{d\zeta_1}{d\theta} + f^1(\Gamma(\theta) + \zeta(\theta)\rho) - f^1(\Gamma(\theta)) \right) + f^2(\Gamma(\theta)) \left(-\rho \frac{d\zeta_2}{d\theta} + f^2(\Gamma(\theta) + \zeta(\theta)\rho) - f^2(\Gamma(\theta)) \right) \right\}, \quad (5.29)$$

and f_2 as

$$f_2(\theta, \rho) = \frac{1}{\mathcal{H}(\Gamma(\theta))} \left\{ f^2(\Gamma(\theta)) \left(\rho \frac{d\zeta_1}{d\theta} f_1(\theta, \rho) + f^1(\Gamma(\theta) + \zeta(\theta)\rho) - f^1(\Gamma(\theta)) \right) + \rho \frac{J_{11}(\Gamma(\theta))f^2(\Gamma(\theta)) - J_{12}(\Gamma(\theta))f^1(\Gamma(\theta))}{\mathcal{H}(\Gamma(\theta))} - f^1(\Gamma(\theta)) \left(\rho \frac{d\zeta_2}{d\theta} f_1(\theta, \rho) + f^2(\Gamma(\theta) + \zeta(\theta)\rho) - f^2(\Gamma(\theta)) \right) + \rho \frac{J_{21}(\Gamma(\theta))f^2(\Gamma(\theta)) - J_{22}f^1(\Gamma(\theta))}{\mathcal{H}(\Gamma(\theta))} \right\}. \quad (5.30)$$

In order to evaluate the functions f_1, f_2 and A for models with dimension larger than two, we need to calculate $d\zeta/d\theta$. Defining by $\nu_i(\theta)$, the direction angles of $\xi(\theta)$, we have that

$$\zeta_i = e_i - \frac{\cos \nu_i}{1 + \cos \nu_1} (e_1 + \xi(\theta)) = e_i - \frac{e_i \cdot \xi(\theta)}{1 + e_1 \cdot \xi(\theta)} (e_1 + \xi(\theta)), \quad i = 2, 3, \dots, n, \quad (5.31)$$

where the index i denotes the column entry of ζ and $x \cdot y$ denotes the dot product between vectors x and y . Defining

$$u_i(\theta) = \frac{e_i \cdot \xi(\theta)}{1 + e_1 \cdot \xi(\theta)}, \quad (5.32)$$

and

$$w_j(\theta) = e_{1,j} + \xi_j(\theta), \quad (5.33)$$

where j denotes the row index, we have

$$\frac{d\zeta_{i,j}}{d\theta} = -u_i \frac{dw_j}{d\theta} - w_j \frac{du_i}{d\theta}. \quad (5.34)$$

By the quotient rule for vectors we find that,

$$\frac{du_i}{d\theta} = \frac{(1 + e_1) (e_i \cdot \frac{d\xi}{d\theta}) - (e_i \cdot \xi(\theta)) (e_1 \cdot \frac{d\xi}{d\theta})}{(1 + e_1 \cdot \xi(\theta))^2}, \quad (5.35)$$

and that

$$\frac{dw_j}{d\theta} = \frac{d\xi_j}{d\theta}. \quad (5.36)$$

Overall, we have that

$$\frac{d\zeta_{i,j}}{d\theta} = -\frac{e_i \cdot \xi(\theta)}{1 + e_1 \cdot \xi(\theta)} \frac{d\xi_j}{d\theta} - (e_{1,j} + \xi_j(\theta)) \left(\frac{(1 + e_1) (e_i \cdot \frac{d\xi}{d\theta}) - (e_i \cdot \xi(\theta)) (e_1 \cdot \frac{d\xi}{d\theta})}{(1 + e_1 \cdot \xi(\theta))^2} \right). \quad (5.37)$$

It is pertinent to consider where this coordinate transformation breaks down, that is, where the determinant of the Jacobian of the transformation $K = \det[\partial x/\partial\theta \ \partial x/\partial\rho]$ vanishes. This never vanishes on-cycle (where $\rho = 0$), but may do so for some $|\rho| = k > 0$. This sets an upper bound on how far away from the limit cycle we can describe the system using these phase-amplitude coordinates. In Fig. 64, we plot the curve along which the transformation breaks down for the ML model. We observe that, for some values of θ , k is relatively smaller. The breakdown occurs where lines of constant θ cross and thus the transformation is not invertible. These values of θ correspond to points along which the orbit has high curvature. This may be seen more clearly in Fig. 65, in which we show the curve along which $K = 0$ in (v, w) space. The colours along the curve indicate the reference point on Γ for that value of θ . We see that the transformation breaks down inside the orbit, near the ‘corners’ of Γ . We note that although this curve has some points which look like they have arbitrarily small ρ , the reference point is actually on the opposite side of Γ so that the transformation breaks down at a finite, but non-zero $|\rho|$. We note that this issue is less problematic in higher dimensional models.

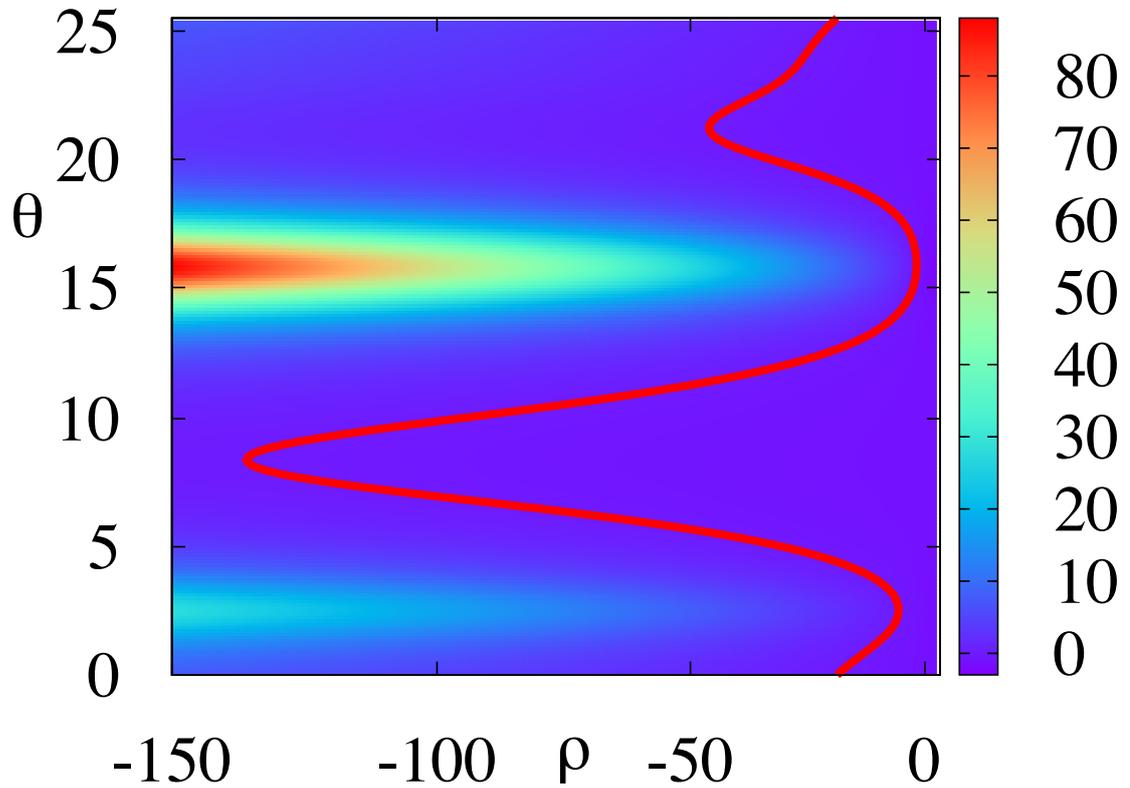


Figure 64: Colours show the value of the determinant, K , of the phase-amplitude transformation for the ML model (2.28)-(2.29). The red curve indicates where $K = 0$, and thus where the coordinate transformation breaks down. Note how all of the values for which this occurs have $\rho < 0$, referring to points on the inside of Γ . Parameter values are as for the homoclinic regime as discussed in Sect. 4.4.2. This red curve is shown in (v, w) coordinates in Fig. 65.

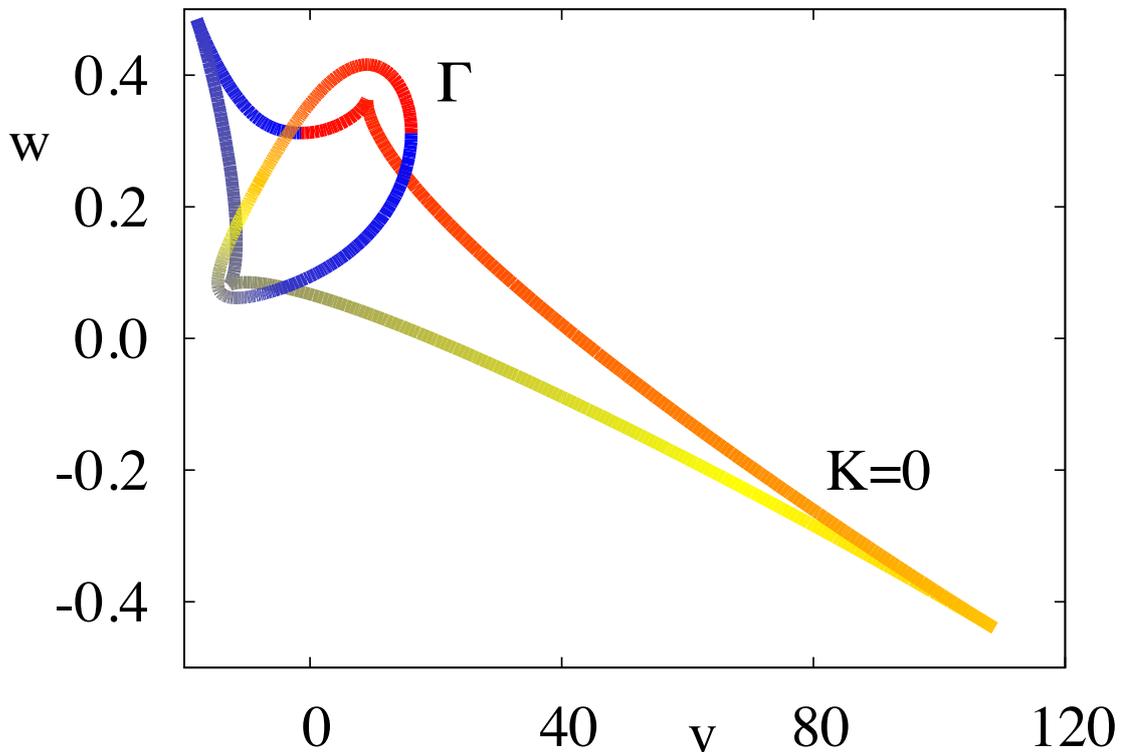


Figure 65: This figure shows the determinant K of the phase-amplitude transformation for the ML model, (2.28)-(2.29), in (v, w) coordinates, along with Γ . The orbit is as in Sect. 4.4.2, with colours specifying the phase along the orbit. The curve $K = 0$ indicates where the transformation breakdown occurs, and is the same as the red curve in Fig. 64. Along this curve, the transformation is not invertible. This curve is also coloured so that the reference point along Γ is identified with the same colour. This means that although the curve $K = 0$ appears to cross the orbit (so that the breakdown would occur for trivial ρ), these points actually have a finite and negative ρ . We see that the breakdown occurs near ‘corners’ of Γ , that is, where the orbit has high curvature. We note that this breakdown occurs when lines of constant θ intersect one another so that the mapping between (θ, ρ) and (v, w) is no longer one-to-one.

5.1.3 Forced systems

We now consider the transformation of forcing terms, allowing ε to be non-zero, so that we now study the system

$$\dot{x} = f(x) + \varepsilon g(x, t). \quad (5.38)$$

Following the same steps as in the derivation of (5.15), we first project (5.38) onto ξ , so that

$$\dot{\theta} = 1 + f_1(\theta, \rho) + \varepsilon h^T(\theta, \rho)g(\Gamma(\theta) + \zeta(\theta)\rho, t), \quad (5.39)$$

with h and f_1 as before. Upon projecting (5.38) onto ζ , the left hand side reads

$$\zeta^T \frac{d\zeta}{d\theta} \rho \left[1 + f_1(\theta, \rho) + \varepsilon h^T(\theta, \rho)g(\Gamma + \zeta\rho, t) \right] + \frac{d\rho}{dt}, \quad (5.40)$$

whereas the right hand side of (5.38) is

$$-\zeta^T f(\Gamma) + \zeta^T Df\zeta\rho - \zeta^T Df\zeta\rho + \zeta^T f(\Gamma + \zeta\rho) + \varepsilon \zeta^T g(\Gamma + \zeta\rho, t). \quad (5.41)$$

Putting the previous two equations together gives

$$\dot{\rho} = A(\theta)\rho + f_2(\theta, \rho) + \varepsilon \zeta^T B(\theta, \rho)g(\Gamma(\theta) + \zeta(\theta)\rho, t), \quad (5.42)$$

where A and f_2 are as before and

$$B(\theta, \rho) = \left[I_n - \frac{d\zeta(\theta)}{d\theta} \rho h(\theta, \rho)^T \right], \quad (5.43)$$

where I_n is the identity matrix in \mathbb{R}^n . Overall, combining (5.39) and (5.42) we arrive at the transformed system:

$$\begin{aligned} \dot{\theta} &= 1 + f_1(\theta, \rho) + \varepsilon h^T(\theta, \rho)g(\Gamma(\theta) + \zeta(\theta, t)\rho, t), \\ \dot{\rho} &= A(\theta)\rho + f_2(\theta, \rho) + \varepsilon \zeta^T B(\theta, \rho)g(\Gamma(\theta) + \zeta(\theta)\rho, t). \end{aligned} \quad (5.44)$$

The functions h and B thus show how inputs to the system $\dot{x} = f(x)$ are transformed to the (θ, ρ) coordinate system. For planar systems, we can show that $\zeta^T B(\theta, \rho) = \zeta^T$.

To see this, let us write $h = (h_1, h_2)$ for some scalar values h_1, h_2 . Thus,

$$\zeta^T B = \zeta^T - \rho \zeta^T \begin{pmatrix} \frac{d\zeta_1}{d\theta} h_1 & \frac{d\zeta_1}{d\theta} h_2 \\ \frac{d\zeta_2}{d\theta} h_1 & \frac{d\zeta_2}{d\theta} h_2 \end{pmatrix} = \zeta^T - \rho \left(h_1 \zeta^T \frac{d\zeta}{d\theta} \quad h_2 \zeta^T \frac{d\zeta}{d\theta} \right) \quad (5.45)$$

We will now show that

$$\xi^T \frac{d\xi}{d\theta} = 0 \implies \zeta^T \frac{d\zeta}{d\theta} = 0, \quad (5.46)$$

so that the second term in (5.45) vanishes. We begin by evaluating $d\xi_i/d\theta$ as

$$\frac{d\xi_i}{d\theta} = \frac{1}{\mathcal{H}(\Gamma(\theta))} \frac{d^2\Gamma_i}{d\theta^2} - \frac{1}{\mathcal{H}(\Gamma(\theta))^3} \frac{d\Gamma_i}{d\theta} \frac{d\Gamma^T}{d\theta} \frac{d^2\Gamma}{d\theta^2}, \quad (5.47)$$

Thus,

$$\xi^T \frac{d\xi}{d\theta} = \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \sum_{i=1}^2 \frac{d\Gamma_i}{d\theta} \frac{d^2\Gamma_i}{d\theta^2} - \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \left(\frac{d\Gamma_i}{d\theta} \right)^2 \frac{d\Gamma^T}{d\theta} \frac{d^2\Gamma}{d\theta^2}. \quad (5.48)$$

After expanding the above equation and dropping a factor of $1/\mathcal{H}(\Gamma(\theta))^2$, we compare coefficients of the sum in $d^2\Gamma_i(\theta)/d\theta^2$, which are

$$\frac{d\Gamma_i}{d\theta} - \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \left\{ \left(\frac{d\Gamma_i}{d\theta} \right)^3 + \frac{d\Gamma_i}{d\theta} \left(\frac{d\Gamma_{j \neq i}}{d\theta} \right)^2 \right\}. \quad (5.49)$$

After taking out a factor of $1/\mathcal{H}(\Gamma(\theta))^2$, we have

$$\begin{aligned} & \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \left\{ \mathcal{H}(\Gamma(\theta))^2 \frac{d\Gamma_i}{d\theta} - \left(\frac{d\Gamma_i}{d\theta} \right)^3 - \frac{d\Gamma_i}{d\theta} \left(\frac{d\Gamma_{j \neq i}}{d\theta} \right)^2 \right\} \\ &= \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \left\{ \left(\left(\frac{d\Gamma_i}{d\theta} \right)^2 + \left(\frac{d\Gamma_{j \neq i}}{d\theta} \right)^2 \right) \frac{d\Gamma_i}{d\theta} - \left(\frac{d\Gamma_i}{d\theta} \right)^3 - \frac{d\Gamma_i}{d\theta} \left(\frac{d\Gamma_{j \neq i}}{d\theta} \right)^2 \right\} \\ &= \frac{1}{\mathcal{H}(\Gamma(\theta))^2} \left\{ \left(\frac{d\Gamma_i}{d\theta} \right)^3 + \frac{d\Gamma_i}{d\theta} \left(\frac{d\Gamma_{j \neq i}}{d\theta} \right)^2 - \left(\frac{d\Gamma_i}{d\theta} \right)^3 - \frac{d\Gamma_i}{d\theta} \left(\frac{d\Gamma_{j \neq i}}{d\theta} \right)^2 \right\} = 0, \end{aligned} \quad (5.50)$$

using the definition of \mathcal{H} . Thus, the dot product $\xi^T d\xi/d\theta = 0$, as is $\zeta^T d\zeta/d\theta$ and we have that $\zeta^T B = \zeta^T$ so that forcing in ρ depends only on θ for planar systems.

5.2 Some neural examples

To demonstrate the application of the above coordinate transformation, we now consider some popular single neuron models.

5.2.1 The ML model

We now return to the ML model and show how A , f_1 and f_2 are manifest in this model. Throughout this chapter, we use the parameter values which put the ML

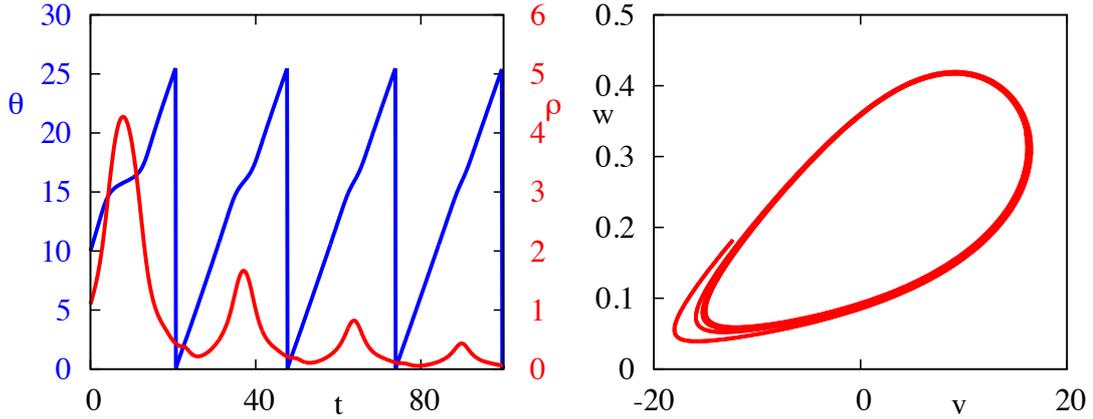


Figure 66: Typical trajectory of the ML model of the transformed system. Left: Time evolution of θ and ρ . Right: Trajectory plotted in the (v, w) phase plane. We see that when ρ has a local maximum, the evolution of θ slows down, corresponding to where trajectories pass near to the saddle.

model near a homoclinic bifurcation, as used in Sect. 4.4.2 (given by Table. 1). As the ML model is planar, ρ is a scalar, as are the functions A and $f_{1,2}$. This allows us to use the moving coordinate system to clearly visualise parts of phase space where trajectories are attracted towards the limit cycle, and parts in which they move away from it, as illustrated in Fig. 66. The functions $f_{1,2}$ and A , evaluated at $\rho = -0.1$ are shown in Fig. 67. The evolution of θ is mostly constant, however we clearly observe portions of the trajectories where this is slowed, along which, $\dot{\rho} \approx 0$. In fact, this corresponds to where trajectories pass near to the saddle, and the dynamics stall. This occurs around $\theta = 15.5$, and in Fig. 67 we see that both $A(\theta)$ and $f_1(\theta, \rho)$ are indeed close to 0. The reduced velocities of trajectories here highlight the importance of considering other phase space structures in forced systems, the details of which are missed in standard phase-only models. Forcing in the presence of such structures may give rise to complex and even chaotic behaviours, as we shall see in Sect. 5.3.

One of the advantages of using (θ, ρ) coordinates, is that we can represent other phase space structures. Because of this, it is easy to tell when trajectories

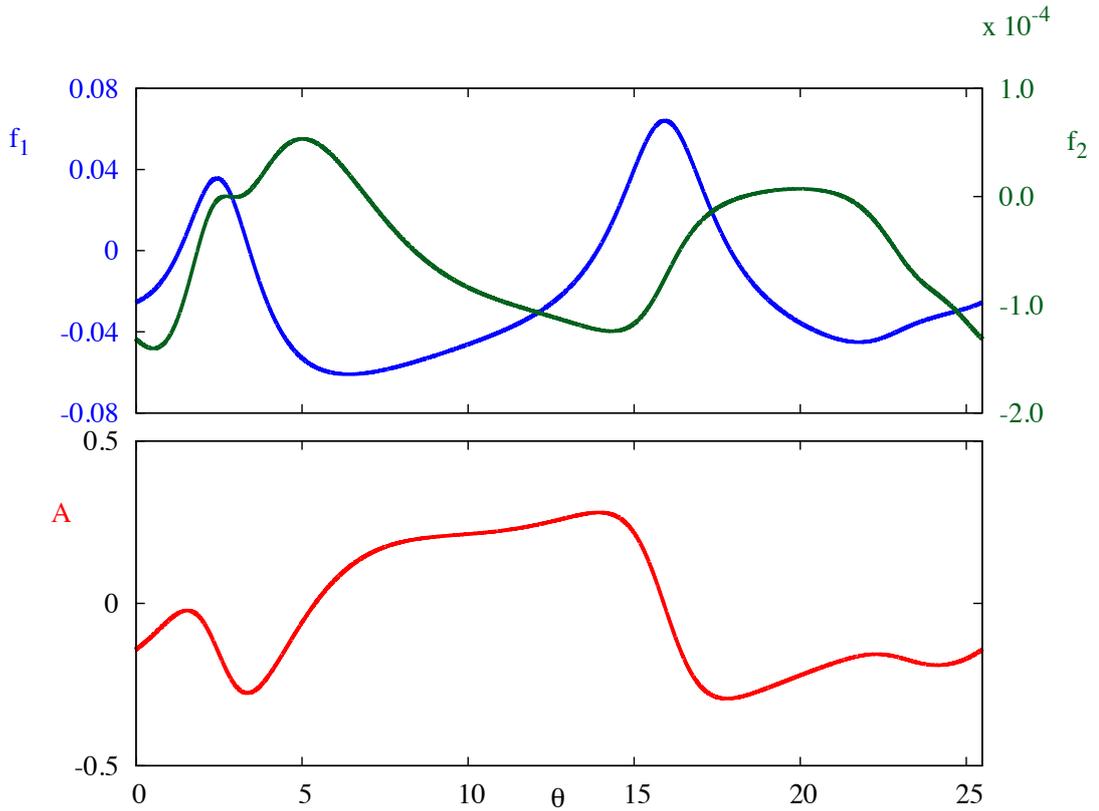


Figure 67: f_1 , f_2 and A for the ML model, evaluated at $\rho = -0.1$. We clearly see the difference in the order of magnitude between f_1 and f_2 for small ρ . Note that, although the average of A over one period is negative, it is positive for a nontrivial interval of θ . This corresponds to movement close to the stable manifold of the saddle.

have exited the basin of attraction of Γ as a result of some external forcing. As isochronal coordinates are only defined in \mathcal{B}_Γ , they are unable to detect these events. In Fig. 68, we plot the relevant phase plane structures for the ML model near a homoclinic bifurcation in (θ, ρ) coordinates, along with the transformed nullclines of the system.

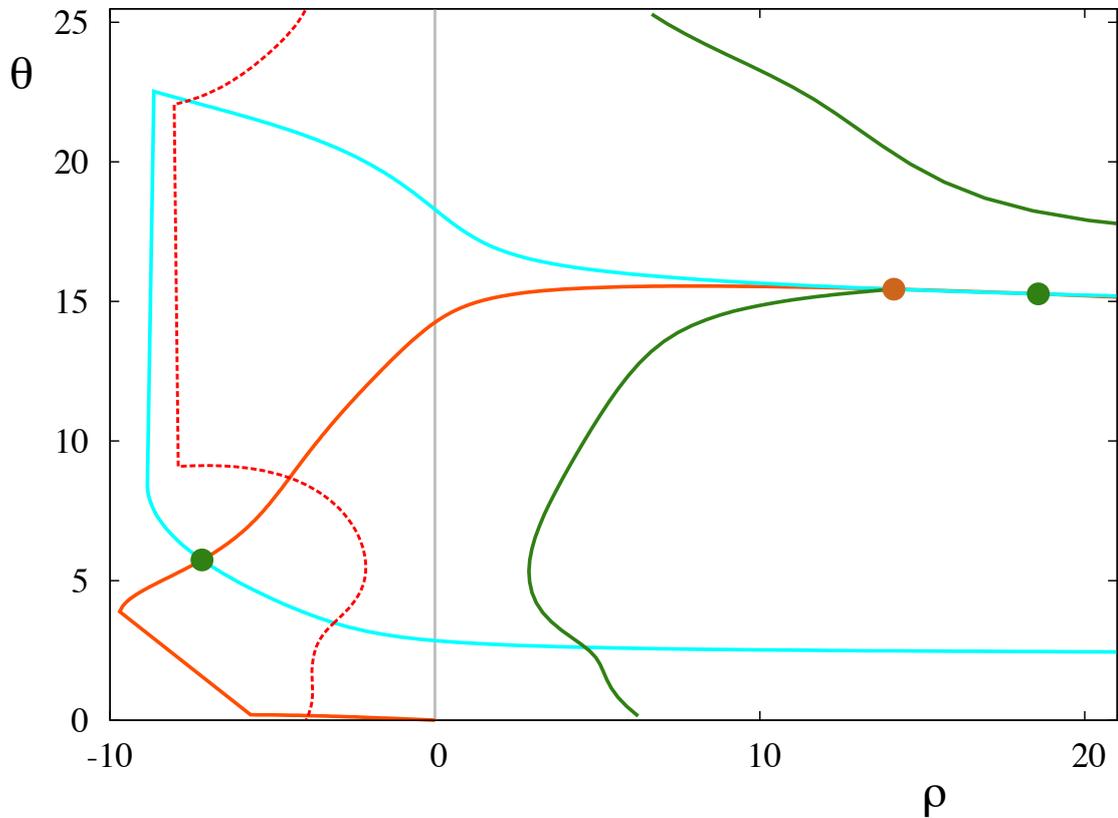


Figure 68: Relevant phase plane structures of the ML model near a homoclinic bifurcation in (θ, ρ) coordinates. The dashed black line represents Γ , the stable limit cycle. The green circles are stable fixed points, whilst the orange circle is the saddle. The stable manifold of the saddle, which acts as a separatrix between Γ and the stable fixed point to the right is shown as a dashed orange line. The unstable limit cycle, which serves as a separatrix between Γ and the fixed point to the left is depicted as a dashed red curve. We have also plotted the v -nullcline as a solid green curve, and the w -nullcline as a solid orange curve. With the exception of the nullclines, these phase plane structures can be seen, plotted in Cartesian coordinates, in the right panel of Fig. 46.

5.2.2 The FitzHugh–Nagumo Model

The phenomenological FitzHugh–Nagumo (FHN) model, [182, 24] is a modification of the van der Pol model [183], designed to capture the mathematical properties of the HH model. This planar model comprises a voltage-like variable v alongside a recovery variable w . The model can be used to describe both the oscillatory and excitable behaviour of nerves. The shape of the nullclines of the reduced HH model, as discussed in Sect. 2.2, are captured by the FHN model by simply making the dynamics for v cubic. The dynamics for w are linear. The model is written as

$$\mu\dot{v} = v(a - v)(v - 1) + I - w, \quad \dot{w} = v - bw. \quad (5.51)$$

Throughout this chapter we use the parameter values $\mu = 0.05$, $a = 0.9$, $I = 1.1$, and $b = 0.5$. An example of a typical trajectory for the FHN model may be seen in Fig. 69. We observe, as in Fig. 66, that there are portions along the trajectory during which the evolution of θ slows, during which ρ is increasing. However, we also note the rapid decay of ρ , so these effects are attenuated by the strong attraction to the limit cycle when μ is small.

Comparison of the ML and FHN model

In Fig. 70, we plot A for the FHN model (the corresponding graph for the ML model is found in the bottom panel of Fig. 67). We observe that, although Γ for the FHN is repelling for some interval of θ , the absolute value of A in these intervals is far smaller than that over intervals where Γ is attracting, so that nearby trajectories are rapidly to it. This is in contrast to the ML model, which has comparable absolute values for A over attracting and repelling intervals. We also note that the size of the intervals over which A is positive are much smaller for the FHN compared with the ML model. Fig. 71, shows a comparison of f_1, f_2 for both the ML and FHN models over a range of ρ . Examining f_2 , which we recall captures the nonlinear part of the attraction to Γ , we observe that it is always negative for the FHN model, whilst it takes both positive and negative values in the ML model. As such, f_2 will cause ρ to decay even faster for trajectories further away from Γ in

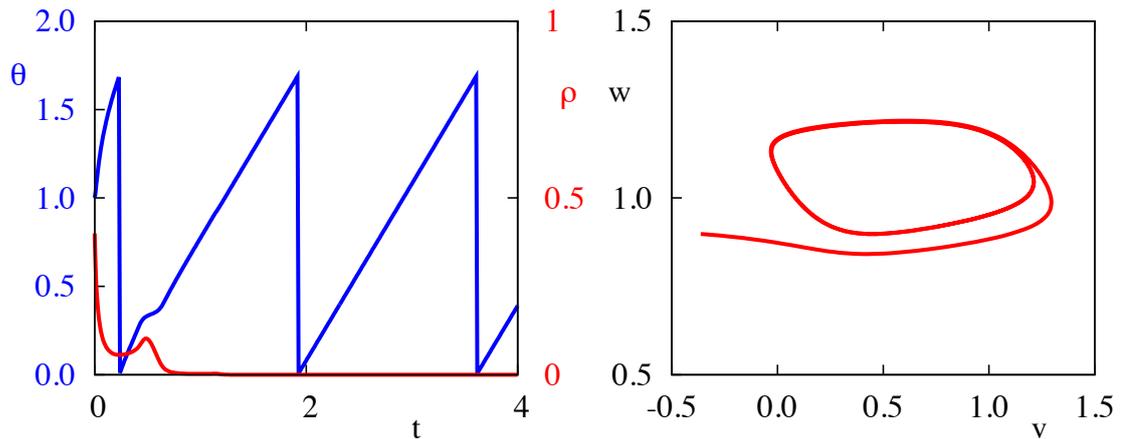


Figure 69: Typical trajectory of the FHN model. Left: Time evolution of θ and ρ . Right: Trajectory plotted in the (v, w) phase plane. We observe that, in contrast to the ML model in the homoclinic regime, the limit cycle is strongly attracting, and so ρ decays to zero very rapidly. Although ρ has a local maximum, at which the evolution θ also slows, as in Fig. 66, this effect is not as pronounced as for the ML case.

the FHN model, whilst it acts in opposition to A for the ML model. This means that $|\rho|$ has to increase for the nonlinear effects, captured by f_2 , to overcome the linear effects, captured by A , in the ML model. The interplay between these two functions reflects the presence of the saddle in the phase space, since trajectories near Γ will be taken away from Γ along the stable manifold of the saddle, before returning to Γ along the saddle's unstable manifold. The structure of f_1 is similar for both models, taking both positive and negative values, indicating that there are regions in the phase plane where the evolution of phase slows down and regions in which it speeds up.

5.2.3 A SNIC bifurcation

In addition to the bifurcation scenarios we considered in the previous chapter, namely the homoclinic and Hopf bifurcations, the ML model also supports a SNIC bifurcation. In this type of bifurcation, a periodic orbit is generated at the point

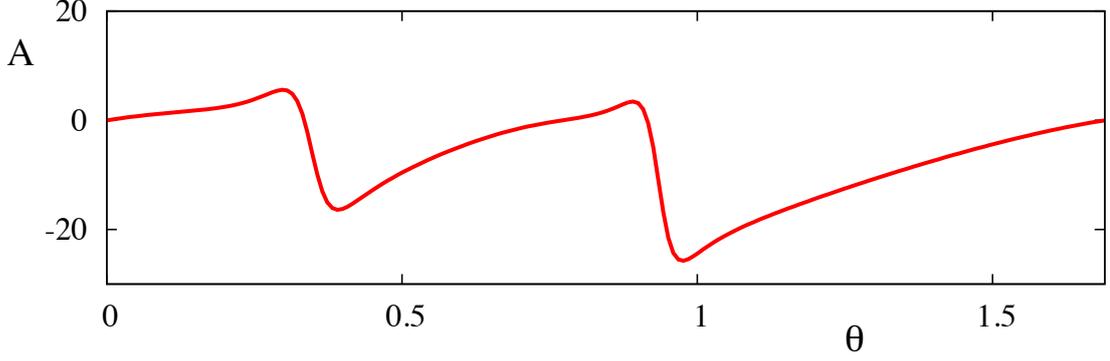


Figure 70: A evaluated for the FHN model, (5.51). The corresponding function for the ML model may be seen in Fig. 67.

where two fixed points collide and annihilate one another. The fixed points may be thought of as existing on such a limit cycle, so that trajectories that would otherwise follow the limit cycle are instead attracted towards the stable fixed point. For parameters values close to, but following the bifurcation, the ‘ghost of the saddle node’ still slows the trajectories in the vicinity where the saddle node bifurcation occurred. This leads to type I excitability, as mentioned in other parts of this thesis, since the frequency of the limit cycle vanishes at the bifurcation point. One of the simplest systems exhibiting a SNIC, [46], is (4.1) with

$$f = \begin{pmatrix} \beta x - my - x(x^2 + y^2) + \frac{y^2}{\sqrt{x^2 + y^2}} \\ mx + \beta y - y(x^2 + y^2) - \frac{xy}{\sqrt{x^2 + y^2}} \end{pmatrix}, \quad (5.52)$$

in which the bifurcation occurs at $m = 1$. For $m > 1$, there exists a limit cycle, Γ , which is available in closed form as

$$\Gamma(\theta) = \sqrt{\beta} (\cos \Omega(\theta), \sin \Omega(\theta)), \quad (5.53)$$

for a function Ω which maps the phase variable $\theta \in [0, \Delta) \mapsto [0, 2\pi)$, where $\Delta = 2\pi/\sqrt{m^2 - 1}$, given by

$$\Omega(\theta) = 2 \arctan \left(\frac{m \sin \pi\theta/\Delta}{\sqrt{m^2 - 1} \cos \pi\theta/\Delta + \sin \pi\theta/\Delta} \right). \quad (5.54)$$

For simplicity, we set $\beta = 1$ hereafter. Using

$$\frac{d\Gamma}{d\theta} = f(\Gamma(\theta)), \quad (5.55)$$

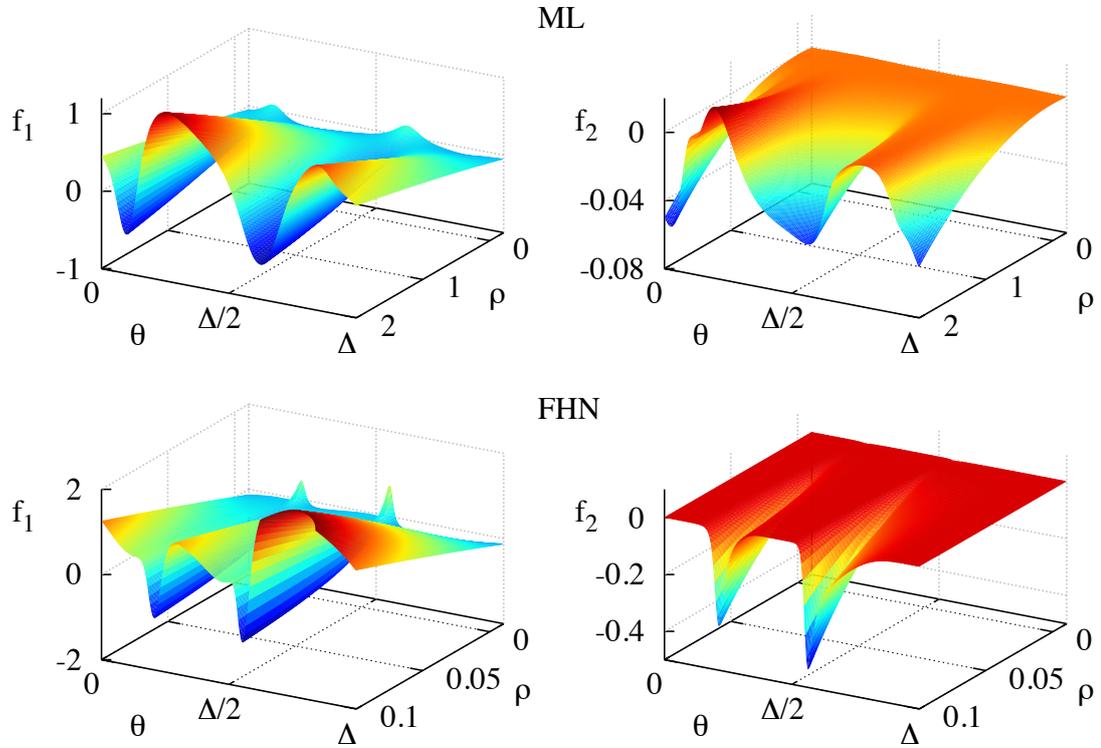


Figure 71: f_1 and f_2 evaluated for the ML model (top) and FHN model (bottom). The structure of f_1 is similar for both models, taking both positive and negative values, indicating that there are regions in the phase plane where the evolution of phase slows down and regions in which it speeds up. Conversely, we note that f_2 is never positive for the FHN model, so that the nonlinear effects captured by f_2 cause trajectories further away from Γ to be attracted faster than those close to it. For the ML model, f_2 acts in opposition to A , so that $|\rho|$ must increase before the nonlinear effects of f_2 overcome the linear ones of A . This interplay reflects the presence of the saddle in the phase plane of the ML model, and in particular describes the effects of the invariant manifolds of the saddle.

we may show that

$$\frac{d\Omega}{d\theta} = m - \sin \Omega(\theta). \quad (5.56)$$

The coordinates for this system are given by

$$\xi(\theta) = \begin{pmatrix} -\sin \Omega(\theta) \\ \cos \Omega(\theta) \end{pmatrix}, \quad \zeta(\theta) = \begin{pmatrix} \cos \Omega(\theta) \\ \sin \Omega(\theta) \end{pmatrix}, \quad \left| \frac{d\Gamma}{d\theta} \right| = \frac{d\Omega}{d\theta} \quad (5.57)$$

Differentiating with respect to θ , we have that

$$\frac{d\zeta}{d\theta} = \frac{d\Omega}{d\theta} \begin{pmatrix} -\sin \Omega(\theta) \\ \cos \Omega(\theta) \end{pmatrix} = \frac{d\Omega}{d\theta} \xi(\theta). \quad (5.58)$$

so that

$$\xi^T \frac{d\zeta}{d\theta} = \frac{d\Omega}{d\theta}. \quad (5.59)$$

This yields

$$h(\theta, \rho) = \left[\frac{d\Omega}{d\theta} (1 + \rho) \right]^{-1} \xi(\theta). \quad (5.60)$$

It can be shown that

$$f(\Gamma + \zeta\rho) - f(\Gamma) = -\rho(1 + \rho)(2 + \rho)\zeta + \rho(m - \sin \Omega(\theta))\xi, \quad (5.61)$$

so that

$$\begin{aligned} f_1(\theta, \rho) &= -h(\theta, \rho)^T \frac{d\zeta}{d\theta} \rho + h(\theta, \rho)^T [f(\Gamma + \zeta\rho) - f(\Gamma)] \\ &= \frac{-\rho(m - \sin \Omega(\theta)) - \rho(1 + \rho)(2 + \rho)\xi^T \zeta + \rho(m - \sin \Omega(\theta))}{(m - \sin \Omega(\theta))(1 + \rho)} = 0. \end{aligned} \quad (5.62)$$

We may also calculate

$$\zeta^T Df(\Gamma)\zeta = -2,$$

so that

$$A(\theta) = \zeta^T \frac{d\zeta}{d\theta} + \zeta^T Df(\Gamma)\zeta = -2, \quad (5.63)$$

since $\zeta^T d\zeta/d\theta = 0$. Finally, since $f_1 = 0$, we have

$$\begin{aligned} f_2(\theta, \rho) &= \zeta^T [f(\Gamma + \zeta\rho) - f(\Gamma)] - \zeta^T Df(\Gamma)\zeta\rho \\ &= -\rho(1 + \rho)(2 + \rho)\zeta^T \zeta + \rho(m - \sin \Omega(\theta))\zeta^T \xi + 2\rho \\ &= -\rho^2(3 + \rho). \end{aligned} \quad (5.64)$$

Substituting, (5.63)-(5.64) into (5.15), we may write the dynamics of the unforced SNIC oscillator as

$$\dot{\theta} = 1, \quad \dot{\rho} = -\rho(2 + 3\rho + \rho^2). \quad (5.65)$$

We note two important features of (5.65). Firstly, since $f_1=0$, there are no shear forces present in the system. Secondly, the limit cycle is attracting for all values of θ , and the strength of this attraction is independent of θ . This means that, in contrast to the case near to a homoclinic bifurcation, we would not expect shear induced chaos for a system near a SNIC, in spite of the similarities between the two bifurcations. Since the ‘sticky’ region of phase space is now along Γ , we do not get the same kind of folds as observed in Fig. 48. That we should not expect shear induced chaos for the SNIC oscillator is reinforced by an inspection of the isochrons of the system in Cartesian coordinates, which are radial lines, emanating from the fixed point at the origin, as shown in Fig. 72. To see why these radial isochrons do not predict shear induced chaos, recall Fig. 45, which shows the case for the linear shear model (4.8). In this figure, we showed the movement of a kick image down the isochrons which, for (4.8), are simply straight lines with negative slope. The shear, in this case, caused a fold to develop as the kick image relaxed back to Γ along these isochrons. Radial isochrons would be represented in such a graph by vertical lines, and so the relaxation of a kick image along these isochrons would not produce any folding. Since no folds are produced, we do not have any shear induced chaos.

In the next example, we show how phase-amplitude coordinates can be used for higher dimensional models. In spite of the increase in dimension, the technique of moving to phase-amplitude coordinates may still be readily applied.

5.2.4 A 4D conductance-based model

The Connor–Stevens (CS) model [184] is built upon the HH formalism and comprises a fast Na^+ current, a delayed K^+ current, a leak current and a transient K^+ current, termed the A-current. The full CS model consists of 6 equations: the membrane potential, the original HH gating variables and an activating and

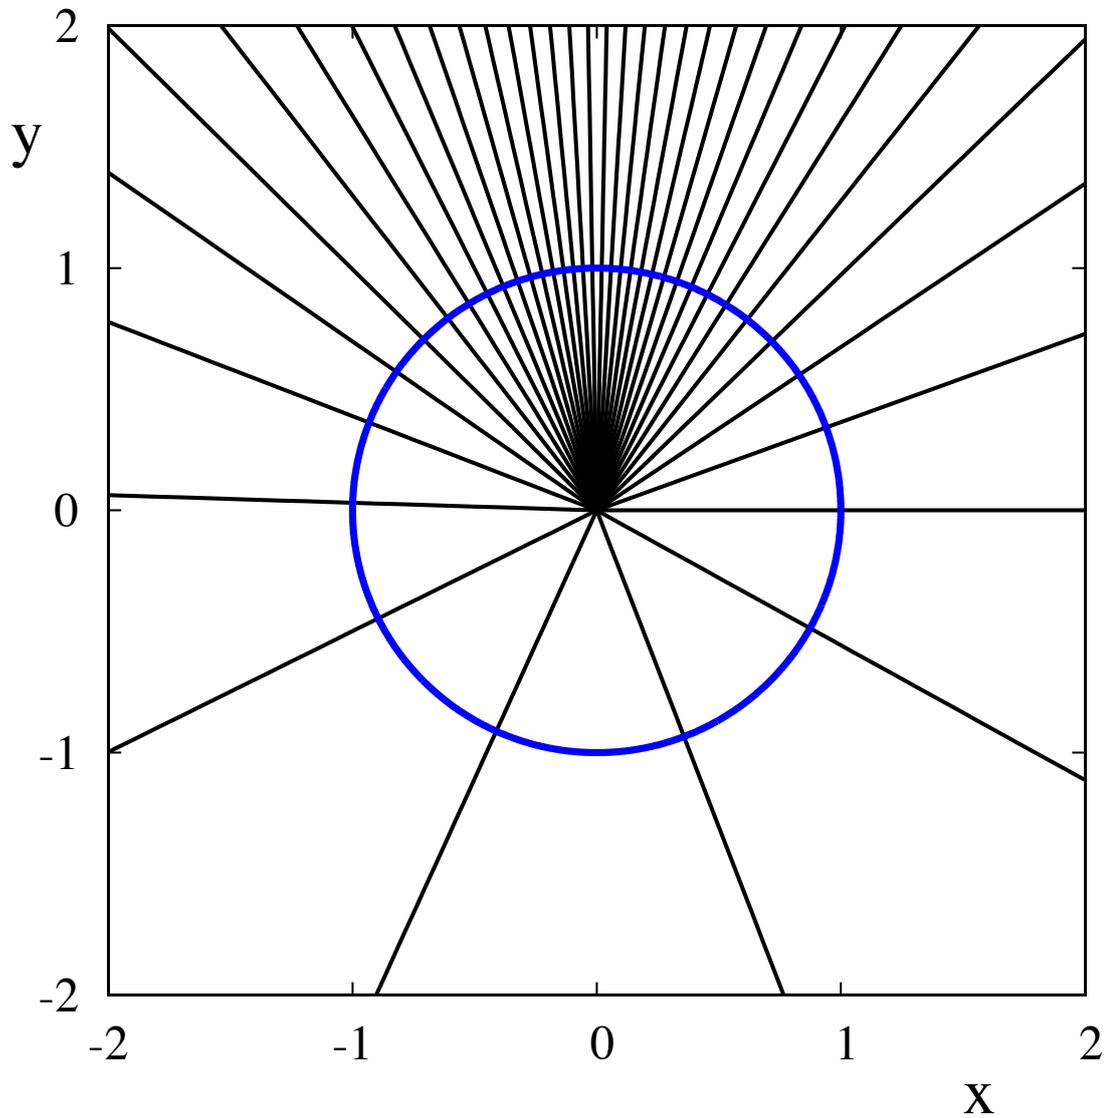


Figure 72: Isochrons of system (5.52). The isochrons are seen to be radial lines, emanating from the fixed point at the origin, and are thus orthogonal to Γ . For system (5.52) in phase-amplitude coordinates, $f_1 = 0$ everywhere and both A and f_2 are negative and independent of θ . This means that, in contrast to the case for a homoclinic bifurcation, we do not expect shear folding or shear induced chaos for a system near a SNIC bifurcation.

inactivating gating variable for the A-current. Using the method of equivalent potentials [21] as done in Sect. 2.2, we may reduce the dimensionality of the system, to include only four variables namely, the membrane potential, v , and equivalent potential u and two gating variables (a, b) . The reduced system is:

$$C\dot{v} = -F(v, u, a, b) + I, \quad \dot{u} = G(v, u), \quad \dot{X} = \frac{X_\infty(v) - X}{\tau_X(v)}, \quad X \in \{a, b\}, \quad (5.66)$$

where

$$F(v, u, a, b) = g_K n_\infty^4(u)(v - v_K) + g_{Na} h_\infty(u) m_\infty^3(v)(v - v_{Na}) + g_L(v - v_L) + g_a a^3 b(v - v_a). \quad (5.67)$$

and

$$G(v, u) = \frac{\left(\frac{\partial F}{\partial h} \left[\frac{h_\infty(v) - h_\infty(u)}{\tau_h(v)} \right] + \frac{\partial F}{\partial n} \left[\frac{n_\infty(v) - n_\infty(u)}{\tau_n(v)} \right] \right)}{\left(\frac{\partial f}{\partial h_\infty} \frac{dh_\infty(u)}{du} + \frac{\partial f}{\partial n_\infty} \frac{dn_\infty(u)}{du} \right)}, \quad (5.68)$$

where $\partial F/\partial h$ and $\partial F/\partial n$ are evaluated at $h = h_\infty(u)$ and $n = n_\infty(u)$. For the gating variables (a, b) , we have

$$a_\infty(v) = \left(\frac{0.0761e^{\frac{v+94.22}{31.84}}}{1 + e^{\frac{v+1.17}{28.93}}} \right)^{\frac{1}{3}}, \quad \tau_a(v) = 0.3632 + \frac{1.158}{1 + e^{\frac{v+55.96}{20.12}}}, \quad (5.69)$$

$$b_\infty(v) = \left(\frac{1}{1 + e^{\frac{v+53.3}{14.54}}} \right)^4, \quad \tau_b(v) = 1.24 + \frac{2.678}{1 + e^{\frac{v+50}{16.027}}}. \quad (5.70)$$

The parameter values we use are $C = 1 \mu\text{F}/\text{cm}^2$, $g_L = 0.3 \text{ mmho}/\text{cm}^2$, $g_K = 36.0 \text{ mmho}/\text{cm}^2$, $g_a = 47.7 \text{ mmho}/\text{cm}^2$, $I = 35.0 \mu\text{A}/\text{cm}^2$, $v_0 = 80.0 \text{ mV}$, $v_a = -75.0 \text{ mV}$, $v_K = -77.0 \text{ mV}$, $v_L = -54.4 \text{ mV}$, and $v_{Na} = 50.0 \text{ mV}$. The solutions to the reduced CS model under the coordinate transformation may be seen in Fig. 73, whilst in Fig. 74, we show how this solution appears in the original coordinates. As for the ML model, θ evolves approximately constantly throughout, though this evolution is sped up close to $\theta = \Delta$. The trajectories of the vector ρ are more complicated, but note that there is regularity in the pattern exhibited, and that this occurs with approximately the same period as the underlying limit cycle. The damping of the amplitude of oscillations in ρ over successive periods represents the overall attraction to the limit cycle, whilst the regular behaviour of ρ represents the specific relaxation to cycle, as shown in Fig. 74.

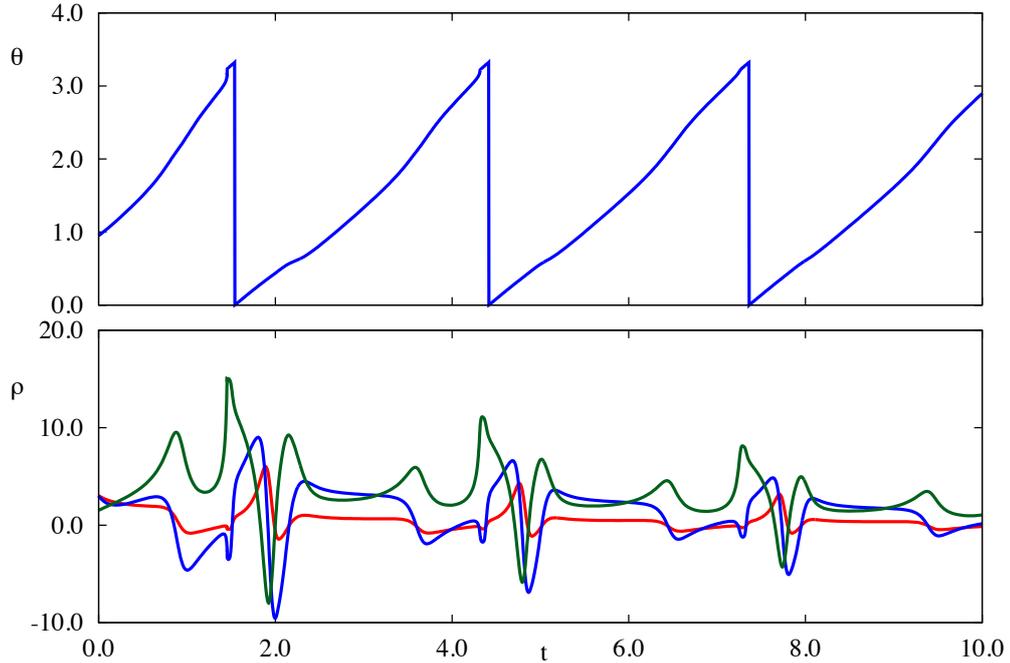


Figure 73: Solution of the transformed CS system. Top: Time evolution of θ . Bottom: Time evolution of ρ coordinates. Upon transforming these coordinates back to the original ones, we arrive at Fig. 74. In this parameter regime the model exhibits type I firing dynamics.

5.3 Pulsatile forcing of phase-amplitude oscillators

We now consider a system with time-dependent forcing, given by (5.38) with

$$g(x, t) = \sum_{n \in \mathbb{Z}} (\delta(t - nT), 0, \dots, 0)^T, \quad (5.71)$$

where δ is the Dirac δ -function. This describes T -periodic *kicks* to the voltage variable. We have already seen in Sect. 4.4.2, that even such a simple forcing paradigm can give rise to rich dynamics [157] in the ML model. The presence of chaos in the homoclinic regime for the ML model means that the response of the neuron is extremely sensitive to the initial phase when the kicks occur. We may use the phase-amplitude coordinates to see more clearly the shear effects that produce folding in the ML model. We now show the analogues of Fig. 48 and Fig. 49 in the

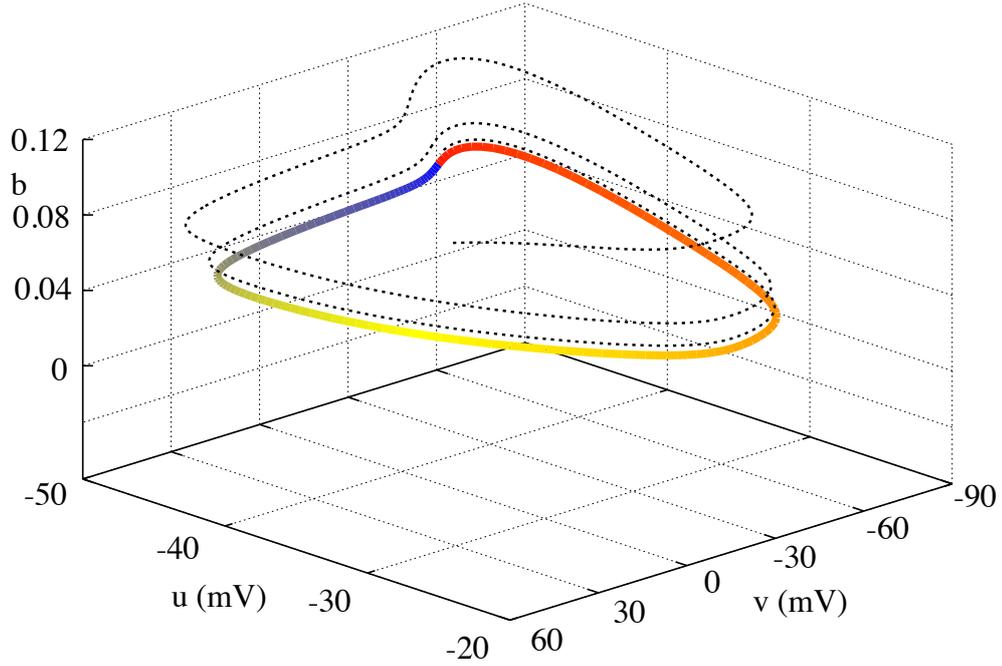


Figure 74: Transformed trajectory in (v, u, b) space of the phase-amplitude description of the reduced CS model. The dotted black line is the phase amplitude solution transformed in the original coordinates, whilst the coloured orbit is the underlying periodic orbit, where the colour corresponds to the phase along the orbit. The solution of the phase-amplitude description of the model in (θ, ρ) coordinates is shown in Fig. 73.

phase amplitude coordinates. As in Fig. 45, we represent Γ by the horizontal grey line, and give it, at time $t = 0$, a kick of the form (5.71) with strength $\varepsilon = -2$ taking it to the black curve in the top left panel. Following this, for Fig. 75, we evolve all of the points on the black curve under (5.15) for the ML model, for the time indicated on the panels. We notice the effect of the repelling parts of the orbit, as ρ increases for some parts of the black curve, before decreasing as all of the points are eventually attracted back to Γ . More importantly, we can see the effect of the ‘stickiness’ of the saddle, which has coordinates $(\theta, \rho) \approx (15.4376, 14.1251)$. Points along the black curve passing near to these points get ‘trapped’, whilst points away from it do not. We can see a fold in the black curve forming at around $t = 10$,

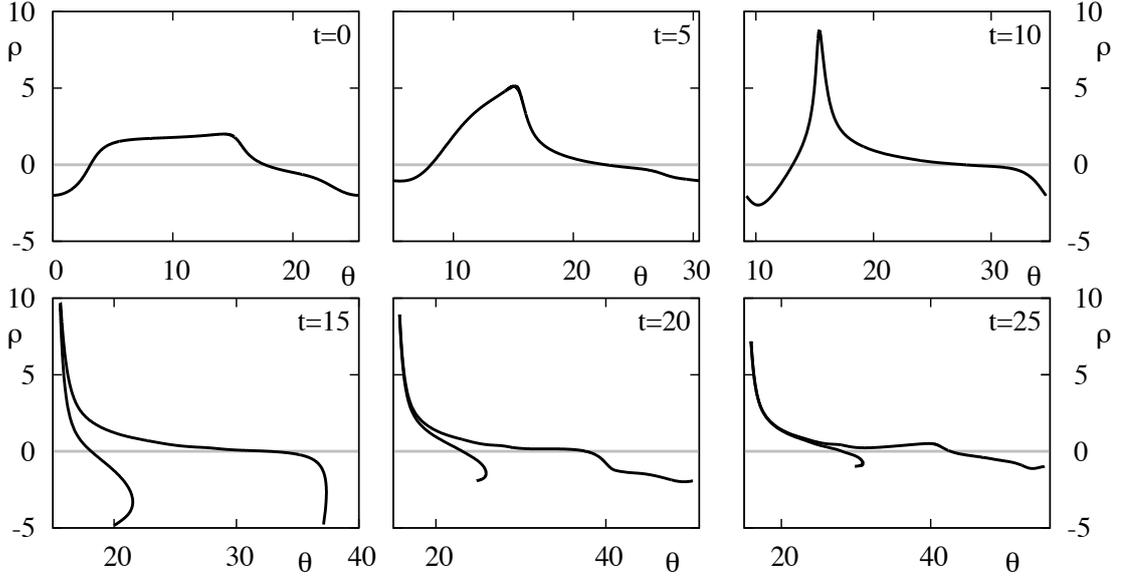


Figure 75: Analogy of Fig. 48 in phase-amplitude coordinates. At $t = 0$, the grey line, representing Γ is given a kick of the form (5.71) with strength $\varepsilon = -2$, taking it to the black curve in the top left panel. We then evolve points on this curve under the flow of the ML model, using (5.15), in (θ, ρ) coordinates for the time indicated in the panels. As in the case for the Cartesian coordinates, we can see a fold developing as time continues. We go on to show how these folds accumulate under repeated kicks in Fig. 76.

near $\theta = 15.5$, which sharpens as time increases. As time continues to increase, all points along the black curve will decay towards Γ and so the fold will not be as prominent as $t \rightarrow \infty$. However, we may repeat the kicks, evolving under (5.44), with $T = 25$ to see the effect of repeated kicks on the system, as shown in Fig. 76, where we plot points just after the kicks. We can clearly see the folds accumulate. As this process continues, it will give rise to the chaos observed in Sect. 4.4.2. We would now like to compare full planar neural models to the linear shear model model (4.8), which we recall here, for convenience

$$\dot{\theta} = 1 + \beta\rho, \quad \dot{\rho} = -\lambda\rho + \varepsilon P(\theta) \sum_{n \in \mathbb{Z}} \delta(t - nT). \quad (5.72)$$

By comparing with the phase-amplitude dynamics described by equations

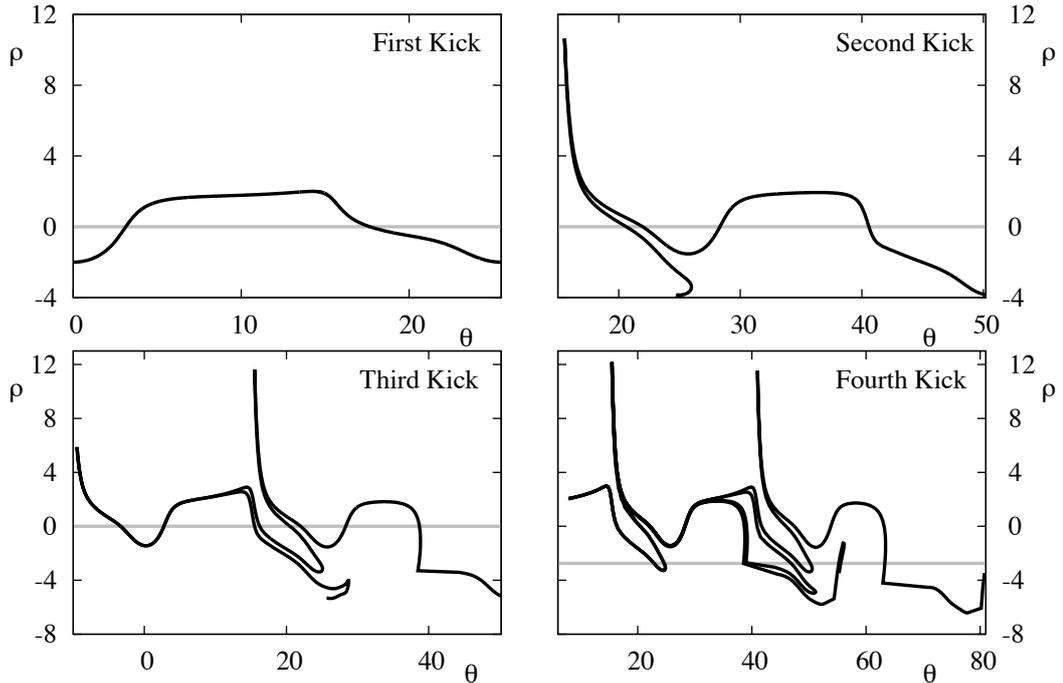


Figure 76: Image of Γ (black curve) under (5.44), with $T = 27$, where the grey line represents Γ . This is the analogy of Fig. 49, so that we are repeating iterations of the map which composes the action of the kick and the unforced flow. We can clearly see folds accumulating under the repeated action of the kicks. As these folds continue to accumulate, a strange attractor forms, leading to chaos, as shown in Sect. 4.4.2.

(5.44), we see that the model of shear considered in (5.72) is a proxy for a more general system, with $f_1(\theta, \rho) \rightarrow \sigma\rho$, $A(\theta) \rightarrow -\lambda$ and $h(\theta, \rho) \rightarrow 0$, and $\zeta_1(\theta) \rightarrow P(\theta)$.

Compared to the phenomenological system (5.72), models written in phase-amplitude coordinates as (5.44) have two main differences. The intrinsic dynamics (without kicks) are nonlinear, and the kick terms appear in both equations for $\dot{\theta}$ and $\dot{\rho}$ (not just $\dot{\rho}$). Simulations of (5.44) for both the FHN and ML models, with $\varepsilon = 0.1$, show that the replacement of $f_1(\theta, \rho)$ by $\sigma\rho$, dropping $f_2(\theta, \rho)$ (which is quadratic in ρ), and setting $A(\theta) = -\lambda$ does not lead to any significant qualitative change in behaviour (for a wide range of $\sigma, \lambda > 0$). We therefore conclude that, at least when the kick amplitude ε is not too large, it is more important to focus on

the form of the forcing in the phase-amplitude coordinates. In what follows, we are interested in discovering the effects of different functional forms of the forcing term multiplying the δ -function, keeping other factors fixed. As examples, we choose those forcing terms given by transforming the FHN and the ML models into phase-amplitude coordinates. To find these functions, we first find the attracting limit cycle solution in the ML model (2.28)-(2.29) and FHN model (5.51) using a periodic boundary value problem solver and set up the orthonormal coordinate system around this limit cycle. Once the coordinate system is established, we evaluate the functions $h(\theta, \rho)$ and $\zeta(\theta)$ (that appear in the system (5.44)). Using the forcing term (5.71), we are only considering perturbations to the voltage component of our system and thus only the first component of h and ζ will make a nontrivial contribution to the dynamics. We define P_1 as the first component of h and P_2 as the first component of ζ . We wish to force each system at the same ratio of the natural frequency of the underlying periodic orbit. To ease comparison between the system with the ML forcing terms and the FHN forcing terms, we rescale $\theta \mapsto \theta/\Delta$ so that $\theta \in [0, 1)$ in what follows. Implementing the above choices leads to

$$\begin{aligned}\dot{\theta} &= 1 + \beta\rho + \varepsilon P_1(\theta, \rho) \sum_{n \in \mathbb{Z}} \delta(t - nT), \\ \dot{\rho} &= -\lambda\rho + \varepsilon P_2(\theta) \sum_{n \in \mathbb{Z}} \delta(t - nT).\end{aligned}\tag{5.73}$$

It is important to emphasise that $P_{1,2}$ are determined by the underlying single neuron model (unlike in the toy model (5.72)). As emphasised in [185], one must take care in the treatment of the state dependent ‘jumps’ caused by the δ -functions in (5.73) to accommodate the discontinuous nature of θ and ρ at the time of the kick. To solve (5.73), we approximate $\delta(t)$ with a normalised square pulse $\delta_\tau(t)$ of the form

$$\delta_\tau(t) = \begin{cases} 0 & t \leq 0, \\ 1/\tau & 0 < t \leq \tau, \\ 0 & t > \tau, \end{cases}\tag{5.74}$$

where $\tau \ll 1$. This means that for $(n-1)T + \tau < t \leq nT$, $n \in \mathbb{Z}$, the dynamics

are governed by the linear system $(\dot{\theta}, \dot{\rho}) = (1 + \beta\rho, -\lambda\rho)$. This can be integrated to obtain the state of the system just before the arrival of the next kick, $(\theta_n, \rho_n) \equiv (\theta(nT), \rho(nT))$, in the form

$$\theta_n = \left[\theta(t) + nT - t + \frac{\beta}{\lambda} \rho(t) (1 - e^{-\lambda(nT-t)}) \right] \bmod 1, \quad (5.75)$$

$$\rho_n = \rho(t) e^{-\lambda(nT-t)}. \quad (5.76)$$

In the interval $nT < t < nT + \tau$ and using (5.74) we now need to solve the system of nonlinear ODEs

$$\dot{\theta} = 1 + \beta\rho + \frac{\varepsilon}{\tau} P_1(\theta, \rho), \quad \dot{\rho} = -\lambda\rho + \frac{\varepsilon}{\tau} P_2(\theta). \quad (5.77)$$

Rescaling time as $t = nT + \tau s$, with $0 \leq s \leq 1$, and writing the solution (θ, ρ) as a regular perturbation expansion in powers of τ as $(\theta(s), \rho(s)) = (\theta_0(s) + \tau\theta_1(s), \rho_0(s) + \tau\rho_1(s)) + \dots$, we find, after collecting terms of leading order in τ that the pair $(\theta_0(s), \rho_0(s))$ is governed by

$$\frac{d\theta_0}{ds} = \varepsilon P_1(\theta_0(s), \rho_0(s)), \quad \frac{d\rho_0}{ds} = \varepsilon P_2(\theta_0(s)), \quad 0 \leq s \leq 1, \quad (5.78)$$

with initial conditions $(\theta_0(0), \rho_0(0)) = (\theta_n, \rho_n)$. The solution $(\theta_0(1), \rho_0(1)) \equiv (\theta_n^+, \rho_n^+)$ (obtained numerically) can then be taken as initial data $(\theta(t), \rho(t)) = (\theta_n^+, \rho_n^+)$ in (5.75)-(5.76) to form the stroboscopic map

$$\theta_{n+1} = \left[\theta_n^+ + T - \frac{\beta}{\lambda} \rho_n^+ (1 - e^{-\lambda T}) \right] \bmod 1, \quad (5.79)$$

$$\rho_{n+1} = \rho_n^+ e^{-\lambda T}. \quad (5.80)$$

Note that this has been constructed using a matched asymptotic expansion, with (5.74), and is valid in the limit $\tau \rightarrow 0$. For weak forcing, where $\varepsilon \ll 1$, $P_{1,2}$ vary slowly through the kick and can be approximated by their values at (θ_n, ρ_n) so that to $O(\varepsilon^2)$

$$\theta_{n+1} = \left[\theta_n + T + \varepsilon P_1(\theta_n, \rho_n) + \frac{\beta}{\lambda} (\rho_n + \varepsilon P_2(\theta_n)) (1 - e^{-\lambda T}) \right] \bmod 1, \quad (5.81)$$

$$\rho_{n+1} = (\rho_n + \varepsilon P_2(\theta_n)) e^{-\lambda T}. \quad (5.82)$$

We choose to work with the full stroboscopic map (5.79)-(5.80), which is particularly useful for comparing and contrasting the behaviour of different planar single neuron models with arbitrary kick strength. As an indication of the presence of chaos in the dynamics resulting from this system, we again turn to LEs, evaluating the MLE of the map (5.79)-(5.80). The evolution of the separation vector, $\delta X = (\delta\theta, \delta\rho)$, between kicks, is given by the map

$$\delta\theta_{n+1} = \delta\theta_n^+ + \frac{\beta}{\lambda}\delta\rho_n^+ (1 - e^{-\lambda T}), \quad \delta\rho_{n+1} = \delta\rho_n^+ e^{-\lambda T}. \quad (5.83)$$

In order to find $(\delta\theta_n^+, \delta\rho_n^+)$, we solve the variational equation associated with (5.78), using $(\delta\theta_n, \delta\rho_n)$ as initial condition and use a matched asymptotic expansion, as before, to piece together solutions at $s = 0$ and $s = 1$.

For the phase-only model, as considered in Chapt. 4, the above becomes more straightforward. During the kicks, the evolution of θ for the model $\dot{\theta} = 1 + \varepsilon R(\theta) \sum_n \delta(t - nT)$ is given by

$$\frac{d\theta}{ds} = \tau + \varepsilon R(\theta), \quad 0 \leq s \leq 1. \quad (5.84)$$

Writing solutions as $\theta(s) = \theta_0(s) + \tau\theta_1(s) + \dots$, and working to first order in τ , we obtain

$$\frac{d\theta_0}{ds} = \varepsilon R(\theta_0). \quad (5.85)$$

Defining θ^- as the value of θ just before the kick, separating variables and integrating, matching solutions at $s = 0$ and $s = 1$, we obtain

$$\int_{\theta^-}^{\theta^+} \frac{d\theta}{R(\theta)} = \varepsilon, \quad (5.86)$$

which yields the implicit jump condition

$$\mathcal{F}(\theta^+) - \mathcal{F}(\theta^-) = \varepsilon, \quad \text{where} \quad \mathcal{F}(\theta) = \int \frac{d\theta}{R(\theta)}. \quad (5.87)$$

If \mathcal{F} is invertible, we may write θ^+ explicitly as

$$\theta^+ = \mathcal{F}^{-1}(\mathcal{F}(\theta^-) + \varepsilon). \quad (5.88)$$

If \mathcal{F} is not invertible, then we may solve (5.87) numerically, choosing an initial guess near $\theta^- + \varepsilon R(\theta^-)$. During a kick, we have, for $\delta\theta$

$$\frac{d\delta\theta}{ds} = \varepsilon R'(\theta)\delta\theta \quad (5.89)$$

Separating variables and matching solutions for this equation yields

$$\int_{\delta\theta^-}^{\delta\theta^+} \frac{d\delta\theta}{\delta\theta} = \varepsilon \int_0^1 R'(\theta(s)) ds = \int_{\theta^-}^{\theta^+} \frac{R'(\theta)}{R(\theta)} d\theta. \quad (5.90)$$

Integrating this equation gives

$$\ln\left(\frac{\delta\theta^+}{\delta\theta^-}\right) = \ln\left(\frac{R(\theta^+)}{R(\theta^-)}\right) \implies \delta\theta^+ = \frac{R(\theta^+)}{R(\theta^-)}\delta\theta^-. \quad (5.91)$$

and so we arrive at (4.16), as used in Sect. 4.4.2. We note that we could equivalently derive this expression using the technique in Sect. 3.13, taking (5.87) as the reset map g .

In Fig. 77, we compare the functions $P_{1,2}$ for both the FHN and the ML models, evaluating the functions at $\rho = 0$, where the response of the system to forcing is most prominent. We note that P_2 for the FHN model is near 0 for a large set of θ , whilst the same is true for P_1 for the ML model. This means that kicks in the FHN model will tend to primarily cause phase shifts, whilst the same kicks in the ML model will primarily cause shifts in amplitude.

We plot, in Fig. 78 A and B, the pair (θ_n, θ_{n+1}) , for (5.79)-(5.80) for the FHN and ML models. For the FHN model, we find that the system has a MLE of $-0.0515 < 0$. For the ML model the MLE is $0.6738 > 0$. This implies that differences in the functional forms of $P_{1,2}$ can help to explain the generation of chaos.

Now that we know the relative contribution of kicks in v to kicks in (θ, ρ) , it is also useful to know where kicks actually occur in terms of θ as this will determine the contribution of a train of kicks to the (θ, ρ) dynamics. In Figs. 78 C and D, we plot the distribution of kicks as a function of θ .

After transients, we observe a 1 : 1 phase-locked state for the FHN model. For a phase-locked state, small perturbations will ultimately decay as the perturbed

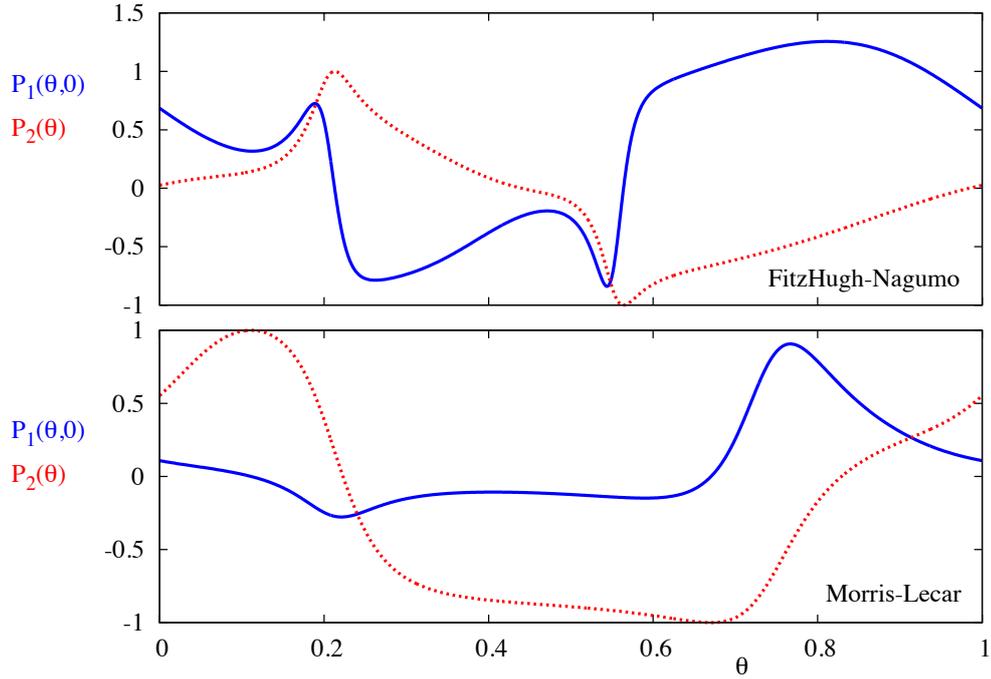


Figure 77: The blue curves show the change in θ under the action of a pulsatile kick in v , whilst the red dashed curves show the change in ρ under the same kick. The top plot is for the FHN model, whilst the bottom plot is for the ML model. We evaluate the effect of the kicks at $\rho_n = 0$, where we observe the largest changes in θ under the action of kicks.

trajectories also end up at the phase-locked state after some transient behaviour. Examining the histogram of the kick phases, we see a grouping of kicks around the region where P_2 is roughly zero. This sharply peaked distribution of kick phases is to be expected for discrete-time systems possessing a negative MLE, since such systems tend to have sinks. Since P_2 is small, kicks will not be felt as much in the ρ variable and so trajectories stay close to the limit cycle. As trajectories do not move far from the limit cycle, there is no possibility of folding, and hence no chaotic behaviour. This helps explain why it is more difficult to generate chaotic responses in the FHN model.

For the ML model, we observe chaotic dynamics around a strange attractor,

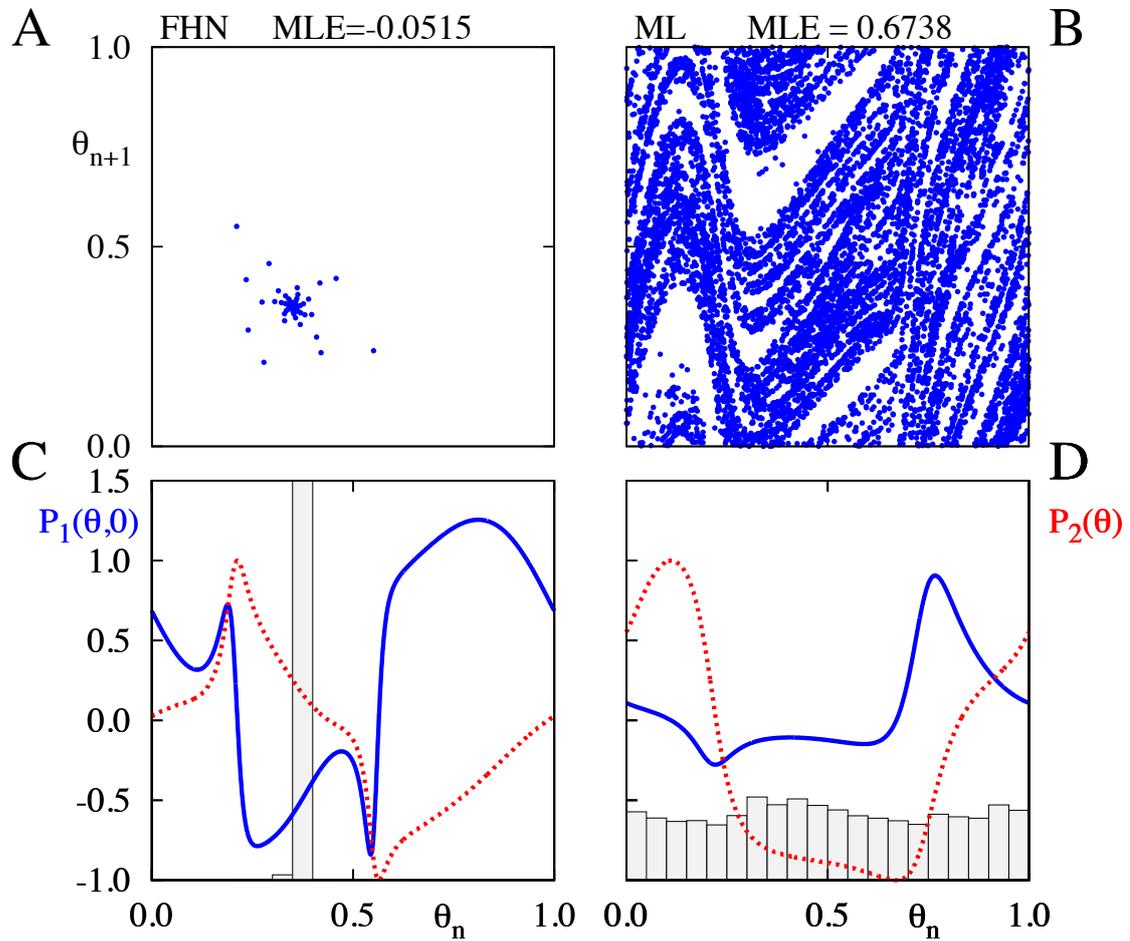


Figure 78: Panel A shows successive iterates of θ in system (5.79)-(5.80) with functions $P_{1,2}$ taken from the FHN model, whilst panel B presents the same iterates but for functions $P_{1,2}$ from the ML model. Panel C shows $P_{1,2}$ for the FHN model, where the bold blue line is P_1 and the red dashed line is P_2 . Superimposed on this panel is a histogram displaying how kicks are distributed in terms of θ alone. Panel D shows the same information, except this time for forcing functions from the ML model. Parameter values are $\beta = 3.0$, $\lambda = 0.1$, $\varepsilon = 0.1$, and $T = 2.0$.

where small perturbations can grow, leading to a positive MLE of 0.6738. This time, the kicks are distributed fairly uniformly across θ , and so, some kicks will take trajectories away from the limit cycle, and we have shear-induced folding and chaotic behaviour.

5.4 Revisiting the two cell ML network

As a final example, we return to the system of two gap junction coupled ML cells, as considered in Sect. 4.9. In the previous chapter, we showed that for weak coupling, with strength $\varepsilon = 0.001$, the phase-only model description of this network predicted a stable anti-phase rhythm between the two cells. The simulation of the system in phase-amplitude coordinates can be seen in Fig. 79, in which the top row shows the evolution of θ for the two cells, and the bottom row shows how this corresponds to the voltage variables in the original coordinates. Using these coordinates, we can shed some light onto the effect of the coupling between the cells. We observe, for both cells, an increase in the rate of evolution of θ (from the unperturbed case, in which $\dot{\theta} = 1$) in the interval $\theta \in (17, 20)$. Recalling that the saddle is at around $\theta \approx 15.5$, this region corresponds to the case where trajectories come closest to the saddle and hence evolve slowly, confirmed by noticing that these points correspond to troughs in the evolution of v , shown in the panels below. The effect of the coupling is to pull trajectories away from this region at a faster rate than the unperturbed case, similar to the release mechanism as studied in [186], though without any network inhibition. As a result of this, the period of the network oscillation is shorter than for an isolated neuron. This mechanism may help to promote the stability of the anti-phase rhythm, since we would not expect this effect for a synchronous network state.

5.5 Discussion

In this chapter, we have used the notion of a moving orthonormal coordinate system around a limit cycle to study dynamics in a neighbourhood around it. This

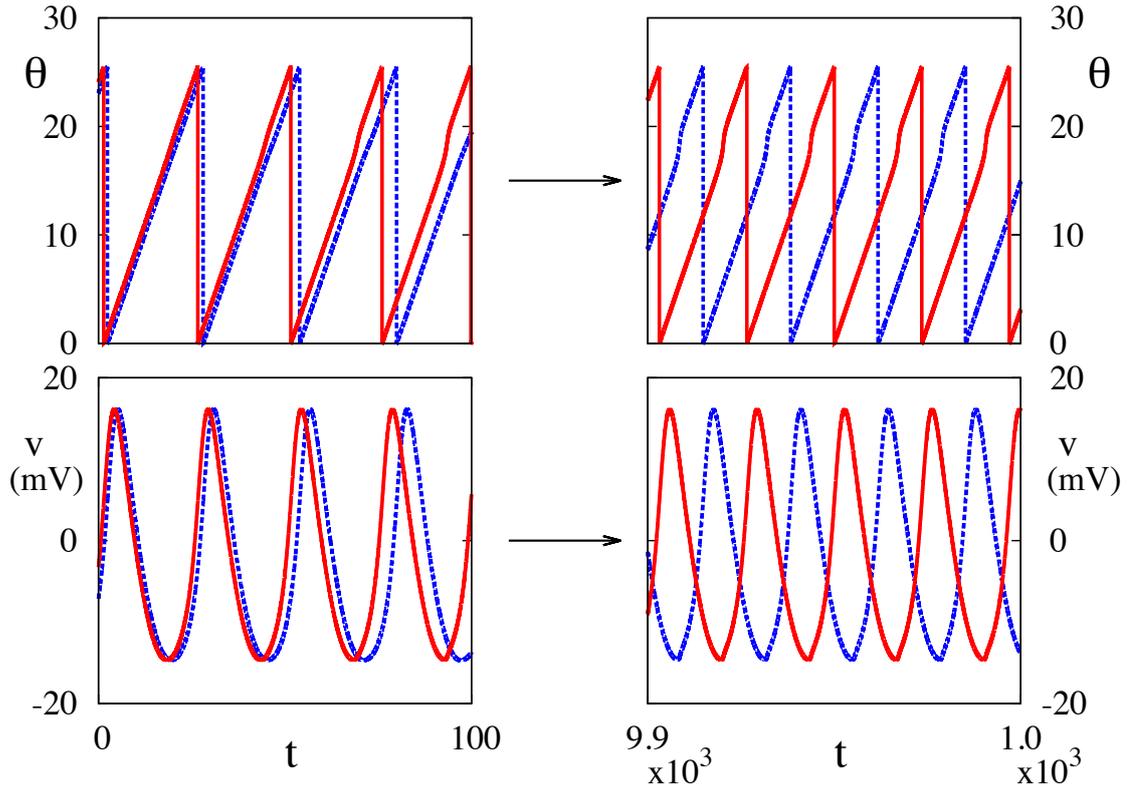


Figure 79: Simulation of a network of two identical ML neurons coupled through gap junctions, using phase-amplitude coordinates (5.44). Using the phase-amplitude coordinates, as shown in the top panels, we see an anti-phase network solution arise. The corresponding solution in v is shown in the panels below. We clearly see how the coupling increases the rate of evolution of θ near the saddle.

phase-amplitude coordinate system can be constructed for any given ODE system supporting a limit cycle. A clear advantage of the transformed description over the original one is that it allows us to gain insight into the effect of time dependent perturbations, using the notion of shear, as we have illustrated by performing case studies of popular neural models, in two and higher dimensions. Whilst this coordinate transformation does not result in any reduction in dimensionality in the system, as is the case with classical phase reduction techniques, it opens up avenues for moving away from the weak coupling limit, where $\varepsilon \rightarrow 0$. Importantly, it emphasises the role of the two functions $P_1(\theta, \rho)$ and $P_2(\theta)$ that provide more infor-

mation about inputs to the system than the iPRC alone. It has been demonstrated that moderately small perturbations can exert remarkable influence on dynamics in the presence of other invariant structures [157], which cannot be captured by a phase-only description. In addition, small perturbations can accumulate if the timescale of the perturbation is shorter than the timescale of attraction back to the limit cycle. This should be given particular consideration in the analysis of neural systems, where oscillators may be connected to thousands of other units, so that small inputs can quickly accumulate.

One natural extension of this work is to move beyond the theory of weakly coupled oscillators to develop a framework for describing neural systems as networks of phase-amplitude units. This has previously been considered for the case of weakly coupled weakly dissipative networks of nonlinear planar oscillators (modelled by small dissipative perturbations of a Hamiltonian oscillator) [187, 188, 189]. It would be interesting to develop these ideas and obtain network descriptions of the following type

$$\begin{aligned}\dot{\theta}_i &= 1 + f_1(\theta_i, \rho_i) + \sum_j w_{ij} H_1(\theta_i, \theta_j, \rho_i, \rho_j), \\ \dot{\rho}_i &= A(\theta_i) \rho_i + \sum_j w_{ij} H_2(\theta_i, \theta_j, \rho_i, \rho_j),\end{aligned}\tag{5.92}$$

with an appropriate identification of the interaction functions $H_{1,2}$ in terms of the biological interaction between neurons and the single neuron functions $P_{1,2}$. Such phase-amplitude network models are ideally suited to describing the behaviour of the mean-field signal in networks of strongly gap junction coupled ML neurons [190, 137], which is known to vary because individual neurons make transitions between cycles of different amplitudes. In addition, such coordinate systems may be useful for studying networks of bursting neurons, which are capable of supporting interesting synchronisation phenomena. In a diffusively coupled network of two square wave bursters, for example, the bursting envelope of the two cells may be synchronised, whilst the spikes within the burst are in anti-phase [191]. Increasing the coupling strength between these two cells, can sometimes promote a switch to a fully synchronous state, in which both the bursting envelope and

spikes within the burst are synchronised. We have shown, in Chapt. 3 and 4 how to use averaging theory to predict the stability of network states for a two cell network. Averaging theory shows what phase shift between the two cells is stable, but it will not capture the subtle effect of anti-phase spikes in a synchronised burst. Furthermore, it can be shown that, for elliptic bursters, with gap junction and synapses, that the degree of synchronisation can change during the active phase of a burst [192]. Whilst phase-only models cannot capture these effects, using phase-amplitude variables will enable us to describe these phenomena in familiar framework of phases. Finally, phase-amplitude equations may enable us to make progress in the analysis of rhythms in heterogeneous networks, in which different units in the network have different parameter values. This is clearly more representative of real neural tissue, which is comprised of many different types of neurons, each with their own morphology.

In the next chapter, we focus on the analysis of stochastically forced oscillators in the phase-amplitude framework.

6 Phase-amplitude descriptions of stochastically forced neural oscillators

Building on work from the previous chapter, the present chapter focuses on the phase-amplitude descriptions of stochastically forced oscillators. Stochastically forced neural oscillators have been receiving a significant amount of interest in recent years, such as in [193, 194, 195], as the community tries to better capture the dynamics of inherently stochastic biological processes. The addition of stochastic forcing terms can have a pronounced effect on oscillatory dynamics, even in cases where the noise strength is weak compared to the deterministic part of the dynamics. In terms of dynamical structures, noise can serve to shift invariant structures around, relative to deterministic ones, and may cause a shift in the critical values of bifurcations [11]. At the level of behaviour, this can cause changes in the frequency of firing of an oscillator, which may have consequences on the synchronisation properties of the network [196]. In extreme cases, it can even quench oscillatory behaviour, or make the system move into a completely unpredictable regime, though this is not typically a problem for weak noise sources. Stochastic forcing can also generate oscillatory activity where the deterministic system is quiescent via stochastic resonance [78], or may enhance entrainment to a deterministic signal through coherence resonance [77], which may help neurons to respond to weak signals [76, 65]. The need to understand the role that noise plays on neural coding means that we need to have a better knowledge of the way that noise sources impact dynamics [197, 198]. For weak noise sources, we can resort to similar techniques as for deterministic oscillators, through the use of isochrons of the asymptotically stable limit cycle [199, 8]. This allows us to form a perturbative expansion of the dynamics, using the noise strength as a parameter, resulting in an equation in phase alone. However, as in the case of Chapt. 4, we can also easily find situations in which these approximations do not always capture all of the necessary dynamical behaviour, even at the single neuron level.

From a modelling perspective, there are a plethora of considerations for stochas-

tic signals: which interpretation of noise is used; whether the noise acts additively or multiplicatively, and if the noise acts multiplicatively, what functional form does the signal take; and what is the structure of the temporal correlations in the noise source? On the simulation side, there are the same considerations as for deterministic systems, namely the trade off between accuracy and speed. This problem is exacerbated in the stochastically forced system, since inclusion of variability in the noise sources results in lower order accuracy of numerical schemes, forcing us to use smaller time steps to achieve the same level of accuracy as for the deterministic system [200]. Additionally, in order to properly simulate sample paths of a noisy system, it is necessary to generate pseudorandom numbers at each time step, further reducing the computational speed of numerical algorithms. We can make use of the Fokker–Planck equation to avoid these issues in simulation [93], and find equations for many quantities of interest associated with the SDE in the form of a PDE, however, in some cases, the complexity in solving the resultant PDE may mean that simulating sample paths is less onerous. Although we only consider the white-noise case, that is where the noise has temporal correlations only at zero lag, we note that there have been recent results in this area for the weak noise case, both by a redefinition of the isochrons to account for the temporal correlations in the noise [201], and through taking ensemble averages to include the correlation function of the noise directly in the governing equation for the phase dynamics [9].

Away from the weak noise limit, another area of interest for stochastically forced oscillators is that of bistability. In this setting, we are interested in cases in which a stable limit cycle coexists with a stable fixed point. If we consider such a system in terms of polar coordinates, this is akin to the classic double well potential, in which we are interested in transitions of a particle with some injected energy between two potential wells [80]. Relevant properties of these transitions, such as exit time distributions, can be computed through the use of the Fokker–Planck equation [93]. In a neural context, the transitions between quiescence and oscillations can be related to up and down states in EEG recordings from the cortex [202].

As in the previous chapter, we assume that the unforced system supports periodic oscillations, and that the forcing may be considered as a perturbation to the unforced system. In Sect. 5.1, we considered a coordinate system to describe the dynamics of a given model around a limit cycle in \mathbb{R}^n . This system had the drawback that it ceased to be invertible at some finite distance away from the limit cycle, and thus broke down. For planar systems, we may overcome this problem by setting up a different coordinate system around the limit cycle. We begin this chapter by reviewing some of the results for weakly forced oscillators, before later relaxing the weak assumption and considering switching between different basins of attraction under forcing.

6.1 Stochastically forced oscillators

Throughout this chapter, we consider systems of the form

$$dx = F(x)dt + \sigma G(x)dW_t, \quad x \in \mathbb{R}^2, \sigma \in \mathbb{R}. \quad (6.1)$$

We assume that the unforced system, with noise strength $\sigma = 0$, has a hyperbolic, attracting limit cycle, Γ , with period Δ , which may be parametrised by $\vartheta \in [0, \Delta)$, with $\dot{\vartheta} = 1$ on Γ . We note that there is only one noise source present in (6.1), which has a state dependence given by G , and that we consider white noise only, that is, where the noise is modelled as increments of a Wiener process.

6.2 Weakly forced systems in isochronal coordinates

We start from (6.1), interpreting the stochastic integral in the Stratonovich sense, so that the ordinary rules of calculus hold. Under this interpretation, we may rewrite this equation in isochronal coordinates, with an asymptotic phase $\vartheta \in [0, \Delta)$ and an amplitude variable, r , which captures the distance along an isochron, so that $r = 0$ corresponds to Γ . If we perform the transformation $(x_1, x_2) \mapsto (\vartheta, r)$, we obtain the system

$$d\vartheta = dt + \sigma h(\vartheta, r) \circ dW_t, \quad dr = f(\vartheta, r)dt + \sigma g(\vartheta, r) \circ dW_t, \quad (6.2)$$

where

$$h(\vartheta, r) = \nabla_x \vartheta \cdot G(x(\vartheta, r)), \quad f(\vartheta, r) = \nabla_x r \cdot F(x(\vartheta, r)), \quad g(\vartheta, r) = \nabla_x r \cdot G(x(\vartheta, r)). \quad (6.3)$$

Following [199], this Stratonovich SDE may be converted to an Itô SDE, yielding

$$\begin{aligned} d\vartheta &= dt + \frac{\sigma^2}{2} \left[\frac{\partial h(\vartheta, r)}{\partial \vartheta} h(\vartheta, r) + \frac{\partial h(\vartheta, r)}{\partial r} g(\vartheta, r) \right] dt + \sigma h(\vartheta, r) dW_t, \\ dr &= f(\vartheta, r) dt + \frac{\sigma^2}{2} \left[\frac{\partial g(\vartheta, r)}{\partial \vartheta} h(\vartheta, r) + \frac{\partial g(\vartheta, r)}{\partial r} g(\vartheta, r) \right] dt + \sigma g(\vartheta, r) dW_t, \end{aligned} \quad (6.4)$$

We will now show, using the Fokker–Planck equation, that the phase-only model approximating (6.4) may be obtained by setting $r = 0$ in the dynamics of ϑ , and setting $dr = 0$. The (forward) Fokker–Planck equation for the evolution of the probability, P , of system (6.4) is

$$\begin{aligned} \frac{\partial P}{\partial t} &= -\frac{\partial}{\partial \vartheta} [\{1 + D(h_\vartheta h + h_r g)\}P] + D \frac{\partial^2 [h^2 P]}{\partial \vartheta^2} \\ &\quad - \frac{\partial}{\partial r} [\{f + D(g_r h + g_r g)\}P] + 2D \frac{\partial^2 [hgP]}{\partial \vartheta \partial r} + D \frac{\partial^2 [g^2 P]}{\partial r^2}, \end{aligned} \quad (6.5)$$

where the subscripts ϑ and r denote partial differentiation and $D = \sigma^2/2$. We assume a periodic boundary condition for ϑ , so that $P(t, 0, r) = P(t, \Delta, r)$. When $D = 0$, the steady state distribution is given by $P_0 = \Delta^{-1} \delta(r)$. For small $D > 0$, the distribution still localises near $r = 0$ and rapidly decreases with increasing $|r|$ because of asymptotic stability of the limit cycle. It may be shown (see [199]) that indeed, P_0 localises near $r = 0$ to a Gaussian distribution for small D .

We now consider the domain $\{(\vartheta, r) \in \mathbb{R}^2; -\varepsilon \leq r \leq \varepsilon\}$ for a small constant ε and the large time behaviour of the system. We are going to approximate the Fokker–Planck equation up to $O(D)$, so that the steady state distribution is no longer given by $P_0 = \Delta^{-1} \delta(r)$. In this case, we may assume that $P_0 = 0$ and $\partial P_0 / \partial r = 0$ at $r = \pm \varepsilon$. Since P converges to P_0 , the localisation property holds for $P(t, \vartheta, r)$ and so $P = 0$ and $\partial P / \partial r = 0$ hold approximately at $r = \pm \varepsilon$.

We are interested in the marginal distribution $Q(t, \vartheta) = \int_{-\varepsilon}^{\varepsilon} P(t, \vartheta, r) dr$. We integrate (6.5) with respect to r , over the interval $[-\varepsilon, \varepsilon]$, upon noting that the terms including the derivative $\frac{\partial}{\partial r}$ disappear on integrating, to obtain

$$\frac{\partial Q}{\partial t} = -\frac{\partial}{\partial \vartheta} \int_{-\varepsilon}^{\varepsilon} (1 + DK_1) Q dr + D \frac{\partial^2}{\partial \vartheta^2} \int_{-\varepsilon}^{\varepsilon} K_2 Q dr, \quad (6.6)$$

where $K_1 = h_\vartheta h + h_r g$ and $K_2 = h^2$. These functions may be expanded in the forms $K_1 = h_\vartheta(\vartheta, 0)h(\vartheta, 0) + h_r(\vartheta, 0)g(\vartheta, 0) + rR_1(\vartheta, r)$ and $K_2 = h(\vartheta, 0)^2 + rR_2(\vartheta, r)$, where R_1 and R_2 are functions of $O(r^0)$ or higher, with respect to r . Defining

$$Z(\vartheta) = h(\vartheta, 0), \quad \text{and} \quad Y(\vartheta) = h_r(\vartheta, 0)g(\vartheta, 0), \quad (6.7)$$

K_1 and K_2 may be further rewritten as $K_1 = Z(\vartheta)Z'(\vartheta) + Y(\vartheta) + rR_1(\vartheta, r)$ and $K_2 = Z(\vartheta)^2 + rR_2(\vartheta, r)$. Consider the integrals $\int_{-\varepsilon}^{\varepsilon} rR_i P dr$, for $i = 1, 2$. Since $P(t, \vartheta, r)$ is approximately $P_0(\vartheta, r)$ and $\lim_{D \rightarrow 0} P_0(\vartheta, r) = \Delta^{-1}\delta(r)$, the profile of $P(t, \vartheta, r)$ may be approximated by $\delta(r)$ in the limit $D \rightarrow 0$. Using this approximation, and noting that rR_i is of order $O(r)$ or higher, we have that $\lim_{D \rightarrow 0} \int_{-\varepsilon}^{\varepsilon} rR_i P dr = 0$, implying that

$$\lim_{D \rightarrow 0} \frac{\partial}{\partial \vartheta} \int_{-\varepsilon}^{\varepsilon} rR_1 P dr = \frac{\partial}{\partial \vartheta} \lim_{D \rightarrow 0} \int_{-\varepsilon}^{\varepsilon} rR_1 P dr = 0, \quad (6.8)$$

and

$$\lim_{D \rightarrow 0} \frac{\partial^2}{\partial \vartheta^2} \int_{-\varepsilon}^{\varepsilon} rR_2 P dr = \frac{\partial^2}{\partial \vartheta^2} \lim_{D \rightarrow 0} \int_{-\varepsilon}^{\varepsilon} rR_2 P dr = 0. \quad (6.9)$$

Thus, we have $D \frac{\partial}{\partial \vartheta} \int_{-\varepsilon}^{\varepsilon} rR_1 P dr = o(D)$ and $D \frac{\partial^2}{\partial \vartheta^2} \int_{-\varepsilon}^{\varepsilon} rR_2 P dr = o(D)$. Substituting the expansions of K_1 and K_2 into (6.6) and using these facts, we obtain the approximate Fokker–Planck equation for Q , up to $O(D)$ as

$$\frac{\partial Q}{\partial t} = -\frac{\partial}{\partial \vartheta} [\{1 + D(ZZ' + Y)\}Q] + D \frac{\partial^2 [Z^2 Q]}{\partial \vartheta^2}. \quad (6.10)$$

The SDE corresponding to (6.10) (in the Stratonovich sense) is

$$d\vartheta = dt + DY(\vartheta)dt + \sigma Z(\vartheta) \circ dW_t, \quad (6.11)$$

which is the stochastic analogue of the phase reduction for the deterministic case. Comparing this with the phase equation in 6.2, we see that there is an extra term DY which appears on making the phase reduction in this stochastic setting.

Let us consider (6.10) with the periodic boundary condition $Q(t, 0) = Q(t, \Delta)$ and consider the steady state solution $Q_0(\vartheta)$, obtained by setting $\partial Q/\partial t = 0$. Upon constructing an asymptotic solution for Q_0 in terms of D up to $O(D)$, we have

$$Q_0(\vartheta) = \frac{1}{\Delta} + \frac{D}{\Delta} [Z(\vartheta)Z'(\vartheta) - Y(\vartheta) + \bar{Y}] + O(D^2). \quad (6.12)$$

where \bar{Y} is the average defined by

$$\bar{Y} = \frac{1}{\Delta} \int_0^\Delta Y(\vartheta) d\vartheta. \quad (6.13)$$

The mean frequency, Ω , of the oscillator is defined by $\Omega = \lim_{T \rightarrow \infty} T^{-1} \int_0^T \dot{\vartheta}(t) dt$. This may be calculated by replacing the time average by an ensemble average: $\Omega = \langle \dot{\vartheta} \rangle$. As there is no correlation between ϑ and W in the Itô interpretation of the noise, we may use the fact that $\langle Z(\vartheta) W_t \rangle = \langle Z(\vartheta) \rangle \langle W_t \rangle = 0$, and so we have that $\Omega = 1 + D \langle Z(\vartheta) Z'(\vartheta) + Y(\vartheta) \rangle$, which, using our expression Q_0 may be evaluated up to $O(D)$ as

$$\Omega = 1 + D\bar{Y} + O(D^2). \quad (6.14)$$

This equation shows the phenomenon of *noise induced frequency shift*. Although Gaussian noise has no intrinsic frequency, it still induces a change in frequency of the oscillator. We thus note the importance of the term DY in (6.11), as without this term, we would observe no noise induced frequency shift. In addition, (6.12) demonstrates the influence of the extra term in terms of the steady state probability distribution around ϑ . We note that we may equivalently derive (6.14) directly from (6.11). Though this technique for performing the phase reduction for the stochastically forced system considered in this section has been demonstrated here in \mathbb{R}^2 , the ideas carry over easily into \mathbb{R}^n .

6.2.1 Differences between the Itô and Stratonovich interpretations of the phase reduction

The original system (6.1) has been considered only in the Stratonovich interpretation thus far. The Stratonovich interpretation implies that the function G is not nonanticipating, that is, it is not statistically independent of the Wiener increment $W(s) - W(t)$ for $t < s$. This is because the stochastic integral requires an evaluation at the mid-point of the time step, rather than the left point, as used in the Itô case. Clearly, if the functions are related to some biological process, they cannot be acausal, that is, they cannot depend on future values, implying that a Stratonovich interpretation is not the best one to use. However, a Stratonovich

interpretation may be used if the noisy process is assumed to be a signal from some other process, not described by the model equations [84]. Even in this case, there may still be an argument for using an Itô interpretation. For example, we may consider the noise process for a neural model to be electrical noise involving mini-postsynaptic potentials. These processes are typically modulated by exocytosis, from a presynaptic cell, of a single vesicle containing neurotransmitters, which evokes a quantal postsynaptic conductance change [203]. The continual barrages of such events gives rise to a noisy signal that may be thought of as a continual approximation of discrete events, in which case an Itô interpretation may be more appropriate than a Stratonovich one [204]. We will now consider the case in which we use the Itô interpretation of the stochastic integral in (6.1). We note that we could simply write the equivalent Itô SDE of (6.11) as

$$d\vartheta = dt + D[Z(\vartheta)Z'(\vartheta) + Y(\vartheta)]dt + \sigma Z(\vartheta)dW_t. \quad (6.15)$$

However, this does not truly capture the differences between the two interpretations in the original variables, since the rules regarding the transformation of variables are different under Itô interpretation. In fact, we must resort to Itô's formula, as discussed in Sect. 2.17, when performing the change of variables.

Starting from (6.1) as before, now using an Itô interpretation of the noise term, the transformation to (ϑ, r) leads to the system

$$d\vartheta = dt + \frac{\sigma^2}{2} \left[h_\vartheta h + h_r g - \sum_i G_i \sum_j \frac{\partial G_j}{\partial x_i} \frac{\partial \vartheta}{\partial x_j} \right] dt + \sigma h(\vartheta, r)dW_t, \quad (6.16)$$

$$dr = f(\vartheta, r)dt + \frac{\sigma^2}{2} \left[g_\vartheta h + g_r g - \sum_i G_i \sum_j \frac{\partial G_j}{\partial x_i} \frac{\partial r}{\partial x_j} \right] dt + \sigma g(\vartheta, r)dW_t, \quad (6.17)$$

where we now see additional terms appearing in (6.16) compared with (6.2), which is derived by first performing the transformation to (ϑ, r) coordinates under the Stratonovich interpretation, and then considering the equivalent Itô SDE in the new coordinates. For convenience, we define

$$\mathcal{G}(\vartheta) = \sum_i G_i \sum_j \frac{\partial G_j}{\partial x_i} \frac{\partial \vartheta}{\partial x_j} \Big|_{r=0} = G^T(x(\vartheta, 0)) \cdot DG^T(x(\vartheta, 0)) \cdot \nabla_x \vartheta \Big|_{r=0}, \quad (6.18)$$

where the T denotes the transpose and DG is the Jacobian of the vector function G with respect to x . Upon performing the phase reduction of (6.16), we have, for $D \ll 1$

$$d\vartheta = (1 + D [Z(\vartheta)Z'(\vartheta) + Y - \mathcal{G}]) dt + \sigma Z(\vartheta) dW_t. \quad (6.19)$$

This has an equivalent Stratonovich SDE:

$$d\vartheta = (1 + D [Y - \mathcal{G}]) dt + \sigma Z(\vartheta) \circ dW_t. \quad (6.20)$$

Here, we can see the presence of extra terms compared with (6.11) in both of the equivalent SDEs. Following the same procedure as before, and considering expansions up to $O(D)$, we may find the steady state probability distribution as

$$Q_0(\vartheta) = \frac{1}{\Delta} + \frac{D}{\Delta} [Z(\vartheta)Z'(\vartheta) - Y(\vartheta) + \bar{Y} + Z(\vartheta) - \bar{\mathcal{G}}] + O(D^2), \quad (6.21)$$

where $\bar{\mathcal{G}} = \Delta^{-1} \int_0^\Delta \mathcal{G} d\vartheta$, and the mean frequency as

$$\Omega = 1 + D(\bar{Y} - \bar{\mathcal{G}}) + O(D^2). \quad (6.22)$$

To highlight the importance of the differences between the interpretations of the stochastic integral, we now go through a numerical example.

6.2.2 Phase distribution and frequency of the SL model

As our model of choice, we study the SL model, so that F in (6.1) is given by (2.42) from Sect. 2.4, with $\lambda = 2, \omega = 1, c = 1$ and $\sigma = 0.2$. The unforced SL model has an attracting limit cycle solution given by $\Gamma(\vartheta) = (\cos(\vartheta), \sin(\vartheta))$, which has period $\Delta = 2\pi$. If we set $x = (1+r) \cos(\vartheta + c \ln(1+r))$ and $y = (1+r) \sin(\vartheta + c \ln(1+r))$, then (ϑ, r) give the desired isochronal coordinates. Since there is no difference between the Itô and Stratonovich interpretations for purely additive noise, we must use multiplicative noise to highlight the difference between the two, and so we take $G = (x, 0)$ to be our forcing function. For the SL model with this form of G , we may calculate the functions Z and Y as

$$Z(\vartheta) = -\cos \vartheta (\sin \vartheta + c \cos \vartheta), \quad Y(\vartheta) = c \cos^2 \vartheta (-\cos 2\vartheta + c \sin 2\vartheta), \quad (6.23)$$

and

$$\mathcal{G}(\vartheta) = -\cos \vartheta (\sin \vartheta + c \cos \vartheta) = Z(\vartheta), \quad (6.24)$$

so that

$$\bar{Y} = -\frac{c}{4}, \quad \bar{\mathcal{G}} = -\frac{c}{2}. \quad (6.25)$$

To compare the analytical predictions for the steady state phase distribution and the mean frequency to those obtained by the solution of the full model, we solve (6.1) with our chosen F and G . Since $\sigma \ll 1$, we may use the numerical scheme for system with weak noise from [92]. Taking y_n as approximations to the true solution x_n at time steps nh , $n = 1, \dots, N$, where $h = T/N$ is the stepsize and T is the total integration length, we have

$$y_{n+1} = y_n + \frac{h}{6}(k_1 + k_2 + k_3 + k_4) + \sigma g(y_n) dW_n \quad (6.26)$$

where $dW_n \sim \sqrt{h}N(0, 1)$ and the k_i , $i = 1, \dots, 4$, are given by the usual Runge-Kutta steps:

$$\begin{aligned} k_1 &= f(x_n), & k_2 &= f(x_n + k_1/2), \\ k_3 &= f(x_n + k_2/2), & k_4 &= f(x_n + k_3). \end{aligned} \quad (6.27)$$

The numerical scheme (6.26) then has a mean square error on the order of $O(h^4 + \sigma h + \sigma^2 h^{1/2})$. This scheme is appropriate for Itô systems. For the Stratonovich case, we must thus first write down the equivalent Itô SDE before (6.26) can be used. The equivalent Itô SDE to the Stratonovich SDE given by (6.1) with our choice of G is

$$dx = \left(F(x) + \frac{G(x)}{2} \right) dt + \sigma G(x) dW_t. \quad (6.28)$$

To obtain the steady state phase distribution, we solve (6.1) and (6.28) using (6.26) to obtain sample paths for 1000 distinct initial conditions and realisations of the noisy process. We take $T = 10^4$, with $N = 10^7$, so that $h = 10^{-4}$, and discard the first $N/2$ time points, treating them as transients. The remaining points are then transformed to (ϑ, r) coordinates using

$$\vartheta = \arctan_2(y, x) - c \log(\sqrt{x^2 + y^2}), \quad (6.29)$$

$$r = \sqrt{x^2 + y^2} - 1, \quad (6.30)$$

where

$$\arctan_2(y, x) = \begin{cases} \arctan(y/x) & x > 0 \\ \arctan(y/x) + \pi & y \geq 0, x < 0 \\ \arctan(y/x) - \pi & y \leq 0, x < 0 \\ \pi/2 & y > 0, x = 0 \\ -\pi/2 & y < 0, x = 0 \end{cases} \quad (6.31)$$

Following this, the ϑ values are then binned into 500 non-overlapping bins to form a histogram, which is then normalised to give an estimate of the marginal probability density Q . To estimate the mean frequency, we use the sample trajectories in Cartesian coordinates and the same method as in Sect. 4.5, using the crossing of y through 0 with $x > 0$ to define a spiking event, and the crossing of x through 0 with $y > 0$ to define a resetting event. Fig. 80 shows a comparison of the analytical steady state distributions under the Itô interpretation (red line) and the Stratonovich interpretation (blue line) along with results from direct numerical simulation, depicted by points with colours matching the lines. Fig. 81 shows the same information, but for the mean frequency of the SL oscillator with increasing D . Whilst the predictions for the phase distribution are qualitatively similar, the same is not true for the mean frequency. In both cases, we see excellent agreement between the analytical and numerical results, but also that the results for the Itô and Stratonovich interpretations differ. An inspection of (6.22) compared with (6.14), noting the values of \bar{Y} and $\bar{\mathcal{G}}$ shows that we expect the opposite behaviour in terms of frequency shift under the two different interpretations. Under the Stratonovich interpretation, the mean frequency falls as D is increased, whilst the opposite is true in the Itô case.

6.3 A new coordinate system for planar systems

It is the purpose of this part of the chapter to provide a general perspective on how to construct planar phase-amplitude coordinates, that allows one to study dynamics in the neighbourhood of a limit cycle for arbitrary vector fields F . In

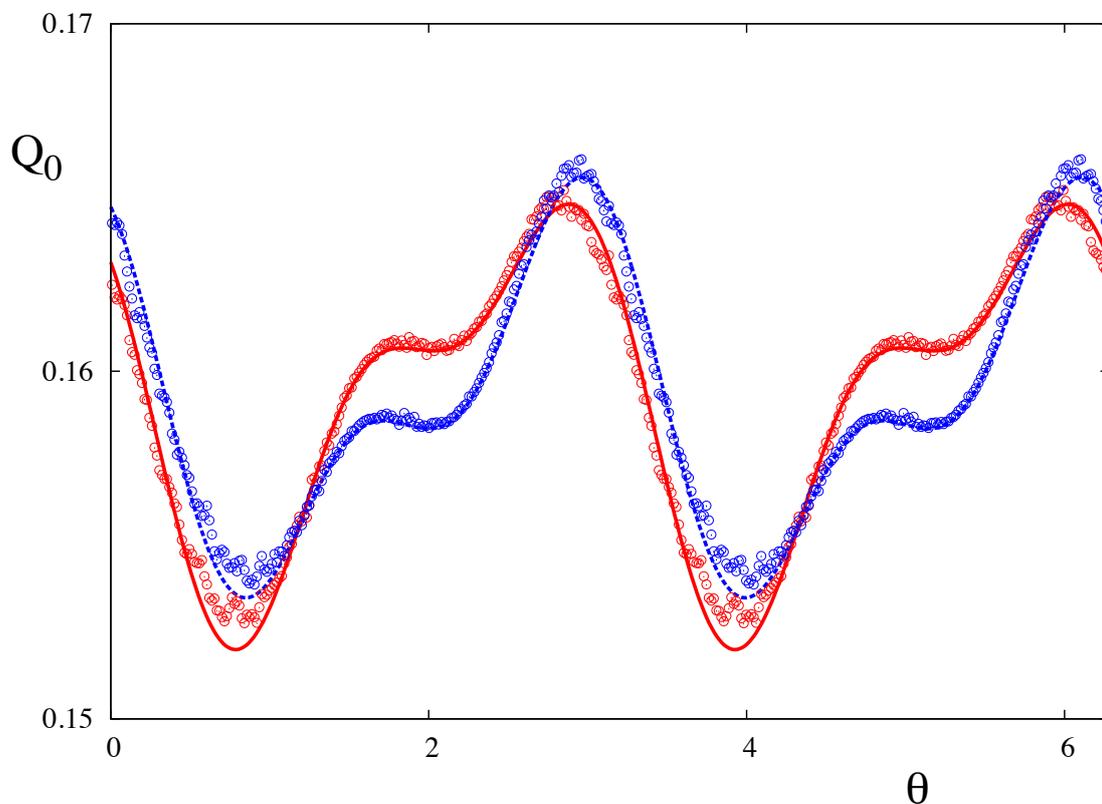


Figure 80: Comparison of steady state phase distributions under the Itô and Stratonovich interpretations of the stochastic integral of (6.1) for the SL model. The red, solid curve and blue, dashed curves correspond for the Itô and Stratonovich cases respectively, where Q_0 is taken from (6.12) and (6.21), whilst the red and blue circles are the numerical approximations found by solving (6.1) and (6.28). We observe excellent agreement between the analytical and numerical approximations, but note the differences between the two interpretations. Although qualitatively similar, we note the quantitative differences between the distributions.

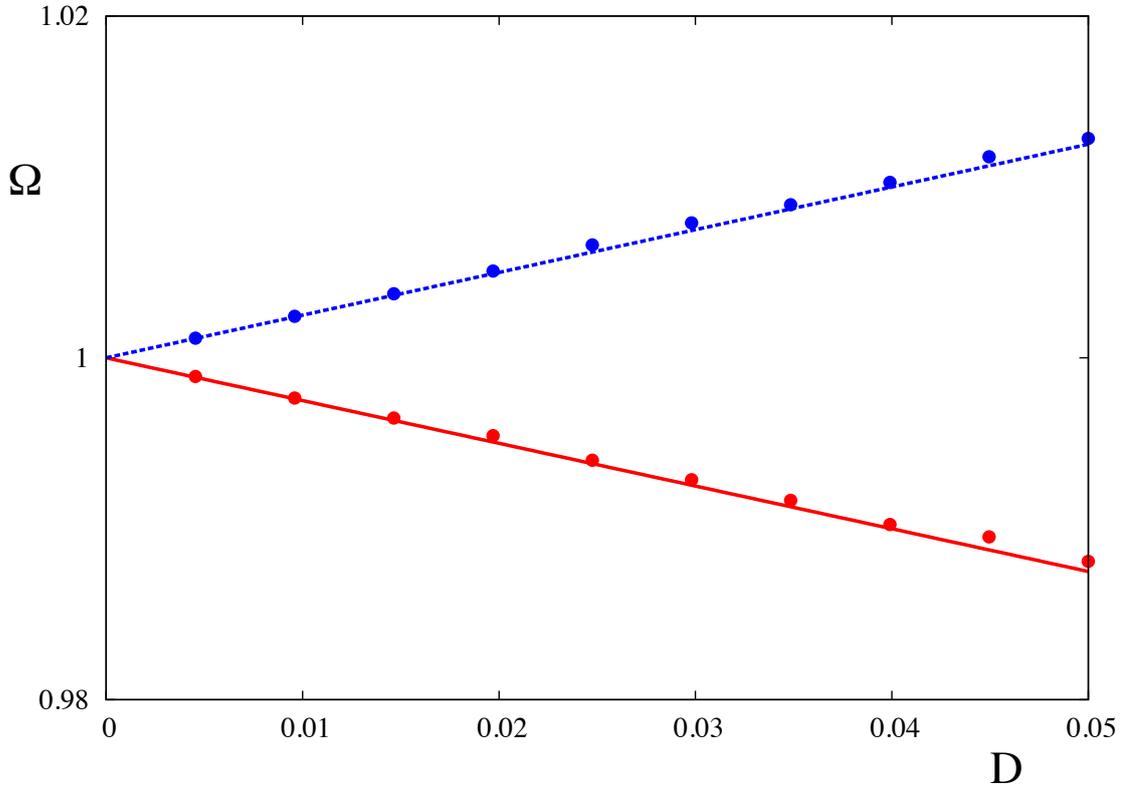


Figure 81: Comparison of mean frequency under the Itô and Stratonovich interpretations of the stochastic integral of (6.1) for the SL model. The red, solid curve and blue, dashed curves correspond for the Itô and Stratonovich cases respectively, where Ω is taken from (6.14) and (6.22), whilst the red and blue circles are the numerical approximations found by solving (6.1) and (6.28). As in Fig. 80, we observe excellent agreement between the analytical and numerical approximations. However, we note that the predictions for the two interpretations are now not qualitatively similar. The Stratonovich interpretation predicts a decrease in Ω as D increases, whilst the prediction is reversed for the Itô case. This can be seen by noting the extra $\bar{\mathcal{G}}$ term in (6.22) compared with (6.14). This example highlights the importance of the interpretation of the stochastic integral when performing phase reductions.

the previous chapter, we used phase-amplitude coordinates, taking the amplitude coordinate, ρ to be the distance along a normal vector to Γ at θ . We note that θ in this coordinate system no longer represents asymptotic phase. Using these coordinates, we were able to describe the effects of invariant structures a finite distance away from Γ , that would be missed by phase-only models. One limitation of this coordinate system is that the transformation ceases to become invertible after some finite distance $|\rho|$ is eclipsed. This problem is associated with the crossing of normals of Γ at distinct values of θ , so that the mapping taking $x \mapsto (\theta, \rho)$ is no longer one-to-one. As such, this is less of a problem for high dimension systems, but may present a serious issue for planar systems. Before moving to consider stochastic oscillators, let us take a step back to consider the deterministic version of (6.1), that is

$$\dot{x} = F(x) + \sigma G(x), \quad x \in \mathbb{R}^2, \sigma \in \mathbb{R}. \quad (6.32)$$

We seek to introduce a phase-amplitude coordinate system that retains some notion of a distance from cycle, but does not suffer from the same coordinate system break down. To do this we introduce coordinates (θ, ρ) such that

$$x(\theta, \rho) = \Gamma(\theta) + \rho J \Gamma'(\theta) + \frac{(\rho J)^2}{2} \Gamma''(\theta) + \dots = \exp \left\{ \rho J \frac{d}{d\theta} \right\} \Gamma(\theta). \quad (6.33)$$

Here, J generates a rotation of $\pi/2$:

$$J = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \quad (6.34)$$

To zero'th order in ρ , (6.33) describes a point on the limit cycle ($x(\theta, 0) = \Gamma(\theta)$). To first order in ρ it defines a point which is on a line orthogonal to the cycle (since $\Gamma'(\theta)$ is tangent to the cycle), which is simply a non-normalised version of (5.2). It is simple to show that

$$\frac{\partial x}{\partial \rho} = J \frac{\partial x}{\partial \theta}. \quad (6.35)$$

These are the Cauchy-Riemann conditions, so that the change of variables between x and (θ, ρ) is locally conformal, and so will not break down anywhere. Given

that for a general nonlinear planar dynamical system, Γ can only be obtained numerically (excluding special cases, such as given by (2.42)), to make practical use of the formal infinite series (6.33) we introduce a Fourier representation for Γ :

$$\Gamma(\theta) = \sum_{n=-\infty}^{\infty} \Gamma_n e^{in\theta}, \quad \Gamma_n = \frac{1}{2\pi} \int_0^{2\pi} d\theta \Gamma(\theta) e^{-in\theta} \in \mathbb{R}^2. \quad (6.36)$$

In this case x , has the representation

$$x(\theta, \rho) = \sum_n e^{2\pi in\theta/\Delta} e^{2\pi in\rho J/\Delta} \Gamma_n, \quad (6.37)$$

which may be implemented numerically using matrix exponential operations and the sum over n may be truncated at some prescribed level of approximation for describing Γ . We note that we may also write x avoiding the use of any complex variables if we so desire, through the use of hyperbolic functions as

$$x(\theta, \rho) = \sum_n \cosh\left(\frac{2\pi n\rho}{\Delta}\right) \left\{ \Gamma_{c,n} \cos\left(\frac{2\pi n\theta}{\Delta}\right) + \Gamma_{s,n} \sin\left(\frac{2\pi n\theta}{\Delta}\right) \right\} + \sinh\left(\frac{2\pi n\rho}{\Delta}\right) \left\{ J\Gamma_{s,n} \cos\left(\frac{2\pi n\theta}{\Delta}\right) - J\Gamma_{c,n} \sin\left(\frac{2\pi n\theta}{\Delta}\right) \right\}, \quad (6.38)$$

where $\Gamma_{(s,c),n} \in \mathbb{R}^2$ are the sine and cosine Fourier coefficients of Γ .

6.3.1 Transformed dynamics

We now provide the dynamics for the evolution of θ and ρ . Using $x = x(\theta, \rho)$ and differentiating with respect to t gives

$$\dot{x} = \frac{\partial x}{\partial \theta} \dot{\theta} + \frac{\partial x}{\partial \rho} \dot{\rho}. \quad (6.39)$$

Projecting this onto the orthogonal vectors $\partial x/\partial \rho$ and $\partial x/\partial \theta$ gives the dynamics

$$\dot{\theta} = \frac{\left(\frac{\partial x}{\partial \theta}\right) \cdot [F + \sigma G]}{\left|\frac{\partial x}{\partial \theta}\right|^2}, \quad \dot{\rho} = \frac{\left(\frac{\partial x}{\partial \rho}\right) \cdot [F + \sigma G]}{\left|\frac{\partial x}{\partial \rho}\right|^2}. \quad (6.40)$$

We may also calculate $\partial x/\partial \theta$ in the useful form:

$$\frac{\partial x}{\partial \theta} = \sum_n \left(\frac{2\pi in}{\Delta}\right) e^{2\pi in\theta/\Delta} e^{2\pi in\rho J/\Delta} \Gamma_n, \quad (6.41)$$

with $\partial x/\partial \rho$ given by (6.35). For $\rho = 0$ we obtain the result

$$\dot{\theta} = 1 + Z(\theta), \quad Z(\theta) = \frac{F(\Gamma(\theta))}{|F(\Gamma(\theta))|^2} \cdot G(\Gamma(\theta)), \quad (6.42)$$

where we have made use of the fact that $F(\Gamma(\theta)) = \Gamma'(\theta)$. We interpret $Z(\theta)$ as the phase response to the perturbation G . For notational simplicity, we shall rewrite (6.40) in the form

$$\dot{\theta} = \Omega(\theta, \rho) + \sigma h(\theta, \rho), \quad \dot{\rho} = f(\theta, \rho) + \sigma g(\theta, \rho), \quad (6.43)$$

where

$$\begin{aligned} \Omega(\theta, \rho) &= \frac{\frac{\partial x}{\partial \theta} \cdot F}{\left| \frac{\partial x}{\partial \theta} \right|^2}, & h(\theta, \rho) &= \frac{\frac{\partial x}{\partial \theta} \cdot G}{\left| \frac{\partial x}{\partial \theta} \right|^2}, \\ f(\theta, \rho) &= \frac{\frac{\partial x}{\partial \rho} \cdot F}{\left| \frac{\partial x}{\partial \rho} \right|^2}, & g(\theta, \rho) &= \frac{\frac{\partial x}{\partial \rho} \cdot G}{\left| \frac{\partial x}{\partial \rho} \right|^2}. \end{aligned} \quad (6.44)$$

The limit cycle (for $\sigma = 0$) is given by $\rho = 0$ with uniform rotation rate $\Omega(\theta, 0) = 1$.

6.3.2 The SL oscillator

To demonstrate the use of the coordinate system (θ, ρ) , we will now write the dynamics of (2.42) in the new coordinate system. The model has a limit cycle given by $\Gamma = (\cos \theta, \sin \theta)^T$, having period $\Delta = 2\pi$. Rather than resorting to the use of Fourier coefficients, we seek to write phase points x using (6.33). To do this, we must take derivatives of Γ . It is easy to show that

$$\Gamma^{(n)} = (-J)^n \Gamma, \quad n = 1, 2, \dots \quad (6.45)$$

where the superscript denotes differentiation with respect to θ , and that

$$J^2 = -\mathbb{I}_2, \quad J^{n+3} = -J^{n+1}, \quad n = 1, 2, \dots \quad (6.46)$$

where \mathbb{I}_2 is the 2×2 identity matrix so that we may write (6.33) as

$$x(\theta, \rho) = \Gamma(\theta) + \rho J \Gamma'(\theta) + \frac{(\rho J)^2}{2} \Gamma''(\theta) + \frac{(\rho J)^3}{6} \Gamma'''(\theta) + \dots \quad (6.47)$$

$$= \Gamma + \rho \Gamma + \frac{\rho^2}{2} \Gamma + \frac{\rho^3}{6} \Gamma + \dots \quad (6.48)$$

$$= \left(1 + \rho + \frac{\rho^2}{2} + \frac{\rho^3}{6} + \dots \right) \Gamma \quad (6.49)$$

$$= e^\rho \Gamma. \quad (6.50)$$

Differentiating this with respect to θ and ρ , we have

$$\frac{\partial x}{\partial \theta} = -e^\rho J \Gamma, \quad \frac{\partial x}{\partial \rho} = e^\rho \Gamma, \quad \left| \frac{\partial x}{\partial \theta} \right|^2 = \left| \frac{\partial x}{\partial \rho} \right|^2 = e^{2\rho}. \quad (6.51)$$

Using these relations, we may go on to find Ω and f as

$$\Omega = 1 + \frac{\lambda c}{2} (1 - e^{2\rho}), \quad f = \frac{\lambda}{2} (1 - e^{2\rho}). \quad (6.52)$$

These two functions may be seen, plotted over ρ , in Fig. 82.

6.3.3 Steady state distribution for weak forcing

We shall now return to consider the stochastic system (6.1). For additive noise, both the Itô and Stratonovich SDE for the phase-amplitude coordinates will be given by

$$d\theta = \Omega(\theta, \rho)dt + \sigma h(\theta, \rho)dW_t, \quad d\rho = f(\theta, \rho)dt + \sigma g(\theta, \rho)dW_t, \quad (6.53)$$

Since we use Fourier coefficients to describe the orbit Γ , it is convenient to also describe P_0 in terms of its Fourier coefficients. Setting

$$Y(\theta) = \frac{\partial h(\theta, 0)}{\partial \theta} g(\theta, 0), \quad (6.54)$$

and adopting a Fourier representation for Q_0 , Z and Y as

$$Q_0(\theta) = \sum_n Q_{0,n} e^{in\theta}, \quad Z(\theta) = \sum_n Z_n e^{in\theta}, \quad Y(\theta) = \sum_n Y_n e^{in\theta}, \quad (6.55)$$

allows us to obtain a set of equations for the unknown amplitudes $Q_{0,l}$ solving (6.10) as

$$-Q_{0,l} + D \sum_{n,m} Z_n Z_m i(l-n) Q_{0,l-(n+m)} - D \sum_n Y_n Q_{0,l-n} = K \delta_{l,0}, \quad l \in \mathbb{Z}, \quad (6.56)$$

for some constant K . For $D = 0$, we have that $Q_{0,0} = K$, and after enforcing normalisation we may set $K = 1/\Delta$. For small D , we may then substitute $Q_{0,0}$ into (6.56) and work to next order in D to obtain an approximation for the remaining amplitudes, $l \neq 0$, in the form

$$Q_{0,l} = \frac{D}{\Delta} \left\{ \sum_{\{n,m \mid n+m=l\}} Z_n Z_m i m - Y_l \right\}. \quad (6.57)$$

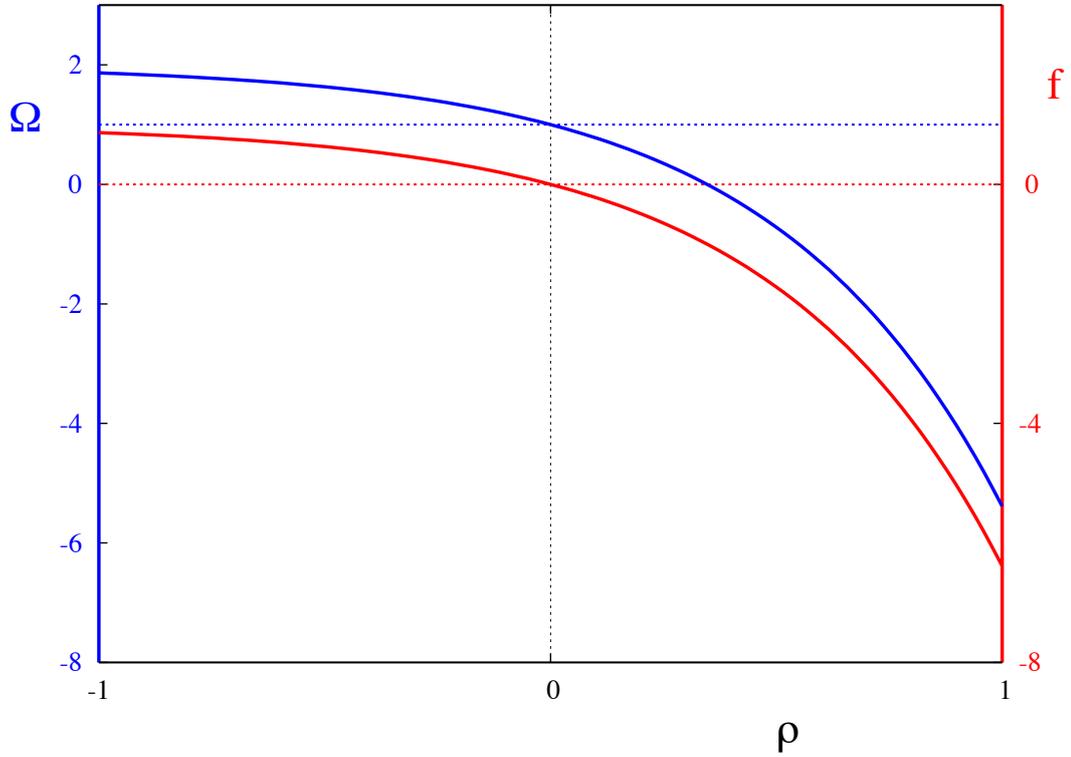


Figure 82: Functions Ω and f for the SL model in the new phase amplitude coordinate system defined by (6.33). Both functions are independent of θ , and have a similar dependence on ρ . The vertical dashed line represents the limit cycle, at $\rho = 0$, whilst the blue and red lines respectively depict the values of Ω and f at $\rho = 0$. As expected, $f(\theta, 0) = 0$, $\partial f/\partial \rho|_{\rho=0} < 0$, reflecting the stability of the limit cycle. Similarly, $\Omega(\theta, 0) = 1$, so that θ evolves with unit speed on Γ , whilst $\partial \Omega/\partial \rho|_{\rho=0} < 0$, suggesting that trajectories with $\rho < 0$ evolve faster (in terms of θ) than those with $\rho > 0$.

Using this we may reconstruct the distribution $Q_0(\theta)$, for small D , from (6.55) as

$$Q_0(\theta) = \frac{1}{\Delta} + \frac{D}{\Delta} [Z(\theta)Z'(\theta) - Y(\theta) + Y_0], \quad Y_0 = \frac{1}{2\pi} \int_0^{2\pi} Y(\theta) d\theta. \quad (6.58)$$

In a similar fashion to (6.14), the mean frequency of the oscillator is defined by

$$\tilde{\omega} = \omega + DY_0 + O(D). \quad (6.59)$$

We may also calculate the phase-diffusion \tilde{D} as

$$\begin{aligned} \tilde{D} &= \int_{-\infty}^{\infty} \langle [d\theta(t+\tau) - \langle d\theta \rangle][d\theta(t) - \langle d\theta \rangle] \rangle d\tau \\ &= \int_{-\infty}^{\infty} \langle Z(\theta(t+\tau))dW_{t+\tau} Z(\theta(t))dW_t \rangle d\tau + O(D^2) \\ &= D \int_0^{2\pi} Z^2(\theta) d\theta + O(D^2), \end{aligned} \quad (6.60)$$

where we use the fact that $\langle dW_t \rangle = 0$ and $\langle dW_s dW_t \rangle = \delta(t-s)$. This recovers a result of Kuramoto [140] as expected.

6.3.4 Steady state distributions for non-weak forcing

As done in Sect. 6.2.2, if we numerically simulate a large number of sample paths over a long time, we can bin these paths to form histograms which, when normalised, will approximate the steady state probability distribution for a given model. An alternative way to calculate the steady state distribution is by solving the Fokker–Planck equation presented in Sect. 2.20.1. The advantages of solving the Fokker–Planck equation over numerical simulation are that we do not need to generate and process large volumes of data, and that the Fokker–Planck equation obeys the normal rules of calculus, as opposed to Itô’s formula. For many one-dimensional systems, we can solve the Fokker–Planck equation explicitly, and even go on to say something more about quantities of interest in the system. For systems possessing more variables, finding explicit solutions is rare, and so we must use numerical techniques to find an approximate solution. In the remainder of this chapter, we will be solving both the Fokker–Planck equation in Cartesian coordinates, in plane polars, and in the new phase-amplitude coordinate system,

(6.40). In the former, the problem set-up can be found in Appx. C, whilst its equivalent for the latter two cases is detailed in Appx. C.1.

In both cases, we need numerical techniques to solve the Fokker–Planck equation. Often used in engineering problems, particularly those involving fluid dynamics, finite element methods are an efficient way of solving PDEs [205]. These methods work by dividing the domain S into a finite number of mesh elements, the topology of which are decided generally on the problem type. The PDE to be solved is then written in its weak form, after having been multiplied by some test function and integrated over the problem domain to deal with higher order derivatives. The equation is then solved over the mesh, writing solutions at the nodal values of the mesh elements as coefficients of some basis functions, such that these nodal values are the exact solution to the PDE at these points [206]. Finite element methods are generally efficient, as the number of elements required for an accurate solution is generally not too large. In addition, the flexibility of the discretisation of the domain allows complex geometries to be considered, although this is not an issue for the present problem, since we are solving on a regular domain. We use the toolbox FEniCS [207]. Written in the Unified Form Language (UFL) Dolfin, FEniCS presents a flexible, easy-to-use toolbox, also possessing extra toolboxes that allow for quick visualisation of solutions and refinement of meshes. FEniCS scripts may be written in C++ or python. In either case, expressions and functions are first converted to C-code by FEniCS to allow for efficient computation of solutions. The details for solving the Fokker–Planck equation in FEniCS can be found in Appx. D.

An alternative to the finite element method is the finite difference method [208]. Instead of the discretising the domain into elements and solving the weak problem, the finite difference method divides the domain by discrete differences, giving rise, for planar systems, to a rectangular lattice on which the solution to the strong form of the problem is approximated. As a comparison, we note that the simplest mesh elements in the finite element scheme, namely Lagrangian elements, are triangular, rather than rectangular. Finite difference methods have the advantage that the

equations need not be converted to the weak form before solving. Additionally, it has been recently demonstrated that finite difference methods may be more accurate than finite element methods in approximating the tails of the probability distribution [209]. In Appx. E, we present the steps we use for solving the Fokker–Planck equation using the finite difference method.

As an example of using the finite difference method to solve the Fokker–Planck equation in the phase-amplitude coordinate system, we return to the SL system (2.42). We have already found Ω and f for this model. It is not difficult to find the functions h and g in this case, and in fact, they turn out to be simple trigonometric functions. However, since we are primarily interested in defining functions through Fourier coefficients along with the underlying vector field in Cartesian coordinates alone, we instead use (6.44), with (6.41). To highlight the difference in the approximation of the steady state distribution for weak noise, to the true distribution, we first find the weak noise approximation using (6.58). This turns out to be a flat distribution $Q_0(\theta) = 1/(2\pi)$, $\forall \theta \in [0, 2\pi)$. As we increase σ , this approximation will fail to capture the nonlinear dependence on σ , and so we expect a non-flat probability distribution for non-weak σ . To test this, we discretise the domain $S = [0, 2\pi) \times [-1, 1]$ and evaluate the functions Ω, f, g and h at the mesh points. Since these functions are independent of σ , we need only evaluate them once, and we can use these values for any given σ . Following this, for $\sigma \in (0, 0.4]$, we use the scheme as described in Appx. E, with the upwind scheme discussed in Appx. E.1, with 1000 mesh points in each of the two axes, to find the steady state distribution in terms of (θ, ρ) . We then numerically integrate out ρ , to form the marginal density $Q(\theta) = \int_{-R}^R P(\theta, \rho) d\rho$. This is then subtracted from Q_0 to give the difference between the weak noise approximation and the solution from the full system. The results are displayed in Fig. 83. For small σ , as expected, we see little difference between the weak approximation and the full distribution. As σ increases, there arises a periodic modulation in the steady state distribution in terms of θ , as the nonlinear dependence of Q on σ grows. The amplitude of the variations in the steady state distribution increases

as σ continues to grow, and the location of the extrema of said distribution is also dependent on σ , but we note that the maximum deviations from the mean seem to be symmetric in terms of their magnitude. Numerical histograms of the steady state distribution were also formed using the algorithm as detailed in Appx. B, and the same integration method as in Sect. 6.2.2. We show a comparison of the steady state distributions approximated by the Fokker–Planck equation and those obtained through forming histograms of sample paths in Fig. 84. In this figure, the black dashed line is the weak noise approximation to the steady state distribution, the blue curve is the solution from the Fokker–Planck equation, whilst the red dots represent the approximation from histograms of sample paths. The top panel shows distributions for $\sigma = 0.01$, and we observe excellent agreement between all three approximations. In the bottom panel, σ has increased to 0.2. Now, we see that the two non-weak approximations agree very well with one another, but that they differ from the flat distribution for the weak noise approximation.

6.4 The subcritical Hopf

In general, the summative effects of noisy perturbations in a given system may not be weak. In this setting, we may expect more exotic behaviour from the system, such as mixed mode oscillations [195]. We will now consider a system which is bistable, and study transitions between the two stable states brought about as a result of the stochastic perturbations. As our model of choice, we use the subcritical Hopf bifurcation presented in [210] as a canonical model for cortical activity in the alpha band (~ 10 Hz). The model is a fifth order Hopf, and is written in polar coordinates (ϕ, r) as

$$\dot{\phi} = \phi_c, \quad \dot{r} = -r^5 + \lambda r^3 + \beta r, \quad (6.61)$$

where $\phi_c > 0$, $\lambda, \beta \in \mathbb{R}$. The ϕ dynamics of (6.61) suggest that the system oscillates with constant frequency ϕ_c . A bifurcation analysis of the r dynamics, as plotted in Fig. 85, shows that, for small β , the system has $r = 0$ as the only stable fixed point. Thus, although ϕ evolves at a constant rate, r stays at the origin, and

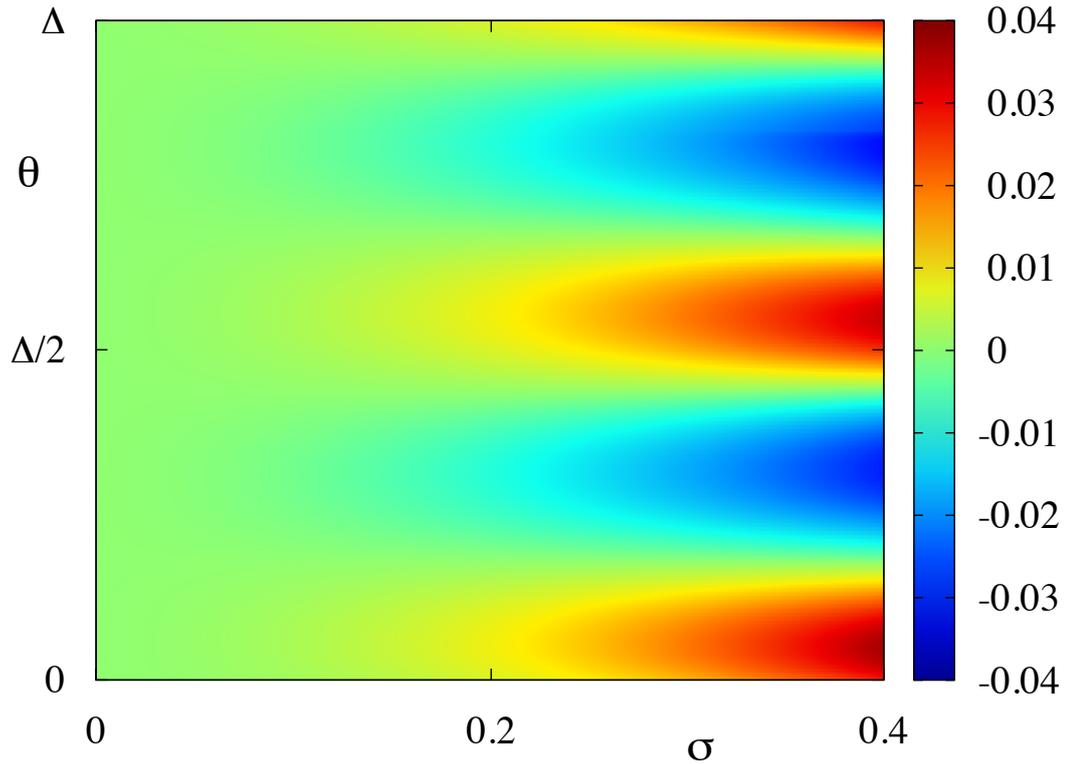


Figure 83: Difference between steady state approximations for the weak noise case of the SL model, with those formed by solving the Fokker–Planck equation in the non-weak case, under variation of σ . The distribution for the weak noise case is flat, with $Q_0 = 1/(2\pi) \forall \theta \in [0, 2\pi)$. For small σ , the two distributions agree well. As σ increases, we observe a periodic modulation around the mean, which grows in amplitude as σ increases. Numerical histograms corresponding to $\sigma = 0.01$ and $\sigma = 0.2$ may be seen in Fig. 84.

so we would not actually observe oscillations in the dynamics (following transients). We refer to this as quiescence. As we increase β , the system undergoes a bifurcation giving rise to bistability. In this bistable region, there are, as expected, five fixed points of the system for $r \in \mathbb{R}$. However, since we are working in polar coordinates, r is restricted to the non-negative real line. Due to the symmetry of the fixed point equation about $r = 0$, two pairs of these fixed points are equivalent, and so we only need consider the region $r > 0$. The quiescent state, with $r = 0$ is still stable, but

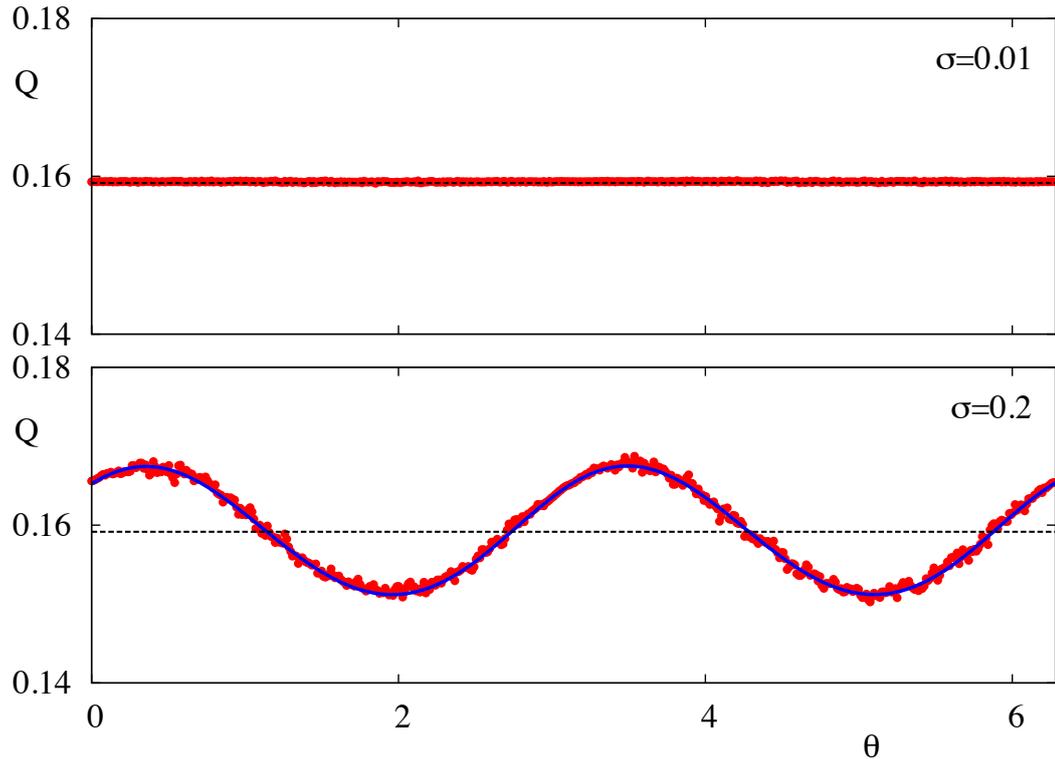


Figure 84: Comparison between steady state distributions approximated by solving the Fokker–Planck equation (blue), and by binning sample paths to form histograms (red). Also shown is the approximation for weak noise (black). In the top panel, $\sigma = 0.01$, and we see excellent agreement between all three approximations. In the bottom panel, we observe periodic modulation in both of the non-weak approximations, but excellent agreement between them.

there is now also another stable solution, depicted by the top black curve. Since r is now nonzero along this branch, we now observe oscillatory activity, and we thus refer to this as the oscillatory state. Separating these two stable solutions, is an unstable branch for an intermediate value of r , shown by the grey line. Where these solutions exist, we shall denote the value of r corresponding to the oscillatory state by r_s and that corresponding to the unstable branch as r_u . As we increase β further still, the system undergoes another bifurcation, and now the fixed point at $r = 0$ becomes unstable, and the only remaining stable solution is the oscillatory

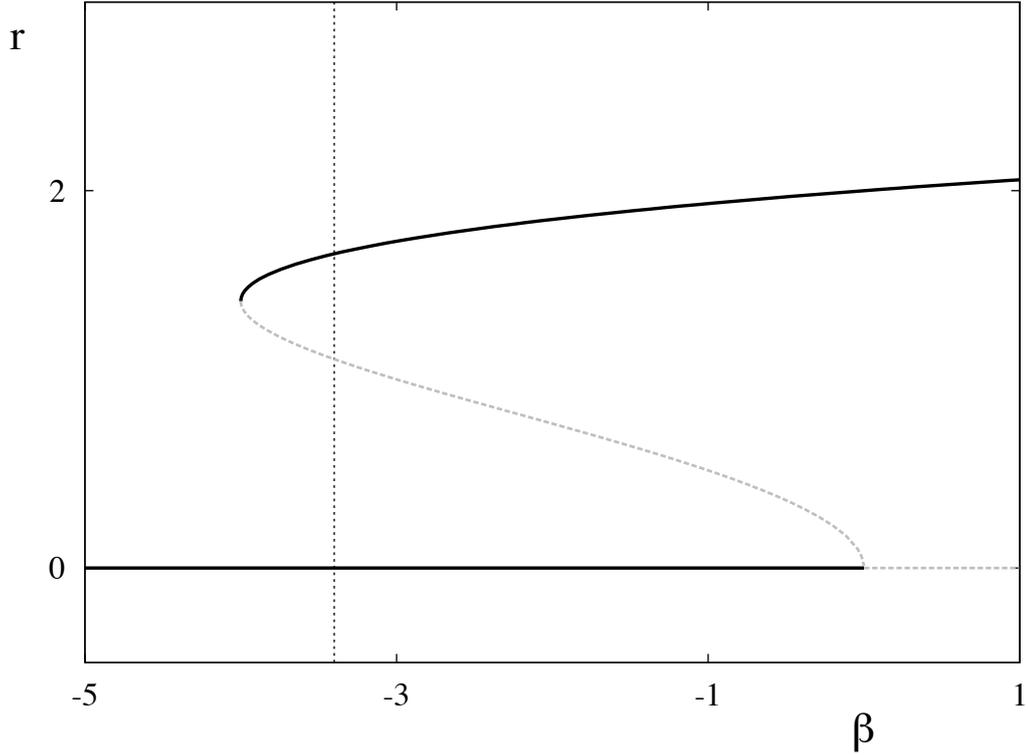


Figure 85: Bifurcation diagram for the r dynamics of (6.61) using β as a control parameter. The solid black lines correspond to stable fixed points, whilst the dashed grey lines are unstable fixed points. We observe a region of bistability in the system. To the left of this bistable region, only the quiescent state is stable, to the right, only the oscillatory state remains. The vertical dotted line shows where we choose the parameter values for our study of this model.

one.

Setting $\phi_c = 10$ ensures that we represent the alpha rhythm, whilst a choice of $\lambda = 4$, $\beta = -3.4$ puts the system in the bistable regime, as shown by the dotted vertical line in Fig. 85. We shall use these parameter values in throughout this section. The authors of [210] then add two white noise processes, one additive and one multiplicative to the r dynamics, leaving the ϕ dynamics as they stand. The SDE for the r dynamics now reads

$$dr = (-r^5 + \lambda r^3 + \beta r)dt + \sigma(1 - \eta)dW_1 + \sigma\eta r dW_2. \quad (6.62)$$

where σ is the noise strength and η is the fraction of noise that is multiplicative whilst $1 - \eta$ is the proportion that is additive. Owing to the fluctuations added by the noise terms, trajectories will now not stay at the fixed points of the r dynamics. Instead, they will fluctuate around these values, and will occasionally cross $r = r_u$ and move towards the other fixed point value. This implies that we need a new definition of the quiescent and oscillatory state. We shall refer to $r < r_u$ as the quiescent state, and $r > r_u$ as the oscillatory state. The authors claim that additive noise alone is not sufficient to replicate the relation between the mean and the fluctuations that they observe in cortical patient data. Clearly, some additive noise is needed, since otherwise the noise would be quenched when trajectories enter the quiescent mode with $r = 0$. The authors perform numerical simulations of (6.62) to obtain histograms representing the steady state probability distribution of the analytic signal, and a distribution of dwell times in either the quiescent or oscillatory modes. We first note that, since there are no perturbations in ϕ , and the governing equations decouple, we can consider the dynamics of r separate to those of ϕ . Hence, the system under consideration is effectively one dimensional, and its dynamics are prescribed by a polynomial function. We may solve the associated forward Fokker–Planck equation to give the steady state distribution explicitly. We may also use the backward Fokker–Planck equation to find the probability distribution of switching times between quiescence and the oscillatory mode. The results from the paper [210] seem to be erroneous, in that the parameter values used for the noise do not correspond to the figures they produce. This can be shown easily by solving the Fokker–Planck equation. We provide the correct results for the model. We note that, since the authors do not include any noisy dynamics in the ϕ dynamics, there is no noise induced frequency shift. In general, a dynamical system representing neural activity is unlikely to have noise acting only in the radial direction. This may be overcome by transforming (6.61) into Cartesian coordinates, resulting in

$$\begin{aligned}\dot{x} &= f_x(x, y) = \beta x - \phi_c y - \lambda x(x^2 + y^2) + x(x^2 + y^2)^2, \\ \dot{y} &= f_y(x, y) = \phi_c x + \beta y - \lambda y(x^2 + y^2) + y(x^2 + y^2)^2.\end{aligned}\tag{6.63}$$

As before, we then use additive white noise in the x dynamics only. Upon transforming this back into polar coordinates, the noise terms will appear in both the ϕ and r equations, and will be multiplicative in both cases. The drawback of including such noise terms is that the equations no longer decouple, so that we cannot find an exact solution for the steady state probability distribution. However, we can use the methods outlined in Appx. D to find a numerical solution instead.

6.4.1 Solution for the 1D case

The forward Fokker–Planck equation for (6.62) is

$$\frac{\partial P}{\partial t} = -\frac{\partial}{\partial r}[(-r^5 + \lambda r^3 + \beta r)P] + \frac{\sigma^2}{2} \frac{\partial^2}{\partial r^2} [((1 - \eta)^2 + \eta^2 r^2)P]. \quad (6.64)$$

The steady state is found by setting $\partial P/\partial t = 0$, whereupon the resulting equation is an ODE. So that we can impose natural boundary conditions at infinity, we shall solve the Fokker–Planck equation on an infinite domain. Suppose that we have the steady state distribution on $r \in (-\infty, \infty)$ as \tilde{P} , and we wish to find the distribution restricting r to the positive real line, denoted by P . Since the equations are symmetric under the transformation $r \mapsto -r$, we have that $P(r) = \tilde{P}(r) + \tilde{P}(-r) = 2\tilde{P}(r)$, $\forall r \in [0, \infty)$. For convenience, we will drop the tilde in what follows. Assuming that P and its derivative vanish at $r = \pm\infty$, we deduce from (6.64)

$$-(-r^5 + \lambda r^3 + \beta r)P + \frac{\sigma^2}{2} \frac{d}{dr} [((1 - \eta)^2 + \eta^2 r^2) P] = 0, \quad (6.65)$$

whence

$$\frac{dP}{dr} = 2 \frac{-r^5 + \lambda r^3 + \beta r - \sigma^2 \eta^2 r}{\sigma^2 [(1 - \eta)^2 + \eta^2 r^2]} P. \quad (6.66)$$

This linear equation can be solved, for example, by separation of variables, resulting in

$$P = C e^\Lambda, \quad (6.67)$$

where

$$\begin{aligned}\Lambda &= \frac{2}{\sigma^2} \int \frac{-r^5 + \lambda r^3 + \beta r - \sigma^2 \eta^2 r}{(1 - \eta^2) + \eta^2 r^2} dr \\ &= \frac{1}{2\sigma^2 \eta^6} (\eta^2 r^2 (2 - \eta^2 (2 - 2\lambda + r^2)) \\ &\quad - 2(1 + 2\eta^6 + \eta^2(\lambda - 2) - \eta^4(\lambda + \beta - 1)) \log(1 + \eta^2(r^2 - 1))),\end{aligned}\tag{6.68}$$

and C is chosen to satisfy the normalisation condition, so that

$$C = \left(\int_0^\infty e^\Lambda dr \right)^{-1}.\tag{6.69}$$

We may find explicit expressions for r_u and r_s as

$$r_u = \sqrt{\frac{1}{2} \left(\lambda - \sqrt{\lambda^2 + 4\beta} \right)}, \quad r_s = \sqrt{\frac{1}{2} \left(\lambda + \sqrt{\lambda^2 + 4\beta} \right)}.\tag{6.70}$$

We note that the condition that these roots are guaranteed to be satisfied by our choice of parameters. In the unforced system, it is clear that the trajectories will exist in only one of the basins of attraction. In the forced system, trajectories are now free to cross between the two basins. Under appropriate conditions, this may lead to a bimodal probability distribution, which will have peaks near $r = 0$ and $r = r_s$. We may find conditions for the existence of a bimodal distribution, as well as the location of these peaks, when they exist, by using (6.66). Solving this equation after setting $dP/dr = 0$ in this equation will give the information we seek. Thus, we must find the roots of the equation

$$r(-r^4 + \lambda r^2 + \beta - \sigma^2 \eta^2) = 0.\tag{6.71}$$

One of the roots is clearly zero, corresponding to the quiescent state. The other roots, for $r > 0$ are given by

$$r_\pm = \sqrt{\frac{1}{2} \left(\lambda \pm \sqrt{\lambda^2 + 4\beta - 4\sigma^2 \eta^2} \right)},\tag{6.72}$$

where r_+ corresponds to a peak of the distribution, whilst r_- corresponds to a trough. Note that, where these roots exist, they do not coincide with the location of stable and unstable limit cycles of the unforced system, which are given by (6.70), except when either σ or η vanish. Thus, the multiplicative nature of the noise acts to

shift the location of the attractors and repellers of the system. The condition for these roots to exist on the real line is that $\lambda^2 + 4\beta > 4\sigma^2\eta^2$. At the point where $4\sigma^2\eta^2 = \lambda^2 + 4\beta$, under the variation of the either of the two noise parameters σ, η , there is a bifurcation, resulting in the probability density over r becoming bimodal. This differs from the deterministic bifurcations (which are called dynamic bifurcations), which give rise to the creation of new attractors. Indeed, no new attractors have been created in the system, rather this is a so-called P -bifurcation, as it is a change in the qualitative nature of the probability distribution. For a review of P -bifurcations, and their differences from dynamic bifurcations, we refer the reader to [211]. The condition for a bimodal distribution suggests only a unimodal distribution around $r = 0$ exists for large values of σ and η , whilst bimodal distributions are possible for smaller values. Recalling that η is the noise parameter that gives the fraction of noise that is applied in a multiplicative way, this suggests that multiplicative noise applied in this way is more likely to lead to quiescence. This is likely due to the form of noise applied to the system; since the contribution of the term $\sigma^2\eta^2r^2$ is small when r is small, trajectories are likely to remain near $r = 0$ when η is close to 1. This tendency to remain near $r = 0$ under such forcing is the likely reason for the deviation in switching times for the purely additive noise case as observed in [210]. In Figs. 86 and 87, we show this P -bifurcation using η and σ as the control parameter respectively. In both figures, the colour corresponds to the probability at a location r . Since we are interested in the shape of the distributions, rather than their quantitative values, and for ease of comparisons, the probability distributions are scaled so that the global maximum is set equal to 1. The black lines show the minima and maxima (away from $r = 0$) in the probability distributions as a function of the control parameter. In both cases, the bifurcations, characterised by the coming together of a peak and trough are clearly shown. The red curve inset shows the bifurcation plot of the switch between a bimodal and unimodal distribution under variation of η and σ simultaneously. Thus, in Fig. 86, we take a horizontal cross-section of the inset, marked by the white dotted line, whilst in Fig. 87, we take a vertical

cross-section. In both cases, the bifurcation point in the main figure is observed to correspond to a crossing of a bifurcation curve in the inset. Although the bifurcation from a bimodal to a unimodal distribution is qualitatively similar in both figures, the behaviour past the bifurcation is not. As the noise moves from being primarily additive to multiplicative, under the increase of η , we observe a tightening of the probability distribution. In fact, beyond the bifurcation point, the distribution is tightly grouped around the remaining peak at $r = 0$. This behaviour is fairly robust across a range of σ , for reasons discussed earlier. Conversely, as the noise strength σ increases, the probability distribution broadens following the bifurcation. Although there now exists no peak corresponding to the stable limit cycle, there is still a nontrivial probability at this point. Fig. 87 is plotted for $\eta = 0.3$, in which the noise is mostly additive. As η increases, the effect of the broadening of the distribution is less pronounced, in which case the distribution remains principally around the peak at $r = 0$. Even though the distribution may be bimodal, it does not guarantee that the probability at the peaks is high. Indeed, an examination of Fig. 87 shows that, for small σ , there is switch between high power in the peak corresponding to the stable limit cycle to the peak corresponding to quiescence. In fact, the probability at $r = r_+$ decreases rapidly under variation of σ , whilst the quiescent peak grows rapidly in this same window. We may consider this further through an examination of how the probability changes under variation in σ . To this end, we may calculate

$$\frac{\partial P}{\partial \sigma} = \frac{\partial C}{\partial \sigma} e^\Lambda + C \frac{\partial \Lambda}{\partial \sigma} e^\Lambda, \quad (6.73)$$

where

$$\begin{aligned} \frac{\partial \Lambda}{\partial \sigma} = & \frac{1}{\eta^6 \sigma^3} 2 \log(\eta^2(r^2 + 1) - 2\eta + 1) (\eta^4(-\beta + \lambda + 1) \\ & + \eta^6 \sigma^2 - 2\eta^3(\lambda + 2) + \eta^2(\lambda + 6) - 4\eta + 1) \\ & + \eta^4 r^4 - 2\eta^2 r^2 (\eta^2(\lambda + 1) - 2\eta + 1) - \frac{2 \log(\eta^2(r^2 + 1) - 2\eta + 1)}{\sigma} \end{aligned} \quad (6.74)$$

and

$$\frac{\partial C}{\partial \sigma} = \left(\int_0^\infty e^\Lambda dr \right)^{-2} \int_0^\infty \frac{\partial \Lambda}{\partial \sigma} e^\Lambda dr. \quad (6.75)$$

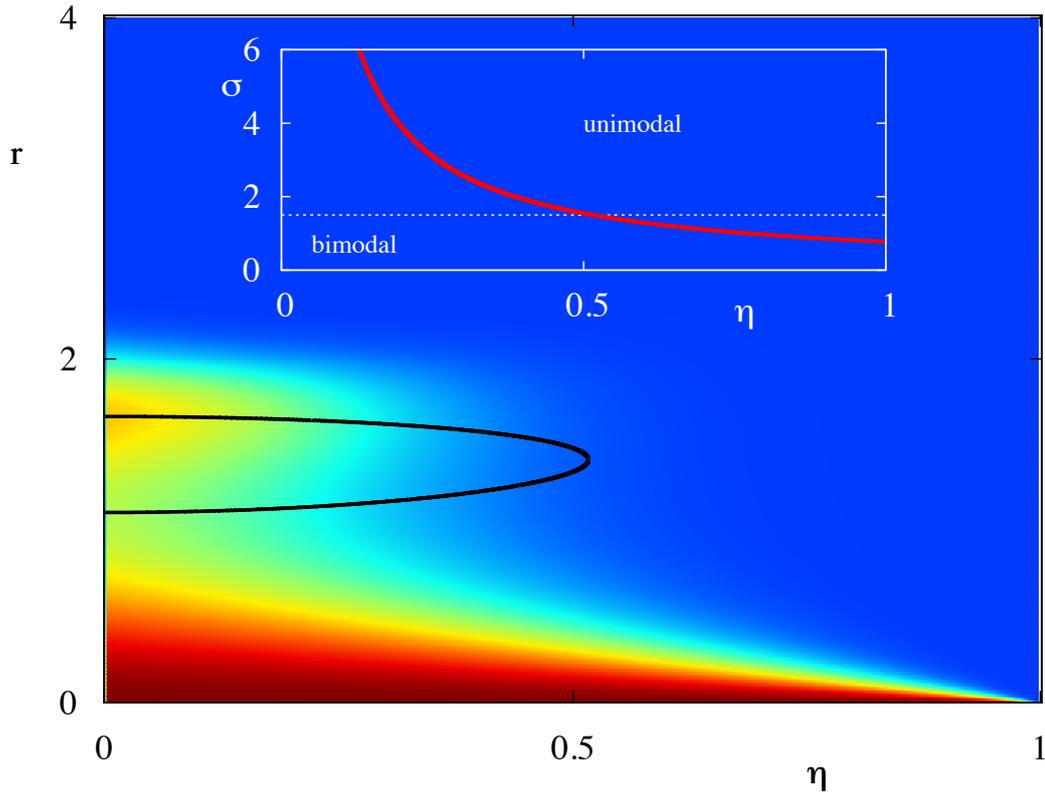


Figure 86: Distribution of probabilities under variation of η . Distributions have been scaled so that the maximum value of the distribution is equal to 1, reflected by the colour. The black lines correspond to the minima and maxima (away from $r = 0$). A P -bifurcation, whereupon the distribution changes from being bimodal to being unimodal can be observed where the black curves meet. The red curve in the inset is the bifurcation diagram showing this transition under the simultaneous variation of η and σ . The main figure thus corresponds to taking a horizontal cross-section, as shown by the dotted white line. Here, $\sigma = 1.5$.

Zeros of (6.73) will tell us when probability begins to accumulate, or to leave the point r . In Fig. 88, we plot both the probability at the peak corresponding to the stable limit cycle, with its derivative shown below. Since the probability distribution is sharply peaked for small σ , rather than plotting $P(r_+)$, we instead plot $L^* = \log(P(r_+))$, and similarly, we plot $dL^*/d\sigma$. We observe that L^* is initially high, however rapidly decreases under the increase in σ . The vertical dotted line

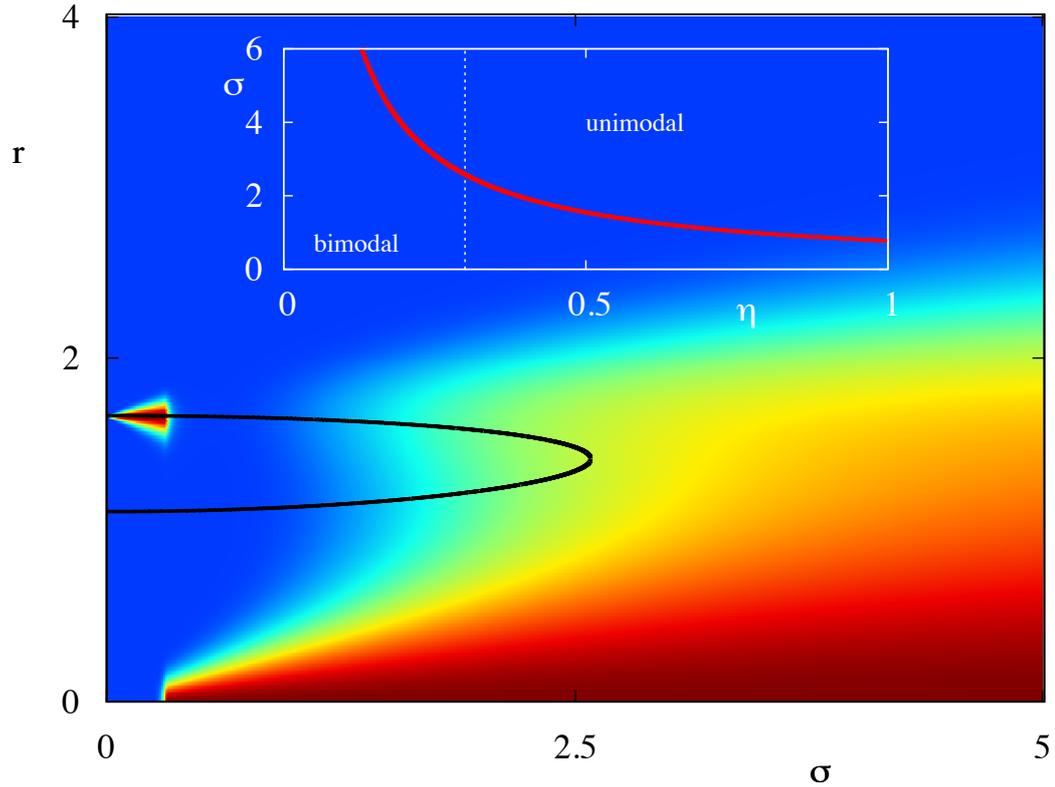


Figure 87: This figure is equivalent to Fig. 86, except using σ as the control parameter. This corresponds to now taking a vertical cross-section through the bifurcation diagram shown in the inset. Here, $\eta = 0.3$.

shows the point at which $dL^*/d\sigma = 0$, after which L^* begins to increase again and no more extrema of the first derivative are found before the P -bifurcation, after which P is unimodal. This predicts that we expect near regular oscillatory activity for very small σ , which is quickly replaced by quiescence under small increases in σ . After this, we would expect intermittent oscillatory and quiescent behaviour, switching between the two modes for intermediate σ , up to the P -bifurcation.

Now that we know where we expect switching between oscillatory activity and quiescence, we wish to know about the expected dwell times in each of these regimes. We may set up the backward Fokker–Planck equations which give the distribution of exit times from a given region (in this case, these regions will be intervals). The first exit time for a point r in a given interval is the time at which

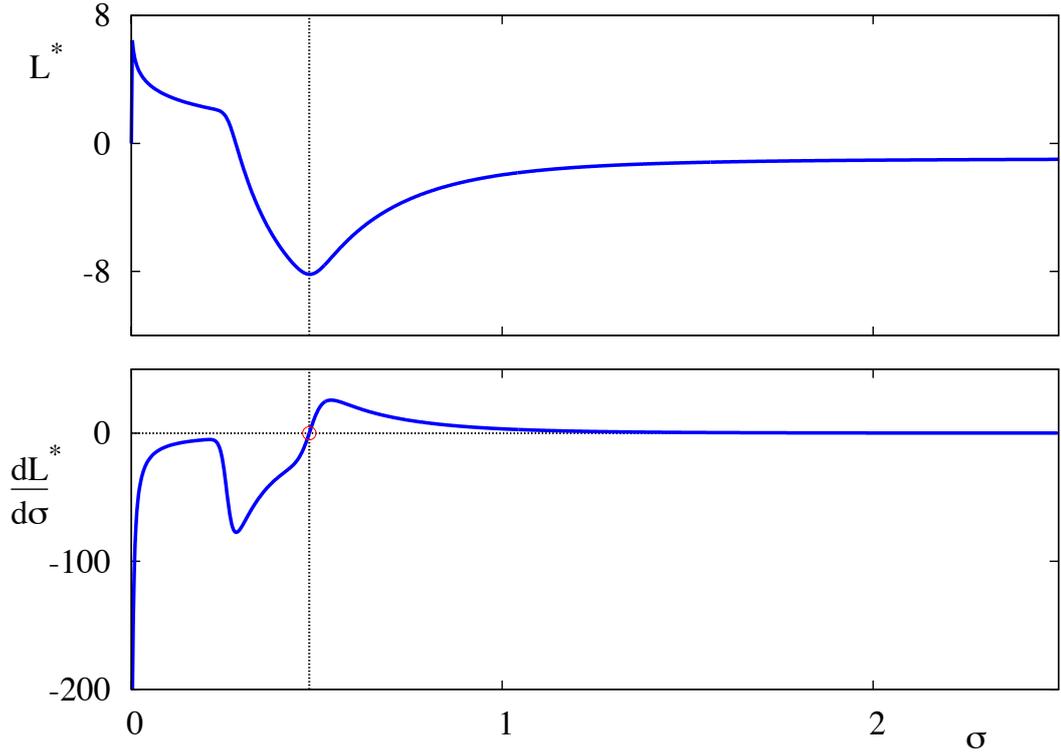


Figure 88: Natural logarithm of the probability at the peak $r = r_+$ under variation of σ (top), along with its derivative with respect to σ (bottom). The vertical dotted line shows the point where this derivative is equal to zero, at which point probability begins to accumulate at this peak. Note the rapid decrease in probability at the peak for small σ . Parameter values are as in Fig. 87.

the true solution, with initial conditions at r leaves the region of interest. Defining $\tau(r)$ to be the first exit time from a given point in the region of interest, we set $P_\tau(r, t) = P(\tau(r) > t)$ to be the probability that τ is larger than t . This enables us to write down the following backward Fokker–Planck equation

$$\frac{\partial P_\tau}{\partial t} = (-r^5 + \lambda r^3 + \beta r) \frac{\partial P_\tau}{\partial r} + \frac{\sigma^2}{2} ((1 - \eta)^2 + \eta^2 r^2) \frac{\partial^2 P_\tau}{\partial r^2}. \quad (6.76)$$

After noting that the cumulative distribution function (CDF) of exit times is given by $1 - P_\tau$, the solution to (6.76), can be used to find the CDF for each spatial point in S . Now, in each of the two firing modes, we expect sample paths to remain near one of the peaks of the steady state probability distribution, making

infrequent transitions to the other peak. This means that the CDF evaluated at the peaks should approximate the behaviour in that firing mode, thus we wish to evaluate the distribution of exit times from $r = 0$ and $r = r_+$ to $r = r_-$, so that our regions of interest are $R_1 = [0, r_-)$ and $R_2 = (r_-, R)$, where $R \gg 1$. It remains to establish boundary and initial conditions for (6.76). Clearly, $P_\tau(r, 0) = 1$, since $\tau(r) > 0, \forall r \in (0, \infty)$, regardless of which peak we assume we start from. If we are interested in the exit from $r = 0$, we solve the problem in R_1 , and use the absorbing boundary condition $P_\tau(r_-, t) = 0$, since points at r_- will exit the region instantaneously, and the reflecting boundary condition $\partial P_\tau(r, t)/\partial r|_{r=0} = 0$. Similarly, if we need the exit time from $r = r_s$, we solve the problem in R_2 and use the conditions $P_\tau(r_-, t) = 0$, and $\partial P_\tau(r, t)/\partial r|_{r=R}$. The probability density function (PDF) of exit times, which is the quantity of interest, is given by the derivative, with respect to t , of the CDF.

In order to solve (6.76), we may simply use the finite difference scheme as described earlier, using a backward Euler step for the time discretisation, so that

$$\frac{\partial P_\tau(r, t)}{\partial t} = \frac{P_\tau(r, t + \delta t) - P_\tau(r, t)}{\delta t} + O(\delta t), \quad (6.77)$$

for $\delta t \ll 1$. This scheme results in a system of linear equations to be solved at each time step, which may be performed using Gaussian elimination, for example. To obtain the PDFs from $P_\tau(0, t)$ and $P_\tau(r_s, t)$, we use cubic splining to take the derivative of the CDFs. Setting $\rho = 0.3$, let us consider the effect of increasing σ on the exit time distributions. In Fig. 89, we show the PDFs for exit from the quiescent mode (bottom) and the oscillatory mode (top). In order to facilitate comparison of distributions across values of σ , we plot the results against $\log t$ rather than t . If P is the probability with respect to t , then $\hat{P} = tP$ will give the probability in terms of $\log t$. As expected, we observe a reduction in the peak of the distribution, as increased noise tends to increase the frequency of switching between the two modes. Additionally, we note the tightening of distributions as σ increases for the exit from quiescence, meaning that long dwell times are less likely as σ increases. Comparing the exit distributions for the two regions, we note that the peaks of the distribution for exit from the oscillatory mode are at lower

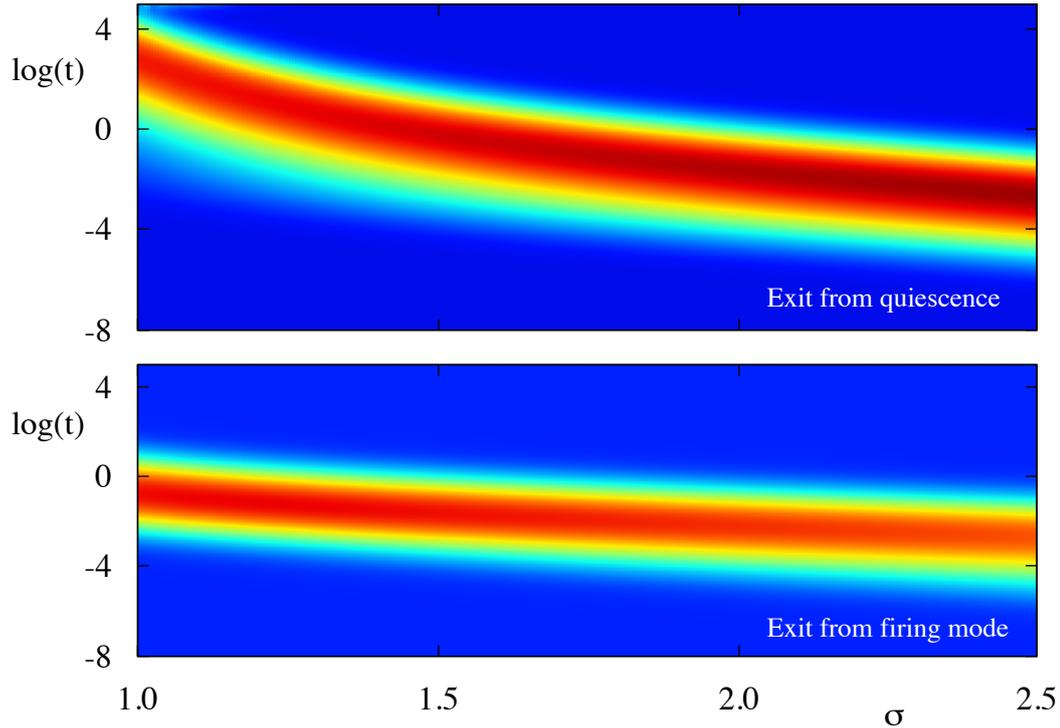


Figure 89: Exit time distributions from quiescence to the oscillatory mode (top) and from oscillatory behaviour to quiescence (bottom). The colour denotes the probability in $\log t$ coordinates. As we increase the noise strength, the peak of the distribution moves to lower values of $\log t$, as expected. We also observe a tightening of the distribution. Note that this figure indicates that, for the chosen value of $\eta = 0.3$, transitions from oscillatory behaviour to quiescence are faster than the converse.

values of t than for the quiescent mode, for small σ , but that they converge as we approach the bifurcation point. This is unsurprising, since $r_s - r_u$ is smaller than r_u , as can be seen in Fig. 85, suggesting that transitions from oscillation to quiescence are ‘easier’ than the converse. This observation suggests that transitions from oscillation to quiescence occur quicker for small noise, suggesting that it is harder to reach an oscillatory mode from quiescence. Let us now turn to the question of how the type of noise affects the distributions, by fixing σ and varying η from 0 up to the bifurcation point. The numerical considerations are exactly

the same as before. We show, in Fig. 90, the analogous results for Fig. 89, now varying η , fixing $\sigma = 1.5$. For the exit distributions from quiescence, we observe that the peak shifts towards higher values as η increases. This is due to the lower contribution of noise sources near $r = 0$. In addition, the distribution broadens as we increase η , suggesting that the transitions are less regular as the noise becomes more multiplicative in nature. There is no discernible difference in the exit time distributions from $r = r_+$, suggesting that transitions from oscillatory behaviour to quiescence are mediated through noise strength, rather than how the noise acts.

6.4.2 The two-dimensional case

As mentioned previously, it is unlikely that the noise will act selectively in the radial coordinate. In order to introduce more natural noise effects, we apply additive noise in the first component to (6.63). This may then be transformed back into polar coordinates, to give the following system

$$d\phi = \phi_c dt - \frac{\sigma}{r} \sin(\phi) dW_t, \quad dr = (-r^5 + \lambda r^3 + \beta r) dt + \sigma \cos(\phi) dW_t. \quad (6.78)$$

We now only have one noise source, but it appears in the dynamics for both ϕ and r . The Fokker–Planck equation for the steady state distribution associated with (6.78) is given by

$$\begin{aligned} 0 = & -\phi_c \frac{\partial P}{\partial \phi} - \frac{\partial}{\partial r} [(-r^5 + \lambda r^3 + \beta r)P] \\ & + \frac{\sigma^2}{2r^2} \frac{\partial^2}{\partial \phi^2} [\sin^2(\phi)P] - \frac{\sigma^2}{2} \frac{\partial^2}{\partial \phi \partial r} \left[\frac{1}{r} \sin(2\phi)P \right] + \frac{\sigma^2}{2} \frac{\partial^2}{\partial r^2} [\cos^2(\phi)P]. \end{aligned} \quad (6.79)$$

Our domain of interest is now the rectangle $S = [0, 2\pi) \times (0, R]$, for $R \gg 1$. For this equation, we need to impose periodic boundary conditions at $\phi = 0$ and $\phi = 2\pi$, since they are identified. To find the steady state distribution P , we must also impose reflecting boundary conditions at $r = 0$ and $r = R$, where $R \gg 1$. Since we are only interested in the distribution in terms of the radial coordinate r , we may integrate out ϕ , to give the marginal density $Q(r) = \int_0^{2\pi} P(\phi, r) d\phi$.

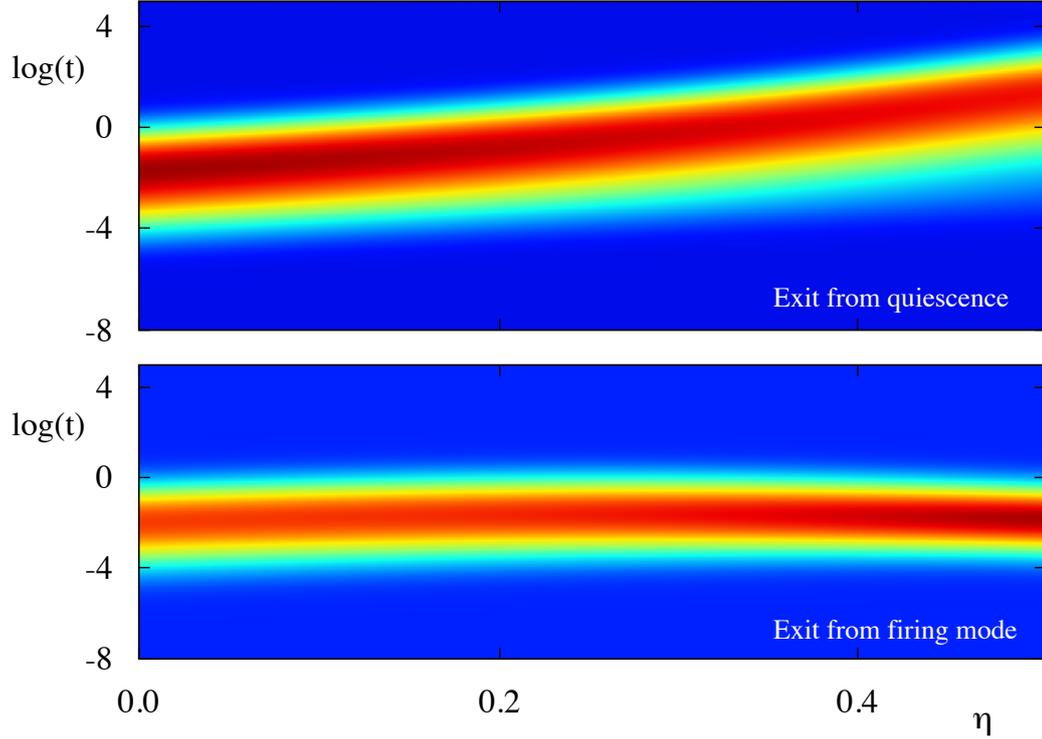


Figure 90: Exit time distributions from quiescence to the oscillatory mode (top) and from oscillatory behaviour to quiescence (bottom). The colour denotes the probability in $\log t$ coordinates. For the transition from quiescence to oscillation, as we increase η , the peak of the distribution moves to higher values of $\log t$, as expected. We also observe broadening of the distribution. There appears to be a very limited dependence of the exit time distribution from $r = r_+$ on η , suggesting that these transitions are mediated through the noise strength σ , rather than through η . In this figure, we set $\sigma = 1.5$.

We use the method outlined in Appx. D to solve (6.79), using the SUPG method from Appx. D.1 where needed. After normalising P , we use numerical integration routines to find Q . In Fig. 91, we show how this marginal density varies with increasing σ . Of course, since we now have only one noise source, we cannot discuss the interplay between additive and multiplicative noise. However, we do observe a transition between a unimodal distribution to a bimodal one for increasing σ , in contrast to the one dimensional case, in which the opposite occurred. As σ is

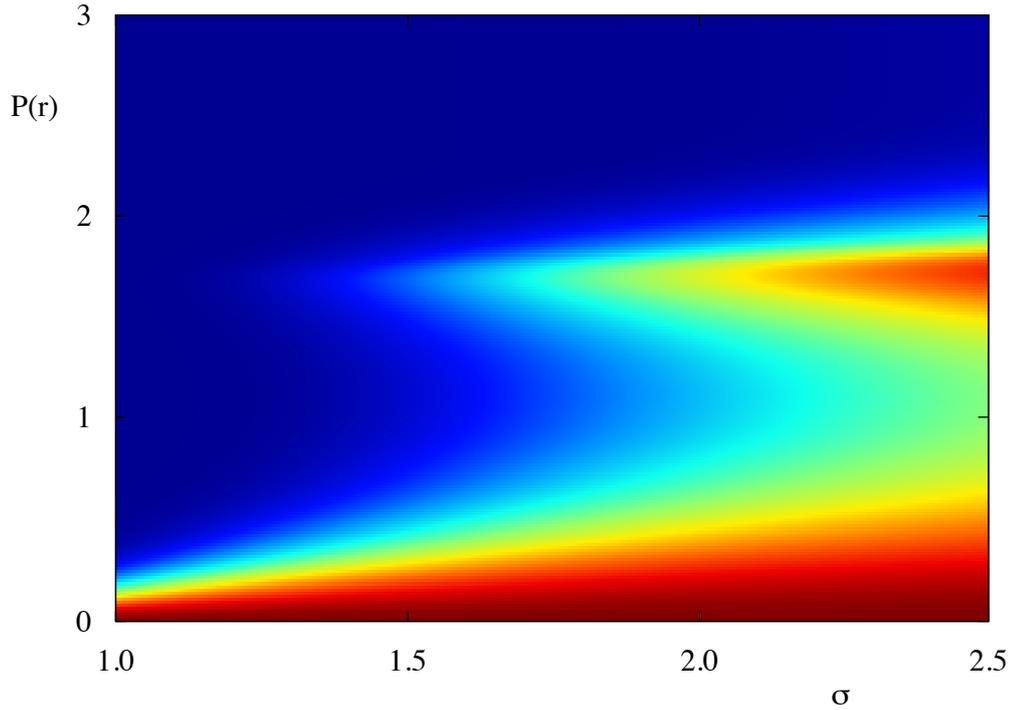


Figure 91: Marginal steady state distribution of (6.78) in r . For ease of comparison, the distributions are scaled so that the peak of the distribution is set equal to 1. As σ increases, we observe a transition from a unimodal to a bimodal distribution, corresponding to bistability between a quiescent mode and an oscillatory one. In contrast to Fig. 87, increasing σ in this case seems to promote bistability, rather than eliminate it. We also note the location of the peaks and troughs in the distribution appear to depend on σ only very weakly, so that the roots of the deterministic equation, given (6.72) may be used to approximate them.

increased further, more probability accumulates in the region corresponding to a stable limit cycle, suggesting that increasing noise in this case supports, rather than quenches oscillatory behaviour. It becomes difficult, in the planar case, to find the switching time distributions, owing to fact that there are now an infinite number of ways to transition between the two modes. The case for switching between quiescence and oscillatory behaviour can still be done by computing the distribution at the origin. However, the same cannot be done for the oscillatory mode, since we would then need to consider the whole limit cycle (rather than just

at the peak in P). Since the point $r = 0$ is not defined in plane polar coordinates, to find the exit time distribution from quiescence to oscillations, we solve the Fokker–Planck equation for (6.63) in Cartesian coordinates, since the two are equivalent once we choose our point of reference. Again, we define $P_\tau(x, t) = P(\tau(x) > t)$, where $\tau(x)$ is the first exit time from a point $x \in S$ out of our domain. For our domain, we choose a disc with radius given by the trough in the marginal distribution Q , so that $S = \{(x, y) | \sqrt{x^2 + y^2} \leq r_u\}$. This circular domain can be handled in FEniCS using standard Lagrangian elements. We note that the boundary of this domain, does not exactly correspond to the limit cycle in the stochastic system, since the multiplicative nature of the noise will serve to move the location of this attractor, and the radial location of the extrema in the steady state distribution will be dependent on ϕ . However, indications from Fig. 92, suggest that this will serve as a good estimate for where to place our boundary. Along this barrier, we impose the Dirichlet boundary condition $P_\tau = 0$, and use the initial condition $P_\tau(x, 0) = 1, \forall x \in S$ as before. The Fokker–Planck equation to be solved is given by

$$\frac{\partial P_\tau}{\partial t} = f_x \frac{\partial P_\tau}{\partial x} + f_y \frac{\partial P_\tau}{\partial y} + \frac{\sigma^2}{2} \frac{\partial^2 P_\tau}{\partial x^2}, \quad (6.80)$$

where f_x, f_y are as in (6.63). In Fig. 92, we see how the distribution changes as we increase σ . As expected, as σ increases, the peak of the distribution shifts to lower values of $\log t$, and there is a tightening of the distribution, although this effect is much less pronounced than for the 1-dimensional model.

Although we cannot find the exit time distribution for the converse transition, we can still ask questions about the spatial distribution of the mean time of this switch. The differential equation for the mean first exit time T , in polar coordinates, is given by

$$\begin{aligned} -1 = \phi_c \frac{dT}{d\phi} + [-r^5 + \lambda r^3 + \beta r] \frac{dT}{dr} \\ + \frac{\sigma^2}{2r^2} \sin^2(\phi) \frac{\partial^2 T}{\partial \phi^2} - \frac{\sigma^2}{2r} \sin(2\phi) \frac{\partial^2 T}{\partial \phi \partial r} + \frac{\sigma^2}{2} \cos^2(\phi) \frac{\partial^2 T}{\partial r^2}. \end{aligned} \quad (6.81)$$

Since we are interested in transitions to the quiescent state, our domain of interest is now $S = [0, 2\pi) \times [r_u, R]$ for $v_R \gg 1$. As before, we set $T = 0$ at the boundary

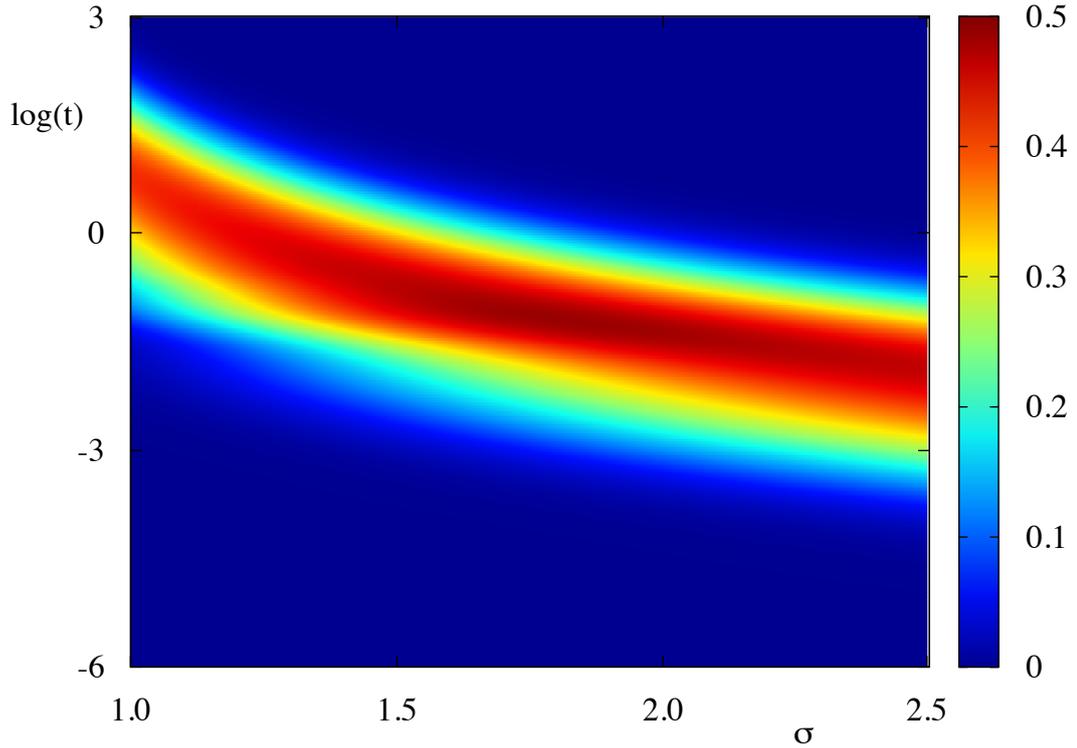


Figure 92: Variation of the exit time distribution from the quiescent mode for the planar Hopf model. The colour represents the probability that the exit from quiescence occurs at a given value of $\log t$. As σ increases, we note a shift in the peak of the distribution to lower values of $\log t$. Simultaneously, there is a tightening of the distribution around this peak, though this effect is much less pronounced than for the one-dimensional model considered in Fig. 89.

$r = r_u$. In addition, we apply a reflecting boundary condition $dT/dr = 0$ at $r = R$ and retain the periodic boundary condition identifying $\phi = 0$ and $\phi = 2\pi$. We note that once the mean first exit time has been found, we may use this to find higher moments of the exit time distribution, though we omit the details here. The solution of (6.81), found through FEniCS may be seen to have an approximately flat distribution across a wide range of σ , indicating that the first passage time to the quiescent state is generally independent of the location within the domain.

We may also wish to know about the effect of the transitions on the probability of firing, by which we mean a passing of trajectories through $\phi = 2\pi$, for $r > r_u$.

Clearly, trajectories from points with $r < r_u$ will escape the quiescent regime with probability 1, so we cannot say much about points here. For trajectories starting from points with $r > r_u$, there are two options: either they will reach the firing threshold before they enter the quiescent mode, or they will not. Using the backward Fokker–Planck equation, we can find the probability that trajectories starting at a given point will reach the firing threshold before they cross $r = r_u$. For this problem, we choose as our domain $S = [0, 2\pi] \times [r_u, R]$. Let us define $P_f(y, t)$ to be the probability that a trajectory starting from a point $y = (\phi, r) \in S$ reaches $\phi = 2\pi$ before crossing $r = r_u$. We may find P_f by solving the equation

$$0 = \phi_c \frac{dP_f}{d\phi} + [-r^5 + \lambda r^3 + \beta r] \frac{dP_f}{dr} + \frac{\sigma^2}{2r^2} \sin^2(\phi) \frac{\partial^2 P_f}{\partial \phi^2} - \frac{\sigma^2}{2r} \sin(2\phi) \frac{\partial^2 P_f}{\partial \phi \partial r} + \frac{\sigma^2}{2} \cos^2(\phi) \frac{\partial^2 P_f}{\partial r^2}. \quad (6.82)$$

We set reflecting boundary conditions at $\phi = 0$ and $r = R$. Since points on $\phi = 2\pi$ will meet our condition with probability 1, we use the boundary condition $P_f(y = (2\pi, r), t) = 1$. Similarly, points with $r = r_u$ will always fail to meet the firing condition, so we set $P_f(y = (\phi, r_u), t) = 0$. The resulting problem can be solved easily in FEniCS, to give Fig. 93, using the functions given in Appx. F. In this figure, the colour represents the spatial profile of P_f , which we plot for three distinct values of σ . In particular, we may ask a question about the probability of repeat firing events, that is, how likely the neuron is to reach threshold again just after being passing through threshold. The points just after threshold correspond to the left hand portion of the plots. The top figure shows P_f for $\sigma = 1$. In this instance, nearly all points have a high probability of firing, so that once the cell has entered an oscillatory regime, it is more likely to reach threshold again before becoming quiescent. As we increase σ , to $\sigma = 2$ in the middle panel and $\sigma = 3$ in the bottom panel, we observe two effects. Firstly, P_f for points just after firing drops significantly, so that repeat firing events without entering the quiescent mode become increasingly unlikely. Secondly, away from $r = r_u$, the domain appears to separate into two distinct regions, separated very roughly at $\phi = \pi$. For points in the region $\phi > \pi$, the drop-off in P_f , under variation of σ ,

is much smaller than for the region $\phi < \pi$, suggesting that, once trajectories have passed $\phi = \pi$, the effect of the noise on P_f is small. This may be due to the location of the peaks of the distribution. In Fig. 94, we plot the probability distribution for (6.79) for $\sigma = 0.2$ in polar coordinates. We see that there are two clear peaks in the distribution along the oscillatory attractor (with $r > 0$), one with $\phi < \pi$ and one with $\phi > \pi$. Between these two peaks, sample paths are more likely to cross $r = r_u$. However, at around $\phi = \pi$, the drift will start to dominate over the diffusion, and sample paths will be attracted towards the peak, away from $r = r_u$. Once they have reached this point, sample paths are then fairly close to the firing threshold at $\phi = 2\pi$, and so the noise is unlikely to have sufficient time to cause sample paths to enter the quiescent mode. This helps explain the split in firing probabilities around $\phi = \pi$. As we continue to increase σ further, the distribution across the whole domain will become flatter, and transitions between quiescence and oscillations will become much more frequent, and so P_f will continue to fall. Concurrently, the exit time distribution from the quiescent state will continue to shift to smaller times, reflecting these faster transitions.

6.5 Discussion

Being biological in nature, neural systems are inherently stochastic. Exactly how detrimental or beneficial noise is to neural computation is still an open topic, but it is clear that mathematical models can be used to capture some aspects of this stochasticity, and may be helpful in providing clues to answer these questions [54, 11]. In this chapter, we have highlighted some of the potential issues when dealing with stochastically forced planar oscillators. For weak forcing, we showed the importance of the interpretation of the stochastic integral. In particular, we demonstrated how the steady state probability distribution changes under the different noise interpretations for a standard oscillatory model. Perhaps of greater consequence, we demonstrated that, for that particular model, increasing the strength of the noise had contrasting effects on the mean frequency of oscillations depending on the interpretation of noise used. In the Stratonovich case, the

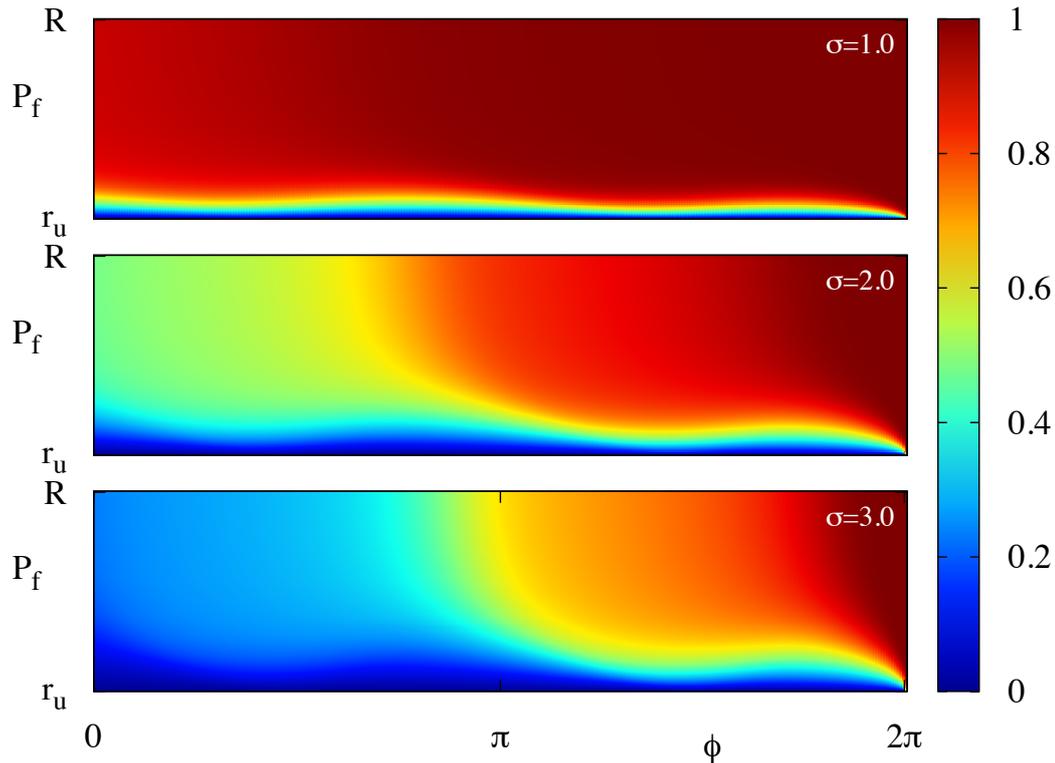


Figure 93: Spatial distribution of the probability of reaching the threshold $\phi = 2\pi$ before entering a quiescent state, denoted P_f , with $r < r_u$, plotted for three distinct values of σ . As σ increases, P_f falls, as expected, but we observe that the domain appears to separate into two distinct regions, around $\phi = \pi$. In each of these regions, the profile of P_f is roughly flat. For $\phi > \pi$, the reduction in P_f as σ increases is much less pronounced, suggesting that the noise has a mitigated effect on P_f in this region.

noise acted to decrease the mean frequency, whilst for the Itô case, the same noise instead increased the mean frequency.

Moving from the single cell to the network level, these effects will likely have an impact on emergent behaviour, and synchronisation properties on the network. The differences between the two increase as we increase the noise strength, so we expect these differences to be exacerbated as we move away from the weak noise limit. In most cases, one would expect a Stratonovich interpretation of the noise to be a good choice for biological system, since we can assume that the noise

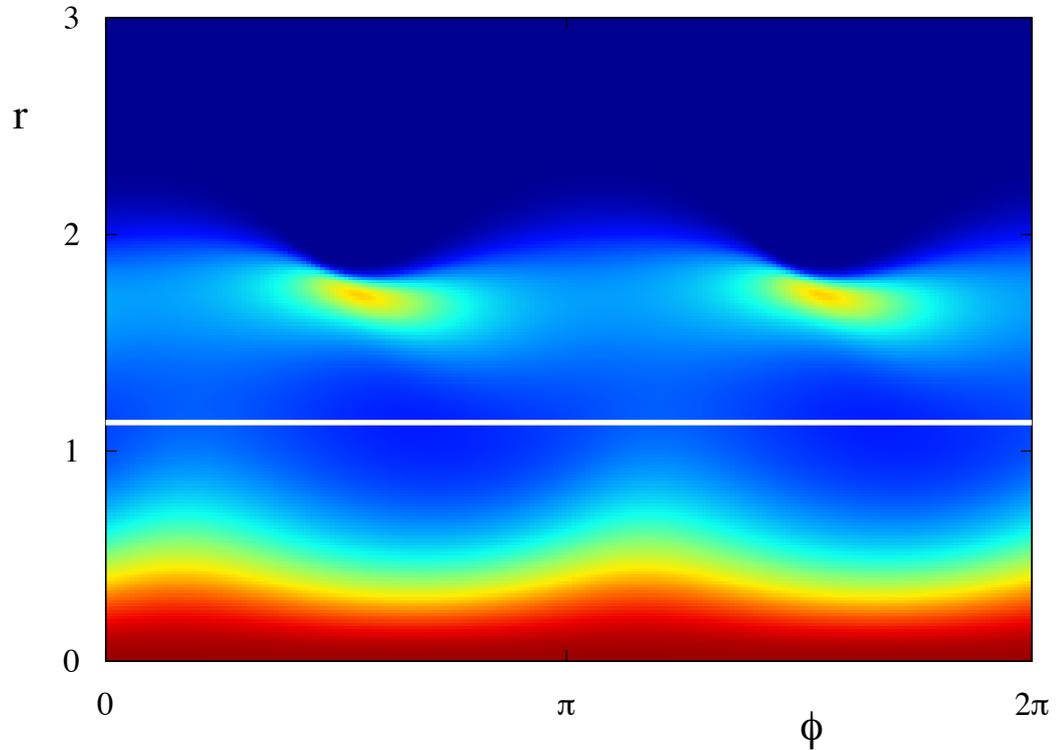


Figure 94: Steady state distribution of (6.79) for $\sigma = 0.2$. Colours represent the probability, whilst the white line is $r = r_u$. As well as the large peak in probability around $r = 0$, we can see two other peaks in the probability distribution in the oscillatory mode. Sample paths will tend towards these peaks as they evolve. Between the peaks, sample paths are likeliest to cross $r = r_u$. Following the second peak, sample paths are nearly at the firing threshold, and so are now likely to fire, even in the presence of noise. This helps to explain why the firing probability in Fig. 93 appears to be split into two distinct regions around $\phi = \pi$.

is, at least in part, due to some underlying process not captured by our model. However, they may also be good reasons for opting for the Itô interpretation. In this chapter, we have primarily focussed on what would generally be regarded as background electrical noise. Such a noise source may be thought of as arising from the sum of many spontaneous bursts of activity, and may not be part of some other process. In fact, the quantal release of neurotransmitter may actually be better modelled as an Itô process. This implies that there is a need for a better dialogue in the community about precisely which of these interpretations is best suited for the study of neural systems. Of course, the answer to this question is likely to be problem dependent; the study of point models differs widely from that of neural field models, and the choice of stochasticity should reflect this. Indeed, the large number of potential sources of randomness in biology in general, from the micro- to the macroscopic level is likely to lead to a diverse choice of mathematical modelling of the system.

Following the work in Chapt. 5, we devised a coordinate system to include both a phase around a limit cycle and an amplitude variable capturing some notion of distance around from the limit cycle. In contrast to the coordinate system from Chapt. 5, the amplitude variable in the new coordinate system does not measure the Euclidean distance from the limit cycle, but instead is proportional to the velocity of the vector field evaluated on-cycle. Although we lose the notion of the Euclidean distance from cycle with the new amplitude variable, the new coordinate system has some nice properties, in particular, that the derivatives of a point x with respect to phase and amplitude coordinates are orthogonal to one another, meaning that the transformation is locally conformal, and does not suffer from the same breakdown condition as the previously studied system. In addition, the dynamics of the phase and amplitude variable may be written compactly, solely in terms of the vector field and Fourier coefficients of the limit cycle. We have shown how to efficiently simulate sample paths of a stochastically forced system in these coordinates to allow for a large number of paths to be computed concurrently. This coordinate system can be easily scaled up to tackle networks of interacting

oscillators, using both phase and amplitude variables, in which coupling between units are stochastic processes, or for studying entrainment of oscillators to a global stochastic signal.

Moving firmly away from the weak limit, we then considered how to analyse a system with bistability between a quiescent and an oscillatory state. Such a model was presented in [210] as a canonical model to explain fluctuations in the alpha band in cortical EEG recordings. We have improved the analysis of this model, firstly, by finding analytical solutions to the model as it is written in the original paper. Secondly, we have used more appropriate parameter values than used by the authors. In fact, we have strong reason to believe that the values quoted in the paper are erroneous. For the one-dimensional case, we have studied the interplay between the noise strength and the way in which the noise acts on the properties of the system. In particular, we have shown how these two factors affect dwell times in each of the stable regions and the location of peaks in the steady state distribution. We have considered the full planar model for the subcritical Hopf, and used a more general form of forcing to examine the effect of background noise on firing probability and dwell times. Finally, we showed the effect that the noise has on the firing probability of points in our domain, and how peaks in the steady state distribution separate regions of high firing probabilities from low ones.

As the field continues to be more interested in stochastic phenomena, and in building larger networks, the need for efficient ways to compute quantities of interest grows. Aside from Monte Carlo simulations of sample paths, one can also resort to solving the Fokker–Planck equation, which can be exploited to find a number of quantities of interest. We have demonstrated two techniques for solving the Fokker–Planck equation: finite element and finite difference methods. Solving the Fokker–Planck equation, particularly for weak noise presents its own set of mathematical challenges. Finding the steady state distribution involves solving a homogenous PDE with natural boundary conditions, for which general solutions are constant. Moreover, any calculation involving small noise terms will give rise to a convection-dominated PDE, which are numerically challenging

to solve. We have presented some ways to tackle these issues; in the first case through the constraining of nodal values, and in the second through the use of upwind schemes. There is still a need to further develop these techniques, both for weak noise, and also to deal with systems possessing more variables [212]. Currently, Fokker–Planck equations up to the fourth dimension can be solved accurately and efficiently [209]. Clearly, if we wish to study networks of interacting oscillators, we need better algorithms for solving these kinds of PDEs. Combining the study of the Fokker–Planck equation with a phase-amplitude coordinate system, such as the one studied in this chapter, will enable us to examine synchronisation properties of small networks away from the weak noise limit, even for systems with other invariant structures in their phase space. This will then pave the way to study complex network rhythms in bistable systems, or networks which have exotic and dynamic synchronisation behaviour, under stochastic forcing. These approaches may also be useful to study the behaviour of coupled bursters, both to study the effects that noise can have of the transitions between the active and silent phase of the burst, and how this is manifest in the network behaviour [213]. Finally, these ideas may also be applied to the study of mixed-mode oscillations, such as those studied in [195], to examine how noise translates into emergent behaviour of the resulting oscillation.

There are clearly many open questions surrounding the role of noise, from a biological, mathematical and modelling perspective. As progress is made in each of these areas, the true role and characteristics of noise in the central nervous system will hopefully come to light.

7 Discussion

In this thesis, we have covered a range of topics surrounding the analysis of neural oscillators. Principally, these works have aimed to use simple models to, firstly, demonstrate the pitfalls in using standard techniques on non-standard problems, and secondly, to develop ways to overcome these issues. Although biophysical models may not capture all of the qualitative features of data, and phenomenological models do not even explicitly include any modelling of the relevant biology, mathematical models can help us to understand the mechanisms of rhythm generation in the central nervous system. The complexity of the models scales up as we couple single cell models together to form network ones. When studying such models, it is not clear which emergent behaviours arise due to the intrinsic dynamics of the oscillators, and which are due to the coupling within the network. One only has to look at examples in [56] to see how network coupling can give rise to interesting behaviour in networks when the individual units are relatively uninteresting. Conversely, in examples of central pattern generators [214, 215], intrinsic dynamics, in conjunction with mutual inhibition, are important in the generation of anti-phase oscillations. By focusing on models of which we have a good understanding at the single cell level, we hope to unravel the contribution from the intrinsic dynamics and coupling at the network level.

In Chapt. 2, we analysed the PWL-IF model as presented in [98]. IF type models are principally used for their computational cheapness; building large coupled networks of IF cells is easy to do, even on personal computers. Since the relevant information for neural processing is thought to be encoded, for the most part, in spike timing, rather than spike shape [32, 33], IF models can be used to study information processing in neural networks. Nevertheless, more recent adaptations of IF models have been fit with considerable accuracy to real data [2], suggesting that, although only phenomenological, these models can capture a wide range of neural behaviour. In our work, we have taken a step away from increasing the complexity of the model, and chose to work in the PWL regime, though we note

that the PWL model we advocate still captures the essence of a spike shape. In this regime, we can construct analytical solutions through the use of matrix exponentials, and by piecing together solutions across the switching manifold, where the governing linear system changes. The advantage of being able to write down solutions in this manner is that we can then probe the dynamics of our system more easily. For the PWL model, this meant we were able to write down exact solutions for the spike time, and return maps for spiking solutions. From here, we could then fully characterise the behaviour of the single neuron in terms of its bifurcation structure. Furthermore, we found structures giving rise to spike adding bifurcations, which are generally difficult to identify in nonlinear models [216, 217], and finally to locate regions where the system is chaotic under constant forcing. Recently, the PWL framework has been used to study canard phenomena in the FHN model [105], and has further been used to examine mixed-mode oscillations for the same model under stochastic forcing [154]. In both examples, explicit solutions allow us to make progress through the identification of separatrices between small and large amplitude oscillations.

The advantages of the PWL paradigm also carry over to the network level. For example, in [98], we constructed networks coupled through gap junctions. In this case, the resultant network dynamics are still linear (since gap junction coupling only adds extra linear currents). We demonstrated how to find the stability of synchronous and asynchronous network solutions, for which the stability conditions are given by explicit formulas.

The non-smooth nature of IF type models gives rise to mathematically interesting dynamics. For the PWL-IF model, there are two regions along which dynamics are non-smooth, namely, the switching manifold, across which the linear dynamics change, and the threshold, at which the reset conditions are applied. At each of these manifolds, we can observe non-smooth bifurcations. Many non-smooth bifurcations have smooth analogues, however, the qualitative behaviour of the attractors borne out of these bifurcations may be markedly different. For some bifurcations, such as the grazing bifurcations, there are no smooth analogues,

and so these give rise to behaviour that cannot be captured by a smooth system. The classification of generic non-smooth bifurcations, and behaviours arising from them, is an open area of research [122]. After finding the non-smooth bifurcations in the PWL-IF model, we demonstrated the failure of the standard method of calculating Lyapunov exponents in non-smooth systems, and calculated an explicit formula to give the true exponents, using techniques originally developed for impacting mechanical oscillators [149]. Using the Lyapunov exponents, we were able to identify regions of chaotic activity near spike adding bifurcations and burst death. Typically, planar systems do not support chaotic behaviour, but the non-smooth mechanisms in IF models can make systems unpredictable. A perturbative analysis around threshold, and the switching manifold may give more insight into the way chaotic behaviour emerges in the PWL-IF model. A more general method for calculating Lyapunov exponents in IF models was recently presented in [218], in which the authors go on to correct results from earlier studies [219, 220] pertaining to networks of IF neurons. That such a paper has recently been published, and has identified errors in quantification of exponents, highlights the need within the computational neuroscience community to better understand the effect that non-smoothness can have on systems.

We have studied the nature of the non-smooth bifurcations at only the single neuron level, but it would be interesting to observe how these are manifest in a network level description. A good place to start to answer these questions is through the study of Arnol'd tongues, which delimit regions of entrainment of oscillations to a given signal, and thus classify the oscillators response. These tongues will then give us some idea about the ability of the network to synchronise, for example, by treating the signal as input from other cells in the network. A study of Arnol'd tongues for a PWL caricature of the adaptive exponential integrate-and-fire model has recently been performed in [151]. It is possible that the study of PWL models will allow for the construction of explicit solutions for the Arnol'd tongues of the PWL-IF, which will aid understanding of entrainment of this model. Of course, the construction of Arnol'd tongues is still principally a single neuron

study, and full network studies will still need to be performed.

Another strength of IF type models is the ease in which plasticity can be introduced to the system, by including dynamic synaptic variables [221], and it is clear that the PWL-IF model could easily be modified for this purpose. The flexibility of planar IF models in capturing a wide range of dynamical behaviour make them a good candidate for studies of heterogeneous networks. Now that we have identified regions of different behaviour, in terms of bursting with distinct spike numbers, alternate kinds of tonic spiking and quiescence, we can analyse networks composed of neurons existing in all of these regimes, extending our analysis of homogenous networks in [98]. One of the key goals of these kinds of analyses is to predict the network behaviour through dimension reduction, which we hope is achievable by the use of a polynomial chaos expansion [222]. As the degree of heterogeneity in the network increases, such methods are expected to fail, and so we must develop new techniques to predict the network level behaviour, which we have observed to support the within-burst synchrony changes as discussed in [192].

There is a growing interest in studying stochastic IF models, both in terms of the analysis of the mathematical models [223, 224], and in their application to real systems in the brain [1, 225]. Stochastic inputs, from a modelling perspective, are better able to describe fluctuations in input currents to a given neuron, as well as being able to capture the effect of background electrical noise. Since, for linear systems, we may be able to write down analytical solutions to the stochastic problem, we are in a good position to study the effect that stochastic perturbations have on entrainment, firing frequency and even qualitative behaviour. In particular, since we have categorised the important dynamical structures, we can study how these are shifted by the stochastic perturbations and the behaviour that will result. Following this, we may build network models in which communication between cells is noisy and study the overall effect of the perturbation compared with the deterministic case. By amalgamating analyses for heterogeneous and stochastic networks, with the potential of including plasticity, we hope to construct ever better models of neural activity which are mathematically tractable.

Moving away from nonsmooth modelling, in Chapt. 4, we conducted a numerical case study of the ML model to highlight the limitations of phase-only models. In particular, we examined the response of the ML model near a homoclinic bifurcation and near a Hopf bifurcation (in parameter space) using parameter values from [15]. Clearly, the power of phase models lies in the reduction in dimension that they offer. Away from bifurcations and for weak forcing, phase reduction methods generally do well at approximating the behaviour of the full network and are used in a wide variety of applications. Moreover, experimental techniques can be used to find objects such as the PRC, and thus there are direct parallels between theory experiment in this paradigm [6]. However, as we showed, there are a number of scenarios in which the phase-only model breaks down. Alarming, the phase description in this case offers no evidence that a breakdown has occurred, meaning that it is difficult to predict at what point perturbations can no longer be classed as weak. This problem is exacerbated at the network level; even though the individual inputs received by a neuron from any other cell may be small, their summative effect over the entire network, and over time, may not be. We have already seen, for the PWL-IF model, how small changes in parameters can cause drastic changes to the resulting dynamics, even as far as driving cells into a chaotic regime. In terms of neural processing, this indicates that the neuron is unreliable, since its response to the same stimulus will differ over trials [13], and this lack of reliability will have wider impacts on the ability of a neural network to encode stimuli [198].

We examined three scenarios in which the phase description could potentially break down, namely, the case in which trajectories leave the basin of attraction of the limit cycle, the effect of other invariant structures in phase space, and dynamical shear. For the case of dynamical shear, we pinpointed an example of the shear induced folding, and resulting chaos as predicted in [153, 71]. Although this kind of behaviour is difficult to predict from studying the equations alone, we provided a simple diagnostic check that can help to identify which systems may be susceptible to these shear forces. Further to this, we examined the ‘trapping’

effect of invariant structures and the consequences that this can have on a single oscillator. In general, phase models systematically under-predict the tails of the distribution for the ISI in these instances as they do not ‘feel’ structures at a small but finite distance away from the limit cycle. In all of the cases studied as part of this chapter, the magnitude of the forcing was small, further highlighting that the notion of ‘weak’ in the term ‘weakly coupled oscillators’, can be misleading, at least in the context of computational neuroscience.

The aforementioned effects are still present at the network level. For a two-cell network, we showed that shear effects still give rise to chaotic dynamics, and that the trapping effect of structures in phase space can lead to erroneous predictions from the phase interaction function. Clearly, studies of larger networks are required to see if the same phenomena are observed at this level. In some cases, the size of the network may actually work to overcome the issues discussed in this chapter, as the summative effects of all inputs helps to ‘regularise’ the network activity. In other cases, the opposite may be true; in the worst case the network dynamics may be completely irregular, and even outside of this possibility, the predictions from phase models may be incorrect. The key message from this body of work is to take care when using phase reductions, as subtle effects can result in markedly different behaviour of the full system compared with the reduced one.

Following the discussion of the limitations of perturbative techniques for neural oscillators in the preceding chapter, in Chapt. 5, an alternate coordinate system is presented, following ideas in [12]. Although the underlying coordinates are not isochronal, and so the phase variable is not equivalent to the asymptotic phase, this coordinate system offers a practical way to study systems using similar ideas to phase models, without falling into the pitfalls highlighted in the case study of the ML model. One of the biggest drawbacks to using isochronal coordinates is that the asymptotic phase away from the limit cycle, apart from in rather restrictive models, cannot be found analytically. For truly weakly coupled networks, this is not an issue, since only dynamics in the neighbourhood of the limit cycle are considered. If we wish to study networks away from this limit, we need some

systematic way of ascribing phase off-cycle. Of course, we can calculate a large number of isochrons and use these to approximate dynamics off-cycle. To this end, there have been a number of recent developments in the computation of isochrons through: the identification of transverse vector field for which the flow traces out the isochrons [46], numerical continuation using the distance from cycle as a control parameter [48] and most recently, with operator theoretic approaches [49].

Despite the advances in efficiency of computation of isochrons, the three methods presented in the previous paragraph all suffer from three problems. Firstly, even the quickest of the three methods still takes a long time to generate the number of isochrons necessary to approximate the dynamics of the full system. Secondly, given that the phase space cannot be foliated fully using these algorithms, the resulting dynamical system would still only be an approximation. Finally, the generalisation to higher dimensions (above planar), although possible, vastly increases both the time needed to compute isochrons and the inaccuracy associated with the resulting dynamical system.

Instead of using isochronal coordinates, we use a moving orthonormal coordinate system, in which one coordinate vector is the normalised tangent vector along the limit cycle, and the other coordinate vectors are set up orthonormally to it. If the system in Cartesian coordinates is in \mathbb{R}^n , the system following the application of the transformation has one phase variable and $n - 1$ amplitude variables, which capture the Euclidean distance along the coordinate vectors. This coordinate system is exact; no approximation is made during the transformation and so the resulting dynamical system mimics the original system precisely. Extension of the coordinate system to systems with higher dimension is trivial, and does not drastically increase computation time. Since all the functions involved for the system are given explicitly, in terms of either the original vector field or the limit cycle, the transformation can readily be applied to any system with a stable limit cycle. One of the drawbacks of the phase models is that it is impossible to identify if and when trajectories leave the basin of attraction of our reference limit cycle. Under the moving orthonormal system, we can represent all of the relevant

invariant structures in phase space, and even represent trajectories that do leave the basin of attraction of the limit cycle.

Although the coordinate system does not capture information about the true asymptotic phase of a trajectory, if the forcing is not too strong, the dynamics will behave similarly to an isochronal system. With the inclusion of amplitude variables, we are now able to capture the shear effects described in Chapt. 4 due to the nonlinear effects of the saddle in the ML system. Through a comparison of the ML and FHN model, both by an examination of the functional forms associated with the coordinate transformation, and by probing the dynamics of a system forced periodically, we showed scenarios in which the phase model is sufficiently good to account for off-cycle dynamics, and scenarios in which the effect of other phase space structures gives rise to chaotic dynamics. The latter of these studies was centred around the examination of the resulting response functions for our phase and amplitude variables. In this study, it was found that where trajectories are not pushed far from cycle, the resulting dynamics are not chaotic, as shear forces rely on trajectories occasionally being taken far from the limit cycle. However, the shape of these response functions is not sufficient to fully characterise the neuron's response to forcing, as we also need to know about the distribution of points where perturbations are applied. A more thorough analysis, using specified response functions, may give a greater insight into which types of model are likeliest to support chaotic behaviour.

The ideas and techniques for the response of a single cell could be applied to networks of interacting oscillators. For weakly coupled oscillators, one can make predictions about the synchronisation properties of the network, through the use of phase interaction functions. After averaging, the interactions between cells in the network can be expressed as a function of the difference between their respective phases alone. An analogue of this, constructing a similar object for a network retaining the information about distance from cycle, is clearly a desirable goal. Even in the weak setting, it allows for a wider range of behaviour compared to the phase-only model. In particular, it will also open avenues to study the synchronisa-

tion of heterogeneous networks, for which the intrinsic dynamics may play more of a role for some network units. Since the transformation is exact, it may also provide some clues for predicting network behaviour away from the weak limit, which is one of the key goals of network studies of coupled oscillators.

One of the more subtle aspects of this study was in the appropriate way to deal with the pulsatile forcing. Typically, these kinds of forcing are modelled through δ -functions, mediated by some state dependence on one or more of the intrinsic variables. We adopt the approach in [226, 185] to correctly integrate the effect of this δ -function. In the second of these references, the authors show that other studies, based on [227], have failed to correctly address the discontinuity associated with the δ -function, leading to markedly different network behaviour compared with the true behaviour. Owing to the deceptive nature of this type of forcing we echo our earlier statement that the computational neuroscience community should be aware of the role that non-smooth dynamics can play.

Building on the ideas presented in the preceding two chapters, Chapt. 6 replaces the deterministic perturbations with their stochastic counterparts. Biological processes are inherently stochastic, and mathematical descriptions can be used to capture this nature. In fact, the HH formalism is based around a continuum approximation of Markovian dynamics [16]. There have been a number of recent developments in the study of stochastically forced oscillators, both from the mathematical and experimental standpoint [193, 8, 7, 52, 53]. The consideration of noise can, at a base level, help mathematical models to fit fluctuations in real data. Further to this, analysing sources and consequences of noise may give a better understanding of how key cellular mechanisms operate, as well as giving insight for efficient neural computation in a noisy environment.

In the studies we performed, we firstly highlighted the need for consistency in the interpretation of stochastic integrals in mathematical models. That different choices of interpretation lead to distinct dynamics is not unknown, however, our studies indicate how drastic the consequences can be. For the SL model, we observed a pronounced difference in the steady state distribution of the system

under the Itô interpretation compared with the Stratonovich one. Furthermore, we saw dynamically opposite behaviour of the mean frequency of the oscillations. This will clearly have an impact on the synchronisation properties at the network level. Progress has been made in recent years in correctly performing phase reductions of stochastically perturbed oscillators [193, 7, 9], but these studies all use integrals of the Stratonovich type. As outlined in Chapt. 2, there are a number of potential sources of noise in the nervous system, and there is the ever present measurement noise to be considered when dealing with real data. Due to the large number of such sources, some of which may more naturally be modelled under Itô type integrals, we feel that the community should take care to observe the differences between the two interpretations.

The types of noise we have considered have essentially modelled electrical noise, by simply adding additive white noise to one of the state variables. Of course, there are many other ways of introducing stochastic perturbations into mathematical models, which will likely give rise to differing dynamics, particularly in lieu of the discussion in the last paragraph. In fact, since the base processes underlying biological rhythms are themselves stochastic in nature, a more natural way to model noise might be to include it in the intrinsic dynamics, rather than treating it as acting perturbatively [228, 229]. This leads to a more general form of random dynamical system, but one that may better be able to describe the nature of the real biological system [230, 231].

Following the establishment of a phase-amplitude coordinate system in Chapt. 5, we presented an alternative system in Chapt. 6. In the alternative system, the amplitude variable does not represent the Euclidean distance from the limit cycle, but is instead proportional to the norm of the tangent vector around the limit cycle. Although we thus lose the notion of Euclidean distance, the resulting coordinate system does not suffer from the same issues regarding coordinate breakdown as the system in Chapt. 5. Furthermore, the coordinate system can compactly be written in terms of the vector field and the Fourier coefficients of the asymptotically stable limit cycle. We showed how to calculate the dynamical equa-

tions for the phase and amplitude variables for the SL model, and went on to show how to simulate large numbers of sample paths in these coordinates. Through the use of these numerical routines, and by solving the forward Fokker–Planck equation, we compared the steady state distribution for infinitesimally weak noise to those for increasing noise strengths. In order to do this, we derived a finite difference scheme, using upwind methods, to address the numerical instabilities associated with convection-diffusion problems in which convection dominates diffusion [208]. For very small noise, the steady state distributions match those for the infinitesimal case, but one does not have to increase the noise strength by too great an amount to observe differences in these distributions. In the case of the SL model, we find a periodic modulation of the steady state distribution as the noise strength increases, compared with the flat distribution predicted for infinitesimally weak noise. In both cases, histograms obtained using our numerical routines support the solution of the Fokker–Planck equation approximated using the finite difference method.

We may go on to extend the use of the coordinate system to the analysis of networks. By solving the Fokker–Planck equation for a network state, we will be able to say something about how much the coupling between units affects the underlying oscillations, and also predict where other phase space structures are likely to play a role in the emergent behaviour [80, 200]. By examining the marginal distribution in phase, we will then be able to make predictions about the synchronisation of the network, and also predict where coupling between units will have the greatest effect on network rhythms [232].

It is clear that the development of easy-to-use toolboxes for solving the Fokker–Planck equation, as well as PDEs in general, using both finite element and finite difference methods will open up the avenue to probe networks of coupled oscillators. For solving the Fokker–Planck equation by the finite element method, we used the FEniCS toolbox, freely available to download [207], and well documented [206] with an active user community. Presently, these methods are only really appropriate for solving problems in low dimensional space [233]. This is unsur-

prising, since they are typically used to solve problems in physics, engineering and theoretical mechanics, in which the dimensionality of the system is constrained by real world geometry. However, such toolboxes are seeing more use in answering questions about biology, for example in the study of impulse propagation in cardiac muscles cells [234]. In particular, meshing routines are becoming ever better at dealing with complex geometry, such as that associated with dendritic trees, so that models can be simulated on meshes that capture biophysical structure.

As toolboxes continue to improve, we will then be able to use the Fokker–Planck equation to make statements about network rhythms without ever having to simulate sample paths. Using the same finite element methods, we concluded the chapter with an analysis of a model, previously studied numerically in [210]. Although the model is purely phenomenological, it has been shown to match statistics with cortical EEG data, and so this model may act as a canonical model for cortical rhythms. We have expanded the analysis in [210] to include analytical results, and to study more general forms of forcing, highlighting how the expected dwell times in each of the stable modes is affected. Our study of the bistable subcritical Hopf showed how solutions to the Fokker–Planck equation can identify the way in which structures in state space affect behaviour under stochastic perturbations, as well as the range of quantities that the Fokker–Planck equation can be used to find, such as exit time distributions and exit probabilities through a firing threshold. The same kind of analyses can be used to study other systems supporting interesting behaviour, such as those with greater degrees or different types of multistability, for example, those giving rise to mixed mode oscillations [195, 235]. They may also be useful in answering questions about stochastic and coherence resonance, or how noise affects transitions between silent and active phases of bursting neurons [213], and finally how phenomena such as array induced coherence resonance occur [236, 237], and their effect on information processing in the brain [238].

8 Appendices

A Numerical integration of the SDEs (4.17) and (4.18) for Chapt. 4

To solve (4.17) and (4.18), we use the stochastic Runge-Kutta scheme as described in [89]. For a general system, of the form

$$dX = f(X)\Delta t + \sigma g(X)dW_t, \quad (\text{A.1})$$

this scheme first divides the integration period into N intervals of size h , whereupon the approximation to the true solution is given by y_n , $n = 1, 2, \dots, N$,

$$y_{n+1} = y_n + \sum_{i=1}^s a_i f(Y_i)h + \sigma \sum_{i=1}^s (b_i^1 \Delta W_n + b_i^2 J_{10}) g(Y_i), \quad (\text{A.2})$$

$$Y_i = y_n + \sum_{j=1}^{i-1} \alpha_{ij} f(Y_j)h + \sigma \sum_{j=1}^{i-1} (\beta_{ij}^1 \Delta W_n + \beta_{ij}^2 J_{10}) g(Y_j), \quad i = 1, \dots, s, \quad (\text{A.3})$$

where $\Delta W_n = h^{1/2}\xi_0$, $J_{10} = 1/2 h^{3/2}(\xi_1 + \xi_2/\sqrt{3})$, $\xi_{0,1,2} \sim N(0, 1)$ and α, β, a and b are constant matrices and vectors which are the solution to a series of linear systems corresponding to a rooted tree expansion of (A.1) as a Taylor series. These expansions are considered in more detail in [239, 87, 89]. One solution to these

linear systems gives the coefficients as

$$\begin{aligned}
\alpha &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad a = \left(\frac{1}{6}, \frac{1}{3}, \frac{1}{3}, \frac{1}{6} \right)^T, \\
\beta^1 &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ -0.7242916356 & 0 & 0 & 0 \\ 0.4237353406 & -0.1994437050 & 0 & 0 \\ -1.578475506 & 0.840100343 & 1.738375163 & 0 \end{bmatrix}, \\
\beta^2 &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 2.702000410 & 0 & 0 & 0 \\ 1.757261649 & 0 & 0 & 0 \\ -2.918524118 & 0 & 0 & 0 \end{bmatrix}, \\
b^1 &= \left(-0.7800788474, 0.07353768240, 1.48520013, 0.2199211524 \right)^T, \\
b^2 &= \left(1.693950844, 1.636107882, -3.024009558, -0.3060491602 \right)^T. \quad (\text{A.4})
\end{aligned}$$

The scheme requires four function evaluations per step, but provides an accuracy of $O(h^2)$ in the strong sense. We use this routines as, although explicit, it gives a good approximation to the true solution, without the added cost of an implicit solver.

To improve the accuracy of our approximation of the spike times we use linear interpolation, as in Sect. 3.8, upon crossing threshold, so that the threshold crossing time t^* is given by

$$t^* = t_{n-1} + h \frac{12.5 - v_{n-1}}{v_n - v_{n-1}}. \quad (\text{A.5})$$

We note that although this will improve the accuracy of our firing estimate, we have no notion of what the stochastic process is doing between t_{n-1} and t_n , so that there will always be some error associated with this step.

B Numerical integration of the stochastic phase-amplitude equations as used in Chapt. 6

We seek an efficient algorithm to integrate (6.53). The first step, given $y = (\theta, \rho)$ is to evaluate x , so that $F(x)$ may be evaluated. The time consuming part of this transformation is the matrix exponential, which needs to be done N times per step. The naïve approach to carrying out the transformation would be to evaluate the matrix exponential for each $n = 1, \dots, N$ and then take sum, following how it is written in (6.37). However, this is cumbersome, particularly when wishing to run many simulations concurrently. It is also worth noting that, aside from the matrix exponential, the other terms in the product (6.37) allow for the simultaneous transformation of a number of values of y . A faster approach would be to perform a large matrix calculation, replacing the sum in (6.37). Such matrix approaches are used widely for other uses, particularly in signal processing.

First, we write the Fourier components of the orbit as a vector, of the form

$$\Gamma_n = (\Gamma_{1,1}, \Gamma_{2,1}, \Gamma_{1,2}, \Gamma_{2,2}, \dots, \Gamma_{1,N}, \Gamma_{2,N})^T. \quad (\text{B.1})$$

We write the Fourier modes in the order $0, 1, 2, \dots, N/2, -(N/2 - 1), \dots, -1$, and construct a $1 \times N$ vector m with the modes as entries. Now, x may be written as the matrix product given as

$$x = \mathcal{L} \mathcal{K} \Gamma_n, \quad (\text{B.2})$$

where

$$\mathcal{L} = e^{2\pi i m \theta / \Delta} \otimes \mathbf{1}_2, \quad (\text{B.3})$$

where $\mathbf{1}_2$ is the 2×2 matrix whose entries are all ones, and \mathcal{K} is a block matrix whose entries are the 2×2 matrix exponentials $e^{2\pi i \rho J / \Delta}$. We now need to find an efficient way to compute \mathcal{K} . If we are only interested in doing one run at a time, we can easily form this matrix. For M concurrent runs, where M is large, the memory limitations in MATLAB mean that we will need to alter the approach slightly. In essence, the method works by noting that the matrix exponential e^A ,

for a generic matrix A , may be written in the form [240]

$$e^A = V \operatorname{diag}(\exp(\Lambda)) V^{-1}, \quad (\text{B.4})$$

where Λ is a vector of the eigenvalues of A and V is a matrix whose columns are the eigenvectors associated with the eigenvalues in Λ . The function $\operatorname{diag}(x)$ takes a vector x and forms a square matrix whose leading diagonal is x with zeros everywhere else. We note that this is different to the approach used by MATLAB, which uses Padé approximations to compute the matrix exponential [240]. In fact, the method using Padé approximations may be quicker for one-off calculations, but we will store the values of V and Λ for use multiple times, producing a speed up factor. We also note that it is important that Λ in the above is a vector. It is common, for example, when diagonalising a matrix, to find a diagonal matrix whose leading diagonal is composed of the eigenvalues of said matrix. To evaluate the matrix exponential, we require a diagonal matrix whose leading diagonal is formed by taking exponentials of the eigenvalues. Writing matrices in such a diagonal form allows for the easy computation of matrix powers. For example, if a matrix A is written as $A = V \operatorname{diag}(\Lambda) V^{-1}$, then A^n is given as $A^n = V \operatorname{diag}(\Lambda)^n V^{-1}$. We may exploit this for our matrix exponential by noting that the terms in the exponent of the matrix exponential, for a given ρ , are just powers of one another. Thus, if we let

$$A = \frac{2\pi i \rho}{\Delta} J, \quad (\text{B.5})$$

and let V and Λ be the eigenvectors and eigenvalues of A , then for each n , we may write the matrix exponential as $e^{nA} = V \operatorname{diag}(\exp(\Lambda))^n V^{-1}$. Furthermore, we also note that the value of ρ also acts as an exponent in the matrix exponential, meaning that we can further simplify matters by letting

$$A = \frac{2\pi i}{\Delta} J, \quad (\text{B.6})$$

The matrix exponential for any ρ and n is then

$$e^{n\rho A} = V \operatorname{diag}(\exp(\Lambda))^{n\rho} V^{-1}. \quad (\text{B.7})$$

In this way, we need only compute V and Λ once, and the remaining calculations are simple matrix multiplications with a few tensor products, which MATLAB is optimised to do. Our results show that, even for large numbers of runs, the norm of differences between matrix exponentials found using this approach and the Padé approximation is $\sim O(10^{-13})$, which is smaller than the error we expect from numerical routines to integrate SDEs of the form (6.1).

Due to the memory handling in MATLAB, the optimal strategy now differs, depending on whether we are doing one run or many at the same time. In this case, many means greater than 50. For one run, we may evaluate the matrix \mathcal{K} easily by first evaluating it as a large block diagonal matrix, and then extracting the relevant entries needed to form \mathcal{K} . This method starts by taking the product

$$\mathbf{V} = \mathbb{I}_N \otimes V, \quad (\text{B.8})$$

where \mathbb{I}_N is the identity matrix in \mathbb{R}^N . This creates a block diagonal matrix with copies of V down the leading block diagonal. In a similar fashion, a $2N \times 1$ vector is formed by taking the product

$$\mathbf{D} = [1, 1]^T \otimes \Lambda. \quad (\text{B.9})$$

We note that these computations are done only once, at the beginning, and then stored for later use. At each time step during the numerical integration, a similar computation is carried out, given a $1 \times 2N$ vector containing copies of the Fourier coefficient index and ρ , giving

$$M = [1, 1] \otimes m\rho. \quad (\text{B.10})$$

We note that this is most efficiently achieved in MATLAB by first computing $m\rho$, then introducing another $2 \times N$ matrix

$$L = \begin{pmatrix} m\rho \\ m\rho \end{pmatrix}, \quad (\text{B.11})$$

We then form a $2N \times 1$ vector M by iteratively stacking columns of L . Now, we calculate the block diagonal matrix \mathcal{Z} as

$$\mathcal{Z} = \mathbf{V} \text{diag}((\mathbf{D}^M)) \mathbf{V}^{-1}, \quad (\text{B.12})$$

where the power acts element-wise on \mathbf{D} . The entries of the $1 \times N$ block vector of 2×2 matrices \mathcal{K} are on the leading diagonal of the block diagonal matrix Z . To extract them in MATLAB, we can use logical indexing by only outputting entries of the $2N \times 2N$ matrix \mathcal{Z} corresponding to

$$\mathbf{I} = \mathbb{I}_N \otimes \mathbf{1}_2. \quad (\text{B.13})$$

This will give the requested entries of \mathcal{Z} as a column vector $z = \mathcal{Z}(\mathbf{I})$. We then form two other indexing vectors

$$\text{id}_1 = \mathbf{1}_{N \times 1} \otimes [1, 0, 1, 0]^T, \quad \text{id}_2 = \mathbf{1}_{N \times 1} \otimes [0, 1, 0, 1]^T, \quad (\text{B.14})$$

where $\mathbf{1}_{N \times 1}$ is simply a N vector whose entries are all ones. Finally, we may then form the $2 \times 2N$ matrix \mathcal{K} as

$$\mathcal{K} = [z(\text{id}_1), z(\text{id}_2)]^T. \quad (\text{B.15})$$

Although it is possible to use the above method for many concurrent runs, by first taking the outer product of m and $\underline{\rho}$, where $\underline{\rho} = (\rho_1, \dots, \rho_M)$, this calculation quickly becomes too large. This is because the process first requires creating a matrix N times larger than the one we actually need to compute. Because of the way MATLAB handles memory in blocks, this makes the process time consuming. For a large number of concurrent runs, we may still use (B.2), but instead of using the above process to form \mathcal{K} , we iterate over the respective entries of m and $\underline{\rho}$. This does not require the use of any sparse, block diagonal matrices and so the whole process is faster. Since,

$$\frac{\partial x}{\partial \theta} = \left(\frac{2\pi i n}{\Delta} \right) \sum_n e^{2\pi i n \theta / \Delta} e^{2\pi i n \rho J / \Delta} u_n, \quad (\text{B.16})$$

It is clear to see that we can use exactly the same process explained above to find $\partial x / \partial \theta$. Thus, we have that

$$\frac{\partial x}{\partial \theta} = \tilde{\mathcal{L}} \mathcal{K} \Gamma_n, \quad (\text{B.17})$$

where

$$\tilde{\mathcal{L}} = \left(\left(\frac{2\pi i m}{\Delta} \right) \bullet e^{2\pi i m \theta / \Delta} \right) \otimes \mathbf{1}_2, \quad (\text{B.18})$$

where \bullet denotes element-wise array multiplication. This expression can then be used to find $\partial x/\partial\rho$. Now that we have found both x and $\partial x/\partial\theta$, it is easy to compute the right hand sides of (6.53).

C Steady state distributions for non-weak noise from Chapt. 6

For a planar system in Cartesian coordinates with white-noise forcing only in the first variable, the Langevin equation is given by (6.1) with $G = (1, 0)^T$. The (forward) Fokker–Planck equation for this equation is

$$\frac{\partial P}{\partial t} = -\frac{\partial[F_1P]}{\partial x_1} - \frac{\partial[F_2P]}{\partial x_2} + \frac{\sigma^2}{2} \frac{\partial^2 P}{\partial x_1^2}, \quad (\text{C.1})$$

where the subscript denotes the component of the vector valued function F . In the general case, the domain, S , of x is unbounded, and so we require the boundary conditions that

$$P \rightarrow 0 \text{ as } x_{1,2} \rightarrow \pm\infty. \quad (\text{C.2})$$

If we are solving these equations numerically, our domain must necessarily be finite, in which case we may use reflecting boundary conditions at the edge of our domain, ∂S , given by the condition

$$\hat{n} \cdot \mathcal{J} = 0, \quad (\text{C.3})$$

where \hat{n} is the outward pointing normal on ∂S and \mathcal{J} is the probability flux, given by

$$\mathcal{J}_i = F_i(x)P - \frac{\sigma^2}{2} \sum_j \frac{\partial[\mathbf{G}_{ij}P]}{\partial x_j}, \quad i = 1, 2, \quad (\text{C.4})$$

where $\mathbf{G} = GG^T$. In all cases which we consider, we choose our domain to be rectangular, that is $S = [l_1, l_2] \times [l_3, l_4]$. Since there is no diffusion in the x_2 direction (since $G_2 = 0$), the reflecting boundary conditions (C.4) on $x_2 = l_3$ and $x_2 = l_4$ imply that $P = 0$ along these edges. These are standard Dirichlet boundary conditions. For the edges $x_1 = l_1$ and $x_1 = l_2$, we then impose the condition that

$$\frac{\sigma^2}{2} \frac{\partial P}{\partial x_1} = F_1(x)P, \quad (\text{C.5})$$

which is a mixed boundary condition. Now that we have established the boundary conditions for the problem, we can complete the statement of the problem by imposing the initial condition

$$P(x, 0) = R(x), \quad (\text{C.6})$$

where R is some probability distribution, satisfying the integral condition

$$\int_S R(x) dx = 1. \quad (\text{C.7})$$

The solution of (C.1) for large times will then give the steady state distribution of (6.1). Rather than solving (C.1), subject to an initial condition, over a long time, we may instead solve directly for the steady state distribution. If we set $\partial P/\partial t = 0$ in (C.1), we arrive at

$$0 = -\frac{\partial[F_1 P]}{\partial x_1} - \frac{\partial[F_2 P]}{\partial x_2} + \frac{\sigma^2}{2} \frac{\partial^2 P}{\partial x_1^2}. \quad (\text{C.8})$$

Since the above equation is homogeneous, attempts to solve it directly will yield only a constant solution $P = p$ for some $p \in \mathbb{R}$, which is clearly not the desired solution. We may tackle this problem in one of two ways. One way is to constrain a point in space to equal some arbitrary value when we solve the equation, and then remove this point from the final solution [209]. We will demonstrate how this may be used later in this chapter. The second method is to change (C.8) from a homogeneous one to an inhomogenous one, through the substitution $Q = P + 1$ as used in [232]. Again, we will consider how this works in practice in due course.

C.1 Fokker–Planck equation for systems in the phase-amplitude coordinates from Chapt. 6

If we now consider system (6.40), we must make a few changes to the problem set up. The forward Fokker–Planck equation for the steady state distribution now reads

$$0 = -\frac{\partial[\Omega P]}{\partial \theta} - \frac{\partial[f P]}{\partial \rho} + \frac{\sigma^2}{2} \frac{\partial^2[h^2 P]}{\partial \theta^2} + \sigma^2 \frac{\partial^2[hg P]}{\partial \theta \partial \rho} + \frac{\sigma^2}{2} \frac{\partial^2[g^2 P]}{\partial \rho^2}. \quad (\text{C.9})$$

We solve this equation on the domain $S = [0, \Delta] \times [-R, R]$, where $R > 0$. Since θ is periodic, we use periodic boundary conditions which identify $\theta = 0$ and $\theta = \Delta$. If R is sufficiently large, relative to the dynamics of (C.9), we may simply use the Dirichlet condition $P = 0$ at $\rho = \pm R$, which is easier to implement numerically than the reflecting boundary condition given by (C.4). We now consider two different approaches for the numerical solution of (C.8) and (C.9).

D Solution of the Fokker–Planck equation by finite element methods

In order to use the FEniCS toolbox to solve (C.1), we first need to write our problem in its corresponding weak form. We consider the equation for the steady state distribution in Cartesian coordinates given by (C.8). We first rewrite the equation in an equivalent form, giving rise to the convection-diffusion equation

$$0 = B(x)P + F(x) \cdot \nabla P + \frac{\sigma^2}{2} \nabla \cdot (A \nabla P). \quad (\text{D.1})$$

where

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad B(x) = \sum_{i=1}^2 \frac{\partial F_i}{\partial x_i}. \quad (\text{D.2})$$

Next, we multiply this equation by a test function ν , and integrate over the domain S to give

$$0 = \int_S \nu B(x)P \, dx + \int_S \nu F(x) \cdot \nabla P \, dx + \frac{\sigma^2}{2} \int_S \nu \nabla \cdot (A \nabla P) \, dx. \quad (\text{D.3})$$

In the current context, we refer to P as being a trial function. Equation (D.3) is supposed to hold for all ν in some function space \widehat{V} . The trial function lies in some, possibly different, function space V . To solve the equation numerically, we must transform the continuous variational problem into a discrete one by introducing finite dimensional test and trial spaces that are subsets of the full spaces. For simplicity, we omit the mathematical details of the workings of the finite element method, and instead proceed to complete the problem set up. The second order derivative in (D.3) is too high an order for the finite element method to be applied directly. To cope with this, we integrate the offending term by parts, using Green's identities so that (D.3) now reads

$$0 = \int_S \nu B(x)P \, dx + \int_S \nu F(x) \cdot \nabla P \, dx - D \int_{\partial S} \nu \hat{n} \cdot (A \nabla P) \, ds + D \int_S \nabla \nu \cdot (A \nabla P) \, dx, \quad (\text{D.4})$$

where ds denotes a boundary integral. We require that ν vanishes on parts of the boundary where P is known. In the present problem, that means that the

boundary integral term vanishes when $x_2 = l_3$ and $x_2 = l_4$. Along the boundary $x_1 = l_1$ and $x_1 = l_2$, we impose the reflecting boundary condition (C.4). Denoting ∂S_V as the union of these two edges, we then arrive at the weak form of our problem, that is, find $P \in V$ such that

$$0 = \int_S \nu B(x) P \, dx + \int_S \nu F(x) \cdot \nabla P \, dx - \int_{\partial S_V} \nu f P \, ds + D \int_S \nabla \nu \cdot (A \nabla P) \, dx. \quad (\text{D.5})$$

This equation is homogenous. In order to solve it to obtain a nontrivial solution, we make the substitution $P = \tilde{P} + 1$, to arrive at

$$\begin{aligned} \int_S \nu B(x) \, dx - \int_{\partial S_V} \nu f \, ds &= \int_S \nu B(x) \tilde{P} \, dx + \int_S \nu F(x) \cdot \nabla \tilde{P} \, dx \\ &\quad - \int_{\partial S_V} \nu f \tilde{P} \, ds + D \int_S \nabla \nu \cdot (A \nabla \tilde{P}) \, dx. \end{aligned} \quad (\text{D.6})$$

Note that the above is written in the form

$$L(\nu) = a(\tilde{P}, \nu), \quad (\text{D.7})$$

We refer to L , which contains only linear terms in ν , as the linear form, whilst a is referred to as the bilinear form for similar reasons. The equation can be coded as written in (D.6) in FEniCS. The program then transforms these forms into a large matrix system, which may then be solved using standard techniques for linear algebra, such as Gaussian elimination. Once a solution for \tilde{P} has been found, it may then be transformed back to P , whereupon the normalisation condition $\int_S P \, dx = 1$ may be applied to give the steady state probability distribution.

D.1 Small noise solutions using the finite element method

Whilst the method described in D is well suited for a large number of problems, solutions become oscillatory or blow up if σ is too small. In short, the smallness of the diffusion compared with the convection term causes numerical instabilities. To overcome this problem, we may use streamline upwind Petrov-Galerkin (SUPG) methods to stabilise the numerical scheme [82]. These methods work by adding ‘artificial diffusion’ in the direction of the convection. Provided that the artificial

diffusion added is not too large relative to the convection, this can stabilise the numerical scheme without greatly compromising the accuracy of the solution. One method to add in this diffusion is to replace the test function ν with $\nu + \delta F \cdot \nabla \nu$. The parameter δ defines exactly how much artificial diffusion we put into the problem. The optimal choice for δ is still an open question, but a common choice is

$$\delta = h(2|F(x)|)^{-1}, \quad (\text{D.8})$$

where h is the size of the cells in the discretisation. Our bilinear form (D.6) can easily be updated to include these extra terms by adding to a , the term $\int_S r \delta F \cdot \nabla \nu dx$ where

$$r = B(x)\tilde{P} - B(x) + F(x) \cdot \nabla \tilde{P} - \frac{\sigma^2}{2} \nabla \cdot (A \nabla \tilde{P}), \quad (\text{D.9})$$

and solving as before [206]. When using the finite element methods, we shall always use standard Lagrangian elements of degree one, which are simply triangles. In the absence of any refinement to our initial discretisation of the domain, all of these triangles will be congruent, so that h is identical across elements.

E Solution of the Fokker–Planck equation using the finite difference method

We shall now demonstrate the use of the finite difference method in solving (C.9). In general, the discretisation of the domain does not have to be regular, but for simplicity, we shall choose our discretisation this way. Thus, we split both the intervals $[0, \Delta]$ and $[-R, R]$ into N subintervals, and define our spatial steps to be

$$\delta\theta = \frac{\Delta}{N}, \quad \delta\rho = \frac{2R}{N}. \quad (\text{E.1})$$

Now, we need to discretise the equation (C.9). We note that we could use the product rule to expose terms which are derivatives of P alone, for example

$$\frac{\partial[\Omega P]}{\partial\theta} = \frac{\partial\Omega}{\partial\theta}P + \Omega\frac{\partial P}{\partial\theta}, \quad (\text{E.2})$$

so that we have the derivative of Ω exactly, rather than taking an approximation. Intuitively, one would expect this to be more accurate than the approximation. However, it was shown in [241] that this leads to inaccuracies, meaning that leaving the partial derivatives in the form as in (C.9) is preferable. For all derivatives, we use centred difference approximations. For example, for an arbitrary function $u(x)$ of one variable, these derivatives would be given by

$$\frac{d[uP]}{dx} = \frac{u_{i+1}P_{i+1} - u_{i-1}P_{i-1}}{2\delta x}, \quad \frac{d^2[uP]}{dx^2} = \frac{u_{i+1}P_{i+1} - 2u_iP_i + u_{i-1}P_{i-1}}{\delta x^2}, \quad (\text{E.3})$$

where the subscript denotes the point index, $u_i = u(x_i)$, and δx is the stepsize. On a 2D lattice, we must account for drift and diffusion in both the θ and ρ directions. We order cells $1, \dots, N^2$, so that points with index $i = 1, \dots, N$ correspond to the edge $\theta = 0$, points with index $i = N \times (N - 1) + 1, \dots, N \times N$ correspond to $\theta = \Delta$. This means that the diffusion in θ will have a similar form to (E.3), but with $i + N$ and $i - N$ replacing $i + 1$ and $i - 1$ respectively. For the mixed derivative, we again use a centred difference approximation, so that,

$$\frac{\partial^2[hgP]}{\partial\theta\partial\rho} = \frac{hg_{i+N+1}P_{i+N+1} - hg_{i+N-1}P_{i+N-1} - hg_{i-N+1}P_{i-N+1} + hg_{i-N-1}P_{i-N-1}}{4\delta\theta\delta\rho}. \quad (\text{E.4})$$

Putting together all of the centred difference approximations, we arrive at the following numerical scheme

$$\begin{aligned}
0 = & -\frac{\Omega_{i+N}P_{i+N} - \Omega_{i-N}P_{i-N}}{2\delta\theta} - \frac{f_{i+1}P_{i+1} - f_{i-1}P_{i-1}}{2\delta\rho} \\
& + \frac{\sigma^2}{2} \frac{h_{i+1}^2 P_{i+1} - 2h_i^2 P_i + h_{i-1}^2 P_{i-1}}{\delta\theta^2} \\
& + \sigma^2 \frac{hg_{i+N+1}P_{i+N+1} - hg_{i+N-1}P_{i+N-1} - hg_{i-N+1}P_{i-N+1} + hg_{i-N-1}P_{i-N-1}}{4\delta\theta\delta\rho} \\
& + \frac{\sigma^2}{2} \frac{g_{i+1}^2 P_{i+1} - 2g_i^2 P_i + g_{i-1}^2 P_{i-1}}{\delta\rho^2}.
\end{aligned} \tag{E.5}$$

This equation will hold at all of the interior points of S , that is, away from the boundaries. Along the boundaries, we must impose the correct boundary conditions. We identify $\theta = 0$ and $\theta = \Delta$, so we need to impose periodic boundary conditions along the top and bottom edges. Without loss of generality we consider the adjustments needed along the boundary with $\theta = 0$, though the adjustments are similar for the boundary with $\theta = \Delta$. Clearly, the points $i - N$ do not exist for $i = 1, \dots, N$. However, for periodic boundary conditions, these correspond to the points along with cell number $i = N \times (N - 1) + 1, \dots, N \times N$. Thus, we can replace (E.5) with

$$\begin{aligned}
0 = & -\frac{\Omega_{i+N}P_{i+N} - \Omega_{i+M}P_{i+M}}{2\delta\theta} - \frac{f_{i+1}P_{i+1} - f_{i-1}P_{i-1}}{2\delta\rho} \\
& + \frac{\sigma^2}{2} \frac{h_{i+1}^2 P_{i+1} - 2h_i^2 P_i + h_{i-1}^2 P_{i-1}}{\delta\theta^2} \\
& + \sigma^2 \frac{hg_{i+N+1}P_{i+N+1} - hg_{i+N-1}P_{i+N-1} - hg_{i+M+1}P_{i+M+1} + hg_{i+M-1}P_{i+M-1}}{4\delta\theta\delta\rho} \\
& + \frac{\sigma^2}{2} \frac{g_{i+1}^2 P_{i+1} - 2g_i^2 P_i + g_{i-1}^2 P_{i-1}}{\delta\rho^2}, \quad i = 2, \dots, N,
\end{aligned} \tag{E.6}$$

where $M = N \times (N - 1)$. The equations for the boundary with $\theta = \Delta$ are similar. For simplicity, we use the Dirichlet boundary condition $P = 0$ at points along the boundaries corresponding to $\rho = \pm R$. To impose this condition, we simply set the contribution to the derivative approximations of any boundary points to zero. Note that this also deals with the ‘corners’ of the boundary. After the imposition

of boundary conditions, we now have a system of N^2 linear equations for the unknown P_i . These equations may be then recast in the form

$$KP = 0, \tag{E.7}$$

where K is a sparse, banded diagonal matrix, of the coefficients of (E.5). In a similar fashion to the homogeneous problem discussed in Appx. D, the solution to (E.7) will be some constant. We can circumvent this problem by constraining one of the points in our domain [209]. To do this, we choose a point i_0 and use $P_{i_0} = 1$ as an additional boundary condition. We thus form the inhomogeneous problem

$$K_r P_r = -b, \tag{E.8}$$

where K_r is the matrix K with the i_0 'th column removed, P_r is the vector P with the i_0 'th element removed and b is the i_0 'th column of K . This equation can now be solved to give a nontrivial solution, which may then be normalised to give the steady state probability distribution. Higher order schemes have been developed to give more accurate solutions to the Fokker–Planck equation, [212], with fewer spatial points, but for simplicity, we will use a second order scheme.

E.1 Small noise problems using the finite difference method

For small noise problems, we have the same issue as discussed in Appx. D.1 in that, since diffusion is dominated by convection, the numerical scheme may give rise to oscillatory solutions, with some negative values, or may blow up. To overcome this problem, we can use upwind schemes as before, to add artificial diffusion in the direction of the convective flow. For the finite difference scheme, this amounts to replacing the centred difference with a one-sided one. For example, the forward spatial derivative in one dimension, at a point x_i , is given by

$$\frac{d[uP]}{dx} = \frac{u_i P_i - u_{i-1} P_{i-1}}{\delta x}, \tag{E.9}$$

whilst the backward spatial derivative is given by

$$\frac{d[uP]}{dx} = \frac{u_{i+1} P_{i+1} - u_i P_i}{\delta x}, \tag{E.10}$$

Note that, for both of the above discretisation schemes, information only comes to x_i from one direction. Clearly, this is not as accurate as the centred approximation, since the error for both the forward and backward approximations is $O(\delta x)$ whilst it is $O(\delta x^2)$ for the centred version, but it has the advantage of stabilising the numerical scheme. The choice of which approximation to use depends on which way information ‘flows’ at the point x_i . If $u > 0$, we expect information to come from the left, so we use the forward derivative. Conversely, if $u < 0$, we expect information to come from the right, so we use the backward derivative. We omit the onerous details of constructing the equations using the upwind scheme here, but it suffices to say that each of the centred difference terms in (E.5) may be replaced with a one-sided alternative, including the mixed-derivative term, which reduces the accuracy of the system by an order of magnitude, but provides more numerical stability. In [241], the authors derive a weighted difference scheme for a 1D system, which guarantees positivity of solutions, avoiding the issues with oscillatory solutions. An extension to the 2D case would not be particularly difficult, though we opt to stick with the simple upwind method for our calculations. Boundary conditions are handled in a similar way as (E.6), taking care to account for which one-sided approximation is used at the point x_i . Solutions to the system are found in exactly the same way as described in Appx. E.

F FEniCS expressions for the Fokker-Planck equation for the phase-amplitude coordinates in Chapt. 6

The equation to be solved in FEniCS is

$$\frac{\partial P_\tau}{\partial t} = C(\phi, r)P_\tau + B(\phi, r) \cdot \nabla P_\tau + \nabla \cdot (A(\phi, r)\nabla P_\tau). \quad (\text{F.1})$$

Writing

$$B = (B_1, B_2), \quad A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \quad (\text{F.2})$$

Expanding the right hand side of this equation we get

$$\begin{aligned} \text{RHS} = & \frac{B_1}{r} \frac{\partial P_\tau}{\partial \phi} + B_2 \frac{\partial P_\tau}{\partial r} + \frac{1}{r^2} \frac{\partial A_{11}}{\partial \phi} \frac{\partial P_\tau}{\partial \phi} + \frac{1}{r^2} A_{11} \frac{\partial^2 P_\tau}{\partial \phi^2} + \frac{1}{r} \frac{\partial A_{12}}{\partial \phi} \frac{\partial^2 P_\tau}{\partial \phi \partial r} \\ & + \frac{1}{r} \frac{\partial A_{21}}{\partial r} \frac{\partial P_\tau}{\partial \phi} + \frac{1}{r} A_{21} \frac{\partial^2 P_\tau}{\partial r \partial \phi} + \frac{1}{r} A_{22} \frac{\partial P_\tau}{\partial r} + \frac{\partial A_{22}}{\partial r} \frac{\partial P_\tau}{\partial r} + A_{22} \frac{\partial^2 P_\tau}{\partial r^2}. \end{aligned} \quad (\text{F.3})$$

Now, we may expand the forward Fokker–Planck equation as

$$\begin{aligned} \frac{\partial P_\tau}{\partial t} = & -\frac{\partial \Omega}{\partial \phi} P_\tau - \Omega \frac{\partial P_\tau}{\partial \phi} - \frac{\partial f}{\partial r} P_\tau - f \frac{\partial P_\tau}{\partial r} \\ & + \frac{\sigma^2}{2} \left(\frac{\partial^2 h^2}{\partial \phi^2} P_\tau + 2 \frac{\partial h^2}{\partial \phi} \frac{\partial P_\tau}{\partial \phi} + h^2 \frac{\partial^2 P_\tau}{\partial \phi^2} + \frac{\partial^2 hg}{\partial \phi \partial r} P_\tau + 2 \frac{\partial hg}{\partial \phi} \frac{\partial P_\tau}{\partial r} + hg \frac{\partial^2 P_\tau}{\partial \phi \partial r} \right. \\ & \left. + \frac{\partial^2 hg}{\partial r \partial \phi} P_\tau + 2 \frac{\partial hg}{\partial r} \frac{\partial P_\tau}{\partial \phi} + hg \frac{\partial^2 P_\tau}{\partial r \partial \phi} + \frac{\partial^2 g^2}{\partial r^2} P_\tau + 2 \frac{\partial g^2}{\partial r} \frac{\partial P_\tau}{\partial r} + g^2 \frac{\partial^2 P_\tau}{\partial r^2} \right), \end{aligned} \quad (\text{F.4})$$

whilst the backward Fokker–Planck equation is simply

$$\frac{\partial P_\tau}{\partial t} = \Omega \frac{\partial P_\tau}{\partial \phi} + f \frac{\partial P_\tau}{\partial r} + \frac{\sigma^2}{2} \left(h^2 \frac{\partial^2 P_\tau}{\partial \phi^2} + hg \frac{\partial^2 P_\tau}{\partial \phi \partial r} + hg \frac{\partial^2 P_\tau}{\partial r \partial \phi} + g^2 \frac{\partial^2 P_\tau}{\partial r^2} \right). \quad (\text{F.5})$$

For both the forward and the backward Fokker–Planck equation, equating second order coefficients leads to

$$A = \begin{pmatrix} h^2 r & hgr \\ hgr & g^2 \end{pmatrix}. \quad (\text{F.6})$$

For the backward Fokker–Planck equation we then find

$$B = \left(\Omega r - \frac{\sigma^2}{2} \frac{\partial h^2}{\partial \phi} r - \frac{\sigma^2}{2} \frac{\partial h g r}{\partial r}, f - \frac{\sigma^2}{2} \frac{\partial h g}{\partial \phi} - \frac{\sigma^2}{2r} g^2 - \frac{\sigma^2}{2} \frac{\partial g^2}{\partial r} \right), \quad (\text{F.7})$$

$$C = 0. \quad (\text{F.8})$$

For the forward Fokker–Planck equation,

$$B = \left(-\Omega r + \sigma^2 r \frac{\partial h^2}{\partial \phi} + \sigma^2 r \frac{\partial h g}{\partial r} - \frac{\sigma^2}{2} - \frac{\sigma^2}{2} \frac{\partial h^2}{\partial \phi} r - \frac{\sigma^2}{2} \frac{\partial h g r}{\partial r} \right), \quad (\text{F.9})$$

$$C = -\frac{\partial \Omega}{\partial \phi} - \frac{\partial f}{\partial r} + \frac{\sigma^2}{2} \left(\frac{\partial^2 h^2}{\partial \phi^2} + \frac{\partial^2 h g}{\partial \phi \partial r} + \frac{\partial^2 h g}{\partial r \partial \phi} + \frac{\partial^2 g^2}{\partial r^2} \right). \quad (\text{F.10})$$

For the system (6.78), we have the functions

$$h = -\frac{1}{r} \sin \phi, \quad g = \cos \phi, \quad \Omega = \phi_c, \quad f = -r^5 + \lambda r^3 + \beta r, \quad (\text{F.11})$$

whence

$$\begin{aligned} \frac{\partial h^2}{\partial \phi} &= \frac{\sin 2\phi}{r^2}, & \frac{\partial h g}{\partial \phi} &= -\frac{\cos 2\phi}{r} \\ \frac{\partial h g}{\partial r} &= \frac{\sin 2\phi}{2r^2}, & \frac{\partial g^2}{\partial r} &= 0, \end{aligned} \quad (\text{F.12})$$

leading to

$$\frac{\partial^2 h^2}{\partial \phi^2} = 2 \frac{\cos 2\phi}{r^2}, \quad \frac{\partial^2 h g}{\partial \phi \partial r} = \frac{\partial^2 h g}{\partial r \partial \phi} = \frac{\cos 2\phi}{r^2}, \quad \frac{\partial g^2}{\partial r^2} = 0. \quad (\text{F.13})$$

Putting all of this together gives rise to the following, for the backward Fokker–Planck equation

$$B = \left(\phi_c r - \frac{\sigma^2}{2r} \sin 2\phi, -r^5 + \lambda r^3 + \beta r - \frac{\sigma^2}{2r} + \frac{\sigma^2}{2r} \cos^2 \phi \right). \quad (\text{F.14})$$

For the forward Fokker–Planck equation, we then have

$$B = \left(-\phi_c r + \frac{\sigma^2}{r} \sin 2\phi, r^5 - \lambda r^3 - \beta r - \frac{\sigma^2}{2r} \cos 2\phi - \frac{\sigma^2}{2r} \cos^2 \phi \right), \quad (\text{F.15})$$

$$C = \frac{2\sigma^2}{r^2} \cos 2\phi + 5r^4 - 3\lambda r^2 - \beta. \quad (\text{F.16})$$

The reflecting boundary condition implies that

$$\hat{n} \cdot (A \nabla P_\tau) = -r^5 + \lambda r^3 + \beta r, \quad (\text{F.17})$$

where $\hat{n} = (0, \pm 1)$.

References

- [1] J. Laudanski, S. Coombes, A. R. Palmer, and C. J. Sumner, “Mode-locked spike trains in responses of ventral cochlear nucleus chopper and onset neurons to periodic stimuli,” *Journal of Neurophysiology*, vol. 103, pp. 1226–37, Mar. 2010.
- [2] L. Badel, S. Lefort, T. K. Berger, C. C. H. Petersen, W. Gerstner, and M. J. E. Richardson, “Extracting nonlinear integrate-and-fire models from experimental data using dynamic I-V curves,” *Biological Cybernetics*, vol. 99, pp. 361–370, 2008.
- [3] M. di Bernardo, C. J. Budd, A. R. Champneys, and P. Kowalczyk, *Piecewise-smooth Dynamical Systems*. Springer, 1st ed., 2008.
- [4] H. Markram, “The Blue Brain Project,” *Nature Reviews Neuroscience*, vol. 7, pp. 153–160, 2006.
- [5] F. C. Hoppensteadt and E. M. Izhikevich, *Weakly connected neural networks*. Springer, 1997.
- [6] N. W. Schultheiss, A. A. Prince, and R. J. Butera, eds., *Phase Response Curves in Neuroscience: Theory, Experiments and Analysis*. Springer, 1st ed., 2012.
- [7] K. Yoshimura and K. Arai, “Phase Reduction of Stochastic Limit Cycle Oscillators,” *Physical Review Letters*, vol. 101, pp. 10–13, Oct. 2008.
- [8] J.-N. Teramae, H. Nakao, and G. B. Ermentrout, “Stochastic Phase Reduction for a General Class of Noisy Limit Cycle Oscillators,” *Physical Review Letters*, vol. 102, pp. 1–4, May 2009.
- [9] D. Goldobin, J.-N. Teramae, H. Nakao, and G. Ermentrout, “Dynamics of Limit-Cycle Oscillators Subject to General Noise,” *Physical Review Letters*, vol. 105, pp. 1–4, Oct. 2010.

- [10] N. Berglund and B. Gentz, *Noise-Induced Phenomena in Slow-Fast Dynamical Systems. A Sample-Paths Approach. (Probability and its Applications)*. Springer, 2006.
- [11] C. R. Laing and G. J. Lord, eds., *Stochastic Methods in Neuroscience*, vol. 1. Oxford University Press, Sept. 2009.
- [12] J. K. Hale, *Ordinary Differential Equations*. John Wiley & Sons, Inc., 1969.
- [13] K. K. Lin and L.-S. Young, “Shear-induced chaos,” *Nonlinearity*, vol. 21, pp. 899–922, May 2008.
- [14] P. Dayan and L. F. Abbott, *Theoretical Neuroscience*. MIT Press, 2001.
- [15] B. Ermentrout and D. Terman, *Mathematical Foundations of Neuroscience*. Springer New York, 1st ed., 2010.
- [16] A. L. Hodgkin and A. F. Huxley, “A quantitative description of membrane current and its application to conduction and excitation in nerve,” *Journal of Physiology*, vol. 117, pp. 500–544, 1952.
- [17] W. H. Nernst, “Die elektromotorische Wirksamkeit der Ionen,” *Zeitschrift für Physikalische Chemie*, vol. 4, pp. 128–181, 1889.
- [18] A. L. Hodgkin and B. Katz, “The effect of sodium ions on the electrical activity of the giant axon of the squid,” *Journal of Physiology*, vol. 108, pp. 37–77, 1949.
- [19] D. E. Goldman, “Potential, impedance and and rectification in membranes,” *Journal of General Physiology*, vol. 27, no. 1, pp. 37–60, 1943.
- [20] C. M. Armstrong, “Voltage-dependent ion channels and their gating,” *Physiological Reviews*, vol. 72, pp. 5–13, Oct. 1992.
- [21] T. B. Kepler, L. F. Abbott, and E. Marder, “Reduction of conductance-based neuron models,” *Biological Cybernetics*, vol. 66, pp. 381–387, 1992.

- [22] C. Morris and H. Lecar, “Voltage oscillations in the barnacle giant muscle fiber,” *Biophysical Journal*, vol. 35, pp. 193–213, July 1981.
- [23] R. FitzHugh, “Impulses and physiological states in theoretical models of nerve membranes,” *Biophysical Journal*, vol. 1182, pp. 445–466, 1961.
- [24] J. Nagumo, “An Active Pulse Transmission Line Simulating Nerve Axon,” *Proceedings of the IEEE*, vol. 117, pp. 2061–2070, 1962.
- [25] N. Brunel and M. C. W. van Rossum, “Lapicque’s 1907 paper: from frogs to integrate-and-fire,” *Biological Cybernetics*, vol. 97, pp. 337–9, Dec. 2007.
- [26] R. B. Stein, “A theoretical analysis of neuronal variability,” *Biophysical Journal*, vol. 5, no. 2, pp. 173–194, 1965.
- [27] B. W. Knight, “The relationship between the firing rate of a single neuron and the level of activity in a population of neurons. Experimental evidence for resonant enhancement in the population response,” *The Journal of general physiology*, vol. 59, pp. 767–78, June 1972.
- [28] B. W. Knight, “Dynamics of encoding in a population of neurons,” *The Journal of General Physiology*, vol. 59, pp. 734–66, June 1972.
- [29] M. A. Schwemmer and T. J. Lewis, “Bistability in a Leaky Integrate-and-Fire Neuron with a Passive Dendrite,” *SIAM Journal on Applied Dynamical Systems*, vol. 11, no. 1, pp. 507–539, 2012.
- [30] E. M. Izhikevich and G. M. Edelman, “Large-scale model of mammalian thalamocortical systems,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 105, pp. 3593–8, Mar. 2008.
- [31] C. Gold, D. A. Henze, C. Koch, and G. Buzsáki, “On the origin of the extracellular action potential waveform: A modeling study,” *Journal of Neurophysiology*, vol. 95, pp. 3113–28, May 2006.

- [32] W. Gerstner and W. M. Kistler, *Spiking Neuron Models*. Cambridge University Press, 1st ed., 2002.
- [33] L. F. Abbott, “Decoding Neuronal Firing And Modeling Neural Networks,” *Quarterly Review of Biophysics*, vol. 27, pp. 291–331, 1994.
- [34] F. Theunissen and J. P. Miller, “Temporal encoding in nervous systems: a rigorous definition,” *Journal of Computational Neuroscience*, vol. 2, pp. 149–62, June 1995.
- [35] E. M. Izhikevich, “Simple model of spiking neurons.,” *IEEE transactions on Neural Networks*, vol. 14, pp. 1569–72, Jan. 2003.
- [36] R. Brette and W. Gerstner, “Adaptive exponential integrate-and-fire model as an effective description of neuronal activity,” *Journal of Neurophysiology*, vol. 94, pp. 3637–42, Nov. 2005.
- [37] J. Touboul and R. Brette, “Spiking Dynamics of Bidimensional Integrate-and-Fire Neurons,” *SIAM Journal on Applied Dynamical Systems*, vol. 8, no. 4, p. 1462, 2009.
- [38] M. Akhmet, *Principles of Discontinuous Dynamical Systems*. Springer, 1st ed., 2010.
- [39] A. C. J. Luo, *Discontinuous Dynamical Systems*. Springer, 1st ed., 2012.
- [40] R. Brette, M. Rudolph, T. Carnevale, M. Hines, D. Beeman, J. M. Bower, M. Diesmann, A. Morrison, P. H. Goodman, F. C. Harris, M. Zirpe, T. Natschläger, D. Pecevski, B. Ermentrout, M. Djurfeldt, A. Lansner, O. Rochel, T. Vieville, E. Muller, A. P. Davison, S. El Boustani, and A. Destexhe, “Simulation of networks of spiking neurons: a review of tools and strategies,” *Journal of Computational Neuroscience*, vol. 23, pp. 349–98, Dec. 2007.
- [41] A. T. Winfree, *The Geometry of Biological Time*. New York, NY: Springer, 2nd ed., 2001.

- [42] A. T. Winfree, “Patterns of Phase Compromise in Biological Cycles,” *Journal of Mathematical Biology*, vol. 1, pp. 73–95, 1974.
- [43] J. Guckenheimer, “Isochrons and Phaseless Sets,” *Journal of Mathematical Biology*, vol. 1, pp. 259–273, 1975.
- [44] I. G. Malkin, *The methods of Lyapunov and Poincare in the theory of nonlinear oscillations*. Editorial URSS, 1st ed., 1949.
- [45] I. G. Malkin, *Some problems in the theory of nonlinear oscillations*. United States Atomic Energy Commission, 1st ed., 1959.
- [46] A. Guillamon and G. Huguet, “A Computational and Geometric Approach to Phase Resetting Curves and Surfaces,” *SIAM Journal on Applied Dynamical Systems*, vol. 8, no. 3, pp. 1005–1042, 2009.
- [47] E. M. Izhikevich, *Dynamical Systems in Neuroscience : The Geometry of Excitability and Bursting*. MIT Press, 1st ed., 2005.
- [48] H. M. Osinga and J. Moehlis, “A continuation method for computing global isochrons,” *SIAM Journal on Applied Dynamical Systems*, vol. 9, no. 4, pp. 1201–1228, 2010.
- [49] A. Mauroy and I. Mezic, “On the use of Fourier averages to compute the global isochrons of (quasi)periodic dynamics,” *Chaos*, vol. 22, p. 033112, Sept. 2012.
- [50] D. Hansel, G. Mato, and C. Meunier, “Synchrony in excitatory neural networks,” *Neural Computation*, vol. 7, pp. 307–337, Mar. 1995.
- [51] T. Netoff, M. A. Schwemmer, and T. J. Lewis, “Experimentally Estimating Phase Response Curves of Neurons : Theoretical and Practical Issues,” in *PRCs in Neuroscience; Theory, Experiment and Analysis* (N. W. Schultheiss, A. A. Prinz, and R. J. Butera, eds.), ch. 5, pp. 95–130, Springer, 1st ed., 2012.

- [52] R. Galán, G. Ermentrout, and N. Urban, “Efficient Estimation of Phase-Resetting Curves in Real Neurons and its Significance for Neural-Network Modeling,” *Physical Review Letters*, vol. 94, pp. 1–4, Apr. 2005.
- [53] M. N. Economo, F. R. Fernandez, and J. A. White, “Dynamic clamp: alteration of response properties and creation of virtual realities in neurophysiology,” *The Journal of Neuroscience*, vol. 30, pp. 2407–2413, Feb. 2010.
- [54] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert, “Noise in the nervous system,” *Nature Reviews. Neuroscience*, vol. 9, pp. 292–303, Apr. 2008.
- [55] L. Glass and M. C. Mackey, *From Clocks to Chaos*. Princeton University Press, 1st ed., 1988.
- [56] S. H. Strogatz, *Nonlinear Dynamics and Chaos: with Applications to Physics, Biology, Chemistry and Engineering*. Perseus Books, 1st ed., 2000.
- [57] H. C. Berg and E. M. Purcell, “Physics of chemoreception,” *Biophysical journal*, vol. 20, pp. 193–219, Nov. 1977.
- [58] W. Bialek and S. Setayeshgar, “Physical limits to biochemical signaling,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, pp. 10040–10045, July 2005.
- [59] A. F. Strassberg and L. J. DeFelice, “Limitations of the Hodgkin-Huxley Formalism: Effects of Single Channel Kinetics on Transmembrane Voltage Dynamics,” *Neural Computation*, vol. 5, pp. 843–855, 1993.
- [60] A. L. Fairhall, G. D. Lewen, W. Bialek, and R. R. de Ruyter Van Steveninck, “Efficiency and ambiguity in an adaptive neural code,” *Nature*, vol. 412, pp. 787–92, Aug. 2001.
- [61] W. Bialek, F. Rieke, R. R. de Ruyter van Steveninck, and D. Warland, “Reading a neural code,” *Science*, vol. 252, pp. 1854–1857, June 1991.

- [62] E. Schneidman, B. Freedman, and I. Segev, “Ion channel stochasticity may be critical in determining the reliability and precision of spike timing,” *Neural Computation*, vol. 10, pp. 1679–1703, Oct. 1998.
- [63] J. A. White, J. T. Rubinstein, and A. R. Kay, “Channel noise in neurons,” *Trends in Neurosciences*, vol. 23, pp. 131–137, Mar. 2000.
- [64] W. Rall, “Time constants and electrotonic length of membrane cylinders and neurons,” *Biophysical Journal*, vol. 9, pp. 1483–1508, Dec. 1969.
- [65] A. A. Faisal and S. B. Laughlin, “Stochastic simulations on the reliability of action potential propagation in thin axons,” *PLoS Computational Biology*, vol. 3, p. e79, May 2007.
- [66] L. Harold and R. Nossal, “Theory of threshold fluctuations in nerves. II Analysis of various sources of membrane noise,” *Biophysical Journal*, vol. 11, pp. 1068–1084, 1971.
- [67] A. Manwani and C. Koch, “Detecting and Estimating Signals in Noisy Cable Structures, I: Neuronal Noise Sources,” *Neural Computation*, vol. 11, pp. 1797–1829, Nov. 1999.
- [68] P. Faure, D. Kaplan, and H. Korn, “Synaptic Efficacy and the Transmission of Complex Firing Patterns Between Neurons Synaptic Efficacy and the Transmission of Complex Firing Patterns Between Neurons,” *Journal of Neurophysiology*, vol. 84, no. 6, pp. 3010–3025, 2000.
- [69] W. Calvin and C. Stevens, “Synaptic noise as a source of variability in the interval between action potentials,” *Science*, vol. 17, no. 155, pp. 842–844, 1967.
- [70] R. Conti, Y. P. Tan, and I. Llano, “Action potential-evoked and ryanodine-sensitive spontaneous Ca^{2+} transients at the presynaptic terminal of a developing CNS inhibitory synapse,” *The Journal of Neuroscience*, vol. 24, pp. 6946–6957, Aug. 2004.

- [71] K. K. Lin, E. Shea-Brown, and L.-S. Young, “Reliability of Coupled Oscillators,” *Journal of Nonlinear Science*, vol. 19, pp. 497–545, Feb. 2009.
- [72] R. Baddeley, P. Hancock, and P. Földiák, eds., *Information Theory and the Brain*. Cambridge University Press, 2000.
- [73] J. von Neuman, “Probabilistic logics and synthesis of reliable organisms from unreliable components,” in *Automata Studies* (C. Shannon and J. McCarthy, eds.), pp. 43–98, Princeton University Press, 1st ed., 1956.
- [74] E. Glowatzki and P. A. Fuchs, “Transmitter release at the hair cell ribbon synapse,” *Nature Neuroscience*, vol. 5, pp. 147–154, Feb. 2002.
- [75] T. L. Adelman, W. Bialek, and R. M. Olberg, “The information content of receptive fields,” *Neuron*, vol. 40, pp. 823–833, Nov. 2003.
- [76] M. D. McDonnell and L. M. Ward, “The benefits of noise in neural systems: bridging theory and experiment,” *Nature Reviews. Neuroscience*, vol. 12, pp. 415–425, 2011.
- [77] R. Benzi, A. Sutera, and A. Vulpiani, “The mechanism of stochastic resonance,” *Journal of Physics A: Mathematical and Theoretical*, vol. 14, pp. 453–457, 1981.
- [78] A. Pikovsky and J. Kurths, “Coherence Resonance in a Noise-Driven Excitable System,” *Physical Review Letters*, vol. 78, pp. 775–778, Feb. 1997.
- [79] J. S. Anderson, I. Lampl, D. C. Gillespie, and D. Ferster, “The contribution of noise to contrast invariance of orientation tuning in cat visual cortex,” *Science*, vol. 290, pp. 1968–1972, Dec. 2000.
- [80] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer, 4th ed., 2009.

- [81] D. T. Gillespie, “Exact Stochastic Simulation of Coupled Chemical Reactions,” *The Journal of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.
- [82] S. P. Brooks, “Markov chain Monte Carlo method and its application,” *Journal of the Royal Statistical Society, Series D*, vol. 47, no. 1, pp. 69–100, 1998.
- [83] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, “Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method,” *The Journal of Chemical Physics*, vol. 119, no. 24, pp. 12784–12794, 2003.
- [84] R. Mannella and P. V. E. McClintock, “Itô versus Stratonovich: 30 Years Later,” *Fluctuation and Noise Letters*, vol. 11, 2012.
- [85] B. Lindner, “Interspike interval statistics of neurons driven by colored noise,” *Physical Review E*, vol. 69, p. 022901, Feb. 2004.
- [86] G. E. Uhlenbeck and L. S. Ornstein, “On the Theory of the Brownian Motion,” *Physical Review*, vol. 36, pp. 823–841, 1930.
- [87] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*. Springer, 1st ed., 1992.
- [88] D. J. Higham, “An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations,” *SIAM Review*, vol. 3, pp. 4863–4869, Oct. 2001.
- [89] K. Burrage and P. M. Burrage, “High strong order explicit Runge-Kutta methods for stochastic ordinary differential equations,” *Applied Numerical Mathematics*, vol. 20, pp. 1–21, 1996.
- [90] E. Buckwar, R. Horváth-Bokor, and R. Winkler, “Asymptotic Mean-Square Stability of Two-Step Methods for Stochastic Ordinary Differential Equations,” *BIT Numerical Mathematics*, vol. 46, no. 2, pp. 261–282, 2006.

- [91] E. Buckwar and R. Winkler, “Multi-step methods for sdes and their application to problems with small noise,” *SIAM Journal on Numerical Analysis*, vol. 44, no. 2, pp. 779–803, 2006.
- [92] G. N. Milstein and M. V. Tretyakov, “Numerical methods in the weak sense for stochastic differential equations with small noise,” *SIAM Journal of Numerical Analysis*, vol. 34, no. 6, pp. 2142–2167, 1997.
- [93] H. Risken, *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer, 3rd ed., 1996.
- [94] J.-N. Teramae, “Robustness of the Noise-Induced Phase Synchronization in a General Class of Limit Cycle Oscillators,” *Physical Review Letters*, vol. 93, p. 204103, Nov. 2004.
- [95] D. Goldobin and A. Pikovsky, “Synchronization of self-sustained oscillators by common white noise,” *Physica A*, vol. 351, pp. 126–132, June 2005.
- [96] H. Nakao, K. Arai, and Y. Kawamura, “Noise-Induced Synchronization and Clustering in Ensembles of Uncoupled Limit-Cycle Oscillators,” *Physical Review Letters*, vol. 98, p. 184101, May 2007.
- [97] C. Ly and G. Bard Ermentrout, “Analytic approximations of statistical quantities and response of noisy oscillators,” *Physica D: Nonlinear Phenomena*, vol. 240, pp. 719–731, Apr. 2011.
- [98] S. Coombes, R. Thul, and K. Wedgwood, “Nonsmooth dynamics in spiking neuron models,” *Physica D: Nonlinear Phenomena*, vol. 241, pp. 2042–2057, May 2012.
- [99] H. P. McKean, “Nagumo’s equation,” *Advances in Mathematics*, vol. 4, pp. 209–223, 1970.
- [100] W. P. Wang, “Multiple impulse solutions to McKean’s caricature of the nerve equation. I. Existence,” *Communications on Pure and Applied Mathematics*, vol. 41, pp. 71–103, 1988.

- [101] W. P. Wang, “Multiple impulse solutions to McKean’s caricature of the nerve equation. II. Stability,” *Communications on Pure and Applied Mathematics*, vol. 41, pp. 997–1025, 1988.
- [102] A. Tonnelier, “The McKean’s caricature of the FitzHugh-Nagumo model I. The space-clamped system,” *SIAM Journal on Applied Mathematics*, vol. 63, pp. 459–484, 2002.
- [103] A. F. Filippov, *Differential Equations with Discontinuous Righthand Sides*. Springer, 1st ed., 1960.
- [104] S. Coombes, “Neuronal networks with gap junctions: A study of piece-wise linear planar neuron models,” *SIAM Journal on Applied Dynamical Systems*, vol. 7, pp. 1101–1129, 2008.
- [105] H. G. Rotstein, S. Coombes, and A.-M. Gheorghe, “Canard-like explosion of limit cycles in two-dimensional piecewise-linear models of FitzHugh-Nagumo type,” *SIAM Journal on Applied Dynamical Systems*, vol. 11, no. 1, pp. 135–180, 2012.
- [106] B. W. Connors and M. J. Gutnick, “Intrinsic firing patterns of diverse neocortical neurons,” *Trends in Neurosciences*, vol. 13, no. 3, pp. 99–104, 1990.
- [107] R. J. Butera, J. Rinzel, and J. C. Smith, “Models of respiratory rhythm generation in the pre-Bötzinger complex. I. Bursting pacemaker neurons,” *Journal of Neurophysiology*, vol. 82, no. 1, pp. 382–397, 1999.
- [108] R. J. Butera, J. Rinzel, and J. C. Smith, “Models of Respiratory Rhythm Generation in the Pre-Bötzinger Complex. II. Populations of Coupled Pacemaker Neurons,” *Journal of Neurophysiology*, vol. 82, pp. 398–415, 1999.
- [109] J. Lisman, “Bursts as a unit of neural information: making unreliable synapses reliable,” *Trends in Neurosciences*, vol. 20, pp. 38–43, Jan. 1997.
- [110] S. M. Sherman, “Tonic and burst firing: dual modes of thalamocortical relay,” *Trends in Neurosciences*, vol. 24, pp. 122–126, Mar. 2001.

- [111] P. Reinagel, D. Godwin, S. M. Sherman, C. Koch, M. Tsanov, E. Chah, N. Wright, S. D. Vann, R. Reilly, T. Jonathan, J. P. Aggleton, and S. M. O. Mara, “Encoding of Visual Information by LGN Bursts Encoding of Visual Information by LGN Bursts,” *Journal of Neurophysiology*, vol. 81, pp. 2558–2569, 1999.
- [112] A.-M. M. Oswald, M. J. Chacron, B. Doiron, J. Bastian, and L. Maler, “Parallel processing of sensory input by bursts and isolated spikes,” *The Journal of Neuroscience*, vol. 24, pp. 4351–4362, May 2004.
- [113] J. Rinzel, “Bursting oscillations in an excitable membrane model,” in *Ordinary and Partial Differential Equations. Proceedings of the Eighth Conference held at Dundee, Scotland, June 2529, 1984* (B. D. Sleeman and R. J. Jarvis, eds.), pp. 304–316, Springer, 1st ed., 1985.
- [114] D. Smith, *Singular-perturbation theory*. Cambridge: Cambridge University Press, 1st ed., 1985.
- [115] N. Fenichel, *Invariant Manifolds for Flows: Persistence and Smoothness*. New York University, 1st ed., 1970.
- [116] E. M. Izhikevich and F. C. Hoppensteadt, “Classification of bursting mappings,” *International Journal of Bifurcation and Chaos*, vol. 14, no. 11, pp. 3847–3854, 2004.
- [117] A. Wagemakers, M. A. F. Sanjuán, J. M. Casado, and K. Aihara, “Building electronic bursters with the Morris-Lecar neuron model,” *International Journal of Bifur*, vol. 16, no. 12, pp. 3617–3630, 2006.
- [118] E. M. Izhikevich, “Synchronization of Elliptic Bursters,” *SIAM Review*, vol. 43, no. 2, pp. 315–344, 2001.
- [119] S. J. Hogan, “On the dynamics of rigid-block motion under harmonic forcing,” *Proceedings of the Royal Society of London. Series A*, vol. 425, pp. 441–476, 1989.

- [120] A. Kaharaman and R. Singh, “Nonlinear dynamics of a spur gear pair,” *Journal of Sound and Vibration*, vol. 142, pp. 49–75, 1990.
- [121] P. C. Tung and S. W. Shaw, “A method for the improvement of impact printer and performance,” *ASME Journal of Vibration, Acoustics, Stress, and Reliability in Design*, vol. 110, no. 528-532, 1988.
- [122] M. di Bernardo, C. Budd, A. R. Champneys, and P. Kowalczyk, *Piecewise-smooth Dynamical Systems: Theory and Applications*. Springer, 2007.
- [123] J. Platkiewicz and R. Brette, “A threshold equation for action potential initiation,” *PLoS Computational Biology*, vol. 6, p. e1000850, 2010.
- [124] P. C. Bressloff and S. Coombes, “Dynamics of strongly coupled spiking neurons,” *Neural Computation*, vol. 12, pp. 91–129, 2000.
- [125] P. E. Latham, B. J. Richmond, P. G. Nelson, and S. Nirenberg, “Intrinsic dynamics in neuronal networks I. Theory,” *Journal of Neurophysiology*, vol. 83, pp. 808–827, 2000.
- [126] N. Fourcaud-Trocme, D. Hansel, C. van Vreeswijk, and N. Brunel, “How spike generation mechanisms determine the neuronal response to fluctuating inputs,” *Journal of Neuroscience*, vol. 23, pp. 11628–11640, 2003.
- [127] E. M. Izhikevich, “Simple model of spiking neurons,” *IEEE Transactions On Neural Networks*, vol. 14, pp. 1569–1572, 2003.
- [128] E. M. Izhikevich, “Hybrid spiking models,” *Philosophical Transactions of the Royal Society of London. Series A*, vol. 368, pp. 5061–5070, 2010.
- [129] J. Touboul, “Importance of the Cutoff Value in the Quadratic Adaptive Integrate-and-Fire Model,” *Neural Computation*, vol. 21, pp. 2114–2122, 2009.

- [130] T. Gröbler, G. Barna, and P. Érdi, “Statistical model of the hippocampal CA3 region I. The single-cell module: bursting model of the pyramidal cell,” *Biological Cybernetics*, vol. 79, pp. 301–308, 1998.
- [131] R. Brette and W. Gerstner, “Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity,” *Journal of Neurophysiology*, vol. 94, pp. 3637–3642, 2005.
- [132] J. Touboul and R. Brette, “Dynamics and bifurcations of the adaptive exponential integrate-and-fire model,” *Biological Cybernetics*, vol. 99, pp. 319–334, 2008.
- [133] J. Touboul, “Bifurcation Analysis of a General Class of Nonlinear Integrate-and-Fire Neurons,” *SIAM Journal on Applied Mathematics*, vol. 68, pp. 1045–1079, 2008.
- [134] J. Karbowski and N. Kopell, “Multispikes and synchronization in a large-scale neural network with delays,” *Neural Computation*, vol. 12, pp. 1573–1606, 2000.
- [135] S. Coombes and M. Zachariou, *Gap junctions and emergent rhythms*, vol. 3 of *Springer Series in Computational Neuroscience*, pp. 77–94. Springer, 2009.
- [136] D. Simpson and J. Meiss, “AndronovHopf bifurcations in planar, piecewise-smooth, continuous flows,” *Physics Letters A*, vol. 371, pp. 213–220, Nov. 2007.
- [137] S. Coombes, “Neuronal Networks with Gap Junctions: A Study of Piecewise Linear Planar Neuron Models,” *SIAM Journal on Applied Dynamical Systems*, vol. 7, no. 3, p. 1101, 2008.
- [138] D. Hansel, G. Mato, C. Meunier, and L. Neltner, “On Numerical Simulations of Integrate-and-Fire Neural Networks,” *Neural Computation*, vol. 10, pp. 467–483, Feb. 1998.

- [139] B. Ermentrout, *Simulating, Analyzing, and Animating Dynamical Systems: A Guide to Xppaut for Researchers and Students*. Society for Industrial & Applied Mathematics, 2002.
- [140] Y. Kuramoto, *Chemical Oscillations, Waves and Turbulence*. Springer-Verlag, 1984.
- [141] G. B. Ermentrout and N. Kopell, “Parabolic Bursting in an Excitable System Couple with a Slow Oscillation,” *SIAM Journal on Applied Mathematics*, vol. 46, no. 2, pp. 233–253, 1986.
- [142] G. S. Medvedev, “Reduction of a model of an excitable cell to a one-dimensional map,” *Physica D*, vol. 202, pp. 37–59, 2005.
- [143] T. R. Chay and J. Keizer, “Minimal model for membrane oscillations in the pancreatic Beta-cell,” *Biophysical Journal*, vol. 42, pp. 181–190, 1983.
- [144] R. Bertram and A. Sherman, *Negative calcium feedback: the road from Chay-Keizer*, pp. 19–48. World Scientific, 2005.
- [145] H. Marotto, “Snap-back repellers imply chaos in \mathbb{R}^n ,” *Journal of Mathematical Analysis and Applications*, vol. 63, pp. 199–223, 1978.
- [146] G. Zheng and A. Tonneier, “Chaotic solutions in the quadratic integrate-and-fire neuron with adaptation,” *Cognitive Neurodynamics*, vol. 3, pp. 197–204, 2009.
- [147] G. Benettin, L. Galgani, and J.-M. Strelcyn, “Kolmogorov entropy and numerical experiments,” *Physical Review A*, vol. 14, no. 6, pp. 2338–2345, 1976.
- [148] I. Shimada and T. Nagashima, “A Numerical Approach to Ergodic Problem of Dissipative Dynamical Systems,” *Progress of Theoretical Physics*, vol. 61, no. 6, pp. 1605–1616, 1979.

- [149] P. C. Müller, “Calculation of Lyapunov exponents for dynamic systems with discontinuities,” *Chaos, Solitons & Fractals*, vol. 5, no. 9, pp. 1671–1681, 1995.
- [150] S. Coombes and P. C. Bressloff, “Mode-locking and Arnold tongues in integrate-and-fire neural oscillators,” *Physical Review E*, vol. 60, pp. 2086–2096, 1999.
- [151] L. Shiau and C. R. Laing, “Periodically forced piecewise linear adaptive exponential integrate-and-fire neuron,” *International Journal of Bifurcation and Chaos (to appear)*.
- [152] S. Coombes, R. Thul, J. Laudanski, A. R. Palmer, and C. J. Sumner, “Neuronal spike-train responses in the presence of threshold noise,” *Frontiers in Life Science*, vol. 5, no. 3-4, pp. 91–105, 2011.
- [153] K. K. Lin, E. Shea-Brown, and L.-S. Young, “Reliability of Layered Neural Oscillator Networks,” *Communications in Mathematical Sciences*, vol. 7, pp. 239–247, 2009.
- [154] D. Simpson and R. Kuske, “Mixed-mode oscillations in a stochastic, piecewise-linear system,” *Physica D: Nonlinear Phenomena*, vol. 240, pp. 1189–1198, July 2011.
- [155] J. Rubin and D. Terman, “Synchronized Activity and Loss of Synchrony Among Heterogeneous Conditional Oscillators,” *SIAM Journal on Applied Dynamical Systems*, vol. 1, no. 1, pp. 146–174, 2002.
- [156] A. Pikovsky, M. Rosenblum, and J. Kurths, *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press, 2001.
- [157] K. K. Lin, K. C. A. Wedgwood, S. Coombes, and L.-S. Young, “Limitations of perturbative techniques in the analysis of rhythms and oscillations,” *Journal of Mathematical Biology*, vol. 66, pp. 139–161, Jan. 2013.

- [158] E. Brown, J. Moehlis, and P. Holmes, “On the Phase Reduction and Response Dynamics of Neural Oscillator Populations,” *Neural Computation*, vol. 16, pp. 673–715, 2004.
- [159] S. A. Oprisan, V. Thirumalai, and C. C. Canavier, “Dynamics from a time series: can we extract the phase resetting curve from a time series?,” *Biophysical Journal*, vol. 84, pp. 2919–2928, May 2003.
- [160] T. Netoff, C. Acker, J. Bettencourt, and J. A. White, “Beyond two-cell networks: experimental measurement of neuronal responses to multiple synaptic inputs,” *Journal of Computational Neuroscience*, vol. 18, pp. 287–295, 2005.
- [161] G. B. Ermentrout, “Losing amplitude and saving phase,” in *Nonlinear Oscillations in Biology* (H. G. Othmer, ed.), Springer-Verlag, 1986.
- [162] Q. Wang and L.-S. Y. Young, “From Invariant Curves to Strange Attractors,” *Communications in Mathematical Physics*, vol. 225, pp. 275–304, Feb. 2002.
- [163] Q. Wang and L.-S. Young, “Strange Attractors in Periodically-Kicked Limit Cycles and Hopf Bifurcations,” *Communications in Mathematical Physics*, vol. 240, pp. 509–529, 2003.
- [164] Q. Wang and L.-S. Young, “Strange Attractors with One Direction of Instability,” *Communications in Mathematical Physics*, vol. 218, pp. 1–97, Apr. 2001.
- [165] Q. Wang, “Toward a Theory of Rank One Attractors,” *Annals of Mathematics*, vol. 167, pp. 349–480, 2008.
- [166] W. Ott and M. Stenlund, “From Limit Cycles to Strange Attractors,” *Communications in Mathematical Physics*, vol. 296, pp. 215–249, 2010.
- [167] J. Guckenheimer and P. Holmes, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. Springer, 1983.

- [168] G. Zaslavsky, “The simplest case of a strange attractor,” *Physics Letters A*, vol. 69, no. 3, pp. 145–147, 1978.
- [169] K. Lu, Q. Wang, and L.-S. Young, “Strange attractors for periodically forced parabolic equations,” *Memoirs of the American Mathematical Society (to appear)*.
- [170] Q. Wang and W. Ott, “Dissipative Homoclinic Loops of Two-Dimensional Maps and Strange Attractors with One Direction of Instability,” *Communications on Pure and Applied Mathematics*, vol. 64, pp. 1439–1496, 2011.
- [171] R. Deville, N. S. Namanchchivaya, and Z. Rapti, “Stability of a stochastic two-dimensional non-Hamiltonian system,” *SIAM Journal on Applied Mathematics*, vol. 71, no. 4, pp. 1458–1475, 2011.
- [172] K. K. Lin, “Entrainment and Chaos in a Pulse-Driven HodgkinHuxley Oscillator,” *SIAM*, vol. 5, no. 2, pp. 179–204, 2006.
- [173] F. Rieke, D. Warland, R. R. de Ruyter van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code*. MIT Press, 2nd ed., 1999.
- [174] S. K. Han, C. Kurrer, and Y. Kuramoto, “Diffusive Interaction Leading to Dephasing of Coupled Neural Oscillators,” *International Journal of Bifurcation and Chaos*, vol. 7, no. 4, pp. 869–876, 1997.
- [175] D. Postnov, S. K. Han, and H. Kook, “Synchronization of diffusively coupled oscillators near the homoclinic bifurcation,” *Physical Review E*, vol. 60, pp. 2799–807, Sept. 1999.
- [176] G. B. Ermentrout and N. Kopell, “Oscillator Death in Systems of Coupled Neural Oscillators,” *SIAM Journal on Applied Mathematics*, vol. 50, no. 1, pp. 125–146, 1990.
- [177] A. Demir and O. Suvak, “Quadratic approximations for the isochrons of oscillators: A general theory, advanced numerical methods and accurate phase

- computations,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 29, pp. 1215–1228, Aug. 2010.
- [178] G. S. Medvedev, “Synchronization of coupled stochastic limit cycle oscillators,” *Physics Letters A*, no. 2, pp. 1–17, 2010.
- [179] S. P. Diliberto, “On systems of ordinary differential equations,” in *Contributions to the Theory of Nonlinear Oscillations, Annals of Mathematical Studies*, vol. 20, pp. 1–38, Princeton University Press, 1950.
- [180] K. C. Wedgwood, K. K. Lin, R. Thul, and S. Coombes, “Phase-amplitude descriptions of neural oscillator models,” *Journal of Mathematical Neuroscience*, vol. 3, pp. 1–22, Jan. 2013.
- [181] G. B. Ermentrout and N. Kopell, “Oscillator death in systems of coupled neural oscillators,” *SIAM Journal on Applied Mathematics*, vol. 50, pp. 125–146, 1990.
- [182] R. Fitzhugh, “Impulses and physiological states in theoretical models of nerve membrane,” *Biophysical Journal*, vol. 1, pp. 445–466, 1961.
- [183] B. van der Pol, “A theory of the amplitude of free and forced triode vibrations,” *Radio Review*, vol. 1, pp. 701–710, 1920.
- [184] J. A. Connor and C. F. Stevens, “Prediction of repetitive firing behaviour from voltage clamp data on an isolated neurone soma,” *Journal of Physiology*, vol. 213, pp. 31–53, 1971.
- [185] A. J. Catlla, D. G. Schaeffer, T. P. Witelski, E. E. Monson, and A. L. Lin, “On Spiking Models for Synaptic Activity and Impulsive Differential Equations,” *SIAM Review*, vol. 50, no. 3, pp. 553–569, 2008.
- [186] X.-J. Wang and J. Rinel, “Alternating and Synchronous Rhythms in Reciprocally Inhibitory Model Neurons,” *Neural Computation*, vol. 4, pp. 84–97, 1992.

- [187] P. Ashwin, “Weak coupling of strongly nonlinear, weakly dissipative identical oscillators,” *Dynamical Systems*, vol. 10, no. 3, pp. 2471–2474, 1989.
- [188] P. Ashwin and G. Dangelmayr, “Isochronicity-induced bifurcations in systems of weakly dissipative coupled oscillators,” *Dynamics and Stability of Systems*, vol. 15, no. 3, pp. 263–286, 2000.
- [189] P. Ashwin and G. Dangelmayr, “Reduced dynamics and symmetric solutions for globally coupled weakly dissipative oscillators,” *Dynamical Systems*, vol. 20, pp. 333–367, Sept. 2005.
- [190] S. K. Han, C. Kurrer, and Y. Kuramoto, “Dephasing and bursting in coupled neural oscillators,” *Physical Review Letters*, vol. 75, pp. 3190–3193, 1995.
- [191] A. Sherman and J. Rinzel, “Rhythmogenic effects of weak electrotonic coupling in neuronal models,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 89, pp. 2471–2474, Mar. 1992.
- [192] A. K. A. Azad and P. Ashwin, “Within-Burst Synchrony Changes for Coupled Elliptic Bursters,” *SIAM Journal on Applied Dynamical Systems*, vol. 9, no. 1, pp. 261–281, 2010.
- [193] B. Ermentrout and D. Saunders, “Phase resetting and coupling of noisy neural oscillators,” *Journal of Computational Neuroscience*, vol. 20, pp. 179–190, Apr. 2006.
- [194] J. P. Gleeson, “Phase Diffusion Due to Low-Frequency Colored Noise,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 53, no. 3, pp. 183–186, 2006.
- [195] W. F. Thompson, R. Kuske, and Y. X. Li, “Stochastic phase dynamics of noise driven synchronization of two conditional coherent oscillators,” *Discrete and Continuous Dynamical Systems*, vol. 32, no. 8, pp. 2971–2996, 2012.

- [196] M. Bonnin and F. Corinto, “Phase Noise and Noise Induced Frequency Shift in Stochastic Nonlinear Oscillators,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. PP, no. 99, pp. 1–12, 2013.
- [197] Z. F. Mainen and T. J. Sejnowski, “Reliability of spike timing in neocortical neurons.,” *Science*, vol. 268, pp. 1503–1506, June 1995.
- [198] M. N. Shadlen and W. T. Newsome, “The variable discharge of cortical neurons: implications for connectivity, computation, and information coding,” *The Journal of Neuroscience*, vol. 18, pp. 3870–3896, May 1998.
- [199] K. Yoshimura and K. Arai, “Phase Reduction of Stochastic Limit Cycle Oscillators,” *Physical Review Letters*, vol. 101, pp. 10–13, Oct. 2008.
- [200] G. N. Milstein and M. V. Tretyakov, *Stochastic Numerics for Mathematical Physics*. Berlin: Springer, 1st ed., 2004.
- [201] K. Yoshimura, *Phase Reduction of Stochastic Limit-Cycle Oscillators*, vol. 3, pp. 59–90. Wiley, 1st ed., 2010.
- [202] D. Holcman and M. Tsodyks, “The emergence of Up and Down states in cortical networks.,” *PLoS Computational Biology*, vol. 2, p. e23, Mar. 2006.
- [203] P. Fatt and B. Katz, “An analysis of the end-plate potential recorded with an intra-cellular electrode,” *Journal of Physiology*, vol. 115, pp. 320–370, 1951.
- [204] P. A. Abrams, “Niche Overlap and Environmental Variability,” *Mathematical Biosciences*, vol. 28, pp. 357–372, 1976.
- [205] R. T. Fenner, *Finite element methods for engineers*. Imperial College Press, 2nd ed., 2013.
- [206] A. Logg, K.-A. Mardal, and G. Wells, eds., *Automated Solution of Differential Equations by the Finite Element Method*, vol. 84 of *Lecture Notes in Computational Science and Engineering*. Springer, 2012.

- [207] “<http://fenicsproject.org/>.”
- [208] J. C. Strikwerda, *Finite Difference Schemes and Partial Differential Equations*. SIAM, 2nd ed., 1997.
- [209] P. Kumar and S. Narayanan, “Solution of Fokker–Planck equation by finite element and finite difference methods for nonlinear systems,” *Sadhana*, vol. 31, pp. 445–461, 2006.
- [210] F. Freyer, J. A. Roberts, P. Ritter, and M. Breakspear, “A Canonical Model of Multistability and Scale-Invariance in Biological Systems,” *PLoS Computational Biology*, vol. 8, p. e1002634, Aug. 2012.
- [211] L. Arnold, *Random Dynamical Systems*. Springer, 1st ed., 2003.
- [212] S. F. J. Wojtkiewicz, *Robust Numerical Solution of the Fokker-Planck-Kolmogorov Equation for Two Dimensional Stochastic Dynamical Systems*. University of Illinois at Urbana-Champaign, 1st ed., 1994.
- [213] N. Berglund, B. Gentz, and C. Kuehn, “Hunting French ducks in a noisy environment,” *Journal of Differential Equations*, vol. 252, pp. 4786–4841, May 2012.
- [214] S. L. Hooper, “Central Pattern Generators,” *Encyclopaedia of Life Sciences*, 2001.
- [215] E. Marder and D. Bucher, “Central pattern generators and the control of rhythmic movements,” *Current Biology*, vol. 11, pp. 986–996, Nov. 2001.
- [216] D. Lino, C. Alan, M. Desroches, and M. Storace, “Codimension-Two Homoclinic Bifurcations Underlying Spike Adding in the Hindmarsh-Rose Bursting,” *SIAM Journal on Applied Dynamical Systems*, vol. 11, no. 3, pp. 939–962, 2012.

- [217] P. Channell, G. Cymbalyuk, and A. Shilnikov, “Origin of Bursting through Homoclinic Spike Adding in a Neuron Model,” *Physical Review Letters*, vol. 98, p. 134101, Mar. 2007.
- [218] F. Bizzarri, A. Brambilla, and G. G. Storti, “Lyapunov exponents computation for hybrid neurons,” *Journal of Computational Neuroscience*, vol. in press, 2013.
- [219] K. Youngtae, “Identification of Dynamical States in Stimulated Izhikevich Neuron Models by Using a 0-1 Test,” *Journal of the Korean Physical Society*, vol. 57, p. 1363, Dec. 2010.
- [220] S. Nobukawa, H. Nishimura, T. Yamanishi, and J.-Q. Liu, “Signal Response Efficiency in Izhikevich Neuron Model,” in *SICE Annual Conference*, no. 6, pp. 1242–1247, 2011.
- [221] M. di Volo, R. Livi, S. Luccioli, A. Politi, and A. Torcini, “Synchronous dynamics in the presence of short-term plasticity,” *Physical Review E*, vol. 87, p. 032801, Mar. 2013.
- [222] C. R. Laing and I. G. Kevrekidis, “Periodically-forced nite networks of heterogeneous globally-coupled oscillators: a low-dimensional approach,” *Physica D: Nonlinear Phenomena*, vol. 237, no. 2, pp. 207–215, 2007.
- [223] H. E. Plesser and W. Gerstner, “Noise in integrate-and-fire neurons: from stochastic input to escape rates,” *Neural Computation*, vol. 12, pp. 367–384, Feb. 2000.
- [224] B. Lindner, “Effects of noise in excitable systems,” *Physics Reports*, vol. 392, pp. 321–424, Mar. 2004.
- [225] A. V. Rangan and L.-S. Y. Young, “A network model of V1 with collaborative activity,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 104, pp. 20167–72, Dec. 2007.

- [226] M. Nedeljkov and M. Oberguggenberger, “Ordinary differential equations with delta function terms,” *Publications de l’Institut Mathématique*, vol. 91, no. 105, pp. 125–135, 2012.
- [227] M. Tsodyks, K. Pawelzik, and H. Markram, “Neural networks with dynamic synapses,” *Neural Computation*, vol. 10, pp. 821–835, May 1998.
- [228] R. Fox and Y. Lu, “Emergent collective behavior in large numbers of globally coupled independently stochastic ion channels,” *Physical Review E*, vol. 49, pp. 3421–3431, Apr. 1994.
- [229] J. H. Goldwyn and E. Shea-Brown, “The what and where of adding channel noise to the Hodgkin-Huxley equations,” *PLoS Computational Biology*, vol. 7, p. e1002247, Nov. 2011.
- [230] M. Güler, “Stochastic Hodgkin-Huxley equations with colored noise terms in the conductances,” *Neural Computation*, vol. 25, pp. 46–74, Jan. 2013.
- [231] R. F. Fox, “Stochastic Versions of the Hodgkin-Huxley Equations,” *Biophysical Journal*, vol. 72, pp. 2068–2074, 1997.
- [232] R. F. Galán, G. B. Ermentrout, and N. N. Urban, “Solving the Fokker-Planck equation with the finite-element method: An example studying stochastic synchronisation of neuronal oscillators,” *Physical Review E*, vol. 76, 2007.
- [233] S. F. J. Wojtkiewicz and L. A. Bergman, “Numerical Solution of High Dimensional Fokker-Planck Equations,” in *8th ASCE Specialty Conference on Probabilistic Mechanics and Structural Reliability*, (Paris), pp. 1–6, 2000.
- [234] A. McCulloch, L. Waldman, J. Rogers, and J. Guccione, “Large-scale finite element analysis of the beating heart,” *Critical Reviews on Biomedical Engineering*, vol. 20, pp. 427–449, 1992.
- [235] N. Berglund and D. Landon, “Mixed-mode oscillations and interspike interval statistics in the stochastic FitzHugh-Nagumo model,” *Nonlinearity*, vol. 25, pp. 2303–2335, Aug. 2012.

- [236] J. F. Lindner, B. K. Meadows, W. L. Ditto, M. E. Inchiosa, and A. R. Bulsara, “Scaling laws for spatiotemporal synchronization and array enhanced stochastic resonance,” *Physical Review E*, vol. 53, no. 3, pp. 2081–2086, 1996.
- [237] C. Zhou, J. Kurths, and B. Hu, “Array-Enhanced Coherence Resonance: Nontrivial Effects of Heterogeneity and Spatial Independence of Noise,” *Physical Review Letters*, vol. 87, p. 098101, Aug. 2001.
- [238] L.-Y. Wu and Z.-H. Liu, “Enhancement of Information Transmission by Array Induced Stochastic Resonance in the Processes of Amplitude and Frequency Modulations,” *Chinese Physics Letters*, vol. 23, pp. 1110–1113, 2006.
- [239] J. Butcher, *The Numerical Analysis of Ordinary Differential Equations*. New York: Wiley, 1st ed., 1987.
- [240] C. Moler and C. V. Loan, “Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later,” *SIAM Review*, vol. 45, no. 1, pp. 801–836, 2003.
- [241] J. S. Chang and G. Cooper, “A Practical Difference Scheme for Fokker-Planck Equations,” *Atomic Energy*, vol. 6, pp. 1–16, 1970.